

Misperception and its Evolutionary Value

by

Lachlan Brumley



Thesis

Submitted by Lachlan Brumley

for fulfillment of the Requirements for the Degree of

Doctor of Philosophy (0190)

Supervisor: Dr. Kevin Korb

Associate Supervisor: Dr. Carlo Kopp

**Clayton School of Information Technology
Monash University**

May, 2014

© Copyright

by

Lachlan Brumley

2014

Under the Copyright Act 1968, this thesis must be used only under the normal conditions of scholarly fair dealing. In particular no results or conclusions should be extracted from it, nor should it be copied or closely paraphrased in whole or in part without the written consent of the author. Proper written acknowledgement should be made for any assistance obtained from this thesis.

I certify that I have made all reasonable efforts to secure copyright permissions for third-party content included in this thesis and have not knowingly added copyright content to my work without the owner's permission.

Contents

Contents	iii
List of Figures	xi
List of Tables	xv
Abstract	xvi
Acknowledgments	xviii
1 Introduction	1
1.1 Why Misperception?	1
1.2 The Main Thesis	2
1.3 Analysing and Understanding Misperception	4
1.3.1 Methods for Analysing Misperception	4
1.3.2 Defining Misperception	6
1.4 Aims and Goals of this Research	7
1.4.1 Theoretical Considerations of Misperception	8
1.4.2 Simulating Evolutionary Benefits of Misperception	9
1.4.3 Restoring and Maintaining Cooperation with Misperception	10
1.4.4 Topics Beyond the Scope of this Thesis	12
1.5 New Research Contributions	13
1.6 A Brief Outline	15
2 Background	17
2.1 Biological and Evolutionary Theory	18
2.1.1 The Basic Evolutionary Mechanisms	18

2.1.2	Evolution and Misperception	21
2.1.3	Summary	26
2.2	Evolutionary Artificial Life Simulation	27
2.3	Game Theory	30
2.3.1	Hypergames	34
2.3.2	The Prisoner’s Dilemma	36
2.3.3	Evolutionary Game Theory	41
2.3.4	Summary	43
2.4	Misperception	44
2.4.1	Perceptual Errors	45
2.4.2	Deception	45
2.4.3	Channel Impairment	46
2.4.4	Self-deception	48
2.4.5	The Simulation of Misperception in Artificial Life Environments	54
2.4.6	Rationality and Misperception	58
2.4.7	Summary	60
2.5	Information Warfare	61
2.5.1	What is meant by “Information”?	62
2.5.2	Information Warfare: Definitions and Operations	67
2.5.3	Shannon’s Communication Theory and Information Warfare	83
2.5.4	Domains of Information Warfare	88
2.5.5	Applications of Information Warfare Theory	91
2.5.6	Summary	94
2.6	Perceptual Cycle Models	95
2.7	Categorising the Sources of Misperception	102
2.7.1	Information Gathering Errors	103
2.7.2	Information Processing Errors	107
2.7.3	Incestuous Amplification	113
2.7.4	Summary	116
2.8	Background Summary	116

3	Research Problems and Methodology	119
3.1	Understanding Misperception	120
3.1.1	The Effects of Misperception on the Decision-making Cycle	122
3.1.2	Misperception and the Orientation Step	124
3.2	Misperception in a Foraging Environment	125
3.2.1	Evolutionary Advantage in Akaishi and Arita's Simulation	126
3.2.2	Behavioural Diversity through Misperception	127
3.3	Misperception and the Iterated Prisoner's Dilemma	129
3.3.1	Using Misperception to Restore Cooperation	129
3.3.2	Asymmetric Misperception	130
3.3.3	Preventing Exploitation of Forgiveness	132
3.4	Unaddressed Research Topics	132
3.5	Summary	135
4	Misperception and the OODA Loop	137
4.1	Boyd's OODA Loop Model	138
4.1.1	A Simplified OODA Loop Model	139
4.2	Information Warfare Attacks and the OODA Loop Model	141
4.2.1	Degradation Attacks	142
4.2.2	Corruption Attacks	145
4.2.3	Denial Attacks	148
4.2.4	Subversion Attacks	151
4.3	Information Gathering Errors	157
4.4	Information Processing Errors	159
4.5	Self-deception	161
4.5.1	Suppressing Unwanted Information	162
4.5.2	Aiding Deception	166
4.6	Misperception and OODA Loop Tempo	168
4.6.1	Comparing Tempi	169
4.6.2	A Mathematical Examination of OODA Loop Tempo	171
4.6.3	Is Under-sampling Misperception?	174
4.7	Summary	175

4.7.1	Locations of Misperception Errors	175
4.7.2	Similarities Between Misperception Sources	177
4.7.3	OODA Loop Tempo	179
4.7.4	Information Flow Through the OODA Loop	179
4.7.5	Protecting Against Misperception	180
4.7.6	The Importance of the Orientation Step	181
5	Misperception and Orientation	183
5.1	Revisiting Boyd’s Model	184
5.2	A Derivative Cognitive Architecture	188
5.3	Comparing these Architectures	192
5.4	Expanding the Orientation Step	194
5.4.1	Functions of the Orientation Step	195
5.4.2	Combining these Functions	198
5.4.3	The Procedure of the Orientation Step	200
5.5	Errors During the Orientation Step	202
5.5.1	Identification	203
5.5.2	Interpretation	204
5.5.3	Aims Derivation	206
5.5.4	Options Generation	206
5.5.5	Effects of Misperception on the Orientation Step	207
5.6	Case Studies Focusing on the Orientation Step	210
5.6.1	Corruption Attacks	210
5.6.2	Subversion Attacks	211
5.6.3	Self-Deception	213
5.6.4	Information Processing Errors	220
5.7	Summary	221
5.7.1	Assessment of the Developed Model	221
5.7.2	The Importance of Accurate Prediction	222
5.7.3	Misperception during the Orientation Sub-steps	223
5.7.4	The Importance of the Orientation Step	224

6	An Evolutionary Benefit from Misperception	227
6.1	Akaishi and Arita’s Simulation	228
6.1.1	Description and Simulation Method	228
6.1.2	Results and Summary	230
6.2	Simulation Design and Implementation	232
6.2.1	Changes from Akaishi and Arita’s Simulation	233
6.2.2	Implementation	237
6.3	Simulation Parameters and Execution	241
6.3.1	Simulation Parameters	241
6.3.2	Parallelism and Cluster Usage	244
6.3.3	Fitness Measure	245
6.4	Results	246
6.4.1	Average Misperception Probability	248
6.4.2	Sub-populations of Misperceiving Agents	251
6.4.3	Average Potential Offspring	253
6.4.4	Source of the Evolutionary Benefit	254
6.4.5	Deadlocking Behaviour and Misperception	256
6.4.6	Comparing the Statistically Significant Results	258
6.5	Summary	258
6.5.1	Average Misperception Probabilities	259
6.5.2	Relative Sizes of Sub-populations	259
6.5.3	Potential Offspring	260
6.5.4	How is Misperception Beneficial?	260
7	Behavioural Diversity and Misperception’s Benefit	263
7.1	Introduction	263
7.2	Method	264
7.2.1	Foraging Methods	264
7.2.2	Simulation Execution and Results	266
7.3	Results	267
7.3.1	Average Error Probability	268
7.3.2	Population Proportions	270

7.3.3	Average Potential Offspring	273
7.4	Conclusion	281
8	Misperception and the Iterated Prisoner’s Dilemma	285
8.1	Noise and the Iterated Prisoner’s Dilemma	286
8.2	Restoring Cooperation with Misperception	288
8.2.1	Simulation Parameters	291
8.2.2	Results	292
8.2.3	Discussion and Conclusions	296
8.3	Evolutionary Value of Asymmetric Misperception	299
8.3.1	Method and Parameters	300
8.3.2	Results	303
8.4	Preventing Exploitation from Punishing Misperception	307
8.4.1	Method and Parameters	309
8.4.2	Results	310
8.5	Limiting Forgiveness to Prevent Exploitation	316
8.5.1	Method	317
8.5.2	Results	318
8.6	The Relationship Between Forgiveness and Punishment	326
8.7	Categorising Player Populations	329
8.7.1	Testing the Categorisation	332
8.7.2	Equilibria and Simulation Stability	333
8.7.3	Summary	334
8.8	Conclusion	335
8.8.1	Restoring Cooperation with Misperception	335
8.8.2	Forgiving and Punishing Misperception	336
8.8.3	Patience Preventing Exploitation	338
8.8.4	The Relationship between Forgiving and Punishing Misperception	338
9	Conclusions	339
9.1	Research Findings	340
9.1.1	Misperception and the OODA Loop	340
9.1.2	Evolutionary Value of Misperception	341

9.1.3	Tit for Tat, Cooperation and Misperception	343
9.2	Overall Conclusions	344
9.2.1	The Relationship Between Misperception and Mutation	344
9.2.2	Behaviours Produced by Misperception	344
9.2.3	A Comparison of Simulated Misperception	346
9.2.4	Preventing Cooperation	347
9.2.5	Behavioural Diversity and Deception	348
9.2.6	An Advantage to Ignorance?	349
9.3	Future Research	350
9.3.1	Beneficial Behaviours Induced by Misperception	350
9.3.2	Misperception of Attribute Values	350
9.3.3	Evolution of Self-Deception	351
9.3.4	Examining Behavioural Diversity	352
9.3.5	Further Research of Foraging and Misperception	353
9.3.6	The Role and Effects of Subversion	354
9.3.7	Relationship between Misperception and Noise	355
9.3.8	Further Misperception and the Iterated Prisoner's Dilemma	356
9.3.9	OODA Loop Tempo and Information Warfare	357
9.4	Summary Conclusion	358
Appendix A Glossary for Evolutionary Foraging Simulations		359
Appendix B Algorithms for the Evolutionary Foraging Simulation		367
B.1	Algorithm for Misperception-affected foraging	367
B.2	Algorithm for Misaction-affected foraging	372
B.3	Algorithm for Reflexive-foraging	372
B.4	Algorithm for Perfect-perception foraging	372
Appendix C Average Misperception Probabilities		375
Appendix D Average Misaction Probabilities		377
Appendix E Additional Data for Misperception and the IPD		379
E.1	Probabilities of Noise-Induced State Change	379

E.2	Algorithm for a Trickster-influenced IPD game	381
E.3	Algorithm for an evolutionary IPD game	381
E.4	Comparing Forgiving and Punishing Misperception	385
E.5	Algorithm to limit Forgiving Misperception	402
E.6	Results from Restricting the Forgiving Misperception Probability	402
E.6.1	Forgiving Misperception Probabilities	403
E.6.2	Punishing Misperception Plots	407
E.6.3	Average Final Individual Score Plots	407
E.7	Algorithm to limit forgiveness	411
E.8	Additional Data from Restricting Forgiveness	411
E.9	Additional Population Data for Restricting Forgiveness	413
E.9.1	Forgiving Misperception Plots	414
E.9.2	Punishing Misperception Plots	414
E.9.3	Average Final Individual Score Plots	418
E.10	Additional Lotka-Volterra Model Fit	420
Vita	423
References	425

List of Figures

2.1	A Third Order Hypergame Between Two Players	34
2.2	A General Communication System (Shannon, 1948)	63
2.3	Mills' Paradox (Kopp, 2010)	79
2.4	Compound and Chained Information Warfare strategies (Kopp, 2005b) . .	82
2.5	Neisser's Perceptual Cycle Model (Neisser, 1976)	96
2.6	Learning or Problem Solving Process (Kolb, Rubin and McIntyre, 1971) . .	98
2.7	Boyd's OODA loop model (Boyd, 1986)	99
2.8	A Necker Cube (Necker, 1832)	108
2.9	A Rubin Vase (Rubin, 1915)	108
4.1	Simplified OODA Loop Model	139
4.2	Two Connected OODA Loop Models	141
4.3	A Degradation Attack mapped into the OODA Loop Model	143
4.4	Animal species that perform Active and Passive Degradation	144
4.5	A Corruption Attack mapped into the OODA Loop Model	146
4.6	Examples of Corruption Attacks used during Operation Fortitude	147
4.7	A Denial Attack mapped into the OODA Loop Model	149
4.8	Denial Attacks against Radar Receivers	150
4.9	A Subversion Attack mapped into the OODA Loop Model (Observation) . .	152
4.10	A Subversion Attack mapped into the OODA Loop Model (Orientation) . .	153
4.11	A Subversion Attack mapped into the OODA Loop Model (Compound) . .	154
4.12	Fake antivirus employing a compound Subversion and Corruption attack . .	156
4.13	An Information Gathering Error mapped into the OODA Loop Model . . .	158
4.14	An Information Processing Error mapped into the OODA Loop Model . . .	160
4.15	Self-Deception reducing anxiety mapped into the OODA Loop Model . . .	163

4.16	Self Deception aiding deception mapped into the OODA Loop Model	166
4.17	Effects of OODA Loop Tempo	170
5.1	Boyd's OODA Loop Model (Boyd, 1986, 1996)	184
5.2	The Expanded Orientation Step	199
5.3	Sources of Misperception and the Expanded Orientation Step	209
5.4	Further examples of Corruption attacks used during Operation Fortitude .	211
5.5	Challenger Disaster	214
5.6	Historical Revisionism and Nikolai Yezhov	219
6.1	Average Misperception Probabilities	249
6.2	Average Misperception Probabilities	250
6.3	Proportions of Agent Populations grouped by Misperception Probability . .	252
6.4	Average Potential Offspring of the Sub-populations	254
6.5	Agents clustered around a Resource Node	257
7.1	Average Misaction Probabilities (Misaction-affected foraging)	269
7.2	Average Misaction Probabilities	270
7.3	Comparison of Average Error Probabilities	271
7.4	Proportions of Agent Populations (Misaction-affected foraging)	272
7.5	Average Potential Offspring (Misaction-affected foraging)	274
7.6	Average Potential Offspring (Reflexive-foraging)	275
7.7	Average Potential Offspring (Perfect-perception foraging)	276
7.8	Average Potential Offspring Comparison	278
7.9	Highest Average Potential Offspring Comparison	279
8.1	States of the Iterated Prisoner's Dilemma game	286
8.2	States and Payoffs of the Prisoner's Dilemma Game	289
8.3	Average Individual Score (1% Noise)	293
8.4	Duration spent in the Iterated Prisoner's Dilemma game states (1% Noise)	294
8.5	Average Individual Score (5% Noise)	295
8.6	Duration spent in the Iterated Prisoner's Dilemma game states (5% Noise)	296
8.7	The Effects of Symmetric and Asymmetric Misperception	298
8.8	Average Player Scores and Misperception Probabilities	304

8.9	Average Player Score and Misperception Probabilities (Run 11)	306
8.10	Average Player Score and Misperception Probabilities (Run 5)	306
8.11	Average Forgiving Misperception Probabilities	311
8.12	Average Punishing Misperception Probabilities	312
8.13	Average Final Individual Score	313
8.14	Optimal Forgiving Misperception Probabilities	315
8.15	Stable Population with both Forgiving and Punishing Misperception.	319
8.16	Stable Population with some evolved Punishing Misperception.	320
8.17	Unstable Population affected by runaway Punishing Misperception.	321
8.18	Average Forgiving Misperception Probability	322
8.19	Average Punishing Misperception Probability	323
8.20	Final Average Score	325
8.21	Predator-prey data plotted against Unrestricted Forgiveness.	328
8.22	Distribution of Populations based upon their Final Average Score	330
8.23	Simulations categorised by the Average Score of their final generations.	332
E.1	States of the Iterated Prisoner's Dilemma game with Probabilities	380
E.2	Average Player Score and Misperception Probabilities (Run 1)	387
E.3	Average Player Score and Misperception Probabilities (Run 2)	388
E.4	Average Player Score and Misperception Probabilities (Run 3)	388
E.5	Average Player Score and Misperception Probabilities (Run 4)	389
E.6	Average Player Score and Misperception Probabilities (Run 5)	389
E.7	Average Player Score and Misperception Probabilities (Run 6)	390
E.8	Average Player Score and Misperception Probabilities (Run 7)	390
E.9	Average Player Score and Misperception Probabilities (Run 8)	391
E.10	Average Player Score and Misperception Probabilities (Run 9)	391
E.11	Average Player Score and Misperception Probabilities (Run 10)	392
E.12	Average Player Score and Misperception Probabilities (Run 11)	392
E.13	Average Player Score and Misperception Probabilities (Run 12)	393
E.14	Average Player Score and Misperception Probabilities (Run 13)	393
E.15	Average Player Score and Misperception Probabilities (Run 14)	394
E.16	Average Player Score and Misperception Probabilities (Run 15)	394

E.17 Average Player Score and Misperception Probabilities (Run 16)	395
E.18 Average Player Score and Misperception Probabilities (Run 17)	395
E.19 Average Player Score and Misperception Probabilities (Run 18)	396
E.20 Average Player Score and Misperception Probabilities (Run 19)	396
E.21 Average Player Score and Misperception Probabilities (Run 20)	397
E.22 Average Player Score and Misperception Probabilities (Run 21)	397
E.23 Average Player Score and Misperception Probabilities (Run 22)	398
E.24 Average Player Score and Misperception Probabilities (Run 23)	398
E.25 Average Player Score and Misperception Probabilities (Run 24)	399
E.26 Average Player Score and Misperception Probabilities (Run 25)	399
E.27 Average Player Score and Misperception Probabilities (Run 26)	400
E.28 Average Player Score and Misperception Probabilities (Run 27)	400
E.29 Average Player Score and Misperception Probabilities (Run 28)	401
E.30 Average Player Score and Misperception Probabilities (Run 29)	401
E.31 Average Player Score and Misperception Probabilities (Run 30)	402
E.32 Average Forgiving Misperception Probabilities (2D Plot)	404
E.33 Average Forgiving Misperception Probabilities (Wire-frame Plot)	405
E.34 Average Punishing Misperception Probabilities (2D Plot)	406
E.35 Average Punishing Misperception Probabilities (Wire-frame Plot)	408
E.36 Average Final Individual Score (2D Plot)	409
E.37 Average Final Individual Score (Wire-frame Plot)	410
E.38 Punishing Misperception evolves and disappears on several occasions.	411
E.39 Somewhat cyclic Forgiving and Punishing Misperception.	413
E.40 Average Forgiving Misperception Probability (2D plot)	414
E.41 Average Forgiving Misperception Probability (Wire-frame Plot)	415
E.42 Average Punishing Misperception Probability (2D plot)	416
E.43 Average Punishing Misperception Probability (Wire-frame Plot)	417
E.44 Average Final Individual Score (2D plot)	418
E.45 Average Final Individual Score (Wire-frame Plot)	419
E.46 Predator-prey data plotted against Forgiveness Count simulation data.	420

List of Tables

2.1	Payoffs for the Prisoner’s Dilemma Game.	37
2.2	Payoffs for the Hawk-Dove Game.	42
2.3	Borden and Kopp’s Information Warfare Strategies	78
2.4	Aims and Canonical strategies of Information Warfare.	78
2.5	Libicki’s categories of Information Warfare and the Canonical Strategies . .	79
2.6	Boundary conditions between the canonical strategies of Information Warfare.	81
4.1	Location of misperception-induced errors in the OODA loop	176
4.2	Relationships between the sources of misperception	178
6.1	Agent Density and Resource Density pairings	242
6.2	Basic Metabolic Rate and Parental Investment Cost pairings	242
7.1	Comparison of the foraging methods’ relative effectiveness	277
8.1	Parameter Sets (Low misaction and misperception probabilities)	292
8.2	Parameter Sets (High misaction and misperception probabilities)	292
8.3	Iterated Prisoner’s Dilemma scores for various game states	303
8.4	Thresholds for categorising the state of a population.	331
8.5	Percentage of the populations within each category.	333
C.1	Average Misperception probabilities for the parameter sets.	376
D.1	Average Misaction probabilities for the parameter sets.	378
E.1	Simulation runs where misperception is ultimately beneficial or harmful. . .	387

Misperception and its Evolutionary Value

Lachlan Brumley
lachlan.brumley@gmail.com
Monash University, 2014

Supervisor: Dr. Kevin Korb
Kevin.Korb@monash.edu
Associate Supervisor: Dr. Carlo Kopp
Carlo.Kopp@monash.edu

Abstract

Misperception has a detrimental effect on an entity's perception of its environment, which can affect its decision-making abilities. To those with a perfect perception of an environment, the decisions and subsequent actions of a misperceiving entity may appear to be deficient in some manner. However, in some cases it is possible that an entity or its population ultimately benefits from the effects of misperception. This thesis aims to study how misperception occurs and better explain the conditions under which it may provide a benefit, evolutionary or otherwise.

Misperception can have many possible causes, which may be divided into two main classes — those caused by inadvertent flaws of the perceiving entity and those caused by the actions of other entities. Deliberate actions that are intended to cause an entity to misperceive are instances of Information Warfare attacks. This work produces a new generalised model that describes how both intentional and unintentional misperception affects entities during their decision-making process. These various sources of misperception are mapped into the Observation Orientation Decision Action (OODA) loop model, revealing how misperception affects the decision-making process in the short and long term. The internal process of the OODA loop is also examined and its Orientation step is expanded to detail how an entity's internal representation of the world is developed and affected by misperception, and also how this representation is used to make decisions. Historical case studies of some misperceptions are also mapped into this model, detailing the causes of such misperceptions.

Previous research has identified some instances where misperception provides some benefit to misperceiving entities. This work also aims to identify what circumstances are required for entities to benefit from misperception and determine whether such a benefit can aid entities in an evolutionary environment. Artificial Life simulations are used to investigate environments where entities may misperceive. A benefit from misperception is demonstrated by a population of entities who evolve a significant level of misperception instead of correct perception. The results of these simulations reveal several different methods by which misperception may benefit both individual entities and their populations in different scenarios.

Acknowledgments

This thesis would not exist without the support and assistance of many people. Firstly, I must thank Carlo Kopp and Kevin Korb for their assistance, mentoring, support and inspiration over the course of my PhD program. I must also thank my parents Neil and Glenda, along with friends and family, for their continued support. The funding for my PhD studies was generously provided by the Daffyd Lewis Trust and the ANZ Trustees, for which I am highly grateful. I must also thank all the fellow academics who have provided useful recommendations and feedback, in particular Alan Dorin, Bernd Meyer and William Hutchinson. Finally, I must thank Emily Pilkington for her proofreading, which has greatly improved the quality of this thesis.

Lachlan Brumley

Monash University

May 2014

Misperception and its Evolutionary Value

Declaration

I declare that this thesis is my own work and has not been submitted in any form for another degree or diploma at any university or other institute of tertiary education. Information derived from the published and unpublished work of others has been acknowledged in the text and a list of references is given.

Lachlan Brumley
Wednesday 28th May, 2014

Chapter 1

Introduction

1.1 Why Misperception?

Misperception, confusion and false belief form a good part of the human condition. Their typical consequences are impaired decision-making and a loss of value, at its worst in the form of a loss of life. When such blunders also prevent people from contributing their genes to future generations we may whimsically commemorate their actions (Northcutt, 2002), while concluding that processes which systematically lead to the misperception of our environment will be selected against during evolution and disappear. While common intuition accepts this, it is not so universally true; there is such a thing as a beneficial misperception.

In the winter of 54BC Julius Caesar's legions were attacked by 60,000 Nervii tribesmen under the command of Ambiorix (Caesar and Handford, 1951). The Nervii had besieged Quintus Cicero's legion, but were drawn away by the arrival of Caesar. Caesar's legions quickly constructed fortifications, with a thoroughly fortified gate and a flimsy looking palisade. This induced the Nervii to attack the palisade, dismantling it with their axes and bare hands. As the Nervii broke through the palisade and charged into the camp, a powerful Roman force threw open the gates and set upon the Nervii from behind. Finding themselves between two strong Roman forces, the Nervii panicked and died in huge numbers. While Ambiorix escaped across the Rhine, most of his people were not so lucky. Their villages and fields were destroyed.

Such misperception was highly beneficial. Not for the misperceiving Gauls, of course, but for the deceptive Romans. Misperceptions of this type are common and easily identified. Rarer still, are instances where misperception benefits the misperceiver. An example from evolutionary biology is demonstrated in how organisms perceive pain over time. Pain evolved to warn organisms of bodily injury, forcing them to take protective actions (Broom, 2001). However, a chronic and intense pain is more likely to debilitate an organism than to help it protect itself. But, in general, over time the body becomes desensitised to a constant source of pain, registering less pain from an injury despite no actual change in the state of the injury. This allows an injured animal to better focus on its environment and organise a response to whatever may have caused the injury, or may follow it. Thus, such misperception (of the state of the injury) is beneficial and can be adaptive, as evidenced by the universality of this kind of desensitisation.

While misperception, confusion and false belief are pervasive, and often detrimental to the interests of their subjects, in some few instances they actually benefit their subjects. This may occur accidentally, as a matter of good luck. It may also be a systematic development, as in the case of natural pain management, where the misperception is regularly beneficial. In these cases we should expect evolution to take advantage of the misperception and fix it in a species' genome. After all, evolution does not respect common intuition and will take advantage of anything that comes to hand, including the systematic sources of false belief.

1.2 The Main Thesis

The connotations of misperception are typically negative; a fault that ultimately causes erroneous behaviour in the misperceiver. While misperception may potentially produce an occasional unintentional benefit for the misperceiver, the idea that repeated misperception may regularly provide a benefit seems counter-intuitive. This thesis argues that not only can misperception provide a benefit, but that in some evolutionary environments a regular benefit will increase the misperceiving entities' opportunities to reproduce and pass on the genes responsible for misperception.

There are many different types of organisms, systems and organisations that are affected by misperception and they will be broadly referred to as **entities**. This term is

intended to describe any potential misperceiver, whether it is a biological organism, a machine system, or a social system such as a government, company or military organisation. Entity will be used to refer to any generic misperceiver, while more specific descriptive labels will be used only when appropriate. A single misperceiving entity in a population will be referred to as an **individual**. While this term typically implies a single human, here it describes any type of distinct atomic entity.

Part of this research examines how misperception provides a benefit to affected entities and who these entities are. Advantageous traits observed in biological organisms increase the reproductive success (fitness) of the entities possessing those traits, thereby ensuring the traits remain in the organism's population (Darwin, 1859). For misperception to function in this manner, the benefits an entity accrues from misperception must outweigh any penalties it suffers. It is not necessary, however, for a misperceiving entity to be the only recipient of any benefits from its misperceptions, as other entities may benefit from its misperception-induced mistakes (Hamilton, 1964; Wynne-Edwards, 1962). In this manner, a population of misperceiving entities may benefit not from their own individual misperceptions, but from those of each other.

An evolutionary benefit occurs when the effects of misperception ultimately contribute to an increase in the entity's reproductive success, which is a measure of how many descendants it has parented. For misperception to provide a benefit it may increase an entity's opportunities to mate, increase the amount of resources it gathers or otherwise aid its survival. In instances where non-biological entities are capable of reproduction, or something akin to it, misperception may provide benefits similar to those hypothesised in biological organisms. However, most non-biological entities, such as human organisations and various machine systems, are incapable of reproduction. In these cases misperception can only provide a competitive benefit, which aids the entity's survival and competitive success in its environment. While such cases do not describe an evolutionary benefit, this thesis will still explore how misperception provides a competitive advantage to such non-reproducing entities.

This thesis does not attempt to argue that misperception is beneficial in all situations. Instead it proposes that there are some circumstances where misperception can increase the reproductive and competitive success of entities. Without attempting to be exhaustive, this work will discuss and examine some of these circumstances. In order to support

the argument that misperception can provide an evolutionary benefit, it is necessary to clarify how misperception affects decision-making entities and identify when its effects are beneficial. Existing research has used simulations to demonstrate several instances where misperception benefits either individual entities or their population collectively (Doran, 1998; Akaishi and Arita, 2002b). While such work somewhat supports the hypothesis that misperception provides an evolutionary benefit, further investigation is necessary to better support it. Furthermore, additional examination of misperception will also provide some insight into the conditions under which misperception provides a benefit, evolutionary or otherwise.

1.3 Analysing and Understanding Misperception

Misperception and any evolutionary benefit it provides may be identified and analysed through several different methods. These methods will be used to clarify what misperception is, determine how and when it arises and how it affects entities. Such analysis can indicate circumstances under which misperception provides a benefit.

1.3.1 Methods for Analysing Misperception

Historical case studies describing instances of misperception can be examined to provide some insight into misperception, along with its immediate and long-term consequences. Many serious disasters have been caused by misperception and these are often well documented by multiple sources. This provides a variety of perspectives from which the beliefs, decisions and actions of a misperceiving entity may be forensically analysed. While these disasters do not demonstrate misperception providing a benefit, they do effectively demonstrate its immediate and long-term effects. The correct analysis of historical cases of misperception is dependent upon the quality and accuracy of the historical sources, with any historical revisionism or misconceptions in the source material likely to impact the analysis. Historical descriptions of misperception also only describe its effects on humans and human social systems. Instances of human misperception are easier to understand, but exclusively focusing on human misperception may not develop generally applicable conclusions.

Theoretical models of decision-making and misperception can be used to explain how occurrences of misperception impact an entity's decisions and subsequent actions. These models may also be tested against historical cases of misperception, of which many examples can be easily identified. This theoretical examination requires a decision-making model that is suitable for studying a diverse variety of entities, in order to ensure that any conclusions are widely applicable. One such model is Boyd's Observation-Oriented-Decision-Action (OODA) loop model (Boyd, 1986, 1996), which is sufficiently general to allow the modelling of the decision-making process of any misperception-affected entities.

Simulating misperception allows the investigation of repeated occurrences of misperception in a scenario, showing its effects on the immediate and long-term behaviour of the affected entities. Artificial Life simulations are commonly used to explore evolutionary hypotheses (Langton, 1995) and allow for a multi-generational study of misperception-affected entities, which is necessary to determine whether misperception is evolutionarily beneficial. Beneficial misperception will increase an affected entity's reproductive success, thereby increasing the likelihood that any genes responsible for misperception will spread to future generations. Therefore, misperception will remain in a simulated population's gene pool when it provides an evolutionary benefit. In competitive, non-evolutionary scenarios, beneficial misperception will increase the competitive success of entities.

These are two distinct approaches for analysing misperception — one in theory and one in simulated practice — and the implicit links between these approaches are the various decision-making processes against which the effects of misperception are analysed. The theoretical approach uses the OODA loop model to assess its vulnerabilities to misperception, while the simulated approach considers misperception's effects upon the evolution of Artificial Life agents and players of game-theoretic competitions. In all cases the entities' decision-making model is a key element in how it is affected by misperception.

The various decision-making models (discussed and compared in Section 2.6) used in these investigations all describe the same general process: entities collect information from their environment, assess that information to update their understanding of their environment, choose a series of actions to achieve their goals and then act to reach those goals. Whether such behaviour is described by an OODA loop, the behaviour of a foraging agent or the actions of a player in a game-theoretic competition, it still describes the same fundamental decision-making process, with the same fundamental vulnerabilities to

misperception. Indeed, the simulations could be modified such that the OODA loop model is explicitly written into each simulated entity's behaviour, but such a change would not significantly alter the underlying structure and behaviour of the simulations.

This common underlying structure of perception and decision-making allows misperception to be examined from several different perspectives – specifically its effects upon individuals, sub-groups and entire populations both in theory and in (simulated) practice.

1.3.2 Defining Misperception

Misperception may occur in many different forms and be caused by many different underlying means. In order to focus upon the instances where misperception which may prove beneficial, it is first necessary to clarify what misperception is and how it ultimately affects an entity's actions. It is also necessary to identify the potential benefits of misperception and who may accrue these benefits. While it is obvious that one entity may benefit from another's misperceptions, it is not obvious that such benefits outweigh their costs. Nor is it obvious that an entity may benefit regularly or reliably from its own misperceptions. Finally, it is necessary to explore these beneficial misperceptions in evolutionary environments. The evolution of such misperception will support the main argument of this thesis.

In order for misperception to provide an evolutionary benefit, it must be heritable. Since genetic problems may cause a diverse range of both perceptual and mental problems, some forms of misperception may be considered a genetic trait (Section 2.7). This allows it to be abstractly modelled as the probability of an error occurring that produces a false belief in the affected entity. This error may affect any activity that is responsible for producing the affected entity's beliefs, including information collection and information analysis. Such inherited flaws provide a source of potentially beneficial misperception, which this thesis aims to study through evolutionary Artificial Life simulations.

Misperception is typically assumed to be caused by some unintentional dysfunction of the affected entity. However, it can also have external causes, such as competing entities using methods such as deception. Such acts are performed by an attacker and are intended to cause the victim to act in a manner that unknowingly benefits the attacker. These actions constitute offensive Information Warfare (Schwartz, 1994), a typically hostile act

that affects the victim's information and information processing apparatus, with the intent of altering its behaviour.

While Information Warfare does cause misperception, such misperceptions are typically intended to produce behaviour that benefits the attacker, and are therefore unlikely to benefit the victim. However, it is possible that beneficial misperception may arise from Information Warfare attacks. This thesis will not explore beneficial misperception arising from Information Warfare, since such attacks are almost always selfish, and any benefit experienced by the victim is usually unintentional or coincidental. This specific benefit received by victims may be better explored in future research.

Despite the low likelihood of Information Warfare producing beneficial misperception, the misperceptions that it does produce are still informative, in that the underlying processes of such misperceptions are similar or identical to unintentional and potentially beneficial misperceptions. Moreover, real-world examples of intentional misperception are numerous and well-described. Therefore, studying how such externally induced misperceptions affect an entity's decision-making process will also reveal how unintentional misperception functions. Accordingly, Information Warfare shall be studied due to its ability to cause misperception, but ignored as the source of any beneficial misperception.

1.4 Aims and Goals of this Research

This research aims to demonstrate that misperception can provide an evolutionary benefit in some circumstances and to investigate how this occurs and under what circumstances it does so. This requires an examination of misperception and its effects upon decision-making entities, which will determine how it may affect an entity's decisions and how this may produce beneficial behaviour. Both intentional and unintentional sources of misperception are examined, as both can affect an entity's decision-making processes in a similar manner. The various types of misperception can be classified by their sources, as well as the types of entities who may be affected by these sources of misperception. It is also important to consider the types of entities who may be affected by misperception. While biological organisms and machine systems may be affected by the same types of misperception, it may be due to very different underlying reasons.

Analysing instances where misperception provides a benefit, evolutionary or non-evolutionary, will suggest other potentially beneficial situations. Comparisons between situations with an identified benefit will reveal some conditions under which misperception provides an evolutionary benefit and suggest other circumstances that may also be beneficial. These results will also provide suggestions for future research into the benefits of misperception.

1.4.1 Theoretical Considerations of Misperception

Once the various potential sources of misperception are categorised, they can be analysed to determine how misperception affects the decision-making cycles of the affected entities. Mapping the various possible occurrences of misperception, both intentional and unintentional, into a model of an entity's decision-making process will identify where it causes errors and how it affects decision-making. Determining which points of the decision-making process are affected by misperception will reveal any similarities between the various sources of misperception, how misperception may be prevented, or how its effects may be mitigated if it is unavoidable. The effects of misperceptions on different types of entities can be examined with a suitably generic model, such as Boyd's OODA loop model (Boyd, 1986, 1996). This includes biological organisms, machine systems, individuals, and social systems such as governments, militaries and companies. Analysing documented historical case studies of misperception in the context of the misperceiving entity's OODA loop illustrates the effects of these misperceptions on the entity's decision-making process. The methodology to be used to study how misperception affects the decision-making process can be found in Section 3.1.1.

The Orientation step of the OODA loop model is acknowledged to be an important part of the decision-making cycle, as it is where new information is analysed and interpreted in the context of the entity's existing beliefs. Since the Orientation step operates upon new information collected from the environment, it is likely to be affected by misperception. However, the various possible sources of misperception errors during the Orientation step are not clear, due to the vague description of the tasks performed during the Orientation step. To overcome this shortcoming, a model of the Orientation step's internal processes will be developed. Such a model will better describe the actions of the Orientation step and allow a more detailed examination of how misperception affects an entity during

this step. This model will be validated against historical and hypothetical examples of misperceptions that are known to cause errors during the Orientation step. The analysis of misperception's effects on the Orientation step will further clarify how misperception affects an entity's decision-making process. The methodology to be used to study the Orientation step is clarified further in Section 3.1.2.

1.4.2 Simulating Evolutionary Benefits of Misperception

Few simulators have investigated whether misperception can be beneficial. Presumably, this is due to the expectation that misperception is always detrimental and that any investigation will only confirm this. However, some previous research identifies instances where misperception does provide a benefit, albeit one that is not always evolutionary. If misperception is beneficial in a non-evolutionary scenario, then an evolutionary variant of that scenario may also exhibit a benefit in some conditions. If an entity's benefit from misperception in a non-evolutionary scenario aids its reproduction or survival, then in an evolutionary scenario it may provide an evolutionary benefit by increasing the entity's reproductive success. A simulation of misperception by Akaishi and Arita (2002b) models a foraging scenario where the population's infrequent misperception increases its foraging success. This benefit is incorrectly described as evolutionary, since the simulated agents were incapable of reproduction. However, in an evolutionary environment, the misperceiving agents' increased resource collection may allow them to parent more offspring, thereby demonstrating an evolutionary advantage of misperception. Section 3.2.1 clarifies the methodology that will be used to implement an evolutionary foraging simulation.

Akaishi and Arita's (2002a) research also discusses the issue of the population's behavioural diversity. Misperception was said to be responsible for increasing the diversity of the simulated agents' behaviour and that this diversity was responsible for the observed benefit. Behavioural diversity describes how likely agents are to act differently to others in their population. Whether or not a population benefits from increased behavioural diversity will depend greatly upon its environment, as increased behavioural diversity is unlikely to benefit entities when conformity is desirable. In the case of drivers selecting which side of the road to drive on, conformity to a common behaviour is highly desirable, as any diversity in behaviour is potentially fatal for the individual and others it encounters.

In the foraging scenario, the members of a population gather resources from locations in the environment. A lack of behavioural diversity can cause the population to over-utilise some resources, while ignoring or under-utilising others. Increased behavioural diversity can aid such a population by encouraging some entities to abandon the over-utilised resources and search for under-utilised resources to exploit, leading to more effective resource collection. This may benefit either the entities abandoning the over-utilised resources, the entities remaining at over-utilised resources, or potentially both groups. The underlying model of this scenario may be identified in many different scenarios where populations compete for access to limited resources. Some examples of possible resources include food, transportation network capacity or marketplace niches.

Misperception is said to be responsible for behavioural diversity in the foraging environment, as it gives individual entities unique false beliefs, which cause them to act differently to each other, thereby increasing the population's behavioural diversity. If increasing the population's behavioural diversity with misperception can provide an evolutionary benefit for foraging entities, then a similar benefit may be achieved with any other phenomenon that increases behavioural diversity. Identifying increased behavioural diversity as the source of an evolutionary benefit will confirm how misperception aids entities in this foraging scenario. Section 3.2.2 further clarifies the methodology for studying how behavioural diversity may affect foraging behaviour.

1.4.3 Restoring and Maintaining Cooperation with Misperception

Game theory is a mechanism for mathematically modelling interactions and competitions between entities, along with the decisions they may make in such situations (Morgenstern and von Neumann, 1947). Traditional game theory assumes that all players have perfect information and rationally select actions to maximise their expected payoff. Study of misperception may therefore be better suited to utilising games of incomplete information, which relax the requirement of perfect information. As in other situations, misperception can affect the players by impairing their understanding of the environment or the situation they find themselves in, causing players difficulty in maximising their expected payoffs during games. This problem becomes worse in cooperative games, where players are better rewarded for performing actions that benefit both themselves and their competitors. In such a situation, misperception by either player will lead to a detrimental outcome for both,

allowing players to be negatively affected by the misperception of others in a cooperative game.

Examining misperception's effects on a cooperative game requires selecting a game that provides numerous interactions, which allows players to adapt to their competitor's behaviours and also provides a suitable sample size for random effects, such as environmental noise and misperception. The game should be competitive, while also allowing cooperation to develop or be disrupted. It should also demonstrate the trade-off between short-term selfishness and long-term cooperation.

The Stag Hunt (Skyrms, 2004) and Prisoner's Dilemma (Poundstone, 1992) games best satisfy these criteria, providing competitive and cooperative games where players choose to Cooperate or Defect. The Prisoner's Dilemma game has stronger incentives against Cooperation than the Stag Hunt game, by rewarding players who Defect against Cooperation. This temptation to Defect also ensures that Cooperation is a riskier alternative in the Prisoner's Dilemma game. As such, the Prisoner's Dilemma game is more competitive and provides a scenario where mutual cooperation is much less likely to develop. Therefore, if misperception does benefit players by encouraging mutual cooperation, this effect is more likely to develop in an Iterated Prisoner's Dilemma game.

The Prisoner's Dilemma game and its iterated variant are popular game theoretic models, which are used to study competition and cooperation between players (Poundstone, 1992). The Iterated Prisoner's Dilemma game models a competitive scenario between two players, where repeated cooperation rewards the players with higher payoffs than other non-cooperative outcomes (Axelrod, 1984). The cooperative game modelled by the Iterated Prisoner's Dilemma game can be identified in various biological, sociological and economic competitions, making it a widely applicable model of cooperation between entities. Such competitions include, but are not limited to: arms races (Majeski, 1994), trade disputes (Goldstein and Krasner, 1984), the examination of predators (Milinski, 1987), denouncements during purges (Grossman, 1994), the over-utilisation of communal resources (Hardin, 1968) and the usage of performance-enhancing drugs in sports (Schneier, 2012). The prevalence of the Iterated Prisoner's Dilemma game ensures that research into the effects of misperception on the game is widely applicable.

Cooperation between players may be disrupted by various sources of noise, including environmental noise or dysfunctions of the players such as misperception. Such noise may

occur in any of the possible environments where this game may be played, although it was often neglected in earlier research. A failure of cooperation between players may also prevent players from resuming cooperation in the future, once their trust has been broken. Nations, animals, machine systems, companies and individual people may refuse to resume disrupted cooperation to their own detriment, due to the loss of trust in the other player. In the worst case, this may lead to a highly detrimental state of hostilities between the entities. In order to avoid these detrimental outcomes, the affected entities need some method to restore cooperation. Section 3.3.1 discusses the methodology used to investigate how misperception may restore cooperation between players in greater detail.

Previous research has identified forgiveness and contrition as methods that can allow players to resume cooperation, thereby ensuring that they do not suffer excessively from disruptions produced by noise (Molander, 1985; Wu and Axelrod, 1995). While misperception is capable of acting as a disruptive source of noise, it may also convince players to resume cooperation. In such a case misperception would function as an unintentional form of forgiveness, allowing players to “forgive” the noise-induced mistakes of others. This benefit may also extend to evolutionary environments, such as one where a population of players who evolve over time compete in an Iterated Prisoner’s Dilemma game. An evolutionary benefit to misperception in this scenario will be demonstrated by a population that evolves misperception to allow its members to cooperate despite the interference from environmental noise. Sections 3.3.2 and 3.3.3 discuss the methodology used to investigate the evolution of misperception that implements forgiveness and the methodology used to prevent the exploitation of the evolved forgiveness.

1.4.4 Topics Beyond the Scope of this Thesis

One research area that this thesis does not attempt to address is the evolutionary value of intentional misperception, as implemented by Information Warfare attacks. Information Warfare attacks offer an advantage in competitive environments, which can increase their implementer’s chances of reproductive success or survival. Kopp and Mills (2002) have studied many different biological organisms that have evolved traits implementing Information Warfare attacks, which supports the argument that Information Warfare provides an evolutionary benefit to an attacker.

Self-deception is another source of misperception that this research examines, however, its evolutionary benefit is not examined. While Self-deception has been previously hypothesised to provide an evolutionary benefit through several methods (Trivers, 1976; Ramachandran, 1995; Van Leeuwen, 2007), this thesis will not attempt to validate those hypotheses. Testing the validity of these hypotheses and their potential evolutionary benefit is a potential task for future research.

1.5 New Research Contributions

This research provides new contributions to several areas. The examination of misperception provides an insight into how it alters the decision-making process and ultimately the behaviour of many types of entities. It also shows that different types of entities are affected by different sources of misperception. Simulations of misperception also demonstrate instances where it can provide an evolutionary benefit. Comparing how misperception affects the population in various simulations also reveals how it provides a benefit in those cases. These discoveries are briefly described below, along with their location in the thesis.

A comparative analysis of different models of Information Warfare: Several definitions of information and Information Warfare are compared in an attempt to clarify what is meant by these terms. These definitions of Information Warfare typically focus on its constituent offensive and defensive actions, however there is some difference in the fields and areas that they describe. This work is described in Section 2.5.

A categorisation of sources of Misperception: There are many potential sources of an entity's misperception and these may be broadly divided into intentional and unintentional sources. Intentional misperception is caused by the actions of others or the entity itself, while unintentional misperception is caused by various dysfunctions of the entity as it collects and interprets information from its environment. This categorisation lists all the possible sources of both intentional and unintentional misperception for different entity types. This work can be found in Section 2.7.1 and Section 2.7.2.

A definition of how Misperception subverts rationality: While a misperceiving entity's behaviour may be described as irrational, the entity may be instead making

rational decisions based upon false beliefs caused by misperception. An affected entity will perform actions that appear irrational, yet were the product of rational decision-making. This is explored in Section 2.4.6.

Models showing how Misperception affects the decision-making process: Examining misperception in terms of the OODA loop model reveals that it causes errors in the Observation or Orientation steps, which lead to the creation of false beliefs later in the Orientation step. These false beliefs may then affect a misperceiving entity's decisions and actions until corrected. This is discussed in Chapter 4, specifically Section 4.2.

Models showing how Misperception affects the Orientation step: The Orientation step is responsible for the analysis and interpretation of information, which is an important element of the OODA loop model. Misperception may cause errors directly during this step or indirectly when existing false beliefs are referenced. Creating a model of the internal process of this step provides further insight into how misperception affects decision-making. It also reveals that Self-deception is not the same as reflexive deception, as some have previously argued. Chapter 5 discusses this work.

Simulations showing the effects of Misperception on foraging entities: When Akaishi and Arita's foraging scenario is altered to contain reproducing agents, misperception does provide an evolutionary benefit to the simulated agents. The increased resource collection by the misperceiving agents allows them to parent more misperceiving offspring. Chapter 6 describes this work.

Simulation showing the effects of behavioural diversity on foraging: Misperception is one method that causes errors in an entity's decision-making cycle, thereby introducing diversity into its behaviour. Random errors may also affect other parts of an agent's decision-making cycle and these can also produce a similar benefit when they increase a foraging population's behavioural diversity in a manner similar to misperception. This is explored in Chapter 7.

Simulations where Misperception maintains Cooperation: Misperception is capable of enabling Tit for Tat players to maintain Cooperation in a noisy environment,

by producing behaviour similar to forgiveness. There are, however, some limitations on the conditions under which this occurs. Chapter 8 explores this problem.

A summary of who ultimately benefits from Misperception: In the cases where misperception has been observed to provide a benefit, this benefit is not always received by the misperceiving entity itself but may instead be received by other entities in the environment. Therefore a collective benefit may arise from the misperceptions of entities in the population. Misperception may spread not by increasing the reproductive success of the misperceiving entity, but instead by increasing the reproductive success of related entities that are likely to also possess the genes responsible for misperception. This is described in Section 9.1.2.

A summary of how Misperception may benefit entities: The existing studies of misperception and those undertaken as part of this thesis identify a common pattern among instances where misperception is beneficial. Misperception provides a method for the affected entities to perform beneficial actions that they either would not choose to perform or are incapable of performing. It is expected that this pattern exists in other instances where misperception is beneficial. Section 9.2.2 discusses this hypothesis.

1.6 A Brief Outline

Existing research in the areas explored by this thesis is discussed and summarised in Chapter 2. Chapter 3 describes the identified shortcomings of the existing research and the methods used to address those shortcomings. Chapter 4 maps the various identified sources of misperception into a model of the decision-making cycle. Chapter 5 extends the model of the most important step of the decision-making cycle and explores how misperception functions within this extended model. In Chapter 6, evolutionary methods are added to a foraging simulation that demonstrates beneficial misperception, thereby demonstrating an evolutionary benefit from misperception. Chapter 7 modifies the foraging behaviours of the evolutionary simulation to explore whether the benefit observed from misperception can be replicated through other means. Chapter 8 examines how misperception may maintain and reinforce cooperation between game players who are competing in a noisy

environment. The conclusions of this thesis are summarised in Chapter 9, along with suggestions for future research.

Chapter 2

Background

Determining whether misperception provides an evolutionary benefit requires examining the evolutionary mechanisms that affect organisms. These mechanisms describe how species survive in and adapt to their environment and how misperception may affect these processes. Studying misperception's effects on evolutionary processes will also require suitable methods for modelling biological processes that misperception may affect. Artificial Life aims to investigate biological processes with various simulation methods to reproduce and study these processes in action. Evolution is one such biological process, which can affect a simulated environment where misperception occurs. Another method for modelling the survival contests between organisms is game theory, which provides a mathematical foundation for creating and analysing games of strategy. Game theory allows the construction and analysis of simple models with rules that approximate the intra-species and inter-species contests that occur in the natural world. Examining and exploring misperception and its boundaries is also necessary, since there are numerous potential sources of misperception which may aid an individual or species in an evolutionary situation. Existing simulations of misperception will be examined, along with the methods utilised for simulating misperception.

Misperception may occur in many different types of information-driven systems, such as biological organisms, social systems and machine systems, affecting their ability to correctly function in their environment. While misperception typically produces unintentional errors that affect these systems, Information Warfare describes methods for intentionally producing similar errors. Since both misperception and Information Warfare may produce

similar effects, it is necessary to understand how entities make decisions in order to examine how both misperception and Information Warfare may affect this process and identify any similarities or differences between them.

Before I can study how evolution and misperception affect each other, I must first cover some important background material. More specifically, what is an evolutionary benefit and how do an organism's traits provide such a benefit?

2.1 Biological and Evolutionary Theory

Biological organisms may possess physical or behavioural traits that cause them to misperceive. Evolutionary theory explains how organisms may benefit from their physical and behavioural traits, as well as how these traits arise and propagate through populations. Discussing the various mechanisms of evolutionary theory will explain how such traits can aid organisms and clarify how misperception may provide an evolutionary benefit. While evolutionary theory is considered in terms of its effects on biological organisms, it may also operate in other competitive environments.

2.1.1 The Basic Evolutionary Mechanisms

The modern synthesis of evolutionary theory uses a combination of concepts to explain how species arise, what enables a species to survive, how a successful organism's descendants inherit their traits and how these traits may change over many generations (Patterson, 1999). The major mechanisms of this theory of evolution are selection, inheritance and variation. Darwin's (1859) theory of natural selection explains why organisms with beneficial traits or behaviours have greater reproductive success.

Darwin observed that there was some naturally occurring variance in the physical traits of the individual members of a sexually reproducing species. These traits may beneficially or detrimentally affect the organism's survivability in its environment. Darwin's theory of natural selection argues that an organism's beneficial physical traits will increase its survivability and chances of reproduction, thereby increasing the likelihood that these traits will exist in future generations. An organism's evolved adaptations to its environment represent beneficial physical traits, which are often said to be adaptive. A species adapts to its environment due to the operation of this selective process over many generations.

Organisms with physical traits poorly suited to their environment are less likely to survive and reproduce; conversely, those with well-suited physical traits are more likely to survive and reproduce. In this manner beneficial traits will increase their representation in future generations, while the representation of detrimental traits will decrease. The neo-Darwinian theory of natural selection, however, requires two other essential mechanisms to operate in the manner described (Patterson, 1999). One is a mechanism of inheritance, which describes how physical traits are passed from parents to their offspring. Another is mutation, which introduces variation in the traits held by a population, as well as introducing new traits into the population.

Explaining inheritance requires a brief extrapolation of the underlying genetic model it operates upon. In classical genetics, genes were considered to be abstract units of inheritance that contained an organism's traits and these traits could be inherited by its offspring (Carlson, 2004). An organism's genome consists of an encoded description of its phenotypic structure, where individual genes or groups of genes describe its physical structure and appearance. A gene can have a number of particular forms, called alleles, depending on the phenotypic structure it describes. As an example, in an organism with a single gene encoding its eye colour, individual organisms could possess alleles for blue eyes, brown eyes or green eyes. Sexual reproduction produces offspring who inherit genes from both of their parents. Modern genetic theory (Everson, 2007) describes the gene in more detail; however, the simpler classical model is sufficient to explain genetic inheritance while avoiding unnecessary complexity.

Genetic inheritance describes how traits are passed from parents to their descendants, and how traits may disappear from a population only to reappear in later generations. Mendel's (1865) experiments investigated the genetic inheritance of traits in peas, showing that when plants with different alleles for a trait were cross-bred, the hybrid offspring would inherit one allele while the other would disappear. However, this allele could later reappear in the hybrid's subsequent descendants. Mendel described the allele that was physically expressed as **dominant**, while the other was said to be **recessive**. Recessive alleles allow a particular trait to remain in the gene pool of a species, without requiring it to be widely expressed in the population. Therefore, recessive alleles can maintain a small amount of genetic variation in a species. This genetic model of inheritance describes how an organism's traits are transmitted to its descendants.

Darwin (1859) noted that his theory of natural selection depends upon an instrument to introduce variation into the population, although he did not explain this mechanism. Natural selection requires something to produce variation in the traits exhibited by a population and also introduce new traits. Since an organism's traits are dictated by its genes, diversity can be introduced by a mechanism that periodically produces new genes or alters existing genes. Sexual reproduction with Mendelian inheritance is capable of producing hybrid offspring with new combinations of genes from both parents, thereby varying the combination of traits inherited by offspring. However, sexual reproduction and inheritance cannot create new alleles and therefore cannot introduce new traits in offspring. Another mechanism is required to create new traits in a population or alter the traits that currently exist: mutation.

In the context of evolution and genetics, a mutation is a change in the elements of an organism's genetic code (Patterson, 1999). Mutations may be caused by copying errors in genetic material during cell division or by exposure to various mutation-inducing substances or conditions. Mutations can be categorised into germ line mutations, which can be transferred to descendants, and somatic mutations, which cannot be transferred to descendants (Patterson, 1999). Evolutionary theory therefore focuses on germ line mutations, as they are heritable; any subsequent discussion of mutation refers solely to germ line mutation.

Traits produced by mutation may be beneficial, deleterious, or have no discernible effect on the organism. As deleterious traits adversely affect the organisms that possess them, natural selection reduces their frequency in the population. Conversely, natural selection will favour any beneficial traits produced by mutation and increase their frequency in the population, accumulating in a species over many generations.

A popular example that demonstrates these evolutionary mechanisms is the case of the peppered moth (Majerus, 1998). The peppered moth lives in England, in a habitat containing many lightly coloured trees and lichen. The moths may be either lightly coloured or darkly coloured¹, with the dark coloured allele dominant and the light coloured allele recessive. The lightly coloured moths are well camouflaged in the lightly coloured habitat from predators, while the dark moths are poorly camouflaged. Due to dark moth's less

¹While Majerus (1998) has noted that the peppered moth can also exhibit intermediary shades between the observed light and dark coloured specimens, that topic is beyond the scope of this discussion, which merely aims to demonstrate the adaptive properties of the two different colours.

effective camouflage, birds are more likely to predate upon them than the light moths. The moth population is therefore mostly light coloured, as the light moths are more likely to avoid predation and reproduce.

During the Industrial Revolution, the introduction of polluting coal-burning factories dramatically changed the peppered moth's habitat, killing the lichen and darkening the white trees with soot. This change transferred the camouflage advantage to the dark moths and increased the likelihood that birds would predate upon the light moths. Now the dark moths were more likely to survive and reproduce, which made them numerically dominant in the population and the light coloured moths the minority. In the peppered moth, only one parent requires the dominant dark allele to produce a dark coloured offspring, while both parents must possess the light allele to produce light offspring. The peppered moths quickly adapted to this environmental change, as the few dark moths were able to easily spread the dark allele to their offspring.

The genetic variation provided by the allele responsible for the dark colouring allowed the moths to adapt to the environmental change caused by pollution. In modern times the amount of pollution has been dramatically reduced, which has somewhat restored the habitat to its initial lightly coloured state. This led to an increase in the frequency of light coloured moths in the population, due to the recessive light allele remaining in the moth population.

For misperception to provide an evolutionary advantage, it must be caused by a heritable trait encoded in one or more of an organism's genes. More importantly, it must also benefit the affected organism in some manner, providing an advantage over other competing organisms. There are several different mechanisms by which traits may benefit organisms and I shall describe these now, to identify possible benefits that misperception may provide in an evolutionary environment.

2.1.2 Evolution and Misperception

To assess whether any trait or behaviour provides an evolutionary benefit, it is helpful to have some measure to assess the benefit or cost of the trait or behaviour. Fitness is a term originally used to describe how suited an organism is to its environment (Sober, 2001); however, it is more accurately used to describe an organism's reproductive success (Ridley, 2004). Fitness comes from both an organism's ability to survive and its ability to

reproduce. An infertile organism that excels at surviving in its environment has a fitness of zero, since it cannot reproduce. Similarly, a fertile organism that is exceedingly poor at surviving in its environment has a low or zero fitness, depending upon the efficacy of its survival skills. If misperception is evolutionarily advantageous for an organism, it must increase that organism's fitness, which will noticeably increase the organism's number of descendants.

There are several different means by which an evolved trait can benefit an organism and increase its fitness. In the most straightforward case, the trait directly benefits the organism that possesses it, as in the earlier example of the peppered moth's camouflage. At first glance, it is unclear how misperception may directly benefit a misperceiver, given that any perceptual errors it causes are more likely detrimental than beneficial. If misperception is to directly benefit an individual misperceiver, its overall benefits must outweigh its negative costs. This will depend upon the likelihood of its positive and negative outcomes, as well as the relative values of potential costs and benefits. Situations in which misperception is likely to provide an individual evolutionary benefit may occur whenever misperception leads to behaviour that has some combination of low negative costs, large benefits, or frequently advantageous outcomes. While it is unlikely that these conditions will commonly occur, there are further methods by which misperception can provide an evolutionary benefit.

The previous discussion of evolution considers its effects on the individuals of a species; however, evolution can also be considered to operate on each individual's genes (Dawkins, 1976). From this perspective, successful genes promote their own propagation by bestowing beneficial traits and behaviours upon their host organism. Genes aim only to increase the chances that their host will survive and reproduce, thereby ensuring their own possible inheritance by the host's offspring. From a gene's perspective, its fitness is determined by its frequency in the population. With this concept, Dawkins describes an organism as a survival machine for its genes.

Examining altruistic behaviour reveals one method in which individually detrimental behaviours may still provide an evolutionary benefit. Hamilton (1963) argues that altruistic behaviour is adaptive if it produces behaviour that, while detrimental to an organism, benefits the organism's kin. This is described by the formula $k > \frac{1}{r}$, where k is the ratio of the benefit to the cost of the altruistic behaviour and r is the relationship co-efficient,

describing how closely the altruist and the beneficiary are related. This model assumes that only one individual performs the detrimental behaviour, which multiple relatives may benefit from. By this formula, in the case of a gene that produces altruistic behaviour towards a sibling ($r = \frac{1}{2}$), the benefit from the behaviour must be more than twice the cost for the behaviour to have an evolutionary benefit. If the beneficiary is a half-brother or half-sister ($r = \frac{1}{4}$), the behaviour must provide a benefit more than four times its cost.

A gene can therefore increase its own evolutionary success if it encodes a trait or behaviour that promotes the survival and reproduction of related organisms, who are likely to share this gene. This is referred to as “inclusive fitness”, where the fitness of a given behaviour or trait comes from both its direct effect on the organism and the indirect effect it has on others who may share the genes responsible for the behaviour or trait (Hamilton, 1964). Price (1970) explains this process, whereby the selection of altruistic behaviour of individuals is selected for within groups of related organisms. Price’s model also permits the opposite behaviour, wherein spiteful behaviour can evolve when individuals are negatively related to their neighbouring population (Hamilton, 1970). Maynard Smith (1964) points out that there are two distinct variants of inclusive fitness, which he refers to as “kin selection” and “group selection”. Kin selection is a special case of inclusive fitness (Hamilton, 1964), where traits are selected that favour the survival of close relatives at the expense of the individual organism. Kin selection limits any benefit to those who are closely related to the organism and are therefore likely to share similar genes, including the gene responsible for such behaviour. Hamilton’s concept of inclusive fitness therefore considers the benefit from the gene’s perspective, instead of the organism’s. Individual selection can be considered a special case of kin selection, where the individual is considered to be a relative that shares 100% of its genes (i.e. $r = 1$).

Since these kin are most likely to share the gene or genes responsible for misperception, any benefit from misperception is likely to benefit other entities with the genes responsible for misperception. If kin selection is to explain misperception’s benefit, the benefit that misperception provides to the entity’s kin must be greater than the penalties incurred by the misperceiving entity and its kin in most instances. This can be calculated by summing the cost or benefit of misperception to each entity, weighting this value by how closely related the entity is to the misperceiver. If this value is greater than zero, then there is an overall benefit (Equation 2.1). Here i represents the i th entity in the population,

b_i is the benefit an entity receives from the misperception, c_i is the cost to the entity of misperception and r_i is the coefficient of relatedness of entity i to the misperceiving entity, which is the probability that it shares the gene or genes responsible for misperception.

$$\sum_i (b_i - c_i)r_i > 0 \tag{2.1}$$

Like kin selection, group selection also promotes the spread of traits that benefit others instead of the individual organisms. Group selection occurs when the organisms in a population all share a trait that incites beneficial behaviour from each other (Wynne-Edwards, 1962). Group selection theory argues that for such adaptations, selection operates at the level of the group, instead of at the level of the individual or the gene. This allows groups or populations that share a collectively beneficial adaptation to become more successful than other groups who lack such adaptations. Unlike kin selection, there is no familial restriction on the beneficiary of the behaviour. Maynard Smith (1964) argues against group selection, on the basis that an adaptation that benefits only the group is unlikely to become fixed in the population in the first place. For a trait to spread throughout a group, it needs to benefit either individuals or their kin. If the trait has spread throughout the population due to individual or kin selection, then group selection is not required to explain any benefit from this trait. While group selection cannot explain how misperception might evolve in a population, it may explain how misperception is beneficial in the specific case of a group in which all members of the group are capable of misperception.

Kin selection and group selection both describe mechanisms by which misperception can provide an evolutionary benefit, despite its expected drawbacks to the misperceiving organism. In such cases, the benefit shared among the other organisms who possess the trait is greater than the cost experienced by the misperceiver. However, in many real world situations, misperception is unlikely to restrict its benefits to others who possess a trait responsible for misperception or even to those who are closely related. Therefore, any misperceptions that benefit kin may also unintentionally aid unrelated organisms. Misperception may also provide a benefit through group selection, if the misperceptions of individual organisms provide a net benefit to others in the group. However, for misperception to benefit a group, it requires all the group members to already have the trait responsible for misperception.

The evolution and spread of group-beneficial traits, especially in social species, may be better understood through multi-level selection (Wilson and Sober, 1994). Multi-level selection states that there are many levels of competition and evolution, which are the genes, cells, organisms and groups of organisms. Competition and evolutionary pressures occur at each of these levels in differing ways, with each level functioning together to maximise the fitness of an organism and its species. Traits that rely upon group selection can thereby evolve and spread within a species, assuming that the trait's benefit at the group level outweighs its effects at the individual organism level. As such, evolution can select for traits that benefit non-kin organisms, providing an explanation for the evolution of social behaviours that benefit groups.

Another method that explains the evolution of detrimental traits in an organism is Darwin's (1871) theory of sexual selection. Sexual selection uses natural selection to explain the physical differences of the different sexes. This occurs in two main forms: inter-sexual competition and intra-sexual competition. Inter-sexual competition involves competition between two organisms of the less limited sex (typically males) competing among themselves for access to members of the limiting sex. The limiting sex is the one that must pay the cost of giving birth to and raising the offspring, while the less limited sex typically has a much lower investment in an offspring. Intra-sexual competition, which is also referred to as 'female choice' or 'mate choice', occurs when males must compete among each other to be selected by females, or vice versa. The existence of such competitions has led to the evolution of traits in the population that specifically aid organisms in these situations. In species with sexual competition, organisms evolve physical traits that aid in their competition with others, like weapons or adornments. Weapons allow organisms to settle territorial disputes through physical conflict, which may prove harmful to any participants; the antlers of deer are one such example. Adornments, such as the peacock's tail or the male lyrebird's song, allow organisms to communicate their relative strength or fitness to each other without resorting to harmful conflicts. Species with inter-sexual competition develop sexual ornaments that are often colourful and highly distinctive. The selecting sex chooses a mate based upon the perceived beauty of its ornament. The bright and colourful plumage of peacocks is a well-known example of inter-sexual competition, which attracts peahens who select a mate based upon the perceived beauty of his plumage.

Sexual selection has also been used by Miller (2001) to explain the evolution of the human mind and some human behaviours where the possible benefits to survival are not obvious. This theory proposes that human behaviours, such as the creation of art or humour, evolved due to selection by both males and females, where these behaviours act as indicators of desirable traits such as intelligence and creativity. These behaviours are therefore analogous to sexual ornaments in humans.

These ornaments are a trade-off for the organism between the reproductive benefit they provide and their detrimental side effects. A peacock is disadvantaged by the time spent preening its tail and the overt nature of its tail, which may attract predators. A peacock also develops its tail at the expense of other beneficial survival aids, such as muscles or claws. Zahavi's (1975) 'Handicap principle' hypothesises that such expensive traits evolve as a method of truthfully signalling the bearer's general health and suitability as a mate. The bearer is able to survive disease, while avoiding starvation and predation until it can reproduce, despite the handicap of its sexual ornamentation. It is highly unlikely that misperception has evolved as a sexual ornament, as it is not a visually observable trait and its existence may only be inferred from observed behaviours under conditions where the behaviours are manifested. In contrast, most ornaments are intentionally highly visible.

2.1.3 Summary

If misperception is to provide an evolutionary advantage, then it must somehow increase the chances of survival and reproduction for organisms possessing the genes responsible for misperception. This may occur if misperception directly increases the fitness of the misperceiver itself, although misperception can also cause organisms to make harmful mistakes, much to their detriment. Therefore, organisms with genes that encode misperception may not directly benefit from misperception unless the benefit to an organism outweighs the costs of its errors. Evolutionary theory explains how new adaptations can evolve in a population and spread throughout that population, if the adaptations aid survival or reproduction. Misperception may benefit individual organisms in some situations, although its effects will often be detrimental, arguing against any evolutionary benefit. However, in situations where misperception is detrimental to individuals, it may still benefit organisms through inclusive fitness.

2.2 Evolutionary Artificial Life Simulation

Misperception is a difficult phenomenon to reliably observe and measure in the physical world. Therefore, observing the behaviour of species in situations where they may misperceive is unlikely to be a suitable experimental method for investigating the evolutionary benefit of misperception. Any examination of misperception requires the ability to determine when misperception has occurred and measure its effects on an evolving population. Artificial Life simulations are a highly suitable tool for examining misperception since they allow one to study evolutionary processes that one cannot otherwise easily observe or measure in the physical world over reasonable time frames (Langton, 1995).

Artificial Life aims to examine biological processes, environments and organisms through the creation and study of artificial systems. These systems may take the form of computer software, robotic agents or chemical compounds. Artificial Life systems allow the study of many processes by simulation, which would otherwise be impractical or impossible. Since evolutionary changes in a species take many generations to appear, the direct observation of evolutionary processes is restricted to species with a very short lifespan, such as the fruit fly *Drosophila melanogaster*. This is one reason why evolutionary processes are often studied with Artificial Life simulations. Organisms modelled in a simulation are called **agents** and simulations focus on the behaviours and interactions of agents in a population over time.

Artificial Life aims to investigate both how life operates and how it may operate in the future (Langton, 1995). Artificial Life may therefore simulate both real and hypothetical biological systems, unlike traditional biology, which is restricted to existing biological systems. Another difference between Artificial Life and traditional biology is that Artificial Life creates a “kind of life” in order to study biological processes.

Bonabeau and Theraulaz (1995) state that the methodology of Artificial Life is “aimed at explaining high-level behaviours from low-level causes” and therefore takes a bottom-up approach in creating models of biological systems, where the rules for a simulation are explicitly defined and coded into the simulation. The agents in the simulation then operate under the restrictions of these rules. Complex behaviours are expected to emerge in the population due to these rules and the interactions of multiple agents with the environment and each other. The simulation’s programmer does not define and enforce

this emergent behaviour; instead it arises as a by-product of the simulated agents' actions and interactions.

The first example of Artificial Life was devised by von Neumann (1951), who developed a complicated self-replicating cellular automaton. As a self-replicating machine, the automaton is analogous to a biological organism. Cellular automata exist in a simulated world consisting of a grid of cells, each of which has an associated state at any discrete time. The state of a cell at any given time t is determined by a function of the states of its neighbouring cells for $t - 1$. Each cell is governed by the same rule for calculating its next state, along with the present state of its neighbours. The cellular automaton created by von Neumann used 29 states to replicate itself. Much simpler cellular automata were proposed in the Game of Life (Berlekamp, Conway and Guy, 1982), where cells have only two states (alive and dead) and three rules determine their state. These rules are as follows:

1. If a live cell has 2 or 3 live neighbours, it stays alive
2. If a dead cell has 3 live neighbours, it becomes live
3. In all other cases the cell stays or becomes dead

The grid environment is initialised with some configuration of alive and dead cells. These simple rules and initial state can lead to the emergence of structured shapes of multiple cells, which develop and migrate across the grid. Static patterns ("still lives"), repeating patterns ("oscillators") or moving patterns ("gliders") may emerge within the environment. Gliders, as moving entities, can transmit information within the grid. The emergence of these patterns in the Game of Life is an example of how simple rules can produce a complex and initially unexpected result. It should be noted that evolutionary Artificial Life simulations are actually complex cellular automata, where the simulated agents' behaviour is dictated by more complex rules and the interactions between agents produces an observable emergent behaviour. Goldstein (1999) defines emergence as "the arising of novel and coherent structures, patterns and properties during the process of self-organization in complex systems".

Another example of emergent behaviour is found in Reynolds' (1987) simulation of the flocking behaviour of birds, fish and insects. Reynolds' agents were called **boids** and modelled virtual birds who flew together in flocks through a three-dimensional environment. A

boid's movement is controlled by a few simple rules, intended to prevent it from colliding with others in its flock. Boids will adjust their headings to avoid collisions, maintain the same velocity as nearby boids and attempt to stay close to nearby flockmates. These three simple rules enable a flock of boids to travel together, with each boid determining its heading and velocity from those of nearby boids in the flock. The boids simulation is similar to a cellular automata, as each boid is independent and its future state depends upon the current state of its neighbours. Unexpected behaviour can also emerge in a flock; for example, when a flock encounters a stationary obstacle, it divides into groups to move around the obstacle and then rejoins into a single flock after navigating around the obstacle. This unexpected behaviour is emergent, as the boids' behaviour rules do not explicitly describe how to move around obstacles. The boid behaviour rules only provide a simple explanation for how the complex structure of a flock can move.

Artificial Life simulations are often used to simulate evolutionary processes in hypothetical environments. One such simulation, *Tierra*, explored the evolution of agents that were sequences of machine instructions competing over access to a computer processor (Ray, 1992). *Tierra* also demonstrated the unexpected evolution of parasitical agents, which utilised the reproductive capabilities of their hosts to reproduce (Adami, 1998). The evolution of an organism's physical structure is another process that cannot be examined in the real world, yet one which has also been successfully studied within Artificial Life simulations (Sims, 1994b,a).

Nature has long inspired solutions to problems and the mechanisms of evolution are no exception, as evolutionary processes may also be adapted to solve problems. Evolution can be considered as a search technique, where the environment is considered a problem and the species a solution to that problem. As evolution causes a species to adapt to its environment, it is improving the survival adaptations, or solutions to the environment, embodied by individuals of the species. This link between evolution and problem solving was examined by Holland (1975), who developed a search mechanism called a Genetic Algorithm to exploit the power of evolutionary processes. Genetic Algorithms are typically used to solve optimisation problems, where the desired solution maximises or minimises one or more variables of the problem. Genetic Algorithms have many potential applications and are typically used to search large solution spaces and optimise known problem solutions (Goldberg, 1989). A Genetic Algorithm creates a population of agents, each of which

represents a potential solution to the problem under investigation. These agents are then subject to an evolutionary competition, where the “better” solutions are permitted to reproduce, improving the quality of the solutions encoded in the agent population.

Since Artificial Life allows the simulation of evolutionary processes under controlled circumstances, it is a suitable tool for investigating misperception’s evolutionary value. Artificial environments may be simulated where the misperceptions made by the agents may be measured and observed, along with the effects of interactions between the misperceiving agents.

2.3 Game Theory

The competitive social and biological contests that occur in evolutionary situations commonly have simple underlying rules, which aids in their analysis. Due to their simplicity and strategic nature, these conflicts can often be modelled mathematically with game theory in order to examine which strategies entities should select and how their opponents may react. These models can also help investigate whether misperception can benefit competitive entities.

The fundamentals of game theory were first formalised by Morgenstern and von Neumann (1947), with the aim of providing a mathematical method for modelling and analysing games of strategy. A game consists of a number of players, the rules for the game and the strategies that the players may select. In a game, each player chooses one of his or her strategies simultaneously, with the resulting combination of strategies referred to as an outcome. Players receive differing payoffs for the various outcomes that may arise due to their choice of strategy. Players are expected to select strategies to cause an outcome that maximises their expected payoff and such selection is described as rational. A game is said to be a zero sum game if one player’s loss always equals its opponent’s gain, otherwise it is a non-zero sum game. An outcome from which no player has an incentive to change its strategy, in order to receive a better payoff, is called an equilibrium (Nash, 1950). It is possible for a game to have zero, one or more outcomes that are equilibria. Since players have no incentive to change their strategy at these equilibria, they are typically considered to be solutions or likely outcomes of a game — assuming rational players. Some strategies may offer a player its best payoff regardless of the strategies selected by its opponents; such a strategy is said to be dominant.

While it is possible to model many situations using traditional game theory, there are some limitations. One is the assumption that all players possess perfect information concerning the game. This information includes who all the players are, what strategies they may select, what outcomes are possible and the payoffs each player receives from each outcome. Another assumption is that the players act in a rational manner, always preferring outcomes with higher payoffs. These assumptions cause some difficulty when attempting to model real world situations, as entities may lack perfect information or act irrationally. Misperception will affect a player's perception of a game and can therefore affect the strategies a player chooses. While a misperceiving player can still act rationally to maximise its expected payoff, since it does not correctly understand the game it is playing, its behaviour will likely fail to maximise its actual payoff.

The applicability of game theoretic models to people can be problematic, as people may lack perfect information or may not make rational decisions. Simon (1957) has stated that people only act rationally enough to "satisfice". By satisficing, people select a strategy that is "good enough", instead of one that is optimal. A truly rational player, however, will always act to optimise its payoff. Even when people have complete information, they may not choose the rational strategy, as Tversky and Kahneman (1988) have observed. If the information is presented in a manner that obscures the dominant strategy, then it is less likely to be selected. The Dollar Auction game (Shubik, 1971) demonstrates such irrational human decision-making, where multiple players bid in an auction of a single dollar. Unlike a typical auction, the players with the two highest bids must both pay after the auction is complete. This modification causes players to bid more aggressively, to avoid losing the money they invested in previous bids. Once the two highest bids total more than one dollar, the players become collective losers. The game then changes into one where the players both attempt to minimise their losses, by winning the auction. At this point the players typically continue bidding against each other, instead of accepting a smaller loss. If the dollar is auctioned for more than its face value, the auctioneer is the true winner.

In some cases, the application of game theoretic analysis to human behaviour can explain behaviour that appears irrational. Kaminski (2004) observed interactions between prisoners and prison guards, analysing their behaviour from a game theoretical context.

Prisoners played games to sort out their social hierarchy or divide resources between themselves. Kaminski observed seemingly irrational acts of self-harm performed by prisoners, including intentionally attempting to contract diseases. If these attempts succeeded, the jailers would transfer the prisoners to the more comfortable prison hospital or reduce their sentence. The prisoners found the benefits of self-injury outweighed its costs, thereby making self-harm a rational strategy.

It is not only prisoners that play games; Berne (1964) has described many psychological games that people play against each other. While not expressed in purely game theoretic terms, Berne's psychological games are similar to traditional games, as they contain players, rules, strategies and payoffs for the players. Berne's players follow a script, which describes the strategies they may select to reach an intended outcome. Unlike traditional game theory, players in Berne's games are not assumed to be rational or to have perfect information and their actions are often driven by ulterior motives. Some of Berne's games also contain an antithetical strategy or antithesis (Berne, 1973), which is a strategy that allows players to leave a losing game they are forced to play, thereby preventing a serious detrimental payoff or attempting to gain a more beneficial payoff. An antithetical strategy allows players to change the rules of the game, presumably in their favour. In most game models, this type of strategy is an extreme action that is not explicitly defined, so players are unlikely to identify it. Wagner, Cheung, Ip and Lee (2005) use Berne's methodology to model an argument between members of a virtual community, identifying the members attempting to use an antithetical strategy (although they refer to it as an anti-game strategy). These players used humour as the antithetical strategy to defuse an argument between others.

A game theoretic analysis of Stalin's Great Terror reveals an excellent historical example of an antithetical strategy (Grossman, 1994). The Terror was a large-scale purge undertaken from 1936 to 1938, to remove people considered to be politically unreliable. Victims were collected and forced through interrogation to confess to "crimes" against the state and to implicate any friends and neighbours who were also "guilty". The Terror was self-sustaining, with those who were interrogated implicating others for interrogation. Grossman describes the antithetical strategy played by a nameless doctor from Kharkov, who complied with his interrogator and listed every single doctor in the city as an accomplice. The interrogator was unhappy with this list, since purging most of Kharkov's

doctors was unacceptable for the state. When the interrogator attempted to force the doctor to remove some names from the list, he refused. The doctor then contacted his interrogator's superiors, to denounce the interrogator as a counter-revolutionary who was protecting enemies of the state. With this act, the doctor implicated his interrogator as an enemy of the state. This change in behaviour was argued to be instrumental in easing and ultimately ending the Terror, as others learned of the doctor's ploy.

Morgenstern and von Neumann (1947) believed that games could be analysed by creating a number of other games that would only exist if any player could choose a strategy with knowledge of its opponent's selected strategy. These new games, called minorant and majorant games, could then be analysed to find equilibria for the original game. Howard (1971) later refined this idea to produce so-called "metagames", where a game is developed for each player that models the player's strategies as if it knows the strategy its opponent has selected. Metagames may be written in a tabular form like traditional games, however size considerations make this impractical for all but the most trivial games. If a player's opponent has a dominant strategy, the player can always select a strategy to maximise its payoff. Such a strategy is called a "sure thing" and the Defect strategy of the Prisoner's Dilemma game is an example (see Section 2.3.2). If each player's sure thing strategy intersects, then that outcome is an equilibrium point for the game. As with Nash equilibria (Nash, 1950) in traditional games, these equilibrium points are stable and represent solutions to the game. A game with no equilibria is considered to be unstable and Howard (1987) states that in such a case, one player will wrongly choose an outcome that does not maximise its payoff.

Fraser and Hipel (1984) propose a slightly different form for game theoretic models, which are intended to improve the suitability of game theory to analyse and resolve conflicts. In these games, each player's strategy consists of a number of options, each of which may or may not be performed. The complete set of possible outcomes for a game is produced by considering all the combinations of options and strategies and then removing any infeasible outcomes. Each player has a preferential ranking of the possible outcomes, which is used to calculate stable outcomes to the conflict. Any outcomes that are stable or equilibria represent solutions to the conflict.

Models of games are commonly displayed as either a matrix of payoffs or a table of options, which can be unintuitive and does not show the relationships between the game's




Figure 2.1: A Third Order Hypergame between players A and B, where each player’s perception is modelled by a Second Order hypergame. In this way each player has a representation of the game it believes it is playing and its own perspective of its opponent’s perspective of the game.

possible outcomes. To overcome this problem, games may be modelled with graphs, where the outcomes of a game become the nodes of a graph and the transitions between outcomes become directed links between the graph nodes (Kilgour, Hipel and Fang, 1987; Fang, Hipel and Kilgour, 1989, 1993). These graphs may be analysed to identify equilibria (Kearns, Littman and Singh, 2001).

2.3.1 Hypergames

Traditional game theory’s requirement of perfect information typically prevents its use in modelling situations where players may misperceive. Hypergames are one method that allow players to possess imperfect information (Bennett, 1977). A hypergame uses multiple games to model each players’ perceptions of the situation. Hypergames may also be nested, so that each player’s perception can itself be a hypergame. This can lead to models containing various levels of perceptions, such as “A’s perception of B’s perception of A’s game”. Figure 2.1 shows a third order hypergame between two players, which consists of eight individual games modelling the various perceptions of the two players.

A first order hypergame uses a single game to model the perception of each player. Each player is unaware that its opponent has a different perception of the situation. In a second order hypergame, each player’s perception is modelled by a first order hypergame and each player knows that its opponent’s perception may differ to its own. A third order game consists of a second order game for each player, allowing a deeper consideration of each player’s perceptions. In general, the complexity of a hypergame is the number of

games required to model all the perceptions of the players. This complexity is given by p^n , where p is the number of players and n is the order of the hypergame. Hypergames are typically limited to the third order to avoid excessive complexity. Most situations being modelled also do not require the additional perception that moving beyond a third order hypergame provides. Methods for analysing hypergames to identify equilibrium states have been developed (Fraser and Hipel, 1984; Takahashi, Fraser and Hipel, 1984; Wang, Hipel and Fraser, 1989) and as with normal games, these states may represent solutions or final states of the situation modelled.

Vane (2006) states that hypergame theory provides a “representation of possibility and likelihood with an embedded representation of error” and discusses the potential applicability of hypergame theoretical methods as tools to aid decision-making under uncertainty. Vane also describes a mathematical measure of the expected utility from a hypergame, called Hypergame Expected Utility, which considers the uncertainty that a game representation is correct and the expected utilities obtainable from the worst-case scenarios of that game. Like Kopp (2003) and Jormakka and Mölsä (2005), Vane also mentions that using a Hypergame analysis to determine a strategy can be expressed in terms of Boyd’s (1986) OODA (Observation, Orientation, Decision, Action) loop model (See Section 2.7). Therefore, the development of the hypergame representation and Hypergame Expected Utility plot encapsulates the information collection and analysis steps of the loop. This plot identifies whether an individual should develop and execute a plan or instead collect and analyse more information to further reduce its uncertainty. During the execution of a selected plan ongoing cycles of information collection and analysis will provide feedback as to whether the model developed and acted upon is accurate.

Hypergames can model situations where information is hidden from one player who is ignorant of the true state of the game it is playing, as shown in a hypergame study of the Fall of France in 1940 (Bennett and Dando, 1979). The Allied armies in France were expecting a German invasion to come either directly at their fortified border, or from the north through Belgium. The German armies, however, perceived a third option of an attack through the Ardennes — a region that was considered by most on both sides to be impassable by armoured vehicles. The model proposed for this game is a first order hypergame, where each player sees a different game, but is unaware that its opponent does not share the same perception. Historically, Germany achieved strategic surprise by

invading through the Ardennes, surprising the Allies with an attack they believed to be impossible (Pitt, 1994).

More complex games with multiple levels of perception may describe situations where players attempt to second guess each other, as in Fraser and Hipel's (1984) third order hypergame model of the deceptions behind the D-Day landings. In this hypergame model, the Allies may invade at either Normandy or Calais, while the Germans may defend at either Normandy or Calais. The model developed shows that the Allies perceive that they have convinced Germany that any invasion at Normandy is a feint and Calais is the real target. The Allies will therefore invade at Normandy, believing it to be less defended. The Allies' perception of the German game has the Germans defending Calais against an Allied invasion at Calais. From the German point of view, Normandy is considered to be a feint and Calais the real invasion target, so Germany will therefore defend Calais. In Germany's perception of the Allies' game, the Allies will invade Calais, believing their feint at Normandy has worked.

Hypergames allow game theoretic methods to model situations where one or more players lacks a perfect understanding of the situation, possibly due to misperception. Hypergames can help to illustrate the importance of correct information and how the outcome of a situation can change when one or more players lack a perfect understanding of the scenario. While hypergames extend traditional game theoretical methods by allowing the modelling of situations with differing perceptions, the additional layers of complexity added by modelling the varied perceptions of all the players complicates both the development and analysis of the model. In situations where there are many players or layers of perceptions, a hypergame model will often be unwieldy to create and analyse. Computer software may be used to aid the analysis of hypergames (Fraser and Hipel, 1984; Brumley, 2004), however these tools do not completely remove the complexity of the situation. Therefore, a hypergame model should only be used when the modeller requires the players to possess varying perceptions of the modelled situation.

2.3.2 The Prisoner's Dilemma

The Prisoner's Dilemma game is a widely used and well understood model for examining cooperation between two players (Tucker, 1950). It models a situation where individually

beneficial selfish behaviours become collectively detrimental when selected by both players. The name derives from a description of such a scenario, where the police capture two men (the prisoners) near the scene of a robbery. The police suspect the two prisoners were guilty, but without a confession they lack sufficient evidence to charge either prisoner. The police separate and interrogate the prisoners, to prevent their communication. Each prisoner is presented with the opportunity to confess to the crime, implicating his accomplice in exchange for his own release. This outcome gives the silent prisoner a long jail term, while the confessing prisoner spends no time in jail. If both prisoners confess and implicate each other, then the police do not know who is most responsible and both prisoners receive moderate jail terms. If both prisoners remain silent, they both receive a short jail term. In this game each prisoner may **Cooperate** by staying silent or **Defect** by implicating his accomplice. Table 2.1 shows the payoff table for a generic dilemma game. The Prisoner's Dilemma game is said to be a 2×2 non-zero sum game, as there are two players who each have two options and one player's gain does not equal its opponent's loss.

	Cooperate	Defect
Cooperate	3, 3	0, 5
Defect	5, 0	1, 1

Table 2.1: Payoffs for the Prisoner's Dilemma Game. Each player's best payoff occurs when he or she Defects against a Cooperating opponent; however, this leads to a sub-optimal outcome if both players choose Defection, which is the Nash equilibrium (Nash, 1950) for this game.

Rational players desire to maximise their payoff by Defecting against a Cooperating player and this payoff is referred to as the **Temptation** to Defect. Players receive the second best payoff when both Cooperate, which is labelled as the **Reward** for mutual cooperation. The second worst payoff is received when both players Defect and this payoff is labelled as **Punishment** for mutual defection. The worst payoff a player can receive is labelled as the **Sucker's** payoff and is received when a player Cooperates while its opponent Defects. Defection is the dominant choice in the Prisoner's Dilemma game, as each player always receives a higher payoff by Defecting, regardless of its opponent's choice. Therefore, mutual Defection is the game's Nash equilibrium (Nash, 1950). This equilibrium reveals the dilemma — while Defection is the dominant and rational choice for both players, mutual Defection provides a lower individual payoff than mutual Cooperation.

However, mutual Cooperation is unstable, since either player can improve its payoff by Defecting. Therefore, mutual Defection is a deficient outcome. Rapoport and Guyer's (1966) taxonomy of the possible 2×2 games places the Prisoner's Dilemma into a category by itself. Among the 78 possible types of 2×2 games, it is the only game with a strongly stable deficient equilibrium, with neither player having any reason to select Cooperation over Defection.

Brams (1977) later examined Rapoport and Guyer's taxonomy and categorised the games based upon the effects of deception on the player's strategies. In these games one player was informed of its opponent's preferences by the opponent, who was able to deceive that player. A deceiver can convince its opponent to select the deceiver's most preferred strategy, unless the potential victim has a dominant strategy and therefore no incentive to alter its strategy. Brams demonstrated that players with a dominant strategy, as in the Prisoner's Dilemma game, cannot be deceived in this manner.

The paradoxical behaviour of the Prisoner's Dilemma game has been observed in real world situations such as Stalin's Great Terror (Grossman, 1994) and cooperation in stickleback fish during predator evaluation (Milinski, 1987). It can also be found in the tragedy of the commons (Hardin, 1968). Those who find themselves in these cases must choose between "rational" selfish behaviour and "irrational" cooperation, just like the prisoners.

The Prisoner's Dilemma also changes substantially when it is played repeatedly by players who remember previous games. This allows players to punish an opponent for past Defections or reward past Cooperation when selecting their strategy. The iterated game is played for a number of rounds and the payoffs totalled from each individual game. The number of rounds is typically concealed from the players, since awareness of the number of rounds can influence their strategies. Luce and Raiffa (1957) have argued that for finitely repeated Prisoner's Dilemma where the players know the number of iterations, cooperation is not rational. This is argued through a kind of reduction by backward induction. Since the players will not meet again in the last round, the Nash Equilibrium is mutual Defection. As mutual defection is now assured in the last round, the penultimate round is now the last round where decision-making occurs. There is now no reason to cooperate on the penultimate round, so defection is again the players' only rational choice. This process continues to the beginning of the game, making mutual Defection the only rational strategy for an Iterated Prisoner's Dilemma game where both players know the

number of iterations. Such detrimental reasoning can be avoided by preventing players from knowing the length of the game, which leaves no apparent final round for the “first” mutual Defection.

Players can maintain mutual cooperation in the Iterated Prisoner’s Dilemma since the potential for punishment in the future deters defection in the present. Players use some strategy to determine their next move, which considers the outcomes from previous iterations of the game. Axelrod (1984) examined strategies for the Iterated Prisoner’s Dilemma in a round robin tournament. Rapoport’s strategy of **Tit for Tat** won, by Cooperating in the first round and then mimicking its opponent’s previous move in all subsequent rounds. Interestingly, Tit for Tat cannot “win” any match by outscoring its opponent and at best may only match an opponent. Axelrod organised a second tournament with some new strategies and it was once again won by Tit for Tat. Analysing the tournament results revealed several common attributes among the more successful strategies: being nice, being forgiving, being non-envious and being retaliatory (Axelrod, 1984). A nice player is not the first to defect. A forgiving player resumes cooperation with an opponent who has previously defected. A non-envious player does not attempt to outscore its opponent. A retaliatory player defects against opponents who defect against it first.

Hofstadter (1986) states that the iterated game is only a true dilemma if the payoffs meet two conditions. The first condition is that the payoff values satisfy the rule $Temptation > Reward > Punishment > Sucker$, which ensures that players may be tempted away from mutual Cooperation. The second condition, $2 \times Reward > Temptation + Sucker$, is only true when the total payoff for cooperative players is greater than that for non-cooperative players.

The threat of future retaliatory Defection allows players of the Iterated Prisoner’s Dilemma to maintain a state of cooperation. Cooperation may also develop between uncooperative players who are capable of learning. As an example, Axelrod (1997) cites the development of cooperation between opposing soldiers in the trenches during World War I. In this situation, Axelrod attributed the soldiers’ wild and ineffectual shooting to a form of cooperation, which was intended to reduce the casualties that soldiers caused to their opponents. When soldiers on both sides reciprocate this behaviour, they all share a lower chance of injury or death. The soldiers could maintain cooperation since both entrenched

armies knew they would face each other in the future, which established trust between both sides. Military leaders on both sides discouraged this behaviour. At first, such cooperation between the soldiers and their enemy seems irrational, as both sides are supposed to fight against their enemy. However, individually soldiers wish to live and by cooperating with a cooperative enemy they may both increase their own chances of survival. Grossman (1995) argues that most people are normally psychologically conditioned to avoid killing, stating that in World Wars I and II only 15%–20% of soldiers would actually fire their guns at the enemy². Grossman claims that the soldiers selfishly avoided killing to avoid causing themselves psychological trauma, whereas Axelrod argues that this observed behaviour is due to cooperation. When soldiers were intentionally firing inaccurately, their actions could be due to a desire to reciprocate inaccurate fire or to avoid psychological harm, either of which may contribute to the observed mutually beneficial outcome.

Noise is one mechanism that can disrupt and prevent cooperation between players in the Iterated Prisoner's Dilemma (Axelrod and Dion, 1988). Here, noise is defined as an error that alters the strategy a player selects. Noise may be caused by the misimplementation of strategies, errors in the communication of strategies or the misperception of performed strategies. Players are typically incapable of differentiating between intentionally selected strategies and those incorrectly selected due to errors. Noise disrupts mutual cooperation by strategies such as Tit for Tat, thereby reducing its performance. Molander (1985) demonstrates that for any significant level of noise, the performance of Tit for Tat players approaches that of two players randomly selecting strategies over a sufficiently long time span. Contrition and forgiveness are two proposed solutions to reduce the impact of noise upon cooperative strategies, such as Tit for Tat, in the Iterated Prisoner's Dilemma (Molander, 1985; Wu and Axelrod, 1995). Contrition alters a player's strategy to prevent retaliation against any defections provoked by its own erroneous Defections. Contrition prevents an unintentional defection caused by noise from echoing between players for multiple turns. A forgiving or generous player responds to some percentage of its opponent's Defections with Cooperation. Wu and Axelrod found that contrition was more effective than forgiveness at maintaining mutual cooperation,

²Engen (2008) states that a review of historical combat reports failed to demonstrate the near-universal reluctance by soldiers to fire upon their enemy that Grossman discussed. Grossman's universal reluctance to kill may instead be a special case, wherein the soldiers feel some compunction against killing an enemy due to some degree of identification with this enemy. While fascinating, further examination of this topic is beyond the scope of this thesis.

since contrition only allows an erroneous Defection to echo once between players, while forgiveness may not immediately correct an erroneous Defection.

The Prisoner's Dilemma game effectively demonstrates a conflict between rational self-interest and collective cooperation, where communication between players is impossible or unreliable. The Iterated Prisoner's Dilemma models situations where repeated interactions provide opportunities for players to establish and maintain long-term cooperation, despite the opportunity for short-term gains from defection. The underlying conflict inherent in the game can both effectively and simply describe many real-world situations, which makes it a highly suitable model for examining various cooperative and competitive scenarios.

2.3.3 Evolutionary Game Theory

Evolutionary game theory applies game theoretical ideas to biology, intending to model the evolution of strategies that players use in biological games (Maynard Smith and Price, 1973). It considers the evolution and propagation of gaming strategies in a population whose behaviour can be abstracted into a game model. Evolutionary game theory focuses on the dynamics of a population of behavioural programs that utilise a particular strategy in their survival game and the distribution of such strategies within a population. A player's strategy is modelled as a trait and can therefore be subjected to evolutionary pressures in the environment. Players of an evolutionary game may play either a pure strategy, where they select one of the actions available to them, or a mixed strategy, where they select from their strategies with some associated probability.

The Evolutionary Stable Strategy is a key concept in evolutionary game theory (Maynard Smith and Price, 1973). A strategy is evolutionarily stable only if it is conserved in a population over a long time span and cannot be replaced by other invading strategies. Evolution affects a strategy in the same manner as any other behaviour or trait. If a new strategy invades a population, its frequency in future generations depends upon its effectiveness compared to the existing strategies. If the invading strategy provides a higher payoff, it is better adapted to the environment and its frequency in the population will increase as it is selected over the existing strategies. If the invading strategy is worse than the existing strategies then the population will not adopt it.

One example of an evolutionary game is the Hawk and Dove game, which was proposed by Maynard Smith (1982) and models a contest between two animals for a resource

(Table 2.2). The animal who obtains this resource increases its fitness by V , while the other receives nothing. The contested resource could be a territory in a favourable habitat or a supply of food. For the contest's loser there is assumed to be adequate space in a less favourable habitat or an alternative supply of food. Whenever two animals meet in a contest, they each choose between two strategies, called Hawk and Dove. Hawk is the more aggressive strategy, with a Hawk player escalating the contest to a physical fight over the resource, until either it is injured or its opponent retreats. A player using the Dove strategy will appear to fight at first, before retreating if its opponent escalates the fight. If both animals choose Hawk and escalate, then the fight continues until one becomes injured and must retreat. Decreasing an animal's fitness by C represents its injury. If both animals play Dove, then both display and share the resource evenly, receiving a payoff of $\frac{1}{2}V$ each. This payoff assumes that displaying has no associated cost and that the resource is divisible.

	Hawk	Dove
Hawk	V or $-C$	$V, 0$
Dove	$0, V$	$\frac{1}{2}V, \frac{1}{2}V$

Table 2.2: Payoffs for the Hawk-Dove Game. When two Hawks contest the resource, the winner receives the resource (V), while the loser is injured and suffers a loss of fitness (C). When two Doves meet, they share the resource.

If both players select Hawk, each player has two potential payoffs. The winner of the fight gains control of the resource and receives the payoff of V , while the loser is injured and receives a payoff of $-C$. Each player has an equal chance to injure its opponent first and win the contest. Therefore the average payoff for Hawk fights is $\frac{1}{2}(V - C)$. It is assumed that the factors governing the player's behaviour are independent of those responsible for its fighting skill.

Players are not restricted to only behaving as pure Hawks or pure Doves; they can also employ mixed strategies where they probabilistically choose between the two actions. In the Hawk-Dove game, this is playing Hawk with probability p and playing Dove with a probability of $1 - p$. For mixed strategies, Maynard Smith (1982) shows that for $V = 2$ and $C = 4$, $p = 0.5$ is an Evolutionarily Stable Strategy. This means that a population where all players use this strategy cannot be successfully invaded by any other strategy.

Hanley, Orbell and Morikawa (2002) argue that the hijackers and passengers on board the aircraft involved in the September 11 attacks played a Hawk-Dove game. In this game,

the hijackers demonstrated that they were playing Hawk and had convinced the passengers to play Dove. Previous hijackings had convinced the passengers that acquiescing to the hijackers and playing Dove was the rational choice. Fighting back against the hijackers by playing Hawk was discouraged by an extremely high value for C , representing death for losing the contest. However, since the hijackers intended to kill the passengers by crashing the aircraft they were not playing the traditional Hawk-Dove game. In this suicidal variant, the passengers face a payoff of $-C$, or certain death, if they play Dove, while Hawk offers either V or $-C$. Hawk is now the dominant strategy, as it offers a chance of living, while Dove guarantees death. In the four hijacked planes, the passengers initially misperceived the situation as a traditional hijacking and rationally played Dove, enabling the hijackers to crash three of the aircraft into their targets. Once news of these attacks reached the last remaining flight, the passengers realised the true nature of the situation and changed their strategy to Hawk by attempting to gain control of the aircraft. Unfortunately the aircraft crashed into a field, likely due to the fight between the passengers and the hijackers (National Commission on Terrorist Attacks Upon the United States, 2004). The high impact of the September 11 attacks and the hijackers' change in strategy is likely to affect the dynamics of future hijackings, with passengers now likely to interpret any hijacking as a suicidal Hawk-Dove type game and act accordingly. A hypergame can model the initial state of the hijacking, with the hijackers playing suicidal Hawk-Dove and the passengers playing normal Hawk-Dove. This hypergame exists until the passengers on the last flight learn that the hijackers are playing suicidal Hawk-Dove, thereby stripping the hijackers of their strategic surprise.

2.3.4 Summary

Game theory provides a mathematical framework for analysing various competitive and cooperative situations and the behaviours of the players engaged in such games. While traditional game theory has difficulty in modelling situations where players suffer from misperception, methods such as hypergames allow situations with various layers of misperception to be modelled and analysed. Evolutionary game theory also provides methods for analysing the distribution of strategies within a population of players and additionally describes how evolutionary mechanisms affect the types of strategies that evolve in a population.

2.4 Misperception

In order to explore the evolutionary value of misperception, it is necessary to define misperception and quantify its effects. This section will explain what is meant by misperception and discuss several potential sources of misperception. Misperception has been noted to affect the decisions made by affected entities, which suggests that it may have some effects upon their abilities to rationally choose behaviours. Therefore, the relationship between misperception and rational decision-making will be discussed.

The Oxford English Dictionary defines misperception as “Wrong or incorrect perception” (Simpson and Weiner, 1989a). From the same source, perception is defined as “The process of becoming aware of physical objects, phenomena, etc., through the senses; an instance of this”, “The mental product or result of perceiving something” and “As a count noun: a direct recognition of something; an intuitive insight; an understanding”. Thus, perception may refer to either the act of perceiving or one or more elements of knowledge produced by the interpretation of perceived information. Perception may describe any sensory input, not just vision. Based upon this definition of perception, the earlier definition of misperception describes both the act of incorrectly perceiving and the production of incorrect knowledge during the processing and interpretation of perceived information. This definition allows for two different types of errors to produce false beliefs: one of perception and one of interpretation. Models of the perceptual cycle show that the tasks of information collection and information processing occur before the creation and manipulation of beliefs (See Section 2.6), allowing errors in these tasks to subsequently produce false beliefs.

This definition indicates some similarity between misperception and deception, since both may give an entity false beliefs. Deception’s false beliefs are intended to induce some behaviour that benefits the deceiver. The misperception errors that may occur are not always unintentional errors of the affected entity. Under some circumstances the entity may desire to adopt false beliefs, as in the case of Self-deception.

In some cases, misperception may occur due to no fault of the entity, but instead due to various types of impairment that affect the communications channel. These types of impairment have several possible causes and may be unavoidable or undetectable by the affected entities. Impairment may block information partially or completely, or alter

its meaning, to produce a false belief when it is perceived by an entity. The actions of other entities may also intentionally or unintentionally cause impairment. Furthermore, false beliefs can affect an entity's future interpretation of information, causing further misperceptions in the future.

Misperception has many potential causes beyond the obvious perceptual error its name implies. In order to effectively study any potential benefits of misperception, it is necessary to analyse these sources of misperception and identify any benefits they may provide. Some existing simulations of misperception demonstrate instances where it may provide a benefit, and these will be examined to determine how this benefit arises.

2.4.1 Perceptual Errors

Errors of perception are the most obvious instance of misperception and occur when information is collected. Entities gather information from their environment with various types of information sensors, which will be utilised in their decision-making process. Potential sensors could include eyes and ears for biological organisms, while electronic systems could include video cameras or microphones. Errors in these sensors will affect the information collected, potentially altering the information or preventing its receipt completely.

Perceptual errors may occur intentionally or unintentionally. An intentional perceptual error is one caused by a competing entity, while an unintentional error is due to some flaw or dysfunction of the individual's sensor. This flaw may be periodic, intermittent or constant. An entity will likely notice the complete loss of information from a sensor, while it may not notice the partial loss of information or the introduction of subtle errors. Damaging an entity's information sensor is classified as an instance of offensive Information Warfare, specifically the canonical strategy of Denial (See Section 2.5.2).

Perceptual Errors affect the flow of environmental information into an entity, reducing the quantity or quality of an entity's collected information. These errors are capable of affecting the entity's behaviour in various manners, which may or may not benefit the entity.

2.4.2 Deception

A deception attack causes misperception in a victim if the victim accepts the deceptive message as a valid message. The deceptive message may induce the victim to hold false

beliefs, which then affect how the victim interprets information in the present and future. Careful planning and implementation greatly increase a deception attack's likelihood of success (Haswell, 1985). If successful, the victim will not immediately realise it has been deceived and it may never discover a particularly successful deception. Deception may be produced by Information Warfare, typically through Corruption attacks (Section 2.5.2). Deception may also be unintentional if an attacker unknowingly communicates incorrect information to the victim, or a deceiver is used as an unwitting proxy for the real attacker.

Deception is widely recognised for its importance in gaining a strategic advantage in competitive situations (Sun Tzu, 1993). While it is possible that a victim could ultimately benefit from its deception, this should only occur if the deceiver intends it; or the deceiver has poorly planned or implemented the deception. While unlikely, the deception could be intended as a beneficial act, which incites the victim to perform beneficial behaviour it would otherwise not perform. Like Perceptual Errors, deception is unlikely to benefit its victims; however, the false beliefs it produces may lead to beneficial behaviour after the initial exploitation by the attacker.

Aged-care workers may become complicit in deceptions that are intended to shield or protect patients or patients' families from harmful information, as Tuckett (2003) observed. The benefits of such deception are similar to those that Ramachandran (1996) (See Section 2.4.4) identifies in Self-deception, with the victim protected from harmful or dissonant information. This well-intentioned deception also somewhat resembles Groupthink (Janis, 1982), with the deceiver as a gatekeeper of the victim's information. Such deception may not be purely selfless, as the deceiver may also benefit from avoiding discussion of certain topics with the victim, indicating a partly selfish motive.

2.4.3 Channel Impairment

Communication channels may be affected by a range of impairments that prevent the effective communication of information, each of which can lead to misperception.

One such impairment is noise, which is commonly described as some form of unwanted signal interference that detrimentally affects communication. However, noise is more precisely defined as undesired signals that were inserted somewhere between transmission and reception (Stallings, 2007), while also being random and possessing specific statistical properties (Held, 1997). Noise affects communication by blocking or distorting messages,

or by changing the meaning of a communicated signal (Weik, 1995). The effects of noise upon electrical and optical communication are well understood (Smillie, 1999), however, such effects are by no means exclusive to these domains. Noise may also be considered from a more general perspective, such as when environmental interference affects communication between biological organisms. Noise can produce effects that are similar to misperception by affecting communications to the degree that either no message can be detected, or that the message received does not match the message originally communicated. A message that has been affected by noise may not be perceived at all or it may be perceived as something that it is not, either of which causes the receiver to possess a false belief. Noise may also mask communicated signals that are less powerful than the noise source, preventing their reception unless the signal is specifically designed for such an environment.

Interference describes anything that intentionally or accidentally alters, modifies or disrupts a signal as it travels through a communications channel from a source to a receiver (Department of Communications, 1986). Typically interference refers to unwanted signals, which are added to a useful signal. Accidental interference can arise from the design and construction of the communications system or from external sources; such as the Sun, which produces electromagnetic interference. Intentional interference occurs when a hostile attacker communicates an interfering signal, intending to prevent a victim from receiving a communication or permanently disabling reception. Radar or radio jamming and other forms of electronic warfare are examples of intentional interference (Skolnik, 1962).

Attenuation describes a loss of energy in the transmitted signal as it travels through a medium (Stallings, 2007). For example, as an electrical signal travels through a wire it expends energy to overcome the resistance of the wire, thereby reducing the strength of the signal. Attenuation may lead to misperception in two different ways. Firstly, attenuation may cause the signal strength to drop below the threshold at which the entity's sensors may detect it, preventing reception of the signal. Secondly, attenuation may cause the signal strength to become near that of the background or environmental noise, at which point the receiving entity cannot accurately differentiate between the signal and the noise, possibly introducing errors.

Dispersion is the physical phenomenon where a wave passes through an inhomogeneous structure and its different frequencies reflect or refract at differing angles (Hecht, 2006). A simple example of optical dispersion occurs when a prism splits white light into a rainbow. Similarly, dispersion may affect light passing through a lens, leading to chromatic aberration where the colours are not focused to the same convergence point. Chromatic aberration produces a distorted image where “fringes” appear along the boundaries separating light and dark areas. Within a fibre optic cable, dispersion can cause the timed pulses of light to become delayed, possibly to the point where pulses overlap with subsequent pulses (Chauvel, 2008). After such dispersion the received signal will not match that sent by the transmitter.

Impairment in digital and analogue communication systems can typically be generalised to other forms of communication. For example, an environment’s ambient aural background noise may affect vocal communication between entities, by masking quiet sounds (low power signals) so that they cannot be detected. Like other errors, impairment can potentially benefit entities if the errors it introduces into the entity’s beliefs later induce behaviour that benefits the entity or others. Of course, it is more likely that these errors will be detrimental. It should be noted that where impairment is a property of the environment, it cannot provide an evolutionary benefit as any trait of the entity would.

2.4.4 Self-deception

Self-deception is another manner by which an entity may introduce false beliefs into its internal model of the environment and thereby cause future misperceptions. While Self-deception suggests that the entity deceives itself, the mechanics of such a behaviour are seemingly paradoxical and require further clarification. The Oxford English Dictionary (Simpson and Weiner, 1989b) defines Self-deception as “The action or fact of deceiving oneself; self-delusion”. This definition suggests that Self-deception is *reflexive*; wherein an entity intentionally uses deception, or some similar process, against itself. Defining Self-deception as merely reflexive deception suggests that an entity may deceive itself, while simultaneously overlooking this deception. While such a definition is seemingly paradoxical, since deception typically requires the deceiver and victim to hold different beliefs, it instead merely represents an inconsistency in the Self-deceiver’s beliefs.

Daniels (1974) criticises the argument for Self-deception as reflexive deception, identifying differences between the deception used in Self-deception and in interpersonal deception. He notes that Self-deception is irrationally accepted and emotionally motivated, such as holding a belief despite contrary evidence and accepting it due to fear or anxiety. Conversely, interpersonal deception is rationally accepted due to supporting evidence, which the victim has no emotional reason to accept.

Attempts to reconcile the Self-deception's perceived paradox instead recast the deception as some form of motivated belief manipulation (Demos, 1960; Fingarette, 1969; Martin, 1986; Mele, 1987). A Self-deceiver therefore intentionally manipulates its beliefs to support or reject a specific belief or beliefs despite contrary information, thereby violating its rationality. However, as with reflexive deception, an entity must fail to notice such manipulation or it would otherwise detect and counteract its Self-deceptive acts. This suggests that an entity requires some unconscious or otherwise concealed system in order to perform Self-deception.

Haight (1980) proposes that since deception requires a distinct separation between deceiver and victim, reflexive deception requires some internal division of the Self-deceiver. A Self-deceiver would therefore consist of multiple distinct entities, who may therefore utilise "interpersonal" deception against each other. Despite Sartre's (1957) protestations to the contrary, such a division of consciousness is well supported. Psychological models proposed by Freud, Jung and others state that all individuals have a conscious and unconscious mind, in which mental actions within the unconscious mind are inaccessible to the individual's conscious mind (Ewen, 2003). Evolutionary psychology proposes that evolutionary mechanisms can explain the development of mental and psychological traits in humans and other species (Tooby and Cosmides, 2005). These traits may be distinct from each other, allowing the mind to be developed in a modular manner, where different traits are attributed to separate modules (Fodor, 1983). Such a modular mind may therefore act Self-deceptively, provided it is motivated to do so.

This clarifies Self-deception as a motivated act by one element of a consciousness that manipulates its beliefs for the perceived benefit of the whole consciousness. While this explains what Self-deception is, it does not justify why it occurs. In humans, Self-deception might be seen as some form of psychological defence. Freud described several psychological defences (Ewen, 2003), where these defences protect against psychological harm caused

by internal or external stimuli. This harm arises due to anxiety, which may be caused by acts or wishes that violate the individual's moral standards; information that signals potential danger in the environment; or information that conflicts with the individual's self-representation. Several of Freud's proposed defence mechanisms are Self-deceptive acts — repression, denial of reality and rationalisation. Through each of these mechanisms, the individual's unconscious mind manipulates its internal representation of the world to produce a preferable representation of reality, and thereby reduce anxiety (Ewen, 2003). An individual uses these acts to exchange psychological relief for a reduced understanding of reality. The anxiety attributed to possessing conflicting beliefs is explained by cognitive dissonance theory (Festinger, 1957). This theory suggests that an individual suffers psychological discomfort from the dissonance produced by possessing conflicting beliefs. Individuals may use various psychological defence mechanisms, which are Self-deceptive, to reduce this dissonance. As such, human Self-deception is seemingly motivated by a desire to reduce the anxiety caused by the possession of conflicting beliefs.

Self-deceptive behaviour can also be observed within organisations. While an organisation cannot Self-deceive, its constituent individuals can Self-deceive and the resultant beliefs can become those of the organisation. While it may seem that an organisation's constituents should identify and prevent Self-deception, there are mechanisms by which organisations can facilitate the collective Self-deception of their members.

Self-deception can infiltrate organisations through their members' desire to maintain internal cohesion. Actions intended to increase internal cohesion in an organisation's actions and beliefs may be Self-deceptive and Janis (1982) calls this **Groupthink**. Groupthink describes how a group of individuals will manipulate their own beliefs and those of other group members, in an attempt to maintain the group's cohesion. The organisation, as a sum of the Self-deceiving individuals, can therefore be considered a Self-deceiver in its own right.

It is worth noting that Self-deception is not the only instance where individuals may unknowingly lie about their abilities. For example, Kruger and Dunning (1999) describe an inability to correctly assess one's own skill at a task, popularly called the "Dunning-Kruger effect", which superficially resembles Self-deception. Kruger and Dunning tested their subjects' logical thinking skills and discovered that those who performed poorly were found to mistakenly believe that they had performed above average. The subjects'

incorrect assessment of their own abilities was attributed not to any Self-deceptive process, but to their lack of knowledge in the relevant area. Without this knowledge the subjects were unable to correctly assess their competence, resulting in the incorrect assessment. Providing the subjects with additional education helped them to improve the accuracy of their self-assessments. Since the subjects in this study are making reasonable self-assessments, albeit from incomplete and incorrect information that was not intentionally degraded or otherwise manipulated, such overconfidence is an example of ignorance and not Self-deception.

Also similar to Self-deception is cognitive conservatism (Edwards, 1968), whereby individuals are reluctant to accept new beliefs that contradict currently held beliefs. However, cognitive conservatism differs from Self-deception in that the beliefs it rejects are considered to be too weak to justify replacing or correcting the established held beliefs, whereas Self-deception rejects beliefs that contradict the individual's preferred beliefs. Confusion between the two mechanisms may still exist when it becomes difficult to determine whether the individual's assessment of weak information is honest or dishonest.

The psychological justifications for Self-deception suggest that it avoids or reduces anxiety caused by possessing beliefs that contradict those desired by the individual or organisation. While common sense suggests that Self-deception is detrimental, its benefits may be worth the cost of inaccurate beliefs. Such a benefit suggests that Self-deception may be an evolved trait.

Trivers (1976; 2000) has asserted that Self-deception is beneficial when it increases the effectiveness of interpersonal deception and that it may have evolved as an aid to deception. When a deceiver communicates a deceptive message to its intended victim, there are signals that can indicate the communicated message's veracity, such as the body language (Allan and Pease, 2006) of the communicating individual. When attempting to deceive another individual, a deceiver may appear noticeably nervous, thereby revealing its deception attempt. Self-deception can allow the deceiver to maintain belief in the deceptive message it wishes to communicate. When a Self-deceiving deceiver attempts to communicate a deceptive message it believes is true, its body language will suggest truthful communication, increasing the likelihood that the victim believes the deceptive information. The deceiver uses Self-deception to mimic an act of unintentional deception, which it may then exploit once its own Self-deception is corrected. Trivers' hypothesis

of Self-deception aiding deception has been demonstrated in a simulation of the Hawk-Dove game that included Self-deception and deception (Byrne and Kurland, 2001). A potential problem with Trivers' theory arises in the deceiver's mechanism for correcting its Self-deceptions. There is the risk that if this mechanism fails, the deceiver will retain the incorrect beliefs and may therefore ultimately suffer the same fate as its victim. While Trivers argues that the incorrect beliefs will be corrected once the individual has deceived its opponent, there is no incentive for an individual to correct an incorrect belief from which it continually benefits. Consider an example of a salesperson who falsely believes that he sells expensive high quality products, despite the cheap wholesale purchase price of these products and the continual return and exchange of broken purchases covered by warranty. While possessing such a Self-deception, the salesman will sell more products and earn a greater commission. Correcting this Self-deception would decrease his earnings and as such, there is no reason or motivation to correct this belief.

Ramachandran (1995; 1996) rejects Trivers' proposal, on the grounds that the act of Self-deception would conceal any knowledge of the deception from the Self-deceiver. Ramachandran argues that the Self-deceiver cannot benefit from a deception attack of which it is unaware. However, the earlier example of the Self-deceiving salesman disproves this argument, as the salesman is unaware that he is deceiving his customers, while benefiting from the increased sales his deceptions provide. It is therefore quite possible that a Self-deceiver may inadvertently benefit from the outcomes of such deceptions. Ramachandran instead argues that individuals and organisations use Self-deception as a defence mechanism, to create a coherent belief structure for themselves that will impose stability on their behaviour and reduce discomfort caused by the possession of conflicting beliefs. Self-deception is therefore argued to be a psychological adaptation that reduces anxiety or distress caused by the possession or awareness of conflicting beliefs. This argument agrees with the psychological explanation that humans have developed Self-deception to reduce cognitive dissonance.

Van Leeuwen (2007) proposes that Self-deception is a "spandrel", which is not an evolved adaptation, but instead an unintended artifact of the individual's cognitive system. A spandrel is a phenotypic feature in an organism that has developed as a side-effect of a true adaptation (Gould and Lewontin, 1978). As a spandrel, individuals have the capacity for Self-deception not because it was selected by evolution, but because Self-deception

is an unavoidable artifact of their cognitive system. However, Van Leeuwen does not attempt to specify which element of the cognitive system might be responsible for Self-deception. The argument for Self-deception as a spandrel conflicts with both Trivers' and Ramachandran's arguments of its potential evolutionary advantage. Furthermore, determining whether Self-deception is a spandrel requires the cognitive system or systems of which it is a spandrel to be identified.

While these theories explaining the evolution of Self-deception do conflict, the benefits they describe are not mutually exclusive. Therefore, Self-deception could conceivably have developed as a spandrel of the human cognitive system, which may also aid an individual's deception of others and sometimes reduce its cognitive dissonance. Further research of the human cognitive system and its implementation of Self-deception is required to determine why Self-deception occurs in humans, although such work is well outside the scope of this thesis.

This discussion of Self-deception has revealed that motivated dishonest behaviour by an individual when faced with contrary beliefs is common among most definitions of Self-deception. Dishonesty occurs as the individual desires to accept a belief as true, despite awareness of its falsity. Self-deception is therefore best described as a variety of behaviours through which the individual manipulates and corrupts its own beliefs in order to support beliefs known to be false, yet desired to be true. Such behaviour requires a divisible conscious, or modular mind, in order to compartmentalise the Self-deceptive acts from other parts of the mind that would quickly thwart any Self-deceptive acts. The false beliefs that Self-deception produces may cause an individual future misperceptions, which may affect it either positively or negatively. Like other sources of misperception, Self-deception may be beneficial if it provides an overall benefit to the individual or to others, similar to altruism. From this point, Self-deception will be assumed to be an intentional manipulation of an individual's beliefs by some internal process, which is concealed from the individual. An individual is therefore unaware of its Self-deceptions, which allows it to potentially benefit from behaviours produced by Self-deception. The assumed intention of Self-deception will be to reduce an individual's cognitive dissonance caused by the possession of contradicting beliefs.

2.4.5 The Simulation of Misperception in Artificial Life Environments

While misperception is a regular occurrence in the real world, it is rarely modelled in Artificial Life simulations. Presumably, this is due to the *a priori* expectation that an individual's misperceptions are always detrimental. Furthermore, any simulation that is not investigating misperception's effects is unlikely to desire its presence. There are some simulations that have demonstrated beneficial instances of misperception, helping support the hypothesis that misperception can provide an evolutionary benefit.

Doran (1998) has identified two instances where misbelief benefits entities. Misbelief is the possession of false beliefs, potentially due to misperception. In both cases the agents' misbeliefs discouraged them from performing detrimental behaviour and this benefit was realised by both the individual agents and their society as a whole. The agents existed in a simulated two-dimensional space, populated with stationary resources that agents could harvest. Parents communicated beliefs to their newborn offspring, allowing offspring to inherit false beliefs.

In the first instance, the simulated environment contained a fatal zone, which immediately killed any agent who entered. Agents were only permitted to move to a resource and harvest it if they believed that they were the closest agent. The agents were permitted to misbelieve in the existence of "pseudo-agents" where none actually existed. These pseudo-agents could deter agents from harvesting resources that were closer to the pseudo-agents. Doran found that many of the agents developed the misbelief of pseudo-agents near resources in the fatal zone, which deterred them from entering the fatal zone. Doran argues that the misbelieving population was fitter, since their society increased its size over the non-misbelievers due to their reluctance to enter the fatal zone. In this simulation misperception benefits any individual agent whose misbelief deters it from entering the fatal zone. These agents may then pass this beneficial misperception to their descendants.

Doran's second experiment investigated the formation of cults among the agents in a similar environment with the fatal zone removed. The agents were now capable of forming friendships with other agents, which enabled them to share information about the location of resources and other agents. Agents were also capable of killing other agents, unless either their target was a friend or the killer and the target shared a common friend. Agents were also now permitted to misbelieve that resources were agents and these were

called “resource-agents”. Agents could both communicate the misbelief of a resource-agent to others and develop friendships with resource-agents. Doran noted that “cults” could form where many agents shared a common misbelief in a resource-agent, who was considered their friend. This friendship prevents the cult-members from killing each other, thereby providing an inclusive fitness benefit to cult-members. While normal friendship groups will disband as the individual agents die, cults do not due to the infinite lifespan of the central resource-agent. In fact resource-agents only really “die” when there are no agents who misbelieve in them. Agents benefit from their misbelief when it produces cults, which collectively benefits all agents with the common misbelief. Cult membership reduces the number of agents who may kill them and thereby permits their society to grow in size.

Doran’s agent cults are somewhat similar to organisations affected by Groupthink (Janis, 1982). The members of the organisation or cult both possess a common outlook on reality that differs to that of external members of their society. Common shared beliefs by the group members, typically of their group’s superiority, are collective narcissism (de Zavala, Cichocka, Eidelson and Jayawickreme, 2009), which enables the group to focus its hostility on any real or perceived external threats.

Doran’s simulations demonstrate that specific misbeliefs held by the members of an evolutionary simulation can provide an evolutionary benefit in specific situations. In both simulations, beneficial misbeliefs increase an agent’s chances of survival and reproduction, allowing it to further spread the beneficial misbeliefs.

Akaishi and Arita hypothesised that misperception could prove to be adaptive if it increases the diversity of a population’s collective beliefs and thereby increases the diversity of the population’s collective behaviour (Akaishi and Arita, 2002a,b). Increased behavioural diversity is hypothesised to reduce direct competition between agents for access to popular resources or locations. An Artificial Life simulation of a two-dimensional grid world populated by agents and resource nodes tested this hypothesis. Agents move through this world gathering resources from resource nodes, all the while maintaining a map of where they believe resource nodes exist. The “fitness” of the population of agents was assessed as the average quantity of resources gathered. Each time an agent observed a location in the environment it could misperceive that location with a constant global probability. Misperception only affected the perception of resources, potentially producing

misperceptions of resource existence or of resource location. An existence misperception causes an agent to see a resource where none existed or no resource where one did exist, while a location misperception causes an agent to correctly perceive a location's contents, but instead store its information at a random location in its resource map. A single misperception may have either a positive or negative effect, possibly leading an agent to a sparsely populated foraging zone or leading it away from known resource nodes into a densely populated area. A misperception may also have no effect, when it alters an agent's beliefs in a manner that produces no change in an agent's behaviour.

Akaishi and Arita explored populations with various global misperception probabilities, discovering that a population with a misperception probability of up to 10% collected more average resources than a population with a 0% misperception probability. Optimal resource collection occurred when the misperception probability was 1%. These results, while counter-intuitive, support the hypothesis that misperception can benefit misperceivers. Akaishi and Arita claimed, however, that this benefit was evolutionary. The simulated system was not evolutionary, as the agents did not reproduce and therefore could not evolve distinct misperception probabilities. While they did identify a "fitness" benefit to misperception, this was strictly defined in terms of resource collection, whereas evolutionary fitness is considered in terms of the quantity and viability of offspring parented. While it is reasonable to assume that increased resource gathering will translate to an increase in offspring parented, the simulation did not prove that misperception led to increased fitness. Akaishi and Arita also attribute this benefit to an increase in the agent population's behavioural diversity. However, they did not determine whether this benefit is exclusive to misperception or whether a similar benefit also arises from other methods that diversify agent behaviour in this environment.

Schermerhorn and Scheutz (2009) investigated the environmental conditions under which communication and memory benefit foraging hive-based agents. In their simulation, a population of agents foraged within an environment and returned collected resources to a single hive. There were two possible resource distributions in their simulated world. One was a random distribution, where the resources were evenly distributed throughout the simulated environment. The other was a clustered distribution, where resources were only located within several clustered areas. The agents were tested with and without memory and communication mechanisms, where memory allowed an agent to recall the

location where it had last gathered food and communication allowed agents at the hive to communicate the known location of food.

Schermerhorn and Scheutz found that in environments with food clusters, communicating agents have an advantage over non-communicating agents if the clusters are sufficiently large. Memory is also beneficial when food is sparsely located — as in the clustered environment — since it allows agents to return to a cluster they have previously encountered. However, in the non-clustered environments with randomly distributed food, communication actually decreases the population’s foraging performance (Scheutz and Schermerhorn, 2008; Schermerhorn and Scheutz, 2009). Communication encourages multiple agents to converge on locations where the first agents to arrive quickly gather all the available resources. As more agents arrive, there are no resources to gather, and they have wasted time and energy travelling to the communicated location. These pointless journeys lower the collective foraging performance of the communicating agent population. This can be considered, in effect, a penalty for foraging at popular locations, which Akaishi and Arita also discussed. In both of these simulations, communication between agents in an environment with randomly distributed resources reinforces the popularity of some resource location beliefs within the population. This reduces the population’s behavioural diversity and thereby reduces its foraging efficiency, as agents are more likely to forage at popular resource locations that are frequently contested.

Akaishi and Arita’s simulation used misperception to deter agents from foraging at popular known locations. While misperception could improve the performance of Schermerhorn and Scheutz’s communicating agents in random environments, a better solution might be to have the communication stop after some time or alter the agents’ reaction to communication. This could be achieved by having agents instead avoid communicated locations, based upon the expectation that any resources at the location will have been collected before the agent can arrive.

The detrimental communication in these simulations essentially encourages agents to compete over popular resource locations. Akaishi and Arita found that misperception-induced randomness somewhat alleviates congestion and resource contention. Randomness can also solve resource contention issues in other domains. For example, contention in computer networking occurs when two Ethernet hosts attempt to transmit a frame concurrently (Peterson and Davie, 2007). Such a collision is detected and resolved by

having both hosts wait a short while before retransmitting their frame. In an attempt to reduce the likelihood of future collisions, both hosts randomly select a wait time and then increase it after any subsequent collision and failed retransmission. Here randomness aims to delay retransmission for a different duration for each host, hopefully allowing for retransmission without a collision. The use of different randomly selected wait times by hosts can be considered analogous to behavioural diversity within a foraging population.

Meyer, Beekman and Dussutour (2008) have also identified noise as the source of an adaptive benefit in foraging ants. Foraging ants selected between two food sources at the ends of a forked path, with one branch shorter than the other. The ants quickly found a path from their nest to the closest food source, via the shorter path. Removing the short path causes the ants to adapt and develop a new path to the more distant food source. However, restoring the shorter path and allowing access to the closer food source does not lead to the ants adjusting their foraging behaviour. However, they found that noisiness in the decision-making of one species of ant, *Pheidole megacephala*, allowed those ants to rediscover the shorter path once it was restored. Without this noise, the ants will not consider the shorter path once it is restored. In this case, noise introduces diversity into the paths taken by members of the ant population, which allows them to rediscover the superior short path. In other situations, noise may not benefit the ants.

These simulations have identified that misperception, or its products, can benefit individuals or populations. Furthermore, communication can either assist or impede misperception, as communication allows a population to consolidate its knowledge and this knowledge itself may be the product of misperception. In the case that communication produces collective behaviour that is detrimental to the population, misperception has been observed to introduce alternative behaviours that can be individually and collectively beneficial. In some cases the benefit from misperception is evolutionary, while in others it aids activities that are likely to benefit evolving organisms. The results of these simulations thereby support the argument that misperception can provide a benefit, albeit in some restricted circumstances.

2.4.6 Rationality and Misperception

The behaviour exhibited by misperceiving entities often appears to be irrational, which suggests that such irrationality is somewhat due to misperception. Rational entities select

actions to maximise their expected payoffs (Russell and Norvig, 2009). This expectation is in turn derived from the entity's understanding of its environment, which is subject to misperception. Irrational behaviours are actions performed by an entity that will provide the entity with a sub-optimal payoff.

Seemingly irrational behaviour can be hypothesised to have three possible causes — irrationality, false beliefs and odd preferences. Irrationality can be inferred if the entity has the requisite beliefs and preferences to determine its optimal action, yet acts in a manner not predicted by these beliefs. False beliefs can be assumed responsible if the action observed can be attributed to one or more beliefs held by the entity that do not match reality. Odd preferences may be inferred in circumstances where the entity's actions show it holds correct beliefs, yet its decisions favour outcomes that it should understand to be sub-optimal. Odd preferences may describe an unconscious predilection for some outcomes. Correctly determining the cause of any seemingly irrational behaviour will be difficult without a good understanding of the observed entity's beliefs. This approach may provide some insight into the erroneous beliefs and reasoning processes of the observed entity; however, it may be unreliable as it depends upon assumptions of the observed entity's beliefs and intentions.

Misperception may alter any aspect of the entity's understanding of its environment, such as the expected consequences of the entity's actions, the actions that the entity may perform, knowledge relating to other competing entities, or methods for analysing information. An entity affected by one or more misperceptions can therefore rationally select an action that then leads to an unexpected and potentially sub-optimal outcome. While the outcome may be sub-optimal for the entity, its decision to perform that action cannot be considered irrational. From the entity's perspective the decision is completely rational and the resulting undesired outcome will be a complete surprise. However, to an outside observer who does not share the same beliefs as the entity, its decisions and actions may appear completely irrational. The inverse of this may also be true, with the entity considering the decisions and actions of an observer, which are based upon different beliefs, to be irrational. The rationality of an entity's behaviour can therefore only be understood in terms of that entity's understanding of itself and its environment. To an outside observer lacking this information, the entity's decisions and actions may appear

irrational. What seems to be irrational behaviour is actually the product of rational decision-making based upon false beliefs.

Misperception does not directly affect the reasoning methods an entity uses to rationally select actions. Instead, it changes the entity's understanding of its environment and the entity then rationally acts upon this incorrect understanding. When an entity makes rational decisions on the basis of incorrect information, its resulting behaviour appears irrational to any observers who lack the same understanding. Entities may also use irrational reasoning methods in conjunction with false beliefs caused by misperception. In these cases, the entity's behaviour is likely to be sub-optimal and appear irrational.

2.4.7 Summary

While misperception is commonly assumed to describe errors of perception, it also encapsulates errors that occur during information analysis. There are many potential sources of misperception, such as perceptual errors, channel impairment, deception and Self-deception, which all occur for different reasons and have different effects on misperceivers. Perception errors are assumed to be due to some fault of the perceiver; however, there are other potential sources of such errors and these need to be categorised and discussed. While I discuss Self-deception's potential benefits, I have mostly neglected its obvious shortcomings. Deliberately corrupting one's own beliefs is likely to have negative consequences when the corrupted information describes elements of the Self-deceiver's environment, such as resources or competitors.

The intentional actions of competitors or friends may potentially cause misperception. Some of these actions can be categorised as Information Warfare attacks, which suggests an implicit link between Information Warfare and misperception. Information Warfare attacks ultimately aim to create one or more specialised false beliefs in the victim, which are intended to alter the victim's behaviour in a manner advantageous to the attacker. Therefore, the relationship between misperception and Information Warfare requires further investigation and is continued in Section 2.7.

Simulations of misperception have demonstrated that entities can benefit from misperception in certain circumstances. This benefit may be achieved through different means, such as discouraging harmful behaviour, encouraging beneficial behaviours or introducing diversity into the behaviour of the affected entities.

While misperception affects an entity's understanding of its environment, it does not impact an entity's rationality. Instead, by altering the entity's beliefs, it changes the framework within which an entity makes rational decisions. In other words, the false beliefs alter the entity's perspective of its environment, and thereby alter the perspective through which it considers the outcomes from its potential actions.

2.5 Information Warfare

Information Warfare encapsulates an array of strategies that allow entities to intentionally create specific misperceptions in their opponents (Schwartau, 1994; Libicki, 1995). All entities, whether they are people, organisations, animals or simulated organisms, are constantly gathering information about their environment and transmitting information through their actions. Entities rely upon information for their survival and are therefore sensitive to events that interfere with the flow of incoming information. Possessing timely and accurate information can convey a competitive advantage over others who lack such information. The survival of entities is also related to the information placed into the environment by their actions. For example, prey who broadcast their presence or location are more likely to be predated upon than those who do not. Therefore, possessing accurate information and the ability to influence how competitors collect information is desirable in any competitive situation, especially if the entity's very survival is at stake. Such an advantage can be obtained by manipulating the information that competitors will receive or preventing competitors from gathering relevant information, while ensuring the integrity of one's own information collection activities. Behaviours such as these represent instances of Information Warfare. While such activities clearly demonstrate an enhanced survivability in biological entities (Kopp and Mills, 2002), Information Warfare's effects on social and technological systems is more commonly studied.

It is widely accepted that the term "Information Warfare" was first used by Thomas Rona in 1976 (Berkowitz and Hahn, 2003) when discussing the advantages of targeting the information and communication systems an opponent depends upon. This general area was later explored by a number of researchers, with the National Defense University establishing a School of Information Warfare Studies in 1992 and the infoWarCon series of conferences launched shortly thereafter. This work spurred further studies of Information

Warfare and led to the development of several similar, yet different, definitions of what actions constitute Information Warfare.

Information Warfare is defined by the United States Air Force as “any action to Deny, Exploit, Corrupt or Destroy the enemy’s information and its functions; protecting ourselves against those actions and exploiting our own military information functions” (Widnall and Fogelman, 1997). Hutchinson and Warren (2001) state that information is both the target of Information Warfare attacks and the weapon utilised to perform such attacks. Denning (1999) defines Information Warfare as consisting of offensive and defensive operations against information resources. These definitions describe information as both a weapon for performing attacks and a target that can be attacked and defended. It is therefore necessary to know what information means in the context of Information Warfare.

2.5.1 What is meant by “Information”?

In the context of Information Warfare, information is described as both a weapon and a target. The definitions and usage of information are somewhat imprecise and vague, which is surprising given the importance of information to Information Warfare. Furthermore, the definitions of information that are used do not always describe easily quantifiable concepts, which complicates measuring the effects of Information Warfare operations.

Information is defined as “Knowledge communicated concerning some particular fact, subject, or event; that of which one is apprised or told; intelligence, news” by the Oxford English Dictionary (1989a). Another definition from the same source states that it is “Separated from, or without the implication of, reference to a person informed: that which inheres in one of two or more alternative sequences, arrangements, etc., that produce different responses in something, and which is capable of being stored in, transferred by, and communicated to inanimate things”. The first definition equates information with news or intelligence regarding a fact, subject or event, while the second describes it as data that can be stored and communicated by machines. Combining these definitions, information provides knowledge of an object, event or phenomenon and can be stored or communicated by people or machines.

The mathematical definition of information comes from Shannon’s (1948) work on communication theory, which proposes an abstract model of communication represented by an information channel between an Information Source and a Destination (Shannon




Figure 2.2: A General Communication System (Shannon, 1948)

and Weaver, 1949). This model (Figure 2.2) generalises communication into five elements – the Information Source, the Transmitter, the Channel, the Receiver and the Destination.

The Information Source selects the Message it desires to send from a set of possible messages. The selected message can take many different forms depending on the communication method utilised. The Message is mapped by the Transmitter into the Signal and then transmitted over the communication channel. The channel is simply the medium that carries the signal and its physical instantiation depends upon the communication method utilised. The channel is non-ideal and as such impairments, which damage the information content carried by the message, may be and most often are introduced during transmission. These impairments may be additive, such as noise or other signals, or may constitute distortion of the signal, which also manifests as damage to the information content of the message. Shannon’s model is primarily concerned with additive white Gaussian noise, due to its prevalence in electronic communications and its mathematical tractability. In Shannon’s model noise is produced by a Noise Source that is connected to the channel. The Receiver collects the signal from the channel and performs the inverse of the Transmitter’s mapping operation, converting the signal back into a Message, which is then passed on to the Destination.

In communication theory, information is a measure of the freedom of choice one has when selecting a message for transmission (Shannon and Weaver, 1949). According to this definition, highly unlikely messages contain much information, while highly predictable

messages contain little information. The Shannon information measure (Shannon entropy), H , of a message comprising of N symbols appearing with probability p_i is (Equation 2.2):

$$H = - \sum_{i=0}^{N-1} p_i \log_2 p_i \quad (2.2)$$

Shannon and Weaver provide a simple demonstration of how the probability of a message's selection affects the amount of information in the transmitted message (Shannon and Weaver, 1949). Suppose there are two possible messages that may be chosen, whose probabilities are p_1 and p_2 . The measure of information, H , is maximised when each message is equally probable; that is, $p_1 = p_2 = \frac{1}{2}$. This occurs when one is equally free to select between the two messages. Should one message become more likely than the other, H will decrease. As a message becomes more probable, the value of H decreases towards zero.

Shannon (1948) also demonstrated that the capacity of a noisy communications channel of a given bandwidth is bounded by the relationship in Equation 2.3, where C is the channel capacity, W is the channel bandwidth, P is the power of the signal and N is the noise power. In the case of electrical or optical communications, the channel capacity may be measured in bits per second, the bandwidth in Hertz and the power and noise measured in Watts. The channel capacity determines how much information a channel can carry and it can be altered by manipulating the channel's bandwidth and signal to noise ratio. Given a fixed channel capacity, information transmission will be maximised if the actual code lengths of the messages are equal to their Shannon information measure, thus, using an "efficient code" (Shannon and Weaver, 1949).

$$C = W \log_2 \left(1 + \frac{P}{N} \right) \quad (2.3)$$

Shannon's definition of information is a purely quantitative measure, determined by the probability of the message's transmission. Therefore, information as defined by Shannon is a property of a communication signal and should not be confused with the semantics (meaning) of the signal. Weaver linked Shannon information to thermodynamic entropy (Sveiby, 1994). Entropy is used in the physical sciences as a measure of the level of organisation and structure (Cammack et al., 2006). A system with high entropy is highly

chaotic or random, while low entropy indicates a well-ordered and predictable system. Following Weaver, Shannon entropy reports how organised the information source is, which determines the rate of information generation (Shannon and Weaver, 1949). Shannon's model relativises information to the prior probability distribution of the receiver. Prior probabilities report the predictability of the message and thus its information content, or "surprise value".

Wiener (1961) defines information in a quantitative manner, similar to Shannon. Wiener describes information from a cybernetic point of view, where it is a property of the signals communicated between the various components of a system, describing the system's state and operations. While Shannon's definition was only applied to communication, Wiener applies his to the control and communication processes in complex mechanical, electronic, biological and organisational systems. Wiener also defines information mathematically in terms of probability, where the probabilities describe the likelihood of a choice between alternatives. Thus, Wiener's definition also relates information and entropy; in particular, Wiener's information measures negative entropy, where the information in a communicated message is a measure of the order, or lack of randomness, in the message. Shannon's measure, on the other hand, was the information source's entropy, describing the uncertainty of the message being transmitted, which can be interpreted as the number of bits needed in an efficiently utilised (noiseless) channel to report the state of the information source. Thus, in Shannon's case, information might be described as a potentiality; a measure of how likely a signal is to occur. In both these definitions information is a property of the communicated signal.

Bateson (1972) instead defines information as a "difference which makes a difference", basing this definition upon Kant's assertion that an object or phenomenon has a potentially infinite number of facts associated with it. Bateson argues that sensory receptors select certain facts from an object or phenomenon, which become information. Bateson suggests that a piece of chalk may be said to have an infinite number of "differences" between itself and the rest of the universe. These differences are mostly useless to an observer; however, a few of these differences are important and convey information to the observer. The filtered subset of important differences for the chalk could include its colour, location, shape and size; but which differences count as information depends upon the perspective of the interested party.

Determining or quantifying the value of an item of information is thus dependent upon the observer, the observer's circumstances, and the time at which the information is acquired. For instance, knowing that stocks will gain in the market before other observers know can yield a higher value than learning this information at the same time as others. Learning such information under circumstances where it cannot be exploited inevitably diminishes its value.

In games of incomplete information, previously discussed in Section 2.3, the value of information in reducing uncertainty can be related directly to the payoff in the game (Morgenstern and von Neumann, 1947; Fraser and Hipel, 1984). If the information results in a high payoff, otherwise denied, the information is of high value. If the game is iterated or comprises multiple turns or steps, the time at which the information is acquired determines the manner in which the value of information changes over time. In this sense, Bateson's representation is a qualitative mapping of what modern game theory tells us indirectly about the context and time variant properties of the value of information. This thesis will not explore the problem of how to quantitatively determine the value of a piece of information, an area well-studied in recent game theory, as that problem is distinct from issues arising from the use of information to gain an advantage in a contest or conflict.

Boisot (1998) provides a different model of information, arguing that what has previously been called "information" can instead be considered three different elements — data, information and knowledge. Entities first observe and make sense of data, converting it to information, which is then understood and incorporated into the entity's knowledge base. Data describes the attributes of objects (while information is a subset of data) produced by the filtering of an entity's perceptual or conceptual processes. These processes are only an abstract description of the actual tasks, as perceptual and cognitive psychological research over recent decades shows that "filtering" or "subsetting" is far too simple a description of what occurs during these tasks. Boisot's definition of information is more psychologically oriented and much broader than the mathematical definitions of Shannon or Wiener.

Definitions of "information" fall into the categories of quantitative or qualitative: the strictly mathematical definitions of Shannon and Wiener versus ordinary language definitions, such as those of the Oxford English Dictionary, Bateson and Boisot. From a mathematical perspective, information is a property of a communicated signal, determined by the probability of that signal. The more likely a signal is, the less information it

has — while the less likely it is, the more surprising its arrival and so the more information it possesses. The informal definitions consider information to be descriptions of some aspects of the world that can be transmitted and manipulated by biological organisms and machines. Of course, quantitative and qualitative definitions are potentially compatible and can be used jointly.

When “information” is used in the context of Information Warfare, it is commonly under its qualitative meaning. For example, when describing Information Warfare against computer systems, information may refer to a computer program, stored data or a message sent between systems. The qualitative definitions of information are, however, vague; leading to conflation with distinct concepts such as knowledge, data and belief. Applying Shannon’s definition of information allows Information Warfare to be studied more rigorously. Whereas the mathematical definitions treat information as a property of a communicated signal, under a qualitative interpretation it is likely to be confused with the semantics of the signal.

In any case, the term “information” as it appears in much of the literature is context sensitive and that context must be interpreted carefully if the meaning of the text is to be read as intended.

2.5.2 Information Warfare: Definitions and Operations

Along with conflicting definitions of information, there are also numerous and often divergent definitions of Information Warfare in current usage, reflecting in part the pervasive nature of the phenomenon and in part the differing perspectives of observers studying the problem. The various definitions of Information Warfare describe actions such as using information as a weapon, targeting information processing infrastructure and protecting one’s own information and information processing infrastructure. In order to clarify what Information Warfare entails, I examine some of the more prominent definitions, along with examples of possible offensive and defensive Information Warfare actions. These examples will reveal the core elements of Information Warfare.

Schwartau (1994) defines Information Warfare in a social context, describing various attacks against information systems and telecommunications networks. Schwartau states that “Information Warfare is an electronic conflict in which information is a strategic asset worthy of conquest or destruction”, a definition covering only offensive actions. The overall goals of Information Warfare attacks are defined as the theft of information, modification of

information, destruction of information and destruction of the information infrastructure, with the ultimate goals of acquiring money and power and generating fear. Schwartau points out that Information Warfare takes advantage of our many modern societies' dependence on information and information systems and is not restricted to governments or government agencies, as is the case with traditional warfare. Schwartau specifies three different classes of Information Warfare attacks, using a taxonomical approach that focuses on the type of target that is attacked.

The first class of operations is **Personal Information Warfare**, in which individuals and their personal details, stored in electronic databases, are the targets. Schwartau describes it as “an attack against an individual’s electronic privacy”, in which the attacker views or manipulates data about the individual stored by various companies and government agencies, which individuals cannot directly protect. Schwartau suggests that attackers could create a false outstanding arrest warrant or supply misinformation to blackmail the individual, although it seems more likely that attackers will steal the individual’s property.

The second class of operations is **Corporate Information Warfare**, in which companies are targeted, typically by their competitors. Schwartau describes industrial espionage, spreading disinformation, leaking confidential information and damaging a company’s information systems as potential examples.

Global Information Warfare is the third class of operations and its victims include industries, political spheres of influence, global economic forces, non-national entities and nations. Typical examples of acts within this category include theft of secrets, denial of technology usage and the destruction of communications infrastructure. Schwartau argues that “it would be stupid for a well-financed and motivated group to *not* attack the technical infrastructure of an adversary”, given the clear vulnerabilities, low risk and large reward of these attacks.

Schwartau’s definition of Information Warfare covers only offensive actions that utilise or affect some sort of electronic information system, implying that Information Warfare is a modern development. Schwartau’s decision to categorise Information Warfare attacks by their intended victim recognises that all victims are not equal and that the motivations for attacking them differ.

Libicki (1995) also defines a taxonomy of Information Warfare attacks, but divides the operations by the environment in which they occur. He gives seven distinct types of operational behaviours that can be categorised as Information Warfare, all of which describe “conflicts that involve the protection, manipulation, degradation and denial of information”. Libicki’s seven types of Information Warfare are: Command and Control Warfare, Information Based Warfare, Electronic Warfare, Psychological Warfare, Hacker Warfare, Economic Information Warfare and Cyberwarfare.

Command and Control Warfare attacks an opponent’s command and communications infrastructure, aiming to degrade its responses to further military action. Command facilities are destroyed to prevent military decision-making, while communications infrastructure is destroyed to prevent the flow of information between decision-makers and the troops implementing those decisions. Libicki points to the effectiveness of Command and Control Warfare by the United States against Iraqi forces as the main reason that the bulk of those forces were ineffectual during the first Gulf War.

Information Based Warfare is the collection and usage of information when planning and implementing military actions. A typical example is using information gained by reconnaissance to assess the effectiveness of previous military attacks or to determine the priority of targets for future strikes — i.e., increasing the situational awareness of the commander.

Electronic Warfare, as described by Libicki, attempts to degrade the physical basis of an opponent’s communications. There are three main targets for such Electronic Warfare attacks, which are radar receivers, communication systems or communicated messages. Anti-radar attacks aim to prevent an opponent’s radar from detecting vehicles, using electronic or physical assaults. Communications systems may be electronically jammed or their physical infrastructure located and destroyed. Cryptography is used to conceal the contents of one’s own communications and to reveal the contents of an opponent’s communications.

Psychological Warfare is defined as the use of information against the human mind, and Libicki divides it into four sub-categories based upon its intended target. Counter-will operations target a country’s national will, aiming to transmit a deceptive message to an entire population. In a military context, messages typically suggest that present and future military attacks are likely to fail. Counter-forces attacks target an opponent’s

military troops, aiming to convince them that fighting is against their best interests. Counter-commander operations intend to confuse and disorientate an opponent's military commanders, detrimentally affecting their decision-making abilities. Cultural conflict targets an opponent's entire culture, attempting to replace their traditions and beliefs with those of the attacker. Libicki states that while cultural conflict has a long history, its implementation is greatly aided by modern technology.

Hacker Warfare consists of attacks against civilian computer networks and systems. Similar attacks against military computer networks are instead considered to be Command and Control Warfare by Libicki. Some aims of Hacker Warfare include the temporary or complete shut-down of computer systems, the introduction of random data errors, the theft of information or services and the injection of false message traffic. Libicki points out that the behaviours he categorises as Hacker Warfare encapsulate much what Schwartau defines as Information Warfare.

Economic Information Warfare is defined as the attempt to control the flow of information between competing nations and societies. An Information Blockade attempts to prevent the real-time transfer of information by methods such as jamming and destruction of equipment, which Libicki argues is difficult to achieve against a determined opponent. Information Imperialism occurs when knowledge-intensive industries become geographically concentrated, which disadvantages those without access to the region. Libicki cites Silicon Valley as an example of Information Imperialism.

Libicki's category of **Cyberwarfare** collects a variety of attacks that are currently unlikely or impossible. However, this term is commonly used by the media to describe acts that Libicki categorises as Hacker Warfare. One of these attacks is information terrorism, a type of computer hacking aimed at exploiting systems to attack individuals, which is similar to Schwartau's Class I Information Warfare. Semantic attacks are another kind of Cyberwarfare, in which computer systems are given seemingly valid information that causes them to produce incorrect output, while appearing to be correct. Another is simulawarfare, an unlikely scenario in which competitors decide to replace conventional warfare with simulated warfare. Gibson-warfare describes an unlikely futuristic conflict between virtual characters inside the system itself. Libicki argues that the current information infrastructure has not developed to the point where these attacks are possible and concedes that it never may in some cases.

Libicki points out that Information Warfare is not a recent development and that some of its varieties, such as Psychological Warfare, have a long history in human conflict. He also notes that new methods of Information Warfare have evolved as technological changes alter the information space. However, while Libicki proposes seven plausibly distinct forms of Information Warfare, some attacks may overlap multiple categories. For example, computer hacking may be considered either Hacker Warfare or Command and Control Warfare, depending on whether the attack targets a civilian or military system. Libicki also somewhat dismisses the effectiveness of performing Information Warfare against non-military targets, such as Hacker Warfare. While such operations do not directly deter military operations, their effects on the civilian population may ultimately reduce political support for military operations and thereby achieve military objectives. Attacks that led to economic losses can also undermine a nation's capability to wage warfare, which may have been the aim of recent Information Warfare attacks against civilian infrastructure in Estonia (Jenik, 2009) and Georgia (Rios, Magalhães, Santos and Jahankhani, 2009).

Libicki (2007) has also explored the use of Information Warfare in Cyberspace, which is defined as any networked computer or communications system. While hostile attacks are the obvious method by which one may conquer cyberspace, Libicki proposes that friendly conquest is also possible. Friendly conquest recognises the power of seduction and develops from mutually beneficial relationships, in which one member becomes dependent upon the information systems or services provided by the other. Friendly conquest differs greatly from other hostile attacks, as it is entered into willingly by the victim in exchange for information or access to information systems that the victim values. The development of social networking sites provide an example of friendly conquest, as users voluntarily exchange personal information in exchange for access to a network that allows them to communicate and socialise with friends.

Widnall and Fogelman (1997) defined Information Warfare for the United States Air Force, describing it in a social context specifically oriented to military operations. Information is said to be produced by perceiving and interpreting phenomena, as in Boisot's definition of information. Information Warfare was defined as "any action to Deny, Exploit, Corrupt or Destroy the enemy's information and its functions; protecting ourselves against those actions and exploiting our own military information functions" (Widnall and

Fogelman, 1997), where acquiring, transmitting, storing or transforming information are information functions. This covers both offensive and defensive Information Warfare.

Widnall and Fogelman detail six types of offensive Information Warfare attacks. Psychological Operations use information to affect the enemy's reasoning and thereby its behaviour. Electronic Warfare denies the enemy accurate information from the environment. Military Deception deceives the enemy as to the attacker's capabilities or intentions. Physical Destruction targets the enemy's information systems for destruction. Security Measures conceal the attacker's military capabilities and intentions from the enemy. Information Attack directly corrupts an opponent's information without visibly changing its physical container. Among these offensive actions only Information Attack and Electronic Warfare are somewhat recent developments, while the others are traditional military operations.

One explicit reason for the United States Air Force's interest in Information Warfare is to enhance its ability to accomplish Air Force missions. Another reason is that the Air Force's dependency on integrated information systems, along with their responsibility for operating the militaries' satellite communication system, makes their information functions a desirable target for attack by opponents. This problem is no longer restricted to the Air Force and other large organisations, as most modern societies have become dependent upon information systems for their daily operations.

Kuehl (1999) provides another military-oriented definition of Information Warfare that considers it in a social context: "Information operations conducted during time of crisis or conflict to achieve or promote specific objectives over a specific adversary or adversaries". Information Operations are said to be "Actions taken to affect adversary information and information systems while defending ones own information and information systems". This implies Information Warfare is a series of offensive and defensive operations that either attack or defend information and information systems, aimed at a specific goal. The requirement that Information Warfare takes place during a crisis or conflict seems to imply that it is the exclusive domain of the military, which disagrees with Schwartz's and Libicki's definitions.

Information Warfare is clearly beneficial when strategically or tactically applied by militaries during war (Widnall and Fogelman, 1997; Kuehl, 1999; Rattray, 2001), however, its usefulness in this role is debatable. Knowledgeable competitors will learn to

expect Information Warfare attacks before and during military operations and attempt to defend against such strikes (Libicki, 2007). Once an entity reveals its Information Warfare capabilities, much of the surprise factor is lost and knowledgeable opponents will increase their defences against similar attacks in the future. During the Millennium Challenge 2002 military exercise, the Red team clearly negated their opponent's Information Warfare attacks (Gladwell, 2005). While the Blue team completely disabled the Red team's radio network, the Red team adopted motorcycle messengers and signalling lights to communicate, suffering only a small reduction in their effectiveness. This example demonstrates Information Warfare as a co-evolutionary race, where attackers locate and exploit flaws in the targets' systems while defenders attempt to correct these flaws as soon as they are observed. This reveals an interesting parallel between Information Warfare attacks and information in that unexpected attacks, like unexpected information, are more valuable than expected attacks or information.

Denning (1999) examines Information Warfare in a social context, this time focusing on information systems and computer security. However, Denning also notes that Information Warfare is not a recent human development, nor restricted to humans for that matter. She defines Information Warfare as offensive and defensive operations that are performed against information resources, which are objects that operate upon information in some manner.

An Information Warfare operation requires at least two players; an offensive player who targets an information resource, and a defensive player who protects the information resource. Players may be individuals or organisations, who may or may not be nation states and may or may not be sponsored by others. Potential offensive players are insiders, hackers, criminals, corporations, governments or terrorists. As every individual and organisation possesses and values information resources, every individual and organisation is a potential Defensive player. Offensive Information Warfare operations aim to increase the value of an information resource to the attacker and decrease its value to the defender. This framework provides a game-theoretical outlook on Information Warfare, where players select offensive and defensive strategies that produce various outcomes with differing payoffs for the players.

Denning states that offensive Information Warfare operations have three overall aims: to increase the availability of the information resource to the attacker; to decrease the

availability of the resource to the defender; and to decrease the integrity of the information resource. These aims closely match Schwartz's stated overall goals for Information Warfare; namely the theft, modification or destruction of information and the destruction of information infrastructure. Stealing or modifying information increases the availability of the information resource to the attacker. Modifying or destroying information decreases the integrity of the information resource. Destroying the information infrastructure decreases the availability of information resources to the defender.

Information resources are protected from Information Warfare attacks by using defensive Information Warfare operations. Denning categorises defensive Information Warfare operations as prevention, deterrence, indication and warning, detection, emergency preparedness and response. Examples include laws and policies that deter various Information Warfare operations, physical security measures that prevent access to information resources, and procedures for dealing with the aftermath of a successful attack.

Denning provides a comprehensive description of both offensive and defensive Information Warfare. Representing Information Warfare in a game-theoretical manner suggests that Information Warfare operations can be analysed with game-theoretical methods. While the examples of Information Warfare focus on computer networks, communication systems and other modern information infrastructure, Denning acknowledges Information Warfare's presence and influence in evolutionary biology.

Borden (1999) combined Widnall and Fogelman's (1997) definition of Information Warfare and Shannon's model of information to define Information Warfare in a social and military context. Borden argues that Widnall and Fogelman's offensive Information Warfare actions — Denial, Exploitation, Corruption and Destruction — are the four main offensive operations of Information Warfare and that any Information Warfare attack may be categorised within one of these strategies.

Degradation replaces destruction and either delays the use of information or damages it partially or completely. Thus, Degradation operates upon the information itself. Examples given by Borden are hiding information from an adversary's collection task and jamming a communications channel, thereby delaying the transmission of messages.

Corruption provides false information for the adversary or corrupts information that the adversary already possesses. Some examples are the use of dummies on the battlefield, spoofing transmissions on the adversary's communications channel and Psychological

Operations performed against the enemy or its allies. The use of dummies and spoofing transmissions attempts to supply corrupt information that an adversary accepts as valid, while Psychological Operations target information already possessed by the target.

Denial is “a direct attack on the means of accomplishment”, meaning anything the adversary uses to collect and process information. This includes destroying or disabling an electro-optic sensor with a High Energy Laser or destroying a computer system that processes information for an adversary with a virus. Denial attacks may either permanently destroy the targeted system or temporarily disable it.

Exploitation is the collection of information directly from the adversary’s own information collection systems. The collected information may help understand the adversary’s point of view.

The overall aims of these strategies match Denning’s three aims of Information Warfare, since Degradation and Denial both reduce the availability of information, while Corruption reduces the integrity of the information and Exploitation increases the availability of information to the attacker.

Independently of Borden, Kopp (2000a) also generates four strategies of offensive Information Warfare. Whereas previous research only described Information Warfare in a human context, Kopp considers Information Warfare in both social and biological systems, arguing that Information Warfare is a basic evolutionary adaptation resulting from competition for survival, which manifests itself in a variety of areas. Kopp derives three types of offensive strategies from Shannon’s model of communication, each of which affects the channel differently, and then adds a fourth strategy that utilises the channel for its communication to the victim. Section 2.5.3 discusses how these attacks affect the communication channel.

Kopp’s **Denial of Information** attack conceals or camouflages information from adversaries, preventing its collection and use. Examples given include insects that blend into their environment, a stealth fighter that uses its shape and radar absorbing material to hide from radar and the use of encryption to hide information from users of a computer system. Denial of Information attacks may be further categorised into either active or passive forms (Kopp, 2006b). Passive attacks are covert and conceal a signal from the victim’s receiver, leaving a victim unlikely to discover the attack. Active attacks blanket

the victim's receiver with noise, preventing it from discerning the signal from the noise. Such attacks are inherently overt, alerting the victim to the attack.

Deception and Mimicry attacks intentionally insert misleading information into a system, which the victim accepts as valid. Kopp's examples include: harmless insects that mimic the appearance of dangerous predators; defensive jamming equipment on an aircraft emitting enemy radar returns with an erroneous position measurement; and techniques used to mask the identity of someone penetrating a network or system. Successful Deception and Mimicry attacks are inherently covert, as they are intended to leave the victim unaware that the information is misleading (Kopp, 2006b).

Disruption and Destruction describes attacks that either disrupt the victim's information system or destroy it outright, in order to prevent or delay the victim from collecting and processing information. Examples given include beetles that spray noxious fluids onto predators to blind them and an electromagnetic pulse weapon that destroys a radar and its supporting communication network. A denial of service attack is an example of disruption in the cyberwar domain. Disruption and Destruction attacks are overt in nature, as the victim will notice the attack's effects on its information receiver (Kopp, 2006b). Existing military terms can further classify these attacks into "hard-kill" attacks that permanently destroy the information sensor and "soft-kill" attacks that temporarily disable the information sensor or system.

Subversion attacks initiate a self-destructive behaviour in the victim's system, caused by information the attacker inserts. An example of this attack in insects is a predatory insect that mimics the appearance of food to lure prey. This deception triggers a self-destructive response from the victim. Deceptive signals that trigger the premature detonation of proximity fuses on guided missiles are an example of subversion in aerial warfare. Logic bombs and viruses are examples of subversive behaviour in cyberwar, where the system uses its own resources to damage itself. Most examples of Subversion employ a Deception and Mimicry attack to first insert the self-destructive signal into the victim (Kopp, 2006b). Subversion can also manipulate how a victim critically assesses information, with an affected victim incorrectly interpreting information in the manner the attacker desires (Kopp, 2006a).

Kopp's categorisation of offensive Information Warfare strategies largely overlaps Borden's. Both describe four canonical offensive Information Warfare strategies, three of

which are, for all intents and purposes, identical. However, these two models converged from very different starting points: Borden's from Widnall and Fogelman's definition and Kopp's from Shannon's information theory.

Kopp's "Denial of Information" attack matches Borden's "Degradation" attack, with both partially or completely concealing information from the victim, thereby reducing the availability of information to the victim. Kopp further categorises these attacks into covert passive attacks and overt active attacks. Borden's analysis also covers the temporary concealment of information, delaying the victim's reception of information, as a method of Degradation.

Borden's "Corruption" strategy and Kopp's "Deception and Mimicry" strategy describe the same behaviour, where a corrupted signal mimics a valid signal and the victim is unable to distinguish between the two. Both attacks aim to reduce the integrity of the information targeted.

The "Denial" and "Disruption and Destruction" strategies also match, with both disabling the victim's apparatus for collecting and processing information temporarily or permanently. Such attacks reduce the availability of information and related processing functions to the victim.

Kopp's "Subversion" strategy lacks any equivalent in Borden's taxonomy, due to Borden's taxonomy following the United States Air Force's convention of folding "Subversion" in the "Denial" strategy (Kopp, 2006b). Subversion attacks aim to decrease the integrity of elements of the victim's decision-making processes, thereby inducing self-destructive actions. On the other hand, Borden's "Exploitation" strategy is not represented in Kopp's taxonomy. Kopp (2003) argues that since Exploitation does not "provide an immediate causal effect in the function of the target", it cannot be an offensive Information Warfare attack. Instead, Exploitation is simply a passive information collection technique.

Borden's and Kopp's models both describe four canonical offensive strategies of Information Warfare that can categorise any offensive Information Warfare attack (Table 2.3). Henceforth, when discussing these strategies the shortest label will often be used to describe the attack.

Denning, Borden and Libicki identified three overall aims for Information Warfare operations, which should be achieved by Borden's and Kopp's four canonical strategies. These aims are to increase the availability of the information resource to the attacker, to

Borden	Kopp
Degradation	Denial of Information
Corruption	Deception and Mimicry
Denial [1]	Disruption and Destruction
Denial [2]	Subversion

Table 2.3: Borden’s and Kopp’s taxonomies of the Information Warfare strategies (Preferred labels in bold)

Effect on Information Resource	Degradation	Corruption	Denial	Subversion*
Increase availability to attacker		✓		✓
Decrease availability to defender	✓		✓	✓
Decrease integrity of resource		✓	✓	✓

Table 2.4: Canonical Information Warfare strategies (Borden and Kopp) and Aims of Information Warfare (Denning).

**Note that Subversion attacks are only capable of meeting these aims if they induce such behaviour in the victim, which is not always the case.*

decrease the availability of the resource to the defender and to decrease the integrity of the information resource. Table 2.4 shows how the four canonical strategies can achieve these three aims. It should be noted that Subversion attacks only achieve an aim when they induce an unintentional behaviour in the defender that happens to achieve that aim.

Libicki’s categories of possible types of Information Warfare can also be compared against the four canonical strategies to identify which attacks each type of Information Warfare employs (Table 2.5). It is worth noting that Information Based Warfare utilises none of the four canonical strategies. Since Information Based Warfare is the same as Exploitation, it is also not an offensive act of Information Warfare. As previously, Subversion attacks only implement one of the types of Information Warfare when Subversion induces the victim to perform an unintended action with whatever effect that type of Information Warfare requires. For example, Subversion will implement Command and Control Warfare if it incites the victim to damage or degrade its own communication systems, such as in a Denial attack.

It is difficult to categorise some attacks into the canonical strategies, specifically when attempting to differentiate between some pairs of attacks — Corruption and Subversion, Subversion and Denial, Degradation and Corruption, and Degradation and Denial. While this problem has been called “Mills’ Paradox” (Kopp, 2010), it is actually a classification problem, where the boundary conditions between attacks must be precisely defined

	Degradation	Corruption	Denial	Subversion*
Command and Control Warfare			✓	✓
Information Based Warfare				
Electronic Warfare	✓		✓	✓
Psychological Warfare	✓	✓		✓
Hacker Warfare		✓	✓	✓
Economic Information Warfare	✓		✓	✓
Cyberwarfare	✓	✓	✓	✓

Table 2.5: Libicki’s categories of Information Warfare and the Canonical Information Warfare strategies that can implement them.

*Note that Subversion attacks are only capable of producing these effects if they induce such behaviour in the victims, which is not always the case.

and rigorously applied. Figure 2.3 presents the classification problems between adjacent strategies and the boundary conditions for differentiating between them.




Figure 2.3: Mills’ Paradox, showing the relationship between the boundary conditions for classifying Information Warfare attacks (Kopp, 2010)

While the opposite attacks in Figure 2.3 can always be easily distinguished from each other, the remaining boundary conditions require careful analysis to determine where one strategy ends and another begins. The criteria for determining the boundary conditions are presented in Table 2.6 and summarised below.

Distinguishing a Subversion attack from a mimicking Corruption attack requires an understanding of how the victim processes the deceptive input message. While both attacks use a deceptive message against the target system, the distinction lies in whether the victim's response is voluntary or involuntary. Corruption affects the victim's perception of external reality, which it then responds to with a voluntary action of some kind. Subversion instead manipulates the victim in some manner, which ultimately induces the victim to perform an involuntary harmful behaviour. Biological cases often prove difficult to separate.

Distinguishing a destructive Subversion attack from a Denial attack can also be problematic, since both superficially produce a hard kill of the victim's information receiver. These two attacks are differentiated by determining whether the attacker or the victim expended energy to destroy the receiver. During a destructive Denial attack, the attacker destroys the victim's receiver via external means, whereas Subversion instead causes the victim to trigger the destructive effect against itself. The victim's role in both attacks is involuntary.

Distinguishing a passive Degradation attack from a mimicking Corruption attack can frequently present difficulties, especially in biological systems. The boundary condition is based upon whether the victim misidentifies the attacker or fails to perceive it at all. Mimicry that is designed to camouflage an attacker against the background noise is a passive degradation attack, since the victim cannot perceive the attacker, whereas Corruption leads the victim to perceive and misidentify the attacker.

Distinguishing an active Degradation attack from a soft kill Denial attack may also be superficially difficult. In both instances the channel has been rendered unusable by an observable attack on the receiver. The boundary condition can be established by determining whether the receiver remains functional or not. An overloading of the receiver to deny its use is quite distinct from a channel that is unusable due to saturation with a jamming signal.

The clarity in differentiating between Degradation and Subversion attacks or Corruption and Denial attacks arises due to the conflicting operation of each of the paired attacks. Degradation requires a functional victim system to achieve its effect, but Subversion results in the destruction or serious functional impairment of the victim. The same

	Degradation	Corruption	Denial	Subversion
Degradation	–	Is effect perceived?	Effect on channel or receiver?	–
Corruption	Is effect perceived?	–	–	Voluntary or involuntary effect?
Denial	Effect on channel or receiver?	–	–	Attacker or victim supplied effect?
Subversion	–	Voluntary or involuntary effect?	Attacker or victim supplied effect?	–

Table 2.6: The boundary conditions for differentiating between the canonical Information Warfare strategies (Kopp, 2010)

dichotomy exists between Corruption and Denial, as a deception cannot be effected if the victim system loses its channel or receiver.

Information Warfare attacks may be combined into a compound Information Warfare strategy (Kopp, 2005b), where the attacks form a partially-ordered directed graph whose precedence is dictated by the graph’s structure. Each individual attack may have multiple predecessor and successor strategies (Figure 2.4), which are designed to shift the victim towards an intended final state. Overall success of a compound strategy depends upon whether the victim’s end state matches that intended by the attacker. Kopp also defines the concept of a chained compound strategy, where an intermediate victim is used to propagate an attack against the final victim. For example, terrorist movements exploit media organisations to spread news of successful terrorist attacks in a chained compound strategy.

Since compound Information Warfare strategies form directed graphs, they can be analysed with graph theory. Graphs can be partitioned into two or more disconnected graphs by cut vertices (Chartrand, 1977). Cut vertex removal corresponds to the failure of an Information Warfare attack, which happens to be a single point of failure within the compound strategy. Attackers may increase the redundancy of their compound strategy to remove any cut vertices, by adding additional functionally identical attacks, analogously to fault-tolerant computer system design. Another advantage of redundancy is that parallel attacks can occur simultaneously, potentially distracting a victim with multiple attacks at




Figure 2.4: Compound and Chained Information Warfare strategies (Kopp, 2005b)

the same time. Defenders who are aware of a compound strategy's structure may attempt to identify cut vertices and focus their defence against those Information Warfare attacks.

In the pursuit of a compound attack, the state of the victim is often crucial. If a predecessor strategy has failed to produce its intended effect, successor attacks in the compound strategy may be ineffective or even counter-productive by betraying the predecessor strategy. In a successful compound attack, the victim's internal state proceeds through a series of intermediate steps, reflecting the successful effects of each node in the compound strategy. This exposes a historically well documented problem in the execution of complex deceptions, as determining the effect of a previously executed strategy may be difficult or impossible.

There are many common elements among the various definitions of Information Warfare. Information Warfare consists of offensive and defensive elements, in which one's own information and information collection and processing functions are protected, while those of competitors are attacked. Information Operations may describe individual acts, while Information Warfare may describe an overall campaign of Information Operations against one or more competitors. Alternatively each instance of Information Warfare may be called an Information Warfare attack.

Information Warfare attacks may be performed by a single entity or a group of entities. From the definitions above, the following users of Information Warfare can be identified: citizens, governments, companies, criminals, countries, non-nation states, political groups and business cartels. Any of these may target their attacks within or across these groups. Criminals may perform Information Warfare attacks against companies, while political groups could attack and defend against attacks from other political groups. Information Warfare capabilities have also influenced the design of machines and systems, with many military examples including stealth aircraft and visual camouflage schemes. Many animal and plant species also utilise Information Warfare to aid their survival and reproduction.

Modern communication networks and computer systems have created a new environment for Information Warfare attacks. Consequently, Information Warfare has often incorrectly been described as a modern development. However, elements of Information Warfare are coextensive with human military conflicts, including chronicled instances of strategic military deception and psychological warfare. This identifies a classical usage of Information Warfare, recently extended to exploit new technology. We can anticipate that this will continue so long as technological advances continue. In biology, there is a clear parallel to co-adaptive arms races between predator and prey species. Narrow definitions of Information Warfare, restricting it to attacks involving computer systems or telecommunications networks or even simply to human activity, miss this broader significance of Information Warfare. A more appropriate concept allows for Information Warfare in any competitive environment where information processing takes place. Identifiable instances of Information Warfare both in human history and in nature support the broad definition (Kopp, 2005a).

The canonical Information Warfare strategies provide a framework for categorising Information Warfare attacks and for identifying the functional similarities between what may appear to be quite different attacks. For example, a camouflaged insect and a camouflaged military vehicle are both utilising the same canonical strategy — Degradation — against potential observers to achieve the same goal — concealment.

2.5.3 Shannon's Communication Theory and Information Warfare

Borden (1999) and Kopp (2000a) both assert that Information Warfare, being based upon the concept of information, should be analysed in terms of Shannon's information theory.

Borden argues that an action a decision-maker performs on data will reduce its uncertainty. Measuring the decision-maker's uncertainty before and after this action will reveal the change in uncertainty, which may be measured in bits.

An example is the case of Paul Revere (Borden, 1999). Revere was awaiting information regarding the movement of British troops, who would move either by land or by sea. This information was then to be forwarded by Revere to the American Revolutionaries. A lookout in a nearby church would observe the British troops and report the method of their approach by showing "one lantern if by land, two if by sea". For Revere, both approaches were equally probable ($p(\textit{land}) = p(\textit{sea}) = \frac{1}{2}$), so he had one bit of uncertainty (Equation 2.4). Revere observed that two lanterns had been lit, which informed him that the British were coming by sea, thereby reducing his uncertainty. Revere's uncertainty after receiving the message is recalculated with $p(\textit{land}) = 0, p(\textit{sea}) = 1$ (with the usual assumption that $0 \log_2 0 = \lim_{x \rightarrow \infty} \frac{1}{x} \log_2 \frac{1}{x} = 0$). As Equation 2.5 shows, Revere's uncertainty was reduced to 0 and so the message reduced 1 bit of uncertainty.

$$\begin{aligned}
 H(X) &= -(p(\textit{land}) \log_2 p(\textit{land}) + p(\textit{sea}) \log_2 p(\textit{sea})) & (2.4) \\
 &= -\left(\frac{1}{2} \log_2 \frac{1}{2} + \frac{1}{2} \log_2 \frac{1}{2}\right) \\
 &= -\left(-\frac{1}{2} + -\frac{1}{2}\right) \\
 &= 1
 \end{aligned}$$

$$\begin{aligned}
 H(X) &= -(p(\textit{land}) \log_2 p(\textit{land}) + p(\textit{sea}) \log_2 p(\textit{sea})) & (2.5) \\
 &= -(0 \log_2 0 + 1 \log_2 1) \\
 &= -(0 + 0) \\
 &= 0
 \end{aligned}$$

Equation 2.4 also demonstrates that when one is equally free to choose between different messages, the amount of information in the transmitted message is maximised. On the other hand, Equation 2.5 shows that when the probability of a message is certain, the message contains no information. As the probability of a message being selected increases, the amount of information it contains decreases. It is clear that in Shannon's definition

of information, the key aspect of a message is its improbability. So, one feature of Information Warfare is the attacker's attempt to reduce its own uncertainty or to increase the uncertainty of its target.

Information Warfare may also be considered in terms of its effects on the capacity of an information channel. Kopp states that "Information Warfare in the most fundamental sense amounts to manipulation of a channel carrying information in order to achieve a specific aim in relation to an opponent engaged in a survival conflict" (Kopp, 2003). Borden (2001) has described Information Warfare as a "battle for bandwidth", in which opponents compete over an information channel's capacity. Therefore, the canonical Information Warfare strategies can be explained in terms of their effects on an information channel (Kopp, 2000a), by examining the effects of the strategies on the terms of Shannon's channel capacity formula (Formula 2.3). An information channel's capacity (C) can be reduced by decreasing its bandwidth (W), decreasing the power of the signal (P) or increasing the noise in the channel (N).

This model makes two assumptions. The first is that the victim receiver can wholly understand and thus decode the messages it receives, which may or may not be true in general. The second is that some repeatable mapping exists between a message, background noise and the quantitative measures of P and N . A basis for establishing such a mapping lies in Shannon entropy, which shows that a message with an entirely predictable content contains no information (Shannon, 1948):

$$I(m) = -\log_2 p(m) \quad (2.6)$$

where $I(m)$ is the information content of the message and $p(m)$ is its probability. As $p(m) \rightarrow 1$, $I(m) \rightarrow 0$. If we define noise in this channel as messages without useful content from the receiver's perspective, this provides a basis for our mapping.

A Degradation attack may render the signal sufficiently noisy that the receiver cannot discern the signal from the background noise in the channel. An active Degradation attack transmits additional noise into the information channel, so that the signal is harder to detect. Injecting much noise into the channel will make $N \gg P$ and thereby force $C \rightarrow 0$. A passive Degradation attack reduces the power of the signal ($P \rightarrow 0$) so that it is too faint to be detected, which also forces $C \rightarrow 0$.

A Corruption attack substitutes a valid signal for a corrupted signal. In terms of Shannon’s formula, the attack replaces $P_{Genuine}$ with $P_{Corrupt}$, while W and N remain unchanged.

Denial attacks disable or destroy transmission links or information receivers, denying the victim the ability to receive information. These attacks reduce the available bandwidth of the channel (W) and in an effective attack $W \rightarrow 0$ or $W = 0$.

Subversion would not normally be described in terms of channel capacity as it does not affect the signal transmitted, the contents of the signal, the information channel or the receiver. Instead Subversion is likened to an attack against the decision-making process, which can be modelled by a Turing machine (Kopp, 2003) and decision processes within victims can themselves be construed as information-bearing channels. The program that controls the operation of a Turing machine is a tape that contains a series of symbols. A Subversion attack against such a machine is performed by covertly altering the symbols on the tape, thereby altering the behaviour of the Turing machine in some manner. This can be considered analogous to Corruption attacks, by replacing the target’s original probability distribution with one that misrepresents the situation.

There are also other deceptive techniques that resemble deception, yet are not Corruption attacks, which Kopp (2006a) has labelled “Deception by Omission”, “Deception by Saturation” and “Deception by Spin”. These three techniques are frequently found in commercial and political product marketing, as well as in many intelligence deceptions. Detailing how these techniques affect an information channel shows how they differ from Corruption attacks.

Deception by Omission occurs when an attacker presents a message, or multiple messages, which appear to be complete, but are not. This technique is a Passive Degradation attack, where $P \rightarrow 0$ for the hidden information, thus reducing its contribution to channel capacity to zero.

Deception by Saturation, which is also known as “flooding” (Kopp, 2006a; Libicki, 2007), arises in two forms; either as an Active Degradation attack, or a soft kill Denial attack. During a saturation attack, the attacker inundates the victim with messages, most of which are redundant or irrelevant, with the aim of saturating the victim’s channel so the victim cannot gather information that might contradict the attacker’s message. Even an alert victim who may have the capacity to find valid messages embedded in a large

volume of redundant messages may be effectively attacked, if the victim does not have the available time to sort through all of the received messages.

The messages sent in a saturation attack may be considered as noise in the channel. Where the victim cannot successfully filter a message from the background noise, for whatever reason, the capacity of the channel will degrade. In terms of Shannon's model for channel capacity, the redundant or information free messages represent noise N , with $N \gg P$ resulting in $C \rightarrow 0$.

In an alternative form of this attack, the victim is capable of distinguishing between real and information-free messages, yet cannot do so in a timely manner. In terms of Shannon's model, this describes a scenario where the bandwidth of the channel is inadequate for the problem, that is $W \ll W_{Required}$. As a result, the channel lacks the capacity to carry the real message, allowing the attack to succeed. Attacks that render the channel unusable by compromising its available bandwidth are classified as soft kill Denial attacks.

Deception by Spin is a form of Subversion attack that is often used in a compound strategy supported by Deception by Omission, or sometimes Deception by Saturation. During a spin attack, an attacker encourages the the victim to assess a fact — possibly unwanted, acknowledged or accepted by the victim — from a perspective that is less damaging to the attacker. This thereby subverts the victim's mechanism for critically assessing the unwanted fact. In information theoretical terms, Deception by Spin is a classical compound Subversion attack, where the victim uses its own internal processing resources to infer false conclusions from the received message.

Analysing Information Warfare as an attack against an information channel provides three distinct targets for Information Warfare attacks, each of which is uniquely associated with one canonical Information Warfare strategy. These attacks either reduce the capacity of the channel or target messages that are inside the channel. Kopp (2006b) argues that since there is one attack that affects each term of the equation and one that affects the underlying probability distribution (Subversion), these four attacks exhaust the possible canonical Information Warfare strategies.

2.5.4 Domains of Information Warfare

The earlier concepts of Information Warfare focus on its usage in the contemporary social or military domain, specifically its targeting of modern information systems, such as computers and communications networks. Information Warfare, however, is neither a modern invention nor unique to humans. The use of information in competitive survival contests is quite general.

The Information Warfare practices of modern militaries have a historical basis (Widnall and Fogelman, 1997). Psychological operations were widely used by the Mongols during their invasions to spread fear about their approaching armies, encouraging “merchants” or survivors to spread word of their attack and strength ahead of their invasion force (Chambers, 1988). Military deception similarly has been used widely throughout history by military leaders (Haswell, 1985; Bose, 2003). Julius Caesar’s rapid forced marches deceived enemies into thinking no imminent military action was possible (Caesar and Handford, 1951). Sun Tzu advocated concealing one’s army — “have a capability, but appear not to” (Sun Tzu, 1993) — as a security measure for military warfare. Cases of such concealment leading to battle success are rampant throughout the history of warfare, disproving the modernity of Information Warfare strategies (Conley, 1988).

Indeed, Sun Tzu states that “All warfare is based on deception” (Sun Tzu, 1963, 1981). Since deception is a key element of the canonical strategies of Information Warfare (especially the Corruption strategy), Sun Tzu’s writings provide a historical basis for the theory of Information Warfare and stress the importance of gathering and protecting information in warfare. Sun Tzu’s statements on Information Warfare include:

Protecting Information: “In making tactical dispositions, the highest pitch you can attain is to conceal them; conceal your dispositions, and you will be safe from the prying of the subtlest spies, from the machinations of the wisest brains.”

Gathering Information: “Thus, what enables the wise sovereign and the good general to strike and conquer, and achieve things beyond the reach of ordinary men is foreknowledge.”

Deception: (1) “Hence, when able to attack, we must seem unable; when using our forces, we must seem inactive; when we are near, we must make the

enemy believe we are far away; when far away, we must make him believe we are near.” (2) “Hold out baits to entice the enemy. Feign disorder and crush him.”

These statements demonstrate that Sun Tzu understood the importance of information and the advantages provided by offensive and defensive Information Warfare.

Kuehl (1999) identifies three distinct target types for Information Operations: hardware, software and “wetware”. Hardware describes the physical devices that form the information collection, communication and information processing systems. Software describes the coded instructions that control the operation of the hardware devices. Wetware refers to the human mind, and Kuehl uses this term to emphasise that the human cortex is as important as software and hardware in modern information systems. Kuehl also declares that Information Operations against wetware have a long history and notes that Sun Tzu’s teachings were infused with the idea that the enemy’s mind was the target that possessed the greatest payoff.

While the majority of descriptions and definitions of Information Warfare focus on its usage by humans and its effects on modern communications, Kopp (2000a) observes that “The fundamental paradigm of IW/IO [Information Warfare/Information Operations] appears to be a basic evolutionary adaptation resulting from competition in the survival game”. This argument is supported by examples of Information Warfare attacks performed by insects. Denning (1999) also acknowledges that biological organisms can perform Information Warfare. Further biological examples of Information Warfare attacks are provided by Kopp and Mills (2002), who describe the Information Warfare strategies employed by numerous species of insects, fish and birds. These examples, along with many others, demonstrate that Information Warfare is a fundamental survival mechanism, which many different animal species have separately evolved over millions of years. As such, Information Warfare is neither a modern artifact nor a uniquely human endeavour, but instead a mechanism that may aid competition in *any* environment.

Information Warfare has likely been misidentified as a recent development due to the recent creation of worldwide telecommunications networks and computer systems, which have provided a new operating environment for Information Warfare. This modern environment contains new targets to attack and new methods to do so. Schwartau (1994) and

Denning (1999) both give examples of these types of attacks. Regardless, the underlying strategies of modern, historical and primeval Information Warfare attacks are all the same.

The growth of telecommunication networks has increased their value to users, in line with Metcalfe's Law (Metcalfe, 1995). Metcalfe's Law, broadly interpreted, states that the value of a telecommunication network increases as the square of the number of devices connected, where the network is used for service delivery or distribution. Therefore, modern telecommunication networks, which allow rapid worldwide communication between individuals and organisations, have become increasingly valuable to their users. Such networks and systems also provide an environment for Information Warfare activities. Each connected device, or its user, is a potential target for Information Warfare attacks performed through that network, along with the network itself. As networks become larger and more valuable, they become much higher value targets for Information Warfare attacks against the network's users, as more targets are affected by the disruption or destruction of the larger network. Larger networks also provide attackers with access to more potential targets.

Information Warfare may also occur in competitive non-military social environments, including politics or product marketing (Kopp, 2006a). In these cases, deception and related forms of Information Warfare are used to promote a group, an idea, or a product to various people, typically among members of the general public.

While nations at war are not restricted in their use of Information Warfare, the legality of Information Warfare attacks between nations at peace is unclear (Komov, Korotkov and Dylevski, 2007), as there are currently no international laws that either ban or regulate how Information Operations may be used or what types of retaliation, if any, are permitted under the United Nations charter. Komov, Korotkov and Dylevski have analysed the current body of international law and concluded that "almost any information operation with a psychological bias, implemented in peacetime with respect to another state, would qualify as intervention in its domestic affairs. Even good intentions, such as the advancement of democracy, cannot justify such operations.". While the current interpretation of international law does not support this argument, Komov, Korotkov and Dylevski believe that it would be advisable for international law to consider offensive acts of Information Warfare between nations as "aggression", which is prohibited by the UN charter. Korotkov (2008) has argued that such aggression occurs whenever a government

promotes ideas on the Internet with the goal of subverting another country's government. This concept is further enshrined in an agreement between member states of the Shanghai Cooperation Organisation (Shanghai Cooperation Organization, 2009), where the "Dissemination of information harmful to social and political, social and economic systems, as well as spiritual, moral and cultural spheres of other states" is considered to be a threat to ensuring international information security, alongside other acts of Information Warfare performed by states, criminals or terrorists.

This legal interpretation is problematic, as it considers attacks that are at worst propaganda to be equivalent to large-scale Information Warfare attacks that destroy or cripple a target nation's military or civilian infrastructure. While the latter should obviously be prohibited, the former may only be discussion and criticism of another nation's actions and only warrants the label of Information Warfare; and possibly a retaliatory response, if it is deceptive. The lack of consensus on a legal definition of Information Warfare is partly due to the differing legal perspectives of the various international observers; what some nations label Information Warfare, others consider freely permitted speech.

In short, Information Warfare is useful across a wide variety of domains, including indeed any domain that offers some competitive advantage to one actor over another, whether in biology, warfare, sports, politics or marketing. Its potential presence should be expected and planned for in all such situations.

2.5.5 Applications of Information Warfare Theory

The study of theoretical concepts of Information Warfare and their application to existing systems has many potential benefits. Information Warfare may occur in any situation where there is communication between two competing entities and may therefore be applied in a wide array of potential areas, some of which are discussed below. An important factor is that the increasing complexity and integration of man-machine systems, typified by networked computing systems, present increasing vulnerabilities to attack. Therefore, a general model for understanding the vulnerabilities of such systems is needed, especially for designers. Information-theoretic models of Information Warfare, such as the canonical strategies, provide such general models.

Islam, Pose and Kopp (2005) consider the security concerns of various wireless ad-hoc networking protocols from an Information Warfare perspective. Potential attacks against

a wireless network are described in terms of canonical and compound Information Warfare strategies, along with the effectiveness of the wireless network protocols in defending against such attacks. This approach allowed Islam, Pose and Kopp to identify a potential vulnerability to Subversion attacks in their proposed networking protocol, providing a focus for future security extensions to the protocol.

Kopp (2005a) has described how government and non-government organisations perform Information Warfare to implement “perception management” campaigns against a victim population. Nazi Germany and the Soviet Union are provided as examples of regimes that thoroughly adhered to Haswell’s (1985) principles of deception in such campaigns. Typically, perception management is performed by authoritarian regimes, where the regime controls the media apparatus and is intent on deceiving its population for its own benefit. This effectively creates a hypergame (Bennett, 1980; Fraser and Hipel, 1984) between the regime and its victim population. Perception management may also be used by regimes or movements to perform deception and propaganda campaigns against the populations of other nations. In that case, foreign media organisations are the initial target of a deception, which compels them to distribute further deceptive messages to the victim population; an example of chained Information Warfare, employing compound strategies of Subversion, Degradation and Corruption. Denial attacks are avoided, since these damage the delivery channel and thereby prevent its future reuse.

The use of the mass media as a conduit for Information Warfare presents a difficult defensive problem. There are three distinct groups who may defend against these attacks — the mass media, the victim populations and the governments of the victim populations. The mass media benefits from their role as a conduit and therefore have no interest in preventing the propagation of the Information Warfare attack against the victim population. Governments of the victim populations, especially democracies, may be powerless to stop such an attack, as legislation often prevents direct control of the mass media, while in any case modern technology simplifies the circumvention of government censorship. The victim population may be the only entity in a position to identify and resist such attacks. However, it is unlikely that all members of the population will be capable of defending themselves against these attacks, allowing the campaign to deceive much of the population.

Political and product marketing is another area in which deception is commonly observed (Kopp, 2006a), where policies, products or services may be marketed despite limitations or deficiencies. Kopp identifies three pseudo-deceptions that are used in this area and describes methods of defending against such attacks. I consider these methods to be **pseudo-deceptions** because they do not use the Corruption strategy typically associated with deception. This allows their use when regulation or legislation prevents Corruption's untruthful communication, to achieve results similar to Corruption.

Broader and more fundamental issues arise when we consider the impact of the use of Information Warfare techniques on various paradigms of conflict, especially in conflicts involving nation states and non-state actors. Deception and propaganda has been a central part of such conflicts for as long as they have existed. Modern communication technologies permit large amounts of data, and so also information, to be communicated and disseminated very rapidly. Most established paradigms of conflict have evolved in environments where the underlying technology base provided no such capability. As a result they exhibit varying levels of sensitivity to the introduction of systems that can transmit or disseminate data and information on a large scale. Numerous cases studies can be found in the impact of networking technologies upon contemporary military systems and organisations (Kopp, 2009).

Kopp (2000b) argues that “a fundamentally different adaptation is required in order to survive and prevail in such an environment. This adaptation is the ability to evolve faster in technology and operational doctrine over potential opponents. Indeed it is worth stating this as an axiom: ‘The player who can evolve technology and doctrine faster than an opponent, all other things being equal, will prevail.’ ”

Investigations of the ideas and applications of Information Warfare may allow potential victims to identify Information Warfare attacks and better defend against them. Knowledge of the structure and properties of an Information Warfare attack also provides knowledge of potential defences against such an attack. The theory of Information Warfare can also be used to analyse existing physical and non-physical security systems and provide insight into their vulnerabilities.

2.5.6 Summary

This discussion has surveyed existing research covering information theoretical models for Information Warfare, and performed a critical analysis of the definitions, models and canonical strategies for Information Warfare.

In the context of Information Warfare, the term “information” commonly refers to the natural language concept of information. However, Information Warfare may also be considered from the point of view of Shannon’s information and communication theory, providing a formal background for understanding how Information Warfare functions.

All definitions of Information Warfare describe it as a combination of offensive actions, performed against an opponent’s information and information processing capabilities, and defensive actions to protect oneself from such attacks. Successful attacks are intended to affect the victim’s decisions and actions. The many potential offensive Information Warfare attacks can be divided into four canonical strategies, which may be applied in any competitive domain.

While much of the Information Warfare literature focuses on its application to computer systems and telecommunications networks, there are other situations where Information Warfare arises. These span technological, biological and social systems. This is evidence that Information Warfare is not a recent human invention, but rather an evolved, general capability that provides a competitive survival advantage in situations where information is of benefit to competing entities. Organisms use Information Warfare to cause their same-species competitors, predators or prey to misperceive, likely inducing actions that benefit the attacker. As such, misperceptions caused by Information Warfare are unlikely to provide an evolutionary benefit to the misperceiver, as this is not their intention.

All of the four canonical Information Warfare strategies are capable of causing a victim to misperceive. Degradation aims to manipulate the information so that it cannot be perceived accurately by the victim, thereby causing a misperception of the concealed information. Corruption inserts false beliefs into the victim, which may cause the victim to misperceive in the future. Denial damages or disables the victim’s information sensor so that it is unable to gather information and may thereby cause misperception. Subversion affects the victim’s behaviour directly and can cause misperception if the Subversion-induced behaviour damages the victim’s information collection or processing capabilities.

2.6 Perceptual Cycle Models

The effects of misperception and false beliefs upon decision-makers can be better understood by examining models of the decision-making process. There are many conceptual models that attempt to describe the information collection and processing behaviours of entities, which may also be labelled as their perceptual cycle. Such models represent the decision-making process in a simple manner, against which the effects of misperception can be studied.

These perceptual cycle models also demonstrate how an entity creates and manipulates its beliefs, and how they instruct present and future decisions. There have been several different models proposed for modelling perception, each created in different domains and for different reasons. These models place the processes of perception and decision-making into an iterated cycle that the entity proceeds through. While each model may only be intended for use in a limited domain, they are often sufficiently general to cover other domains.

Neisser (1976) proposes the Perceptual Cycle (Figure 2.5) as a framework for the psychology of perception. This model consists of three elements — Exploration, the Object and the Schema — and shows their relationship during an entity's perceptual cycle. Exploration is the act of interacting with the environment and collecting information. The Object represents the phenomenon or element of the environment observed by the entity, while the Schema represents the entity's accumulated beliefs.

In Neisser's model, the Schema directs Exploration, with the entity's existing beliefs influencing how and where it will gather information from the environment. An entity will therefore focus its information collection on an Object that its Schema indicates is important. Exploration samples the available information about the Object, with the entity collecting information that may describe the Object's physical properties or its current state. Analysing and interpreting the sampled information modifies the entity's Schema, thereby updating the entity's beliefs in some manner. The entity's updated Schema will now direct its future decisions and actions, as it will influence the entity's future Exploration of its environment. In Neisser's framework, actions are performed by the entity during its Exploration step.




Figure 2.5: Neisser's Perceptual Cycle Model (Neisser, 1976). An entity's interactions with its environment modifies its understanding of the environment, which then shapes its future interactions with the environment.

Neisser's framework recognises that an entity's existing beliefs determine how it collects and interprets new information. Since the collected beliefs drive the entity's Exploration, these beliefs can and will bias how it gathers and interprets information. Due to the cyclic nature of the framework, an entity's existing false beliefs can lead it to incorrectly explore its environment and gather incorrect information into its Schema, thereby compounding the errors from false beliefs. However, future Exploration may also sample correct information, allowing the entity to correct its Schema.

Norman's (1990) Human Action cycle describes the process undertaken by individuals as they perform tasks with a computer and intends to aid the design and analysis of user interfaces for computer software. The cycle begins with Goal Formation, where individuals determine what they wish to do with the computer. This is followed by the Execution phase, where individuals translate this goal into a set of ordered tasks and perform them. The final phase is Evaluation, where individuals perceive the results of their actions, interpret the outcomes and compare the actual result to their desired result. While intended to focus on human-computer interaction, Norman's model shares its structure with other perception cycle models, encapsulating a cycle of information collection, decision-making and action. As in the other models, individuals possess beliefs that direct their decisions and information interpretation, which is updated through repeated interaction with the computer system.

Artificial Intelligence and Artificial Life simulations often use a similar model for the decision-making processes of simulated agents (Russell and Norvig, 2009). During such a simulation, an agent gathers information with its sensors and determines the state of its environment. The agent then decides upon a suitable action, given its understanding of the environment. Finally the agent implements this action, altering the environment in some manner it may perceive in future iterations of the cycle. This model contains three main elements found within the other perceptual cycle models — perception, decision-making and action.

The scientific method embodies the perceptual cycle of a scientist, and even of the scientific community as a whole (Zumdahl, 2007); wherein a scientist or the scientific community updates and refines his or her or their knowledge through iterations of observation, hypothesis, prediction and experimentation. Firstly, observations provide the scientists with information about some phenomenon in the environment. These observations are shaped, to some degree, by the scientists' current understanding of both their environment and the phenomenon. The scientists produce Hypotheses that explain the observed phenomenon, given their current scientific understanding. These hypotheses allow the scientists to predict new phenomena, which they may then test with experiments. These experiments produce further observable phenomena, starting a new iteration of this cycle. When considering this process in terms of the scientific community, publishing experimental results communicates these new ideas among the community and thereby affects future cycles.

It is worth noting that the act of observation is not just an act of collecting information, but an experience (Hanson, 1958), which is affected by the observer's prior knowledge and experiences. Essentially, what observers see is conditioned by their prior knowledge, which directs how they understand a perceived phenomenon or object. This conditioning biases observation, preventing neutral observation. As such, two individuals with different knowledge can observe the same phenomenon, yet develop different understandings of it.

When an experiment's results disprove a hypothesis, scientists must develop a new hypothesis to explain the observed phenomena and then test it through further experimentation. By confirming a hypothesis, scientists increase and update their body of knowledge, which then allows them to develop and test further hypotheses. Repeated iterations of this cycle update and refine the scientists' body of knowledge, with confirmed hypotheses

adding new knowledge or refining existing knowledge and disproved hypotheses removing false beliefs. Like the other models, the scientific method also describes a feedback loop where the scientists' experiments influence their environment in some manner that they later observe.

Similar to the Scientific Method, Kolb et al.'s (1971) model of the learning or problem-solving process describes an individual's learning as a continual cycle through a four-step process (Figure 2.6). In this process, individuals observe and reflect upon their concrete experiences, allowing them to develop abstract concepts and generalisations, which they then test in new situations. They also state that an individual's existing knowledge produces needs or goals, which direct the individual's learning.




Figure 2.6: A model of the Learning or Problem Solving Process (Kolb et al., 1971).

Kolb et al. (1971) apply their model to identify an individual's learning style, thereby potentially identifying flaws in the individual's learning style that the individual may correct. However, this knowledge is also valuable to the individual's competitors, who can potentially use it to exploit the individual. For example, if an individual is known to learn best through hands-on interaction and experimentation, a knowledgeable competitor could benefit from restricting or preventing the individual's attempts to learn in such a manner. However, a competitor must first analyse its intended victim to effectively limit its learning, since the victim is unlikely to publicly advertise any such flaws. Also, a competitor should conceal such exploitation and its preparation.

The Observation Orientation Decision Action (OODA) loop model (Boyd, 1986, 1996) is another method of modelling an entity's decision-making and action cycles. It was originally developed to model the decision-making process of fighter pilots; however, its generality allows the modelling of many decision-making entities. The OODA loop model

is commonly used to model the decision-making process in both military and business strategy (Thompson, 1995).

Boyd's OODA loop is a four step cyclic model that describes an entity's information gathering, decision-making and actions, with earlier behaviour providing feedback to the current analysis and decision activities (Figure 2.7). The model breaks the continuous act of perception and its subsequent decision-making into four discrete steps, which are accurate for many entities or systems. The model can be adapted to systems that are not discrete. A typical and implicit assumption in the OODA loop model is that it involves players in a competitive game, which may be one of either complete or incomplete information. In either circumstance, what information these players perceive, and how they understand or misunderstand it, determines the players' subsequent actions or moves and the game's eventual outcomes and payoffs.




Figure 2.7: Boyd's OODA (Observation Orientation Decision Action) loop model (Boyd, 1986). Note the feedback from present Decisions and Actions to future Observations, as well as the control that Orientation has over Observation and Action.

The loop begins with the **Observation** step, where the entity collects information about the state of its environment. It may collect this information with any sensors it possesses, such as its eyes or ears. Where and how the entity gathers information is guided by its beliefs.

During the **Orientation** step, the entity combines the gathered information with its stored beliefs about the environment and itself, which may include previous experiences, cultural traditions, genetic heritage, and analysis and synthesis methods. The entity uses

all of these elements to update its internal model of its environment. The internal model represents the entity's current understanding of the state of its environment and as a product of the entity's perceptions, beliefs and information processing abilities, this model may not accurately match reality. Boyd (1987) states that the Orientation step is the "schwerpunkt" (focal point or emphasis) of the OODA loop model, as an entity's Orientation determines how it will interact with the environment; affecting how it Observes, Decides and Acts.

During the **Decision** step an entity considers its potential actions and the expected outcomes of these actions. If the entity is a rational decision-maker, it will select the action or actions that it believes will lead to its most preferred outcomes. The possible actions and their expected outcomes that the entity develops are entirely a product of the entity's model of its environment. Therefore, an entity will only consider actions it believes are possible and the outcomes it believes those actions will have. The entity's beliefs therefore constrain its decision-making.

During the **Action** step, the entity performs its selected action or actions, changing the state of the environment. The entity can observe these changes in future OODA loop cycles, along with any changes caused by other entities. This functions as a feedback loop between the entity and its environment.

Boyd argues that in a conflict, the side that operates at a faster tempo through its OODA loop, while denying such an ability to its opponent, will have an advantage. This is referred to as "operating inside an opponent's OODA loop" (Boyd, 1987). Faster progress through the OODA loop makes one appear ambiguous and therefore unpredictable to an opponent, which may generate confusion and disorder in the opponent (Boyd, 1986). This leads an opponent to misunderstand its reality, thereby provoking incorrect responses. While an entity operating inside an opponent's OODA loop can cause the opponent to misperceive, the entity may have difficulty accurately predicting such misperception, therefore limiting its exploitation.

A common element of these perceptual cycle models is that the entities possess an internal model of their current understanding of the state of their environment. Entities create and update this model as they interact with their environment and it affects how they interpret information and make decisions. Since an entity must interpret any information in terms of its understanding of its environment, entities with different understandings

can develop different interpretations of an observed phenomenon. This is consistent with Jauss's (1982) Reader Reception Theory, which states that a reader's prior life experiences affect how it processes and interprets elements of a story, which determines how it understands the story.

Hanson (1958) proposes a similar model of perception, arguing that what an entity perceives is not the complete information its senses gather, but instead a filtered version of this information. An entity creates a filter from its existing preconceptions and therefore "sees" not what it visually perceives, but instead what it expects to see. This argument is supported by optical illusions of an ambiguous image, where the entity's preconceptions affect how it resolves the image. For example, consider what is "seen" by two astronomers, one from the 13th century and the other from the 20th century, while observing a sunrise (Hanson, 1970). Hanson argues that the two astronomers will perceive two different things due to their differing astronomical beliefs, even though they witness the same phenomenon and the same visual information reaches their eyes. The 13th century astronomer believes in a geocentric model of the universe, where the Sun revolves around the Earth and sees the Sun rising in the morning. However, the 20th century astronomer believes in a heliocentric model of the universe, where the Earth orbits the Sun, and therefore sees the Sun made visible by the Earth's rotation. All of these perceptual cycle models demonstrate how understanding shapes an entity's perception.

These models also all implement a feedback cycle between the entity and its environment. An entity's actions will affect its environment in some manner, which it can then observe during future perceptual cycle iterations. This feedback can help entities to correct their incorrect beliefs when differences arise between the entities' understanding of their environment and their observations of its actual form.

The models also do not require entities to develop or maintain a correct understanding of their environments, which allows them to model misperception and its effects on the entities' decision-making processes. Misperception can cause an entity to possess false beliefs, which will affect both its current and future decision-making and reasoning. Future iterations of the perceptual cycle may reinforce, propagate or correct false beliefs. Entities learn about their environment by observing and interacting with it. Since false beliefs can affect an entity's decision-making processes, misperception is capable of affecting an entity until it is corrected. Different types of misperception can affect the decision-making

process in different ways; however, these effects have not been thoroughly examined to date.

Of the perceptual cycle models discussed, Boyd's OODA loop model is the most suitable for examining the decision-making processes of a variety of entities and the effects of misperception on their decision-making processes. While the elements Boyd described as part of the Orientation step are specific to human information analysis and interpretation, we may ignore any inappropriate elements when considering the Orientation of non-human or non-biological entities. Boyd also considers an entity's decision-making speed important in competitive situations. By Observing, Orienting, Deciding and Acting faster, an entity can influence the environment before its opponent. This changes the state of the environment that the opponent must interact with, possibly altering the effects of the opponent's actions. Faster decision-making creates and widens the gulf of understanding between the faster and slower entities, causing the slower entity difficulty in correctly responding to the actions of the faster. The lack of understanding exhibited by the confused entity is a product of its misperception. While this property can exist in other perceptual cycle models, it is commonly only considered in conjunction with the OODA loop model. Due to these useful differences, I will use the OODA loop model for any subsequent analysis of an entity's perceptual cycle.

2.7 Categorising the Sources of Misperception

The term misperception covers a wide range of phenomena, which may introduce a variety of errors into an entity's understanding of its environment. Some types of entities, such as biological organisms, machine systems and human organisations, may suffer from some types of misperception while ignoring others. I previously defined misperception as errors affecting either perception or interpretation. External entities or some dysfunction of the entity may cause these errors. As such, I will categorise the sources of misperception based upon how they affect an entity and whether or not they are intentional.

As the perceptual cycle models have demonstrated, an entity's beliefs affect how it interprets newly acquired information. If these beliefs are false, an entity may not correctly interpret new information and create further false beliefs. In the worst case, this process

can function in a repeated feedback loop, where an entity's false beliefs lead to incorrect interpretations that produce further false beliefs.

It is worth noting that this categorisation is not intended to be an exhaustive and in-depth study of misperception's possible sources, but is instead intended to only provide a broad overview.

2.7.1 Information Gathering Errors

Perception, as the act of gathering new information from the environment, is the obvious element of an entity's perceptual cycle where errors may lead to misperception. It is also the first part of the perceptual cycle where misperception may occur. These errors commonly occur due to flaws of the entity's information sensor, of which there are a great variety. Sensors possessed by biological organisms include eyes, ears and noses, while machine systems commonly possess mechanical, electro-optical or electrical sensors that measure properties of the environment that are important to the system. In some instances a machine system's sensors may duplicate the functionality of a biological systems' sensors, such as video cameras and microphones, while other sensors may have no biological analogues, like radar. The two main classifiers for information gathering errors are whether or not the error was caused intentionally and whether the affected entity is a biological organism or a machine system. Further distinctions may be made based upon the flaw that affected the information sensor.

Unintentional Errors

Some unintentional perception errors are caused by flaws in the sensors or mechanisms an entity uses to gather information from its environment. Sensors may fail either completely or partially. A sensor that has completely failed will provide no valid information, while a partially failed sensor may only provide valid information intermittently or provide information that is always incorrect, yet only by a small amount. Biological, electrical and mechanical sensors differ in the types of flaws that affect their collection of information. Competitors may exploit unintentional errors by using them to conceal information from an affected entity. Competitors may also unknowingly benefit from an affected entity's sensor limitations.

Natural Deterioration of information sensors occurs in biological organisms as the organism ages and the effectiveness of its sensors declines. Macular degeneration is an example of this deterioration and typically occurs in elderly people (de Jong, 2006). Machine sensors often consist of many components, which may fail due to corrosion, wear out, or a lack of maintenance. Components commonly fail either early in their lifespan or much later in their lifespan (Bazovsky, 2004). After such a failure, the sensor will collect either incorrect information or no information at all, causing the entity difficulty in correctly identifying environmental elements and phenomena.

Accidental Damage describes any event in which the information sensor is unintentionally damaged to the degree that it is partially or completely disabled. Partially damaged sensors may return incorrect or lower quality information, subsequently causing the entity difficulty in correctly understanding aspects of its environment. Accidental damage to the areas of the brain responsible for the various senses can produce similar effects to sensor damage (Doty, Yousem, Pham, Kreshak, Geckle and Lee, 1997). Machine sensors may be accidentally damaged during construction, maintenance, repair or usage of the sensor or its associated system. Sudden physical impact or exposure to dangerous conditions or substances — such as water, electromagnetic fields or rapid changes in temperature — are some mechanisms that may damage machine sensors. Other potential dangers to sensors include abrasion caused by wind-driven particles and biological attack from micro-organisms. Accidental damage to the entity’s information sensors produces the same effects as a Denial attack.

Sensor Limitations are caused by inherent flaws or features of the sensor, which render it unable to detect information correctly or at all in some circumstances. These limitations are typically dependent upon properties of the information itself and have typically always existed in the entity’s sensor. Bats, dolphins, whales and porpoises all use high frequency sound for navigation (Hughes, 2001). Humans cannot hear these sounds due to the limitations of our ears. These species are not attempting to intentionally exploit the limitations of human hearing and therefore cannot intentionally cause misperception. Design limitations can also prevent mechanical and electronic sensors from correctly perceiving or operating in some conditions, allowing a sensor to only collect valid information under a limited range of conditions. As an example, a common household scale designed to weigh people cannot correctly weigh objects or organisms much heavier or lighter than

a person. It is possible to intentionally exploit the known limitations of information sensors by manipulating an information signal so that the sensor cannot detect it. Such exploitation is an example of a Degradation attack.

Inherited Genetic Conditions are limited to biological organisms and occur when the affected organism possesses a deleterious mutation that impairs the correct development or functioning of an information sensor. Albinism is one such condition (Witkop, Quevedo, Fitzpatrick and King, 1989) and it may cause astigmatism and light hypersensitivity, which reduces the organism's effectiveness in gathering visual information.

Illnesses and Diseases are another potential cause of sensor flaws that are unique to biological organisms. Illnesses and diseases can produce symptoms that degrade or disable an organism's information sensors. Affected organisms may suffer from a complete or partial loss of sensory input while they are ill, with the sensor functionality possibly restored after the organism recovers. For example, humans may suffer a temporary loss of hearing caused by infections of the ear (Griffith, 1995).

Misdirected perception occurs when an entity's information collection tasks are directed away from helpful information. Perceptual models show that an entity's existing beliefs direct how and where it collects information. If these beliefs are incorrect, an entity may ignore relevant sources of information, focus on information sources that are irrelevant, or otherwise adjust or deploy its sensors in a manner that prevents the collection of helpful or relevant information. The entity will not recognise such shortcomings as it collects information. Misperception occurs because the entity unknowingly avoids or disregards potentially helpful information. Feedback within the entity's perceptual cycle allows false beliefs to misdirect the entity's perception. The offensive actions of competitors or any unintentional error may cause false beliefs that lead to misdirected perception. As an example, consider a scientist who is convinced by a deceptive spin attack that a certain academic journal is of poor quality and therefore contains no useful research. This belief will direct the scientist's information collection away from this journal and any potentially useful information that it contains.

Intentional Errors

Unlike the many unintentional errors that may affect an entity while it perceives its environment, there are few intentional errors that affect an entity's information collection.

Intentional errors are caused by the Information Warfare attacks of competing entities and produce similar effects to unintentional errors, specifically Denial attacks and Degradation attacks.

Denial Attacks are performed against an entity's information sensor, either disabling it temporarily or destroying it completely. For example, one species of stick insect sprays an irritating fluid into the eyes and nose of a potential predator (Kopp and Mills, 2002), which temporarily prevents the predator from perceiving the stick insect, allowing it to escape. Denial attacks are highly overt, with the victim likely to notice the attack. However, in some instances an attacker may perform a covert Denial attack. Such an attack attempts to disguise the sensor's damage as an expected type of accidental damage or reliability failure the sensor could experience, while also concealing any evidence of the attacker's involvement in such an attack.

Degradation Attacks are another attack performed against an information channel to prevent an entity from observing some crucial signal in its environment and may be either active or passive.

During an active attack, the attacker floods the environment with a noise-like signal, so that the victim's information sensor cannot detect the signal that the attacker wishes to conceal. Barrage radar jamming is an example of an active Degradation attack, as it floods the frequencies used by a tracking radar with noise (Van Brunt, 1978). This noise hides the radar return of a vehicle the attacker wishes to conceal, which is commonly an aircraft.

A passive attack alters a signal so that the victim cannot detect it among the environment's existing background noise. In a biological context, the evolved camouflage markings of many species perform passive Degradation attacks. The markings and shape of such animals often make them visually indistinguishable from the background of their environment, as in the case of the peppered moth (Majerus, 1998) and the polar bear (Dawkins, 1986).

Subversion Attacks degrade a victim's information collection processes by manipulating how the victim performs these tasks. An attacker may manipulate the victim's assessment of the relevancy, validity or importance of an information source. If the victim believes that an information source is incorrect or irrelevant then it may ignore the source, thereby degrading its information collection tasks.

2.7.2 Information Processing Errors

Various errors during the processing and interpretation of gathered information can also cause misperception. Such errors affect the various processing operations the entity performs on the information, causing it to analyse the information incorrectly. Therefore the entity will develop incorrect conclusions, which will likely affect its future behaviour.

Unintentional Errors

There are many possible unintentional errors that may occur as an entity processes information. An unintentional information processing error occurs when an entity develops an incorrect understanding of its environment, despite collecting correct information. This may occur if the apparatus an entity uses to process information inherited or acquired a flaw, or the methods an entity uses to process information suffered some error. In any case, the entity desires to correctly process the information, yet is incapable of doing so.

Misidentification is one type of information processing flaw that affects an entity as it attempts to recognise and identify the various elements within the information gathered by its information sensors. There are far too many potential flaws that may occur in these systems to give a complete overview, so this section instead aims to discuss a small number of examples. While this discussion focuses on misidentification in biological systems, machine systems may also possess these flaws and exhibit the same or similar problems.

Many biological organisms rely heavily upon their sense of vision. Researchers have identified various flaws of the human visual system, labelling some of these flaws optical illusions (Wade, 1982). Optical illusions are information processing errors where the brain incorrectly interprets the visual information it perceived, which is attributed to biases of the human visual processing system. Optical illusions may cause people to see hidden objects in an image, movement where none exists or an image as something it is not. The perception of the shape and orientation of objects is also commonly distorted by optical illusions. These illusions are caused by the architecture of the human information system and identify a flaw in the cognitive systems responsible for their interpretation. The illusions are not errors of perception as the information is collected correctly from the environment.

Ambiguous images are one type of optical illusion and occur when an image has multiple different interpretations, of which any interpretation is valid (Heuer, 1999). When

observing an ambiguous image one will “see” one interpretation of the image, which may then appear to switch between its different possible interpretations. Examples of such an illusion are the Necker Cube (Figure 2.8), which shows a wire-frame cube with an ambiguous orientation, and the Rubin Vase (Figure 2.9) (Brain, 2000), which one may interpret as either the silhouette of a pair of faces or a vase.




Figure 2.8: A Necker Cube (Necker, 1832). The ambiguous cube (a) may be interpreted as it is viewed from either slightly above (b) or slightly below (c).




Figure 2.9: A Rubin Vase (Rubin, 1915). The image may be interpreted as either the silhouette of a white vase or a pair of black faces.

Such optical illusions may not exclusively affect humans. Ruxton (2002) proposes that the zebra’s stripes exploit a cognitive flaw in one of the zebra’s predators, providing an evolutionary advantage. The stripes of a herd of zebra are hypothesised to dazzle lions, preventing a lion from easily identifying a single zebra among a herd and thereby reducing the lion’s ability to hunt zebra, which relies upon targeting individuals. As in the case of human optical illusions, such a flaw, if it exists, is caused by a limitation of the lion’s visual processing systems.

While these illusions demonstrate flaws and biases that exist in the parts of the brain responsible for processing visual information, illusions are not restricted to vision and may affect other senses (Deutsch, 1980; Ramachandran and Hirstein, 1998). Attackers may exploit these flaws to gain an advantage, as in the case of the zebras and the lions. The zebras are performing a Degradation attack against the lions, which exploits the limitations of the lions' visual processing system.

False Beliefs can also cause misperception, as an entity uses its current beliefs to analyse new information. These beliefs may describe the state of the environment, phenomena in the environment, other entities in the environment, relationships between other elements of the environment or methods for processing new information. Some of this information may be physically encoded into the entity, while other elements may be acquired and manipulated by the entity. Misperception may occur when an entity analyses new information with erroneous beliefs or if it uses the beliefs in an unsuitable manner. Any previous information gathering or processing errors may have caused these false beliefs. If an entity utilises false beliefs to analyse new information, the entity may interpret the information incorrectly. For example, consider a person who mistakenly believes that large pharmaceutical companies frequently produce dangerous drugs and are only motivated by profit. A person holding such a belief may avoid any drugs produced by those companies and pursue alternative treatments for any ailments, while disregarding information from those drug companies.

Incorrect Information Processing Methods may also cause an entity to misperceive. These misperceptions occur when an entity's information processing methods are flawed in some manner and produce incorrect results. While superficially similar to false beliefs, incorrect information processing methods are distinctly different. Misperceptions due to false beliefs occur after the correct interpretation of false information; whereas incorrect information processing methods perform invalid operations upon valid information, thereby producing misperceptions. These incorrect methods may have been learned or physically encoded into the entity and there are many different ways in which a processing method may become incorrect. If the method was learned by the entity, it may have been incorrectly remembered or taught. If the processing method is an unchangeable element of the entity, it may always have been incorrect. Biological organisms may be

born with such flaws, while machine systems may have these flaws designed and built into them.

These processing errors may also arise if the entity misuses its processing methods or beliefs in situations where they are unsuitable. Such an error occurs as the entity is unaware of the reason that its processing methods or knowledge are unsuitable. As an example, Joe, a civil engineer, could produce an erroneous set of plans for a wooden bridge, by using various engineering methods that were only suitable for stronger steel and concrete bridges. Joe's plans would therefore contain potentially dangerous flaws, of which he is unaware due to the unsuitable application of his engineering knowledge.

Physical Damage can also cause an entity to misperceive when its elements responsible for information processing are damaged. In complex biological organisms, the brain is the organ mostly responsible for analysing and interpreting information, as well as storing beliefs. Electronic systems typically use computer hardware and software for their information processing systems, with the software controlling the system's behaviour. In either biological organisms or machine systems, damage to the physical elements of the information processing systems may degrade an entity's capability to correctly process information, thereby causing misperception. Non-persistent errors, such as soft errors that cause flipped bits in digital systems, also produce similar effects. Physical damage may alter or destroy either the stored processes or the physical devices used for information processing, thereby detrimentally affecting information processing in the future.

Biological organisms may suffer from various illnesses, diseases and conditions that affect the neurological systems responsible for information processing. Humans and other animals may suffer from delirium, which reduces a victim's awareness of its environment and changes its cognition in a manner that may impair problem solving or memory recollection (Hodges, 1994). Aging may also cause an acquired intellectual deterioration, such as dementia or Alzheimer's disease (Bennett and Aggarwal, 2004). In machine systems, aging can degrade the components of a system to the point where they may introduce subtle errors during the system's operation, such as well known degradation effects in semiconductor components, which may affect the system's storage devices or processing hardware. If the processing hardware is affected by component failure, information processing may produce incorrect results. Damage to the storage mechanism may corrupt or destroy the system's information processing methods or other data, which will cause

processing errors should the system attempt to use the affected processing methods or data.

The **Communication of Incorrect Beliefs** can also cause misperception. Entities may unknowingly communicate false beliefs to others, effectively performing an unintentional Corruption attack. If the unintentional Corruption attack succeeds, the victim may then share this incorrect belief with others. Such an error only affects entities who can communicate. Mimetic theory (Dawkins, 1976; Lynch, 1996) explains how beliefs may spread throughout a population. Entities communicate these beliefs not because they are correct, but because they are pleasing. For example, an ideology may spread throughout a population — despite any false beliefs it imparts — if there is a real or perceived benefit for its adherents.

A communicated false belief negatively affects an entity’s information processing and analysis, in the same manner as any other false belief the entity possesses. Like a Corruption attack, it fails if the victim does not accept the attacker’s false belief. This leaves the victim to determine whether or not the Corruption attack was intentional and how it should respond. The victim may decide to retaliate against what it perceives as an attack or to warn the communicating entity of its error. Any such warning may benefit the entity who unknowingly communicates incorrect beliefs, as it has effectively outsourced its own error detection methods to its victim. This benefit, however, depends upon the other entities being capable of recognising the communicated error; being trustworthy enough to warn the entity; and successfully communicating a warning to the entity that the entity then understands.

Entities may also communicate incorrect beliefs if they are indifferent to the truth or correctness of their statements. Frankfurt (2005) describes such communication as “bullshit”, which is intended not to deceive, but distract or misdirect. As such, a victim receives a bluff instead of an overt lie, which superficially appears reasonable, but fails under closer scrutiny. Since the entity did not intend the bluff to be completely truthful it can communicate incorrect beliefs and thereby unintentionally cause misperception.

Intentional Errors

Only Corruption attacks and Self-deception can intentionally cause errors as an entity processes information. Corruption attacks are initiated by a competing entity, who aims

to cause its victim to misperceive. Self-deception is unique among the sources of misperception, since it is an intentional attempt by an entity to cause its own misperception.

A **Corruption Attack** is one of the canonical strategies of Information Warfare during which an attacker places a corrupted signal into the victim's environment. The corrupted signal mimics a signal the victim would expect to find. The victim observes and then interprets this signal, incorrectly accepting the corrupted information as valid information and analysing it accordingly. This gives the entity incorrect beliefs, which will affect its decision-making process in some manner as planned by the attacker.

The corrupted information provided to the victim may take many different forms, depending upon the type of entity attacked and the aim of the corruption attack. Corruption is widely used by animal species, who typically mimic the appearance of other species to aid their own survival (Kopp and Mills, 2002). While such Corruption attacks may only have immediate short-term effects, others may last much longer. For example, the education given to the students of a totalitarian regime (Kim, 1969; Wertsch, 1999) will affect how these students act and interpret information, likely for much of their lives. Therefore, such indoctrination represents a compound strategy employing both Corruption and Subversion.

Subversion Attacks can also cause misperception when an attacker manipulates the victim's information processing capabilities. After such an attack, the victim misuses its own information processing capabilities to derive false conclusions about the state of its environment. While marketing and advertising may often employ such spin attacks (Kopp, 2006a), more severe examples are cases of religious or ideological indoctrination. The previous example of Corruption also describes how indoctrination by totalitarian regimes affects how people interpret and analyse information for much of their lives, thereby subverting their ability to accurately assess information from their government, or any other source.

Self-deception is a method by which an entity may intentionally cause itself to suffer a misperception. A Self-deceiving entity intentionally alters either its new information or its existing beliefs, to produce a new desired belief or reinforce an existing belief. Self-deception is assumed to be unique to humans and human organisations, since other entities presumably lack the required motivations or cognitive mechanisms. While other biological organisms may Self-deceive, there are no documented studies identifying this

behaviour. Current electronic systems are also rejected as potential Self-deceivers due to both a lack of cognitive power and motivation — it is highly unlikely such a system would be intentionally designed with the capability for Self-deception. However, if future Artificial Intelligence systems can fully mimic human cognitive systems, then such systems could Self-deceive.

While the hypothesised motives for Self-deception in humans and human organisations differ (Trivers, 1976; Ramachandran, 1996), either motivation leads an entity to intentionally manipulate its own beliefs, as it deems appropriate, while processing and analysing new information. The entity intentionally decides to possess false beliefs, which are constructed to produce or support a desired belief over an unwanted belief. These false beliefs demonstrate that a misperception has occurred. They may also affect an entity's future interpretations and analyses, causing further misperceptions.

2.7.3 Incestuous Amplification

The models of the perceptual cycle demonstrate that an entity's stored beliefs can affect it when it processes or collects information. False beliefs, caused by previous misperceptions, are also a potential source of misperception. They are therefore capable of affecting how new information is analysed, or even whether it is gathered at all, which can lead an entity to develop further false beliefs. Misperception can therefore function in a positive feedback loop, where each misperception in turn leads to further misperceptions. Spinney (2013) labels this process as **incestuous amplification**, noting its prevalence within the Pentagon. As the gulf between the entity's perceived and actual environment increases, the entity is more likely to misperceive as it analyses new information; potentially leading to further misperceptions, which may then lead to further misperceptions, ad infinitum.

Such incestuous amplification can produce a framework of interdependent false beliefs from an initial handful of false beliefs. These beliefs could be modelled as a graph, with the false beliefs connected by their dependencies. While either intentional or unintentional misperceptions may contribute to incestuous amplification, it does not fit exclusively in either category. Incestuous amplification may affect any type of entity.

An example of incestuous amplification can be found in the events that led the Soviet Union and the United States close to a nuclear war in 1983 (Walker, 1993; Fischer, 1997; Gaddis, 2005; Hughes-Wilson, 2006). This example describes the incestuous amplification

of misperceptions by the governments and intelligence agencies of the Soviet Union and the United States. The amplification begins with an initial false belief held by the Soviets, which dramatically alters the way in which the actions of the United States are interpreted.

In early 1981, the United States began a PSYOPS (Psychological Operations) campaign against the Soviet Union, probably intended to unnerve and confuse the Soviet Union (Hughes-Wilson, 2006, p 292). During the campaign, air and naval forces participated in probing missions near the Soviet borders to observe their defensive capabilities and responses. This act probably initiated the scare, as the Soviets incorrectly determined that these missions were collecting intelligence specifically for a future nuclear first strike. This belief led the Soviets to issue an intelligence alert, redirecting their intelligence agencies to gather any information of a nuclear first strike. If the Soviets detected indications of a potential first strike, they could pre-empt it with either conventional or nuclear attacks.

The scare worsened in 1982, as the United States was about to deploy nuclear Pershing II and Tomahawk missiles in Europe, giving them the capability to attack the Soviets' in-theatre nuclear facilities (Hughes-Wilson, 2006, p 293). The deployment of these missiles reinforced the Soviets' belief that the United States intended a nuclear strike. The United States then proposed the development of the Strategic Defense Initiative (SDI), a defensive system that would protect the United States from nuclear ballistic missiles (Walker, 1993). While intended purely to protect the United States against a nuclear attack, the Soviets' existing beliefs convinced them that such a system was only necessary to defend against their retaliation to the United States' nuclear first strike.

Soon after this the Soviet Union shot down a South Korean airliner, flight KAL 007, that had crossed into its airspace (Fischer, 1997). The late Soviet response to this incident claimed that the airliner was testing Soviet air defences for the United States, who were therefore responsible for the deaths of the passengers on board KAL 007. This belief may have been due to the earlier PSYOPS campaign, in which United States' aircraft had regularly penetrated Soviet airspace, or it may have been an attempt to justify their actions. In response to the incident, the United States increased its defence spending, which it justified by increased Soviet aggression. The Soviets' expectation of a nuclear first strike led them to interpret this increased expenditure as an indication that the United States intended to launch a nuclear attack.

The crisis worsened as the United States began a planned military exercise, codenamed ABLE ARCHER 83. Part of this exercise had the United States simulate the preparations for a nuclear attack against the Soviet Union. Given the existing Soviet belief of an aggressive United States, the ABLE ARCHER exercise was incorrectly interpreted as a cover for an expected nuclear strike. Two further unrelated events also reinforced the Soviet belief of an imminent nuclear attack (Walker, 1993, p 276). Firstly, the United States put its forces in the Middle East on alert, following the bombing of a Marine barracks in Lebanon. Secondly, the United States exchanged encrypted communications with Britain before invading Grenada, a member of the British Commonwealth. Once again, the Soviets interpreted somewhat innocent actions as evidence of an upcoming nuclear attack, which reinforced the belief that ABLE ARCHER would be a nuclear attack against the Soviet Union. The Soviets responded by readying their nuclear armed aircraft in Poland and East Germany, to retaliate against the expected attack. Later, as the exercise was under way, Soviet intelligence sent a mistaken warning that the United States bases were on alert (Walker, 1993, p 277), which was interpreted to mean that an attack was imminent. The crisis finally ended when the exercise concluded without the nuclear attack the Soviets expected.

The earlier probing missions were intended as muscle-flexing exercises by the United States, which the Soviets incorrectly interpreted as the preparations for a nuclear first strike attack. This false belief affected how and where the Soviets subsequently gathered information and then affected how they interpreted such information, which eventually caused them expect a nuclear attack and prepare accordingly. These subsequent incorrect interpretations also acted to reinforce the false belief that the United States intended to launch a nuclear attack. In this example, the feedback arises from the information processing errors caused by the initial false belief. None of the Soviets' misperceptions were intentionally caused by the United States, but instead they were caused by the incorrect interpretation by the Soviets. This scenario may also be considered as a hypergame, where the Soviet Union's game models an expected nuclear attack and their responses to it, while the United States' game models what appeared to be typical Cold War posturing by the Soviet Union.

2.7.4 Summary

Various errors that occur as an entity gathers or processes information from its environment may cause misperception. These errors may have either intentional or unintentional causes; with the unintentional causes attributed to various accidents, flaws and dysfunctions, while the intentional errors are produced by the Information Warfare attacks of competitors, except in the case of Self-deception. Competing entities may exploit unintentional misperceptions if they know of the flaw responsible for the misperception. If an entity is known to possess a certain incorrect belief, an attacker can provide information that exploits this belief to provoke behaviour desired by the attacker. As the attacker and victim possess different understandings of the situation, a hypergame can model their competitive interaction.

Incestuous amplification describes a feedback cycle where misperceptions produce false beliefs, which in turn produce further misperceptions. This cycle leads an entity to produce a network of interdependent incorrect beliefs, likely affecting both its current and future behaviour while also potentially causing future misperceptions.

2.8 Background Summary

In order to address the question of whether misperception can produce an evolutionary benefit, it is necessary to first understand a diverse amount of existing research. The most important of these is evolutionary theory, which describes how organisms evolve and reproduce, along with the role their physical and behavioural traits play in this process. The evolutionary study of misperception will assume that an entity's misperception is caused by some trait that the entity possesses and that this trait is encoded into the entity's chromosome. If misperception is to provide an evolutionary benefit, there are several ways in which misperception can potentially benefit entities individually or collectively.

Artificial Life or game theory can both model the interactions of misperception-affected entities. Artificial Life offers a method of simulating biological processes using computer simulations, while game theory provides a mathematical framework for analysing strategic games between entities. Both of these methods can model situations where misperception may occur regularly and affect an entity's behaviour and interactions with others of its population.

While previous research has identified limited situations where misperception has potentially provided an evolutionary benefit, there are few studies that explicitly focus on misperception or its potential advantages. These few studies have indicated that there are some instances where misperception can benefit entities in competitive situations, including some that incorporate evolution. This supports the hypothesis that misperception can provide an evolutionary benefit; however, the scarcity of existing research does little to identify the conditions under which this occurs.

Studying the potential sources of misperception also identifies a link between misperception and Information Warfare. Information Warfare consists of numerous competitive behaviours that are intended to intentionally cause others to misperceive. While such misperceptions are not intended to benefit the misperceiving entity, they do share the same effects as other unintentional sources of misperception. As such, an attacker can disguise its Information Warfare attacks as unintentional misperceptions, which is highly advantageous. While misperception can have both intentional and unintentional sources, only unintentional misperceptions are expected to benefit a misperceiver. The relationship between these intentional and unintentional misperceptions is interesting, especially their ability to produce identical effects. Further study of these similarities will help to further understand both Information Warfare and unintentional misperception.

There are various models that describe the information collection and decision-making cycle of entities. In all of these models, an entity proceeds through a cycle of information collection, processing, decision-making and acting. During this process the entity manipulates and references a store of beliefs that guide each of these actions. Information Warfare attacks and misperception may both affect these beliefs, creating false beliefs that may affect an entity's behaviour or help create further false beliefs. As an entity acquires more false beliefs, its understanding of reality increasingly diverges from actual reality, commonly to the entity's detriment. This suggests that an entity's continued misperception will cause it to suffer when its false beliefs strongly direct its decision-making process. Entities are unlikely to evolutionarily benefit from such severe misperception.

Examining the existing literature of the study and simulation of misperception reveals several shortcomings, which this thesis aims to address. One such shortcoming is a detailed study of misperception, determining how it affects an entity's perceptual and decision-making processes in general. It is not currently clear how the decision-making

process is affected by Information Warfare and unintentional misperception. Current perceptual cycle models do not attempt to describe how misperception affects their “normal” operation. Neither is it clear what steps may be taken, if any, to minimise the harmful effects of misperception on the decision-making process.

Another shortcoming is the scarcity of Artificial Life simulations investigating any potential advantages of misperception, especially in an evolutionary environment. This is likely due to the expectation that misperception is either detrimental or an undesirable nuisance. In the few simulations that have studied misperception’s benefits, the simulations have not always modelled an evolutionary environment. Furthermore, it is not clear whether alternative mechanisms for introducing errors into an entity’s perceptions or decisions can functionally replace misperception.

While existing research has examined how misperception disrupts cooperation between entities, in some scenarios misperception may nevertheless convince otherwise non-cooperative entities to cooperate. Therefore, exploring the potential situations under which this may occur is worthwhile.

The next chapter will further explore these shortcomings, proposing several avenues of research to determine whether misperception can provide an evolutionary benefit and the circumstances under which it may do so.

Chapter 3

Research Problems and Methodology

Examining the literature discussing misperception and its related topics reveals an overall paucity of research that specifically studies misperception. In the existing research there are several distinct gaps that may be explored. These gaps have been identified in the existing research into misperception covering areas such as Information Warfare and Artificial Life. This research aims to address some of the existing shortcomings in the study of misperception, and also to attempt to connect the pre-existing studies of misperception and explain the implications of this work.

One such shortcoming is the lack of a detailed examination of misperception and its underlying mechanisms. While typically assumed to be an error that affects perception, misperception also covers errors that affect the interpretation of information. Examining the various possible sources of misperception has identified a dichotomy between intentional and unintentional sources of misperception, where intentional misperception is typically caused by Information Warfare attacks and unintentional misperception is caused by various errors of the affected entity. This clarifies what misperception entails, but it has not revealed how it affects an entity's decision-making process or its immediate and long-term effects upon entities. Study of misperception's effects upon an entity's perceptual process will clarify how it ultimately affects that entity's decisions and actions. In the case where these errors are due to Information Warfare attacks, insight into this process can aid attackers in their utilisation of attacks and may also aid victims in defending against attacks.

Some prior simulations have identified situations where misperception is beneficial and some of these situations are evolutionary. Such simulations support the hypothesis that misperception may provide an evolutionary benefit and have demonstrated instances where misperceptions may discourage harmful behaviour or encourage beneficial behaviour. Akaishi and Arita's study of misperception in a foraging simulation makes two interesting and original claims that should be investigated further. Both of these hypotheses will be explored, as they support the main argument of this thesis; that misperception can provide an evolutionary benefit.

The Iterated Prisoner's Dilemma is another situation where misperception may be beneficial. Misperception, or any other source of noise, in the Iterated Prisoner's Dilemma has been demonstrated to disrupt cooperation between players utilising the Tit for Tat strategy, thereby reducing their scores. However, in the case that those players are unable to maintain or develop Cooperation, misperception may also cause these players to Cooperate, thus benefiting both players. In such cases misperception functions similarly to forgiveness. Therefore, misperception that implements forgiveness is another potential method by which it may prove beneficial, especially in an evolutionary environment.

3.1 Understanding Misperception

It should be reiterated that this thesis is not arguing that misperception is universally beneficial for misperceiving entities. Indeed, there are many occasions where misperception is clearly detrimental and undesirable. Instead, this thesis argues that there are some situations where misperception can provide an enduring evolutionary benefit. Such an analysis of misperception requires its potential effects on the decision-making processes of affected entities to be examined, in order to identify how it affects entities and under what conditions its effects may be advantageous.

Misperception and some of its sources were discussed previously in Section 2.4, along with several previous simulations that had demonstrated beneficial misperception. It was also stated that misperception does not affect an entity's rationality during its decision-making, but instead manipulates the information upon which rational decisions are made. The various potential sources of misperception were categorised in Section 2.7. These

sources may be categorised by whether or not they are intentional, or upon whether misperception initially impairs either the collection or the analysis of the information.

Since this thesis aims to investigate how misperception may affect a variety of different types of entities, it will be necessary to describe the structure of a generic entity. This generic entity is an abstracted description of any type of entity that could conceivably be affected by misperception. Such an entity is assumed to possess some information sensors for collecting information from its environment and some limbs or actuators for interacting with its environment. These entities are also assumed to possess a mechanism that stores the beliefs they develop from analysing information collected from their environment. While these beliefs will describe the entity's understanding of its environment, there is no requirement that they must always be true. An entity will decide how to behave in its environment based upon the beliefs it holds. This entity model can also be considered an abstraction of the Artificial Life agents created in the simulations. This generic entity model will be used throughout this thesis when considering the effects of misperception upon entities. Specific types of entities, such as humans, animals, human organisations or machine systems, can be considered as specific instantiations of this generic entity, with their own unique sensors, actuators and methods of storing beliefs and making decisions. This abstract model of an entity is also intended to be compatible with many of the decision-making cycles that were previously discussed.

To further explain how misperception affects entities, the various sources of misperception will be mapped into a generic model of the decision-making cycle — Boyd's OODA loop model. This will reveal how the various possible types of misperception may affect an entity's decision-making process. This work will also identify any similarities in the processes by which misperception affects entities. The Orientation step of the OODA loop model is the point in which entities will maintain their beliefs, updating existing beliefs and creating new ones. Since misperception will affect these beliefs, the Orientation step is affected by misperception in some manner. The internal process of the Orientation step is vaguely defined in Boyd's model. To better explain how misperception affects this step of the OODA loop, a procedural model of the Orientation step's internal processes will be developed. Sources of misperception that primarily affect the Orientation step will then be examined in terms of this new model to determine how misperception affects entities

during their Orientation step. When misperception is examined in this manner, its immediate effects on the affected entity are considered. The potential long-term effects of the misperception may also be determined, as well as its potential effects on future decisions and interactions with other entities to some degree.

3.1.1 The Effects of Misperception on the Decision-making Cycle

Misperception may occur due to a variety of sources, which may be either intentional or unintentional. The existing studies of misperception do not explain how these different sources affect an entity's decision-making cycle. Determining how misperception may affect an entity's decision-making process will provide further insight into misperception's short-term and long-term effects. It also provides a procedural and chronological demonstration of how the distinct sources of misperception map into an entity's decision-making cycle. Such a process will highlight any vulnerabilities within an affected entity, providing possible clues as to how an entity may protect itself from misperception or minimise the impact of misperception upon its decision-making. Potentially, these insights may also be applied in reverse, allowing hostile entities to better increase the success rate of misperceptions targeted against their opponents.

In order to better clarify the procedural elements of the various sources of misperception, it will be necessary to model each of these potential sources in terms of how they affect each element of the decision-making cycle. Boyd's OODA loop model will be used to provide a discrete representation of an entity's decision-making process. In the case of intentional sources of misperception, the decision-making processes of both the attacker and the victim will be examined.

The model of the OODA loop that was discussed previously (Section 2.6, Figure 2.7) describes the transition between the four states of the loop and the control feedback between these steps. Here I am more interested in the flow of information between these steps, as it will detail how information moves within the OODA loop model as the entity acts upon it in various ways. This requires a minor modification to the OODA loop model to display the flow of information through decision-making entities, instead of the model's control and feedback processes. The addition of the information flow to the OODA loop model and the removal of the model's control and feedback processes is a

simple modification. This altered model will then be used to examine the various sources of misperception.

The sources of misperception to be examined in this manner are the four canonical Information Warfare attacks (Degradation, Corruption, Denial and Subversion), Information Gathering Errors, Information Processing Errors and Self-deception. The Information Gathering Errors and Information Processing Errors broadly encapsulate the unintentional sources of misperception. Given the competing hypotheses that explain Self-deception's underlying motives, multiple models of Self-deception will be considered. This mapping will take the various actions of each misperception and then fit them into the location of the entity's OODA loop in which they occur, to produce a procedural description of each misperception's effects upon an entity's decision-making process.

Once the various sources of misperception are mapped into the OODA loop model, they can be compared to identify any similarities. Given a description of how these misperceptions affect an entity, it will also be possible to determine how the effects of these misperceptions may be mitigated. These possible methods will describe general methods by which entities may attempt to protect themselves from misperception. Entities who desire to intentionally cause misperception through Information Warfare attacks can also consider these methods as defensive counter-measures that their attack must overcome in order to succeed. Attackers may also use the knowledge of a victim's vulnerabilities to misperception in order to better increase the chances of a successfully causing the victim to misperceive.

Historical and hypothetical examples of misperception will also be used to demonstrate how each type of misperception maps into the OODA loop model. Such examples describe how a specific instance of misperception occurred and often describe both the immediate and long-term effects of that misperception. Historical case studies are preferable to hypothetical examples, due to their concreteness.

Boyd has also stated that operating at a faster tempo through the OODA loop can produce misperception in a competing slower opponent. The method in which this misperception is caused will be identified and contrasted against the other sources of misperception. This will be achieved by concurrently comparing the OODA loop sequences of two entities, one of which has a slower tempo than the other. The various tasks of each entity's OODA loop will be compared as they occur and the differences in the environment

created by the players will be analysed. This will determine how such misperception is caused and where it fits into the existing categorisation of misperception.

3.1.2 Misperception and the Orientation Step

The Orientation step is a highly important element of Boyd's OODA loop model, where an entity interprets newly gathered information in the context of its existing beliefs. The entity then updates its beliefs with the interpreted information. Misperception aims to give entities false beliefs, and the Orientation step is where entities create and update their beliefs. Therefore, the Orientation step is the point where an entity may acquire false beliefs from misperception. The Orientation step describes a complex process with a fairly simple abstract description. Boyd states that through a process of analysis and synthesis, new information from the environment is somehow understood and combined with an entity's existing beliefs, in order to update the entity's beliefs. The details of the processes that occur during the Orientation step are not explained in great detail. Examining the internal processes of the Orientation step will further clarify how misperception produces errors during this step.

I will examine the Orientation step by developing a model that explains its internal processes. This model will break the Orientation step down into its procedural sub-steps, thereby describing the various operations of the Orientation step. My simplified model of the OODA loop will be modified to display these sub-steps and their information flows within the model. The sub-steps themselves will also be analysed to determine what types of errors can occur during each sub-step and whether such errors are misperception or not. This model will be tested by mapping the sources of misperception found to affect an entity's Orientation step into the extended OODA loop model. Mapping will convert misperception's actions into sub-actions that will then fit into their respective Orientation sub-steps. This will describe the misperception's errors during the Orientation step, specifically the sub-steps during which misperception affects an entity. Comparing misperception's effects upon an entity's Orientation step will identify similarities between different sources of misperception. Corruption, Self-deception and Information Processing Errors will be mapped into the model of the OODA loop and its expanded Orientation step, since they all primarily produce errors during the affected entity's Orientation step.

The model will be tested by studying historical and hypothetical case studies of misperception. This will provide further insight into the Orientation step model and also the processes of misperception that primarily affect the Orientation step. Historical case studies, typically of disasters or accidents, are suitable for examining the effects of misperception for numerous reasons. Disasters, mistakes and other such failures that are frequently attributed to misperception often have much documented historical discussion and analysis. With hindsight it is also easy to identify the misperceptions, along with the false beliefs, flawed reasoning and erroneous actions that are its hallmarks. Historical analysis of disasters often considers the effects of any errors, such as misperception, and their role in the disaster. They also provide “real world” examples of misperception, which are often more interesting than purely hypothetical examples. Hypothetical examples will only be used to describe misperception in the event that suitable historical examples cannot be identified.

3.2 Misperception in a Foraging Environment

The examinations of misperception’s effects on an entity’s decision-making process will focus on single instances of misperception, describing the effects of misperception on a single entity. Artificial Life simulations allow the effects of misperception on an entire population to be modelled, showing how it affects the population as a whole through the various interactions between the affected entities. It also allows the simulated entities to be repeatedly influenced by misperception and the results of continued misperception observed over time. In this manner many more instances of misperception can be observed and the compounded effects of such errors within a population observed and quantified over time.

While misperception has been demonstrated in some situations to benefit either affected entities or their populations, there are very few simulations that have investigated these benefits. Doran’s simulations modelled specific misbeliefs in the simulated agents and demonstrated that in two different evolutionary environments, false beliefs benefited the agents through two different mechanisms. The simulation work later performed by Akaishi and Arita also demonstrated a benefit from misperception, although this benefit was not evolutionary as they claimed.

While Akaishi and Arita's simulation work does demonstrate instances of beneficial misperception and suggests a mechanism for how this benefit is provided, they do not show that this result is producible in an evolutionary situation, nor do they prove that the behavioural diversity caused by the entities' misperceptions is the underlying source of the observed benefit. This work aims to address these shortcomings and provide further evidence in favour of the hypothesis that misperception can provide an evolutionary benefit. This work will also investigate the extent to which behavioural diversity benefits the foraging agents.

3.2.1 Evolutionary Advantage in Akaishi and Arita's Simulation

Akaishi and Arita claimed that their simulation demonstrated an evolutionary benefit from misperception. However, this claim was not supported by their work, as their simulation did not model an evolutionary environment. The agents that were modelled did not reproduce and misperception was not a heritable trait of the agents that could vary among the agent population. They instead assumed that the increased quantity of resources collected by a population of misperceiving agents would translate to increased reproductive success for those agents, thereby potentially providing an evolutionary benefit. While this may be a reasonable assumption, it was not supported by their results. Furthermore, using a fixed misperception probability for the entire population only identifies whether there is a group-wide benefit to misperception. Allowing the agents to evolve differing misperception probabilities would help reveal whether the benefits of misperception extend to individual agents or their kin, as well as identifying the optimal misperception probability for a given environment.

The hypothesis that misperception is providing an evolutionary benefit will be tested by converting the original foraging scenario proposed by Akaishi and Arita into a similar scenario that occurs in an evolutionary environment. This simulation will attempt to closely approximate the evolutionary processes that are found in the real world. Since I could not obtain a copy of Akaishi and Arita's simulation, I will instead re-implement their simulation, adapting it to include the required evolutionary methods. The simulation describes a simple two-dimensional foraging world, where agents will maintain a model of their perception of the environment. The agents will move around this environment,

gathering resources from fixed resource nodes. Misperception will affect an entity's perception of the resource nodes in its environment. In this evolutionary environment, an agent's gathered resources do not measure its fitness, but its ability to reproduce.

Over the length of the simulation, the population of agents will evolve distinct probabilities of misperception that are optimal for their environment. The hypothesis that misperception can provide an evolutionary benefit in this foraging scenario can be tested in several ways. If misperception is beneficial, then the population will evolve to a stable state where the average misperception probabilities of the agent population are significantly different to zero misperception. This shows that the agents in the population are misperceiving and that the misperceiving agents have not died out due to evolutionary pressure against this trait. Furthermore, beneficial misperception implies that misperceiving agents should parent more offspring than non-misperceiving agents, which will also be measurable.

3.2.2 Behavioural Diversity through Misperception

Akaishi and Arita also claim that the observed benefit from misperception is due to an increase in the behavioural diversity of the population. This claim is not directly supported by their results, which do not show any measure of the level of behavioural diversity in the agent population, other than the population's global misperception probability. This diversity in behaviour is said to be the underlying source of the benefit that was also attributed to misperception. In this manner misperception is one means of increasing a population's behavioural diversity and thereby providing a benefit. Based on this assumption, it is hypothesised that other sources of behavioural diversity will also benefit the agent population in a manner similar to misperception. This hypothesis will be tested by determining whether various other methods of introducing behavioural diversity into the agent population are also evolutionarily beneficial.

Behavioural diversity defines some abstract measurement of how differently the various members of a population act. Misperception increases a population's behavioural diversity by altering an entity's understanding of its environment, which subsequently affects its actions. Diversity thereby increases as agents are affected by various differing misperceptions. Behavioural diversity therefore occurs at the population level, due to the misperceptions of the individual agents in the population.

Akaishi and Arita's existing simulation already compared the resource gathering capabilities of agent populations with varying global misperception probabilities. Their results indicate that populations with some level of misperception performed better than populations with no misperception. Since the populations affected by misperception will have more behavioural diversity than those with none, Akaishi and Arita's work supports the hypothesis for a non-evolutionary environment. This work will test whether the hypothesis also holds true in an evolutionary environment. It will also further explore the hypothesis by altering the existing simulation and replacing misperception with other methods that may produce what Akaishi and Arita have described as behavioural diversity. If these methods also produce an observable advantage similar to that exhibited by misperception, then this diversity is the underlying source of misperception's benefit.

One alternative source of behavioural diversity is misaction. While misperception affects the agent's understanding of its environment, misaction affects its ability to correctly perform its intended actions. In this foraging simulation, misaction will affect the movement of the agents and agents will inherit their misaction probability from their parents. Another method of introducing behavioural diversity is to replace the simple resource location and decision-making processes of the agents with purely random movement. These agents will move randomly about the environment, only gathering resources when they encounter resource nodes. This method will introduce a large amount of diversity into the population's behaviour, which may not benefit the agents if its effects are too severe. A system with little behavioural diversity will be created to act as a baseline for comparison with the other foraging methods. This system will simulate a population of agents who are unable to misperceive and whose offspring will also be unable to misperceive. Such agents will be unaffected in any way by misperception, leading to very little behavioural diversity.

If increasing a population's behavioural diversity is beneficial as hypothesised, then the simulations with an increased diversity should exhibit increased resource collection and reproductive success. Similarly, the simulation with decreased behavioural diversity should exhibit less foraging and reproductive success than the other simulations.

3.3 Misperception and the Iterated Prisoner's Dilemma

The Iterated Prisoner's Dilemma is a well-known social dilemma that demonstrates the struggle between cooperation and selfishness that players face in a competitive environment. Noise disrupts the cooperation of players in the Iterated Prisoner's Dilemma, reducing the payoffs both players receive. If players are using retaliatory or unforgiving strategies, once noise disrupts their cooperation they are often unlikely to resume cooperating.

The effects of noise upon players in an Iterated Prisoner's Dilemma contest has been explored by Axelrod and Dion (1988). They state that noise may be modelled in one of three ways: as an error that alters the communication of the move played; as the incorrect implementation of a selected action by a player; or as a perceptual error when a player observes its opponent's previous move. The first type of noise is due to some element of the communication system, while the second and third are errors of the players. This definition of noise is specific to the Iterated Prisoner's Dilemma and differs to Shannon's (1948) widely used mathematical definition. When discussing the Iterated Prisoner's Dilemma, noise will collectively refer to all three of these errors, while a specific error will be referred to by its name (misperception, mis-implementation or communication error).

Existing research often focuses on the use of the Tit for Tat strategy during the Iterated Prisoner's Dilemma because it is both simple and yet effective at establishing and maintaining mutual Cooperation, thereby providing the best long-term payoff for the game. The Tit for Tat strategy has also been identified in various real world situations that are similar to the Iterated Prisoner's Dilemma. For these reasons this research of the Iterated Prisoner's Dilemma will focus exclusively on players using the Tit for Tat strategy.

3.3.1 Using Misperception to Restore Cooperation

When two Tit for Tat players compete in an Iterated Prisoner's Dilemma game in a noisy environment, any noise will disrupt the established cooperation between the two players and thereby reduce their scores. The Tit for Tat players are unable to recognise this noise as not being an intentional move by their opponent and are therefore unable to resume mutual Cooperation, unless subsequent noise changes the game to this state. Noise is

therefore generally detrimental to these players. However, there are some situations where selective limited noise may benefit these players.

Consider an iterated game between two Tit for Tat players, where one or both players manages to mis-implement their initial move. This game starts, and will continue indefinitely, in a non-cooperative state. A possible cause for this state might be that a third-party has managed to influence a player's initial move in some manner, such as through an Information Warfare attack. This third-party presumably benefits from this act in some way that is outside the game's payoff system. While Tit for Tat is a forgiving strategy, it will not resume cooperation until its opponent does so first. Playing against itself, neither Tit for Tat player will do anything other than reciprocate the earlier defection, to both their detriment. In the worst case scenario for these players, the initial errors are the only instance of noise that affects them, especially if it is intentionally caused by an external third-party. This act locks the Tit for Tat players into a sub-optimal state, from which misperception may represent a potential solution.

Misperception may be able to benefit Tit for Tat players of the Iterated Prisoner's Dilemma in some noisy environments. In this environment noise may be provided by either misperception, misaction or a combination of the two. There may also be no sources of noise present. The Tit for Tat players may also find themselves affected by a third-party, who forces one of the players to Defect initially. Misperception is hypothesised to aid such players if it can reliably return the game to a state of mutual cooperation. This hypothesis will be tested with an Iterated Prisoner's Dilemma game between two Tit for Tat players, where players affected by beneficial noise should receive a higher score than those who are not.

3.3.2 Asymmetric Misperception

Forgiveness is another solution to the effects of noise on reciprocating strategies, such as Tit for Tat. A forgiving player is one who Cooperates against some percentage of the Defections it observes. Forgiveness may also be considered as a specialised misperception, in which some percentage of the Defections a player observes are misperceived as Cooperation and provoke unwarranted reciprocal Cooperation. The other type of misperception that may occur during the Iterated Prisoner's Dilemma is the misperception of Cooperation as Defection. I define such misperception as punishing misperception, as it will

punish cooperative players, disrupting Cooperation between players and thereby reducing the payoffs they both will receive.

Existing research into the effects of misperception on the Iterated Prisoner's Dilemma has focused on misperception that simply toggles the observed previous move of an opponent. Such misperception functions as a combination of both forgiving and punishing misperception. Misperception may instead be modelled by dividing it into both punishing and forgiving misperception, each of which occurs with a separate probability. While forgiveness has been considered as a solution to the problem of noise, its suitability in the face of potential exploitation has not been investigated. This suitability may be further reduced in an evolutionary environment, where players may adapt and begin to exploit those who are excessively forgiving.

In an Iterated Prisoner's Dilemma game played by Tit for Tat players in a noisy environment, forgiving misperception is expected to help the players to Cooperate. Since it will maintain or create a state of mutual Cooperation, forgiving misperception will be beneficial. However, in any environment where such forgiveness is likely to occur, punishing misperception can also be beneficial. An unprovoked Defection caused by environmental noise is indistinguishable from one caused by punishing misperception and both are equally likely to be forgiven by a player. Therefore, punishing misperception may exploit the forgiveness of others by Defecting against a Cooperating player and then being forgiven for this transgression. In such a case, the player who Defects will receive a higher payoff that iteration, at the expense of its opponent. However, if the unprovoked Defection is not quickly forgiven, then both players will suffer with lower payoffs until they resume Cooperation.

While forgiveness has been argued by Axelrod to benefit players of the Iterated Prisoner's Dilemma, it is not clear whether this benefit is obtainable when faced with punishing misperception that will exploit it. It is hypothesised that there is an optimal forgiving misperception probability that benefits the players in a noisy environment, whilst also being small enough that it is not detrimentally exploited by punishing misperception. Such an optimal value may best be identified by creating an Iterated Prisoner's Dilemma tournament that operates within an evolutionary framework. The players will have distinct probabilities for forgiving and punishing misperception, which will evolve within the player population as multiple generations compete in the tournament.

If misperception is beneficial then the players should evolve a stable forgiving misperception probability to counteract the noise in their environment. When forgiving misperception is beneficial in this noisy environment it will increase the players' average score, close to the score received by cooperating players in a noise-free environment. The punishing misperception probability of the population should also be low if forgiving misperception does not occur frequently enough to be reliably exploited.

3.3.3 Preventing Exploitation of Forgiveness

While the separation of misperception into forgiving and punishing misperception may allow players to evolve forgiving misperception to counteract environmental noise, punishing misperception may subsequently evolve in the agent population to exploit excessively forgiving players. This problem arises as the players cannot differentiate between Defections caused by noise and Defections caused by exploitation. Since the evolution of exploitative punishing misperception relies upon some level of forgiving misperception, there presumably exists some optimal level of forgiving misperception that is both high enough to counteract environmental noise and low enough to avoid the evolution of exploitation within the population.

This optimal level will be identified by globally restricting the forgiving misperception probability that the players may evolve. Placing an upper bound on this value will allow the potential values of this variable to be explored and the optimal level of forgiving misperception thereby identified. The optimal level of forgiveness will be identified from the players' scores and evolved punishing misperception probability. Populations where forgiving misperception is beneficial will show scores indicative of mutual cooperation that is only slightly affected by noise, along with little or no evolved punishing misperception. Populations where forgiving misperception is exploited will have demonstrably lower scores than the cooperative populations and significant levels of punishing misperception.

3.4 Unaddressed Research Topics

There are several identified shortcomings in the existing body of research that will not be addressed in this thesis, yet are suitable topics to address in future research.

Most Information Warfare research focuses on its application to computer systems, telecommunications networks and other modern technological systems. However, Kopp and Mills have argued that Information Warfare is actually an evolved biological survival mechanism, to which the instances observable in modern systems are but a specific sub-type. As a survival mechanism, Information Warfare may manifest itself as an evolved trait in biological organisms, which provides an evolutionary benefit to organisms. While it is not necessary to demonstrate the evolutionary benefit of the various Information Warfare strategies, given the numerous examples found in biological systems, the biological underpinnings of Information Warfare could be better reinforced through further research. Such research would help to broaden the amount of Information Warfare literature that examines it from a evolutionary and biological context, instead of from a human social context.

In most of the simulations and models of misperception previously discussed, misperception occurs in a highly restricted scope, especially when it is compared to real world examples of misperception. Typically misperception is modelled as a Boolean perceptual error, which occurs with some probability. In these models misperception either occurs or it does not, providing a representation that fails to fully capture the nuanced errors that it may produce in perceptual systems. In many instances, an entity's beliefs do not describe a Boolean representation of some aspect of the environment, but a property of the environment that may have many potential values. Therefore, the degree to which an entity may misperceive these values is highly variable. As an example, a woman driving a car at 100 km/h who believes the car's speed to be 101 km/h is unlikely to be greatly affected by such a minor misperception. However, if she believed the car to be travelling at 75 km/h, then this greater misperception is likely to have a more significant effect on her behaviour. In another example, misperception could affect how two entities in a contest each perceive their relative strengths. Such misperception may produce over-confidence or over-cautiousness in the entities, both of which can impact their chosen behaviours. In these examples the effects of misperception are linked to its scale, instead of its occurrence. While this thesis will not investigate models with such scalar types of misperception, the development of these models of misperception would be a suitable topic for future research. Creating a model where an entity's misperceptions may have various levels of severity would allow the development of models and simulations that explore the different

effects that frequent minor misperceptions and rare major misperceptions may have upon an entity's behaviour.

Another aspect of misperception that this thesis will not investigate is the degree of strength with which an entity holds its beliefs. Modelling confidence would require a more complex structured system of beliefs for the entities, where the analysis of new information creates new beliefs and alters the degrees of confidence of other beliefs held by the entity. The more certain an entity believes something is true, the less likely it is to alter that belief and the more likely that belief will affect its behaviour. Conversely, beliefs in which an entity has little confidence are more likely to be altered by new analyses of information and less likely to affect an entity's behaviour. Misperception may also affect an entity's certainty in some of its other beliefs as it creates false beliefs. For example, when modelling Self-deception, a Self-deceiver could determine which beliefs it will manipulate to reduce cognitive dissonance, based upon their certainty. Presumably an entity would desire to alter beliefs that it is least certain are true, as it is less "invested" in the correctness of these beliefs. Assuming such a model of an entity's system of beliefs is created, the changes in an entity's certainty after an Information Warfare attack or any other misperception would provide a metric for assessing the severity of the error, and in the case of an Information Warfare attack the effectiveness of the attack. The development of such a metric would be a useful goal to pursue in future research of misperception or Information Warfare specifically.

These two methods of modelling misperception could be also combined, allowing an entity's beliefs to be incorrect by a varying factor and to also have a certainty to which it accepts these beliefs as true. The more certain an entity is that its beliefs are correct, the less likely it is to change them faced with contradictory information. This model would allow a much more realistic representation of misperception; however, the complexity of this model may be a drawback. Such a model should be tested in other circumstances where misperception has been identified to provide a benefit, evolutionary or otherwise. Entities possessing a more complex and realistic model of misperception should also benefit in the same situations as entities with simple models of misperception.

The current models of misperception used to simulate its effects on agents are fairly simple and, due to this, unrealistic. This research does not aim to address the shortcomings of these models; however, changes to these models are a suitable area for future

investigation. Using more complex models of misperception will also add complexity to the entities modelled, which may be undesirable in some cases. In the case of a simple game like the Iterated Prisoner's Dilemma, while such a complex model of misperception could be added to the players, it is unlikely to greatly affect the game's dynamics due to its underlying simplicity. The Boolean representations of misperception that have been used previously are fairly restrictive in that they do not permit entities to suffer from minor misperceptions. However, increasing the number of objects and phenomena that can be misperceived by simulated agents will also increase the complexity of analysing misperception's effect on the simulation. The development and investigation of more complex and realistic models of misperception is an interesting area for future research into misperception and its effects. These proposed models would be useful to further investigate misperception and to quantify its effects.

3.5 Summary

One research problem that this thesis will investigate is the effects of misperception upon an entity's decision-making process. This will be achieved by mapping the intentional and unintentional sources of misperception into Boyd's OODA loop model to determine how these instances of misperception function. This work will detail how misperception occurs and provide insight into how it may be prevented or its effects minimised. Modelling and analysing the Orientation step's internal processes will support this investigation of misperception. Where suitable, historical case studies of various types of misperception will be studied to further clarify how misperception occurs and what its effects may be.

Previous Artificial Life simulations by Akaishi and Arita have shown an advantage to misperception in a foraging environment. However, their research makes several claims that are not directly supported by their results. The claim that misperception is responsible for producing beneficial behavioural diversity is interesting, as it describes a new mechanism by which misperception, or any other source of noise, may benefit entities. If behavioural diversity is responsible for the benefit observed from misperception, other methods that create behavioural diversity should display a similar benefit. Testing both these claims requires the development of an evolutionary foraging simulation that closely matches that described by Akaishi and Arita.

Noise in the Iterated Prisoner's Dilemma typically disrupts Cooperation between players and reduces the payoffs they receive. This research aims to investigate whether misperception may aid these players by restoring a state of mutual Cooperation. When misperception is modelled with forgiving and punishing components, it is proposed that a population of players in a noisy environment will evolve an optimal forgiving misperception probability. This will occur despite the potential for exploitation from punishing misperception.

One final problem that this research work will address in its conclusions is to connect the findings from this study of misperception and explain its general and specific implications. This will include clarifying the relationship between misperception and Information Warfare and identifying the similarities and differences between the two. The Artificial Life simulations of misperception are intended to demonstrate situations where misperception can provide an evolutionary benefit. It is hoped that by comparing these situations and those from previous research it will be possible to identify some elements that are common among situations where misperception is beneficial. This will begin to formalise a theory of how misperception may be beneficial and describe the circumstances under which this is true.

Chapter 4

Misperception and the OODA Loop

As previously discussed, misperception may be either intentional or unintentional, with both producing similar effects in misperceiving entities. Mapping operations that cause misperception into Boyd's OODA loop model will better explain how these misperceptions occur. Moreover, this will also reveal which steps of the OODA loop are vulnerable to misperception. The sources of misperception that will be examined are the canonical Information Warfare strategies, Information Gathering Errors, Information Processing Errors and Self-Deception. These sources of misperception will be examined with historical and hypothetical examples of their effects on the OODA loop of affected entities. The feedback cycle that exists in the OODA loop model will also demonstrate any long-term effects produced by these misperceptions.

By identifying the methods by which Information Warfare attacks affect victims, it may be possible to determine methods that allow potential victims to identify attacks, prevent attacks or reduce the effectiveness of attacks against themselves. This is also true for unintentional misperceptions, which would allow entities to identify how their own errors are caused, what effects these errors produce, and finally, suggest methods that may reduce the effects of such errors or prevent them entirely.

While the OODA loop model has already been described, this discussion will use a simplified version of the loop that shows both the change in states and the flow of information through the loop. Detailing the flow of information through an entity will better

illustrate how misperception affects and disrupts the flow and processing of information in an affected entity.

4.1 Boyd's OODA Loop Model

There are many possible models that describe the perceptual cycle of an entity, which were described previously in Section 2.6. These models all described a similar process, in which an entity collects environmental information, analyses that information and then determines what actions it will perform. Each of these models share many commonalities and they can therefore be used interchangeably. This research will use Boyd's OODA loop model to represent the perceptual cycle of entities.

One reason for selecting the OODA loop model is its generality, which allows the model to be applied across decision-making entities from various backgrounds. Decision-making entities may range from simple entities to organisations or even to coalitions of complex organisations. A second reason for selecting this model is that it makes a distinction between information collection and processing. While in practice it may sometimes be difficult to differentiate between these two tasks, it is advantageous to refer to them separately. Another reason for selecting the OODA loop model is its concept of perceptual tempo. Boyd proposes that entities who proceed through their OODA loop faster than their opponents have an advantage, in that they can affect the environment before their competitors can do so. When performed correctly, this is said to degrade the slower opponent's ability to operate in its environment and cause the faster entity to appear ambiguous to the slower entity.

It should be noted that while an entity's decision-making process typically occurs continuously in the real world, in the OODA loop model the decision-making process is played out discretely. The OODA loop model consists of four discrete states that a decision-making entity passes through in repeated cycles. Events that occur in the real world may be mapped into the OODA loop model by assigning elements of the continuous time into each state.

Examining the effects of misperception on the OODA loop model first requires the model to be altered, to better demonstrate the internal flow of information as an entity proceeds through its OODA loop. This informational flow is not clearly shown in Boyd's




Figure 4.1: A simplified version of Boyd's OODA loop model with the transitions between steps marked with solid arrows and the information flows marked with dotted arrows. The entity's beliefs are stored internally and may affect all steps of the model, while all steps of the model may update an entity's beliefs.

diagram and may be confused with internal control feedback. Therefore a simplified variant of Boyd's model is required, with the feedback and control information replaced with the information flow.

4.1.1 A Simplified OODA Loop Model

Boyd's OODA loop model shows that feedback from the Decision and Action steps can affect future Observation steps, while the Orientation step, via the entity's beliefs, can guide the performance of the Observation, Decision and Action steps. A modified version of the OODA loop model is shown in Figure 4.1, which simplifies Boyd's original OODA loop model (See Figure 2.7) by displaying only the four main steps and the flow of information as the entity proceeds through the loop. This conceals the extra control and feedback signals, somewhat simplifying the model.

Like the standard model, the simplified model arranges the four main steps of the loop in a cycle. Observation leads to Orientation, which leads to Decision, which leads to Action, which leads to Observation again in the next iteration of the loop. These transitions between states are represented by solid arrows on the diagram. Displaying the flow of information through this model requires two abstracted information storage sources to be added. One is the environment, which is a source of information for the entity. Entities also place information into the environment as a result of actions performed during their Action step. The second information storage location is actually part of the entity itself and contains the entity's stored beliefs. This represents an abstraction of the entity's

memory and information storage mechanisms and is accessible in some manner from all the steps of the loop. This allows the beliefs to describe stored information for humans, organisations, biological organisms, machine systems or any other type of entity. In this manner, information flows into an entity from its environment and is then manipulated and stored as part of the entity's beliefs. These beliefs then affect the entity's Decision and Action steps, thereby affecting the information that an entity places back into its environment.

An entity's beliefs contain an internal representation of its environment. During the Observation step, newly gathered information is temporarily stored in preparation for its analysis and interpretation during the Orientation step. This alters the entity's beliefs, thereby updating the entity's internal representation of its environment. The beliefs then determine which options an entity will select during its Decision step, and this selection is also stored within the entity. The entity collects the selected option during its Action step, along with any pertinent beliefs required for the implementation of that option. It then performs the option, which changes the state of the environment and, therefore, the information it contains. These changes can be perceived during future loop iterations, demonstrating the feedback between past actions and the present state of the environment. Feedback also allows an entity to assess the effectiveness of its actions during future loop iterations.

The entity's beliefs are an important element of the model and affect each of the four OODA loop steps. An entity's beliefs are not required to be accurate, which permits it to develop and maintain false or inaccurate beliefs. These false beliefs represent flaws in an entity's understanding of some aspect of itself or its environment and, just like any other belief, may affect the entity at any point of its OODA loop cycle. An entity's beliefs direct its Observation to some degree, affecting what information it gathers and where it gathers this information. An entity that ignores important objects and phenomena in its environment or focuses on those that are unimportant is likely to develop a flawed model of its environment. The majority of the entity's information usage occurs during the Orientation step, where it uses its beliefs to analyse and understand new information, synthesise new beliefs from this information and update its beliefs when appropriate. The Decision step uses an entity's beliefs to select its future course of action and any false beliefs are likely to misdirect the possible actions considered or misrepresent the potential




Figure 4.2: A simplified OODA loop model showing how the interactions between two entities may only take place through the environment. Each entity's actions alter the environment, which may be later observed by both entities.

outcomes of those actions. An entity's beliefs may also describe how it should perform certain actions during its Action step.

Any interaction between entities within this model, including communication, occurs through the environment. One entity's actions manipulate the environment and this may be perceived by other entities during their Observation steps. Figure 4.2 demonstrates how this exchange occurs between two entities — although there is no limit, theoretical or otherwise, to how many entities may interact in such a manner.

These models will be used to detail the flow of information within a misperceiving entity while it proceeds through its OODA loop. I will now map various types of misperception into these models to determine how they affect an entity's decision-making process and behaviour. The points in an entity's OODA loop at which misperception causes errors will also be identified to demonstrate potential vulnerabilities of the loop. While the initial and subsequent effects of errors may occur during different OODA loop iterations, here they are assumed to occur in the same iteration for simplicity and brevity.

4.2 Information Warfare Attacks and the OODA Loop Model

With the exception of Self-deception, the main sources of intentional misperception are Information Warfare attacks. Borden and Kopp have categorised the various possible offensive actions of Information Warfare into four canonical strategies: Degradation, Corruption, Denial and Subversion. As already discussed, Borden's Exploitation is not an offensive Information Warfare attack and therefore does not interfere with the victim's

perceptual cycle in any way. In terms of the OODA loop model, Exploitation occurs during an attacker's Observation step, as it gathers information from sensors connected in parallel to the victim's.

Mapping the four canonical strategies into the simplified OODA loop model will demonstrate how each causes misperception to occur in the victim's decision-making process. These attacks show a pair of communicating entities, where the left entity is the attacker and the right entity the victim. Each attack commences during the OODA loop of the attacker and then continues during the next iteration of the victim, detailing the effects of the attack.

Since these attacks typically aim to give the victim false beliefs, the final state of a successful attack can be considered to be a hypergame. This can be modelled as a first order hypergame, where the attacker and victim both possess different beliefs as to the state of the environment, with the victim unaware of the attacker's intentions and actions.

4.2.1 Degradation Attacks

The canonical Information Warfare attack of Degradation describes an attack where a signal or message is altered in some manner that renders it imperceptible to the victim. This attack has two forms: an overt form of which the victim is aware, and a covert form, which the victim will not perceive. The overt form is likened to "jamming" and is actively performed by an attacker who floods an information channel with noise, noise-like signals or messages devoid of information content in order to conceal a signal from the victim. The covert form may be either an attribute of the entity, such as the camouflage used by many animal species, or it may be actively emitted by the entity, as in the case of some covert forms of communication. In either case, Degradation conceals the signal within the background noise of the information channel, cloaking it from the victim. Figure 4.3 shows how both the active and passive forms of Degradation affect the victim's decision-making process.

An active Degradation attack begins during the attacker's OODA loop cycle. The attacker performs this attack during its Action step, by emitting a large amount of noise into the environment. This noise is intended to drown out a signal in the environment the attacker wishes to conceal and continues for as long as the attacker wishes to conceal the signal. At this time, the victim commences its Observation step and collects information




Figure 4.3: A Degradation attack mapped into the OODA loop model. The attacker's signal is indistinguishable from the environmental background noise, preventing the victim from perceiving it.

from the environment. If the attack is successful the noise will conceal the signal, thereby preventing the victim's information sensor from observing it. Otherwise, the victim will perceive both the signal and the noise that was supposed to conceal it. A successful active Degradation attack therefore affects the victim's Observation step, preventing the perception of a signal. The perceived noise is then passed onto the victim's Orientation step. From the analysis of this noise, the victim can only ascertain that it is being attacked and will update its beliefs accordingly. During its Decision step, the victim will select its strategy, knowing that it is being attacked but unaware of the presumably important signal concealed by the noise. In the victim's Action step it may respond to the noise; however, it cannot respond in any manner to the contents of the concealed signal, of which it is unaware.

Active Degradation is employed by many cephalopods, who produce a cloud of ink to cover their escape from potential predators (Hanlon and Messenger, 1996). As a specific example, consider an instance of a seal attempting to predate upon a cuttlefish (Figure 4.4a). The cuttlefish releases ink during its Action step, clouding the water within its vicinity. It then jets rapidly away from the seal. During the seal's Observation step it attempts to observe the cuttlefish it was pursuing; however, the seal now sees only the cloud of ink, which masks the cuttlefish's escape. After updating its beliefs during the

Orientation step, the seal now knows that the cuttlefish it was pursuing is nearby, but it does not know where the cuttlefish was fleeing. Given this lack of knowledge, the seal cannot select or implement options that target the fleeing cuttlefish during its Decision and Action step. Instead it is faced with the options of abandoning the chase, or moving through the ink cloud and attempting to rediscover the cuttlefish. The escaping cuttlefish may also use a compound strategy to help escape the seal, by using its adaptive camouflage after inking to conceal itself from the seal. This camouflage may take two different forms, either mimicking inedible or inanimate objects (Corruption) or blending into the background surface of the seabed (passive Degradation).



(a) The Paintpot Cuttlefish (*Metasepia tullbergi*). ©まっちゃん, <http://photozou.jp/photo/show/78680/27521450>



(b) Female Peppered Moth (*Biston betularia*). ©Donald Hobern, <http://www.fotopedia.com/items/flickr-2938017418>

Figure 4.4: Examples of animals species that perform active and passive Degradation. Cuttlefish (a) use active Degradation to escape from predators. A cuttlefish will emit a cloud of ink to conceal its escape from a nearby predator, before quickly jetting away. The camouflage of the Peppered Moth (b) is an example of passive Degradation, which conceals the moth from its predators when it rests upon a similar background.

A passive Degradation attack modifies a signal emitted or reflected by the attacker so that it cannot be easily discerned from the environment's background noise. Potential victims are affected as they attempt to observe this signal in the environment during their Observation step. If the Degradation attack succeeds during the victim's Observation step, the victim cannot differentiate the attacker's signal from the background noise and is therefore unable to perceive it. The victim then passes this information onto its Orientation step, where it updates its beliefs from information that is correct — except for the existence of the concealed object. The victim's choice of options during its Decision step is constrained, as it is unaware of the attacker's signal and the concealed object. This

then affects the victim's Action step, as any action it performs is ignorant of the concealed object.

For an example of a passive degradation attack, reconsider the Peppered Moth (Figure 4.4b) discussed in Section 2.1.1. Both the lightly and darkly coloured phenotypes use passive degradation to conceal themselves within their distinct environments. When a predator, such as a bird, is searching for prey during its Observation step, it is unable to detect moths that are resting against surfaces that resemble the colouring of the moth. Here the moth's camouflage renders it indistinguishable from the background appearance of its environment. The bird then updates its beliefs during its Orientation step and since it did not see the moths, it cannot store beliefs of the existence or location of these moths. Subsequently, during the bird's Decision and Action steps it cannot select and perform actions against the hidden moths, such as preying upon them.

In both passive and active Degradation attacks, three distinct errors are observed. The first occurs during the victim's Observation step, as the victim is unable to correctly gather some important information from its environment. This causes the second error during the Orientation step, where inaccuracies induced by the first error are introduced into the victim's beliefs. The third error occurs during the Action step, as the victim performs actions that were selected based upon its false beliefs. Of these three errors, the last two are dependent upon the first and therefore the first error during the Observation step is deemed responsible for the misperception and its effects. Such an attack fails if the victim manages to perceive the important information during its Observation step, despite the attempts to conceal it. In the case of failure, the victim updates its beliefs with a correct understanding of what the attacker failed to conceal during Orientation. These beliefs should give it a better understanding of its environment for its subsequent Decision and Action steps.

4.2.2 Corruption Attacks

Corruption is the act of transmitting false information to a victim with the intention that this information will induce a specialised misperception in the victim. This misperception is intended to ultimately produce behaviour by the victim that is beneficial to the attacker. A successful Corruption attack mapped into the OODA loop model is shown in Figure 4.5.




Figure 4.5: A Corruption attack mapped into the OODA loop model. The attack causes the victim to adopt false beliefs, which ultimately cause exploitable behavioural errors in the victim.

A Corruption attack begins with the attacker proceeding through its own OODA loop cycle to determine the specific details of its attack. During its Action step, the attacker places a corrupted signal into the environment for the victim to perceive. This signal mimics a valid signal that the victim might expect in the environment. As the victim commences its next OODA loop cycle it collects this signal during its Observation step and passes it along to the Orientation step to be interpreted in terms of the victim's beliefs. The attack is successful if the victim accepts the corrupted signal from the attacker as a valid signal and updates its beliefs accordingly. The false beliefs will first affect the victim's Decision step, causing it to select a course of action that actually benefits the attacker. This action will subsequently be performed in the Action step, producing some change in the environment, which may leave the victim open to later exploitation by the attacker or benefit the attacker in some other manner.

The process of a Corruption attack mapped into the OODA loop model may be better explained by a historical example. Haswell (1979) described the series of strategic deceptions used during Operation Fortitude by the Allies to aid the D-Day landings in France. One aim of these deceptions was to convince the Germans that the invasion would be aimed at Norway and Calais, instead of Normandy. This was achieved by using numerous Corruption attacks to convince the Germans that a fictitious army, the First United States

Army Group (FUSAG), was preparing to land at Calais. Some of these attacks included radio transmissions that mimicked those a real army would transmit, the deployment of decoy tanks (Figure 4.6a), landing craft (Figure 4.6b) and armoured vehicles, and the appointment of General George Patton to lead FUSAG. Another false army group was created in Scotland to convince the Germans that Norway was also going to be invaded by the Allies. These deceptions were intended to keep German troops defending Calais and Norway, so they could not reinforce the German defence at Normandy when the invasion occurred.



(a) An inflatable decoy Sherman tank. ©Imperial War Museum (H 42531), <http://www.iwm.org.uk/collections/item/object/205201879>



(b) Dummy landing craft. ©Imperial War Museum (H 42527), <http://www.iwm.org.uk/collections/item/object/205201876>

Figure 4.6: Operation Fortitude consisted of many Corruption attacks, some of which created the fictitious FUSAG that was to invade Calais. These attacks mimicked the forces and materiel required for FUSAG's invasion, in part by deploying decoy tanks (a) and decoy landing craft (b). Further examples of such Corruption can be found in Figures 5.4a and 5.4b

In this example, Germany's various intelligence services and military leaders will be considered its actual decision-makers. During each OODA loop cycle the Germans collected deceptive information from various sources, such as radio receivers and aerial reconnaissance flights, during their Observation step. During their Orientation step, each item of corrupted intelligence was assessed and interpreted in terms of the beliefs the Germans already possessed. As each Corruption attack succeeded, it first created and then reinforced the belief that Calais and Norway were the real targets of the invasion, while any other attacks were a feint to divert the German defenders. Due to these beliefs, Germany decided that it should expect invasions at Calais and Norway and kept many of its troops in reserve to defend Calais from FUSAG's expected invasion. This Corruption attack succeeded so well that the false beliefs it created influenced Germany's Actions for

several days, well after the landings at Normandy. One of the reasons for the success of this deception plan was its multi-channel support (Haswell, 1985), where different sources each provided a corruption attack that was coherent within the overall plan, so that each Corruption attack supported and reinforced future Corruption attacks.

Corruption attacks generate two distinct errors. The first occurs during the Orientation step, where the deceptive signal is incorrectly identified as a real signal and the victim's beliefs updated accordingly. The second error occurs during the victim's Action step, when it performs actions it selected based upon false beliefs. No error occurs during the Observation step, as the deceptive signal was gathered correctly from the environment. A Corruption attack fails if the victim rejects the new information during its Orientation step, believing it to be erroneous or deceptive. Depending upon the specific details of the attack, the victim may become aware of the attempted attack and possibly identify the attacker. The victim could even turn the attack to its advantage, by performing its own Corruption attack in response; feigning a successful deception in order to then exploit the attacker. Such a response forces the situation into one with varying levels of perceptions and understandings, as in a hypergame. A Corruption attack may be considered a failure by the attacker if it creates a false belief in the victim that does not produce the intended exploitable behaviours.

4.2.3 Denial Attacks

Denial is the act of impairing, disabling or destroying the victim's information receivers, to degrade the victim's information collection or prevent it entirely. Reducing the quantity or quality of information the victim receives leaves it less able to accurately understand its environment. Such a lack of understanding may cause the victim to act unwisely or affect its interpretation of other information in the future.

Figure 4.7 shows a Denial attack mapped into the OODA loop model. A Denial attack begins when the attacker decides to perform such an attack and selects a sensor to impair, disable or destroy. During its Action step, the attacker impairs, disables or destroys the sensor. This action is typically overt and the victim is therefore likely to be aware that it has been attacked. During the victim's next Observation step, it receives little or no information from its affected sensor. Any information that is collected, along with the belief that its sensor has been damaged, is passed along to the Orientation step. At




Figure 4.7: A Denial attack mapped into the OODA loop model. Damage to the victim's information sensor prevents the victim from correctly understanding its environment, ultimately affecting its selected actions.

this point the victim has less information from its environment to update its beliefs than usual. With less available information, the victim is more likely to incorrectly interpret the current state of its environment and thereby develop false beliefs. The victim may also attempt to use its existing beliefs to determine what happened to its sensor and possibly assess whether it can identify an attacker. With less information to update its understanding of the environment, the victim will likely be ignorant of some element of its environment during its Decision step, which prevents it from choosing some potential actions. During its Action step it will perform its selected course of action, which will be ignorant of any recent changes to its environment. This course of action may benefit the attacker in some manner.

The destruction of some German radar emplacements by the Allies before the D-Day invasion, such as those shown in Figure 4.8, demonstrates a Denial attack. Here the Allies destroyed some German radar receivers to conceal their invasion from the Germans (Haswell, 1979; Holt, 2004). Interestingly, some radars were left operational so that they could be later exploited by some of the Corruption attacks discussed previously (Section 4.2.2).

The Allies first proceed through their OODA loop cycle, where they identified Germany's radar emplacements and decided to destroy them, which was achieved through



(a) Freya radar installation at Auderville, France. ©Imperial War Museum(C 5477), <http://www.iwm.org.uk/collections/item/object/205022365>



(b) Würzburg-Reise parabolic radar antenna. ©Sebastian Ballard, <http://www.geograph.org.uk/photo/769562>

Figure 4.8: Many German radar emplacements in France were destroyed before the D-Day invasion in Denial attacks. Both the Freya (a) and Würzburg (b) radar systems were targeted by such attacks, which were intended to hide the Allied invasion forces approaching France from the Germans.

bombing raids. After the targeted radars were destroyed, the Germans received no radar information from the destroyed receivers during their Observation step. However, the Germans were of course aware that their radar receivers had been destroyed. During their Orientation step, the Germans now had no information regarding the location of any nearby aircraft or ships and were only aware that some of their radar receivers had been destroyed. Without functioning radar receivers the Germans could not detect any aircraft or ships in some areas, which limited their defensive actions against any Allied aircraft or ships in those areas. The Germans' possible options were therefore reduced during their Decision step, with possible choices including replacing or repairing their radar receivers or communicating news of the attack to those responsible for overall defence of France. The Germans could not choose to take any actions against the aircraft or ships they could no longer detect. Germany's selected actions were then performed during its Action step. In this instance the Denial attack was overt, as it is impossible to conceal the destruction of their radar receivers from the Germans. With less functioning radar receivers, the Germans were unable to detect any Allied aircraft or ships in the affected region during their Observation step. They were therefore prevented from performing many actions against any Allied aircraft or ships until their radars were replaced.

Errors in the victim's OODA loop caused by a Denial attack can be identified at three distinct points. The first occurs during the Observation step, when the victim

cannot correctly gather information from its environment due to its disabled or damaged sensors. This then causes the second error, where the victim develops false beliefs from the incomplete information caused by the first error. The third error is the victim performing actions that are likely detrimental, due to the false beliefs it now possesses. The first error during the Observation step is responsible for causing the second and third errors and is therefore responsible for the misperception. An attacker may take advantage of a victim's disabled sensors by altering the environment in a manner that the victim cannot detect. The victim therefore cannot respond correctly to the environment, as it does not correctly understand it. Any false beliefs that the victim develops while its sensors are affected are unlikely to be corrected until it regains full functionality of its sensors, and even then it may not be possible to do so.

4.2.4 Subversion Attacks

The canonical Information Warfare strategy of Subversion communicates a signal to a victim that induces some form of involuntary self-destructive behaviour. In order to fool the victim, the subversive signal is often concealed with the aid of a Corruption attack, with the signal mimicking a signal that the victim expects. While the other canonical strategies attempt to manipulate information in some manner, Subversion differs in that it ultimately targets how the victim implements the various processes that make up its OODA loop. It is worth noting that not all Subversion attacks will produce misperception.

There are two distinct methods by which Subversion may cause a victim to misperceive. The first method is through spin or misdirection, where the victim's information sensors or information processing capabilities are manipulated and induced to function in a sometimes subtly different manner, thereby reducing the victim's understanding of its environment. The second method employs Subversion in a compound strategy, where a successful Subversion attack induces the victim to perform an involuntary self-targeted attack. For example, an attacker might use Subversion to induce the victim to damage its own information sensors, thereby implementing a self-targeted Denial attack.

Given that Subversion may target the underlying processes of any step of the victim's OODA loop, subtle manipulation of either a victim's Observation or Orientation step can cause misperception.

A Subversion attack that affects the victim's Observation step is shown in Figure 4.9. In this example, the attacker conveys a subversive signal to the victim during the attacker's Action step, which the victim gathers during its Observation step. This signal provides information which covertly manipulates how the victim collects new information, potentially affecting where, how or why the victim gathers information from some sources. In effect, the victim has been induced to create a blind-spot within its Observation step. During the victim's subsequent Observation steps it will now fail to fully observe information within that blind-spot, thereby causing it to misperceive. This lack of complete or accurate information causes the victim to develop false beliefs during its Orientation step, which then lead it to choose and perform actions during its Decision and Action steps that favour the attacker. Note that this differs from a Denial attack since the victim, and not the attacker, alters its information collection processes.




Figure 4.9: A Subversion attack mapped into the OODA loop model and targeted against the victim's Observation step. The victim gathers the subversive signal, inducing the victim to alter the process of its Observation step. This change causes a subsequent Observation step to collect incomplete information, which then causes the victim's Orientation step to develop false beliefs, leading the victim to act in a manner that benefits the attacker.

A Subversion attack that targets the victim's Orientation step, as in a spin attack, is shown in Figure 4.10. In this example, the attacker conveys a subversive signal to the victim during the attacker's Action step, which the victim gathers during its Observation step. The signal is then passed on to the victim's Orientation step, where the signal induces the victim to manipulate the methods it uses to interpret and analyse information.

During the Orientation step of subsequent OODA loop iterations, the victim will use the interpretation methods manipulated by the attacker to analyse and interpret information, ultimately interpreting some aspect of its environment in a manner that is favourable to the attacker. This causes the victim to develop false beliefs, which affect its Decision and Action steps and cause it to perform Actions that benefit the attacker. This attack differs from Corruption in that the victim’s own interpretation of information produces the false belief, whereas a Corruption attack has the attacker simply provide the false belief.




Figure 4.10: A Subversion attack mapped into the OODA loop model and targeted against the victim’s Orientation step. The victim gathers the subversive signal, which induces the victim to alter the processes of its Orientation step. This change causes a subsequent Orientation step to incorrectly analyse information and thereby develop false beliefs. These false beliefs ultimately lead the victim to act in a manner that benefits the attacker.

Subversion may also cause misperception when it induces an involuntary action that is a self-targeted Information Warfare attack. Figure 4.11 maps such a Subversion attack into the OODA loop model. Subversion begins during the attacker’s Action step, as it puts the subversive signal into the environment where the victim may perceive it. To conceal the true nature of the subversive signal, Corruption is often used. During the victim’s Observation step, it observes the subversive signal and passes it along to be analysed during its Orientation step. When the victim interprets the signal, it is likely to accept the signal as valid due to the Corruption attack, and it will therefore treat this signal as valid and use it as appropriate. Unlike a traditional Corruption attack, the subversive signal is not intended to overtly alter the victim’s beliefs, thereby affecting its Decisions and Actions. Instead the signal induces the victim to alter some aspect of how it collects

information, interprets information, makes decisions or performs actions. Such a change may then affect any step of the victim's OODA loop in the future. To cause misperception, the victim's altered processes then induce it to perform a self-targeted Information Warfare attack.




Figure 4.11: A compound Subversion attack mapped into the OODA loop model. The victim collects the subversive signal during its Observation step and interprets it during its Orientation step, which induces the victim to involuntarily manipulate some process of its OODA loop. In this case, the Action step is manipulated, causing the victim to perform an involuntary self-targeted Information Warfare attack during its Action step.

Subversion can be used to induce a victim to perform self-targeted Degradation, Corruption or Denial attacks, as well as further Subversion attacks. While this discussion only briefly examines the canonical strategy of Subversion, compound Information Warfare strategies and the real-world usage of Subversion within compound Information Warfare strategies, a more detailed analysis of these topics has been undertaken by Kopp (2006a).

A historical example of a Subversion attack against a victim's Orientation step is the Green Ball campaign organised by Edward Bernays in 1934 to promote Lucky Strike cigarettes (Tye, 1998). At the time the green packaging used by Lucky Strike was considered by women to be unfashionable, as they felt the green packaging clashed with their clothes. Since Lucky Strike would not change their packaging, Bernays sought to change the prevailing fashion culture. As part of his campaign, Bernays organised a green-themed event for prominent members of society in New York City, with the underlying aim of changing the public acceptance of green-coloured fashion. Bernays succeeded and

by the end of the year green was much more fashionable and Lucky Strike cigarettes had become a fashion accessory.

In terms of the OODA loop model, the prominent members of society were exposed to Bernays' subversive message during their Observation step, when they attended the ball. During their Orientation step, widespread exposure to the colour green caused the attendees to associate the colour green with fashion. This new belief does not initially change the behaviour of these prominent members during their Decision and Action steps. Instead, during subsequent OODA loop iterations when former attendees of the ball interpret green clothes during their Orientation step, they were more likely to interpret those clothes as fashionable. This more favourable interpretation would then guide their Decision and Action steps, leading them to involuntarily purchase, design and wear green clothes more frequently. As such fashion choices became widespread, green fashion became favoured, thereby addressing the problem of Lucky Strike's packaging clashing with popular fashion.

Subversion's other method for causing misperception is through a compound strategy that induces an involuntary self-targeted Information Warfare attack. The self-targeted attack may implement any of the canonical Information Warfare strategies, although the following examples only describe Degradation and Corruption.

Sony BMG used a compound Subversion and passive Degradation attack against customers who purchased Sony music CDs (Kerr, 2007). Sony placed malicious software onto some of its music CDs, which installed copy-protection software containing a rootkit onto the customers' Windows PCs. The copy-protection software restricts access to the contents of the CDs, using the rootkit to conceal the copy-protection software's files, registry entries and processes from Windows. Concealing the rootkit from the user was a Degradation attack and not a Corruption attack since the effects of the attack (concealing specific files, registry entries and processes) were not perceived by the victim.

Another example of compound Subversion attacks causing misperception is frequently observed in various types of rogue security programs (Symantec Corporation, 2009). Rogue security software is initially installed, either knowingly or unknowingly, onto a victim's computer system as part of a Subversion attack. Through this attack the software then induces the victim's system to produce various fake error messages, indicating that the system may have been compromised by malicious software or otherwise had its performance

degraded in some manner. Figure 4.12 shows an example of a fake error message. These error messages are crafted to mimic genuine warnings from legitimate security software and are a Corruption attack, which aims to convince the victim to purchase fake security software to address the non-existent problem. If the Corruption attack is successful, the victim will purchase further software to correct any perceived problems and then install this software. This software may then perform further attacks against the victim's system, possibly removing or disabling existing security tools or silently installing other malicious programs. In this compound attack, the malicious software uses a Subversion attack to display the fake error messages, effecting a Corruption attack against the victim; which, if successful, may lead to further Subversion attacks.




Figure 4.12: A fake virus warning that attempts to convince a user to install a fake antivirus program. Once installed this program performs a compound Subversion and Corruption attack against the user, displaying further fake errors and warnings that try to convince the user to pay money to repair the non-existent problems.

These examples demonstrate historical examples of Subversion causing misperception in two different manners. In the first, Subversion alters how the victim implements tasks during its Observation or Orientation steps, which then affects the victim during subsequent Observation or Orientation steps and thereby causes misperception. While such attacks resemble either Corruption or Denial attacks — since they lead to errors in the Observation or Orientation step — they differ in that the victim causes the error and not

the attacker. In the second case, Subversion may alter how the victim implements any step of its OODA loop, which then induces an involuntary self-targeted Information Warfare attack. This self-targeted attack causes the victim to misperceive. As such, Subversion can ultimately produce misperception by initially inducing errors in any step of the victim's OODA loop. However, these are considered to be minor errors, while the major errors are instead those that actually cause misperception. In the case of Subversion these major errors may occur during either the Observation step or the Orientation step, due to the victim's manipulated implementation of these steps or from the self-targeted attack produced by Subversion.

Since this thesis is concerned with misperception, instances of Subversion that do not cause misperception will not be examined any further.

4.3 Information Gathering Errors

There are various types of errors that may occur as Information Gathering Errors. Such an error typically leads the affected entity to develop an incorrect understanding of its environment. These errors are not due to the hostile actions of an attacker and are instead caused by an existing flaw or dysfunction of the information sensor, or the processing functions utilised by that sensor. A variety of causes of such flaws are categorised in Section 2.7.1.

Figure 4.13 shows how an Information Gathering Error will affect an entity's OODA loop model. During the victim's Observation step, flaws in its information sensors reduce the quality or quantity of information it collects from the environment. The information that is collected is then forwarded to the Orientation step, where it is analysed and interpreted in terms of the entity's current beliefs. The entity must attempt to determine the current state of the environment, using its existing beliefs and the incomplete or incorrect information. If the victim is aware of its sensor flaws, its analysis methods may be able counteract the effects of these flaws to some degree. Updating its beliefs with information that is possibly incorrect is likely to create false beliefs. These false beliefs are likely to further affect the entity's Decision step, affecting the entity's consideration of its possible options. Finally during its Action step, the entity will perform its selected option,

which was likely dependent upon false beliefs produced by errors the entity's sensor flaws produced.




Figure 4.13: An Information Gathering Error mapped into the OODA loop model. The entity's sensor is unable to correctly gather information, affecting its understanding of its environment and ultimately its actions.

Since no attacker is responsible for the information sensor damage, these errors are completely unintentional. For an example of an Information Processing Error, consider the weather radar a female farmer uses to detect precipitation over his or her farm. However, the farmer has poorly maintained her weather radar, allowing the the physical hardware to deteriorate to the point that it fails completely.

In this instance, the farmer receives no information from the failed weather radar during her Observation step. During her Orientation step, the farmer analyses the lack of information and updates her beliefs to show that she is now unaware of the presence or severity of any precipitation within the area covered by her failed weather radar. She may also be unsure why the radar failed, potentially attributing the failure to malicious acts if she considers such behaviour possible. The farmer then must choose which options she should perform in her Decision step; however, she cannot select any potential actions that require knowledge of precipitation near her farm that the weather radar would have provided. During her Action step the farmer performs her selected Action, which was chosen based upon possibly false beliefs created by incomplete or incorrect information.

An Information Gathering Error consists of three distinct errors. The first occurs during the Observation step, as the entity's sensors gather incomplete information. The second error occurs during the Orientation step, as the entity creates false beliefs from the incomplete information. The third error occurs during the Action step, as the entity performs actions derived from the false beliefs, which are likely to be detrimental. The basic form of this error shares the same characteristics of both Degradation and Denial attacks. As in those cases, the first error is responsible for the two subsequent errors.

4.4 Information Processing Errors

Information Processing Errors occur during an entity's attempts to analyse the newly collected information from the environment. These errors are unintentional and have a variety of potential causes, which are discussed in Section 2.7.2. Such errors in human cognition are extensively catalogued. For example, conservative updating (Edwards, 1968) affects how decision-makers update their beliefs in the face of new information, which may lead them to adjust their beliefs by a lesser magnitude than the new evidence warrants, or ignore it outright. Other biases include those of representativeness, availability and anchoring (Tversky and Kahneman, 1974), which affect how a person assesses the probability of an event occurring or the size of an estimate. Similarly, Gigerenzer (2002) discusses the human propensity for statistical "innumeracy", which he describes as the inability to reason about uncertainty and risk. People commonly categorise individuals, groups and ideas based upon stereotypes, which may or may not accurately represent reality (McGarty, Yzerbyt and Spears, 2002). Heuer (1999) discusses how to identify and prevent such biases, specifically in the field of intelligence analysis, where such biases can (and do) widely influence international relations.

In general, Information Processing Errors occur when an entity collects correct information from the environment but erroneously analyses it, thereby creating false beliefs. These analysis errors represent an error of understanding; the entity has gathered correct information but its analysis has produced the wrong conclusions. These errors are unintentional and produced by the entity's own shortcomings in analysing and interpreting the new information it collects.




Figure 4.14: An Information Processing Error mapped into the OODA loop model. The entity correctly gathers information but is unable to determine its true meaning and instead creates a false belief. This false belief ultimately degrades the entity’s understanding of its environment and thereby affects its Decision and Action steps.

Figure 4.14 maps an Information Processing Error into the OODA loop model of a single entity. This process begins in the entity’s Observation step, where it gathers correct information from its environment. This is passed along to the Orientation step, to be analysed in the context of the entity’s existing beliefs. At this point a processing error occurs, causing the entity to produce one or more false beliefs from its analysis of the new information. The errors responsible for the incorrect interpretation are due to shortcomings of the entity’s analysis methods or existing beliefs. During its Decision step the entity uses its beliefs to assess and select the options it may perform; however, its choices may be affected by the false beliefs caused by the faulty interpretation. The selected option is then performed during the Action step, possibly having effects that the entity did not anticipate.

The Gambler’s fallacy (Piattelli-Palmarini, 1994) is another example of an Information Processing Error, in which a person incorrectly reasons that a small sample is as representative of its population as a large sample. When a person applies this reasoning to the toss of a coin, he or she may interpret a series of previous Heads as evidence that the coin is “due” to come up Tails. Consider how this fallacy affects the decisions of a man playing

roulette, who has observed that the previous ten spins have all come up Red. The gambler first observes the results from the previous roulette games during his Observation step. During his Orientation step, the gambler analyses these results and due to the fallacy he incorrectly reasons that the small sample of results that were predominantly Red indicates that the future spins are much more likely to be Black. This false belief then impacts his Decision step, causing the man to determine that he should bet on Black, since it is a more likely outcome than Red. The man then bets on Black during his Action step. This bet is not in the gambler's favour, as Red and Black are both equally likely outcomes of the next spin.

During an Information Processing Error there are two distinct points in the OODA loop model where errors occur. During the Orientation step the entity incorrectly analyses the new information and its existing beliefs to develop false beliefs. Later, during the Actions step the entity performs actions that were selected due to those false beliefs. The error during the Orientation step is responsible for the subsequent error during the Action step. Information Processing Errors occur exclusively during the Orientation step and affect the entity's assessment of new information, causing it to create false beliefs and thereby degrade its understanding of its environment. These false beliefs may affect both future and current iterations through the OODA loop, just as they do when caused by other errors or Information Warfare attacks. False beliefs may therefore generate further false beliefs, allowing misperception to amplify itself across many OODA loop iterations.

4.5 Self-deception

Self-deception has been previously defined as the intentional dishonest manipulation of an entity's beliefs, which then benefits the entity in some manner. There are two different hypotheses attempting to explain why entities perform Self-deception. One states that a Self-deceiver benefits directly from suppressing unwanted information, while the other argues that Self-deception may be used indirectly as an aid to deception. Since these two methods differ in their aims and implementations, they will be separately mapped into the OODA loop model.

4.5.1 Suppressing Unwanted Information

It has been proposed that Self-deception may be used by entities in a manner akin to a psychological defence mechanism — to conceal or manipulate any unwanted information. Information may be unwanted if its analysis produces beliefs that conflict with existing beliefs that the entity possesses. Cognitive dissonance theory (Festinger, 1957) states that the possession of contradictory beliefs produces anxiety and discomfort in the entity, which is alleviated by removing the dissonance between those beliefs. New beliefs may be unwanted for many reasons. They might identify weaknesses and shortcomings in the entity, which it cannot accept, or disprove fundamental beliefs that are crucial to the entity’s understanding of itself, its fellow entities or its environment.

An entity may use Self-deception to suppress or manipulate contradictory new beliefs that conflict with its existing beliefs. When faced with contradictory beliefs, the entity should attempt to honestly correct its own beliefs if they are known to be false. Instead, Self-deception allows an entity to retain its current beliefs, while rejecting its new beliefs and the new information responsible for those beliefs¹. Self-deception thereby allows the entity to suppress the unwanted beliefs and act as if it never encountered the contradictory information. Andrews (2004) states that “Decisions to ignore key intelligence because it did not fit a reality consistent with a particular Belief System led to the iconic destruction of the World Trade Centre”. In this case, possible Self-deception by the various American intelligence agencies led them to reject gathered intelligence that suggested a future attack, as it conflicted with their beliefs. Had the beliefs created by this intelligence instead been accepted, the attacks may have been thwarted.

Figure 4.15 demonstrates how this type of Self-deception is mapped into the OODA loop model. The entity initially collects new information from its environment during its Observation step. This new information is then passed along to the Orientation step, where it will be analysed, creating new beliefs and updating existing beliefs. However, the new and existing beliefs are recognised to be contradictory, which causes the entity anxiety, as it cannot reconcile the new beliefs and its existing beliefs. To reduce this anxiety, the entity Self-deceptively manipulates its beliefs in some manner and thereby removes the

¹Entities may also be affected by various mechanisms that produce similar effects to Self-deception — such as cognitive conservatism (Edwards, 1968), whereby entities are extremely reluctant to accept new beliefs that contradict currently held beliefs. Determining which mechanism or combination of mechanisms is actually responsible for a given cognitive error requires a detailed analysis of that specific instance.




Figure 4.15: Mapping Self-deception that reduces anxiety into the OODA loop model. The entity Self-deceives to manipulate its own beliefs, which reduces the anxiety caused by possessing contradictory beliefs.

conflict. This manipulation will remove any conflicts, but is likely to produce false beliefs. The entity now enters its Decision step and determines what actions it should take based upon what are likely to be false beliefs. The entity then implements this selection during its Action step. Furthermore, the false beliefs created by Self-deception can affect the entity during all future iterations of its OODA loop.

An entity may manipulate its beliefs in several different ways to remove any conflicts. The entity could erase or conceal the observed information that is responsible for the new conflicting beliefs, preventing others from observing it in the environment. This constitutes denying some aspect of reality and is somewhat similar to the psychological defence mechanism of repression (Ewen, 2003). Alternatively, the entity could alter the gathered information itself, so that its interpretation of this information does not produce contradictory beliefs. This could include acknowledging that an event occurred but instead altering some of its attributes, such as its causes, its outcomes, or its relative importance to those responsible for the event. In effect, this constitutes a spin attack by the entity against itself (Kopp, 2006a). An entity may also alter its interpretation of the new information, producing non-contradictory beliefs that the information does not support. Entities may also choose to ignore certain sources of information that they believe would conflict with

their existing beliefs, preventing information collection from such sources and essentially turning a blind eye to those sources. In any of these instances, the beliefs manipulated by Self-deception during an entity's Orientation step are likely to be false beliefs, since they will not accurately represent reality.

An example of this type of Self-deception performed by an organisation can be found in the case of a nation's government, which wishes to believe that it respects its citizens rights, while also repressing these citizens. To remove any conflicts between these two beliefs, the government may classify any documents that refer to the repression. These classified documents would describe various acts that contradict the image the state wishes to portray both internally to its citizens and externally to other nations. Possible acts that the nation may wish to conceal could include political corruption, espionage, repression and the torture of its citizens. Several historical examples of such Self-deceptive behaviour as performed by nations and organisations have been described by Van Evera (2002) and Kopp (2005a). Another example of such Self-deception can be found in fraud victims, who often attempt to conceal their victim-hood, since it implies gullibility, stupidity or even dishonest intentions if the fraud required the victim to participate in illicit activities. Such a Self-deception can both prevent guilt and shame felt by the victim and help to conceal evidence of the fraud from others, who may criticise and reinforce the negative emotions felt by the victim (Titus, 1999). Interestingly, in certain cases such Self-deception may also aid the deception of those charged with detecting and preventing such fraud. This prevents the Self-deceiving victim from being punished for its acts, the act of which would likely lead them to suffer further from their negative emotions.

In the case of the Self-deceiving government, it perceives reports that indicate abuses during its Observation step. These reports are then forwarded to its Orientation step, where they are analysed and used to update the government's beliefs. The beliefs created from these reports cannot be integrated by the government, as they conflict with the government's existing beliefs. In order to prevent this conflict, one part of the government secretly classifies the reports, effectively concealing them from itself, or at least parts of itself. Classification prevents the documents from being analysed and thereby prevents the development or spread of conflicting beliefs. This allows the government (in general) to maintain the belief that it treats its citizens fairly, despite possessing evidence to the contrary. During the government's subsequent Decision and Action steps, it will select

and perform options that indicate it is unaware of any reports of abuses of its citizens. The government may therefore maintain the desired false belief that it upholds its citizen's rights. In this example, the government's actions function akin to psychological repression; where part of the government, acting as the unconscious mind, manipulates the beliefs that are accessed by the majority of the government, or the conscious mind. Examples of such behaviour can be found in the Ministry of Truth in Orwell's (1949) *Nineteen Eighty-Four* or the Soviet Union's revisionism of its official history (King, 1997), which is explored further in Section 5.6.3.

Self-deception is sometimes observed in fraud victims, preventing from recognising that they have been defrauded and avoiding the painful admission of losing money to their family or friends (Office of Fair Trading, 2006). Instead the victims rationalise their failure to receive their payment from the fraud as being due to some other failure of the fraudulent scheme. When such a victim conceals their fraud with Self-deception, it first gathers information during its Observation step indicating that it has not actually won any money in a lottery and has lost the money it paid as a fee to receive this prize-money. During its Orientation step, the victim's beliefs describing its loss of money conflicts with its beliefs that it will soon receive its prize-money, causing the victim anxiety. Since these beliefs conflict, the victim Self-deceptively manipulates its beliefs to conceal evidence of the fraud from itself, thereby reducing its anxiety. The victim thereby maintains the belief that it is not gullible, avoiding feelings of shame or guilt, and possibly retaining a degree of trust in the fraudster. During the victim's Decision and Action step, its decisions and actions will indicate that it does not believe itself to be a fraud victim. Such actions might include destroying any physical evidence of the fraud, as this evidence conflicts with the victim's desired beliefs. The victim's actions will not include reporting the crime to the police or to its friends or family, since these acts require the victim to acknowledge that it was defrauded.

Two errors are observed when this form of Self-deception is mapped into the OODA loop model. The first occurs during the Orientation step, where the entity dishonestly manipulates its own beliefs to reduce its dissonance. The second occurs during the Action step, as the entity performs actions that have been affected by the false beliefs. The location of these errors match those of a Corruption attack, which supports the existing argument that Self-deception is reflexive deception. Self-deception used in this way can

fail if the Self-deceiver is unable to retain the false belief, possibly because new information conflicts with the belief in such a way that it cannot be easily resolved. At some point the unwanted belief is reintroduced to the Self-deceiver's beliefs and causes the Self-deceiver anxiety that may not be easily remedied.

4.5.2 Aiding Deception

Trivers has proposed that Self-deception can aid the performance of deception, and indeed has evolved to do so, by preventing the transmission of unconscious signals that intimate deceptiveness. These signals could include an entity's tone of voice or other aspects of its behaviour, which may collectively be referred to as an entity's "body language" (Allan and Pease, 2006). During an attempted deception these signals may reveal the attacker's untruthfulness to the victim. The deceiver wishes to conceal these signals if possible, to increase the chance of the deception succeeding. This is an example of multi-channel support for deception, where multiple information sources all send mutually supportive and reinforcing deceptive information to the victim. In terms of the canonical Information Warfare strategies, Self-deception aiding deception consists of a compound strategy of one or more Degradation or Corruption attacks, which are dependent on one or more reflexive Degradation or Corruption attacks.




Figure 4.16: Self-deception aiding Deception mapped into the OODA loop model. The attacker uses Self-deception to believe its own lie, thereby increasing the chance that it will fool its victim.

A deception attack can be modelled by two linked OODA loop models and Figure 4.16 shows how this process may be mapped into the OODA loop model. This example also considers the government that mistreats its citizens, which was discussed previously. However, in this instance the bad government wishes to convince other governments that it does not torture or mistreat its citizens, in order to convince them to sign a trade agreement. The deceiving government implements its Self-deception by concealing reports of abuses against its citizens and releasing internal propaganda that demonstrates to its citizens the various good acts it performs. This is achieved by manipulating its own beliefs during its Orientation step in order to facilitate the creation of the propaganda during its subsequent Action step. These actions aim to convince the government that it treats its citizens well; so that its actions, which are observed by other governments, support this claim.

Now the attacking government can use deceptive propaganda and arguments against the other nations. During its Action step it places further deceptive propaganda that extols its wonderful treatment of its citizens into the world where other nations may perceive it. Along with this propaganda, the other nations may perceive the behaviours of the government, which show it treating its citizens well. Both the statements and actions of the Self-deceiving government are perceived during the other nations' Observation step. During their Orientation step, the statements and actions are compared to the nations' existing beliefs. Since the attacking government appears to be acting truthfully, the nations will believe its statements and update their beliefs accordingly. Since the deception has succeeded, the nations are more likely to sign the trade agreement. During the Decision step the nations will choose to sign the treaty and finally do so during the Action step. After this agreement has been signed, the attacker may remove the Self-deceptive beliefs during a later Orientation step and return to mistreating its citizens. However, in this example, the attacking government may instead wish to retain its Self-deceptive propaganda and combine it with more discrete methods of mistreating its citizens.

In Trivers' description of Self-deception aiding deception, it is assumed that the attacker removes the Self-deceptive beliefs it adopted after the deception has succeeded. However, as in the previous example, the attacker may not need to correct these errors in order to benefit from the deception attack. Self-deception aiding deception may fail in two different ways. In the first case, the deception may fail despite the attacker's supporting

body language. In this case, the intended victim rejects the attacker's information and may become aware of the attacker's hostile intentions. In the second case, the attacker is unable to remove the Self-deceptive beliefs from its model of the world and is unaware of them. The attacker therefore possesses the same false beliefs as its victim, which may render it unable to exploit the victim after the deception attack; or worse, cause it to perform the same potentially self-destructive act as the victim.

In this case the Self-deceiver is only affected by one error, which occurs during its Orientation step as it adopts the false belief. There is no error during the attacker's Action step, as its actions are intended to support the attacker as it communicates the corrupted signal. Therefore, while these actions are impacted by its false beliefs, this impact is intentional.

4.6 Misperception and OODA Loop Tempo

Boyd argued that the speed at which an entity proceeds through its OODA loop is also important to its decision-making process in a competitive environment. Boyd defined the concept of operating "inside an adversary's OODA loop" (Boyd, 1986), as the ability to transition through the loop faster than an opponent. This is intended to cause the faster decision-maker to appear ambiguous and unpredictable to the slower opponent, creating confusion within the slower entity that leads to an inaccurate understanding of reality. This ambiguity that Boyd describes may manifest itself as an inability of the slower entity to fully understand the faster entity's actions, motivations and beliefs. In other words, when a faster entity operates inside an opponent's OODA loop, it can affect the slower entity in a manner similar to misperception. Advantageous examples of operating inside an opponent's OODA loop have been documented by Thompson (1995), describing how businesses may benefit from a faster decision-making tempo than their competitors.

Boyd states that mismatches in decision-making tempo can cause confusion, and potentially even paralysis, within the slower entity. Confusion may reduce the slower entity's tempo further, widening the mismatch in tempo and thereby leading to further confusion. This process can operate as a positive feedback loop, continuing until the slower entity's tempo is so low that it is effectively paralysed by an inability to make timely decisions. The creation of such disparities may or may not be intentional, depending upon whether or

not the faster entity is actively leveraging its faster tempo to induce exploitable behaviours in its slower opponent. Such acts are covert, as an entity is unlikely to recognise their occurrence, source or effects. So far this work has not considered the effects of differing tempi and instead assumed, for simplicity, that both entities iterate through their OODA loops at similar speeds.

4.6.1 Comparing Tempi

The OODA loop models an entity's decision-making process in a discrete manner. As the faster entity becomes ambiguous and unpredictable to the slower entity, it is, in effect, reducing the slower entity's understanding of its environment. This produces a similar effect to a Corruption attack, but through different means. While a Corruption attack directly reduces the accuracy of an opponent's beliefs, operating at a faster tempo instead allows the faster entity to more frequently change its own state and that of the environment, thereby creating a mismatch between the opponent's beliefs and reality. The slower entity's beliefs remain the same, but the physical environment they describe has been altered by the faster entity's actions, thereby causing the slower entity to possess false beliefs. These false beliefs may then affect the slower entity during its future OODA loop cycles, as the false beliefs may prompt the slower entity to act unwisely or even begin to suffer from incestuous amplification.

Figure 4.17 demonstrates how operating at a faster tempo through the OODA loop can appear to cause misperception in the slower entity. This can be better explained with the example of a fight between a cat and a dog, where the cat has a slightly faster OODA loop tempo than the dog. In this example it is assumed that there are no other sources of misperception to further complicate issues and that transitioning through the loop takes a constant time for both animals. Initially both animals start their OODA loops at the same time. The faster cat completes its OODA loop first and performs its first action (F1), altering the state of the environment in some manner, possibly by moving itself. When the slower dog completes its cycle and performs its first action (S1), the environment it is acting upon does not match the environment it believed existed during its earlier Observation and Orientation steps. In this case the cat is not where the dog expected it to be when the dog acts, which would complicate any of the dog's actions that rely upon correct knowledge of the cat's location. The dog is unaware of this difference when




Figure 4.17: Differences in the time it takes competing entities to proceed through their OODA loop can cause effects similar to misperception. Here the repeated interactions of two entities with different iteration speeds are compared. The faster entity completes its Action step before the slower entity, allowing it to change the state of the environment from what the slower entity expects to encounter during its own Action step. Note the use of colour to show how each entity’s actions alter the state of the environment and how this change may occur between the time an entity’s Observation and Action steps, potentially rendering the entity’s beliefs inaccurate.

it finally acts. During the dog’s second Orientation step, it needs to correctly understand the actions previously performed by the cat (F1) and how these actions affected the actions that the dog attempted to perform (S1). However, by the time the dog has updated its beliefs and begun to plan its actions, the cat has already acted again (F2), once again changing the state of the environment. When the slower dog performs its actions (S2), it must again interact with an environment that does not match the one described by its beliefs. This process is continually repeated, where the faster cat’s actions alter the environment from the state that the slower dog expects to affect with its own actions.

In this example, the two OODA loop tempi are fairly similar. However, if the slower entity’s tempo is much slower than its faster opponent’s, the opponent can perform multiple cycles in the time it takes the slower entity to complete a single cycle. While the faster entity is arguably responsible for the slower entity’s confusion, this is only true if the faster entity is aware of the advantages of a faster tempo and it deliberately operates at a higher tempo to confuse the slower opponent. If the faster entity is unaware that its faster decision-making tempo confuses a slower opponent, then the effect cannot be

considered intentional. As well as degrading a slower opponent's decision-making cycle, operating at a faster tempo gives the faster entity more opportunities to update its beliefs to better reflect the true state of the world.

The example also demonstrates that the faster entity can suffer the same problems as the slower entity. Consider the cat's second and third OODA loop cycles, where in both cases the dog's Actions can occur between the cat's Observation and Action. Therefore, in this example, the fast cat can also be affected by the disparity between reality and its outdated beliefs — just like the slow dog.

4.6.2 A Mathematical Examination of OODA Loop Tempo

Boyd's concept of operating inside an opponent's OODA loop can also be demonstrated mathematically with Nyquist's sampling theorem (Nyquist, 1928; Shannon, 1949). According to this theorem, a time variant process (the state of the environment) needs to be sampled at least twice as frequently as its highest frequency component (changes in the environment) in order to capture all the information it contains (Equation 4.1). Here, f_s is the sampling frequency and f_{max} is the highest frequency component.

$$f_s \geq 2f_{max} \quad (4.1)$$

An entity's sampling frequency, f_s , represents how frequently it samples its environment, whereas f_{max} represents the frequency at which its opponent interacts with the environment. If we assume the sampled variable is the state of the environment being acted upon by two entities alone, then the highest frequency at which an environmental change can be effected is determined by the reciprocal of the entity's OODA loop tempo (Equation 4.2).

$$f_{entity} = \frac{1}{t_{OODA}} \quad (4.2)$$

Translated into the terms of the OODA loop model, the slower entity fails to gather all the possible information if the duration between its Observations, or samplings, of the environment do not occur at greater than twice the frequency at which the environment changes, or its highest frequency component. If this condition is not true, then the slower

entity is considered to be losing information, as its environment is changing faster than it can be sampled due to the faster actions of its opponent.

Consider the case where two roughly equivalent competing entities, a and b , both possess identical OODA loop iteration speeds, such that $t_{OODA}[a] = t_{OODA}[b]$.

For a to avoid under-sampling its environment, specifically the actions of b , a must satisfy Nyquist's equation.

$$\begin{aligned} f_s &\geq 2f_{max} \\ f_s[a] &\geq 2f_{max}[b] \\ \frac{1}{t_{OODA}[a]} &\geq \frac{2}{t_{OODA}[b]} \end{aligned}$$

Since a and b share the same OODA loop tempo $t_{OODA}[a] = t_{OODA}[b]$. Substitution gives:

$$\frac{1}{t_{OODA}[a]} \not\geq \frac{2}{t_{OODA}[a]}$$

Since this equality does not hold, a is not fully sampling its environment for b 's actions. Conversely, since a and b are interchangeable:

$$\frac{1}{t_{OODA}[b]} \not\geq \frac{2}{t_{OODA}[b]}$$

Therefore, in the case of two entities with identical OODA loop iteration speeds, both are under-sampling their environment and losing information that describes their opponents' actions.

For a more concrete example, consider the aforementioned fast cat and the slow dog. Given that $t_{OODA}[cat] = 3$ and $t_{OODA}[dog] = 4$:

$$\begin{array}{ll} f_s[cat] \geq 2f_{max}[dog] & f_s[dog] \geq 2f_{max}[cat] \\ \frac{1}{t_{OODA}[cat]} \geq 2\frac{1}{t_{OODA}[dog]} & \frac{1}{t_{OODA}[dog]} \geq 2\frac{1}{t_{OODA}[cat]} \\ \frac{1}{3} \not\geq \frac{2}{4} & \frac{1}{4} \not\geq \frac{2}{3} \end{array}$$

Unsurprisingly the slow dog is under-sampling the fast cat. However, the faster cat is also under-sampling the slower dog.

The amount of information potentially lost due to under-sampling depends upon the relative differences between the two entities' tempi. The less often a slower entity observes its environment, the less opportunities it has to observe the faster entity's actions and the effects of those actions. As a slower entity's tempo decreases relative to a faster entity's tempo, the faster entity can perform more actions that the slower entity cannot observe. Since the slower entity does not observe these actions in a timely manner, or possibly at all, the slower entity potentially loses information due to its under-sampling of the environment.

Consider the example of a car manufacturer whose factory places completed cars in an external parking lot, where they await delivery to the car dealerships. The number and type of parked cars fluctuates quickly as the factory manufactures more cars and completed cars are then delivered. A reasonable operational tempo for this factory might be one hour, with cars entering or leaving the parking lot hourly. A competing car manufacturer regularly sends an employee to count the number of parked cars, so that they may estimate the factory's productivity and sales numbers. If the employee observes the parking lot daily, then he or she only observes the state of the lot once on that day. However during a standard working day of eight hours, the number of cars parked in the lot might change up to eight times. If the competitor instead counts the parked cars weekly, then during business hours the state of the parking lot might change 40 times, with the majority of these changes unobserved. If the competitor only views the parking lot monthly, then the contents of the lot might change 160 times between observations. Each of the manufacturer's actions that isn't observed by the competitor represents lost information, since the competitor is unaware of the changes made to the factory's stock levels. As the difference between the OODA loop tempi of the factory and the competitor increases, the amount of information potentially lost by the competitor due to under-sampling increases.

The under-sampling of information that is identified by Nyquist's theorem in entities with mismatched OODA loop iteration speeds also exists between entities with similar OODA loop iteration speeds. This means that absenting any other misperceptions, the entities' similar operational tempi will lead to the under-sampling of each others' actions, thereby allowing both entities' beliefs to inaccurately describe their environment.

In a military context, the basic level of mutual confusion about an opponent's true state makes such inaccurate beliefs part of the "fog of war" (von Clausewitz, 1969) that,

along with the other sources of misperception, create confusion on the battlefield between opponents. Applying Nyquist's sampling theorem to entities with similar OODA loop tempi thereby provides what has been termed a mathematical proof of Clausewitz (Kopp, 2011), where the equal or similar operating tempo of both entities leads to mutual under-sampling, thereby losing information and contributing to the fog of war. This is of importance for those who collect large amounts of information to help minimise the effects of the fog of war, as the collected information still needs to be analysed and interpreted in a timely manner to avoid under-sampling. If not, then under-sampling may lead the faster opponent to 'lose' vital information to the fog of war, despite its advantages in both tempo and information collection.

4.6.3 Is Under-sampling Misperception?

Both misperception and under-sampling share the same end result, wherein entities possess inaccurate beliefs and these beliefs dictate the performance of actions that are, to some degree, sub-optimal for their environment. While the final results of the two effects are seemingly identical, there are notable differences in how they occur. Misperception occurs due to some procedural error, either intentional or unintentional, initially in the Observation or Orientation steps that ultimately affects the actions that an entity performs. Misperception affects an entity's understanding of its environment, leading the entity to decide upon actions that are likely unsuitable for the environment.

Under-sampling occurs because an entity's OODA loop tempo is too slow for its environment. Like misperception, under-sampling may be intentional or unintentional, although it may be much more difficult to recognise intentionality than with misperception. Under-sampling also may be attributed to errors during the Observation or Orientation step. However, these errors cannot be generalised to one location or the other, except in specific examples. Under-sampling may occur due to the lack of timely information collection during the Observation step. Under-sampling may also occur in the Orientation step if an entity is unable to fully interpret its newly collected information in a timely manner. One or both of these problems may contribute to under-sampling.

Under-sampling is therefore a special instance of either an Information Gathering Error or an Information Processing Error. An entity may be affected by either under-sampling

error; since both occur due to the entity's slow OODA loop tempo, they may occur together. An under-sampling Information Gathering Error occurs when the state of the environment changes too quickly for the entity to fully perceive during its Observation step, due to the entity's slower tempo. An under-sampling Information Processing Error occurs when the entity's slower OODA loop tempo prevents it from interpreting the perceived state of the environment before the environment changes. These mismatches between an entity's perceived and actual environment do represent misperception, albeit in a different form. Therefore, under-sampling may be categorised as a source of misperception, specifically either Information Gathering Errors or Information Processing Errors.

4.7 Summary

Examining misperception in terms of Boyd's OODA loop model has demonstrated how various types of misperception affect entities and how these types of misperception may be produced. Misperception produces errors early in the OODA loop model, producing false beliefs that then affect the subsequent steps of an entity's OODA loop. The flow of information through the OODA loop model demonstrates that an entity's beliefs are highly important to its actions. It has also suggested some general solutions that may allow entities to protect their decision-making processes from some types of misperception. The speed at which an entity proceeds through its OODA loop may also produce effects similar to those of misperception. Further examination of this topic is beyond the scope of this thesis.

4.7.1 Locations of Misperception Errors

Table 4.1 shows the various types of misperception and the location in the OODA loop at which they first cause a major error. With the exception of Subversion, these initial errors occur during either the Observation step or the Orientation step. In the case of Subversion, minor errors affect one step of the OODA loop before the major errors cause misperception during the Observation or Orientation step, either through a manipulated implementation of the victim's Observation or Orientation step or a self-targeted Information Warfare attack that affects either of those steps.

In all cases, if misperception is to create false beliefs, it must do this no later than the Orientation step. These misperceptions are caused either by affecting the entity's perception of its environment or by affecting the beliefs that detail its understanding of its environment. The locations of these errors also match the definitions given earlier for misperception, as either an incorrect instance of perception (an error during Observation) or incorrect belief (produced by an error during Orientation).

Source of Misperception	Major Error Location
Degradation	Observation
Corruption	Orientation
Denial	Observation
Subversion	Observation, Orientation
Information Gathering Error	Observation
Information Processing Error	Orientation
Self-Deception aiding Deception	Orientation
Self-Deception reducing Cognitive Dissonance	Orientation

Table 4.1: The location within the OODA loop model where misperception first induces major errors.

The canonical strategies of Degradation, Denial and Subversion, along with Information Gathering Errors all lead to errors during the Observation step. These errors reduce the quality or quantity of information that the affected entity collects from the environment. With less available information, the entity then creates false beliefs during its Orientation step. The production of these false beliefs are considered an Observation error and not an Orientation error, as the interpretation tasks during the Orientation step are performed correctly while the information gathering tasks of the Observation step are affected by errors. These types of errors also highlight an important relationship between the steps of the OODA loop model. Due to the dependencies between the steps, errors in one step are capable of affecting later steps in both current and subsequent OODA loop iterations. Here errors during the Observation step lead to the creation of false beliefs during the Orientation step and these false beliefs may then affect both the Decision and Action steps.

Corruption attacks, some Subversion attacks, Information Processing Errors and Self-deception all produce errors during the Orientation step, while not causing errors in the Observation step. These errors all create false beliefs, which then go on to affect an entity's Decision step and then its Action step. In these cases the false beliefs an entity develops may arise due to intentional or unintentional means.

Subversion is the odd strategy out, since it aims to manipulate the entity's behaviour instead of the information channel. When this effect is directed at the processes behind the victim's Observation or Orientation steps, then Subversion can cause the victim to misperceive. Similarly, when Subversion induces an entity to perform a self-targeted Information Warfare attack, the self-targeted attack (implementing Degradation, Corruption or Denial) ultimately affects either the entity's Observation or Orientation step.

From the perspective of misperception, the Decision and Action steps of the OODA loop are less interesting than the Observation and Orientation steps. For all of misperception's causes it was observed that no major errors occurred during the Decision and Action steps. Instead, these steps deal with the outcomes of misperception, as it uses an entity's beliefs to select and perform its actions. In all cases, the Decision step operates correctly upon false beliefs; developing actions for the entity that may not benefit the entity. If errors do occur in the Decision or Action step, they do not represent instances of misperception. In the case of Subversion attacks that affect the victim's Decision or Action steps, such attacks will affect the victim's behaviour but will not cause it to misperceive, unless they trigger a self-targeted Information Warfare attack that does so.

The harmful actions performed during a misperceiving entity's Action step do not describe mistakes of the relevant entity during that step, but instead describe the overall product of that OODA loop iteration. In cases of misperception, during the entity's Action step the entity performs some action that it would not have selected to perform had it not been affected by some form of misperception.

4.7.2 Similarities Between Misperception Sources

The procedural examination of misperception's intentional and unintentional sources in terms of the OODA loop model has demonstrated some of the similarities previously described. Table 4.2 displays these similarities, with the columns organised by the cause and the rows organised by the equivalent effects on the same step of the OODA loop.

Degradation, Denial and some Subversion attacks all affect an entity's Observation step initially and are initiated by an external attacker. While Information Gathering Errors also initially affect the Observation step, they are unintentional errors. These three misperception-causing methods have equivalent effects; they all restrict the quality or

Error During	Intentional		Unintentional
	<i>External Cause</i>	<i>Internal Cause</i>	
Observation	Degradation, Denial, Subversion		Inform. Gathering Error
Orientation	Corruption, Subversion	Self-deception	Inform. Processing Error

Table 4.2: Relationships between the various sources of misperception. Sources with equivalent effects are listed in the same row.

quantity of information an entity can gather from its environment, which then causes the entity to develop false beliefs during its Orientation step.

Corruption attacks and some Subversion attacks are similar to Self-deception and Information Processing Errors. These misperception-causing methods are equivalent since they all affect how an entity understands its environment during its Orientation step. However, each of these has a different root cause — Corruption attacks and Subversion attacks are caused by external attackers, Self-deception is initiated by the entity and Information Processing Errors are unintentional.

All of the canonical Information Warfare attacks have matching unintentional sources, which can produce similar errors in an entity. While the unintentional errors can be broadly categorised into those affecting information collection or information processing, Subversion attacks that cause misperception may resemble errors within either category. In general, the various unintentional errors that produce the same outcomes as Subversion may be collectively labelled as incompetence, with respect to the task at hand. For example, incompetence at information collection can produce similar effects to a Subversion attack that targets a victim’s information collection task. Furthermore, Kruger and Dunning (1999) argue “that the skills that engender competence in a particular domain are often the very same skills necessary to evaluate competence in that domain”. As such, an incompetent entity’s assessment of anything related to areas of its incompetence may also be deficient.

This study distinguishes between what are effectively two different types of misperception, based upon which part of the OODA loop is initially affected. One affects the Observation step and may be caused by Degradation attacks, Denial attacks, some Subversion attacks and Information Gathering Errors. The second affects the Orientation step and may be caused by Corruption attacks, some Subversion attacks, Self-deception and Information Gathering Errors. When Subversion attacks implement a self-targeted

variant of the other canonical strategies, they may affect either the Orientation step or Observation step, depending upon which of the canonical strategies they implement. This dichotomy between the misperception-causing methods concurs with the earlier definition of misperception as either an incorrect act of perception (errors during Observation) or the production of false beliefs while interpreting information (errors during Orientation).

4.7.3 OODA Loop Tempo

An opponent with a slower OODA loop lacks sufficient information to correctly understand a faster opponent's behaviour, causing it to unwittingly possess false beliefs since it lacks sufficient timely and correct information about its opponent. Such beliefs become false through no error of the entity; yet they become outdated due to changes in the environment between the entity's collection of information and its performance of actions based upon that information. Such under-sampling is different to misperception, as it is not produced by errors within the perceiving entity.

When an entity's OODA loop tempo is considered in terms of Nyquist's sampling theorem, it suggests that an entity must have a much higher OODA loop tempo than its opponent to avoid under-sampling. More specifically, the faster entity's OODA loop tempo should be more than twice that of its opponent, in order to gather information frequently enough to avoid missing information about the opponent's behaviour.

4.7.4 Information Flow Through the OODA Loop

Examining the flow of information through the OODA loop model reveals how information moves through an entity and how it is operated upon during this cycle. Information flows into an entity through its Observation step, then into its Orientation step, where it is analysed to produce new or updated beliefs. An entity's beliefs in turn affect the flow of information into an entity, by directing where and how it gathers information during its Observation step and by affecting how the information is analysed during its Orientation step. An entity may develop false beliefs due to errors that occurred during its Observation or Orientation steps. False beliefs may only be corrected during the Orientation step, when an entity has analysed new information to develop a belief to replace the false belief.

An entity's beliefs act as a store of analysed information that is updated during each iteration of the OODA loop. Misperception will cause an entity to develop false beliefs.

Since an entity will retain false beliefs until they are corrected and its beliefs can affect both the collection and interpretation of new information, an entity's false beliefs can operate in a self-reinforcing manner, producing a positive feedback loop. Given an incorrect understanding of its environment, an entity may use its extant false beliefs to produce more false beliefs during its Orientation step. Since the Orientation step is the only place where these errors may be corrected, this process is unlikely to stop unless the entity becomes aware that it possesses false beliefs and begins to make corrections. This process has previously been identified as incestuous amplification and it is perpetuated by the continued feedback from an entity's false beliefs.

Information also flows out of entities into their environment, through the Actions they perform, which are governed by their beliefs and Decisions. An entity's actions may convey information about its existence, location, state or behaviour. Since an entity's actions are a product of its beliefs, its actions will also convey some information about an entity's beliefs.

4.7.5 Protecting Against Misperception

Misperception, whether intentional or unintentional, produces false beliefs and its effects may therefore be lessened or prevented by avoiding the creation or possession of false beliefs. This can be achieved by correcting false beliefs when they are identified and by preventing misperception from introducing new false beliefs whenever possible. Due to the generic model of misperception examined, these solutions may be applicable to a broad variety of entities. However, some or all of these proposed solutions may be impossible for some entities to implement.

Errors during the Observation step decrease the quality and quantity of information received. Entities may prevent or mitigate such misperceptions by performing various actions that increase the quality and quantity of information. The effects of damaged or destroyed information sensors can be mitigated by redundant information sensors, which will allow information gathering to continue after some sensors are affected. An entity may also benefit from improving the quality of its information sensors. If an entity is aware that its information sensors are flawed or damaged, then it should replace them. Entities should also attempt to protect the integrity of their information sensors, as some sensors cannot be easily replaced or duplicated.

Misperception may also cause errors during the Orientation step, which then produce false beliefs. Entities may prevent the development of false beliefs by establishing the veracity of new information, especially if it contradicts their existing beliefs. They should also attempt to gather further information to help determine whether this new information is truthful, deceptive, simple noise or the product of another entity's misperceptions. The trustworthiness of information sources should also be considered when assessing the truthfulness of their information. Such scepticism may protect an entity from intentional or unintentional Corruption attacks, or from some Information Processing Errors. Entities should also be aware that their beliefs may be false and they should not assume themselves to be infallible. Self-deception is more difficult to protect against, as an entity will likely be unaware of its Self-deceptive behaviour. Instead, other friendly entities may be required to identify potential acts of Self-deception in each other and then warn the affected entity. Such behaviour is seen in the academic peer-review system.

When entities can recognise that their operational tempo may be leading to under-sampling of their environment, there are several possible solutions they may employ to minimise the effects of outdated information. The obvious solution is for the entity to increase its tempo if possible, thereby attempting to avoid under-sampling. However, this is only a suitable solution if the entity can increase its decision-making tempo without causing other errors in its decision-making process. A better potential solution is the selection of actions that emphasise the potential unreliability of the entity's current beliefs. This approach is effectively risk management, with the entity assessing potential problems and detrimental outcomes to determine how the problems posed by such events may be mitigated or avoided. Such actions should emphasise flexibility and adaptability to an environment that may potentially change in a rapid and unexpected manner during the entity's decision-making process. The chosen actions wherever possible should be easily changed, reversed or otherwise altered to adapt to the environment as it is encountered during the entity's Action step, instead of what was expected during the Orientation and Decision steps.

4.7.6 The Importance of the Orientation Step

The misperceptions produced by errors during the Orientation step are represented in the OODA loop as false beliefs. These false beliefs may affect all the steps of an entity's

OODA loop. The Orientation step is responsible for the maintenance of the entity's beliefs, performing both the analysis of new information to update its beliefs and the synthesis of new beliefs from existing beliefs. Due to these roles, an entity creates and maintains false beliefs in its Orientation step. Since all aspects of an entity's decision-making cycle are affected by its beliefs and these beliefs are managed during the Orientation step, the Orientation step is arguably the most crucial part of the OODA loop model.

Since the Orientation step manages an entity's beliefs, any attempts to create false beliefs must either occur prior to this step or during it. This places such attempts in either the Observation step or the Orientation step of the OODA loop. The Observation step and the potential attacks against it are fairly straightforward, while the Orientation step is much more complex. Errors that occur during the Orientation step are caused by either existing false beliefs or flaws with the analysis methods used to produce new beliefs. Due to the importance of the Orientation step, a model of the internal processes of the Orientation step is a crucial tool in further analysing the errors that cause misperception during the Orientation step. This is undertaken in the next chapter.

Chapter 5

Misperception and Orientation

The Orientation step was previously described as the most important step of the OODA loop, due its role in maintaining an entity's stored beliefs. Therefore, misperception that affects the Orientation step may create false beliefs, which then affect other elements of the OODA loop model. This chapter further examines the Orientation step of the OODA loop, since an entity's beliefs affect all the steps of the loop. Boyd did not specifically describe the internal processes of the Orientation step, possibly to avoid creating a model that was overly specific. However, it is possible to produce a general model of the internal processes of the Orientation step that remains applicable to a variety of entities.

The previous chapter examined a simplified model of the OODA loop and focused on how it was affected by misperception. This chapter develops a model of the underlying processes within the Orientation step and then determines how misperception affects these processes. However, in order to do so it is first necessary to examine the underlying cognitive architecture implicitly described by the OODA loop model and relate this to existing models.

The internal processes of the Orientation step will be examined and then used to model the internal procedure and flow of information within the Orientation step. The Orientation step was previously shown to be the location where Corruption attacks, Information Processing Errors and Self-deception all may cause errors in the entity's decision-making cycle. The expanded model that is developed will then be tested and validated by examining how these errors affect the newly developed model. Historical case studies will then be mapped into the developed model, to demonstrate how documented decision-making errors affected the decision-maker's OODA loops. Where suitable, these case studies and

examples will refer to examples previously described, in order to reinforce the effects of these studies through the presentation of a familiar example.

5.1 Revisiting Boyd's Model

Boyd's OODA loop model presents a discrete cyclical model of the decision-making process (Figure 5.1). It describes the repeated process of information collection, analysis, decision-making and action performed by an entity in a competitive environment. Learning is achieved through internal feedback within the steps and iterations of the loop, where the results of previous actions can inform present and future decisions. While the model was originally developed specifically to model the decision-making processes of fighter pilots during aerial combat, it is sufficiently generic to model the decision-making process of any biological entity, mechanical system, software agent, or human organisation. The OODA loop model assumes that the decision-making entity is in some competitive environment or situation, competing against others for its survival; however it has a much wider application under the commonly accepted analogy of 'Nature as Opponent'.

The OODA loop model is a decision-making process that attempts to understand its environment. Testing the results of this process provides an "analytical/synthetic feedback loop for comprehending, shaping, and adapting to that world." (Boyd, 1992)




Figure 5.1: Boyd's OODA (Observation Orientation Decision Action) loop model (Boyd, 1986, 1996). Note the feedback from present Decisions and Actions to future Observations, as well as the control that Orientation exerts on Observation and Action.

An entity begins an iteration of its OODA loop in the **Observation** step, where it collects information from its environment. Note that here information is used in a broader and less precise manner, to encapsulate both the data and noise within the various signals received from the environment. This information is obtained from its own interactions with the environment, unfolding circumstances and other outside information. This collection of information is directed by the entity's knowledge, as shown by the control feedback from the Orientation step.

The collected information is fed forward to the **Orientation** step, where the entity interprets the information in the terms of its existing knowledge, beliefs and analysis methods. The collected information is also used to develop new beliefs that further update the entity's understanding of its environment. The beliefs that are developed and refined during the Orientation step are able to direct the subsequent behaviour of the Observation and Action steps, controlling some elements of how the entity performs these tasks; hence the importance of the Orientation step.

It should also be noted that in some instances there may be a functional overlap between the Observation and Orientation steps, when an entity's information sensor performs some processing of the newly gathered information before passing it on to the Orientation step for further analysis and processing. For example, in an electronic system, the information sensor may process and filter newly gathered information as part of its Observation step. It then passes this processed information along for further analysis during the Orientation step.

An entity selects between a set of possible actions during its **Decision** step, based upon some consideration of the outcomes these actions are expected to produce. This consideration is based upon the entity's stored knowledge and understanding, to select a course of action for the entity that it finds satisfactory.

During the **Action** step, the entity performs its chosen actions, using its existing knowledge to govern how those actions are performed. An entity's potential actions are limited by both its physical or simulated body, its beliefs of its own capabilities and the state of its environment. For example, consider the differences in routine car maintenance performed by a trained mechanic and an untrained person, given their different knowledge and skills. A trained mechanic will complete the maintenance faster and to a higher

standard than an untrained person, given their superior knowledge and training for the task.

The Orientation, Decision and Action steps all provide feedback to the Observation step, with the aim of directing the entity's information collection during its next Observation step. As an example, if an agent's Orientation step identifies some operational shortcoming in the preceding Observation step, then its feedback can direct the subsequent Observation step to address this shortcoming.

Boyd has also identified similarities between the OODA loop model and the scientific method. Both models define an iterative process where a decision-maker repeatedly interprets new information, develops hypotheses and then tests these hypotheses through experimentation — the results of which provide further information for interpretation. This is not surprising; both models effectively describe an iterative learning system, whereby an entity can develop and refine an understanding of its environment through the repeated application of rational thought during interactions with its environment. In the OODA loop model, the Decision step represents the point at which hypotheses are formed from existing knowledge and beliefs. These are then tested in the Action step, when the entity interacts with its environment, thereby performing experiments. The outcomes of this interaction provide new information, in the form of experimental results, that is then collected during the Observation step. The Orientation step then interprets and analyses the experimental results, updating the stored knowledge, which is then used to generate further hypotheses.

As discussed in the previous chapter, the speed, or tempo, at which an entity proceeds through its OODA loop is highly important for competitive entities (Section 4.6). Since this chapter focuses solely on the Orientation step, OODA loop tempo will not be explored in this chapter.

Despite being developed to model rational decision-making, the OODA loop model does not impose rationality upon the decision-making entity. The OODA loop model only specifies the structure of the decision-making process; it does not enforce “correctness” upon the entity's knowledge and beliefs or rationality on the entity's decisions, where rationality could depend on either the true state of the entity's environment, or merely its understanding of the environment. It also does not specify what criteria entities use when deciding between various possible actions available to them.

The Orientation step consists of the interactions and relationships between five elements of the decision-making entity: **Cultural Traditions**, **Genetic Heritage**, **Analysis and Synthesis**, **New Information** and **Previous Experience**. These elements encompass the various types of knowledge, beliefs, skills and cognitive processes that represent the entity's understanding of its environment, and also allow it to further analyse and interpret new information from its environment. The Orientation step allows the analysis and synthesis of new and old information, both collected from the environment, developed from previous analyses and produced by cultural and genetic heritage. Collectively these elements represent an entity's beliefs, essentially its understanding of its own state and the state of its environment.

Cultural Traditions represents the various ideas, methods and beliefs that the entity has derived from its culture. This details what an entity's society or predecessors have learned about the environment and how they typically interact with the environment and each other.

Genetic Heritage represents the knowledge, methods and skills that are incorporated in the entity's genotype, such as its brain and associated psychology.

Analysis and Synthesis represents methods for pulling things apart (Analysis) and then putting them back together (Synthesis) in new combinations, to identify relationships between seemingly unrelated ideas and actions (Boyd, 1987). This process develops new concepts that can further the entity's understanding of its environment or provide insights that yield a competitive advantage.

New Information is the newly gathered information collected from the environment by an entity, which will be examined and analysed to refine the entity's understanding of the current state of its environment. The new information details the observed outcomes of the entity's previous actions, observations of the environment and observations of the outcomes of competing entities' actions.

Previous Experience represents a store of collected knowledge of previous actions, outcomes, phenomena and their relationships to each other. This details what the entity has learned about itself, its competitors and its environment.

While the elements of Cultural Traditions and Genetic Heritage appear to limit the applicability of the OODA loop to biological entities — or even more specifically humans and human organisations — machine systems and software agents can still be understood in terms of this model. For such systems, Cultural Traditions and Genetic Heritage may instead represent functions of the systems, artefacts of the design and programming of the systems, as well as some of the Cultural Traditions and Genetic Heritage of the systems' creators and designers.

Boyd states that Genetic Heritage, Cultural Traditions and Previous Experience provide an “implicit repertoire of psychophysical skills shaped by environments and changes that have been previously experienced” (Boyd, 1996). Analysis and Synthesis, along with New Information, are required to evolve new repertoires to deal with unfamiliar phenomena or unforeseen change. This allows the entity to learn from and adapt to new or unexpected experiences, which will then become part of its Previous Experience. The relationships between the elements are required for the analysis and synthesis to take place.

Analysis and Synthesis allow the creation of new and novel ideas by analysing existing ideas and then using the insights from this process to synthesise new ideas. Boyd (1987) clarifies this concept with an example of designing a snowmobile; analysing four existing ideas and then combining elements of these ideas in a novel manner. The four ideas are a skier on a mountain, a bicycle, a boat with an outboard motor and a tank or other tracked vehicle. Analysing these ideas breaks each of them down into their main functional components, of which Boyd selects one from each idea. These are the skier's skis, the bicycle's handlebars, the boat's motor and the tank's treads. Combining these specific elements in a new form synthesises the design of a snowmobile.

5.2 A Derivative Cognitive Architecture

As discussed previously (Section 2.6), various fields have analysed the process of rational decision-making, producing numerous methods that can be used to represent and analyse the decision-making methods. One widely used model of cognitive architecture is the agent model, which is widely used in the study of Artificial Intelligence and Artificial Life. Since the agent model is widely used in various fields it is worth exploring it in further detail and comparing it to the OODA loop model.

The agent model of cognitive architecture details the decision-making processes of simple and complex agents. An agent is simply an entity that collects information with its sensors, uses this information to decide what it should do and then implements its chosen actions with its actuators. This process is continually repeated by the agent, with the agent's sensors exposed to **stimuli** that represent events and phenomena within the environment. Information collected by an agent's sensors is called a **percept** and an agent's decision-making process acts to map percepts onto suitable behaviours for the agent.

Russell and Norvig (2009) define several main types of cognitive architectures used to design intelligent agents: simple and model-based reflex agents, goal-based agents, utility-based agents and learning agents.

Reflex agents are one of the simplest types of agents, in that their behaviour is governed by a set of rules. These rules describe a Boolean function of stimuli and its associated action. When the agent's sensors detect a recognised stimulus, the programmed actions are performed by the agent. The action selected by a reflex agent is determined by its current percept and (possibly) a model containing a representation of the state of its environment. This model is developed and maintained from the agent's interpretations of its current and previous percepts.

Goal-based agents maintain a model of the state of the environment and decide which actions they should perform based upon the expected results of actions that will achieve their stored goals. Percepts are used by a goal-based agent to model the state of its environment, which is used to predict what the likely effects of the agent's possible actions will be. These possible actions can then be assessed to determine which actions will satisfy the agent's goals. Some goals may not be easily satisfied and may require a complex series of actions to satisfy them. In such cases, the agent will need to possess and utilise searching and planning methods that can develop a sequence of actions that will achieve the agent's goal.

Utility-based agents base their decisions upon the agent's preference for one potential outcome over others, where the agent assigns each outcome a corresponding numerical value representing the outcome's degree of usefulness to the agent. Compared to goal-based agents, where a potential sequence of actions may or may not satisfy an agent's goals, utility-based agents instead compare their preferences for various possible outcomes

and select actions with higher utilities. A utility function allows an agent to deal with conflicting goals, since the utility function specifies an appropriate trade-off between the goals. Furthermore, in situations where it is uncertain whether any of the agent's goals can be satisfied, utility allows the likelihood of success to be considered, along with the importance of the goals. A utility-based agent can therefore select a set of actions to maximise its expected utility when it is uncertain of the state of its environment and possesses conflicting goals.

A **learning agent** uses percepts to make predictions about its environment, which it then tests through its actions. Newer percepts detail the results of previous experiments, which are used to update the agent's knowledge and thereby guide its future actions. A learning agent consists of four conceptual components: a **performance element**, a **learning element**, a **critic** and a **problem generator**. The performance element selects the agent's external actions, interprets percepts and passes this knowledge to the learning element. The learning element makes improvements to the agent's performance element, using feedback from the critic. The critic tells the learning element how effective the agent's actions are, given an external performance standard. As percepts do not communicate the success of an agent's actions, the critic instead uses the performance standard to interpret success from the percepts. The problem generator suggests new exploratory behaviour, which may produce actions that are sub-optimal in the short-term, but lead to the development of better actions in the future. As a learning agent interacts with its environment, it will learn from those interactions and then modify its behaviour over time.

These agent types may also be combined to form hybrids. Learning, in particular, can be added to reflexive, goal-based or utility-based agents, with learning processes improving their reflex rules, goals or utilities respectively. This allows agents to better adapt to changes in their environments, or to better deal with unexpected phenomena. Simple reflex rules can be added to other types of agents, providing fast and simple heuristics that can produce suitable timely responses to a given stimulus faster than other cognitive methods. For example, the recognition heuristic (Goldstein and Gigerenzer, 2002) shows how decision-makers with limited information can make more accurate decisions than those with more information in certain circumstances. Similarly, Gladwell (2005) has explored the benefits of trained and learned reflexive behaviours in humans, as well as their potential problems and shortcomings.

Agents are designed to make decisions that are rational, which means that the agent will always choose actions that it believes will maximise its utility. However, such rational selection is not always observed in humans (Tversky and Kahneman, 1988). Simon (1957) has described human behaviour that “satisfices”: where the decision-maker selects an action that leads to an outcome that is merely satisfactory, but not optimal. Satisficing behaviour can also be utilised in designed agents, assuming that the designer finds the trade-offs involving optimal behaviour and satisfactory behaviour suitable.

As mentioned above, the agents are all assumed to possess perfect information about their environment and they are all capable of making decisions in this certain environment. But how can they do so in an uncertain environment? And, how can such uncertainty be represented? In most non-trivial scenarios, uncertainty is inevitable, resulting from both internal and external factors. Ramsey (1926) stated that this uncertainty in beliefs can be represented with probabilities. This approach allows an agent who cannot ascertain the definite truth or falsity of a fact to instead use the probability of this fact being true as its belief.

Utility theory allows outcomes to be ranked by their expected benefit to the agent, with the agent preferring outcomes that yield a higher utility. Given a set of possible actions and their expected outcomes, utility theory indicates which actions the agent should perform to maximise its benefit.

To select between uncertain outcomes, such as those typically involved in games of chance, utility theory can be combined with probability theory. Probabilities allow for uncertainty, such as that introduced by random or unpredictable phenomena, to be incorporated into utility theory, allowing agents to assess the utility of uncertain outcomes (Morgenstern and von Neumann, 1947; Raiffa, 1968). This uses an expected utility, which is the product of an event’s probability and its utility. Expected utility allows agents to make decisions when they are uncertain as to whether or not a given event will occur, but possess knowledge of the likelihood that it will occur.

The agent model provides a cognitive architecture whose underlying elements are built upon a solid mathematical foundation that was derived from efforts to analyse and create rational decision-making agents. When applied to human decision-makers these efforts are not descriptive — in that they describe how a person ought to act, and do not always predict how a person will act.

5.3 Comparing these Architectures

Despite the structural similarities imposed by being feedback loops for interactive decision-making, these two cognitive architectures set forth two very different models of the decision-making process. This arises due to their differing sources and purposes. The OODA loop model was developed to model the decision-making process of a fighter-pilot in a time critical competitive environment, and then later generalised to any entity in an environment where timely actions are highly desirable. As a result, it is a generalised model that may be adapted to represent most decision-making entities. The agent model is even more abstract and is intended to model the decision-making processes of an abstracted rational entity, which could be a person, a simulated intelligent agent or an animal. This allows its usage in fields such as Artificial Intelligence, Artificial Life, cognitive psychology, biology and economics. Its generality has allowed for much research into how such entities make decisions, including the design and implementation of systems based on this model that can make decisions. The agent model has been developed from research into decision-making in a number of fields, and as such is flexible and descriptive in the behaviours and entities it can model.

While both models are abstract, the agent model is more abstract than the OODA loop model. Indeed, given the highly abstract nature of the agent model, the OODA loop model can be considered as a specialised instantiation of the agent model, where the steps of the loop specify the internal activities of the agents. These activities of the OODA loop could be implemented with the various elements from learning agents, goal-based agents, utility-based agents or even reflex agents.

For example, within the Orientation step, the Synthesis and Analysis element is similar in function to a learning agent's problem generator, given that both are responsible for introducing novelty. To achieve this, both models use the decision-maker's existing knowledge to develop exploratory actions whose results are intended to provide new information for analysis in future decision-making iterations. The critic and learning element of a learning agent would also be subsumed into the Synthesis and Analysis element of the Orientation step, although these functions could also be partly performed by the Cultural Traditions and Genetic Heritage elements. An agent's newly collected percepts are equivalent to the Orientation step's New Information element.

Both models treat the concept of decision-making tempo differently. The OODA loop model explicitly focuses on the idea of tempo and the competitive advantages of introducing and maintaining mismatches between an entity's decision-making tempo and that of its opponent. While decision-making tempo and competitive advantage are also important to the agent model, tempo is typically not its primary focus, since it is more general. Instead, the agent model tends to focus more upon the agent's decision-making process, and attempts at improving the quality of an agent's decisions.

Neither model requires that the entities modelled are rational. Agents may or may not make rational decisions, depending upon the underlying actors and behaviours that they are modelling. While the OODA loop model does not explicitly mention rationality, its operators are implicitly assumed to make mostly rational decisions, given the assumption that their understanding of their environment is fairly accurate. At the same time, operating at a faster tempo is intended to indirectly disrupt and prevent an opponent's rational thought.

Comparing these two cognitive architectures reveals that their major differences are due to their differing purposes. While both models can be used to model and examine the processes of rational (or irrational) decision-making, there is an apparent dichotomy in their typical usage. The agent model typically focuses on modelling and examining the processes of decision-making, often to better explain human decision-making or create artificial agents capable of making rational decisions. In contrast, the OODA loop model examines the decision-making process in order to better understand and facilitate the disruption and prevention of an opponent's rational thought, which can then be exploited for a competitive advantage.

The various types of agents and their design elements provide some insight into the actions that may constitute the Orientation step's tasks. A learning agent's critic assesses the effectiveness of the agent's current behaviour and suggests that the Orientation step should be capable of testing whether the entity's previous and current actions succeed or fail. Previous errors or sub-optimal behaviours may also be corrected, assuming that they can be correctly identified by the entity. The problem generator suggests that some mechanism is required to suggest new behaviours, which would partially represent the Synthesis already undertaken by Orientation step. The mechanism by which utility-based agents calculate and assign value to possible future outcomes could also be utilised during

the Orientation step. This would provide an entity to assess the perceived benefit of any new behaviours it has conceived.

While the agent model cognitive architecture is preferred in many fields of research due to its strong mathematical foundations, the OODA loop model is an inherently better tool for exploring the various flaws and attacks against the decision-making process and the effects of such actions in competitive environments. However, the developed designs of agents will be useful in clarifying the processes of the Orientation step.

5.4 Expanding the Orientation Step

Now I develop a cognitive model derived from Boyd's OODA loop model, specifically detailing the internal processes of the Orientation step.

Boyd's description of the Orientation step details a combination of new information, cultural traditions, previous experience and genetic heritage that are combined through a process of analysis and synthesis. Collectively these five elements describe an entity's beliefs and the mechanisms used to maintain them. This procedure somehow updates an entity's understanding of the current state of its environment. One method for determining the internal process of the Orientation step is to apply each of these elements in turn to any newly gathered information. However, this procedural approach assumes that there should be some ordering of these elements during this process. Any ordering may, however, require making some assumptions about the relative relationships of these elements. Furthermore, assuming an ordering for these elements may remove the generality of the OODA loop model, as different types of entities will likely find different elements to be more or less important than others. For example, while genetic heritage and cultural traditions may be largely irrelevant to a mechanical system, they are highly important to most biological organisms.

Another method of expanding the Orientation step is to list all the functions that are performed during Orientation and then group the topically related functions into sub-steps. These sub-steps can then be ordered chronologically to provide a sequence that describes the internal process of the Orientation step. Once the sub-steps have been listed, the flows of information between them can be identified. Identifying the flow of information through the Orientation step will demonstrate how the entity processes its collected information

and where errors may occur during this process. If all types of entities can perform the tasks used to develop the sub-steps then this approach will produce a general model of the Orientation step.

5.4.1 Functions of the Orientation Step

Breaking the Orientation step down into its constituent sub-steps first requires identifying the functions that are performed during its execution. These are then grouped to produce the sub-steps of the Orientation step. The Orientation step has two main tasks, which are:

1. Updating an entity's internal representation of its environment, as described in its beliefs.
2. Updating and maintaining the overall strategy that governs the entity's behaviour, based upon its current understanding of its environment.

The Orientation step begins with arrival of newly gathered information from the Observation step, which is ready for analysis and interpretation. During the Orientation step an entity must take its recently collected information and determine what this information means, understanding the information in terms of its existing beliefs. Once an entity has an understanding of the current state of the environment, it needs to determine what it may do in the future and how this will likely affect the environment. This knowledge provides the input for the entity's Decision step, where it selects between the potential outcomes it anticipates.

Updating the Internal Representation

The Orientation step's first function is to update an entity's internal representation of its environment, using the newly collected information from the Observation step. Doing so firstly requires the entity to examine its newly gathered information and determine what it has actually observed. This act of identification or classification allows the entity to use its existing knowledge and understanding to recognise the various elements of its environment that it may have just observed.

There are many possible elements of the environment that an entity may observe during its Observation step. These elements can be broadly categorised as either objects or

phenomena (processes) and from this point these terms will be used. During its Orientation step, an entity needs to determine what objects and phenomena it has just observed, identifying and classifying the objects and phenomena based upon its existing beliefs. Entities may observe the attributes of objects and phenomena, as well as the relationships between them. Furthermore, if an object or phenomenon is similar to something that has been previously observed, the entity should recognise this similarity. An entity needs to be able to remember previously observed objects and phenomena along with their inter-relationships. This information is remembered by the entity as it updates its beliefs and creates new beliefs. The entity's newly created beliefs may not accurately describe its environment and they may need to be updated in the future as the entity's understanding improves.

The second function of updating the entity's understanding of its environment is the analysis and interpretation of the entity's new information. An entity is able to determine the implications of its gathered information, thereby understanding the meaning of such information. The entity will use analysis and synthesis methods from its beliefs to understand the new information and determine its implications. These could include determining what caused a phenomenon to occur or determining how an object came to be in its present location or state. Such methods rely upon the relationships an entity believes to exist between identified objects and phenomena. Such assessment may also lead to the production of new beliefs, with entities using the interpretations of new information to produce new beliefs and update existing beliefs. These may then be used in present and future iterations of the OODA loop both inside and outside the Orientation step.

Once an entity has interpreted its new information, the implications of this new information and finished updating its beliefs, the entity will now possess an updated representation of what it believes is the current state of its environment. With an updated understanding of its environment, an entity can now begin to plan its next behaviours.

Updating the Strategy

The second major activity that an entity performs during its Orientation step is the continued development of the strategy that ultimately drives its Actions. Such a strategy requires an entity to first determine what it wants or needs to do or achieve, leading to

aims or goals that it should strive to implement. An entity can then develop potential actions that it believes may achieve its aims.

An entity's behaviour is governed by some form of reasoning, in which its beliefs causally determine its actions. Internally, an entity may set aims that direct its behaviour in some manner to perform some task that it believes is beneficial or necessary. An entity needs to be able to produce new aims, which may be either short or long-term goals that it should strive to complete. Entities must also assess these aims, to determine whether their actions are achieving their aims or not. Aims that become impossible or infeasible will also need to be removed or corrected. The creation and assessment of these aims also depends upon the entity's beliefs, as an incorrect assessment of the environment may lead the entity to develop impossible or foolish aims or reject aims falsely understood to be impossible. In a learning agent the critic element performs a similar function, assessing the effectiveness of the agent's actions.

An entity needs to be able to select actions that will enable it to achieve its aims. Therefore an entity needs to formulate potential actions and then predict their expected results. Developing new actions depends heavily upon the entity's beliefs; requiring the entity to predict the future state of its environment, given its current understanding of its environment, its own proposed actions and the potential effects of these actions. These potential actions and their expected outcomes are the final output of the Orientation step and are provided to the Decision step. The problem generator of a learning agent performs a similar function, in that it suggests new behaviours to explore an environment.

In order to assess its potential actions, an entity needs to know the likely consequences of those actions, which requires it to possess some type of predictive ability. An entity possesses and develops some expectation of the future state of its environment, based upon its existing beliefs. An entity develops expectations from its existing beliefs, which allow it to extrapolate likely future events and the possible environmental changes from these events. However, the entity may become confused if its expectations greatly differ from its observed reality, possibly producing cognitive dissonance. For example, a company could find that its trade secrets have been discovered by its competitors, which could imply that someone inside the company is untrustworthy and has leaked this information. Such a belief would conflict with the company's belief in the trustworthiness of its employees. This conflict of beliefs could be a source of cognitive dissonance and require the entity

to accommodate the differences between the expected and observed environment, either through Self-deception or some other mechanism. This process partially overlaps the tasks of a learning agent's performance element and learning element, in that it interprets new information and updates the agent's current beliefs.

Many different types of entities may perform these actions. However, they do assume some cognitive complexity on the part of the entity; which simpler entities, such as some machine systems or biological organisms, may not possess. Such entities may therefore be incapable of tasks such as prediction and aims manipulation.

5.4.2 Combining these Functions

Breaking down the necessary functions of the Orientation step reveals two main tasks, each of which consists of two minor tasks. When ordered procedurally, these four activities model the internal processes of the Orientation step. These four activities, or sub-steps, are Identification, Interpretation, Aims Derivation and Options Generation.

1. **Identification** encapsulates recognising known objects, phenomena and their relationships from the new information gathered from the environment. This then allows an entity to retrieve and employ existing beliefs that are linked to the identified elements. As noted earlier, in some cases such Identification occurs earlier within the Observation step.
2. **Interpretation** is the analysis of the new information and its identified constituent elements with known processing methods. This allows an entity to produce new beliefs and update its existing beliefs, which represent its understanding of the present state of its environment. An entity's beliefs affect all the steps of its OODA loop cycle and more accurate beliefs allow an entity to better predict the future state of the environment.
3. **Aims Derivation** creates new aims and maintains existing aims, based upon an entity's beliefs. If an entity successfully achieves an aim it is removed. Likewise, if an entity finds an aim to be impossible or infeasible, it is also removed. New aims are created to guide the entity's actions, based upon its understanding of its environment. An entity's aims are stored and manipulated in much the same manner as its beliefs.

4. **Options Generation** determines the potential options that will allow the entity to achieve its aims and then predicts the expected outcomes of these options, based upon the entity's beliefs. The entity stores these options and outcomes as beliefs for evaluation during the Decision step.

It is important to realise that all of these Orientation sub-steps operate within the context of the entity's understanding of its environment. Consequently, an entity is not required to accurately understand its environment to any degree. The accuracy of any predictions an entity makes are completely dependent upon how accurately the entity perceives its environment and predicts the future state of its environment.




Figure 5.2: The Simplified OODA loop model with the internal processes and information flow of the Orientation step expanded. Each step and sub-step of the loop requires access to the entity's stored beliefs in order to govern its activities or update those beliefs.

These four activities describe the sub-steps of the Orientation step. Each of these sub-steps requires access to the entity's existing beliefs, which constitute the entity's model of its environment. The Orientation step commences with the newly gathered information

and the entity's beliefs as input. It will then proceed through these sub-steps sequentially, updating the entity's beliefs and then finally preparing a list of suitable options and expected outcomes for the Decision step to select between. This model of the sub-steps of the Orientation is displayed in Figure 5.2, displaying the sub-steps within the simplified model of the OODA loop. The procedural elements are displayed with solid arrows, while the information flows to and from the entity's stored beliefs are shown with dashed arrows.

The functions of the four Orientation sub-steps can also be considered as a synthesis of functions from the other perceptual cycle models (see Section 2.6). The functions of Identification can also be seen in how Neisser's Perceptual Cycle model samples an Object. Interpretation is a function that is also modelled in learning agents, the scientific method and Neisser's Perceptual Cycle's modification of its Schema. Aims Derivation is also exhibited in the goal formation modelled in Norman's Human Action cycle. A learning agent's critic element assesses the effectiveness of an agent in achieving its aims. Options Generation is also performed somewhat by a learning agent's performance element and problem generator.

5.4.3 The Procedure of the Orientation Step

To further clarify how this model of the Orientation step functions, consider how each of the sub-steps is performed by an entity during its Orientation step.

An entity commences its Orientation step in the Identification sub-step, comparing its newly gathered information to its beliefs. This allows the entity to recognise known objects and phenomena, allowing it to retrieve any other beliefs associated with that object or phenomenon. If the entity does not recognise the object or phenomenon, it can create new beliefs to refine during future OODA loop iterations. Subsequent steps further analyse the beliefs related to identified objects and phenomena. During the Identification sub-step the entity is essentially asking itself the question "What is it?" with regard to the various objects and phenomena it has observed and the relationships between them.

Next is the Interpretation sub-step, where the entity compares objects and phenomena to its internal representation of its environment. This allows the entity to update its representation of the environment to match its observed current state. The entity can also identify how its environment has changed since its internal representation was last updated during its previous OODA loop iteration. The entity analyses any changes in the state

of the environment to determine the causes of the changes, updating its beliefs with this information. After determining the present state of its environment, the entity determines what the observed objects, phenomena and relationships imply. An entity learns during this sub-step by analysing and interpreting the newly gathered information to update its beliefs. While updating its beliefs an entity may correct any false beliefs it possesses or introduce new false beliefs. When an entity interprets its newly gathered information it can be considered to be asking itself “What does this mean?”.

During the Aims Derivation sub-step the entity assesses whether it has achieved its aims and then produces new aims. In order to assess its aims, the entity compares its belief of the current state of the environment with the outcomes required to satisfy its current aims. If the environment matches the state that the entity intended, then the entity has achieved the relevant aim. If the state of the environment does not match that required to satisfy the entity’s aim, then the entity has not met its aim. This may occur because the entity’s work towards that aim is incomplete or the aim itself is impossible and cannot be achieved. If the aim is impossible, then the entity needs to update its aims, assuming it realises its error. This sub-step is also where the entity generates new aims to guide its future behaviour. An entity’s aims may be considered as a tiered model of higher order aims and a lower order set of aims, where the lower order aims are more concrete tasks that are intended to help advance the higher order aims. For example, the higher order aims of a biological organism could be survival and reproduction, which are achieved by the lower order aims of avoiding predators, consuming food and searching for a mate. The entity’s higher order aims will direct the development of its lower order aims, which will ultimately direct its behaviour. In this sub-step an entity may be considered to be asking itself “What should I be doing?” and “Am I succeeding at what I am doing?”.

Finally, during the Options Generation sub-step the entity uses both its aims and the updated representation of its environment to determine what options it may perform in the future and predict the expected outcomes of these options. Each option is an action that the entity believes it can perform, while the outcomes are the predicted consequences of those options. The entity’s options and outcomes are intended to further its aims in some manner. These options and their associated outcomes are the final output of the Orientation step and provide the Decision step with a number of potential options to assess and select between. In the Options Generation sub-step the entity can be considered to

be asking itself two questions: “What actions can I perform?” and “What are the likely consequences of these actions?”.

After proceeding through these four sub-steps of its Orientation step, the entity has potentially identified any changes in its environment and the likely causes of these changes, predicted what further changes may occur in the future, assessed the achievement of its aims and determined what its future aims should be, and determined its possible options and the expected outcomes of these options. The Orientation step updates an entity’s beliefs to provide the Decision step with a list of potential options the entity believes it can perform, along with the expected outcomes of these options. The entity is then expected to assess these options and outcomes, using a variety of comparison methods to select an option whose outcomes will best achieve some of its aims and then perform this option during its Action step. This paints the Orientation step as an introspective model, where an entity questions itself about the structure of its reality while determining how best to interact with this reality.

5.5 Errors During the Orientation Step

The previous discussion of the sub-steps of the Orientation step demonstrates how an entity successfully proceeds through this step. However, various errors may occur during each of the four sub-steps, possibly due to faults of the entity itself or interference from external sources. Any errors that occur may introduce errors into the entity’s beliefs, which can then affect other steps of its OODA loop. Corruption attacks, some Subversion attacks, Information Processing Errors and Self-deception all produce misperception that affects the Orientation step, which will therefore affect these sub-steps. Misperception may also occur when false beliefs are created due to a lack of accurate information, possibly due to Degradation attacks, Denial attacks, some Subversion attacks or Information Gathering Errors. Any creation of a false belief is an instance of misperception and identifying where these errors occur provides further insight into misperception. Some of the potential errors that may affect these sub-steps will be described, along with their potential sources and effects.

5.5.1 Identification

An entity's identification of an object or phenomenon may fail for a variety of reasons. Such identification errors lead the entity to develop an incorrect understanding of its environment, which may then affect how the entity subsequently interacts with its environment.

An entity may fail to correctly identify objects and phenomena, due to a lack of understanding of the observed object or phenomenon. There are several causes for this lack of understanding. One is ignorance, as an entity is likely to have difficulty correctly identifying objects and phenomena it has never previously encountered or learned about. A Degradation attack could conceal an object or phenomenon, thereby preventing an entity from correctly identifying the object or phenomenon. Incomplete information may be provided by damaged information sensors, which were potentially affected by a Denial attack or an Information Gathering Error. A Subversion attack may also degrade or reduce the available information an entity collects. Lacking information from a sensor, an entity may be unable to correctly identify the objects or phenomena in its environment. Such an error is found in the story of the three blind men who misidentify an elephant. Each blind man is permitted to touch a different part of the elephant (its legs, its trunk and its side) and from this information each identifies the elephant differently (as a tree, a snake and a wall respectively). If the men were not blind, or in this case used their auditory or olfactory senses, they would be able to correctly identify the elephant.

A Corruption attack is another potential source of errors during the Identification sub-step. Such an attack uses mimicry to disguise an object or phenomenon as a different object or phenomenon. During a successful attack the affected entity incorrectly identifies the object or phenomenon as whatever it mimics. The entity gathers the corrupted signal during its Observation step and then forwards it onto the Orientation step. During the Identification sub-step, the corrupted signal is compared to existing objects and phenomena. Since the corrupted information mimics a valid signal that the entity recognises, the entity incorrectly identifies the object described in the corrupted information as the object or phenomenon it is mimicking. Consider the Corruption attack performed by the various species that mimic the venomous Coral snake (Brodie, 1993), including the Scarlet King snake. Such attacks are targeted against predators such as birds, who associate the

patterning on the non-venomous Scarlet King snake with the Coral snake. A bird that observes a Scarlet King snake and is successfully deceived will incorrectly identify the snake as a Coral snake. Thus its Interpretation will determine the implications of attempting to predate upon this venomous snake and its following sub-steps will proceed on the belief that the snake is actually venomous and should not be attacked.

Subversion attacks may also cause errors in an entity's Identification sub-step, in a similar method to a Corruption attack. Such an attack manipulates the processes that the victim entity uses to identify objects or phenomena in its environment, which may lead to the incorrect identification of observed objects or phenomena. Any incorrect identification is then forwarded to the Interpretation sub-step, likely causing further errors to arise.

5.5.2 Interpretation

While analysing new information the entity may, either accidentally or intentionally, incorrectly interpret this information. Incorrect interpretation can be produced by incorrectly applying analysis methods to the information gathered or by using analysis methods that do not produce valid results. An example of an error produced by incorrectly applied analysis methods can be found in an engineer studying a building's structure. The engineer may incorrectly use analysis methods that are only suitable for brick buildings to analyse the structure of a steel and concrete building. Such an error may give the engineer a false belief regarding the attributes of the brick building. An entity might use an inappropriate analysis method because it incorrectly believes that the method is appropriate; possibly due to ignorance, an inability to learn, a Corruption attack or a previous Subversion attack. An erroneous analysis method could be a mathematical function intended to calculate the average value of a series of numbers, which omits some numbers from the calculation. An entity may possess erroneous analysis methods due to its ignorance or incompetence, or as the result of a Corruption or Subversion attack that has targeted the entity's analysis methods. In both of these types of analysis errors, the results of the interpretation are likely to be incorrect and will lead to the creation of false beliefs, which may affect future behaviour.

Subversion attacks may target and manipulate the information or methods that an entity uses to analyse new information. Such attacks alter the Interpretation of new and current beliefs the entity holds. For example, an automotive company who has just

initiated a large safety recall of its vehicles might employ an advertising campaign to convince potential customers that it cares about their safety. This Subversion attack intends to alter how potential customers perceive the car company's attitude to safety. Potential customers who were affected by this attack may interpret news of the automotive recall as evidence that the car company cares about the safety of its customers, whereas those unaffected by the attack may interpret the recall as evidence that the company's cars are unsafe.

Entities may suffer from anxiety during their Interpretation sub-step, as cognitive dissonance develops between their newly analysed information and their existing beliefs. This dissonance can be reduced by Self-deception, which irrationally re-interprets the new information in a manner that supports the belief that the entity desires to possess. During this re-interpretation either the new information or the existing beliefs are manipulated in some manner to remove the conflict between the two. This allows the entity to deal with dissonant information by ignoring it, manipulating it into an acceptable form or manipulating its existing beliefs. Self-deception cannot occur earlier in the Identification sub-step, as while an entity can identify dissonant objects and phenomena, it cannot determine why they are dissonant or how to reduce the dissonance until the Interpretation sub-step.

Self-deception that is used to aid deception operates in much the same manner, except that beliefs that may lead to indications of untruthfulness are manipulated during the Interpretation step, with the intention of restoring those beliefs during a later Interpretation step after successfully deceiving the entity's target.

The mechanisms that an entity uses to access its stored beliefs may fail. This may prevent an entity from accurately retrieving beliefs or storing beliefs. When retrieving its beliefs an entity may instead receive false beliefs, which are then used while analysing new information. The entity may then produce false beliefs that it then stores. An entity may also fail to correctly store any new beliefs it has created. An error during the storage of the new belief could replace it with a false belief instead. In humans such errors would be attributed to forgetfulness or poor memory. Equivalent failures may also occur in human organisations, machine systems and some other biological entities.

Under-sampling may lead to errors during Interpretation when the entity's environment changes state while the entity is still interpreting information collected during the

environment's previous state. The entity will then develop false beliefs regarding the state of the environment, which will then affect its decision-making process until its next Observation step, when it observes the new state of the environment.

5.5.3 Aims Derivation

During the Aims Derivation sub-step an entity may fail to correctly assess its success or failure in achieving its aims. An entity may not recognise that it has successfully achieved an aim, possibly due to an incorrect understanding of either the aim or the environment. Alternatively, an entity may not recognise its failure to achieve an aim, instead incorrectly determining that it has succeeded. Such errors are caused by possessing false beliefs describing the conditions that determine success or failure in achieving the aim. These false beliefs could be due to earlier errors of the entity or hostile actions of the entity's competitors.

An entity may develop incorrect aims that are impossible to achieve or difficult to correctly assess. In such a case, the entity's false beliefs are responsible for this aim that it cannot achieve. Similarly, an entity may too vaguely define an aim, therefore leaving it unable to assess whether or not it achieves this aim. Earlier errors of the entity, Information Warfare attacks or a combination of the two may produce this error.

A Subversion attack that targets an entity's Aims Derivation could manipulate the entity's methods for assessing its success or failure in achieving its aims or those methods responsible for developing aims. Such a Subversion attack could convince an entity to leave a task unfinished or attempt an impossible task. Such an attack would not directly cause misperception, but would affect an entity's behaviour.

5.5.4 Options Generation

An entity may fail to recognise potential options that it may perform during its Options Generation sub-step. Due to an incorrect or incomplete understanding of its environment, an entity may disregard an option it believes to be impossible or unsuitable for achieving a desired outcome. An entity may also be unaware of an option in some instances. An option that is not recognised cannot be considered during the Decision step and therefore restricts the actions that the entity may perform.

Another type of error that can affect an entity during the Options Generation sub-step occurs when the entity generates options that it mistakenly believes are possible to perform. Should the entity select these options during its Decision step, it will encounter difficulties during its Action step when it attempts to perform them.

During the Options Generation sub-step an entity may also predict inaccurate outcomes for its possible options. False beliefs lead the entity to develop an unrealistic prediction of the outcome for a given option. This can be considered to be wishful thinking on the part of the entity, as the outcome it expects from that option will not occur. An example of such an error could be found in the expectations of a country that has the option of invading a hostile country to depose its dictator. The attacking country's false beliefs may lead it to incorrectly determine that such an invasion would be followed by a quick transition to a friendly democracy. The option and its erroneous outcome may be chosen during the Decision step if it is assessed to be the most suitable or beneficial course of action.

Subversion attacks against an entity's Options Generation sub-step manipulate the methods an entity uses to generate options and determine their expected outcomes. Such manipulations could cause an entity to ignore options that an attacker abhors or to add options that an attacker desires. Furthermore, an entity's mechanisms for determining likely outcomes may also be manipulated by Subversion, causing an entity to develop incorrect predictions of the consequences of its actions. In either case, an attacker manipulates an entity's Options Generation sub-step in order to cause the entity to involuntarily alter its behaviour. Such an attack will not directly cause misperception, but will affect an entity's behaviour.

With the exception of Subversion, all of the errors that affect the Options Generation sub-step are caused by false beliefs, which affect either the options the entity believes it can perform or the resultant outcomes of these options. These false beliefs may have been produced by the Information Warfare attacks of competitors, other unintentional misperception-causing errors or a combination of these errors.

5.5.5 Effects of Misperception on the Orientation Step

As demonstrated in the previous chapter, misperception occurs in two forms. One form impairs the quantity or quality of information that an entity collects, while the other

represents an error of understanding. In either case, the affected entity produces false beliefs during its Orientation step, which then affect the remainder of its current OODA loop and also subsequent decision-making cycles. In the model of the Orientation step developed, the errors previously described as misperception occur in the Identification and Interpretation sub-steps. The Aims Derivation and Options Generation sub-steps rely upon the entity's current beliefs to produce further beliefs, which then govern the entity's eventual behaviour. Misperception affects this model by cultivating false beliefs in the Identification and Interpretation sub-steps, which then affect the subsequent Aims Derivation and Options Generation sub-steps. A Subversion attack can affect any sub-step of the Orientation step. However, in order to directly cause misperception, Subversion must affect either the Identification or Interpretation sub-steps. Figure 5.3 shows where the different sources of misperception affect this expanded OODA loop model.

Errors of the Observation step, such as Degradation attacks, Denial attacks, some Subversion attacks and Information Gathering Errors, cause the entity to produce false beliefs during its Orientation step. If misperception affects the Observation step, the affected entity collects incomplete information from its environment, which it then passes along to its Identification sub-step. During Identification, the entity may fail to correctly identify elements within the gathered information or it may be unable to identify elements that its information sensors did not correctly detect. During the Interpretation sub-step the newly collected information is analysed. Since the information is incomplete and does not fully describe all the elements of the environment, the entity will create false beliefs when it attempts to interpret the information. These false beliefs will then affect the remainder of the entity's OODA loop and any future iterations of its OODA loop.

This model of the Orientation step also reveals a difference between Corruption attacks, Information Processing Errors and Self-deception. Corruption attacks affect an entity during its Identification step, when the corrupted signal is incorrectly identified as the signal it mimics, which then leads the entity to develop false beliefs. However, Information Processing Errors and Self-deception do not affect an entity until its Interpretation sub-step, where they lead the entity to develop false beliefs. This difference between Corruption and Self-deception argues that Self-deception cannot be considered to be reflexive Corruption, since they function differently in this model.




Figure 5.3: The various sources of misperception mapped onto the expanded OODA loop model. Note that while Subversion may affect all steps and sub-steps of the OODA loop, it only directly causes misperception during the Observation step, Identification sub-step and Interpretation sub-step. During the other steps and sub-steps it may instead induce self-targeted attacks, which then cause misperception in the Observation step, Identification sub-step and Interpretation sub-step.

General incompetence within a field may also contribute to various errors during any sub-step of the Orientation step. The Dunning-Kruger effect notes the correlation between incompetence in an area and the inability to accurately assess one's actions and capabilities within that area. For example, consider the effects of an entity's incompetence in basket weaving. This incompetence might prevent the entity from accurately identifying tools used to weave baskets, understanding the elements of complex basket weaving processes, accurately assessing and deriving Aims associated with basket weaving or accurately determining the outcomes of actions involved in weaving or using a basket. In any case, incompetence begets ignorance, and causes errors that may affect any step or sub-step of the OODA loop.

Now I will map historical case studies of Corruption attacks, some Subversion attacks, Information Processing Errors and Self-deception into this model of the Orientation step to further test the model. This will also demonstrate how those instances of misperception occurred and how they affected their respective entities.

5.6 Case Studies Focusing on the Orientation Step

How Information Warfare and other unintentional sources of misperception may affect an entity's decision-making process was discussed previously in Chapter 4. It was demonstrated that Corruption attacks, Subversion attacks, Self-deception and Information Processing Errors cause errors during the Orientation step. In order to test and validate this model of the Orientation step, examples of these types of misperception will be mapped into this model.

5.6.1 Corruption Attacks

An earlier example described how the overall deception plan used by the Allies before the D-Day landings operated against Germany's OODA loop (Section 4.2.2) (Haswell, 1979, 1985). Now I will revisit this example, focusing on the actions that occurred within Germany's Orientation step. This will demonstrate that Corruption attacks cause errors in the Identification sub-step, as previously argued. As before, Germany's intelligence services and military leaders are considered to be its decision-makers.

In the case of Operation Fortitude, the various deceptions the Allies performed were part of a compound Information Warfare attack that consisted of two complementary deception plans, each consisting of numerous Corruption attacks. One deception plan mimicked the preparations for an invasion at Norway, while the other mimicked the build-up of forces for an invasion at Calais. Figures 5.4a and 5.4b show examples of the decoy vehicles used to mimic an invasion force. These plans were intended to hold German troops away from Normandy, who would instead mistakenly reinforce Calais and Norway. Other parts of the larger overall invasion plan, Operation Bodyguard, also acted in concert with Operation Fortitude to distract the Germans from Normandy.

The German intelligence services gathered information that the Allies had permitted them to gather during their Observation step, such as aerial reconnaissance photographs and reports from turned spies. This information from their Observation step was then passed along to the Identification sub-step of the Orientation step, where the German intelligence services incorrectly identified the various deceptive elements (such as inflatable decoy vehicles, empty troop encampments and fake radio transmissions) as evidence of a forthcoming invasion. During the Interpretation sub-step, the deceptive elements were



(a) An inflatable decoy lorry. ©Imperial War Museum (H 42530), <http://www.iwm.org.uk/collections/item/object/205201877>



(b) Dummy aircraft. ©The National Archives (AIR 20/4349 (Oct 1943)), <http://www.nationalarchives.gov.uk/battles/dday/popup/deception.htm>

Figure 5.4: Operation Fortitude consisted of many Corruption attacks, some of which mimicked the forces and materiel required for the invasion of Calais, in part by deploying lorries (a) and decoy aircraft (b). Further examples of such Corruption can be found in Figures 4.6a and 4.6b.

analysed and interpreted by the German intelligence services. With the false belief that the Allies were assembling an invasion force in south-east England, the German intelligence services concluded that Calais and Norway would be the targets of the invasion. This interpretation was correct from Germany's perspective, given the false beliefs created by the Allies' deceptions. The interpretation established and reinforced the belief that the invasion would occur at Calais and Norway. Germany's subsequent Aims Derivation and Options Generation was now guided by the belief that Calais and Norway were the invasion targets. Germany therefore developed aims to reinforce those expected invasion targets. Those aims led to options and outcomes in which large amounts of troops were held in reserve at Calais and Norway to defend against those expected invasions. By selecting and performing these options, Germany made the Allies' invasion at Normandy much easier.

The Corruption attacks, performed by the mimicry of an entire army group, affected the German intelligence services during their Identification sub-step. This error created and reinforced false beliefs, which then affected Germany's following sub-steps. The subsequent steps of the OODA loop were then correctly performed with these false beliefs.

5.6.2 Subversion Attacks

The spin used to help market and promote "biosolids" in the United States (Stauber and Rampton, 1995) demonstrates an example of Subversion. In 1972, the passage of the

Clean Water Act mandated that waste water treatment facilities remove up to 85% of the pollutants from sewage. While this law dramatically improved the standards for waste water treatment, it also raised the problem of what to do with the pollutants removed from the waste water. These pollutants form a thick toxic sludge, containing heavy metals, toxic chemicals, bacteria, viruses, pharmaceutical waste (Don, 2013) and other pollutants.

In an attempt to solve the growing problem of sludge disposal, the Environmental Protection Agency (EPA) and the treatment plant operators turned to the cheapest and easiest solution — spreading the sludge on farmland as fertiliser. In 1991 the federation of treatment plant operators renamed their sludge as “biosolids”, intending to alter the perception of their waste. By 1992, the EPA had altered its regulations so that biosolids could be used as fertiliser and funded public relations campaigns to educate the public about the beneficial uses of biosolids. Such positive reporting on the benefits of sludge often used scientists and scientific reports to tout the safety of biosolids, while presenting positive stories into the media where biosolids were to be used.

Such methods represent a Subversion attack, which shall now be mapped into the OODA loop model and expanded Orientation step, by considering the case of a farmer subjected to such an attack.

The farmer begins his Observation step by reading a newspaper story touting the safety and advantages of using biosolids as fertiliser. During the Identification sub-step of his Orientation step, the farmer identifies the media story and its scientific references as seemingly valid. During the farmer’s Interpretation sub-step, the information from the newspaper story is interpreted and accepted as true by the farmer, since he accepts the idea that biosolids are a safe fertiliser. The farmer then continues through the rest of his OODA loop after reading the story, having manipulated his methods for accurately assessing the merits of biosolids.

Later, the farmer meets with some fertiliser salespeople who try to sell him various types of fertiliser, including biosolids. During the farmers Observation step he gathers all the relevant price and quality information on the various types of fertiliser, which is passed along to the Identification sub-step of his Orientation step. During the Identification sub-step, the farmer recalls the fertiliser types and their relevant characteristics, including those of the biosolids. During his Interpretation sub-step, the farmer applies the previous positive information about biosolids with his need for fertiliser and determines

that biosolids would be a suitable fertiliser for his fields. During his Aims Derivation and Options Generation sub-steps the farmer considers whether he should use biosolids to fertilise his farm. The farmer then chooses to purchase biosolids during his Decision step and signs the contract during his Action step.

This Subversion attack convinces the general public, including the farmer, that biosolids are a safe and effective fertiliser product, ultimately convincing the farmer to purchase and use biosolids to fertilise his farm. However, the farmer has actually been induced to pay someone else to dump toxic waste on his land.

5.6.3 Self-Deception

Since the majority of documented examples of Self-deception describe its effects on humans, these examples will only consider human Self-deception. These examples focus on three different instances of Self-deception. The first considers Self-deception performed by individual humans to reduce anxiety caused by unacceptable beliefs. The second describes how Self-deception that aids deception produces errors in the Orientation step. The third discusses an organisational example of Self-deception used to conceal unacceptable beliefs, which can also aid deception against others. While Self-deception can arguably benefit Self-deceiving entities, these three examples all demonstrate instances where this benefit was ultimately outweighed by the Self-deception's negative outcomes.

Self-deception and Challenger

The decision-making processes that led to the Challenger disaster are a good example of Self-deception. Self-deception was performed by the managers of the company responsible for building the shuttle's solid rocket boosters. These managers approved the launch despite warnings from their engineers that the booster rockets could fail due to the forecast low temperature for the launch.

On January 28th 1986, the space shuttle Challenger disintegrated shortly after take-off (Lewis, 1988) (Figure 5.5). This was caused by a failure of the shuttle's solid rocket boosters, which was later attributed to a failure of the O-rings within the boosters due to the low temperature at the time of the launch. Data from earlier launches had indicated that the O-rings sustained more damage during low-temperature launches and were therefore more likely to fail. Engineers at both NASA and Morton Thiokol, the makers of the solid



(a) Ice on the Challenger launch gantry. ©NASA, <http://grin.hq.nasa.gov/ABSTRACTS/GPN-2000-001348.html>



(b) Plume of exhaust leaking from Challenger's left booster. ©NASA, <http://grin.hq.nasa.gov/ABSTRACTS/GPN-2000-001425.html>

Figure 5.5: Low temperatures before the Challenger launch (a) led to the failure of the O-rings in the shuttle's solid rocket boosters. The failure of the O-rings on the left booster as the shuttle climbed is clearly shown as a large plume (b) above the left solid rocket booster's exhaust nozzle. Exhaust gases from the booster escaped through this point and cut the struts attaching the booster to the main tank, which eventually caused the shuttle to break apart.

rocket boosters, were aware of the damage to the O-rings caused during earlier launches and that it was a critical problem that could lead to the loss of the shuttle. The night before the launch Morton Thiokol and NASA were at an impasse over whether it was safe to proceed with the launch, due to a weather forecast predicting a temperature colder than all earlier launches. NASA believed it was safe to launch while Morton Thiokol's engineers disagreed, as the data from earlier launches showed that O-ring damage increased as the launch temperature decreased. The forecast temperature was colder than any earlier launch and this worried the engineers, as it was outside their experience. Morton Thiokol's management, however, incorrectly interpreted the damage to the O-rings as evidence of a safety factor in the design (Starbuck and Milliken, 1988), which convinced them the launch was safe and authorised NASA to proceed.

While the shuttle disaster was ultimately an engineering failure, Self-deception on the part of Morton Thiokol's management team led them to approve the launch despite awareness of the potential danger raised by their engineers (Trivers, 2011). This "safety factor"

interpretation of the data was preferred by the managers as it resolved their impasse with NASA, who were a valued customer, and also prevented them from acknowledging flaws in their solid rocket booster design, which would have required a lengthy and expensive redesign to fix. To compound this error, the fact that Morton Thiokol had initially argued against the launch was not communicated to those at NASA responsible for authorising the launch. According to the Rogers Commission Report into the disaster, NASA management would have been unlikely to launch the shuttle if they had been aware of the concerns raised by Morton Thiokol's engineers about the effects of low temperatures on the O-rings (Rogers Commission, 1986).

The behaviour of Morton Thiokol's managers will now be explained in terms of the OODA loop model with an expanded Orientation step. This behaviour is categorised as Self-deception, since Morton Thiokol's management wanted to believe that it was safe to launch the shuttle, despite awareness of evidence from their engineers to the contrary. The warnings from the engineers were observed during the managers' Observation step and passed along for Identification, where they were correctly identified as warnings about the safety of the O-rings. These warnings were passed along to the Interpretation sub-step to determine their significance. During Interpretation, the consequences of delaying the launch due to the temperature and potentially redesigning the solid rocket booster joints was found to be dissonant with the beliefs reinforced by earlier launches, which indicated that the boosters were safe. The managers Self-deceptively interpreted the warnings from the engineers as an indication that the O-rings provided an adequate margin of safety. Once the managers had Self-deceptively "discovered" this safety margin they had no reason to prevent the launch. Subsequently their Aims Derivation and Options Generation focused on authorising Challenger's launch, since they believed that it was safe to do so. The managers then selected such an action during their Decision step and performed it during their Action step.

The managers' Self-deception afflicted Orientation step can be contrasted against that of the engineers, who did not Self-deceive. The engineers correctly interpreted the meaning of the O-ring data from earlier launches to indicate that launching at a temperature lower than any earlier launch would likely produce even greater damage to the O-rings, possibly causing a catastrophic failure. This belief guided them to recommend the launch be delayed.

In this example, Self-deception affected the managers' Interpretation sub-step, when they analysed the information from their engineers and decided to interpret the warnings as evidence of a non-existent safety factor. This could not have occurred in the Identification sub-step, as the managers would have had no justification for manipulating their beliefs before they had determined the implications of the engineers' report.

Self-deception causing an Air crash

Another instance where Self-deception led to a disaster can be found in the case of the crash of Air Florida Flight 90 in 1982. Shortly after take-off during a snowstorm, the aircraft crashed into the Potomac River in Washington DC. A build-up of ice on the aircraft's wing caused the aircraft to stall and crash. The effects of the ice on the wings was made worse due to the pilots' failure to use the aircraft's anti-icing equipment. The ice build-up also caused a thrust sensor in the engine to over-report the thrust produced by the engines. This caused the engines to produce less thrust than the aircraft's instruments indicated during the take-off. The pilot convinced the co-pilot to continue the take-off, despite his awareness of the ice build-up on the wings that would decrease the lift produced by the wings. When the co-pilot informed the pilot that the thrust sensor readings were inconsistent with the other instruments, the pilot ignored this warning instead of aborting the take-off. The pilot also convinced the co-pilot that the readings were acceptable.

Trivers and Newton (1982) argued that the willingness to disregard the inconsistency between the thrust sensors and other instruments was Self-deception, on the basis of the discussion between the pilot and co-pilot that was captured on the cockpit voice recorder. During take-off the co-pilot was in charge and under the instruction of the pilot. Based upon the cockpit voice recordings the pilot was described as an "overconfident Self-deceiver", while the co-pilot was said to be "timid but reality-oriented" (Trivers, 2002). According to the accident report, both pilots had little winter flying experience (National Transport Safety Bureau, 1982) and this inexperience may have led the pilots to overestimate their winter flying abilities and to be unaware of the true danger posed by the conditions. The cockpit voice recordings of the take-off reveal the co-pilot repeatedly questioning the instrument readings. The pilot ignores these warnings at first, then later tells the co-pilot that the instrument readings are correct. The pilot does not want to believe the instruments are incorrect, as this would require him to abort the take-off and

delay the flight. The flight was already late due to the poor weather and the pilot did not want to further delay the aircraft's departure, which would worsen the airline's timetable and potentially cause the pilot to be criticised by his management for his tardiness. Instead the pilot uses Self-deception to convince himself that it is safe to continue the take-off and then unintentionally deceives the co-pilot that it is safe to continue with the take-off (Trivers, 2011).

When the pilot's behaviour is mapped into the OODA loop model and the expanded Orientation step, the pilot receives the unacceptable contradictory information from the co-pilot during his Observation step. During the Identification sub-step of the Orientation step, the pilot correctly identifies the warnings from the co-pilot, which are then passed along for Interpretation. Here the warnings are analysed and found to contradict with the pilot's beliefs that it is safe to continue the take-off and his desire to avoid any delays. The pilot may have disregarded the co-pilot's warnings because he believed he was more experienced, which would help him to rationalise ignoring the warnings. The pilot then Self-deceptively interprets the airspeed indicator reading as correct, which permits him to disregard the co-pilot's warnings and believe that it is safe to continue the take-off. The pilot therefore continues his Aims Derivation and Options Generation sub-steps by producing options to continue the take-off, which he predicts to have an outcome of a successful take-off. During the Decision and Action steps the pilot selects and implements this option, continuing the take-off and communicating to the co-pilot that there is nothing wrong with the instruments. The pilot's actions therefore function as deception against the co-pilot, who defers to the pilot's seniority and accepts his decision to continue the take-off.

In this example the pilot desires to take-off and rejects the information that suggests that the co-pilot should abort the take-off. This manipulation occurs during the Interpretation sub-step, where the pilot determines the consequences for aborting the take-off and their contradiction with his existing beliefs. Furthermore, after this Self-deception the pilot's behaviour and statements deceptively convince the co-pilot that the sensor readings were acceptable and that it was safe to continue the take-off. In this case, the pilot's usage of Self-deception to aid deception was initially advantageous, as it prevented the co-pilot from aborting the take-off. However, it prevented the co-pilot from determining the true

state of the aircraft, which would likely have convinced him to abort the take-off, thereby preventing the crash.

Historical Revisionism in the Soviet Union

Another example of Self-deceptive behaviour can be found in the Soviet Union's historical revisionism. This example of Self-deception by an organisation demonstrates the use of Self-deception to reduce dissonance from conflicting beliefs.

The official history of the Soviet Union was frequently rewritten, in order to conceal mistakes from the population and remove references to people who were no longer favoured by the regime (King, 1997). For example, works by denounced politicians and writers were removed from libraries and official archives, while school history books and other official publications were frequently revised to remove references to such people. King (1997) has documented many instances where individuals were removed from official photographs and paintings after they fell from favour with the regime. This behaviour by the state developed a flexible version of history and is well described in a joke — “What is more uncertain than the future? The history of the Soviet Union” (Holden, 1994). Such revisionism, when taken to its extreme, leads to a situation similar to that described in Orwell's (1949) *Nineteen Eighty-Four*, where the official history is constantly manipulated to suit the immediate short-term needs of the ruling party. Stalin's purges of the late 1930's led to the murder of his political opponents, along with their removal from all forms of pictorial existence (Figure 5.6). At this time, Stalin was glorified as “the only true friend, comrade and successor to Lenin” (King, 1997). To maintain this lie, Stalin was then edited into places he had never been and his revolutionary contemporaries removed, while many documents and materials in the official archives were also destroyed (McCauley, 2003). One notable effect of the constant acts of revisionism was the cancelling of the 1988 Russian high school history exams. The Soviet government deemed this action necessary, as it was forced to acknowledge that much of the history contained in the high school history books was untrue and inaccurate (Wertsch, 1999).

In this example, the Soviet Union as a nation is the Self-deceiver. It commonly desired to accept false beliefs, despite awareness of documentary evidence disproving those beliefs. To enable such Self-deception, the various official documents and records were destroyed or manipulated to support the Soviet Union's desired beliefs. This appears to be an



(a) Kliment Voroshilov, Vyacheslav Molotov, Josef Stalin and Nikolai Yezhov on 22 April 1937. http://commons.wikimedia.org/wiki/File:Voroshilov,_Molotov,_Stalin,_with_Nikolai_Yezhov.jpg



(b) The revised photograph, used after Yezhov's execution in 1939. http://commons.wikimedia.org/wiki/File:The_Commissar_Vanishes_2.jpg

Figure 5.6: Nikolai Yezhov was head of the NKVD during Stalin's Great Terror and was responsible for the purges at Stalin's order. During the purges Yezhov is seen with Stalin (a), yet Yezhov is removed (b) after Stalin purged him for his role in the Great Terror.

example of Self-deception used to reduce cognitive dissonance, which was produced by contradictions between the official and actual histories of the Soviet Union. A nation's Self-deception produces false historic narratives that can build unity within its citizens (Trivers, 2011), institutionalising Groupthink by the nation's population.

During its Observation step, the Soviet Union would collect new information that described its recent actions, including purging "unreliable" citizens. During the Identification sub-step, a purged citizen might be recognised as a previously influential luminary of the state. In the Interpretation stage this new information is analysed and found to be dissonant with the official history of the Soviet Union, which shows the purged citizen as a loyal patriot, who had offended the regime or was a political competitor to Stalin. To reduce this dissonance, the official history and the documents that describe it are manipulated to remove references to the citizen. The subsequent Aims Derivation and Options Generation sub-steps are then performed normally by the Soviet Union, with the beliefs in the updated official history used to affect this behaviour. Finally, the Decision and Action steps operate normally while referring to the updated official history.

In this example cognitive dissonance caused by the contradictions between the official history and recently observed actions is reduced by manipulating the official history to remove the contradiction. While this Self-deception was intended to reduce cognitive

dissonance, it also permitted the deception of others observing the Soviet Union's official history. Such deception could occur when historical information is communicated to individuals or organisations within or outside its borders.

5.6.4 Information Processing Errors

A previous example of an Information Processing Error describes the actions of a man affected by the Gambler's fallacy (Section 4.4). The fallacy affects his perception of the likelihood of a roulette wheel being "due" to come up Black after the 10 previous spins have all come up Red. His incorrect understanding of the true statistical nature of the roulette wheel's results affects his behaviour, potentially to his detriment. Now consider his behaviour specifically during his Orientation step.

During the man's Observation step, he correctly observes the results of the 10 previous games of roulette, where the wheel has come up Red in each of those games. The information passes to the man's Identification sub-step, where he correctly identifies the results of the previous 10 games as being statistically unlikely. During his Interpretation sub-step, the man analyses the likelihood of such a sequence of Red outcomes. However, due to the Gambler's fallacy, the man incorrectly determines that the wheel is "due" to come up Black soon and therefore Black is a more likely future result than Red. Since the man believes that the next spin of the roulette wheel is much more likely to come up Black, he develops aims to bet on Black during his Aims Derivation step. He turns this aim into an option during his Options Derivation sub-step. During his Decision and Action steps, the man now chooses to bet heavily on Black and does so. However, contrary to the man's beliefs, the bet is not in his favour, as Red and Black are both equally likely outcomes on the next spin of the roulette wheel.

This error occurs when the man analyses the likelihood of the unlikely outcome during his Interpretation sub-step. The Gambler's fallacy causes him to incorrectly reason that the small sample of 10 Red outcomes indicates that the next spin is much more likely to be Black, to even out the 10 consecutive Red spins. However, the Interpretation sub-step is not the only place where processing errors may occur. A Processing Error could also occur if the victim incorrectly identifies the objects or phenomena, leading the victim to apply an incorrect processing method or use the information incorrectly. In this example, such an error could arise if the man misidentifies the roulette wheel as "rigged" in some

manner that is more likely to come up Red than Black. In such a case, the error occurs during the Identification sub-step, which then causes the man to incorrectly determine that Red is a more likely outcome than Black during his Interpretation sub-step. This false belief would also elicit a similar behaviour where the man bets heavily on Red, based upon his flawed reasoning that the roulette wheel is much more likely to come up Red on the next spin.

5.7 Summary

The OODA loop model was compared to another cognitive architecture, the agent model, to determine whether it might be a suitable replacement for the OODA loop model. While the agent model is often functionally similar to the OODA loop model, the OODA loop focuses on competitive environments where competing entities possess incomplete information, making it a preferable choice for examining misperception.

To better explain how misperception affects an entity's Orientation step, a model of the internal processes of the Orientation step was developed. This model demonstrates both the internal procedure of the Orientation step, along with the flow of information that takes place during this process. Developing this model and examining how misperception affects its processes has helped to reinforce how much entities depend upon their beliefs during their decision-making process.

5.7.1 Assessment of the Developed Model

One goal in the development of this model was to maintain the generality of the OODA loop model, which enables it to be applied to many different types of decision-making entities. As the Orientation step operates upon the entity's stored beliefs, an expanded model of the Orientation step is mainly useful when considering entities that update and manipulate their beliefs frequently. These entities may be humans, human organisations, machine systems or other cognitively complex biological organisms. As the historical case studies have shown, this model is suitable for modelling the Orientation step of humans and human organisations.

The model of the Orientation step arguably requires entities to have a higher cognitive complexity. This model may not be applicable to simpler entities, due to their simple

decision-making processes and limited information processing capabilities. For example, entities whose behaviour is governed by simple rules are not capable of analysing their aims or determining which options will produce an optimal outcome for them. The developed model of the Orientation step assumes that the entities examined possess sufficient cognitive capabilities to perform each of the sub-steps. For this reason the expanded model of the Orientation is unlikely to be useful for modelling the OODA loops of less cognitively complex entities.

This model was also developed to demonstrate the flow of information inside the entity as it progresses through its Orientation step. The model describes the entity's Orientation step as it begins with the newly observed information and concludes as the entity produces a number of options and related outcomes, which serve as input for the Decision step. Each of the internal sub-steps requires access to the entity's stored beliefs. An entity's Identification and Interpretation sub-steps update its understanding of its environment, which it stores in its beliefs. The entity's existing beliefs also act as input for these steps. The observed information enters these steps, where the entity processes, analyses and manipulates it in various ways, to update its beliefs. These beliefs are then the input to the entity's Aims Derivation and Options Generation sub-steps, which assess the entity's past behaviour and determine its future behaviour. The entity's updated beliefs will then affect later steps of the current OODA loop iteration, along with future loop iterations. An entity's beliefs form a centralised storage location for the beliefs, created by the entity as it analyses and interprets its collected information. The effects of various misperceptions will accumulate in this storage location as the entity produces false beliefs. These beliefs, both true and false, govern an entity's behaviour during all the steps of its OODA loop.

5.7.2 The Importance of Accurate Prediction

One important aspect of Orientation step this research demonstrates is the importance of an entity's ability to accurately predict the future state of its environment. While the OODA loop model implicitly suggests this importance, expanding the Orientation step better demonstrates how an entity's actions ultimately shape its predictive abilities. Among other things, an entity uses these predictive abilities to estimate how the environment will change, how its actions will affect the environment and how its competitors will likely act in the future.

An entity's predictions are made based upon its current beliefs; specifically, its current understanding of its environment and the various rules and restrictions that govern the environment and how the entity interacts with it. Predictions are important as they ultimately drive the entity's actions, by attaching predictions to proposed actions and then assessing the desirability of the predicted outcome occurring. In this model, the Options Generation sub-step uses an entity's predictive capabilities and beliefs to list potential options and their expected resultant outcomes. Its final output is a list of options and their predicted outcomes, which will then be assessed to select between the potential options that were considered.

Any sources of misperception, whether intentional or unintentional, will ultimately lead to the introduction of false beliefs during an entity's Interpretation sub-step. These false beliefs will be used by the entity to evaluate what the future state of the environment may be and what the likely outcomes of its actions will be. Degrading an entity's ability to predict the future state of its environment will cause it to select and perform actions based upon an incorrect understanding of its environment, which are likely to yield sub-optimal outcomes.

5.7.3 Misperception during the Orientation Sub-steps

Analysing how misperception functions in the expanded Orientation step reveals that misperception arises due to errors within either the Identification or Interpretation sub-steps, causing the entity to produce false beliefs. These false beliefs may then affect the entity during its Aims Derivation and Options Generation sub-steps, or during any subsequent element of its OODA loop. The historical case studies detail how these errors affect the beliefs of an individual or an organisation and how these false beliefs affect behaviour.

Self-deception was examined from both an individual and organisational perspective, detailing its capability to reduce cognitive dissonance and to aid the deception of others. Self-deception causes errors during the entity's Interpretation step, when the entity deliberately alters its own beliefs. In the case of a Corruption attack, the errors occur during the entity's Identification sub-step, as it mistakes the corrupted information for the valid information it mimics. Furthermore, mapping Self-deception and Corruption attacks into the expanded Orientation step demonstrates that Self-deception differs from Corruption,

implying that Self-deception is not reflexive Corruption. Self-deception requires that the newly gathered information is understood to conflict with the entity's existing beliefs, which must occur during the Interpretation step. Corruption, however, causes an error during the Identification sub-step before the entity can determine that the information conflicts with its existing beliefs.

Information Processing Errors may affect either the Identification or Interpretation sub-steps. This causes an affected entity to develop false beliefs, which may affect it during the remainder of its current OODA loop iteration and also during future iterations.

In both the Challenger and Air Florida Flight 90 disasters, the Self-deceivers chose not to abort ultimately catastrophic procedures, despite an awareness of evidence identifying potentially dangerous behaviour. In both cases the Self-deceivers considered and rejected the abort option due to its significant perceived penalties, while ignoring the potential penalties of proceeding and failing. This suggests that despite the warnings, the Self-deceivers did not consider failure to be a probable outcome, possibly as they had no knowledge of previous or similar failures. This oversight may in turn be due to either further Self-deception or ignorance. It also indicates that the Self-deceivers' risk assessment and risk management mechanisms were degraded by the Self-deceptions or that they were otherwise incapable of correctly assessing the negative potential outcomes of their actions. Such a failure of judgement (ignoring the possibility of failure) is not unique to Self-deceiving entities and may occur in any instance in which an entity lacks the required beliefs or assessment algorithms to correctly determine the risks and consequences of its actions.

5.7.4 The Importance of the Orientation Step

In studying the OODA loop and expanding its Orientation step, it has become clear that an entity's beliefs are of central importance to its decision-making abilities. An entity's beliefs also shape how it will interpret new information and ultimately how it will act. The Orientation step is important as it is where the entity uses information from its environment to update its beliefs. Boyd has previously described the Orientation step as the "schwerpunkt" or focal point of the OODA loop, due to its importance in the decision-making cycle. It is also highly vulnerable to errors, as earlier errors in either the Observation step or a previous cycles' Orientation step can introduce false beliefs into an

entity. These errors may also be caused by Information Warfare attacks. In most cases, an entity's beliefs can therefore be considered the ultimate target of such attacks, with the attacker manipulating the victim's beliefs in order to affect its subsequent behaviour.

False beliefs, especially those created intentionally by competitors, will likely produce detrimental behaviour in an entity. However, in some circumstances, specific types of misperception may produce beliefs that lead to behaviours that benefit either the misperceiving entity or other entities in its population. As such, misperception may be able to provide an evolutionary benefit and I will now explore this hypothesis.

Chapter 6

An Evolutionary Benefit from Misperception

There have been few attempts to investigate the beneficial effects of misperception through Artificial Life simulation, likely due to the assumption that misperception is always detrimental. There are, however, several methods by which misperception may provide an evolutionary benefit. Simulations performed by Doran (1998) demonstrated cases where misperception could be supported in an evolutionary environment through individual selection and kin selection. In these simulations misperception benefited agents by either deterring individually harmful behaviour or preventing hostile actions against those who shared a common misbelief.

Another Artificial Life simulation that demonstrated a benefit from misperception was created by Akaishi and Arita (2002a), who simulated a foraging environment populated by misperceiving agents. While the simulation did not model evolutionary processes, the benefit observed was argued to be adaptive, as it increased the quantity of resources harvested by the agents. In an evolutionary environment, increased resource collection should correlate with increased reproductive success. This hypothesis shall be investigated by creating an evolutionary simulation that closely models the foraging scenario described by Akaishi and Arita, with the aim of determining whether foraging agents receive an evolutionary benefit from misperception.

6.1 Akaishi and Arita's Simulation

Akaishi and Arita created a foraging scenario where the simulated agents gathered resources from locations in the environment. These agents developed a map of their surrounding environment to aid their foraging and the accuracy of an individual agent's map is affected by its misperceptions. The false beliefs caused by misperception affected the agents' foraging behaviour and this effect was found to benefit the agents.

This work follows the work of Akaishi and Arita instead of Doran's due to the more general manner in which misperception may benefit agents in Akaishi and Arita's simulated environment. Doran's simulations demonstrated a benefit from misperception only for specific types of misperception, which were those misperceptions that either discouraged movement or the specific shared incorrect belief of the true status of a resource node.

6.1.1 Description and Simulation Method

Akaishi and Arita's simulation models a two-dimensional non-torus grid world, populated by agents and resource nodes. The agents move around this environment collecting resources from the fixed resource nodes. Both agents and resource nodes are randomly distributed throughout the environment when the simulation is initialised. As agents move and observe their environment, they develop a map of the locations that they believe contain resource nodes. These beliefs are affected by misperception; which, in this simulation, is implemented as two different types of errors, each of which occurs with some probability when an agent observes its environment. Each agent uses its resource map to plan its movements between locations that it believes contain resource nodes.

The simulation executes for a fixed number of turns. During each turn, the agents activate in a random order, allowing each agent to iterate through its action cycle. This cycle possesses the same structure as the various models of the perceptual cycle previously described (see Section 2.6). A foraging agent's action cycle begins with an agent observing the cells within its visual range, which in this simulation are the eight cells adjacent to the agent. An agent may misperceive each cell it observes, determined by a global misperception probability shared by all members of the population. This misperception may affect an agent's understanding of the location or existence of resources at a given location with equal probability. Agents update their resource map with the observed

information, which updates the locations where they believe there are resource nodes. Once it has updated its resource map, an agent determines which is the closest resource node to its current location and makes this location its intended target. The agent then begins to move towards this location. If an agent is unaware of any resource nodes, it will instead move in a random direction. An agent's movement is constrained by a maximum movement speed, which limits how far it may move each turn. An agent stops moving once it has travelled the distance given by its maximum speed or arrived at its intended destination. If the agent's location contains a resource node then the agent harvests the resources from this node, reducing the quantity of resources held by the node to zero. This completes an agent's activation for the current turn. Once all the agents have moved, the current simulation turn ends and the next begins. This process continues until the simulation has run for the required duration.

Agents in Akaishi and Arita's simulation could move horizontally or vertically between locations, but not diagonally. Since only one agent may occupy a location at once, an agent may find its movements obstructed by the presence of another agent. The agents have no method of resolving or avoiding these obstructions (Akaishi, 2004) and when an agent's movement is obstructed it will (unrealistically) wait until for the obstruction to move out of its way. Misperception may alter an agent's beliefs, causing it to decide to move in a different direction and thereby reduce the impact of obstructions caused by adjacency between agents.

When an agent misperceives the existence of a resource node, its belief in the contents of the observed cell is affected. This causes it to observe an empty cell as containing a resource node and observe a cell containing a resource node as empty. A location misperception causes an agent to correctly observe the contents of a cell, but to incorrectly observe the cell's location. The agent creates a false belief by recording the contents of this cell at a random location in its resource map. Since each observation overwrites the existing information in an agent's map, new observations may correct false beliefs produced by previous misperceptions, while new misperceptions may also replace true beliefs with false beliefs. These false beliefs affect the locations that are believed to contain resources and thereby affect an agent's behaviour.

When an agent consumes a resource from its environment, the resource node is depleted and the agent's internal energy store is increased by the same amount. After the

resource node has remained empty for one turn, its available resources regenerate to a maximum amount, which defaulted to one resource unit. Akaishi and Arita incorrectly described the quantity of resources collected by an agent as its fitness and the population's fitness as the average quantity of resources gathered by the population. Since their simulation does not describe an evolutionary scenario, these values do not describe fitness, but instead the foraging success of the agents. Misperception benefits the agents if it increases their individual or collective success and Akaishi and Arita attributed this benefit to misperception's effects on the population's behavioural diversity.

Akaishi and Arita also investigated the effects of misperception during communication upon these foraging agents. Communication between agents allowed adjacent agents to exchange a belief about the locations of resources during the beginning of their turn. Misperception that affected the agent's communications was called indirect misperception and again occurred with a global probability. Communication was implemented by Akaishi and Arita to test the hypothesis that it reduced a population's behavioural diversity and thereby reduced the resources collected by the population. Indirect misperception was also hypothesised to increase this diversity and benefit the population in the same manner as direct misperception.

6.1.2 Results and Summary

Akaishi and Arita measured the quantity of resources the population collected and used this to determine whether misperception was beneficial. Their results show that populations of agents with global misperception probabilities of up to 10% gathered more resources on average than a population with 0% misperception. The agents' resource gathering was maximised with a global misperception probability of 1%. These results support the hypothesis that misperception can be beneficial. In this scenario, the benefit is attributed to an increase in the diversity of the beliefs held by the agent population, which causes an increase in the diversity of the population's behaviour. Behavioural diversity may encourage agents to avoid popular locations, with the agents' misperception potentially directing them into previously unexplored areas, thereby enlarging their search area. In this way misperception may benefit both the misperceiving agent, by guiding it away from popular foraging areas to less popular ones; and the non-misperceiving agents,

by reducing the competition encountered by the remaining agents in the popular foraging area.

A benefit of misperception that Akaishi and Arita do not discuss in this scenario is that of deadlock resolution. When two agents are mutually obstructing each other, each will wait for the other to move and such behaviour leaves both agents deadlocked. In other situations, one agent may find its movements obstructed by another agent, who may be obstructed by another agent and also cannot move. Misperception by either agent may alter its planned movement, causing it to move in a different direction and breaking up these deadlocks. Misperception may therefore benefit both the misperceiving agent and other deadlocked agents when it removes such obstructions and allows agents to resume their foraging.

Akaishi and Arita also observed beneficial misperception in the simulations with communicating agents, whose communications were also affected by misperception. Communication reduces the diversity of the agents' collective beliefs, as it allows the agents to share their beliefs. Indirect misperception that affects the communication between agents also creates false beliefs, thereby increasing the diversity of the population's beliefs and its behavioural diversity. In a population with only indirect misperception, an increase in resource collection was observed for indirect misperception probabilities up to 40%. When both direct and indirect misperception were simultaneously enabled, the population gathered more resources on average than when direct misperception was the only type of misperception. This demonstrates that the source of the misperception is not important when both sources can produce the same types of beneficial behaviour. Enabling both sources of misperception merely serves to introduce further diversity into the agent population's behaviour.

Akaishi and Arita's simulation demonstrated a foraging scenario where misperception during observation or communication could benefit agents. The benefit observed from misperception was incorrectly described as an evolutionary benefit, despite the lack of evolutionary mechanisms, such as reproduction, in their simulation. While they did identify a "fitness" benefit from misperception, this benefit was defined in terms of increased resource gathering; whereas fitness is usually defined in terms of reproductive success, such as a measure of the number of descendants. Nevertheless, in an environment where

reproduction does occur, increased resource collection by misperceiving agents should aid their reproductive success and therefore increase their fitness.

In an evolutionary environment where the probability of misperception is a phenotypical trait of the agents, the population of agents will eventually evolve to express probabilities of misperception that are optimal for their environment. If misperception does provide an evolutionary benefit, then a population of misperceiving foraging agents will evolve to a stable state with a misperception probability that is significantly different to 0%. Each agent's misperception probability may differ from the other agents' misperception probabilities, which may allow sub-populations with different misperception probabilities to develop. Agents with lower misperception probabilities may also receive an additional advantage from the more frequent mistakes of other agents with higher misperception probabilities. Akaishi and Arita's simulation modelled populations with a uniform misperception probability, which does not allow such exploitation to occur. Since Akaishi and Arita's simulation has shown that misperception does benefit these foraging agents, it seems reasonable that misperception is also beneficial in a similar evolutionary environment. Such a population should evolve towards an optimal probability of misperception that is stable among the population.

The benefit observed by Akaishi and Arita was attributed to the behavioural diversity misperception produced in the agent population. However, they did not further explore the hypothesis that behavioural diversity is responsible for the benefit to the foraging agents. I explore this hypothesis further in Chapter 7, by comparing the effects of different methods of producing behavioural diversity in the foraging agent population.

6.2 Simulation Design and Implementation

Confirming the hypothesis that misperception can provide an evolutionary benefit required the creation of an evolutionary simulation that modelled Akaishi and Arita's foraging scenario. The changes to the scenario were mainly limited to the behaviours of the agents themselves and consisted of adding mechanisms that allowed the agents to reproduce, consume gathered resources, and die from either old age or starvation. These modifications to the simulation allow the fitness of an agent or the population to be described in terms

of its number of descendants and thereby allow the evolutionary benefit of misperception to be analysed.

Much of the design and function of this new simulation remained the same as Akaishi and Arita's, with the changes to the scenario mostly affecting the agents. It was also decided to retain the agents' lack of deadlock resolution methods. While this did lead the agents to act in an unrealistic manner that is not found in biological systems, it did allow misperception to benefit agents in a manner similar to that identified in the original simulation.

6.2.1 Changes from Akaishi and Arita's Simulation

Akaishi and Arita's foraging scenario was modified to include sexual reproduction and death. Collected resources were consumed by agents to avoid starvation and to reproduce. In this manner an agent's reproductive success was governed by its foraging success, with agents that were less successful at foraging facing limited or no opportunities for reproduction. With this change, an agent's fitness is determined by the number of offspring it parented. The simulated agents possessed a chromosome encoding their individual misperception probability. Evolutionary mechanisms added to this simulation operated upon the misperception probabilities of the simulated agent population, through competitive foraging and reproduction. These changes were intended to model the reproductive mechanisms of sexually reproducing biological organisms (Appendix A contains a glossary of the elements of this simulation).

Individual Misperception Probabilities

Each agent possesses an individual misperception probability instead of sharing a single global misperception probability, which allows the individual misperception probabilities to compete against the others in the agent population. In an evolutionary system, the misperception probability becomes a heritable trait of the agents, which enables the agent population to evolve an optimal probability of misperception for their environment over time. Due to differing misperception probabilities, agents with lower misperception probabilities may benefit by exploiting the frequent mistakes of those with higher misperception probabilities, while those with higher misperception probabilities will likely develop more false beliefs and produce more diverse behaviour. The simulation was initialised with

agents having unique randomly generated misperception probabilities. These probabilities were low, as Akaishi and Arita's results suggested that lower misperception probabilities were more likely to prove beneficial.

Sexual Reproduction

The addition of sexual reproduction to the agents was the largest change made to the conditions of the foraging scenario. Sexual reproduction could occur between two adjacent agents, producing a new agent at an empty location adjacent to one of the parents. In order to reproduce an agent must be able to pay a fixed **Parental Investment Cost** from its surplus of gathered resources, which is an adjustable simulation parameter. Both potential parents must pay this cost equally. The agents are genderless, with any agent permitted to reproduce with any other, assuming that both could afford the parental investment cost. This differs from most real biological systems, where the reproducing organisms are of different genders and one gender commonly has a much larger parental investment than the other. Implementing a more realistic reproductive system would further complicate the foraging scenario, while providing further unnecessary variables to be controlled. For this reason the simpler system was used.

Each agent possesses a simple chromosome, containing its misperception probability and its mutation rate. The probability that an agent will misperceive each location it observes is dictated by its misperception probability. The mutation rate is the probability that the agent's misperception probability is mutated during reproduction. A crossover operation with mutation will create the new offspring's chromosome from those of its parents. This is a greatly simplified model of how sexual reproduction occurs in biological organisms and is suitable for this simulation.

Resource Consumption

Resource collection and storage by agents was the same as in Akaishi and Arita's foraging simulation. However, agents were required to consume, or metabolise, a quantity of their gathered resources each turn in order to stay alive. As in other biological systems, agents who have no resources to consume starve and eventually die. This change provides a mechanism to remove poorly foraging agents and introduces another type of evolutionary competition to the foraging scenario.

The amount of resources an agent consumes each turn is the **Basic Metabolic Rate** and this value is an adjustable simulation parameter. Higher metabolic rates require an agent to consume more resources each turn, thereby putting pressure on its foraging skills and reducing the amount of excess resources available to spend reproducing.

Death

Death is another element of evolutionary biological systems added to the foraging scenario. One cause of death that was simulated is starvation, which occurs when an agent lacks sufficient gathered resources to metabolise, per the Basic Metabolic Rate. Agents whose internal store of resources fall below zero die due to starvation. Starvation provides strong selective pressure against agents with poor foraging skills.

In addition to death by starvation, agents were also removed from the simulated population due to old age. Death by old age occurred when the agent's lifespan exceeds a determined simulation parameter, which was common to all agents. While giving all the agents an identical maximum lifespan is not completely realistic, it does prevent any agents with a slightly longer lifespan from having a small advantage over those with a shorter lifespan. The slightly longer lifespan of an agent could allow it more opportunities to increase its foraging or reproductive success, thereby complicating the direct comparison of the simulated agents. Death by old age therefore prevented extremely long-lived agents from existing in the simulation, whose longer lifespans could have provided them more opportunities to forage and reproduce than younger agents.

The agents' lifespan can also affect the population's evolutionary speed, as organisms with shorter lifespans will see a single set of parental genes producing less generations of offspring. At first, this seems advantageous for the simulation. However, shorter lifespans allow less opportunities for agents to misperceive, either positively or negatively. The agent's lifespan therefore represents a trade-off between evolution speed and opportunities for misperception. Since this simulation examines misperception, the agent population's lifespan will be configured to allow misperception opportunities.

Agent Density Cap

While the population size in Akaishi and Arita's simulation was fixed, the population size of an evolutionary environment should fluctuate over time as the organisms die and reproduce. However, such realistic behaviour conflicts with the fixed population of Akaishi and

Arita's foraging simulation. Their simulation permitted only one agent to occupy a location at once, which permitted agents to unintentionally obstruct each others' movements. Such obstructions may be resolved if the obstructed agents suffer from a misperception that alters their movements. If mutual obstruction is a potential problem for misperception to solve, then it is an important element of the foraging scenario that should remain in this simulation. However, if the agent population becomes too large then the agents are likely to have difficulties foraging without becoming obstructed.

To limit, but not prevent such obstruction from occurring, while retaining the one agent per location restriction, an agent density cap was introduced in an attempt to limit overcrowding. This cap places an upper bound on the number of agents that may exist in the simulation at one time and is an adjustable parameter of the simulation, the **Agent Density**. The cap allows agents to reproduce up to a point, while preventing the environment from becoming too densely populated.

When the population cap was reached, the agents were prevented from reproducing despite accumulating sufficient resources to pay the parental investment costs. Therefore any measure of an agent's actual offspring was an unreliable measure of its fitness, as the Agent Density cap can prevent reproduction. Instead, **potential offspring** was used as a measure of fitness, describing how many offspring an agent could afford to parent with the resources it had collected at the time it died. This value was zero for agents that starved to death, as they had no excess resources. Potential offspring are discussed further in Section 6.3.3.

Agent Communication

As Akaishi and Arita have noted, the communication of correct information reduces the diversity in the population's beliefs caused by misperception. Therefore, error-free communication is likely to counteract the effects of misperception to some degree. Since this simulation aimed to investigate the potential benefit of misperception to the foraging agent population, any mechanism that reduced the effects of misperception was undesirable.

Communication may also create further complications when it occurs between agents with highly disparate misperception probabilities. In such an event, each agent will receive beliefs that their own perception is unlikely to develop; a frequently misperceiving agent

will receive true beliefs, while the rarely misperceiving agent will receive false beliefs. These beliefs are the opposite of those that the agent will likely obtain from its own observations.

If misperception is beneficial, then such communication will also aid those who are not misperceivers, while also increasing diversity among that population's beliefs. While communication would further spread the effects of misperception within the agent population, it confused the effects of an agent's misperceptions on its behaviours. Therefore communication between agents was not implemented in this simulation. However, under different circumstances, the combined effects of misperception and communication may provide an interesting topic for future research.

6.2.2 Implementation

The simulation of the evolutionary foraging scenario was written in Java and makes use of the Repast toolkit (Repast Development Team, 2006). The toolkit provides an easy interface to run simulations manually and graph results in real time, along with methods to easily create graphical representations of the simulated environment. The graphical output of the simulation's state during execution was initially helpful; however, it was later disabled to reduce the computation time required by the simulation. The results of the simulation were then written to text files for further processing and subsequent analysis. Random number generation was performed by the Colt library (CERN, 2006), which is included with Repast.

The following discussion broadly describes how the simulation operates, focusing on the behaviour of the misperceiving agents (Appendix B describes the underlying algorithms and equations that fully detail the behaviour of the simulation and its agents).

The various input parameters that the simulation requires were passed in directly via the command-line when the simulation was executed. The four main simulation parameters explored were the **Agent Density**, the **Resource Density**, the **Basic Metabolic Rate** and the **Parental Investment Cost**. The other variable parameters of the simulation were the simulation world dimensions, the duration of the simulation, the initial agent population size, the maximum agent lifespan, the maximum agent speed, the regeneration delay of resource nodes, the meta-mutation rate and the mutation standard deviation. A brief explanation of the effects of these parameters now follows (Labels used

to describe parameters in parentheses match those used in Appendix B, which describes the simulation's operation).

- The total area of the simulation environment was determined by a single value, providing both the X and Y dimensions of the square grid environment (*GridSize*).
- The maximum agent density controlled the maximum number of agents that may exist in the environment, given as a percentage of the total environment area. This restricted the maximum number of agents that may simultaneously exist in the simulated environment (*AgentDensity*).
- Similarly, the resource density controlled how many resource nodes were created in the environment and was also given as a percentage of the total environment area (*ResourceDensity*).
- The basic metabolic rate defined how many resources an agent must consume each turn from its internal stockpile to stay alive (*BasicMetabolicRate*).
- The parental investment cost determined the amount of resources an agent must spend in order to reproduce (*ParentalInvestmentCost*).
- The duration of the simulation is the number of turns it was executed (*Turns*).
- The initial agent population determines how many agents were created in the environment when the simulation was initialised. This value may be set lower than the agent density cap (*InitialPopulationSize*).
- The maximum agent lifespan controls how long agents lived before dying of old age (*AgeLimit*).
- The maximum agent speed functioned as it does in Akaishi and Arita's simulation, describing the maximum number of locations an agent may move through each turn (*MaximumSpeed*).
- The resource node regeneration delay was how many turns a resource node waits before beginning to regenerate its resources (*RegenRate*).
- The meta-mutation rate parameter is the probability with which a new offspring's mutation probability was mutated during reproduction (*MetaMutationRate*).

- The mutation standard deviation parameter was the standard deviation of the normally-distributed random number generator used to alter a new offspring's misperception probability when it was mutated. This parameter therefore affected the scale of mutations to the misperception probability (*MutationSigma*).

The simulation environment was initialised with a random population of agents whose size was given by the initial agent population size. These agents had a randomly assigned misperception probability that was uniformly distributed between 0.0 and 0.1. The simulation then began executing for the desired number of turns (Algorithms 1 and 2 in Appendix B detail this process).

During each turn of the simulation, the agents were activated in a random order to iterate through their action cycle, ensuring that if there is a benefit to earlier or later activation, it was not consistently received by the same agents. This cycle differed little from Akaishi and Arita's original agent cycle (Algorithm 3 in Appendix B describes this process in detail). Each agent observed its environment and then updated its beliefs of where resource nodes exist. Misperception of a resource node's location or existence may affect these observations. The agent then determined which observed resource node was its target and began to move towards this location. The agent's movement for this turn ended either when it arrived at its intended location or it had travelled its maximum distance for this turn. After moving an agent consumed some of its stored resources, the amount given by the basic metabolic rate. Next the agent would attempt to reproduce with any other adjacent agents who were also capable of reproducing. An agent was deemed capable of reproducing if it could afford to pay the parental investment cost. If an agent looking to reproduce had a choice between multiple potential mates, then it would select the agent with the most resources – that is, the one it considered the most suitable. Note that this agent-derived understanding of “fitness” allows agents to quickly and simply assess which potential mates are more suitable, whereas later discussions of fitness (see Section 6.3.3) focus on a different value.

During reproduction the parental investment cost is deducted from each parent's resources and a new offspring is produced by a single-point crossover with mutation (Algorithm 4 in Appendix B describes this process in detail).

At the end of its movement cycle, an agent's age and stored resources were tested to determine whether it should die of old age or starvation. If so, the agent's details were written to the data logs and it was removed from the simulation environment.

After each turn completed, data from the current agent population was written to the data log files. The simulation continued in this manner until it had executed for the desired duration.

Output from the simulation was written into plain text files, which recorded relevant information from the life of the agents in the simulation. For each agent its misperception probability, mutation probability, agent lifespan, resources collected, birth time and parents were logged in one file. This provided a census of the agent population as the simulation executes, which can be analysed in greater depth once the execution of the simulation was complete. Another series of log files recorded snapshots of the current agent population and their attributes at a series of equally spaced intervals, which provided a simple measure of how the agent population was progressing in the simulation.

These text files were analysed and processed by external scripts, which averaged data over many simulation iterations or over the length of a single simulation. The scripts calculated the agents' average misperception probabilities and potential offspring. The scripts also combined data from several simulation runs, providing a limited statistical analysis of the data gathered. Such processing required the complete data from all the simulation runs and therefore could not be performed by the simulation itself.

This simulation aimed to test the hypothesis that misperception provides an evolutionary benefit to the foraging agents. This benefit should manifest itself in several different manners, each of which can be observed in the simulation's output. One indication of beneficial misperception is the development of a stable misperception probability in the agent population that is distinct from 0%. Akaishi and Arita showed that misperception probabilities of up to 10% yielded more resources for the agent population than no misperception and this should be observable as the evolution of a similar misperception probability. Furthermore, in the case that misperception is not beneficial, the agent population should evolve towards a misperception probability that is very close to 0%. Another indication of whether or not misperception provides an evolutionary benefit is the proportion of the agent population with misperception probabilities that are distinct from 0%, as the existence of a significant quantity of misperceiving agents suggests a benefit

from misperception. If the majority of the agent population evolves to misperceive, then this also supports the hypothesis of beneficial misperception. Another indication of beneficial misperception is an increase in the number of potential offspring that misperceiving agents may parent. Beneficial misperception should allow misperceiving agents to gather more resources and thereby have more potential offspring than non-misperceiving agents. Comparing the potential offspring of misperceiving and non-misperceiving agents should reveal whether such misperception was advantageous.

6.3 Simulation Parameters and Execution

There were many possible parameters that might have been explored by this simulation and which might have impacted the effects of misperception on the evolving population. It was simply not feasible to investigate every possible combination of parameters, due to the amount of simulation time this would have required and the quantity of data produced for analysis. Furthermore, the results for many of these parameter sets may be repetitive, and ultimately uninformative, if that parameter had little effect on the agent population. Therefore it was necessary to investigate misperception's effects in this simulation for a specific, limited range of parameters.

6.3.1 Simulation Parameters

Section 6.2.2 listed 11 different simulation parameters and described their effects on the simulation itself. Four of those are the main parameters that were focused upon: Agent Density, Resource Density, Basic Metabolic Rate and Parental Investment Cost. The remaining seven lesser parameters had a single constant value selected for the simulations. The reasoning behind the values selected for these parameters follows.

Major Parameters

The four main parameters investigated with this simulation were the **Agent Density**, **Resource Density**, **Basic Metabolic Rate** and **Parental Investment Cost**. However, reducing the variable parameters of the simulation to just these four parameters still left too many potential combinations of parameters to thoroughly investigate. These values were therefore combined into pairs and several pairs of values were then examined.

Agent Density	Resource Density
30%	25%
30%	10%
30%	5%
25%	15%
20%	10%
15%	10%
15%	5%
10%	5%
5%	5%

Table 6.1: Agent Density and Resource Density pairings

Basic Metabolic Rate	Parental Investment Cost
0.2	25
0.15	100
0.1	50
0.05	500

Table 6.2: Basic Metabolic Rate and Parental Investment Cost pairings

Agent Density and Resource Density were paired together, as they determine the level of competition between agents for resources (Table 6.1).

If misperception provides an evolutionary benefit, then it is expected that this benefit will be more noticeable in situations where there is greater competition between agents for resources. This would occur when the agents outnumbered the resource nodes that existed in the environment. It should be noted that Akaishi and Arita's original simulation used an Agent Density of 6% and a Resource Density of 5%. This simulation examined higher agent densities in order to investigate the effects of misperception in more competitive and congested situations. The pairings of agent and resource densities focused on values that allow for a more densely populated environment than that examined by Akaishi and Arita, along with a greater disparity between the amount of agents and resources. This disparity helped to increase the competition between agents for resources.

The Basic Metabolic Rate and Parental Investment Cost were also paired together since they both affect the agents' cost of living (Table 6.2). These pairings were selected to pair a high basic metabolic rate with a low parental investment cost and a low basic metabolic rate with a high parental investment cost to determine what effect more extreme values have. These extreme pairings were contrasted with two moderate pairs of values. A high basic metabolic rate or parental investment cost was expected to test the foraging abilities of the agents, increasing the selective pressure for agents that were effective foragers and against those that were not. High and low values were paired together to avoid creating pairings that were either too difficult or too easy for agents to survive in.

When combined, the nine resource density pairings and four cost of living pairings provided 36 different combinations of parameters that were explored. In order to obtain sufficient data for a statistical analysis, it was necessary to simulate each parameter set multiple times. Each of the 36 parameter sets was simulated 50 times, which required the execution of 1,800 simulations in total. While this was a large number of simulations, the total execution time of the simulations could be dramatically reduced through the execution of multiple simulations simultaneously, since each simulation is completely independent from any other. This was achieved by running the simulations on a computer cluster with many nodes that were capable of executing the simulations in parallel (see Section 6.3.2).

Minor Parameters

While Akaishi and Arita used a 50×50 grid for their simulations, this simulation instead examined a smaller 20×20 grid environment due to the much longer execution time required for larger grid sizes. Since the simulation aims to demonstrate an evolved benefit from misperception, it is expected that such a benefit is more likely to develop in more competitive environments. A more competitive environment has a higher ratio of agents to resource nodes, ensuring that there is more competition for access to those resource nodes. Since the agent and resource numbers are controlled by densities, the actual size of the grid environment does not affect the amount of competition between agents and different environment sizes are interchangeable.

The simulation's execution time is primarily dependent on the number of agents simulated, which was derived from the agent density and the size of the environment. As the agent and resource densities are percentages of the size of the environment, larger environments will contain many more agents, which increases the simulation's execution time too much. For example, an agent density cap of 30% in a 50×50 environment permits up to 750 agents simultaneously, while a 20×20 environment only permits up to 120 agents. The larger simulation world therefore had an execution time that was approximately six times longer. Since the quantity of agents in the simulation directly affects its execution time and since the selected parameter sets intentionally focus on higher agent densities, it was therefore necessary to reduce the size of the environment to reduce the simulation execution time.

The lifespan of the agents was set to 5,000 turns, to allow successful foraging agents sufficient time to gather enough surplus resources to reproduce, along with sufficient opportunities to misperceive. The duration of the simulation was set to 3,000,000 turns, to allow for sufficient time for the population of agents to converge on an optimal stable misperception probability. Initial experimentation indicated that 3,000,000 turns was a sufficiently long time-span for the simulation to reach a stable state.

In order for the simulation to remain as close as possible to Akaishi and Arita's version, the agents' maximum movement speed was set to three. This allowed agents to move through up to three locations during their action cycle. While the usage of a smaller environment might also suggest lowering the agent's maximum speed, due to the discrete nature of locations within the simulated environment, lowering the movement speed ensures that agents are more likely to observe the same locations during consecutive turns, which provides them with more opportunities to correct previous misperceptions, thereby reducing the impact of misperception upon their behaviour. The parameters for the resource nodes also remained the same as those used by Akaishi and Arita. The resource nodes stored a maximum of one resource unit and began to regenerate this amount one turn after it was harvested, at a rate of one resource regenerated per turn.

The meta-mutation rate was set to 5%, which ensured that a new offspring's mutation rate has a 5% chance of being mutated during reproduction. The mutation standard deviation, which affects the size of the mutations that affect a new offspring's misperception probability, was set to 0.02. Since the size of the mutations was normally distributed, there was a 95% chance that a mutated misperception probability was no more than ± 0.04 from the original value.

6.3.2 Parallelism and Cluster Usage

The execution time of each simulation run was predominately affected by the number of agents in the simulation environment, since every turn the simulation must activate and move each agent. For a given world size, simulations with higher agent densities — and therefore more agents — take longer to execute than those with lower agent densities. In some cases the high agent density simulations could take six or more hours to complete. This problem was further compounded by the 50 iterations per parameter set required

to collect statistically meaningful data. While some attempts to optimise the simulation code were made, there was a limit to how much execution time such optimisation saved.

The solution to this problem was to execute the simulations in parallel, which was performed on the Monash University computer cluster. This cluster consists of 52 nodes, each with a 3 Gigahertz Pentium 4 processor. With this system, assuming that all the cluster nodes are operational, the 50 iterations of each parameter set may be executed in parallel, reducing the execution time approximately to that of a single simulation. In actual practice, approximately 30 cluster nodes were operational at a given time, which still provided a significant reduction in the execution time of the simulations, compared to that of a single processor system.

Management and coordination of the cluster nodes to execute the simulation is performed by the EnFuzion software (Axceleon, Inc, 2003), a commercial derivative of Nimrod (Abramson, Sobic, Giddy and Hall, 1995). The EnFuzion software greatly aided the implementation of parametric simulations and reduced their total execution time by executing multiple simulations in parallel.

EnFuzion was provided with a run file that described all the parameters and their possible values, which it then used to enumerate the complete set of parameter value combinations simulated. From this set of parameter values, each parameter set was then allocated by EnFuzion to an available cluster node and a simulation was executed with those parameters. When complete, the results were copied back to a specified location. Once all the simulations were complete, the results were collated and analysed.

6.3.3 Fitness Measure

Fitness describes an organism's capability to produce descendants. Fitter organisms are better adapted to their environment and therefore more capable of producing offspring, who are also well adapted to this environment. In their simulation, Akaishi and Arita considered resource collection to be a crude measure of fitness. In this evolutionary simulation, an agent's number of offspring (children) or second-generational offspring (grandchildren) are considered to describe its fitness.

However, as noted earlier, the population cap prevented reproduction whenever there was no available space for the new offspring. This affected an offspring-based measure of an agent's fitness, as an agent who is capable of reproducing may be prevented from

doing so by the simulation rules. For two agents with comparable quantities of gathered resources, one agent could therefore appear “fitter” simply because it was lucky enough to have been permitted to reproduce when the population was smaller.

To avoid this problem, fitness in the simulation is defined as the number of **potential offspring** an agent can produce (Equation 6.1).

$$\text{Potential Offspring} = \frac{\text{Total Resources Gathered} - (\text{Age} \times \text{BMR})}{\text{Parental Investment Cost}} \quad (6.1)$$

This value calculated how much of the resources collected by an agent were metabolised during its lifetime to survive, and subtracting these from the total gathered during its lifetime determines the surplus resources the agent could have allocated towards reproduction. This surplus is then divided by the Parental Investment Cost to determine the potential offspring that the agent could have parented if circumstances had allowed it. This value is always greater than or equal to zero.

6.4 Results

The simulation was executed on the computer cluster for 50 iterations of each of the 36 previously described parameter sets. The results from these simulations were analysed for several different indications of a benefit from misperception in this evolutionary environment. One was the average misperception probability that evolved within the agent population for each parameter set. This value would be optimal for the population in the environment and a misperception probability that was significantly different from 0% would support the hypothesis that misperception was beneficial. Another indication of a benefit was a comparison of the proportions of the population considered to be misperceivers against those that were not. If misperception was beneficial, then a significant proportion of the population would have misperception probabilities that were not 0%. Another indication of a benefit from misperception was the quantity of potential offspring produced by misperceiving and non-misperceiving agent sub-populations. If misperception did provide an evolutionary benefit to the foraging agents then their more effective foraging would have increased their potential offspring.

In order for misperception to benefit agents, misperception needs to occur frequently enough to affect the agent’s beliefs and thereby affect its actions. Therefore, such a

determination must be made on the basis of the agent's misperception probability. While it is easy to identify agents with misperception probabilities of 0%, some criterion is required to distinguish between those whose misperception probability is non-zero. However, merely observing agents with non-zero misperception probabilities does not necessarily indicate beneficial misperception. Given the role of mutation in this simulation such agents may be the offspring of agents with 0% misperception probabilities and their existence therefore does not prove the hypothesis of beneficial misperception. Of course, it is also possible that such agents are the offspring of agents with higher misperception probabilities that mutated close to 0%. Therefore agents with non-zero misperception probabilities can be sub-divided into two groups: those whose misperception probability is extremely unlikely to have mutated from 0% and those whose misperception probability may have mutated from 0%. Of these two sub-populations, the performance of the former will indicate a benefit to misperception.

During reproduction, the misperception probability of new offspring was potentially mutated by adding a small normally distributed delta value, which had a mean of 0.0 and a controllable standard deviation (σ), given by the mutation standard deviation parameter. Therefore 95% of the mutated misperception probabilities were within $\pm 1.96\sigma$ of the original misperception probability inherited from one parent. This range of misperception probabilities will be referred to as the **mutation range**. In all the simulations the standard deviation was 0.02, which gives a mutation range of 0.0392 or 3.92%. One agent is unlikely to have inherited its misperception probability from another if the difference between their two misperception probabilities is greater than the mutation range. This determination allows the agents with non-zero misperception probabilities to be sub-divided into two categories. Agents were considered to be misperceivers if their misperception probability was outside this mutation range from 0%, i.e. greater than 3.92%. Agents whose misperception probability was 0% are non-misperceivers, while those whose misperception probabilities were between these two values were considered to be partial misperceivers. These agents misperceive, but they may be the mutated offspring of agents who were not misperceivers. Hence their existence neither supports nor disproves the hypothesis of an evolved benefit from misperception.


6.4.1 Average Misperception Probability

To examine the stable value of misperception that the population evolved, the population's average misperception probability from the final turn of the simulation was examined. The average misperception probabilities were calculated for each of the parameter sets and plotted in Figure 6.1. The average misperception probability is plotted on the vertical axis, while the two pairings of density values and cost of living are plotted on the horizontal and depth axes respectively. Interesting parameter sets have been numerically labelled throughout the results discussed here and in Chapter 7 in order to clarify the discussion. Several points on this plot are marked with dots and numbered, with those numbered 1–7 identifying parameter sets where the average misperception probability is above the mutation range of 3.92%. This is especially true for parameter sets 1 and 2. Parameter sets 8 and 9 are also labelled, due to the existence of substantial sub-populations of misperceivers, which shall be discussed later. Parameter set 1 shows a high average misperception probability of nearly 10%.

The lowest average misperception probability here is 1.19% and the majority of the average misperception probabilities are between 1% and 3%. Akaishi and Arita's previous results suggested that the optimal misperception probability was 1% and while the average misperception probabilities are not 0%, the low but non-zero misperception probability populations may be agents who are the mutated offspring of agents with 0% misperception probabilities (Table C.1 in Appendix C contains the data values for this plot).

The majority of the points where the average misperception probability was above the mutation range were clustered along a pair of raised ridges that run across the edges of these graphs. The values along these ridges were not all above the mutation range. One ridge runs across the parameter sets where the agent density and resource density were both 5%, while the other runs across the points where the basic metabolic rate was 0.05 and the parental investment cost was 500. The agent and resource densities of 5% were the closest pair to the values used by Akaishi and Arita in their simulation. The higher misperception probabilities across the rear of the graph occur with the higher parental investment cost, which increases the selective pressure for more effective foragers.

While these results suggest a benefit from misperception, it is possible that the average misperception probabilities are largely affected by agents who fall within the mutation range. Such agents may therefore be the mutated offspring of agents with misperception



Agent Density, Resource Density Pairs

Figure 6.1: The Average Misperception Probabilities for each parameter set. The points marked with a circle and numbered 1–7 indicate parameter sets whose average misperception probability was above the mutation range, potentially indicating a benefit from misperception. The points marked with a triangle and numbered as 8, 9 and 10 indicate three parameter sets with average misperception probabilities below the mutation range that may be of subsequent interest. Misperception probabilities above the mutation range (i.e. greater than 3.92%) are coloured to show which parameter sets are evolving higher misperception probabilities.

probabilities of 0%, meaning that their existence is not due to any benefits of misperception. To determine whether the identified populations of misperceiving agents are statistically distinct from the agents whose misperception probabilities fall under the mutation range, the 95% confidence interval of the average misperception probabilities was computed and this data is shown in Figure 6.2.

Of the 7 numbered parameter sets with average misperception probabilities above the mutation range, only 1 and 2 are statistically distinct from the mutation range. A one-sample t-test also confirms that the average misperception probabilities of these two parameter sets are statistically distinct from the mutation range. This confirms that the bulk of the agents in these populations had misperception probabilities significantly different to 0%, therefore indicating a benefit from misperception. However, these two parameter sets differ greatly in their parameters, as 1 simulated an environment with high agent (0.3) and resource (0.25) densities and a low Parental Investment Cost (50), while 2




Figure 6.2: The Average Misperception Probabilities and 95% Confidence Intervals for each parameter set, compared to the mutation range. Of the numbered parameter sets, only 1 and 2 differ from the mutation range in a manner that is statistically significant. Hence the vast majority of those populations are agents that were not the mutated offspring of agents whose Misperception Probability was 0.0%.

simulated an environment with low agent (0.05) and resource (0.05) densities and a very high Parental Investment Cost (500).

Of the parameter sets numbered 3–7 the lower range of their confidence intervals fall beneath the mutation range, indicating that some proportion of those agent populations may not be the result of beneficial misperception. In these cases the suggested benefit of misperception is not statistically significant. The parameter sets whose agents have confidence intervals entirely within the mutation range exist as some combination of the mutated offspring of agents with misperception probabilities of 0% and the offspring of agents benefiting from small misperception probabilities. However, with the current simulation data and mutation range it is impossible to clearly distinguish between these two sub-populations.

Were misperception to provide no benefit at all, Figure 6.1 would show a very flat plot where the average misperception probabilities are never significantly above the mutation range. Furthermore, the parameter sets whose results are not statistically significant suggest the possibility of a benefit from misperception, albeit one that is overshadowed in those instances by misperception probabilities of 0%. Overall, the average misperception

probabilities observed in the agent populations suggest that misperception can provide an evolutionary benefit.

6.4.2 Sub-populations of Misperceiving Agents

Another method of assessing the benefit of misperception is to compare the relative sizes of the misperceiving sub-population and the non-misperceiving sub-population. This was achieved by dividing the population into three sub-populations based upon their misperception probabilities. One sub-population contains agents whose misperception probability is 0%; another contains those whose misperception probability is within the mutation range; and the last contains agents whose misperception probability above the mutation range. If misperception is beneficial, then there should be many agents whose misperception probabilities are above the mutation range.

The total agent population from all 36 parameter sets was divided into these three sub-populations, as shown in Figure 6.3. Here a benefit from misperception should manifest as an increased proportion of agents with misperception probabilities that are above the mutation range. The previously labelled parameter sets are again marked here, to identify the effects of higher misperception probabilities on the relative proportions of the populations. These sets show larger proportions of the population with misperception probabilities that are above the mutation range, but most of the parameter sets did not previously demonstrate statistically significant results. Also common to these parameter sets is the smaller percentage of agents whose misperception probability was 0%, which was commonly less than 10% of the population. These results for the numbered parameter sets support the hypothesis that misperception is beneficial.

Parameter set 8 was labelled here due to its larger misperceiving sub-population and small non-misperceiving sub-population. While this population's average misperception probability was not above the mutation range, the proportions of its sub-populations are similar to that of parameter set 7, indicating a similar benefit from misperception.

The two parameter sets with average misperception probabilities that are statistically distinct from the mutation range unsurprisingly have the highest proportion of agents with misperception probabilities above the mutation range. Parameter set 1 has more than 80% of its agents above the mutation range, while 2 has more than 60% of its agents above the mutation range. They also show the smallest number of agents with 0% misperception

probabilities and with misperception probabilities within the mutation range. This result clearly demonstrates the benefits of misperception in those parameter sets.




Figure 6.3: Proportions of agent populations grouped by misperception probability. The parameter sets with misperception probabilities that were above the mutation range (1–7) here demonstrate that a large proportion of their populations have misperception probabilities in that sub-population. Parameter set 8 was included with these other parameter sets as its population proportions are similar to those of parameter sets 4, 5, 6 and 7, indicating a similar benefit from misperception.

Also of interest is the sizable proportion of the population whose misperception probabilities were within the mutation range, which typically appeared to consist of approximately 50% of the population. While these agents did not have misperception probabilities that were above the mutation range, they did have a non-zero probability of misperception; which would have affected their behaviour, albeit to a lesser degree. While the proportions of agents with smaller misperception probabilities does not definitively prove that misperception is beneficial, it does disprove the argument that 0% is the optimal misperception probability for the foraging agents.

If misperception does not benefit foraging agents, then the proportions of misperception probabilities in the agent populations should differ to those shown in Figure 6.3. Specifically, the non-misperceiving sub-populations should be very large, the minor misperceiving sub-populations should be small, while the major misperceiving sub-populations should be very small or non-existent. Such results were not seen for any of the parameter sets, suggesting that 0% is not the evolutionarily optimal misperception probability in this foraging scenario for any of the explored parameter sets.

6.4.3 Average Potential Offspring

Another indication of a benefit from misperception in this simulated environment is evidence that misperceiving agents are more effective foragers, which permits them to parent more offspring. Since the agent density cap may prevent reproduction, agents that were effective foragers in the simulated environment should have more potential offspring than those that were not. As explained previously, an agent's potential offspring is a measure of the agent's fitness, calculating how many offspring an agent could afford to parent from its surplus resources. This value can only be calculated at the end of the agent's life, when the values for the agent's total resources gathered and its age have finalised (see Equation 6.1).

Once again the populations from the 36 parameter sets were divided into sub-populations based upon their misperception probabilities (Figure 6.4). The labelled parameter sets that suggest a benefit from misperception are again identified. Previously the labelled parameter sets have suggested a benefit from misperception, albeit not always at a statistically significant level. In the majority of parameter sets, there is little difference between the average potential offspring of the agents in the different groups. Here all three groups of the population each averaged about two potential offspring. Given the measured confidence intervals there is no discernible advantage or disadvantage in terms of potential offspring to misperception by the three sub-populations in those instances.

In six parameter sets the highly misperceiving sub-populations are significantly fitter than the other sub-populations (parameter sets 1, 2, 3, 5, 9 and 10). Parameter set 9 and 10 were added to the labelled parameter sets here due to the statistically significant difference in the average potential offspring received by their misperceiving sub-populations. The difference in average potential offspring between the sub-populations is statistically significant, since the confidence interval for the agents above the mutation range does not overlap that of the non-misperceiving sub-population.

Misperception's benefit may not be restricted to the highly misperceiving sub-population. Parameter sets 1 and 3 both demonstrate a significant increase in the potential offspring for their other two sub-populations, suggesting that those agent sub-populations were also benefiting from the misperception of other agents. In those parameter sets agents other than the misperceivers may be benefiting from the misperception, although it is




Figure 6.4: Average potential offspring of the sub-populations with 95% confidence intervals. The previously labelled parameter sets show several instances where the misperceiving sub-populations' average significantly more potential offspring than those that do not misperceive. Parameter set 9 and 10 are labelled here due to the statistically significant advantage in potential offspring that their misperceiving sub-population possess.

also possible that those parameter sets also yield more potential offspring among all sub-populations.

If misperception does not provide an evolutionary benefit and detrimentally impacts the foraging capabilities of the agents, then the two misperceiving sub-populations should have less potential offspring than the 0% misperception sub-population. This behaviour was not observed in any of the parameter sets investigated. Instead, the results show either nearly equivalent potential offspring between the sub-populations for the parameter sets, or increased potential offspring of the misperceiving sub-populations in some parameter sets. This disproves the argument that misperception is detrimental to the foraging and reproductive success of the agents in this foraging scenario. Since the misperceiving sub-populations were as fit as the non-misperceiving sub-populations, it seems that the misperception-affected agent sub-populations were still as capable of foraging as effectively as the non-misperceiving sub-populations.

6.4.4 Source of the Evolutionary Benefit

One potential source of misperception's evolutionary benefit is an individual benefit received by a misperceiving agent when it avoided a deadlock or a densely foraged area and headed to an area that was sparsely populated. For this method to be the sole source of

the benefit misperception provides, then it must outweigh the individual cost imposed by any detrimental misperceptions. These misperceptions may direct the agent away from resources or into congested areas. To produce an individual net benefit, the individual benefit from misperceptions must outweigh the penalties imposed by misperception. This requires misperception to either benefit agents more frequently than it harms them or to provide significant benefits with only minor penalties. When misperception occurs it is more likely that it will either provide no effect or a detrimental effect than provide an immediate benefit, since the conditions required for an individual benefit are much more situational. The individual penalties imposed by misperception are also likely to be larger than the benefits it provides. Since neither of these conditions were met reliably in this simulation, it is unlikely that individual selection alone is responsible for propagating misperception in this foraging scenario. It may be possible for a few lucky misperceiving agents to be selected in such a manner; however, they are more likely a rare exception.

Akaishi and Arita's interpretation of the benefit misperception provides to these foraging agents states that misperception increases the population's behavioural diversity. Agent behaviour is diversified by convincing agents to forage away from popular resources nodes, which have more competition for access and may attract clusters of agents. This may lead to increased competition over the limited resource and in severe cases a "traffic jam", where agents prevent or restrict each others movements. Misperception may misdirect agents away from resource nodes, thereby reducing congestion around the resource and restoring access to the resource by others. This is an altruistic act, with the agent giving up its chance to gather the resource for an increased chance for its relatives to gather the resource. This behaviour describes an inclusive fitness (kin selection) advantage to the agent's misperception. If the misperceiving agent's misperception probability is not too high, it will not significantly affect the agent's search for a new, less congested resource node.

Since new offspring are born adjacent to one of their parents, there is an increased likelihood that agents who forage in the same region are related to each other. The increased probability of relatedness between nearby agents thereby increases the chances that any misperception will benefit relatives. However, despite this increased chance of relatedness, there is no guarantee that the non-individual benefit from misperception will accrue to the misperceiver's kin. In such cases, agents who are unrelated to the misperceiver can

benefit from the reduced competition for resources in the area the misperceiver leaves. Such cases will provide selective pressure against misperception in the population. This does not prevent kin selection from operating in this foraging scenario; however, it does decrease the overall benefit that misperception can provide through kin selection, as some quantity of misperceptions will not provide a net benefit to related misperceiving agents.

6.4.5 Deadlocking Behaviour and Misperception

Since agents are capable of clustering around resource nodes and locations may only contain one agent at a time, congestion occurs around resource nodes in the simulation, as in Akaishi and Arita's earlier simulation. This can and does lead to situations where agents may obstruct each others' intended movements. In an example, one agent may intend to move onto a resource node, while that node is occupied by an agent intending to move off the node. This congestion, especially around resource nodes, is detrimental to the agents as it delays both their movement and resource collection. At some point congestion changes from slowing agent movement to preventing it altogether, and this is called deadlock. Akaishi and Arita have stated in communications (Akaishi, 2004) that there was no deadlock resolution method available to the agents, making this behaviour an intentional element of the simulation. Deadlock is much more likely to occur in simulations with a higher agent density. In the simulation, agents will wait indefinitely for obstructions to clear, with an obstructed agent potentially waiting until it dies of old age. This behaviour is highly unrealistic, as biological organisms do not contest resources indefinitely, instead preferring to abandon impossible movements or unobtainable resources. The lack of deadlock resolution behaviour in the agents and the restriction of one agent per location were retained in order to more closely model the foraging scenario described by Akaishi and Arita. One reason for this is that misperception is a potential solution to deadlocked agents, since it may alter their beliefs and cause movements that release the deadlock.

A deadlock scenario may involve one agent who is on a resource node and one or more others who are adjacent to the node and wish to move onto it. A simple example of such a situation is shown in Figure 6.5, where each agent is prevented from moving by another agent. Without misperception one agent is trapped on the resource, while the other three will wait indefinitely for it to leave the resource. Misperception can resolve this deadlock.

If the agent on the resource misperceives it may decide to move in a different direction that is not obstructed by other agents, such as South or West. This misperception will altruistically benefit the waiting agents, as the resource is now available to them. If the agent to the East of the resource should misperceive and then move away from its current location, then this will enable the agent on the resource node to vacate the node. This leaves the node available for the agent to the North to move onto it. Should the agent to the North-East of the resource node misperceive then it may benefit individually from the misperception, if its misperception guides it away from the deadlock. There is also the potential for such a misperception to provide an altruistic benefit, as it frees up locations near the resource node for others who may be attempting to access the resource node.




Figure 6.5: The agents (A) clustered around a resource node (shaded centre location) are mutually deadlocked, with each agent's intended movements prevented by another agent.

If the simulation was modified so that agents no longer obstruct each other, deadlocking would be prevented. However, a similar problem arises in the allocation of the resource to the agents who are present. A resource node's popularity determines how many agents will be present at a given time attempting to harvest it. Increased competition for a node will reduce any share that an agent will receive from an even distribution of the resources and also reduces an agent's chances of securing the resource through a contest. In such a scenario, misperception may still benefit agents when it convinces them to forage in areas with less competition for resources. This hypothesis could be explored in future research, which would investigate whether misperception would still benefit agents in this more realistic foraging environment. Such a scenario was not explored here, as the aim was to test Akaishi and Arita's hypothesis that their foraging scenario exhibited an evolutionary benefit of misperception.

6.4.6 Comparing the Statistically Significant Results

Of the parameter sets examined, only sets 1 and 2 both demonstrate a statistically significant benefit from misperception in both the average misperception probability and the average potential offspring. Given only two data points and the difference in their parameter sets it is difficult to speculate upon how the parameters of the simulation allowed the evolution of beneficial misperception. I hypothesise that the simulation's high agent density provided a situation where misperception helped affected agents to better avoid deadlock, while the low parental investment cost ensured that it was easier for them to reproduce, although this cannot be confirmed from the simulation results.

There were six parameter sets (1, 2, 3, 5, 9, 10) that showed a statistically significant benefit to misperception based upon average potential offspring. The parameters (agent density, resource density, basic metabolic rate, parental investment cost) of each of these parameter sets are:

1 $[0.3, 0.25, 0.1, 50]$

2 $[0.05, 0.05, 0.05, 500]$

3 $[0.05, 0.05, 0.1, 50]$

5 $[0.05, 0.05, 0.15, 100]$

9 $[0.05, 0.05, 0.2, 25]$

10 $[0.15, 0.1, 0.1, 50]$

Four of the six sets share the same agent and resource density pairing of 0.05 and 0.05. With the lowest agent density there would be much less deadlocking between agents for misperception to overcome, but misperception may have encouraged agents to better explore their environment, thereby encountering more of the sparsely located resource nodes to aid their survival and reproduction. This pair of values also closely matches that originally used by Akaishi and Arita in their simulation.

6.5 Summary

If misperception is always detrimental, negatively impacting the individual or collective survival of entities, then evolutionary pressure should remove misperceiving entities from

such a population. Based upon the expected negative individual outcomes of misperception, conventional wisdom supports the argument that misperception is detrimental to the individual. Even if misperception is only a small handicap to individual entities, over time selective pressures should reduce its representation in the population. However, the results of this simulation suggest that misperception is sometimes providing an evolutionary benefit, which is statistically significant for some parameter sets. Furthermore, these results do not match those that would be expected if misperception was detrimental. This supports Akaishi and Arita's claim that misperception can provide an evolutionary benefit.

6.5.1 Average Misperception Probabilities

One indication of an evolutionary benefit from misperception is the average misperception probabilities evolved by the agent populations. In two parameter sets the agent populations evolved stable misperception probabilities that were statistically different from the mutation range. Several other parameter sets evolved misperception probabilities that were not statistically significant but such results do suggest a lesser benefit in those instances. The populations for the majority of the parameter sets produced misperception probabilities that were typically between 1% and 3%. While not statistically significant, these values do not support the argument that misperception is harmful to the agent populations

6.5.2 Relative Sizes of Sub-populations

Dividing the agents into sub-populations based upon their misperception probability enabled the identification of the relative proportions of the three agent sub-populations. In the parameter sets where the population evolved a statistically significant non-zero misperception probability, the sub-population of agents with misperception probabilities above the mutation range made up sizable proportions of the population. In these two instances, the non-misperceiving sub-population was also very small. This also supports the argument for beneficial misperception, as if misperception is either ineffective or detrimental, then it should not be so well represented among the agent populations.

6.5.3 Potential Offspring

Calculating the potential offspring of the sub-populations also provided support for the hypothesis of an evolutionary benefit to misperception. In six parameter sets, the misperceiving sub-populations had a statistically significant advantage in the amount of potential offspring over the non-misperceiving sub-populations. In many of the other parameter sets there was little difference in the potential offspring of the three sub-populations. While this similarity between sub-populations does not support the argument of a benefit from misperception, it does argue against any penalty incurred by the misperceiving populations, as they had as many potential offspring as the non-misperceiving sub-population. In the cases where it appears there is neither an advantage nor disadvantage to misperception, it may be that these beneficial effects are not noticeable in the agent population. Since each of an agent's misperceptions may have either a positive, negative or neutral impact on its behaviour, misperception cannot have no effect on a population containing misperceiving agents. What may actually be occurring is that the advantages and disadvantages of misperception, as seen affecting the quantity of potential offspring, are approximately equal and cancel each other out.

6.5.4 How is Misperception Beneficial?

Akaishi and Arita claimed that misperception's benefit comes from increasing the diversity of the population's beliefs and behaviour. This suggests that here misperception benefits other agents rather than the misperceiving agent. Since misperception is unable to consistently provide an individual benefit in the evolutionary foraging scenario, misperception must benefit other agents affected by the misperception. The agents receiving the external benefit from misperception may be kin of the misperceiving agents, which allows the altruistic act of leaving a congested area to provide a benefit through inclusive fitness. If the agents who benefit from the misperception are not related to the misperceiving agent, then misperception is not increasing its inclusive fitness. This benefit to non-kin agents may be argued to provide a further benefit through group selection or multi-level selection, however these simulations cannot confirm this hypothesis. Misperception may also sometimes contribute an individual benefit in cases where the misperceiver is directed to a new foraging area with less competition for resources than its previous area.

If increased behavioural diversity is actually responsible for the benefits these misperceiving agents received, then it follows that any other mechanism that introduces behavioural diversity should similarly benefit a population of foraging agents. The next chapter explores this hypothesis further.

Chapter 7

Behavioural Diversity and Misperception's Benefit

7.1 Introduction

Akaishi and Arita hypothesised that the benefit from misperception could be adaptive if it increased the diversity of the population's collective beliefs and thereby increased the diversity of the agent's behaviour. Misperception by individual agents produces new erroneous beliefs, thereby introducing diversity into the population's collective beliefs and, likely, introducing diversity into the behaviour of agents within the population.

In the environment modelled in Akaishi and Arita's foraging simulation, increasing the behavioural diversity of the population is believed to reduce the direct competition between agents for access to popular locations or resources. This benefit may be received either by agents whose misperception guides them away from a popular resource or by agents at a popular resource that a misperceiving agent avoids. These benefits may be acquired simultaneously. If misperception benefits agents specifically by increasing their behavioural diversity, then this benefit can also be achieved by any other mechanism that introduces similar diversity into the behaviour of the agent population.

In this chapter the hypothesis that misperception produces behavioural diversity, which is beneficial to the foraging agents, is tested by modifying the evolutionary simulation of Chapter 6, to replace misperception as the underlying source of any behavioural diversity. Also from Akaishi and Arita's hypothesis, any methods that reduce the behavioural diversity of the population should reduce the fitness of the population. Such changes will

affect the agents' foraging behaviour, allowing their effects on the agent populations to be compared to those of misperception.

7.2 Method

If beneficial behavioural diversity is the product of the foraging agents' misperceptions, then other mechanisms that introduce this diversity into the agents' behaviour should also provide a similar benefit to misperception. Since misperception is modelled as a random error that affects an agent's understanding of its environment, other types of random errors that produce similar effects may be substituted for it and their effects compared against those of misperception.

The behavioural diversity hypothesis will be tested by altering the foraging behaviour of the agent population and observing the resultant effects upon the population's fitness. These changes will be implemented by altering the agents' foraging behaviour with several different mechanisms that are intended to affect the behavioural diversity. Akaishi and Arita's hypothesis argues that these mechanisms will provide a noticeable advantage or disadvantage, depending upon whether the behavioural diversity is increased or decreased.

The following discussion details the terms used to describe aspects of this simulation (Appendix A contains a condensed glossary of these terms).

7.2.1 Foraging Methods

The simulation from Chapter 6 models a population of misperceiving agents that forage for resources in their environment. The misperceiving agents are said to suffer from **misperception-affected foraging**, where each misperception affects an agent's belief of the location or existence of a resource node.

The existing simulation was modified to introduce three new foraging behaviours to the agents, which will alter the way in which they gather resources from their environment. These changes are intended to affect the behavioural diversity of the agent population, thereby providing an effect similar to that observed with misperception. Each of these alterations to the foraging methods are intended to produce behavioural changes in the individual agents.

Misaction-affected foraging

Random errors in an agent's behaviour are one method through which behavioural diversity may be introduced into the agent population. Misaction causes an agent to fail to correctly perform its planned action. Altering agents by replacing misperception with misaction produces agents that suffer from **misaction-affected foraging**. This change moves the random error from the start of the agent's perceptual activities to the end of its decision and action cycle. Unlike misperception, each misaction is temporary and will only affect an agent for a single turn, allowing the agent to recover from its misactions in the future.

Misaction-affected foraging will be implemented in the simulation by removing misperception and using the random error provided by misperception to instead trigger misaction when the agent attempts to move. When such an error occurs, an agent will select a random unobstructed direction other than its intended direction and then move in this direction. If the agent cannot misact by moving in a different direction, it will instead misact by remaining at its current location. The value in the agents' chromosome for the misperception probability will be reused for the misaction probability (Appendix B.2 describes the behaviour of these agents in greater detail).

Reflexive-foraging

Behavioural diversity may also be introduced by any other method that randomly influences an agent's actions. In the most extreme case, an agent would perform a random walk around its environment, eating only when it encounters resources within its visual range. This behaviour completely removes the deterministic decision-making cycle previously used by the agents and replaces it with a highly stochastic behaviour. Such a change would be expected to produce an extremely high level of behavioural diversity. Implementing and simulating such a high level of behavioural diversity will demonstrate whether too much behavioural diversity in this foraging environment is detrimental for the agent population.

This foraging behaviour, called **reflexive-foraging**, will completely replace the decision-making processes of the affected agents with a simpler behaviour of random movement. The agents will only move in a non-random direction when they are moving to a resource node they have observed within their immediate vicinity. These agents will still possess a

value for misperception in their chromosome; however it will not influence their behaviour and therefore have no selective pressures operating upon it. Therefore in a population of reflexive-foraging agents any measure of their misperception probability is meaningless (Appendix B.3 describes the behaviour of these agents in greater detail).

Perfect-perception foraging

If behavioural diversity is the underlying source of the benefit observed in Akaishi and Arita's simulation, then reducing the behavioural diversity should produce a measurable reduction in the fitness of the agent population. This may be achieved by having a population of agents who do not misperceive. This method is called **perfect-perception foraging** and has the agents utilise the same decision-making methods as the misperceiving or misacting agents. This makes them a control population for comparison against the populations exhibiting behavioural diversity. According to the hypothesis that behavioural diversity is the source of the benefit observed in the misperceiving agents, it is expected that a population of agents with perfect perception and no diversity should perform comparatively worse, or at least no better, than the other foraging methods.

The perfect-perception foraging method mostly uses the same foraging behaviour as the misperception and misaction-affected agents. The major difference is that all the agents will have a misperception probability of 0%, causing their behaviour to be nearly deterministic. This is achieved by initialising the simulation with a population of agents whose misperception probability is 0%, and preventing the misperception probability of any new offspring from mutating away from 0%. Due to this modification, measurement of the misperception probability is also meaningless here, as all the agents share the same probability of 0% (Appendix B.4 describes the behaviour of these agents in greater detail).

7.2.2 Simulation Execution and Results

These simulations also used the same simulation parameters as those detailed in the previous chapter, allowing a direct comparison of the various foraging methods. The foraging method used by a simulated population was controlled by a parameter of the modified simulation.

These simulations were also executed for 50 iterations for each of the 36 parameter sets, repeated across each of the three new foraging methods. In total, 5,400 simulations were

executed, with the computer cluster and EnFuzion software used to perform and manage the simulations. Usage of the cluster dramatically reduced the time spent performing the simulations, which allowed the simulations to explore a larger parameter space. Without access to the cluster, the explored parameter space would have been greatly reduced, potentially leading to the collection of a non-representative set of results.

Since the modified foraging methods were the only modification to the newer simulation, the same simulation data was again collected. This included a continuous log of the agents in the simulation, which detailed their lives, total resources collected and their amount of direct and second-generation offspring. Regular snapshots of the agent population were also collected, describing in detail the agent population at that point in the simulation's execution.

Once all the simulations were complete, several different processing scripts were run on the collected raw data. These scripts collated and processed the relevant information for analysis from the data produced by each iteration of the simulation. Values such as the potential offspring for individual agents were calculated and then averaged across all the iterations for each parameter set. These final results were then placed into tabulated data files and visualised using `gnuplot`, an open source plotting and graphing program.

7.3 Results

The simulation results produced by the new foraging methods were collected and processed, before being graphed and analysed in the same manner as those from the misperception-affected foraging simulations. The effectiveness of the four different foraging methods can therefore be compared directly, by studying three values across the foraging methods and parameter sets. One is the average error probability, given by either the average misperception probability or the average misaction probability, depending upon the population's foraging method. As discussed previously, this value is meaningless for both the reflexive-foraging and perfect-perception agents and is therefore ignored for those foraging methods.

Another comparison point is the relative proportions of the agent sub-populations, given by the agents' misperception or misaction probability. This comparison point is also unsuitable for both the reflexive and perfect-perception foraging agents, since the average error probabilities of these foraging populations are meaningless.

Finally the average potential offspring for the new foraging methods is examined. This computed value describes the amount of offspring an agent could have parented, given optimal conditions in the simulation. The average potential offspring were calculated and compared for all the foraging methods, allowing a comparison between all the foraging methods.

7.3.1 Average Error Probability

The average error probability collectively describes the errors that affect the misaction-affected foraging agents and the misperception-affected foraging agents. Due to the similarity between these two foraging methods, the same analysis methods may be applied to populations affected by either method. The agents' average misperception or misaction probabilities may be collectively described as their average error probability. Since this value is meaningless for reflexive-foraging agents and perfect-perception foraging agents, misaction-affected foraging is the only foraging method analysed here.

As in the case of misperception-affected foraging, if an agent population evolves an average misaction probability that is above the mutation range then this may indicate a benefit to misaction, albeit one that may not be statistically significant. Figure 7.1 shows the average misaction probability of the misaction-affected foraging populations. In this graph the parameter sets where the average misaction probability is above the mutation range are identified with hollow diamonds. The graph also labels the misperception-affected foraging parameter sets previously identified in Figure 6.1, using numbered circles to identify those parameter sets. As in Chapter 6 parameter sets of interest are numerically labelled to simplify their identification and discussion (Table D.1 in Appendix D contains the data values for this plot).

There is an overlap between the evolution of misaction and misperception probabilities in these two groups, as of the nine points where the average misaction probability is above the mutation range, six correspond to points where the average misperception probability was above the mutation range. The overlap of the parameter sets where misaction and misperception probabilities are above the mutation range is likely due to the similar effects of misperception and misaction, as both were implemented as random errors that affected the behaviour of the agents. Four of these overlapping parameter sets occur in the parameter sets with a low Basic Metabolic Rate (0.05) and a high Offspring




Figure 7.1: The Average Misaction Probabilities plotted for each parameter set. There are several parameter sets where the population evolves a misaction above the mutation range and these peaks have been marked with hollow diamonds. The coloured region of the plot demonstrates parameter sets where misaction may be beneficial, although not at a statistically significant level. The numbered dots identify the parameter sets that previously demonstrated significant levels of misperception. The partial overlap between the beneficial parameter sets for misaction and misperception indicates that they provide a similar benefit, likely due to their behavioural similarity.

Cost (500), suggesting that agents benefit more from both misaction and misperception in environments where reproduction is expensive.

There are also three other parameter sets where the average misaction probability was above the mutation range but the average misperception probability was not and these are labelled as 10, 11 and 12. These parameter sets all have relatively high metabolic rates, which require agents to expend more of their resources on survival, leaving less resources available to parent offspring. This result is also somewhat reflected in the misperception-affected foraging simulations, where the two latter parameter sets also demonstrate a minor increase in their populations' average misperception probabilities.

In these nine identified parameter sets the population of agents has evolved an average misaction probability above the mutation range, suggesting an overall advantage to the errors produced by infrequent misaction in those instances, albeit one that is not statistically significant.




Figure 7.2: The Average Misaction Probabilities and 95% Confidence Intervals for each parameter set, compared to the mutation range. Parameter sets 1, 2, 6, 7 and 11 all show a statistically significant misaction probability, indicating that misaction is significantly beneficial in these parameter sets.

Figure 7.2 shows the average misaction probabilities with their 95% confidence interval. Of the numbered parameter sets 1, 2, 6, 7 and 11 all have average misaction probabilities that are significantly different from the mutation range. A one-sample t-test confirms that the difference from the mutation range is statistically significant for these parameter sets. As in the case of misperception-affected foraging, the majority of the parameter sets do not demonstrate a statistically significant benefit of misaction and only parameter sets 1 and 2 both show such a benefit for both foraging behaviours.

Figure 7.3 contrasts the error probabilities obtained by misperception-affected foraging and misaction-affected foraging behaviours. Given the similarity in how these two foraging methods are implemented, it is unsurprising that there is no statistical difference between the evolved error probabilities for many of the parameter sets. The correlation coefficient between the average misperception and misaction probabilities is $r = 0.51411$.

7.3.2 Population Proportions

Since the average error probability is a meaningless measure for populations that use either reflexive-foraging or perfect-perception foraging, dividing the population into groups based upon this value would also be meaningless. Therefore only the population proportions for misaction-affected foraging were measured and compared to those developed by






Figure 7.3: The Average Misperception and Misaction Probabilities and 95% Confidence Intervals for each parameter set.

misperception-affected foraging. As in the previous chapter, the agent population was divided into three sub-populations — agents with a 0% misaction probability, agents whose misaction probability is below the mutation range, and agents whose misaction probability is above the mutation range. Figure 7.4 shows the proportion of the misaction-affected foraging populations that are in each of the three groups, and also compares these results to those of the misperception-affected foraging populations.

In Figure 7.4a, agents with misaction probabilities that are above the mutation range make up substantial proportions of the population for several parameter sets. These parameter sets match those where the average misaction probability was above the mutation range. The six labelled parameter sets (1, 2, 4, 5, 6, 7) previously demonstrated average error probabilities above the mutation range in populations using both foraging methods. Unsurprisingly, these six parameter sets had large sub-populations of agents whose misaction probability was above the mutation range. Also the parameter sets 11, 12 and 13, which previously demonstrated high average misaction probabilities, also displayed large sub-populations of agents with misaction probabilities above the mutation range. These results are similar to those observed for misperception-affected foraging, where average misperception probabilities above the mutation range corresponded to a large proportion of the population having misperception probabilities above the mutation range.



(a) Population proportions of Misaction Probabilities (Misaction-affected foraging)



(b) Comparison of population proportions between Misaction-affected foraging (first) and Misperception-affected foraging (second)

Figure 7.4: Population proportions of misaction-affected agents (Figure 7.4a) and a comparison of the sub-population proportions for misaction-affected foraging and misperception-affected foraging (Figure 7.4b). In both graphs the numbered parameter sets are those labelled due to indications of benefits of misperception or misaction; the parameter sets marked with circles indicate instances where both the average misaction and misperception probabilities were above the mutation range.

Figure 7.4b compares the size of the sub-populations in the misaction-affected foraging populations and misperception-affected foraging populations, with the first column for a parameter set showing the results for misaction and the second those of misperception. It should be noted that the values compared in this graph are percentages of the populations and therefore the actual population sizes likely differ.

In the parameter sets where there is arguably little or no benefit from misperception or misaction, the proportion of agents with misaction probabilities of 0% is higher than the proportion of agents with misperception probabilities of 0%. The high proportions of agents who did not misact suggests that there is a stronger disadvantage to misaction for those parameter sets than there is for misperception.

One possible explanation for the greater proportion of agents unaffected by misaction lies in the differences between how misaction and misperception may affect an agent's behaviour. A misaction error will definitely immediately affect an agent's behaviour and may be corrected on the agent's next turn. However, a misperception error may not immediately affect an agent's behaviour and, indeed, may not even affect an agent's behaviour at all. For example, a location misperception could cause an agent to believe a nearby empty location is actually located far away from its current location, and such a misperception is unlikely to affect an agent's behaviour. Therefore, it follows that a lower probability of misaction should produce an equivalent amount of behavioural diversity as a higher misperception probability.

7.3.3 Average Potential Offspring

The average potential offspring is the main data point for comparing and analysing the effectiveness of the foraging methods. Since this value was previously calculated for each of the three agent sub-populations of misperception-affected foraging agents, the misaction-affected foraging agents were also divided into these sub-populations. For the reflexive-foraging agents and perfect-perception foraging agents the average potential offspring is calculated for the entire agent population. Each foraging method was analysed individually, before the four foraging methods were compared against each other.

Misaction-affected foraging

As with the populations of misperception-affected foraging agents, the misaction-affected foraging agents (Figure 7.5) also show situations where the agents whose misaction probabilities are above the mutation range have significantly more potential offspring than the non-misacting agents within the populations. Of the previously labelled parameter sets, all with the exception of 4 show a significant benefit in the sub-population of misacting

agents, specifically significantly more potential offspring. Many other parameter sets also shared this result and are labelled with #’s.




Figure 7.5: Average potential offspring (misaction-affected foraging) with 95% confidence intervals. The parameter sets with average misaction probabilities above the mutation range show many instances where the misacting sub-populations average significantly more potential offspring than those that do not misact. Parameter sets labelled with numbers or “#” contain misacting agents with significantly higher average potential offspring than non-misacting agents (except for 4).

Of the six parameter sets with significantly more average potential offspring from misperception-affected foraging, all six also demonstrate the same benefit with misaction-affected foraging. However, there are many further parameter sets that also received significantly more average potential offspring when misacting. While this significant benefit was prevalent among the misaction-affected foraging populations, many populations of misperception-affected foraging agents demonstrated no statistically significant difference between the average potential offspring of the three groups of agents. This difference implies that misaction as implemented in this simulation offers a greater benefit than misperception, which was expressed as more effective foraging by the misacting agents that would allow them to potentially parent more offspring.

Reflexive-foraging

The average potential offspring of the reflexive-foraging agents for all the parameter sets is shown in Figure 7.6. The most noticeable effect of this foraging method upon the agent populations is in the higher potential offspring measured across many of the parameter sets.

Indeed, the lowest value of potential offspring received by these agents for some parameter sets matches those that misaction and misperception receive for many parameter sets. Overall the average potential offspring often exhibits a regular pattern across the groups of agent and resource density pairings. While the fact that high and low values for the average potential offspring do correspond to the parental investment costs of the parameter sets is unsurprising, given how the parental investment cost is calculated, this outcome was not clearly observed in the misperception or misaction-affected foraging populations.




Figure 7.6: Average potential offspring (reflexive-foraging) with 95% confidence intervals. The majority of the simulations possessed average potential offspring values between 2 and 6 for the majority of the parameter sets. For this specific foraging environment, reflexive-foraging is surprisingly effective. The regular repeating shows that for each agent and resource density pairing the average potential offspring decreased as the parental investment cost increased. This trend was not clearly demonstrated by the other foraging methods.

Akaishi and Arita hypothesised that the underlying benefit in the foraging scenario arises from the population's increased behavioural diversity. Since the reflexive-foraging agents utilise a random foraging behaviour, they should exhibit much more behavioural diversity than any other foraging method. Therefore, this increased behavioural diversity is presumably responsible for the increased potential offspring of the reflexive-foraging agents. The misperception and misaction-affected foraging methods produce less behavioural diversity than the reflexive-foraging agents, since their behaviour is mainly determined by their knowledge of resource nodes in the environment and only secondarily by their randomly-influenced misperceptions or misactions.

Conventional wisdom suggests that the random methods employed by reflexive-foraging agents should be less effective than those of the misperception-affected and misaction-affected foraging agents. However, reflexive-foraging agents are less likely to be affected by congestion and deadlocking, since their foraging behaviour is not driven by stored beliefs of resource node locations. Instead, reflexive agents are only drawn to resource nodes within their perception range, while they otherwise move in a random walk. This difference means that reflexive-foraging agents are much less likely to encounter other agents waiting at a resource node.

Perfect-perception foraging

The average potential offspring values calculated for populations of agents with perfect-perception foraging is shown in Figure 7.7. Of the 36 parameter sets, 15 had an average potential offspring close to 2.0, while 11 had noticeably higher average potential offspring values that indicated some possible advantage to perfect-perception foraging in those instances. Most of these identified advantageous parameter sets had small to moderate agent densities, which would in turn present less opportunities for those agents to obstruct each others' movements. As such, these perfect-perception foraging agents may have had more potential offspring due to the lower likelihood of congestion in the less-densely populated scenarios.




Figure 7.7: Average potential offspring (perfect-perception foraging) with 95% confidence intervals. While some parameter sets show higher potential offspring than other foraging methods, these tend to occur for parameter sets with lower parental investment costs.

For some parameter sets the perfect-perception foraging agents outperformed the reflexive-foraging populations. However, for a number of parameter sets, the average potential offspring is comparable to that of the misperception-affected and misaction-affected foraging populations, having either slightly more or slightly less potential offspring. When compared to the four foraging methods, perfect-perception foraging has the least potential offspring in a majority of cases. According to Akaishi and Arita's hypothesis, this is due to the lack of behavioural diversity within the populations of perfect-perception foraging agents. Conventional wisdom, however, suggests that perfect-perception foraging should be the fittest foraging method of those examined here.

Comparing the foraging methods

Figure 7.8 directly compares each foraging method's average potential offspring across all the parameter sets, with vertical axes adjusted to share a common scale.

Figure 7.9 overlays and contrasts the average potential offspring data from each foraging method. Note that this directly compares sub-populations of misperception-affected foraging and misaction-affected foraging agents against the complete populations of perfect-perception foraging agents and not their respective sub-populations of agents with 0% misperception or misaction probabilities. Table 7.1 uses this data to tally how frequently each foraging behaviour's average potential offspring was significantly greater than that of another foraging method.

	Misperception	Misaction	Reflexive	Perfect
Misperception	-	0	2	5
Misaction	13	-	4	21
Reflexive	23	9	-	23
Perfect	9	2	2	-

Table 7.1: Number of parameter sets where one foraging method (column) yields significantly more potential offspring than another foraging method (row).

Misperception-affected foraging fails to yield significantly more potential offspring than misaction-affected foraging for any parameter sets. It yields significantly more potential offspring than reflexive-foraging for 2 parameter sets and perfect-perception foraging for 5 parameter sets. All but one of these parameter sets has a low parental investment cost. Three of those parameter sets had the highest possible agent density of 0.3. Under




Figure 7.8: Average potential offspring comparison for the four foraging methods (top to bottom: Misperception, Misaction, Reflexive, Perfect), with the vertical axes adjusted to a common scale. The parameter set labels are all referenced by the common key at the bottom.




Figure 7.9: A direct comparison of the average potential offspring from each foraging method. For misperception and misaction-affected foraging the potential offspring value comes from the sub-populations that were misperceiving or misacting. Only two parameter sets show perfect-perception foraging to yield more potential offspring than the methods that increase behavioural diversity. The majority of the parameter sets show that significantly more potential offspring were obtained by agent populations affected by some source of behavioural diversity.

such conditions, any source of behavioural diversity that encourages exploration, thereby directing agents away from congested areas, would be beneficial.

Misaction-affected foraging yields significantly more potential offspring than misperception-affected foraging in 13 parameter sets and this occurs in parameter sets with varied agent densities, resource densities and parental investment costs. As discussed already in Section 7.3.2, this advantage of misaction over misperception may arise as it will immediately affect an agent’s behaviour and introduce behavioural diversity. The similarities between these two foraging methods is likely due to similarities in their implementation and effect on individual agent behaviour. Of the two methods, misaction more clearly demonstrates a consistent advantage within its populations, with misacting agents often having significantly more potential offspring than non-misacting agents.

Misaction-affected foraging yields significantly more potential offspring than reflexive-foraging in 4 parameter sets, most with low parental investment costs. Conversely, reflexive-foraging yields significantly more potential offspring than misaction-affected foraging in 9 parameter sets. Of the two, reflexive-foraging’s significantly higher potential offspring may be due to the increased behavioural diversity derived from its random behaviour.

Misaction-affected foraging yields significantly more potential offspring than perfect-perception foraging in 21 parameter sets, while perfect-perception foraging only had a significant advantage over misaction-affected foraging in 2 parameter sets. As such, misaction produces a significant benefit when compared against perfect-perception foraging, which supports the hypothesis of behavioural diversity increasing agent fitness in this foraging scenario.

Reflexive-foraging had a significantly higher average potential offspring than than misperception-affected foraging in 23 parameter sets, misaction-affected foraging in 9 parameter sets and perfect-perception foraging in 23 parameter sets. In each case the parameter sets have varied agent densities, resource densities and parental investment costs. Perfect-perception foraging only yields significantly more potential offspring than reflexive-foraging for two parameter sets. Therefore, reflexive-foraging is significantly fitter than perfect-perception foraging. This significant benefit may be due to the increased behavioural diversity reflexive-foraging provides, and its capabilities to minimise congestion in the simulated environments. The significant observed advantages of reflexive-foraging against the guided foraging methods suggests that the simulated foraging scenario does not adequately describe biological foraging scenarios or behaviours, where such random behaviour would likely be much less advantageous, or even maladaptive.

Perfect-perception foraging yielded significantly more potential offspring than all of the foraging methods simultaneously in only 2 parameter sets and yielded more potential offspring than misperception-affected foraging in an additional 7 parameter sets. These results suggest that there is no significant benefit to perfect-perception foraging in the simulated foraging environment.

There are four parameter sets where all the erroneous foraging methods have significantly more potential offspring than perfect-perception foraging, three of which were previously labelled as 1, 4 and 10. These are:

- $[0.3, 0.25, 0.2, 25]$ (-)
- $[0.3, 0.25, 0.1, 50]$ (1)
- $[0.15, 0.1, 0.1, 50]$ (10)
- $[0.3, 0.05, 0.05, 500]$ (4)

Three of these parameter sets have very high agent densities, which ensures that a high number of agents compete for access to resources or potentially obstruct each others' movements, providing more occasions for any source of behavioural diversity to be beneficial.

Parameter set 1 also demonstrates another benefit of misperception and misaction in this foraging environment. Comparing the average potential offspring of the agent sub-populations with 0% misperception probabilities and 0% misaction probabilities to the population of perfect-perception foraging agents shows a statistically significant difference. In this parameter set, the agents with perfect-perception who co-existed with misperceiving or misacting agents significantly outperformed those in a homogeneous population of perfect-perception foraging agents. This demonstrates that individuals or kin do not solely benefit from the behavioural diversity caused by misperception or misaction. Group selection or multi-level selection may explain the selective pressure for misperception or misaction in this instance.

7.4 Conclusion

These simulations show that like the misperception-affected agents, the misaction-affected and reflexive-foraging agents can all benefit from the increase in behavioural diversity introduced by these mechanisms for some parameter sets. This supports the hypothesis of Akaishi and Arita, which attributes the observed benefit of misperception to the behavioural diversity it produces in the agent population. In some circumstances, different mechanisms for introducing behavioural diversity were found to provide an even greater benefit than that previously observed from misperception.

The benefit received by misaction-affected foraging agents is very similar to that of misperception-affected foraging. However, misaction-affected foraging demonstrated many more instances where the misaction-affected foraging agents had significantly more potential offspring than the misperception-affected foraging agents in comparable circumstances. For many parameter sets, misaction-affected foraging agents were significantly fitter, by measure of average potential offspring. Misaction-affected foraging evolves a stable level

of misaction for some parameter sets, which are similar to those in which misperception-affected foraging populations evolve a stable probability. As with misperception-affect foraging, misaction-affected foraging is argued to benefit agents when it occurs infrequently.

Despite the similarities between the implementations of misperception and misaction-affected foraging, the subtle differences in their immediate and long-term effects upon the simulated agents leads to distinctly different results. Therefore, misaction may also be an interesting topic to examine within future research into misperception.

The results for the reflexive-foraging populations demonstrated that random behaviour provides a significantly greater average fitness for its population than misperception-affected foraging for many parameter sets and for misaction-affected foraging populations in some parameter sets. While this is in line with the hypothesised advantage of behavioural diversity, it seems to defy conventional wisdom that random movement should prove better than planned movement. However due to their random movement patterns, the reflexive-foraging agents are less likely to be affected by congestion caused by deadlock, which explains the fitness of these agents.

If reflexive-foraging does benefit mainly from its ability to better avoid deadlocks and congestion in the foraging scenario, then such random movement or behaviour is unlikely to provide a benefit in different scenarios. Further research, however, is required to test this hypothesis. Such research would also fulfil the more general goal of investigating the effects of misperception in other scenarios and environments.

Conventional wisdom suggests that perfect-perception foraging should be the fittest of all foraging methods, yet the results obtained here disprove this for many parameter sets. Both misperception-affected and misaction-affected foraging provide a significant benefit, both to the affected agents themselves and the sub-populations of perfect-perception agents directly competing against them. Reflexive-foraging agents also often had significantly more potential offspring than the perfect-perception foraging agents, supporting Akaiishi and Arita's hypothesis that any source of behavioural diversity will benefit the agents. For the majority of parameter sets studied, perfect-perception foraging produced less potential offspring than a foraging method with greater behavioural diversity.

While increased behavioural diversity is argued to provide an evolutionary benefit in this foraging scenario, there are many other potential situations where behavioural diversity is unlikely to prove beneficial. These would include situations where uniform

behaviour by entities is highly desirable and different behaviour is severely detrimental to the individual and/or population. In such situations the entities of a group expect all members to exhibit similar behaviour, such as in a flock of birds or a squad of soldiers in a combat situation.

Chapter 8

Misperception and the Iterated Prisoner's Dilemma

The underlying model of the Iterated Prisoner's Dilemma occurs in many diverse, real-world interactions. The Tit for Tat strategy provides a maximal continual benefit to both players, relying upon reciprocal cooperation between players (Axelrod, 1984). In an evolutionary environment, the Tit for Tat strategy can quickly become the common strategy due to the advantages cooperating players receive (Axelrod and Dion, 1988). Furthermore, the Tit for Tat strategy's reciprocal cooperation and punishment has been proposed to explain various interactions in both nature and human social systems (Goldstein and Krasner, 1984; Lombardo, 1985; Milinski, 1987).

Noise, whether it is caused by a player's misperception or some other source, disrupts cooperation between players of the Tit for Tat strategy. Tit for Tat players who are affected by noise can end up locked into a retaliatory cycle, where an unintended Defection echoes back and forth between the players. While Axelrod argued that forgiveness and contrition may aid Tit for Tat in a noisy environment, in some instances players may be unable to implement either action. Forgiveness may not be possible when the players do not trust each other, while contrition requires Tit for Tat players to recognise their own errors and then abstain from retaliating against any Defections these errors provoke. In instances where these methods are unavailable or impossible, misperception may be able to restore a state of beneficial cooperation between Tit for Tat players.

8.1 Noise and the Iterated Prisoner's Dilemma

The Iterated Prisoner's Dilemma game has four possible states. When both players use the Tit for Tat strategy, the game effectively has only three states (Figure 8.1) — **Mutual Cooperation**, **Alternating Cooperation and Defection** and **Mutual Defection**. In a noisy environment, noise acts as a random source of error and affects how a player perceives the outcome of the previous round of the iterated game. The game begins in the state of mutual Cooperation, where it will stay until noise causes an error and the game changes state.




Figure 8.1: States of the Iterated Prisoner's Dilemma game as played by Tit for Tat players. The Tit for Tat strategy will keep the game in one of these three states until noise causes a player to make a mistake and changes the game's state.

Noise in the Iterated Prisoner's Dilemma is commonly understood to transpose the Cooperate and Defect strategies and may be implemented as a misperception of an opponent's previous move or a misimplementation of a player's move. In either case, an error will disrupt the mutual Cooperation between players. When both of those players are utilising the Tit for Tat strategy, the error will start a cycle of alternating Cooperation and Defection as the Defection is echoed back and forth. This represents a state change in the game, from mutual Cooperation to Alternating Cooperation and Defection. This state change, and any others, may only be caused by noise.

A drawback of the Tit for Tat strategy is that it performs poorly in noisy environments, as noise prevents it from accurately reciprocating its opponent's previous move. Tit for

Tat players have no direct way of returning to the optimal state of mutual Cooperation, as neither can ignore or forgive the Defection introduced by the noise. Axelrod (1984) has previously examined the effects of forgiveness on players in the Iterated Prisoner's Dilemma, noting that forgiveness can help players to maintain Cooperation in a noisy environment. In some instances misperception may be able to simulate forgiveness, thereby benefiting the Tit for Tat players. Misperception may also be able to implement forgiveness in situations where the players are incapable of forgiveness or will not choose to forgive their opponent. For example, one player may perceive an unwarranted Defection that disrupts the mutual Cooperation between the players as an intentional, hostile act, which prevents the victim player from forgiving its opponent. The hypothesis that misperception can benefit Tit for Tat players by functionally mimicking forgiveness shall be examined. Confirming this hypothesis will demonstrate that causing unintentional forgiveness is one mechanism by which misperception may benefit misperceiving Tit for Tat players.

Misperception in the Iterated Prisoner's Dilemma may either change Defection to Cooperation or Cooperation to Defection. The former may be described as forgiveness and will therefore benefit Tit for Tat players, while the latter may be interpreted as unwarranted punishment and disrupts the Cooperation of Tit for Tat players. The misperception experienced by the Tit for Tat players may be separated into these two opposite forms. If both types of misperception are a heritable trait of the Tit for Tat players, then it is expected that misperception that functions as forgiveness should evolve in a noisy evolutionary environment. This can be tested in an evolutionary variant of the Iterated Prisoner's Dilemma, populated by Tit for Tat players. While forgiveness may benefit the population in a noisy environment, forgiving unwarranted defections may be exploited, thereby increasing the adaptive value of this Punishing Misperception. In such an environment players may evolve an optimal probability of forgiveness, which is sufficient to maintain mutual Cooperation while avoiding excessive forgiveness that is potentially exploitable. If a population in an evolutionary environment can evolve a stable level of Forgiving Misperception, then this supports the hypothesis that misperception can provide an evolutionary benefit by identifying a situation where this occurs.

8.2 Restoring Cooperation with Misperception

Noise affects Tit for Tat players by disrupting their mutual cooperation; changing one player's Cooperation to a Defection, which may then echo back and forth between the players endlessly. Further instances of random noise will cause the game to move randomly between the three possible game states. The Tit for Tat strategy is incapable of recovering to a state of mutual Cooperation in this scenario. In the worst case example, a single instance of noise would affect the first move of the iterated game, locking the players into a state of Alternating Cooperation and Defection for the length of the game. In a more typical instance, noise occurs for the duration of the game at random intervals. In either instance the Tit for Tat players' cooperation is disrupted, leading to a sub-optimal payoff. Misperception is hypothesised to be able to correct the errors caused by noise, restoring the Tit for Tat players to a cooperative state and thereby increasing their payoffs.

This hypothesis will be tested with a simple Iterated Prisoner's Dilemma game played between two Tit for Tat players. The effects of misperception on both a standard noisy environment and in a hypothetical worst case situation will both be examined. Noise in the iterated game is provided by the players' misaction, which causes them to unknowingly perform the wrong move. As a player is unaware of its misaction, it cannot adapt its behaviour to correct the error misaction introduces. Misperception affects the correct observation of an opponent's previous move. In the case of either a misperception or misaction, the move being observed or performed is replaced with its opposite. The two errors are independent, allowing either one or both to simultaneously affect the players. These two errors also affect the extremities of the player's OODA loop, with misperception occurring during the Observation step and misaction occurring during the Action step.

The payoffs for the outcomes of the Prisoner's Dilemma game used in this experiment match those originally used by Axelrod (1984) (Figure 8.2). These values meet Hofstadter's (1986) criteria for a dilemma (see Section 2.3.2), in that the Reward payoff from mutual Cooperation is greater than the average payoff from alternating between the Temptation and Sucker payoffs.

Axelrod (1984) has shown that noise prevents Tit for Tat players from maintaining a state of mutual cooperation and thereby lowers their score. However, this assumes that the players begin the game in a state of mutual Cooperation. If this is not true, then

		Player	
		Cooperate	Defect
Opponent	Cooperate	Reward 3	Temptation 5
	Defect	Sucker 0	Punishment 1

Figure 8.2: The states and payoffs of the Prisoner's Dilemma Game as seen from one player's perspective. The states and payoffs are the same for both players when their positions are reversed. The players are Rewarded for mutual Cooperation and Punished for mutual Defection. Temptation describes the desire to exploit cooperative players for an individual benefit, while Sucker describes the victim of such exploitation.

noise that convinces the players to begin Cooperation may prove beneficial. While there are many reasons that players may be unable to begin cooperating, this model assumes that the initial uncooperative state of the game is due to an external third party, called the Trickster and named from the character archetype. Beginning the iterated game in a non-cooperative state is considered to be a worst-case scenario for the Tit for Tat players. The Trickster has no influence upon the iterated game, except to change one player's initial Cooperation to a Defection, placing the iterated game in an initial state of Alternating Cooperation and Defection. The Trickster is assumed to have some desire to initially force the players from a state of mutual cooperation; however, this desire depends upon the Trickster's identity and the real-world situation that the Iterated Prisoner's Dilemma models.

The hypothesis is tested with a simple Iterated Prisoner's Dilemma game, which is played by two Tit for Tat players. During this game the players may be affected by a combination of three different sources of noise — misperception, misaction and the Trickster. Misperception and misaction will both occur randomly during a game with separate constant probabilities, while the Trickster's initial influence may be enabled or disabled. Misperception will toggle the move a player believes its opponent previously played, while misaction will toggle the move a player has chosen to play. Toggling a move switches it from Cooperation to Defection or from Defection to Cooperation. When the

Trickster affects a game, one Tit for Tat player misimplements its initial Cooperation move, thereby placing the game in a state of Alternating Cooperation and Defection. Interference by the Trickster acts as a worst case scenario for the Tit for Tat players, initialising the game in a non-cooperative state. Games that are not affected by the Trickster allow the players to begin in a state of mutual Cooperation. The iterated game is played for a fixed number of turns. During each game the players observe the previous move of their opponent, use the Tit for Tat strategy to determine their response to the previous iteration, and then act by performing their chosen strategy. Noise may affect the players' Observation and Action steps. Once both players have selected a strategy, their payoffs are awarded based upon the outcome of the game and added to their total scores (Appendix E.2 describes the algorithm for this simulation).

The Iterated Prisoner's Dilemma simulation was written in Java. The default Java random number generator was replaced with the random number generator included in the Colt library (CERN, 2006). The Colt library was selected as it provides several different pseudo-random number generators, which implement uniformly and normally distributed random number generators. The Colt random number generator also has a much longer period, producing a longer sequence of pseudo-random numbers before repeating itself. However, it was not expected that the simulations would be executed for long enough to exceed the limits of the default Java random number generator. Each execution required little time due to the simplicity of the simulation.

If misperception is beneficial for these players, then this benefit will be observed in two main results from the iterated game. The obvious indication of the effectiveness of misperception at restoring mutual Cooperation is the players' total scores. Higher total scores in simulations with misperception than those with no misperception indicate that there is some benefit to misperception. The players' total scores may also be compared to the score that they would have obtained in a noiseless environment where mutual Cooperation was maintained for the entire game. If misperception is implementing beneficial forgiveness, then the scores received by players in a noisy environment should also approach that of mutual Cooperation in a noiseless environment. Another indication of misperception's effectiveness is the amount of time the players spend in each of the three possible game states. If misperception benefits players who are affected by the Trickster, then these

players should spend more time in a state of mutual Cooperation than players who are affected by the Trickster and unaffected by misperception or misaction.

8.2.1 Simulation Parameters

For such a simple simulation there are only a handful of parameters to control. Two fixed parameters were the number of game iterations to be played and the number of times each parameter set was simulated. The Iterated Prisoner's Dilemma games were run for 10,000 iterations, to give the players a long history of interaction in which a sufficient number of misperceptions and misactions could occur. Each parameter set was simulated 100 times in order to generate a sufficiently large sample of results.

The three variable parameters to investigate are the misaction probability, the misperception probability and the Trickster's interference. The misaction probability controls how frequently a player will misimplement its intended move, the misperception probability determines how frequently a player misperceives its opponent's move and the Trickster flag controls whether or not the Trickster affects the first move of the game. The effects of these parameters upon the game is more clearly explained by the algorithm describing this simulation (listed in Appendix E.2). Both 0% and 1% are suitable values of misaction and misperception probabilities to investigate, providing four combinations of parameters to investigate. These combinations model situations where the players may be affected by no sources of noise, one source of noise (either misperception or misaction) or both sources of noise. During a second set of experiments, the 1% noise probability is increased to 5%, which provides another four combinations of parameters to simulate. Testing a secondary noise probability may provide a rough indication of how much noise is required to help players affected by the Trickster.

The final parameter to control was the influence of the Trickster, which may be disabled or enabled. This provides eight parameter sets focusing on the lower misperception and misaction probability (Table 8.1) and another eight focusing on the higher misperception and misaction probability (Table 8.2). This produced 16 different parameter sets, each of which was simulated 100 times, for 1,600 simulations in total.

The simple Iterated Prisoner's Dilemma simulation created requires little computational time or power to execute. However, it was decided to again utilise the resources of the computer cluster and EnFuzion software (Axceleon, Inc, 2003) to aid in the execution

Misaction Probability	Misperception Probability	Trickster
0.0	0.0	Disabled
0.0	0.0	Enabled
0.0	0.01	Disabled
0.0	0.01	Enabled
0.01	0.0	Disabled
0.01	0.0	Enabled
0.01	0.01	Disabled
0.01	0.01	Enabled

Table 8.1: Parameter Sets (Low misaction and misperception probabilities)

Misaction Probability	Misperception Probability	Trickster
0.0	0.0	Disabled
0.0	0.0	Enabled
0.0	0.05	Disabled
0.0	0.05	Enabled
0.05	0.0	Disabled
0.05	0.0	Enabled
0.05	0.05	Disabled
0.05	0.05	Enabled

Table 8.2: Parameter Sets (High misaction and misperception probabilities)

of this parametric simulation. While the simulation does not require the computational power of the computer cluster, it was much easier to use the EnFuzion software to manage the execution of the large number of parametric simulations and the subsequent collection of the results.

8.2.2 Results

Low Misaction and Misperception Probabilities

Figure 8.3 shows the average individual score of the players of the eight parameter sets with low misaction and misperception probabilities, with the Trickster enabled and disabled. This value is the average of each player's individual total score from the iterated game. The error bars shown here cover the 95% confidence interval of the average individual scores, which are averaged across the total individual player scores from the 100 simulated instances of each parameter set. The four combinations of misaction and misperception probabilities are listed across the horizontal axis. The first pair of values show the Tit for Tat players in a noiseless environment. When the Trickster does not affect the Tit for Tat players, they receive the maximum payoff possible from maintained mutual Cooperation, which is a score of 30,000 after 10,000 game iterations. When the Trickster affects these players, they find themselves in a state of Alternating Cooperation and Defection, with an average payoff that is always 25,000. These two values exhibit no variation as with no noise the players are always locked into the game's initial state by their Tit for Tat strategy.

The average scores received in the three cases with either misaction, misperception or both all show similar average scores and confidence intervals. In these cases a player's




Figure 8.3: Average Individual Score received by Tit for Tat players with 1% noise. Note that with either misperception or misaction as a source of noise, players receive a lower score than when the Trickster influences a game in a noiseless environment. There is also little difference between the scores received with different sources of noise.

average payoff is approximately 22,500, while the majority of the scores fall between 21,000 and 24,000. The average score obtained when affected by misperception or misaction is less than that when there is no noise and the Trickster forces the game into an uncooperative state. There is also no difference between the cases with and without the Trickster's initial influence. The effects of many misactions and misperceptions during the game seem to dwarf the Trickster's initial effect on the game. These results do not appear to support the hypothesis that misperception could restore cooperation between Tit for Tat players affected by the Trickster.

Comparing the duration that the players spent in each of the game's three possible states reveals why there was no benefit from misperception (Figure 8.4). In the two parameter sets with no noise (0% misaction and 0% misperception) the players spend the entire game locked into their initial state, which is mutual Cooperation without the Trickster and alternating Cooperation and Defection with the Trickster. The remaining six parameter sets show that the players spent approximately 2,500 turns in states of Mutual Cooperation and Mutual Defection and approximately 5,000 turns in the two alternating states. Since there are two alternating states, the players are spending an approximately equal time in all of the four possible game states when they are affected by

misperception or misaction. There is no indication that misaction or misperception help restore cooperation in these instances.




Figure 8.4: The duration spent in the three states of the Iterated Prisoner’s Dilemma by Tit for Tat players with 1% noise probabilities. Note that with either misperception or misaction as a source of noise, players spend an equal number of turns in each of the states (The values for the two states of Alternating Defection and Cooperation are combined).

These results demonstrate that the introduction of noise through either misperception or misaction does not restore Cooperation to Tit for Tat players initially affected by a Trickster, but instead disrupts it. This produces an outcome where the players spend equal amounts of time in each game state, which has a worse average score than the Trickster-affected game with no noise. In this situation misperception is clearly detrimental to the Tit for Tat players.

High Misaction and Misperception Probability

The second group of parameters increased the misaction and misperception probabilities from 1% to 5% and the results from those simulations are shown in Figure 8.5. These results are very similar to those obtained with the lower misaction and misperception probabilities (Figure 8.3). In the six parameter sets where the Tit for Tat players are affected by either misaction or misperception, the average score is once again approximately 22,500. As previously, the average scores for the parameter sets with any noise are less than that of the Alternating Cooperation and Defection caused by the Trickster, showing that misperception or misaction are not benefiting players affected by the Trickster. Once

again the Trickster's initial effect on the game is indiscernible in simulations affected by misaction or misperception.




Figure 8.5: The Average Score received by Tit for Tat players with 5% noise. Note that as with 1% probabilities for either misperception or misaction, players receive a lower score than when the Trickster influences a game in a noiseless environment.

The average number of turns the players spent in each of the four game states are approximately the same as those obtained when the noise probabilities were only 1% (Figure 8.6). When noise caused by misaction or misperception affects the Tit for Tat players, they spend an even amount of turns in each of the four states. Enabling or disabling the Trickster in the simulations affected by noise has no noticeable effect on the players' scores. If misperception was restoring Cooperation between these players, then they should spend more time mutually Cooperating than they do mutually Defecting.

The only notable difference between the simulations with 1% and 5% noise probability (Figure 8.3 and Figure 8.5 respectively) is that the average scores from the 5% noise probability has a smaller confidence interval. This may be due to the more frequent occurrences of misaction and misperception producing more transitions by players between game states, which brings the duration players spend in each state of the individual games closer to the uniform average. This reduces the range of the scores received.

As in the simulations with the lower noise probability, misperception provides a worse score for Tit for Tat players than the alternating state that the Trickster's interference causes. This disproves the earlier hypothesis that it could benefit players affected by the Trickster.




Figure 8.6: The duration spent in the three states of the Iterated Prisoner’s Dilemma by Tit for Tat players with 5% noise. As in the case with 1% misperception or misaction probabilities, players spend an equal number of turns in each of the states, indicating random transitions between game states (The values for the two states of Alternating Defection and Cooperation are combined).

8.2.3 Discussion and Conclusions

The results collected do not show noise from either misaction or misperception providing the hypothesised benefit to players forced by the Trickster into a sub-optimal uncooperative initial state. While noise may help to restore cooperation, it is just as likely to disrupt it or force the players to a state of mutual Defection. Furthermore, even if misperception does restore the players to a cooperative state, any future noise will again disrupt the re-established mutual Cooperation. While the Tit for Tat players are capable of maintaining the state of the game by reciprocating actions, noise causes the players to change state. Once there is sufficient noise, which here is shown to be a 1% probability of either misaction or misperception, the players begin to uniformly randomly move between states. As such the likelihood of each state being the outcome for that turn is equal and therefore the game spends an equal amount of time in each of the possible states, which yields a score approaching that received by players uniformly randomly selecting their strategy.

These results confirm Molander’s (1985) statement that in a sufficiently noisy environment the Tit for Tat strategy performs similarly to the random selection of strategies by players. In the case of players selecting strategies at random, the players will change the game’s state almost every turn. Noise-affected Tit for Tat players however only change

states whenever noise occurs. Over a sufficiently long game, both Random players and noise-affected Tit for Tat players will spend approximately the same duration in all four states. The only difference between the two strategies is that the Random players will make many more transitions between states than the noise-affected Tit for Tat players. During each game there is a 75% chance that the Random players will select moves that alter the state of the game. In a noise-affected game with a single noise source that has a probability of 5%, there is a 9.75% chance that noise will affect at least one player and induce a state change in the game (for further details see Appendix E.1).

The average payoff received by players randomly selecting their strategy is 2.25 per turn (Equation 8.1). Over a game with 10,000 iterations the Random player will receive a score of approximately 22,500, which matches the average score previously recorded for noise-affected Tit for Tat players.

$$\begin{aligned}
 \frac{1}{4} \times T + \frac{1}{4} \times R + \frac{1}{4} \times P + \frac{1}{4} \times S &= \frac{1}{4} \times 5 + \frac{1}{4} \times 3 + \frac{1}{4} \times 1 + \frac{1}{4} \times 0 & (8.1) \\
 &= \frac{5}{4} + \frac{3}{4} + \frac{1}{4} \\
 &= 2.25
 \end{aligned}$$

This is less than the 2.5 per turn that the players would receive in an alternating state with no noise, where the players cycle between scores of 5 and 0 each game. Hence, the introduction of noise as simulated here to counteract the Trickster provides a worse payoff than Alternating Cooperation and Defection. The only way by which noise could appear to provide a benefit is if the Trickster forces the initial state of the game to be mutual Defection. In such a case the average payoff per turn for the Trickster-affected players without noise would be 1.0, which happens to be less than the score obtained from the noise-affected Tit for Tat strategy. However, any such benefit would not be due to the restoration of mutual Cooperation, but instead due to the series of random game state transitions providing a higher score than mutual Defection.

If misperception is to restore Cooperation between Tit for Tat players, then it needs to be limited in some manner; allowing Defection to be misperceived as Cooperation, but preventing the misperception of Cooperation as Defection. Symmetric noise, which affects both Cooperation and Defection evenly, does not benefit Tit for Tat players, regardless of

the presence of any Trickster. Restricting misperception so that it is only possible to misperceive Defection as Cooperation is more likely to be beneficial, as it prevents the players from suffering from misperceptions that move the game to a worse state. The simulation was altered so that the players could only misperceive Defection as Cooperation and the hypothesis tested by performing further simulations. Misaction was retained as a source of noise, while the Trickster was removed due to its ineffectiveness. In this simulation misaction produces noise, while asymmetric misperception may forgive Defections, thereby restoring mutual Cooperation.

This simulation was tested with noise probabilities of 1% and 5% and compared to the previous results (Figure 8.7). With noise probabilities of 1% and 5%, asymmetric misperception gives the Tit for Tat players an average score that is slightly greater than 25,000. This is greater than the average score received for players stuck in an alternating state (exactly 25,000) and that received by players in a noisy environment (approximately 22,500). Here asymmetric misperception allows the Tit for Tat players to resume Cooperation after misaction disrupts it, benefiting the Tit for Tat players. This scenario does demonstrate that misperception can beneficially restore mutual Cooperation between Tit for Tat players, albeit with asymmetric misperception that only acts as forgiveness.




Figure 8.7: The average scores from asymmetric misperception in a noisy environment compared to the previous results with symmetric misperception. In both cases where misperception only functions as forgiveness, the players receive a greater score than that received from a Trickster-affected game in a noiseless environment or in a noisy environment with symmetric misperception.

Forgiveness has previously been demonstrated to benefit Tit for Tat players in a noisy environment (Axelrod and Dion, 1988). Asymmetric misperception that acts as forgiveness provides a benefit similar to intentional forgiveness, although it may occur in situations where intentional forgiveness may not. Conversely it is expected that asymmetric misperception of Cooperation as Defection would harm Tit for Tat players, regardless of whether or not there are other noise sources in the environment. Dividing misperception into two separate types here provides a new model of misperception in the Iterated Prisoner's Dilemma, which will now be investigated.

8.3 Evolutionary Value of Asymmetric Misperception

Misperception in the Iterated Prisoner's Dilemma can be divided into two separate forms, which can be made to occur with separate probabilities. Misperception that changes a Defection to a Cooperation is called a **Forgiving Misperception** and allows the two Tit for Tat players to move closer to their optimal state of mutual Cooperation. The misperception of Cooperation as Defection is called **Punishing Misperception**, as the opposing player is believed to have performed an unwarranted Defection that must be reciprocated. Generosity is a possible descriptive name for the trait that causes Forgiving Misperception, while paranoia may aptly describe the trait embodied by Punishing Misperception.

These two separate misperception probabilities will effectively alter the strategy of affected Tit for Tat players in different ways. A Forgiving Misperception probability of 1.0 will cause every Defection to be misperceived as Cooperation, which causes the player to act as though its strategy is Always Cooperate. A Punishing Misperception probability of 1.0 will cause every Cooperation to be misperceived as Defection, which causes a player to act as if its strategy is Always Defect. A misperception probability of 1.0 for both produces a behaviour that functions as the opposite of Tit for Tat, Cooperating against Defection and Defecting against Cooperation.

Separating misperception in this manner is hypothesised to allow a population of Tit for Tat players in a noisy environment to evolve optimal probabilities of forgiveness and punishment. Symmetric misperception did not benefit Tit for Tat players in a noisy environment, as it performed similarly to players selecting strategies at random. Splitting

misperception into two types allows for the differentiation between Forgiving Misperception that can beneficially restore mutual Cooperation and detrimental misperception that disrupts mutual Cooperation. Forgiveness is, however, potentially exploitable by players who utilise Punishing Misperception, as their extra unwarranted Defections provide an individual benefit at the cost of the forgiving player.

In an evolutionary situation, Forgiving Misperception could evolve due to the individual and collective benefit it provides in noisy environments. Punishing misperception may also evolve to exploit the forgiving players. Since Forgiving Misperception will restore mutual cooperation and Punishing Misperception will disrupt mutual cooperation, it would seem that Forgiving Misperception would prove beneficial and spread throughout the population, while Punishing Misperception would die out. However, if forgiveness is widespread in a population of Tit for Tat players, punishment becomes advantageous and may invade the player population. In a population of competing Tit for Tat players, the relative benefit of these two types of misperception will depend upon the frequency of the other type in the population. Punishing Misperception will only benefit players if there is sufficient Forgiving Misperception for it to exploit.

It is hypothesised that a population of Tit for Tat players in a noisy evolutionary environment could evolve a stable level of Forgiving Misperception, indicating a benefit to misperception. This benefit would evolve despite the detrimental impact that Punishing Misperception could have on the individual players. Here the Tit for Tat players would evolve a probability of Forgiving Misperception suitable for their environment, while also evolving a Punishing Misperception probability very close to 0%. This hypothesis is tested through simulation, with an evolutionary Iterated Prisoner's Dilemma game between misperceiving players using Tit for Tat strategy.

8.3.1 Method and Parameters

This hypothesis will be tested by simulating a population of Tit for Tat players that compete in an evolutionary Iterated Prisoner's Dilemma tournament. In such a tournament a population of players compete in iterated games against all other members of their simulated population. Between tournaments the population is adjusted through reproduction between members of the population. Punishing and Forgiving Misperception probabilities

will be a heritable trait of the players, allowing the population to evolve optimal probabilities for both values over many generations. If the population of Tit for Tat players can evolve a stable level of Forgiving Misperception in a noisy environment, then this indicates a potential benefit of misperception. Such a benefit would be identified by measuring the scores received by players and the Forgiving Misperception probability they evolve. This score should be greater than that received from random strategy selection by the players and close to that of mutually cooperating players (Appendix E.3 further details the operation of the evolutionary tournament simulation).

Each player has a chromosome storing its Forgiving and Punishing Misperception probabilities, which are values between 0.0 and 1.0. Punishing Misperception may occur whenever a player observes its opponent Cooperating in the previous game iteration, while Forgiving Misperception may occur whenever a player observes its opponent Defecting in the previous game iteration. Players use the Tit for Tat strategy to determine which move they will select, based upon the observed previous move of their opponent, which may be misperceived. After both players have selected their strategy during a game, the payoff from the resulting outcome is added to their total score. A player's score is considered the only measure of its fitness in this environment.

Environmental noise is again modelled by the players' misactions, with each player in the simulation having the same probability of misacting. Misaction is symmetric and will cause an affected player to implement the opposite of its intended move. This misaction cannot be detected or prevented by the players and is intended to model the regular unavoidable mistakes that can affect an organism. However, the Trickster has been removed as a source of noise, since it was previously demonstrated that regular random noise has a greater effect than the Trickster's single initial error.

The population is initialised with players whose misperception probabilities are randomly generated. These values are taken from a normal distribution with a mean of 5% and a standard deviation that was set as a simulation parameter (later fixed at 2%). During this evolutionary Iterated Prisoner's Dilemma tournament, each player competes in an iterated game of a fixed duration against all the other players in the population. All players are then ranked by their total score from all the games they played. A player's score is the main measure of its fitness, with the highest scoring player assumed to be the fittest and the lowest scoring player the least fit. A player's misperception probabilities

will affect its behaviour and determine how it responds to noise during its competitions against other players.

Reproduction occurs after each generational tournament has concluded, between the fittest player and a player selected at random from the population. The least fit player is removed from the simulation and replaced by the new offspring. The next generational tournament between the updated population then begins. Reproduction utilises a crossover operation of the parents' chromosomes to produce the new offspring's chromosome, which is then mutated. Mutation applies a small randomly generated mutation value to each value in the offspring's chromosome, which alters the Punishing and Forgiving Misperception probabilities of the new offspring slightly. Mutations are prevented from increasing or decreasing a misperception probability beyond its minimum or maximum value of 0.0 and 1.0 respectively.

There are several adjustable parameters that control important aspects of each simulation run and where possible these were assigned a suitable constant value.

Misaction Probability: The global probability that a player will misimplement a move during any game, corresponding to noise affecting the game. A misaction probability of 1% was selected, as it allows a direct comparison with the previous simulations.

Population Size: The fixed number of players competing in the Iterated Prisoner's Dilemma tournament each generation. A population of size n requires $(n^2 - n)$ contests between players, making the population size a major influence on the simulation's execution time. The population size for the simulations was set to 25, which means that the population competes in 600 Iterated Prisoner's Dilemma games each generation.

Tournament Length: The number of rounds players compete for in the Iterated Prisoner's Dilemma game. This parameter was set to 200 rounds.

Generations: The number of generational tournaments during the simulation's execution. This parameter was set to 10,000.

Mutation Standard Deviation: The standard deviation of the normally distributed random number generator that generates mutations during reproduction. A value of 1% was selected for this parameter.

Since this simulation is focused on identifying any optimal Forgiving Misperception, it will record the Forgiving and Punishing Misperception probabilities of the player population and their total scores. This data may then be averaged across individual populations or multiple runs of the simulation. If the population evolves a stable level of Forgiving Misperception, this will be observable in a noticeable average Forgiving Misperception probability that exists over many generations. Such a level of Forgiving Misperception should also ensure that the players' scores are greater than those received from random strategy selection and close to those received for continual mutual Cooperation. Such results will support the hypothesis that Forgiving Misperception can provide an evolutionary benefit to the Tit for Tat players in a noisy environment.

8.3.2 Results

The simulation was executed with the selected parameters 30 times to produce a sufficient statistical sample. The scores received by a player during each generational competition can be contrasted with those that would be received from various pure hypothetical outcomes of the games (Table 8.3). These outcomes assume that the two players compete in a tournament match where the game stays in the same state for the entire game, which is impossible in the noisy environment the simulation studies. However, these values provide several benchmarks against which the average scores obtained by populations may be compared to determine how effective they are at maintaining mutual Cooperation. These scores are calculated by taking the average per-turn payoff a single player receives for those outcomes and multiplying them by the number of Prisoner's Dilemma games that a player competes in during a generation, which in this case is $24 \times 200 = 4,800$.

Game State	Score
Mutual Cooperation	$24 \times 200 \times 3 = 14,400$
Alternating Cooperation	$24 \times 200 \times 2.5 = 12,000$
Random Strategy Selection	$24 \times 200 \times 2.25 = 10,800$
Mutual Defection	$24 \times 200 \times 1 = 4,800$

Table 8.3: Iterated Prisoner's Dilemma scores for various game states

Figure 8.8 shows the average individual score plotted against the left vertical axis and average misperception probabilities plotted against the right vertical axis for the 30 simulation iterations. The area under the average score is shaded in these plots to differentiate between the score and the misperception probabilities. The average individual

score is calculated from the total scores of each population of players and averaging this value across the 30 simulations (Appendix E contains plots of the average scores over time within the population of each simulation).




Figure 8.8: The average player score and average evolved misperception probabilities from 30 simulations plotted over time. Forgiving Misperception quickly evolves to maintain Cooperation. Punishing Misperception later evolves to exploit the higher levels of Forgiving Misperception in the populations, which reduces the average population score.

The populations of the simulation begin with initial Forgiving and Punishing Misperception probabilities that are approximately 5%. From this point the population evolves a high Forgiving Misperception probability and a near-zero Punishing Misperception probability, which increases the population's total score close to that received from Mutual Cooperation. This demonstrates the hypothesised benefit Forgiving Misperception can provide from returning Tit for Tat players to a cooperative state in a noisy environment. After approximately 1,000 turns, the populations started to evolve higher Punishing Misperception probabilities. The average Punishing Misperception probability continued to rise, decreasing both the average player scores and the Forgiving Misperception probability. After 10,000 generations, the average population's score has been reduced to less than if the players were randomly selecting strategies each turn. This is due to the high

probability of Punishing Misperception, which evolved in the population to take advantage of the Forgiving Misperception.

Of the 30 runs of the simulation, only five had stable populations of players who were mostly mutually cooperating with substantial levels of Forgiving Misperception and low Punishing Misperception probabilities (Appendix E.4 contains plots of each iteration's individual results). Of the 30 runs simulated, 25 populations show instances where Punishing Misperception has evolved to exploit the Forgiving Misperception, reducing the population's score. These populations appear to evolve higher probabilities of Punishing Misperception, which reduces their scores below the score obtained in a noisy environment with no misperception. Therefore, in the majority of cases, misperception is not providing an overall benefit for the players. In some cases, such high Punishing Misperception probabilities are evolved that the average score is much less than the score non-misperceiving players would receive in a noisy environment. Therefore, misperception cannot be argued to provide a consistent benefit for these players in a noisy environment.

Run 11 (Figure 8.9) of the simulations showed a population that managed to maintain a high Forgiving Misperception probability, despite the evolution of Punishing Misperception in several cycles. In this iteration the average score that the players receive is often close to the theoretical collective maximum obtainable from continued Mutual Cooperation and is also typically greater than the average score obtainable from random strategy selection. The relationship between the evolved Punishing Misperception probability and the population's total score is clearly seen in this graph — peaks in the average Punishing Misperception probability correlate to falls in the population's total score. While this cyclic behaviour of Punishing Misperception suggests that the exploitation by Punishing Misperception will not always drive Forgiving Misperception from the population, the average data from all the simulation iterations indicates otherwise. Different behaviour can be seen in Run 5 (Figure 8.10), where the population evolved a high Punishing Misperception probability to exploit the Forgiving Misperception in the population. Unlike Run 11, the population of Run 5 could not maintain any Forgiving Misperception probability, which ultimately led to a stable state where the players received scores similar to those of Mutual Defection. In this instance Forgiving Misperception was exploited and unable to benefit the Tit for Tat players, producing a situation worse than if the players were randomly selecting strategies.




Figure 8.9: Average score and average evolved misperception probabilities for Run 11 plotted over time. In this iteration there are several cycles where Punishing Misperception evolves and then dies out in the population. Forgiving Misperception is beneficial in this iteration as the population typically maintains an average score indicative of mutual Cooperation.




Figure 8.10: Average score and average evolved misperception probabilities for Run 5 plotted over time. In this iteration runaway Punishing Misperception evolves and Forgiving Misperception nearly disappears. The extremely high Punishing Misperception ensures that the players' behaviour becomes similar to the Always Defect strategy, with its resulting poor payoffs.

Based on the trends observed in the simulated populations, it is expected that increasing the number of generations studied would have produced more cases where the population evolved to a stable state with an extremely high Punishing Misperception probability. Forgiving Misperception does provide a benefit in this evolutionary simulation; however, once it reaches a sufficient threshold it is likely to be exploited by Punishing Misperception. In the simulation developed here this threshold is an average Forgiving Misperception probability of approximately 30%. This exploitation easily outweighs any benefit from Forgiving Misperception and leads to a state where the players perform worse than if they were competing in a noisy environment with no misperception. As

the Punishing Misperception probability increases, the selection pressure against forgiveness also increases, causing it to decline in the population. The invasion of populations with Forgiving Misperception by Punishing Misperception demonstrates that high Forgiving Misperception probabilities are not an Evolutionarily Stable Strategy. Furthermore the inability of most forgiving populations to avoid invasion by Punishing Misperception implies that Punishing Misperception is an Evolutionarily Stable Strategy.

The evolution of the high probabilities of Punishing Misperception is an example of a ‘tragedy of the commons’ situation (Hardin, 1968), wherein players act in their own best interests to the eventual detriment of the entire population when such behaviour is universal. While in some cases Forgiving Misperception does benefit the Tit for Tat players in their noisy environment, this benefit does not universally evolve. In some instances, runaway Punishing Misperception may evolve, demonstrating that Forgiving Misperception can be invaded by Punishing Misperception and is therefore not an Evolutionarily Stable Strategy. Hence, it cannot be argued that Forgiving Misperception is universally beneficial in the evolutionary Iterated Prisoner’s Dilemma scenario examined. For Forgiving Misperception to benefit these Tit for Tat players, some mechanism to prevent the evolution of excessive levels of Punishing Misperception is required.

8.4 Preventing Exploitation from Punishing Misperception

High levels of Forgiving Misperception are not Evolutionarily Stable, as they provide an excellent environment for the emergence of Punishing Misperception. This most often leads to a runaway process, where the population’s Forgiving Misperception probability dramatically declines to prevent its exploitation by Punishing Misperception. Simultaneously, the Punishing Misperception probabilities of the population increase to exploit any forgiving players in the population. This leads to the Evolutionarily Stable state where the population has evolved a very low Forgiving Misperception probability and a very high Punishing Misperception probability. This effectively changes the strategy of the players from Tit for Tat to Always Defect. This state is highly sub-optimal, as players receive an average individual payoff of 1.0 per turn, much lower than either Alternating Cooperation and Defection (2.5 per turn) or random behaviour (2.25 per turn). If there are too many

players with high Punishing Misperception probabilities, Forgiving Misperception cannot prevail.

A possible solution to this problem is to limit the maximum Forgiving Misperception probability that players may evolve. If the population is prevented from evolving a high Forgiving Misperception probability then there will be less opportunities for Punishing Misperception to exploit. This produces less selective pressure for Punishing Misperception, thereby preventing it from invading the forgiving player population. It is hypothesised that an upper bound on Forgiving Misperception may prevent the development of a population that is easily invaded by Punishing Misperception, while still allowing the population to benefit from Forgiving Misperception. An optimal value for Forgiving Misperception is one that maximises the benefit from restoring and maintaining Cooperation, while also not permitting excessive exploitation from Punishing Misperception. Varying the upper bound of Forgiving Misperception and measuring the evolved Forgiving and Punishing Misperception probabilities of the population should identify the optimal upper bound for Forgiving Misperception. Up to this limit the player populations should benefit from Forgiving Misperception, helping maintain mutual Cooperation. Beyond this upper bound, the population may develop Forgiving Misperception probabilities that invite the evolution of high Punishing Misperception probabilities and reduce the scores of the players.

The alternative to restricting a population's Forgiving Misperception probabilities is to restrict the Punishing Misperception probabilities that it may evolve. If Punishing Misperception is restricted, whenever there is sufficient Forgiving Misperception to exploit the population will likely evolve the maximum permitted level of Punishing Misperception. Restricting Punishing Misperception treats the symptoms but not the underlying cause, which is the benefit from exploiting players with high Forgiving Misperception probabilities. Limiting the population's Forgiving Misperception probabilities will instead solve the cause of this problem and it should be possible to do this while the population still benefits from the misperception.

8.4.1 Method and Parameters

Forgiveness can aid Tit for Tat players in a noisy environment when Forgiving Misperception is limited in the population, thereby preventing the evolution of Punishing Misperception to exploit the population's forgiveness. Forgiveness will be restricted by adding an upper bound to the player's Forgiving Misperception probability, which will become a variable parameter of the simulation. During the execution of the simulation no players will be permitted to develop a Forgiving Misperception probability greater than this upper bound (Appendix E.5 details how the mutation that affects the Tit for Tat players does not allow Forgiving Misperception probabilities above the upper bound to evolve within the population).

The simulation parameters were reused from the previous section, with the exception of the misaction probability. Seven different values for the misaction probability were investigated — 0.0, 0.005, 0.01, 0.02, 0.03, 0.05 and 0.1. These values are low noise probabilities, as in an excessively noisy environment the population of Tit for Tat players will randomly move between game states and any effects from misperception, beneficial or detrimental, will be concealed. For the Forgiving Misperception probability limit, 24 values were investigated — 0.05, 0.1, 0.2, 0.25, 0.26, 0.27, 0.28, 0.29, 0.3, 0.31, 0.32, 0.33, 0.34, 0.35, 0.36, 0.37, 0.38, 0.39, 0.4, 0.45, 0.5, 1.0. There is an increased focus on Forgiving Misperception limits between 0.25 and 0.4 since previous simulations identified that high probabilities of Punishing Misperception did not evolve until the Forgiving Misperception probabilities reached this range. A finer examination of the Forgiving Misperception probabilities within this range could thus reveal the exact point at which Forgiving Misperception switches from beneficial to detrimental. These two parameters produce 168 different parameter sets, which will each be simulated 30 times. This gives a total of 5,040 simulations. The EnFuzion software and computer cluster will again be used to aid in the execution of this parametric simulation and to reduce its total execution time.

Since this simulation collected the same data and utilised the same game parameters as the previous version, the average scores for the Tit for Tat populations can be compared to the scores listed in Table 8.3. Comparing the scores to those values will indicate whether the population is attempting to maintain a state of mutual Cooperation or not. Once again, a score that is close to that of mutual Cooperation indicates that the population is benefiting from Forgiving Misperception. Furthermore, studying the change in the

average player scores as the Forgiving Misperception probability upper bound is increased will indicate at what point Punishing Misperception evolves in the population. This will identify the optimal upper bound value that allows the players to benefit from forgiveness, while avoiding the evolution of punishment in their population.

8.4.2 Results

The effect of adding an upper bound to the player's Forgiving Misperception probability can be seen by examining the average Forgiving Misperception probability as its upper bound is increased. This value is the average Forgiving Misperception probability that evolved in the final generation of each simulated population. Previous results indicated that with a Forgiving Misperception probability of approximately 30%, Punishing Misperception became highly beneficial and evolved in the player population to exploit forgiving players.

Figure 8.11 shows the average Forgiving Misperception probability of the final simulated generation plotted against the Forgiving Misperception probability upper bound and misaction probabilities. The gradual slope shows that the average Forgiving Misperception probability is limited by the upper bound until it approaches approximately 30%, the transition point previously observed. Beyond this point, the average Forgiving Misperception probability declines, as Punishing Misperception evolves to exploit it. The average Forgiving Misperception probability does not rise as quickly when the misaction probability is 0% because in such a case there are no errors for forgiveness to correct and therefore no incentive for players to evolve forgiveness.

The sharp decline observed in the average Forgiving Misperception probability correlates with a sharp increase in the average Punishing Misperception, as shown in Figure 8.12. This increase clearly demonstrates that a sufficiently high probability of Forgiving Misperception in the player population increases the evolutionary benefit of Punishing Misperception. This rise is much smaller when the misaction probability is 0%; however, in such cases there is also a lower average Forgiving Misperception probability. Since the average Forgiving Misperception probability is lower, the average Punishing Misperception probability is lower, as Punishment is more beneficial when its unprovoked Defections are forgiven.




Figure 8.11: The average Forgiving Misperception probabilities of the final generations, plotted for the various misaction probabilities and Forgiving Misperception Probability Upper Bounds. The sharp decrease in the Forgiving Misperception probability near the Forgiving Misperception probability upper bound of 30% is due to exploitation by Punishing Misperception. The misaction probability of 0% produces less selective pressure for Forgiving Misperception, thereby lowering the population’s average Forgiving Misperception probability in those simulations (Appendix E.6.1 shows other plots of this data).




Figure 8.12: The average Punishing Misperception Probabilities, plotted for the various misaction probabilities and Forgiving Misperception Probability Upper Bounds. The Punishing Misperception probability is very low until the Forgiving Misperception probability upper bound reaches approximately 30%, where it steeply increases. This increase matches the sharp decrease in the Forgiving Misperception probability, since Punishing Misperception is exploiting Forgiving Misperception. When the misaction probability is 0%, there is less evolutionary pressure for the development of Forgiving Misperception, ultimately resulting in a much lower average Punishing Misperception probability (Appendix E.6.2 shows other plots of this data).




Figure 8.13: The average final individual score of the Tit for Tat populations, plotted for the various misaction probabilities and Forgiving Misperception Probability Upper Bounds. The decrease in the population’s score occurs at the same Forgiving Misperception Probability Upper Bound as the increase in Punishing Misperception and the decrease in Forgiving Misperception, showing the effects of Punishing Misperception’s exploitation on the population’s average score. The simulations with a misaction probability of 0% have a higher score since those instances had much lower average Punishing Misperception probabilities than the simulations with higher misaction probabilities (Appendix E.6.3 shows other plots of this data).

The total score of the player population also demonstrates the effects of the upper bound applied to the players' Forgiving Misperception probabilities. The average scores from the final generations of the Tit for Tat players are shown in Figure 8.13. Higher scores indicate that Forgiving Misperception is maintaining mutual Cooperation between the players despite the noise in the environment, while lower scores indicate a lack of cooperation. Here the score received by the player populations reaches a peak at approximately 30%, before dropping rapidly. Beyond this threshold Punishing Misperception increases and the population's average score decreases accordingly. When the misaction probability is 0%, the average score also decreases somewhat before increasing again. In these cases the average scores are higher as there is no noise from misaction to affect the players.

There is a regular curve before the sharp decline in player scores caused by the increase in Punishing Misperception. This curve shows that when there is some probability of noise from misaction, the average score will increase along with the Forgiving Misperception probability upper bound. However, this trend ceases once the optimal forgiveness threshold is exceeded, at which point Punishing Misperception evolves and the population's score drops dramatically. When the Forgiving Misperception cap is 0.0, preventing Forgiving Misperception, the average score is reduced due to the effects of noise that cannot be corrected. With the highest misaction probability of 0.1 the average score is reduced to approximately 10,800, which is the score obtained from random behaviour. As small amounts of Forgiving Misperception are permitted, the score gradually rises up to the threshold near the 30% Forgiving Misperception probability. Beyond this threshold the score drops below that of random behaviour, due to the evolution of exploitative Punishing Misperception.

When there is no noise from misaction the population typically has a higher average score. This increase aligns with lower Punishing and Forgiving Misperception probabilities when there is no misaction. With no noise-induced errors to correct, there is less selective pressure upon Forgiving Misperception and this reduction in Forgiving Misperception leads to a reduction in Punishing Misperception.

Figure 8.14 shows the Forgiving Misperception upper bound values that produced the highest average scores and the threshold at which the average scores rapidly decrease due to the evolution of exploitative Punishing Misperception. The threshold before the average player score decreases may be considered to be close to the optimal value for the Forgiving




Figure 8.14: Comparison of the highest Forgiving Misperception probability upper bound where the player population maintains Cooperation (beneficial threshold) and the Forgiving Misperception probability upper bound at which the highest average score was obtained. The optimal Forgiving Misperception probability upper bound that maximises the population’s score is typically less than the highest value for which Cooperation is maintained, suggesting that exploitative Punishing Misperception has already begun to evolve in the latter case and is reducing the average score.

Misperception upper bound in that situation and can be called the **beneficial threshold**. The beneficial threshold is the highest Forgiving Misperception probability upper bound that permits beneficial forgiveness without allowing excessive Punishing Misperception to invade the population.

The earlier simulations showed that with Forgiving Misperception probabilities of approximately 30% or higher, Punishing Misperception evolves in the player population to exploit the forgiveness. The beneficial threshold is close to the point where Punishing Misperception was earlier observed to begin to evolve and degrade the population’s performance. The optimal population score was typically obtained with a Forgiving Misperception probability upper bound less than the beneficial threshold and this suggests that Punishing Misperception has begun to emerge at this point. The optimal scores are very close to those obtained from mutual Cooperation, indicating that these players are benefiting from the forgiveness caused by misperception.

Adding an upper bound to the Forgiving Misperception probability has demonstrated that Forgiving Misperception can provide a benefit to Tit for Tat players competing in the Iterated Prisoner’s Dilemma in a noisy environment. This benefit requires that Forgiving Misperception be restricted to avoid creating excessive selective pressure for Punishing Misperception. Beyond a certain threshold of forgiveness, Punishing Misperception rapidly

increases in the population, discouraging forgiveness. When Punishing Misperception evolves in an extremely forgiving population, it initially benefits the punishing player at the expense of the forgiving and cooperative players. However, as Punishing Misperception becomes more prevalent it decreases the amount of forgiveness and cooperation in the population. In some cases, it has been observed that the population can evolve Punishing Misperception probabilities that are very close to 1.0, which effectively alters the players' strategy from Tit for Tat to Always Defect. This represents an Evolutionarily Stable State, since Forgiving Misperception cannot successfully invade such a population as it will quickly become exploited.

8.5 Limiting Forgiveness to Prevent Exploitation

While a global limit on the Forgiving Misperception probability did prevent the evolution of exploitative Punishing Misperception, this was a rather severe and highly unrealistic solution to the problem. A more biologically realistic approach that should also yield similar results is to limit how many times a player may forgive an opponent during an iterated game. Once this limit has been reached, the player may no longer forgive its opponent, even if its Forgiving Misperception probability would initiate forgiveness. This eventually shows players that certain opponents cannot be trusted to cooperate, by implementing a measure of *patience* that a player has with an opponent. While a machine learning heuristic might enable the players to better identify exploitation, such a heuristic would likely be overly complicated when contrasted against the simple Tit for Tat IPD players and their forgiveness count. In the interests of maintaining simplicity, such an approach was ultimately rejected. Once a player's patience is exhausted, it may no longer misperceive Defection as Cooperation, allowing it to respond to retaliate against Defection. Such behaviour should prevent exploitation, provided that a reasonable limit on a player's patience is imposed.

Implementing the forgiveness count attempts to limit the benefit of Punishing Misperception by reducing the amount of times that a player can benefit from Defection against a forgiving opponent. This change to the simulation should prevent the widespread evolution of exploitative Punishing Misperception within the population, which ultimately drives Forgiving Misperception from the player population. However, unlike the global

restriction of forgiveness, it is also expected that this change should allow the population to evolve to a stable state where both Punishing and Forgiving Misperception can coexist in the population without runaway Punishing Misperception eventually driving Forgiving Misperception out of the population. This change should also reveal the impact that the existence of a stable sub-population of Punishing Misperception has upon the Forgiving Misperception in the population.

While it would be possible to make the forgiveness count a unique element of each player's chromosome, and therefore subject to evolution, this would introduce a second variable into the simulation that would affect each players' likelihood of forgiving Defections. This change would in turn complicate the analysis of an individual player's success or failure — a player's high score may be due to its Forgiving Misperception probability or its forgiveness count or the combination of the two. It also complicates the measurement of the Forgiving Misperception within a population, since a high Forgiving Misperception probability may not indicate a highly forgiving population; and, additionally, will complicate the comparison between the effects of a forgiveness count and the upper bound on the Forgiving Misperception probability. Since the point of this change is not to identify the optimal forgiveness count evolved by the players, but to determine whether stable populations with both Forgiving and Punishing Misperception can coexist, the forgiveness count will be the same population-wide.

8.5.1 Method

The existing Iterated Prisoner's Dilemma parametric simulation was modified to implement a limit on how many times a player could forgive its current opponent in each iterated game, which is the same for all players in the population. This forgiveness count takes the role of the Upper Bound on Forgiving Misperception, but instead aims to indirectly limit Punishing Misperception by allowing eventual retaliation against exploitative Punishing Misperception. At the beginning of each contest, both players' forgiveness count is initialised with the global parameter. During the contest, whenever Forgiving Misperception would cause a player to forgive a Defection, the player's forgiveness count is first checked to determine whether the Defection may be forgiven. If the forgiving player's forgiveness count is greater than zero, forgiveness causes the Defection to be misperceived as Cooperation and decrements the forgiveness count by one. If the forgiveness count is zero,

Forgiving Misperception does not occur and the perceived Defection is correctly perceived, thereby allowing Defection against the opponent. Once a player's forgiveness count has reached zero, it is as if the player has exhausted its patience with its current opponent (Appendix E.7 details how the forgiveness count was implemented in the simulation).

The values of the forgiveness count examined were 5, 10, 15, 20, 25, 30, 35, 40 and 45. These values were run for all of the misaction probabilities previously studied (0.005, 0.01, 0.02, 0.03 0.05, 0.1), allowing the efficacy of the forgiveness count to be compared to that of the Forgiving Misperception probability upper bound. As previously, each parameter set was run 30 times to provide a suitable sample size for the population data, which necessitated running 1,620 ($9 \times 6 \times 30$) individual simulations.

8.5.2 Results

The benefits of limiting forgiveness should manifest in much the same manner as the Forgiving Misperception probability upper bound and should also produce similar effects. Lower forgiveness counts should allow the populations to maintain their cooperation during contests, which will be observable in the Forgiving Misperception probabilities and scores received by the players. Higher forgiveness counts should permit limited exploitative behaviour, allowing Punishing Misperception to invade player populations when such exploitation reaches the point that it is beneficial. This invasion should reduce the population's score.

Individual Populations

Examining the results from individual simulation runs reveals how the forgiveness count can provide functionally equivalent results to those obtained with the upper bound on the Forgiving Misperception probability. While a wide range of forgiveness values were used as simulation parameters, here forgiveness counts from 20 to 35 and a misaction probability of 0.05 are the focus; these values allow Punishing Misperception to develop in the population, but do not allow it to dominate the population at the expense of Forgiving Misperception.

These plots are not intended to be representative of all runs for that combination of forgiveness count and misaction probability, but they have been selected to demonstrate

three distinct behaviours (Appendix E.8 contains additional plots detailing the behaviours within the populations).

Figure 8.15 shows a stable population that has evolved both Forgiving and Punishing Misperception. In this population, a low level of Punishing Misperception evolves, which rises and falls in step with Forgiving Misperception. The scores obtained by this population indicate that the population is approaching the theoretical maximum obtained from mutual Cooperation, with several small decreases corresponding to rises in the average Punishing Misperception probability. This result shows that a forgiveness count of 20 does allow coexistence between Forgiving Misperception and Punishing Misperception, but in this case Punishing Misperception is prevented from dominating the population.




Figure 8.15: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Forgiveness Count = 20, Run 21:** Stable population with both Forgiving and Punishing Misperception, where Punishing Misperception is quite low and the population is fairly cooperative.

Figure 8.16 shows a stable population with a forgiveness count of 30, where both Forgiving and Punishing Misperception evolve to coexist in the population. After approximately 1,000 generations, Punishing Misperception evolves in the population and stays stable at approximately 20%. This suggests that 20% is the optimal Punishing Misperception probability for a forgiveness count of 30 and a misaction probability of 0.05. In the earlier simulations where an upper bound restricted Forgiving Misperception, such a high level of Punishing Misperception would typically lead to runaway Punishing Misperception (see Figure 8.10). The population also experiences a lower average score of approximately 12,000, which is greater than the 10,800 that would be obtained by random strategy selection, yet equal to that obtained from Alternating Cooperation and Defection.




Figure 8.16: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Forgiveness Count = 30, Run 9:** A population where a stable Punishing Misperception probability has evolved, while the forgiveness count prevents exploitative Punishing Misperception from evolving.

This population demonstrates a threshold between the populations where Forgiving Misperception allows the players to maintain mutual Cooperation and the populations where runaway Punishing Misperception has evolved. The equilibrium state that has evolved in this population exists between those two extremes, with an increased forgiveness count allowing runaway Punishing Misperception to evolve and a lower forgiveness count better discouraging Punishing Misperception from evolving the population.

Figure 8.17 shows a population with a forgiveness count of 35, where runaway Punishing Misperception has evolved. This has occurred because the higher forgiveness count permits players to use Punishing Misperception to successfully exploit their opponent's forgiveness more often. Once the average Punishing Misperception probability becomes sufficiently high and the average Forgiving Misperception probability begins to decrease, there is no disincentive to Punishing Misperception and it evolves towards 100%, at which point the players are effectively playing Always Defect. This is beyond the threshold at which Forgiving Misperception is beneficial, as the higher forgiveness count allows sufficient Forgiving Misperception to permit widespread exploitation from Punishing Misperception.

By using the forgiveness count to limit the opportunities for Punishing Misperception to exploit forgiving players, it has been possible for Punishing and Forgiving Misperception to coexist in a stable population. However, populations where Forgiving and Punishing Misperception coexist will have a lower average score, due to Punishing Misperception's disruption of the cooperation between the players.

All Populations

The effects of the forgiveness count can be better understood by examining how it affects the various different parameter sets over the complete 30 runs. The results obtained with the forgiveness count are comparable to those previously obtained with the Forgiving Misperception probability upper bound (Section 8.4).

Figure 8.18 shows the average forgiving misperception probability as measured in the final generation (generation 10,000) of each simulation. This shows that the average Forgiving Misperception probability gradually declines as the forgiveness count is increased, allowing players to forgive more and opening the door for exploitative Punishing Misperception. When compared with the results from the Forgiving Misperception probability upper bound (Figure 8.11), the forgiveness count allows for a gradual decrease in the population's average Forgiving Misperception probability, as opposed to the sharp drop exhibited when the upper bound exceeds the optimal Forgiving Misperception probability.

The average Punishing Misperception probabilities are shown in Figure 8.19 and clearly show that the average Punishing Misperception probability increases in the population as the forgiveness count increases. Punishing Misperception begins to noticeably evolve in the population once the forgiveness count reaches 25. While the forgiveness count is 25 or less Punishing Misperception is not highly adaptive, as there is insufficient forgiveness for Punishing Misperceivers to regularly exploit. The small benefit that can be obtained




Figure 8.17: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Forgiveness Count = 35, Run 14:** An unstable population affected by runaway Punishing Misperception. Forgiving Misperception is completely driven from the population, while Punishing Misperception becomes ubiquitous.




Figure 8.18: Average Forgiving Misperception Probability, plotted for the various misaction probabilities and Forgiving Misperception Probability Upper Bounds. The average Forgiving Misperception Probability gradually decreases as the forgiveness count is increased (Appendix E.9.1 shows other plots of this data).




Figure 8.19: The Average Punishing Misperception Probability, plotted for the various misaction probabilities and Forgiving Misperception Probability Upper Bounds. There is a gradual increase in the average Punishing Misperception probability as the forgiveness count is increased, which permits more Punishing Misperception to be forgiven (Appendix E.9.2 shows other plots of this data).

in these situations is demonstrated by the low Punishing Misperception probabilities observed up to this point. Once the forgiveness count is increased beyond 25, Punishing Misperception becomes more beneficial and increases its representation in the populations. When the Punishing Misperception probabilities are compared to those obtained with the Forgiving Misperception probability upper bound (Figure 8.12), the increase in the Punishing Misperception probability is much more gradual.

Figure 8.20 shows the average final scores received by the player populations, which initially shows the average final score gradually increasing to the optimal level where players are able to cooperate, despite the environmental noise and lower levels of Punishing Misperception. As the forgiveness count increases, exploitation becomes more likely, which in turn causes the average final scores to decrease. The decrease in the average final scores corresponds with the increase in the average Punishing Misperception probabilities.

When compared against the graph of the scores obtained with the Forgiving Misperception probability upper bound (Figure 8.13), the similarities are striking. Both graphs show very similar increases in the average final score as the upper bound and forgiveness count is increased up to its optimal level. The major difference is in the decline of the average final score as Punishing Misperception increases in the population. With the forgiveness count, the decline is much more gradual, corresponding to the gradual increase in the average Punishing Misperception probability.

Summary

Restricting how many times a player may forgive its opponent is another method that can prevent exploitative Punishing Misperception from dominating cooperative populations who utilise Forgiving Misperception. However, it also allows populations to evolve stable probabilities of both Forgiving and Punishing Misperception simultaneously, which was not possible when an upper bound restricts the Forgiving Misperception probability. It is also a more realistic solution to the exploitation produced by Punishing Misperception.

These results suggest that the optimal forgiveness count for maintaining mutual cooperation and its corresponding high score surrounds 15 and 20. However, if Punishing Misperception must be represented in the population, then the optimal forgiveness count surrounds 25 and 30; allowing both Forgiving and Punishing Misperception to coexist without excessive exploitation of Forgiving Misperception driving it out of the population.




Figure 8.20: The Final Average Score, plotted for the various misaction probabilities and Forgiving Misperception Probability Upper Bounds. As the forgiveness count increases, the average score increases up to the theoretical limit available for mutual cooperation, before gradually decreasing once Punishing Misperception evolves to exploit the forgiving players (Appendix E.9.3 shows other plots of this data).

The downside to this coexistence is that players receive lower scores, due to the greater prevalence of exploitation as a result of Punishing Misperception. Beyond a forgiveness count of 30, there is excessive forgiveness that may be exploited, leading to runaway Punishing Misperception.

8.6 The Relationship Between Forgiveness and Punishment

In the previous sections, data from some simulated populations suggested that the relationship between Forgiving Misperception and Punishing Misperception sometimes appeared to inversely correlate, with increases and decreases in the probabilities of each correlating to increases and decreases in the other. In some cases this interaction expressed cyclic properties, which in many instances are quasi-periodic. This quasi-periodic relationship appears to be similar to that exhibited between populations of predators and prey. The Lotka-Volterra (Lotka, 1920; Volterra, 1928; Edelstein-Keshet, 1988) equations describe a deterministic model of the cyclic relationship between a population of predators and prey, where the increases and decreases of the size of the prey and predator populations drive corresponding changes in each other.

The dynamics of the populations are given by the linear equations describing the relationship between predators (Equation 8.2) and prey (Equation 8.3) (DeRoos, 2002), where b represents the birth rate for the prey, d represents the death rate for the predators and a represents the interaction constant between the predators and prey.

$$N_{Prey}[t + 1] = N_{Prey}[t] + bN_{Prey}[t] - aN_{Prey}[t]N_{Predator}[t] \quad (8.2)$$

$$N_{Predator}[t + 1] = N_{Predator}[t] - dN_{Predator}[t] + aN_{Predator}[t]N_{Prey}[t] \quad (8.3)$$

Since Punishing Misperception can exploit Forgiving Misperception, Forgiving Misperception could be compared to prey behaviour and Punishing Misperception could be compared to predator behaviour. While the interactions within some populations of misperceiving Iterated Prisoner's Dilemma players may resemble a predator-prey relationship,

there are several notable differences between the simulated populations of players and a typical biological population of predator and prey organisms. These are:

1. Punishing Misperception (the predator) doesn't die out within a population after Forgiving Misperception (the prey) dies out. Instead the Punishing Misperception probability typically increases towards 100%.
2. Punishing Misperception 'predates' upon others with Punishing Misperception. In a conventional predator-prey model, predators only consume prey.
3. A conventional predator-prey model assumes that the prey population experiences exponential growth in the absence of predators. Absent Punishing Misperception, Forgiving Misperception typically will not rise beyond 50% and cannot exceed 100%.
4. The measurement of the expression of two misperception probabilities, and the range of possible probabilities, in a fixed population is different to the measurement of the size of two populations.

The predator-prey equations describe an interspecies competition between populations of predator and prey species, while the simulation describes an intra-species competition between two different behaviours. While the simulation does not directly model a predator-prey relationship, it may be that the relationship between Punishing and Forgiving Misperception can sometimes approximate this relationship. While the predator-prey model does not exactly describe the simulated environment, the underlying relationship between the two misperception behaviours may still be considered in terms of these models to determine to what degree the relationships are similar.

The predator-prey equations provide a deterministic model against which the simulation data that exhibits cyclical behaviour can be compared. One instance of a cyclical relationship between Forgiving and Punishing Misperception was obtained from Run 11 from Section 8.3 (Figure 8.9), which was obtained when there was no restriction on Forgiving Misperception and the misaction probability was 0.01. For each of these populations, a greedy search was used to optimise a set of coefficients for a , b and d that produced a predator-prey model that best fit the simulation data. The simulation data can then be directly compared to the predator-prey model.




Figure 8.21: Predator-prey data (from coefficients $a = 0.000326$, $b = 0.001614$ and $d = 0.008739$) plotted against a quasi-periodic average Punishing and Forgiving Misperception probabilities, which were obtained with unrestricted Forgiving Misperception. Peaks in both the average Forgiving and Punishing Misperception probabilities only very loosely match the periodicity and amplitude of the predator-prey relationship.

The search identified the set of coefficients of $a = 0.000326$, $b = 0.0016147$ and $d = 0.008739$ as a good match for the unrestricted Forgiving Misperception simulation data. Figure 8.21 shows the predator and prey relationship calculated from these coefficients plotted against the simulation data with unrestricted Forgiving Misperception. While the peaks and troughs of the simulation data roughly approximates those of the predator-prey populations, the predator-prey model is otherwise a poor fit for the highly stochastic simulation data. A statistically significant test result shows that the model almost certainly did not generate the data.

Another population that exhibited a cyclical relationship was identified; however, the predator-prey model also poorly fits this population's recorded simulation data (Appendix E.10 contains this comparison).

Since the majority of the observed simulations do not even establish cyclic interactions between Forgiving and Punishing Misperception, it cannot be argued that the average Forgiving and Punishing Misperception probabilities universally exhibit a predator-prey relationship. Analysis of data sets that do exhibit cyclic interactions finds only a loose correlation to the predator-prey model. At best there are some connected offset fluctuations between Forgiving and Punishing Misperception, where increases or decreases in Forgiving Misperception are followed by increases or decreases in Punishing Misperception, but

these fluctuations lack well-defined and consistent periodicity. Furthermore the statistical insignificance of the comparison between the simulation data and the predicted results disproves the hypothesis that these two simulation populations exhibit a predator-prey relationship.

However, there are plenty of instances where the simulated populations exhibit offset fluctuations in Forgiving and Punishing Misperception probabilities (this is clearly exhibited in the individual simulation runs in Appendix E.4). Given the prevalence of the offset fluctuations with little to no periodicity, it could be argued that these fluctuations are the only part of the predator-prey relationship modelled within the simulation. However, offset fluctuations are disturbed by runaway Punishing Misperception, which causes a decrease in the Forgiving Misperception probability that is not followed by a decrease in the Punishing Misperception probability. Therefore the simulated populations are at best capable of exhibiting only part of the predator-prey model.

8.7 Categorising Player Populations

While viewing how an individual population's average score and misperception probabilities change over time can easily reveal whether the population has collapsed into a state of runaway Punishing Misperception or is cooperating due to its Forgiving Misperception, there are no simple heuristics for classifying the simulated populations that do not neatly match these two circumstances. In order to label such populations, it is necessary to develop a categorisation scheme that describes the the states of the simulation. Given that the main evolutionary measure of success within the simulations is a player's score, the population's average score provides a simple, single value by which to categorise the populations. While the Forgiving and Punishing Misperception probabilities could also be used to categorise the populations, the relationship between the two probabilities requires the analysis of two linked variables, while the average score is a single variable.

This categorisation was developed by determining how the final average scores from the simulated populations were distributed within the range of possible scores. This approach was applied to the data from the simulations in Section 8.5, which were those that utilised the Forgiveness Count to control potentially excessive forgiveness. To identify how the average final scores of the simulated populations were distributed, the average final scores

were mapped into buckets of varying sizes. Figure 8.22 displays this data for buckets of size 400 and 100. The peaks on this plot indicate a higher frequency of simulations achieving final average scores within the score range dictated by that bucket, where higher frequencies may indicate stable equilibria within the simulation.




Figure 8.22: Distribution of the Final Average Score of the simulated populations from all simulation parameter combinations, with frequencies counted for bucket sizes of 400 and 100. The larger bucket size of 400 shows that the bulk of the simulated populations have higher final average scores that are indicative of some degree of mutual cooperation, which was aided and enabled by Forgiving Misperception. However, there is also a tail of values across the middle range of lower scores, along with a small peak towards the lowest average final scores. With the smaller bucket size of 100, the finer sorting method reveals several potential peaks within the data, indicating that there may be several equilibria within the simulations' outcomes.

These plots identify at least two peaks at the higher scoring end (at about 13,500+ and 12,600 to 13,200), depending on the bucket size. Similarly, there are one or two peaks at the lower end of the scores, both below 6,400. Since these peaks display a higher frequency of simulations ending with scores within their range, it seems likely that they cover possible equilibria for the simulation. Based upon these scores, thresholds for distinguishing between the categorisations could be set at 13,500, 12,400, 10,800 and 6,400. Labelling these categories based upon how much cooperation they likely exhibit yields the categorisation scheme in Table 8.4.

Simulation State	Score Threshold
Major Cooperation	Score > 13,500
Moderate Cooperation	Score > 12,400
Minor Cooperation	Score > 10,800
Little to no Cooperation	Score > 6,400
No Cooperation	Score \leq 6,400

Table 8.4: Thresholds for categorising the state of a population.

The threshold for Major Cooperation was selected due to the peaks near this value on the bucket plots. Similarly, the threshold for Moderate Cooperation was selected due to the higher frequency of simulations with average final scores close to 12,600. The threshold for Minor Cooperation was selected to be 10,800, due to the fact that this value represents a score equivalent to random strategy selection. A score of 10,800 is equivalent to random strategy selection, which might occur if the Tit for Tat players were affected by environmental noise but did not otherwise misperceive in any way. In order to be beneficial, misperception must yield a higher average score than that obtained from random strategy selection due to environmental noise. A score below this threshold indicates that the population is performing worse than random strategy selection and, therefore, their misperceptions are collectively detrimental. The populations exhibiting Little to no Cooperation have a second threshold set at 6,400, due to the two smaller peaks in the population counts identified at or below this point. This final category attempts to differentiate between populations that have succumbed to runaway Punishing Misperception and those that have not (but appear likely to do so in the future).




Figure 8.23: Numbers of simulated populations categorised by the average score, computed over their final 2,500 generations and plotted for the simulation parameters pairings of Misaction Probability and Forgiveness Count. As noted previously, each parameter set was run 30 times, yielding the instances within which each categorisation was counted.

8.7.1 Testing the Categorisation

This categorisation was then applied against the latter generations of the simulated populations. The analysis underlying the categorisation focused exclusively on the final average score from the populations. While the final average score does simply demonstrate the final state of the population, it is not always an entirely accurate indicator of a simulated population's results. To guard against such misconceptions, this test instead examines the population's average score over its final 2,500 generations, which corresponds to the last quarter of the simulation's duration. This range was selected because a population is more likely to have either reached some stable equilibrium or be progressing towards it by the end of the simulation. From this final duration the population's average score was then calculated and compared against the categories, with their totals recorded. These results are shown in Figure 8.23.

The strategies are colour coded as blue, green, yellow, orange and red, which are mapped onto the score categories from highest to lowest. The cooler green and blue were selected to demonstrate instances of beneficial misperception, where the average score indicates cooperative players. Yellow was mapped to Minor cooperation, where players are benefiting from their Forgiving Misperception to some degree, but their mutual

cooperation is more regularly disrupted by Punishing Misperception and the environmental noise produced by Misaction. The warmer colours of orange and red were mapped to the instances with lower scores, which indicate little or no cooperation between players. As stated previously, misperception in these populations is clearly detrimental.

Table 8.5 lists the percentages of the simulated populations that fit into these categories, showing that more than 75% of the populations received average scores over their last 2,500 generations that indicated some degree of beneficial mutual cooperation, produced by Forgiving Misperception.

Simulation State	Percentage
Major Cooperation	30.86%
Moderate Cooperation	28.27%
Minor Cooperation	18.64%
Little to no Cooperation	15.25%
No Cooperation	6.98%

Table 8.5: The percentage of the simulated populations for all parameter sets that fall into each of the five categories. Note that more than 75% of the populations are benefiting from misperception, possessing scores that exhibit some degree of mutual cooperation.

As the Misaction Probability increases, the average score decreases since cooperation between players is more likely to be disrupted. This trend can be seen in the plot moving from right to left along the horizontal axis, with the instances of Major Cooperation (blue) gradually replaced by Moderate Cooperation (green). This transition visually displays the gradual effects of the increased environmental noise upon the cooperating players.

8.7.2 Equilibria and Simulation Stability

This work suggests the existence of at least two equilibria within the simulation: runaway Punishing Misperception and mutual cooperation aided by Forgiving Misperception. However, it is currently unclear whether these are the only two equilibria, or whether each equilibrium remains stable over longer simulation durations. For reference, Figure 8.15 and Figure 8.17 respectively demonstrate populations exhibiting these two equilibria.

Classifying the simulations by their final average score yields at least two peaks within the populations (Figure 8.22), each clustered around the simulation's two equilibria. The larger peak contains over 75% of the simulation populations, with average final player scores greater than 10,800 indicating mutual cooperation enabled by Forgiving Misperception. The smaller peak occurs for the lower scores, indicating a number of populations

affected by runaway Punishing Misperception. However, there is also a sizeable number of populations between these two equilibria, suggesting that not all the populations have stabilised at an equilibrium. This suggests that a longer simulation duration may be necessary to investigate whether there are any intermediary equilibria.

Furthermore, it is not clear whether there is a common point at which the simulations begin to experience runaway Punishing Misperception or at what point the populations stabilise at that equilibrium. From the simulated populations presented previously in this chapter and also in Appendix E.4, it is clear that runaway Punishing Misperception typically begins to develop in the populations after 2,000 generations and typically directs those populations to that equilibrium at some point after approximately 5,500 generations. A superficial examination of these transitions within the 1,620 simulations discussed within Section 8.5 did not identify a specific cause for these collapses or a common time when they occur. Further research to address this question is desirable, but beyond the scope of this work.

8.7.3 Summary

A more detailed examination of the trends exhibited by the populations of misperceiving Tit for Tat players has loosely categorised the populations based upon their average final score. Refining this categorisation revealed that further research is required to better clarify both the categorisations and the equilibria of the populations. While interesting, such work does not directly address the main question of this thesis; specifically whether misperception can provide an evolutionary benefit. Instead such work would further clarify the known conditions under which misperception provides an evolutionary benefit.

Categorising the state of the player populations by their average scores provides a quick indicator of their “success”, from which their levels of Forgiving and Punishing Misperception can also be roughly estimated. Further research into the equilibria of the simulation would likely provide improved heuristics for identifying and categorising the behaviour exhibited by the populations, subsequently improving the accuracy of the developed categories.

Looking at both the categorisation of simulation populations, their equilibria and their stability suggests that a duration of 10,000 generations for the evolutionary simulations is insufficient in some circumstances. While this value was adequate for demonstrating

the evolutionary benefit of Forgiving Misperception, it seems to be insufficient for determining the long-term prospects of Forgiving Misperception, specifically in the face of runaway Punishing Misperception. For the simulation duration of 10,000 generations that was tested, the middle categories of Minor Cooperation and Little to no Cooperation encompassed about one third of the populations. If these categories are not equilibria, then utilising a longer duration may allow a larger percentage of the simulations to move towards the equilibria of Major Cooperation or No Cooperation. A longer duration may also help clarify the conditions required for the evolution of the occasional quasi-periodic relationship observed between Forgiving and Punishing Misperception.

8.8 Conclusion

The usage of misperception as a method of maintaining or restoring Cooperation in a noisy Iterated Prisoner's Dilemma game initially seems counter-intuitive; however the results obtained here have shown that it can benefit players in some limited cases. Specific types of misperception can restore Cooperation between Tit for Tat players in situations where they are otherwise incapable of resuming Cooperation. When misperception provides this benefit it produces a similar outcome to forgiveness although by a very different means, essentially relying on one player to make an error that affects its actions. Due to the widespread applicability of the Iterated Prisoner's Dilemma and the Tit for Tat strategy, any insights revealed here may well be applied in many other domains where ongoing cooperation is valued.

8.8.1 Restoring Cooperation with Misperception

Standard misperception did not benefit the Tit for Tat players in a noisy environment, but instead acted as a secondary source of noise in conjunction with the players' misactions. While misperception was capable of restoring the Tit for Tat players to a cooperative state, it was just as likely to disrupt cooperation or initiate mutual Defection. The noise introduced by either misperception or misaction actually randomly switches the Tit for Tat players between game states, yielding an average payoff similar to that received from the random selection of strategies. Therefore, it does not achieve the desired aim of restoring the disrupted Cooperation between the players.

The Trickster also proved to have little effect on the players in the simulation. The Trickster's initial influence forces Tit for Tat players in a noiseless environment into a sub-optimal state of Alternating Cooperation and Defection. However, in a noisy environment its effects are quickly overshadowed by the many errors of either misaction or misperception. Furthermore, the payoffs that players receive in this alternating state are actually greater than those received when misperception is enabled, showing that misperception is an unsuitable solution to the problem posed by the Trickster. While the Trickster as implemented in this simulation was ultimately ineffectual, the concept of a malevolent external party that may affect the players of a game is an interesting avenue for future research.

Symmetric misperception does not benefit Tit for Tat players by restoring mutual cooperation, since it is also likely to disrupt any cooperation. However, a benefit was observed in a noisy environment when misperception was restricted to "forgive" Defections caused by the noise, thereby restoring Cooperation. Forgiveness is a known method of counteracting noise in the Iterated Prisoner's Dilemma and these results show that misperception is beneficial when it causes unintentional forgiveness. Unintentional forgiveness can be a suitable method of maintaining or resuming cooperation in situations where the players are not otherwise forgiving. Such conditions commonly exist between nations, companies or individuals who are mutually hostile or competitive.

8.8.2 Forgiving and Punishing Misperception

When misperception is divided into Punishing and Forgiving Misperception, it was hypothesised that the player population would evolve a high level of Forgiving Misperception and a low level of Punishing Misperception to deal with a noisy environment. However, while the populations did evolve high Forgiving Misperception probabilities, these probabilities made them vulnerable to Punishing Misperception, which evolved in the population at the expense of forgiveness. A forgiving population provides an excellent opportunity for Punishing Misperception to exploit the forgiving players, which provides Punishing Misperception with an evolutionary advantage over the forgiving population. Punishing Misperception evolves in the population because it is initially individually beneficial, yet this benefit disappears once it becomes widespread throughout the population. As Punishing Misperception becomes widespread, Forgiving Misperception is reduced in the

population due to its inability to compete. After some time, it may even die out in the population, leaving a population of Tit for Tat players in a noisy environment who often respond to Cooperation with Defection. While Forgiving Misperception does benefit the player populations, its potential exploitation by Punishing Misperception results in payoffs to the Tit for Tat players that are worse than that of noise by itself. This exploitation and runaway evolution of Punishing Misperception commonly occurs once the population's Forgiving Misperception is above 30%. The exploitative nature of Punishing Misperception indicates that a high Forgiving Misperception probability is not Evolutionarily Stable, as Punishing Misperception will easily invade it.

The evolution of exploitative Punishing Misperception can be prevented by limiting the Forgiving Misperception probabilities of the player population. Adding an upper bound to Forgiving Misperception allowed players to forgive errors, while also avoiding a situation where Punishing Misperception may evolve. This was possible with a Forgiving Misperception probability upper bound of up to 33% in some cases. Beyond this point there was a sharp decline in the scores of the Tit for Tat population, where Punishing Misperception had evolved in the population to exploit the high forgiveness probabilities. Below this threshold, the populations of Tit for Tat players have very low Punishing Misperception probabilities and are able to maintain mutual cooperation. Beyond this threshold, Punishing Misperception prevents mutual cooperation and the population's score decreases below that received by players who randomly select their strategy as if affected by noise. This threshold represents the transition point between beneficial Forgiving Misperception and exploited Forgiving Misperception, which is ultimately detrimental. The highest scores were obtained with Forgiving Misperception probability upper bounds that were less than highest upper bound where the players maintained mutual Cooperation. The threshold where Punishing Misperception evolves to exploit Forgiving Misperception shows that for misperception to benefit the population of a noisy evolutionary environment it must stay below this level. Forgiveness is also exploitable in many other situations and limiting the frequency of forgiveness in those situations is a possible method of preventing such exploitation from becoming commonplace.

8.8.3 Patience Preventing Exploitation

Another method that allows Forgiving Misperception to avoid exploitation from Punishing Misperception was to give the players a level of patience with defecting opponents. This was implemented as a forgiveness count that would allow a player to forgive up to a certain number of Defections per player. Once this count was exhausted, a previously forgiving player was prevented from forgiving any further Defections. This provided a more realistic and biologically grounded solution to the problem of excessive exploitation via Punishing Misperception. This change also permitted Forgiving and Punishing Misperception to coexist in a stable population.

This alteration to the simulation produced similar results to the instances with the upper bound restricting Forgiving Misperception. A lower forgiveness count permits Forgiving Misperception to maintain a mutually cooperative state, despite environmental noise and occasional Defections due to Punishing Misperception. Higher forgiveness counts above approximately 25–30 allow exploitative Punishing Misperception to evolve in the population, detrimentally reducing the population's average score and at higher levels, leading to the evolution of runaway Punishing Misperception.

8.8.4 The Relationship between Forgiving and Punishing Misperception

The occasionally cyclic relationship between Punishing and Forgiving Misperception initially appears to model a predator-prey relationship, with Punishing Misperception as the predator and Forgiving Misperception as the prey. However, analysing populations that best exhibited this cyclic relationship found that the similarities were statistically insignificant. The two instances examined do exhibit a dependency of Punishing Misperception upon Forgiving Misperception, which superficially resembles a predator-prey relationship when the Forgiving Misperception probability fluctuates in a cyclical manner. This behaviour was only observed within a few simulated populations, indicating that the evolutionary Iterated Prisoners Dilemma game with asymmetric misperception does not model a predator-prey relationship between Forgiving and Punishing Misperception.

Chapter 9

Conclusions

This thesis has explored the hypothesis that, in some evolutionary circumstances, misperception benefits a misperceiver or its population. For a hypothetical example of beneficial misperception, consider Caesar's deception of the Gauls besieging his camp from Chapter 1 (Caesar and Handford, 1951). While the Gauls' misperception was clearly detrimental, what might have happened if it was not? For example, what might have happened if some of the attacking Gauls misperceived a weakness in the seemingly reinforced gates? This belief would have likely induced the misperceiving Gauls to direct their attacks against the camp's gates, potentially allowing the Gauls to break into the Roman camp, pre-empting and negating the successful Roman counter-attack. This hypothetical misperception benefits the misperceiving Gauls and similar cases of beneficial misperception can likely be constructed for many historical events and conflicts. This thesis aims to better explore instances where such benefits arise.

Analysis of misperception and its various sources reveals that misperception, whether beneficial, neutral or detrimental, causes errors as an entity collects and processes information. Many of these effects occur when entities interpret information. If misperception is assumed to be detrimental, and therefore undesirable, then such misperceptions may be recognised and possibly prevented or counteracted. Results from the evolutionary simulations support the hypothesis that misperception can provide an evolutionary benefit, although such a benefit is not universally observed. This evolutionary benefit indicates misperception's potential advantage to biological systems, while the situations examined by these experiments may also model behaviours found in social systems. Future research

work is required to identify other instances where this benefit occurs and further identify the common properties of such instances.

9.1 Research Findings

9.1.1 Misperception and the OODA Loop

Mapping the intentional and unintentional sources of misperception into the cognitive model, using the OODA loop construct (Boyd, 1996), reveals that misperception primarily affects an entity's information collection and processing steps. Through various processes misperception causes an entity to produce incorrect beliefs, which the affected entity stores and uses in current and future iterations of its decision-making cycle. An entity will retain these incorrect beliefs until new information causes it to correct these beliefs or replace those incorrect beliefs with new incorrect beliefs.

Examining the sources of misperception also revealed that the canonical Information Warfare strategies (Borden, 1999; Kopp, 2000a) have outcomes that are analogous to some of the unintentional errors arising from misperception. Each of the four canonical Information Warfare strategies can, when successful, produce the same outcome as an unintentional dysfunction of the misperceiving entity. This similarity allows such errors to be misidentified; allowing intentional attacks to be camouflaged as unintentional errors of the attacker, with the attacker mimicking victim-hood or incompetence to further conceal the attack from the defender. Conversely, an entity may also incorrectly attribute its own unintentional errors to the non-existent hostile Information Warfare attacks of a perceived or real enemy, thereby creating and reinforcing incorrect beliefs. Consider how the inability of a broken security camera to capture footage of terrorists placing a bomb may be interpreted by paranoid conspiracy theorists. Due to the conspiracy theorists' existing beliefs, they incorrectly incorporate the cause of the camera's failure into their conspiracy narrative, while the camera's actual failure may be due to petty vandalism or a failure of the camera or its connected network.

Applying Nyquist's (1928) sampling theorem to the OODA loop model suggests that an entity needs to operate at a tempo at least twice as fast as its opponent, in order to accurately sample the opponent's actions (Kopp, 2011). In any conflict between two evenly matched competitors, the inability of both competitors to simultaneously operate at a

tempo that satisfies the sampling theorem inevitably produces a mutual misunderstanding of each other's actions and intentions. This insight encapsulates the idea of the fog of war.

Developing a model of the internal process of the OODA loop's Orientation step further clarifies how entities analyse new information to create and maintain their beliefs. Errors during an entity's Observation step may provide new incorrect information, which will cause further problems during the Orientation step. The Orientation step acts as a filter between new information the entity collects and its beliefs. Misperception produces errors in an entity's Orientation step either by affecting the information received or by affecting the analysis and interpretation of new information. Therefore such errors occur during the Orientation step's Identification and Interpretation sub-steps. While errors may potentially occur at any sub-step of the Orientation step, those that produce misperception are limited to these early sub-steps where the entity creates, manipulates and stores its beliefs.

This model also disproves the hypothesis that Self-deception is a self-targeted Corruption attack. Corruption attacks cause errors during the Identification sub-step, when information is incorrectly identified as whatever it is mimicking. Self-deception instead causes errors during the Interpretation sub-step, where Self-deceiving entities manipulate their own beliefs or the new information to maintain cohesion between their beliefs.

Since an entity's current beliefs affect how it collects and analyses information, any incorrect beliefs it holds may cause it to develop further incorrect beliefs. This process is called incestuous amplification, and describes a repeated process where an entity's false beliefs drive future misperceptions, leading the entity to develop more false beliefs. Each iteration of this process moves the entity's understanding of its reality further from its actual state. Information Warfare attacks may also attempt to create the beliefs that initiate or reinforce incestuous amplification in a victim.

9.1.2 Evolutionary Value of Misperception

Re-implementing the foraging simulation described by Akaishi and Arita (2002a) as an evolutionary simulation proved their hypothesis that misperception can provide an evolutionary benefit to foraging agents. In such a foraging environment, competition over resource nodes strongly impacts the foraging success of agents and misperception beneficially alters the foraging patterns of affected agents. The results obtained here demonstrate that

statistically significant non-zero misperception probabilities can evolve in their foraging scenario. Misperception's benefit to these foraging agents may accrue to either the misperceiving agents, their kin, unrelated agents, or any combination of these simultaneously. Animals competing over food or companies competing over markets are real-world scenarios that are analogous to the simulated environment, and these scenarios may demonstrate similar behaviours.

Akaishi and Arita (2002a) attributed misperception's benefit to the increased diversity it produced in the behaviour of the agent population. Behavioural diversity occurs when agents abstain from foraging in the same locations. Any agents foraging in the popular area benefit from the decreased competition for resources, while the misperceiving agent may benefit if its misperception directs it to a less competitive area. Misperception may also prove beneficial by resolving deadlocks between agents obstructing each others' movement, however such a benefit is due to the design of the simulation and would therefore be unlikely to aid most biological organisms. In these simulations increased foraging success directly translates to increased reproductive success, as more resources allow an agent to parent more offspring.

Other methods of introducing behavioural diversity to the agent population, such as misaction and random movement, also produced benefits similar to those of misperception. These foraging methods, along with misperception, provided an advantage in some cases over populations of agents unaffected by misperception. This suggests that in this foraging environment, sources of behavioural diversity can provide some benefit to the agent population. Further research is required to determine whether this hypothesis holds true for other foraging environments.

The simulated environment was quite safe for the agents, as their lives were only threatened by starvation and old age. Typically the penalties for misperception or misaction were not severe. However, in a more competitive environment, where predation occurs or incorrect beliefs have fatal consequences, misperception may be less likely to evolve in the agent population, even if it can sometimes benefit individuals or groups. When there is more pressure to avoid fatal mistakes, the benefit to misperception or misaction may disappear completely. Again, further research is required to determine whether the observed benefit from misperception extends to situations where an individual's behavioural diversity may produce a highly detrimental outcome.

9.1.3 Tit for Tat, Cooperation and Misperception

In the noise-affected Iterated Prisoner's Dilemma game (Wu and Axelrod, 1995) played by Tit for Tat players (Axelrod, 1984), symmetric misperception does not allow the players to maintain Cooperation, since it is just as likely to disrupt Cooperation as it is to restore it. Symmetric misperception eventually produces an outcome equivalent to random strategy selection by the affected Tit for Tat players. In order to benefit these players, symmetric misperception must either be externally guided or restricted in its occurrence.

When asymmetric misperception is used to counteract the effects of noise, Forgiving Misperception maintains mutual Cooperation between players, while Punishing Misperception disrupts it. The simulations demonstrated that Tit for Tat players will evolve Forgiving Misperception to maintain mutual Cooperation in a noisy evolutionary environment. However, once forgiveness reaches a certain threshold, Punishing Misperception evolves in the player population to exploit the forgiving players. Such exploitation increases an individual player's score at the expense of other players, but detrimentally affects the population if it becomes widespread, yielding average payoffs worse than those received from random strategy selection. An especially important observation from these experiments is that misperception does benefit the players when it acts as noise-correcting forgiveness, but excessive forgiveness is unstable and therefore highly vulnerable to exploitation by Punishing Misperception.

Punishing Misperception evolves to the detriment of the player population when the average Forgiving Misperception probability is approximately 30%. Up to this point, Forgiving Misperception maintains mutual Cooperation between the players, despite the disruptions caused by noise. However, beyond this point the population evolves high Punishing Misperception probabilities, to their ultimate collective detriment. Limiting the Tit for Tat player population's Forgiving Misperception probability to below this threshold allows the population to benefit from forgiveness, while avoiding the evolution of high levels of exploitative Punishing Misperception. Similarly, restricting how often Forgiving Misperception may "forgive" a defecting opponent also avoids the evolution of exploitative Punishing Misperception, while the player population benefits from the forgiveness.

9.2 Overall Conclusions

9.2.1 The Relationship Between Misperception and Mutation

Mutation introduces variation into populations of biological organisms, thereby introducing changes — potentially beneficial, neutral or detrimental — into species over time (Patterson, 1999). Selective pressures then determine whether these changes replicate within a population or die out (Darwin, 1859). Misperception, not unlike mutation, introduces variation into an entity’s beliefs, ideas or memes. While most of these variations may be detrimental or neutral in their effect, causing them to die out within the population, a few beliefs may be highly effective and spread within the population. In a similar manner misperception introduces incorrect beliefs to entities that they would not have otherwise rationally developed. If these beliefs provide a net benefit, then selective pressures will ensure that such beliefs, or the mechanisms responsible for their creation, will spread throughout a population.

In other words, *misperception is to belief as mutation is to gene*. Both processes promote typically undirected change within beliefs and genes, and the consequences of these changes are subject to selective pressures, which allows the propagation of changes that are somehow beneficial. However, whether such changes are beneficial, detrimental or neutral is, of course, situational.

9.2.2 Behaviours Produced by Misperception

The previously observed benefits of misperception arose from a variety of sources. In Doran’s (1998) simulations, misperception produced a beneficial friendship cult and deterred foraging in a fatal area. Meyer et al. (2008) also demonstrated that noise, acting as a cause of misperception, can aid foraging ants by encouraging some variation in their foraging behaviours. In Akaishi and Arita’s simulation of a foraging scenario, misperception led to behavioural diversity in the agent population. Misperception caused both forgiveness and punishment in the Iterated Prisoner’s Dilemma simulations. In each of these cases misperception was beneficial when it caused an entity to perform individually or collectively beneficial actions that the affected entity would not or could not otherwise perform. This commonality suggests that misperception is also beneficial in other situations where

it can regularly induce seemingly irrational behaviours that are actually individually or collectively beneficial.

These scenarios that demonstrate beneficial misperception fall into two classes. The first describes the effects of beneficial misperception upon resource collection, where misperception may direct individuals to higher concentrations of resources or drive individuals away from popular contested resources. As described, a population's behavioural diversity aids the communal collection and distribution of resources. The second class describes beneficial misperception within competitive social environments, where misperception encourages behaviour that benefits the community (or distinct groups within the community) over behaviour that benefits the individual and harms the community.

Another potential class encapsulates those scenarios where an external entity's Information Warfare attacks produce beneficial misperception within the victim population for the victim's benefit. The benefit from such attacks may or may not be intentional. Furthermore, such attacks may be discouraged by the structure or societal norms of the affected population. For example, the use of intentional deception against others in human society can have legal, ethical or other ramifications.

When examining Information Warfare attacks from this perspective there are two variables to consider; the attacker's motivations, which may be either altruistic or selfish, and the outcome for the victim population, which may be either collectively beneficial or collectively detrimental. These variables describe four possible general outcomes of an Information Warfare attack. Given that failure is always possible, the actual outcome may not match that intended by the attacker. Therefore, such failure may cause altruistic actions to become selfish, beneficial outcomes to become detrimental, or vice versa.

While Information Warfare attacks are unlikely to unintentionally cause misperceptions that benefit their victims, the fields of advertising and public relations contain examples of deliberately beneficial attacks. The wartime usage of propaganda by nations against their civilian populations provides many such examples. Such attacks typically intended to induce collectively beneficial behaviours, deter detrimental behaviours and increase cohesion within their respective victim populations. However, such propaganda is also likely unethical or dishonest, and it may also deliver immediate or subsequent unintended consequences. Increasing cohesion within the population during wartime may increase

nationalistic feelings, yet after the war it may fuel the development of racist or hard-line nationalist groups.

9.2.3 A Comparison of Simulated Misperception

In the foraging simulation, misperception provided a benefit when it was infrequent. Excessive misperception did not evolve in the agent populations, as it presumably prevents effective foraging by the agents. The Tit for Tat players benefited from much higher probabilities of Forgiving Misperception — of up to 30% — before Punishing Misperception evolved to exploit such forgiveness. The Tit for Tat players evolve higher Forgiving Misperception probabilities, as an increased probability of forgiveness will reduce the duration that the players spend in a non-Cooperative state. Punishing Misperception also provided an individual benefit to Tit for Tat players against forgiving opponents, as punishment increased their score while reducing their opponent's score. However, when Punishing Misperception becomes common in the agent population, Forgiving Misperception disappears from the population, thereby decreasing the opportunities to benefit from punishment. The frequency and scale of misperception that evolves within a population depends upon the properties of the simulation.

Asymmetric misperception in the Iterated Prisoner's Dilemma can benefit the players in two different ways. Forgiveness allows players to maintain a cooperative state despite any environmental noise, benefiting both the forgiving player and its opponents. Punishment benefits individual players when they successfully exploit the forgiveness of other players; however, this benefit disappears as Punishing Misperception becomes common among the players. The more complex agent interactions in the foraging simulation allowed misperception to benefit the agents in several different ways. The foraging agents also had more variety in their misperceptions, being able to misperceive either the location or the existence of resource nodes. Misperception could convince an agent to forage in a less densely populated area to its own benefit, convince a deadlocked agent to head in a different direction and thereby end a deadlock or convince an agent to abandon a foraging area to its kin. These benefits may accrue to either the misperceiving entity itself or others in its population. Kin selection (Hamilton, 1964), group selection (Wynne-Edwards, 1962) or multi-level selection (Wilson and Sober, 1994) may explain how such "altruistic" behaviour may evolve in these populations, despite its potentially detrimental outcomes.

Misperception in these simulations was modelled in the traditional manner of a random unintentional error that occurs during the affected entity's Observation step. Both simulations also implemented misaction, which was modelled as a random error during the entity's Action step, to affect the entity's behaviour in a similar manner to misperception. In the foraging simulation, misperception and misaction both introduced behavioural diversity into the agent population. Misaction may only affect an entity's behaviour immediately; however, the consequences of any erroneous behaviours may affect the entity for a long time. The incorrect beliefs created by misperception may affect entities long after any misperception occurs. This long-term effect of misperception did not occur in the Iterated Prisoner's Dilemma games, as the Tit for Tat players only rely on information gathered from the previous game to determine their current actions. Hence, misperception can only have an immediate effect on the players. This is likely why the Iterated Prisoner's Dilemma player populations evolved higher probabilities of misperception.

9.2.4 Preventing Cooperation

Results from the simulation of the Iterated Prisoner's Dilemma demonstrated how mutual Cooperation by Tit for Tat players is easily disrupted by various sources of noise, which may include misaction or misperception on the part of the players. Hostile entities external to the game, such as the Trickster, may perform Information Warfare attacks to introduce singular or continual errors, disrupting cooperation between players. It is also possible that some players could desire to disrupt Cooperation between their competitors. There are many situations where entities may benefit from disrupting cooperation between others by inducing errors similar to those caused by misperception or misaction. Such intentional acts may also be disguised as unintentional errors to further aid their success. For example, predators may benefit from disrupting cooperation between prey, when such cooperation affects the predator's hunting success. Corporations may also benefit from disrupting any cooperation between their competitors or customers. Similarly, nations may benefit from disrupting and preventing cooperation between other hostile nations, which might otherwise seek to ally against them. Criminals may wish to disrupt cooperation between rival criminals or law enforcement agencies. Indeed, the Prisoner's Dilemma game describes such a scenario, as the police act to prevent cooperation between the two arrested prisoners.

An extremely effective disruption attack may even cause previously cooperative entities to become hostile towards each other. Hostile entities are unlikely to maintain or resume cooperation, which greatly benefits an attacker. Hostility will also prevent the victims from allying against the attacker. Any Information Warfare attack could potentially disrupt cooperation between players, although covert attacks would be desirable in circumstances where the attacker must interact with the defenders, to conceal the attack and avoid any retaliation for it. Unintentional dysfunctions of an entity may also trigger retaliation against other entities for perceived disruptions that either did not occur or were not intentional. Affected entities may be able to restore cooperation with forgiveness, although this may not be possible in some circumstances. Humans and human organisations use various laws and social conventions to encourage or enforce cooperation and to discourage hostile disruptions of cooperation. Biological organisms also often rely upon social conventions to maintain cooperation, as their chances of survival and reproduction often greatly depend upon inter-group cooperation.

9.2.5 Behavioural Diversity and Deception

In the foraging simulation, misperception created incorrect beliefs that deterred agents from clustering in popular foraging regions, potentially benefiting the misperceiver, its kin and other unrelated agents. Deception, produced by agents performing Information Warfare attacks against others in popular foraging areas, could also create incorrect beliefs, with the intention of driving competing agents from the area. Any of the four canonical Information Warfare strategies could produce this behaviour, although Corruption seems the most suitable to chase an agent from a foraging area and deter its return, while also avoiding any other harm to the agent. If successful, such an attack would benefit the attacking agent and possibly the victim agent, if it is directed to a better foraging area. In this specific case, Corruption may deter some agents from foraging in popular regions and redirect them to other areas. Indeed, Akaishi and Arita's (2002a) simulation demonstrated that if misperception affected communication between the foraging agents, thereby unintentionally deceiving other agents, there was a small benefit to the agent population.

The main difference between these two sources of misperception is that the Information Warfare attacks have a planned effect, while unintentional misperception errors produce

a random outcome. The planned nature of the Information Warfare attacks may cause them to introduce less behavioural diversity in the agent population than unintentional misperception. While it is expected that Information Warfare could also provide an evolutionary benefit to the foraging agents, it is possible that its planned nature may impact its performance as a source of behavioural diversity. If Information Warfare attacks are regularly beneficial to their victims, then the victims should not evolve abilities to prevent such attacks or react with hostility. Furthermore, failed attacks that are identified or fail to produce the desired false belief in the victim could reduce the ability of the attacker to successfully attack that victim in the future. It is expected that as with misperception, Information Warfare may only be beneficial as a method of introducing behavioural diversity when it is infrequent.

9.2.6 An Advantage to Ignorance?

The study of misperception and its effects has demonstrated that there is a potential benefit to the intentional or unintentional possession and development of incorrect beliefs. Various types of misperception, affecting the entity's Observation or Orientation steps, may produce these incorrect beliefs. Ignorant is an apt label for an entity who possesses a number of incorrect beliefs, as the entity does not correctly understand its environment and it is likely unaware of its ignorance. Such an entity is likely to exhibit the "Dunning-Kruger effect" (Kruger and Dunning, 1999).

It should be noted that as with misperception, ignorance is unlikely to generally prove beneficial and that in many cases such ignorance would be detrimental, or possibly even fatal to the entity. However, in circumstances where possessing incorrect beliefs is advantageous and desirable, methods such as Self-deception and Groupthink may allow individuals and organisations to retain these beliefs in the face of contradictory evidence. This suggests another possible benefit that entities may receive from Self-deception — a benefit from the retention or development of beneficial incorrect beliefs, which would otherwise be rejected by rational thought processes. Like the hypotheses for beneficial Self-deception proposed by Trivers (1976) and Ramachandran (1996), such a benefit is entirely situational.

Ignorance, or the various processes responsible for it, are likely to produce and reinforce behavioural diversity, which may prove beneficial to the affected entity, its kin, its community or any combination of these.

9.3 Future Research

Analysing these research findings has revealed several worthy areas for future research of misperception and its effects. These areas include further exploring circumstances under which misperception is beneficial, simulating more complicated models of misperception and studying the evolutionary value of Self-deception. Some of this proposed research will likely reinforce the main hypothesis that misperception can provide an evolutionary benefit.

9.3.1 Beneficial Behaviours Induced by Misperception

This research has demonstrated that when an entity benefits from misperception, misperception induces the affected entity to perform some beneficial behaviour. This behaviour is one that either the misperceiving entity is otherwise incapable of performing or does not wish to perform, such as forgiveness or movement away from resources. Therefore, it is hypothesised that any benefit attributed to misperception arises due to the beneficial behaviour it induces in the affected entities. This benefit is evolutionary in any instance where it aids the reproductive success of entities. A potential avenue for future research is to identify and investigate further situations where misperception induces advantageous behaviour. Various types of altruistic behaviour are one candidate for such study; the misperceiving entity may not desire such behaviour, yet the entity or other members of its population may ultimately benefit from such behaviours. Furthermore, altruistic behaviours are also likely to be easier to evolve than other kinds of misperception. Other interesting behaviours are those that may yield negative immediate consequences and long-term benefits. Any such identified benefit from misperception in an evolutionary scenario would further support the hypothesis that misperception can provide an evolutionary benefit.

9.3.2 Misperception of Attribute Values

As noted in Section 3.4, previous simulations have modelled misperception in a simplistic manner, where its effects are limited to only one aspect of the simulation. This is unlike the real world, where entities may concurrently suffer from a variety of misperceptions, all of which may affect their understanding of any aspect of their environment. While more

complex models of misperception are more realistic, it is unclear whether the observed benefit is exhibited in more complex models.

Such a model of misperception may instead focus on the extent to which misperception alters an entity's beliefs. In such a case, misperception determines not whether an entity is right or wrong, but instead to what degree its understanding of its environment differs from reality. An affected entity misperceives the scale of attributes that describe elements of their environment, such as the value of resources or the relative strength of competing entities. Misperception will affect these entities' behaviour by altering their understanding of the benefits and penalties their actions may incur. Any comparative analysis by the entity of the cost and benefit of its potential actions is now affected to some degree by its misperception.

A simple investigation of this concept could model an environment where the entities compete over resources and may misperceive various attributes of themselves, other entities and the environment. Entities would decide whether or not to contest resources based upon their perception of their own strength, their opponent's strength, the value of the resource and the expected penalty should they lose. With these values an entity can easily determine its expected payoff from any contest and use this information to determine when it should compete and when it should not. This model of conflict may be applied to many types of competition, including inter-species and intra-species contests, military conflict between nations, and economic competition between businesses. If misperception does provide an evolutionary benefit in this scenario, then it will do so by altering an entity's desire to contest resources in a beneficial manner. In the described scenario, misperception may lead to overly aggressive entities, who fight when it is strategically unwise, and unnecessarily passive entities, who avoid most contests, even those they were capable of winning. Both aggressive and passive strategies have strengths and weaknesses, which may contribute to an advantage in some situations. Similarly the interaction between sub-populations utilising aggressive and passive strategies may also yield interesting results.

9.3.3 Evolution of Self-Deception

Self-deception is a potential cause of intentional misperception, which is commonly assumed to be unique to humans and human organisations. Several hypotheses have been

proposed to explain how Self-deception is beneficial, thereby justifying Self-deception's evolutionary development. Ramachandran (1996) suggests that Self-deception conceals and manipulates beliefs that are dissonant with the entity's existing beliefs, to reduce any psychological discomfort they may cause. Trivers (1976) suggests that Self-deception is used to aid the successful performance of deception. Contrary to these arguments, Van Leeuwen (2007) hypothesises that Self-deception is an unintended artifact of the human cognitive system and any benefit it provides is coincidental. As discussed earlier, another potential benefit from Self-deception may arise from its usage to aid the possession or development of beneficial incorrect beliefs.

These hypotheses could be tested in evolutionary simulations, which would investigate whether Self-deception can evolve to benefit agents in the methods hypothesised. Such work would help to explain the possible contributions of Self-deception to human and animal decision-making. Furthermore, since Self-deception causes misperception, any scenarios that show an evolutionary benefit from Self-deception also support the hypothesis that misperception is evolutionarily beneficial.

9.3.4 Examining Behavioural Diversity

Since introducing behavioural diversity benefits the agents in the simulated foraging scenario, it may also provide a similar benefit to entities in other comparable situations. Conceivably this benefit could be demonstrated in any other situation where there is competition over a limited resource, assuming that misperception or misaction can direct agents away from popular resources to less popular resources. This increased diversity may also potentially provide better utilisation of the overall network of resources, which may further benefit the population.

Real-world examples of the benefits of behavioural diversity may be found in various types of networks, where random processes can introduce behavioural diversity that may subsequently reduce congestion and thereby improve network throughput. Two potential types of networks for this research are congested transportation networks and routing algorithms for communication networks, although any other area also covered by the field of queuing theory would be suitable.

In the specific case of peak-hour travel through a congested transportation network, commuters are contesting access to elements of the transportation network, including

roads, train stations, intersections, bus and tram stops. Here it is assumed that commuters are attempting to minimise the duration of their commute; however commuters may desire to optimise other variables, including cost, comfort, ease and environmental impact.

When minimising their travel duration, commuters will select the transportation route and method, or combination thereof, which allows them to spend the least time travelling. Commuters with similar destinations and departure points will select similar transport methods and routes, causing congestion that detrimentally impacts each others' travel time. Introducing behavioural diversity into the population of commuters may allow them to either alter their route to avoid congestion or alter their departure time to avoid peak times. While such an act is often chosen in a selfish desire to reduce an individual commuter's travel duration, it also provides a small altruistic benefit to those commuters on the routes and transportation methods that were not selected. In some cases it is possible that increasing the behavioural diversity of the commuting population through misperception or misaction could crudely increase the throughput of the transportation network.

9.3.5 Further Research of Foraging and Misperception

The generic model of a foraging scenario described by Akaishi and Arita can be considered a proxy for many scenarios that involve competition over limited resources. On this basis, it seems likely that the benefits attributed to misperception should also accrue in other similar scenarios. Such scenarios should model instances where agents compete over access to resources from several sources. This hypothesis could be tested further by investigating the effects of misperception in different foraging scenarios. Misperception could affect how agents value resources, where agents perceive resources or how agents judge the costs of directly contesting access to resources. Such research would attempt to further explore the conditions under which misperception aids foraging agents and identify other foraging environments where this benefit arises.

A specific real-world example of this model of resource collection and competition avoidance is described by market stratification, where multiple companies competing for access to customers within a market may instead choose to specialise into different niches within this market. Such specialisation can directly benefit a company, while also indirectly benefiting any competitors in the market niches that the company no longer services. This

benefit is both individual and collective, with both companies benefiting from avoiding competition, thereby reducing their own costs and increasing their potential access to resources. In a simulated environment, where agents model businesses competing for access to customers, misperception could provoke agents to adjust which markets they compete in or the resources they allocate towards competition, potentially to their individual or collective benefit. However, misperception will also affect how businesses perceive the actual value of resources, which may also guide them into unwise competition over resources of little real value.

Alternatively, further research into the biological underpinnings of the foraging scenario may instead investigate the methods in which misperception evolves in the affected populations. Kin selection (Hamilton, 1964), individual selection (Darwin, 1859), group selection (Wynne-Edwards, 1962) and multi-level selection (Wilson and Sober, 1994) are all possible explanations for how misperception evolves within an agent population. Detailed examination of how misperception evolves within populations of foraging agents may clarify which of these selection mechanisms are responsible for the evolution of misperception.

9.3.6 The Role and Effects of Subversion

Deception by the canonical strategy of Subversion stands aside from the other three canonical strategies, as it is the only strategy that targets the victim's behavioural functions (Kopp, 2000a). As such it can alter how a victim both acts and thinks, potentially affecting how the victim understands reality and thereby building a permanent misperception *into* the victim.

Existing work in Section 4.2.4 examines the role of Subversion in creating misperception and altering victim behaviour, either independently or as an element of a compound Information Warfare attack. Given the various ways in which Subversion may occur, and its varied effects, there is much scope for further work examining the usage of Subversion attacks in the real world.

Furthermore, it would be advantageous to further study the implementation of Subversion across the various domains where it occurs, since it is the most complex of the canonical strategies. Within biological systems Subversion is often demonstrated through parasitism, where the affected victim typically fulfils part of the parasite's life cycle. This

usually includes the victim performing harmful or even suicidal actions to aid the attacker's reproductive cycle. Similarly, in computer systems and networks, a Subverted system will act for the attacker and not the operator. Human social systems also suffer analogous attacks, where an organisation may be subject to Subversion attacks from individuals or groups of individuals within part of the organisation or from external attackers. In either instance, the effectiveness of the social system or organisation is reduced, with its resources or actions directed to benefit the attacker, effectively parasitising the organisation. The typical aim of such attacks is to gain or redirect resources, whether financial, material or intangible — such as fame or power.

The further study of parasitism as an instance of Subversion should focus upon known case studies within biological systems, human social systems and computer systems. Such research may also reveal defensive countermeasures against such parasitism, and Subversion attacks in general.

9.3.7 Relationship between Misperception and Noise

The simulations created and discussed in this thesis demonstrate that misperception, misaction and other sources of behavioural noise can all produce similar results and are, in some instances, interchangeable. The benefit attributed to misperception may therefore be better considered a benefit that accrues in some situations from modifying one's behaviour without justifying evidence. Noise can also affect an entity in such a manner and may thereby provide a similar benefit.

However, there are some differences between misperception, misaction and noise. Misperception and misaction both represent internal sources of error that are linked to the entity's perception and actions respectively, whereas noise is external to an effected entity. Furthermore, the errors produced by misperception and misaction are unlikely to be purely random, unlike some sources of environmental noise.

The similarities and differences between these potential influences on an entity's behaviour could be explored via agent-based simulation, along with the immediate and second-order impacts of misperception and noise upon the entities.

When sources of behavioural diversity, such as misperception and environmental noise, are equivalent, it is expected that if one is advantageous in a given circumstance, then replacement with the other should also prove beneficial. Increased environmental noise

may therefore be a potential replacement for beneficial misperception. Conversely, in situations where misperception is harmful, environmental noise should also prove harmful.

A useful outcome of such research might be the development of methods for individuals and organisations to identify dysfunction due to various forms of noise or misperception. In scenarios where misperception is undesirable, if entities can apply these identification methods in a timely and accurate manner, then they may be extended to form effective countermeasures against noise, the entity's own perceptual and cognitive errors, or other entities' Information Warfare attacks.

9.3.8 Further Misperception and the Iterated Prisoner's Dilemma

There are many potential questions that further research into misperception's effects on the Iterated Prisoner's Dilemma might explore.

One such area is the examination of different player strategies to determine whether misperception remains beneficial for player strategies other than Tit for Tat. Any such replacement strategy would likely need to fulfil some, if not all, of Axelrod's (1984) criteria for successful player strategies. If populations using a different player strategy also benefit from misperception, then this benefit is more universal than previously suspected.

Another topic for future research work is the identification of equilibrium states within the player populations in the asymmetric misperception-affected Iterated Prisoner's Dilemma. The current work suggests the existence of at least two equilibria: runaway Punishing Misperception and mutual cooperation aided by Forgiving Misperception. However, it is currently unclear whether these are the only two equilibria, or whether each equilibrium remains stable over longer simulation durations.

A related topic is the study of the effects of Forgiving and Punishing Misperception upon populations of Iterated Prisoner's Dilemma game players for many more generations. Simulations of 10,000 generations have shown that not all populations evolve towards either of the two main equilibria. Extending the duration of the simulations might reveal whether populations can evolve to a stable state between the two extreme equilibria or whether populations inevitably collapse to the extreme equilibria. Such work might also identify which conditions cause runaway Punishing Misperception in the population.

In some rare instances, the simulated populations demonstrated a quasi-periodic relationship between Forgiving and Punishing Misperception, which superficially resembled a

predator-prey relationship. While not an exact match for a predator-prey relationship, the observed quasi-periodicity might be examined further, as another potential stable equilibrium of the simulation. Such work should attempt to identify further quasi-periodic relationships, quantify the conditions under which they arise and determine whether such states are stable over the length of the simulation.

The heuristics for categorising player populations is another area that could benefit from clarification. The current methodology categorises a population based solely upon its average scores over many of its latter generations, which can be ambiguous when examining populations between the two extremes of the categorisation scheme. Further study of the behaviour of simulated populations over many generations could develop improved heuristics that more accurately describe the overall behaviour of a simulated population, given the trends demonstrated by its average score, Forgiving Misperception probability and Punishing Misperception probability.

9.3.9 OODA Loop Tempo and Information Warfare

In two competing entities, the differences in which they iterate through their OODA loops can lead both entities to develop inaccurate perceptions of each other and their environment, due to under-sampling. Nyquist's (1928) sampling theorem suggests a minimum sampling rate, or OODA loop tempo, of at least twice an opponent's tempo to avoid under-sampling. However, given that in most environments competing entities are likely to possess similar tempi, such entities are likely to under-sample and thereby develop misperceptions.

Given this potential source of misperception, any actions that either increase an entity's OODA loop tempo or decrease its opponent's may impact the opponent's ability to collect enough information to accurately determine its opponent's actions. Such actions may encapsulate Information Warfare attacks and may affect the entire loop or its individual steps.

Further research will be required to better examine the effects that manipulating an entity's OODA loop iteration speed can have upon its decision-making processes, and also how these effects may be combined to aid or enable Information Warfare attacks. However, while possessing a faster OODA loop tempo has been documented as beneficial in the real world, the applicability of Nyquist's theory to real world examples may be

difficult to quantify, specifically in systems that do not operate with a discrete fixed tempo. Furthermore, this may be complicated by the fact that both the OODA loop model and Nyquist's sampling theory assume a discrete model of time, while the real world operates on a continuous model of time. Under such conditions, the duration of each Observation, Orientation, Decision and Action step also becomes important.

9.4 Summary Conclusion

Misperception can occur in many diverse forms and affect many different types of entities, influencing their beliefs and actions. In some cases these altered actions can benefit the affected entity or other entities in its environment and even increase the reproductive success of the affected entities, despite the assumption that such actions should be harmful. This research has explored such misperception and demonstrated evidence supporting the hypothesis that misperception can provide an evolutionary benefit, albeit in restricted conditions. Future research should aim to further investigate further situations under which this hypothesis is true, by exploring other misperception-affected scenarios.

Appendix A

Glossary for Evolutionary Foraging Simulations

Appendix A contains a glossary of terms used to discuss Akaishi and Arita's foraging simulation, as well as the simulations conducted in Chapter 6 and Chapter 7.

Action Cycle

The action cycle is the sequence of steps that each agent goes through when it is activated each turn. The cycle consists of three main stages, which encapsulate an agent's perception, decision and action behaviours. During the first stage the agent perceives the surrounding environment and updates its model of the environment. During the second stage the agent decides which location it should move to. During the third stage the agent moves itself towards its target location.

Agents

The agents in this simulation are simple foraging agents who move around their environment, gathering resources and reproducing. The agents can misperceive their local environment, which can affect their present or future behaviour.

Agent Density

Agent Density is a measure of the occupancy of the simulated environment by agents. It is the maximum percentage of the locations in the simulated environment that may contain agents and as such it constrains the quantity of agents that can exist at any

time. The maximum number of agents that exist in the simulation at any time is given by $\text{Agent Density} \times \text{World Size}^2$. The actual number of agents in the environment will fluctuate between 0 and the maximum as new agents are born and old agents die. Agent Density is a simulation parameter that is directly accessible to the experimenter.

Average Misperception Probability

The averaged misperception probability for all the agents created during the duration of the simulation or restricted to the current population of agents.

Basic Metabolic Rate

How much energy an agent consumes each turn from its internal stockpile gathered from resource nodes. An agent with zero energy dies of starvation and is removed from the environment. This parameter is directly accessible to the experimenter.

Deadlock

Deadlock describes any situation where at least one agent obstructs the planned movement of another. How an agent is affected by deadlock depends mainly on its foraging strategy.

The agents utilise a simplistic and non-realistic resolution method, which consists of waiting for obstructions to move out of their planned path. However, if multiple agents are mutually obstructing each other, the deadlock is not resolved until one agent decides to move in a different direction. This change can only be initiated by errors (such as misperception or misaction) or the selection of a new random direction to move in (as performed in Reflexive-foraging). Deadlocks can grow and shrink, since obstructed agents may in turn obstruct other agents and random errors may guide agents away from obstructions.

Environment

The simulated world occupied by the agents. It is a two-dimensional, non-toroidal square grid, the dimensions of which the experimenter controls. The environment contains a number of agents and resource nodes, also controlled by the experimenter. Only one agent may occupy a location at a time.

Foraging Behaviours

Foraging behaviours govern how agents are affected by errors (if at all) and how agents navigate the environment to access resource nodes. These foraging behaviours are set for the entire population of the simulation. The four different foraging behaviours investigated in the simulations were Misperception-affected Foraging, Misaction-affected Foraging, Reflexive-foraging and Perfect-perception Foraging.

Misaction

Misaction is the incorrect performance of an agent's intended action. In this simulation, misaction is implemented with the agent either moving in a different random direction or failing to move at all.

Misaction Probability

The chance that an agent will misact. An agent with a 0% misaction probability will never misact, while an agent with a 100% misaction probability will misact whenever it attempts to move. Agents inherit their misaction probability from one of their parents, with the possibility for this value to mutate during reproduction.

Misaction-affected Foraging

Misaction-affected foraging operates similarly to Misperception-affected foraging, except that it causes errors when the agent acts, instead of when it perceives. These agents move around and perceive their environment, creating a model of it as they go. Unaffected by misperception, the agents' model of their environment is accurate. However, whenever agents attempt to move in a direction, they may incorrectly perform their movement, instead moving in a different direction or failing to move altogether.

Misperception

Misperception is the incorrect perception of resource nodes in the environment by the agents. As in Akaishi and Arita's original simulation, misperception occurred in one of two forms, with an equal probability. One is a misperception of existence, where an agent misperceives the contents of an adjacent location as its opposite. The other is

a misperception of location, where an agent misperceives the contents of an adjacent location to exist at a random location.

Misperception Mutation Delta

The delta value that is added to a new offspring's Misperception Probability by mutation. The delta value comes from a normal distribution, whose standard deviation is a variable that can be manipulated by the experimenter to increase or decrease the severity of mutations.

Misperception Probability

The chance that an agent will misperceive whenever it perceives its environment. An agent with a 0% misperception probability will never misperceive, while an agent with a 100% misperception probability will misperceive during every perception. Agents inherit their misperception probability from one of their parents, with the possibility for this value to mutate during reproduction.

Misperception-affected Foraging

This is the standard foraging method, as initially proposed by Akaishi and Arita. Agents move around and perceive their environment, while developing an internal representation of their environment. Agents use this model to plan movements between known resource nodes. Agents may misperceive their environment, thereby storing incorrect information in their own models, which may then affect how they plan their movements.

Mutation

A mutation is a random occurrence that can alter the chromosome of new offspring. Mutation adds random variation to the gene pool of the simulated agents. Mutation occurs with a random chance whenever agents reproduce. In this simulation mutation affects the new agent's Misperception Probability or Misaction Probability, by adding a delta value that adjusts it either upward or downward by a random amount.

Mutation Probability

The mutation probability determines the probability with which the simulation mutates a newly created offspring's inherited Misperception Probability. If a mutation is to occur, then a random delta value is added to the new offspring's Misperception Probability. Otherwise the new offspring inherits its Misperception Probability directly from one of its parents.

Mutation Range

Due to the use of a normal distribution to produce the misperception mutation deltas, 95% of the mutated misperception probabilities will be within $\pm 1.96\sigma$ of the original misperception probability, where σ is the standard deviation of the Misperception Mutation Delta. The range of misperception probabilities that this calculation encompasses is the Mutation Range.

With the Mutation Range it is possible to determine whether two agents with different misperception probabilities are possibly related. If the difference between the two misperception probabilities is greater than the Mutation Range, then the two agents are probably not related. If the difference is less than the Mutation Range then the two agents may be related, although this is not certain. The Mutation Range is also used to differentiate between sub-populations of misperceiving agents, based upon their misperception probability.

Parental Investment Cost

How much energy an agent must spend in order to reproduce. This cost must be paid by both parents. This simulation parameter is directly accessible to the experimenter.

Perfect-perception Foraging

In a Perfect-perception foraging, the agents use a modified variant of Misperception-affected foraging. Perfect-perception foraging differs in that the agents of the initial simulation population have their Misperception Probability set to 0% and their Misperception Mutation Delta set to 0.0. This ensures that any agents will never misperceive and that the offspring of any agents will also have a 0% Misperception Probability.

Perfect Perceiver

A Perfect Perceiver is an agent whose Misperception Probability is 0% and therefore cannot misperceive. Perfect Perceivers entirely populate Perfect-perception foraging simulations, but they may also evolve in simulations with Misperception-affected foraging. Similarly, in simulations with Misaction-affected foraging, an agent whose misaction probability is 0% can be considered a Perfect Perceiver, since their behaviour is unaffected by random errors.

Potential Offspring

An agent's Potential Offspring is a measure of how many offspring it could parent from its surplus energy. Surplus energy is energy that was not consumed by the agent to stay alive. Potential Offspring is only calculated after an agent has died, either of old age or starvation. The formula for this value is:

$$\text{Potential Offspring} = \frac{\text{Total Resources Gathered} - (\text{Age} \times \text{Basic Metabolic Rate})}{\text{Parental Investment Cost}}$$

An agent with 0 Potential Offspring is one that could not gather enough resources to avoid starvation, let alone afford to reproduce.

Reflexive-foraging

Reflexive-foraging has the agents behave in a simple manner that is unaffected by misperception or misaction. Instead, these agents randomly move around the environment, until they perceive a resource in their immediate vicinity. Then they attempt to move into that location and consume the resource. Agents using Reflexive-foraging are unaffected by misperception. Due to the random movement of the agents, there is a large degree of variation in the actions of the agents, even though they share the same behaviour.

Resource Density

Resource Density is a measure of the food per area in the simulation environment. It is a parameter that determines the percentage of locations in the environment that will contain resource nodes. The total number of resource nodes that are created in an environment is

given by $\text{Resource Density} \times \text{World Size}^2$. The number of resources present in the environment remains constant during the simulation and is determined by the resource density and the dimensions of the environment. On initialisation these nodes are created and randomly distributed in the environment. This simulation parameter is directly accessible to the experimenter.

Resource Nodes

Resource nodes are locations in the simulated environment where agents can find and gather resources. Gathered resources are removed from the node for a fixed number of turns, until they regenerate. Gathered resources provide the agents with energy to stay alive and to reproduce.

World Size

The World Size controls the dimensions of the square grid that is the environment. This parameter is directly accessible to the experimenter.

Appendix B

Algorithms for the Evolutionary Foraging Simulation

B.1 Algorithm for Misperception-affected foraging

Algorithms 1 – 4 describe the complete behaviour of the foraging simulation and its misperceiving agents.

Algorithm 1 describes the behaviour of the evolutionary simulation that closely approximates the agent behaviour of Akaishi and Arita’s original simulation.

`RandomBoolean` generates uniformly distributed Boolean variables (True or False). `RandomProbability` generates random numbers from a uniform distribution between 0.0 and 1.0. It is used for testing whether probabilistic events, such as misperceptions or mutations, occur.

Algorithm 2 describes the initialisation process of the agent simulation, detailing how the environment, agents and resources are created and initialised for the beginning of the simulation.

Here, `RandomProbability` selects a random value between 0.0 and 0.1 from a uniform distribution. `RandomEmptyLocation` repeatedly selects random locations within the environment, returning the location of the first empty location it selects. Random locations are identified by two random uniformly distributed integers between 1 and *GridSize*.

Algorithm 3 describes the activation behaviour of a foraging agent during each turn of the simulation. Each agent follows this activation procedure once during each turn of the simulation.

Algorithm 1 Simulation behaviour for Misperception-affected foraging

```

AgentDensity ← simulation parameter
ResourceDensity ← simulation parameter
ParentalInvestmentCost ← simulation parameter
BasicMetabolicRate ← simulation parameter
RegenRate ← 1
MaximumSpeed ← 3
AgeLimit ← 5000
MetaMutationRate ← 0.05
MutationSigma ← 0.02
Turns ← 3000000
Environment, Population ← initialised (See Algorithm 2)
for Turns iterations do
  Randomise order of agents in Population ▷ Randomise agent order to avoid biases
  for all Agent in Population do
    Activate Agent (See Algorithm 3)
  end for
  for all ResourceNode in ResourceNodes do
    if ResourceNode's Delay > 0 then                                ▷ Resource is regenerating
      Delay ← Delay - 1
      if Delay = 0 then
        ResourceNode's Quantity ← 1
      end if
    end if
  end for
end for

```

RandomLocation selects a random location within the environment, using two random uniformly distributed integers between 1 and *GridSize* to produce random coordinates. RandomDirection selects a random direction from the four possible directions an Agent may move (North, East, South and West).

Algorithm 4 describes the single-point crossover operation that produces the chromosome for a new agent. Here *MisperceptionProbability* and *MutationRate* refer to those values of the new offspring, unless specifically stated otherwise. *Chromosome[MispProb]* and *Chromosome[MuteRate]* refer to an Agent's *MisperceptionProbability* and *MutationRate* respectively, explicitly noting that these variables are element's of the Agent's Chromosome.

RandomBoolean generates random Boolean values from a uniform distribution. RandomProbability generates random numbers from a uniform distribution between 0.0 and 1.0. Here it is used for testing whether mutations occur. RandomDelta generates random floating point values that are normally distributed with a $\mu = 0.0$ and $\sigma = \textit{MutationSigma} = 0.02$. These values are used to mutate an agent's *MisperceptionProbability*

Algorithm 2 Initialisation

```

Population  $\leftarrow$  Empty
PopulationSize  $\leftarrow$  0
GridSize  $\leftarrow$  20
InitialPopulationSize  $\leftarrow$  50
AgentPopulationCap  $\leftarrow$  GridSize  $\times$  GridSize  $\times$  AgentDensity
Environment  $\leftarrow$  Empty GridSize  $\times$  GridSize grid
for GridSize  $\times$  GridSize  $\times$  ResourceDensity iterations do
  NewResourceNode  $\leftarrow$  initialised
  NewResourceNode's Delay  $\leftarrow$  0
  NewResourceNode's Quantity  $\leftarrow$  1
  RandomEmptyLocation in Environment  $\leftarrow$  new NewResourceNode
  ResourceNodes  $\leftarrow$  ResourceNodes + NewResourceNode
end for
for InitialPopulationSize iterations do
  NewAgent  $\leftarrow$  initialised
  Age  $\leftarrow$  0
  Resources  $\leftarrow$  5 ▷ Resources given to initial population
  ResourceMap  $\leftarrow$  Empty
  Destination  $\leftarrow$  Empty
  MisperceptionProbability  $\leftarrow$  RandomProbability
  MutationRate  $\leftarrow$  RandomProbability
  Population  $\leftarrow$  Population + NewAgent
  PopulationSize  $\leftarrow$  PopulationSize + 1
  Location  $\leftarrow$  RandomEmptyLocation
  Location in Environment  $\leftarrow$  NewAgent
  CurrentLocation  $\leftarrow$  Location
end for

```

or its *MutationRate*. *RandomAdjacentEmptyLocation* selects a random location that is empty and adjacent to one of the new offspring's parents.

Algorithm 3 Agent Activation for each turn of the simulation

```

AdjacentCells ← observe CurrentLocation
for all Cell in AdjacentCells do
  Location ← Cell's location
  Contents ← Cell's contents                                ▷ Resource Node or Empty
  if RandomProbability < Agent's MisperceptionProbability then
    if RandomBoolean = True then
      Toggle Contents                                       ▷ Existence Misperception
    else
      Location ← RandomLocation                             ▷ Location Misperception
    end if
  end if
  ResourceMap[Location] ← Contents
end for
Destination ← closest known Resource Node from ResourceMap
Steps ← MaximumSpeed
while Steps > 0 and CurrentLocation ≠ Destination do
  Steps ← Steps - 1
  if Destination = None then                               ▷ Agent knows of no Resource Nodes
    Direction ← RandomDirection
  else
    Direction ← determine direction from Destination
  end if
  if Cell in Direction does not contain an Agent then     ▷ Movement unobstructed
    CurrentLocation ← CurrentLocation + 1 square in Direction
    Environment ← move Agent to CurrentLocation
  end if
end while
if CurrentLocation contains ResourceNode then           ▷ Attempt to consume resources
  if ResourceNode's Quantity > 0 then                   ▷ Resource node has resources to harvest
    Resources ← Resources + ResourceNode's Quantity
    ResourceNode's Quantity ← 0
    ResourceNode's Delay ← RegenRate                       ▷ Set regeneration timer
  end if
end if
Age ← Age + 1
Resources ← Resources - BasicMetabolicRate
OtherAgent ← adjacent neighbour who can reproduce with most resources
if PopulationSize < AgentPopulationCap and OtherAgent exists then
  Resources ← Resources - ParentalInvestmentCost
  OtherAgent's Resources ← OtherAgent's Resources - ParentalInvestmentCost
  NewAgent ← one point crossover of Agent and OtherAgent (See Algorithm 4)
  Population ← Population + NewAgent
  PopulationSize ← PopulationSize + 1
end if
if Age > AgeLimit or Resources < 0 then
  Population ← Population - Agent
  PopulationSize ← PopulationSize - 1
end if

```

Algorithm 4 Single-point crossover method and mutation for misperceiving agents

```

Parent1, Parent2 ← parent agents
NewOffspring ← new uninitialised agent
NewChromosome ← None
if RandomBoolean = True then
    Head ← Parent1, Tail ← Parent2
else
    Head ← Parent2, Tail ← Parent1
end if
CutPoint ← random integer between 0 and 2
if CutPoint = 0 then
    NewChromosome[MispProb] ← Tail's Chromosome[MispProb]
    NewChromosome[MuteRate] ← Tail's Chromosome[MuteRate]
else if CutPoint = 1 then
    NewChromosome[MispProb] ← Head's Chromosome[MispProb]
    NewChromosome[MuteRate] ← Tail's Chromosome[MuteRate]
else if CutPoint = 2 then
    NewChromosome[MispProb] ← Head's Chromosome[MispProb]
    NewChromosome[MuteRate] ← Head's Chromosome[MuteRate]
end if
NewOffspring's Chromosome ← NewChromosome
if RandomProbability ≤ NewOffspring's MutationRate then
    Delta ← RandomDelta
    if (MisperceptionProbability = 0.0 AND Delta < 0.0) OR
    (MisperceptionProbability = 1.0 AND Delta > 0.0) then
        Delta ← Delta * -1.0                                ▷ Prevent mutation beyond bounds
    end if
    MisperceptionProbability ← MisperceptionProbability + Delta
    if MisperceptionProbability < 0.0 then
        MisperceptionProbability ← 0.0
    else if MisperceptionProbability > 1.0 then
        MisperceptionProbability ← 1.0
    end if
end if
if RandomProbability ≤ MetaMutationRate then
    Delta ← RandomDelta
    if (MutationRate = 0.0 AND Delta < 0.0) OR (MutationRate = 1.0 AND Delta >
    0.0) then
        Delta ← Delta * -1.0                                ▷ Prevent mutation beyond bounds
    end if
    MutationRate ← MutationRate + Delta
    if MutationRate < 0.0 then
        MutationRate ← 0.0
    else if MutationRate > 1.0 then
        MutationRate ← 1.0
    end if
end if
Age ← 0
Resources ← 5
ResourceMap ← Empty, Destination ← Empty
Location ← RandomAdjacentEmptyLocation
Location in Environment ← NewOffspring                                ▷ Update the environment map
CurrentLocation ← Location

```

B.2 Algorithm for Misaction-affected foraging

Algorithm 5 describes the behaviour of a Misaction-affected foraging agent. This behaviour is very similar to that of a Misperception-affected foraging agent documented by Algorithm 3, with misaction substituted for misperception.

RandomAlternateDirection considers the Agent's current *Direction* and randomly selects between the three other directions that the agent may move. For example, if *Direction* = North, then RandomAlternateDirection would select between East, South and West.

B.3 Algorithm for Reflexive-foraging

Algorithm 6 describes the foraging behaviour of a reflexive agent. Such agents do not maintain a Resource Map and are unaffected by Misperception and Misaction. Instead their movements are a random walk until they encounter a Resource Node in their immediate vicinity.

B.4 Algorithm for Perfect-perception foraging

The Perfect-perception foraging agents fully implement the agent simulation described in Section B.1 for Misperception-affected foraging agents, but with two minor modifications to implement their specific behaviour.

The first modification affects the creation of the simulation's initial population (See Algorithm 2), whose *MisperceptionProbability*'s are initialised to 0.0, instead of a random value.

The second modification affects occurs as the simulation is initialised. *MutationSigma* is set to 0.0 and the random number generator used to produce the mutations during reproduction (See Algorithm 4) is instead initialised with $\mu = 0.0$ and $\sigma = 0.0$. Therefore, the *Delta* values that it produces will all be 0.0, ensuring that mutation cannot alter the *MisperceptionProbability*s and *MutationRates* that new offspring inherit from their parents.

With these two modifications, all agents will have *MisperceptionProbability*s of 0.0 and will therefore be unaffected by misperception.

Algorithm 5 Misaction-affected foraging Agent Activation

```

AdjacentCells ← observe CurrentLocation
for all Cell in AdjacentCells do
  Location ← Cell's location
  Contents ← Cell's contents ▷ Resource Node or Empty
  ResourceMap[Location] ← Contents
end for
Destination ← closest known Resource Node from ResourceMap
Steps ← MaximumSpeed
while Steps > 0 and CurrentLocation ≠ Destination do
  Steps ← Steps - 1
  if Destination = None then ▷ Agent knows of no Resource Nodes
    Direction ← RandomDirection
  else
    if RandomProbability < Agent's MisactionProbability then
      Direction ← RandomAlternateDirection
    else
      Direction ← determine direction from Destination
    end if
  end if
  if Cell in Direction does not contain an Agent then ▷ Movement unobstructed
    CurrentLocation ← CurrentLocation + 1 square in Direction
    Environment ← move Agent to CurrentLocation
  end if
end while
if CurrentLocation contains ResourceNode then
  if ResourceNode's Quantity > 0 then ▷ Resource node has resources to harvest
    Resources ← Resources + ResourceNode's Quantity
    ResourceNode's Quantity ← 0
    ResourceNode's Delay ← RegenRate ▷ Set regeneration timer
  end if
end if
Age ← Age + 1
Resources ← Resources - BasicMetabolicRate
OtherAgent ← test for adjacent neighbour who can reproduce
if PopulationSize < AgentPopulationCap and OtherAgent exists then
  Resources ← Resources - ParentalInvestmentCost
  OtherAgent's Resources ← OtherAgent's Resources - ParentalInvestmentCost
  NewAgent ← one point crossover of Agent and OtherAgent (See Algorithm 4)
  Population ← Population + NewAgent
  PopulationSize ← PopulationSize + 1
end if
if Age > AgeLimit or Resources < 0 then
  Population ← Population - Agent
  PopulationSize ← PopulationSize - 1
end if

```

Algorithm 6 Misaction-affected foraging Agent Activation

```

PotentialDestinations  $\leftarrow$  Empty
AdjacentCells  $\leftarrow$  observe CurrentLocation
for all Cell in AdjacentCells do
  Location  $\leftarrow$  Cell's location
  Contents  $\leftarrow$  Cell's contents ▷ Resource Node or Empty
  if Contents = Resource Node then
    PotentialDestinations  $\leftarrow$  PotentialDestinations + Location
  end if
end for
if PotentialDestinations = None then
  Destination  $\leftarrow$  None
else
  Destination  $\leftarrow$  random selection from PotentialDestinations
end if
Steps  $\leftarrow$  MaximumSpeed
while Steps > 0 and CurrentLocation  $\neq$  Destination do
  Steps  $\leftarrow$  Steps - 1
  if Destination = None then ▷ Agent perceived no adjacent Resource Nodes
    Direction  $\leftarrow$  RandomDirection
  else
    Direction  $\leftarrow$  determine direction from Destination
  end if
  if Cell in Direction does not contain an Agent then ▷ Movement unobstructed
    CurrentLocation  $\leftarrow$  CurrentLocation + 1 square in Direction
    Environment  $\leftarrow$  move Agent to CurrentLocation
  end if
end while
if CurrentLocation contains ResourceNode then
  if ResourceNode's Quantity > 0 then ▷ Resource node has resources to harvest
    Resources  $\leftarrow$  Resources + ResourceNode's Quantity
    ResourceNode's Quantity  $\leftarrow$  0
    ResourceNode's Delay  $\leftarrow$  RegenRate ▷ Set regeneration timer
  end if
end if
Age  $\leftarrow$  Age + 1
Resources  $\leftarrow$  Resources - BasicMetabolicRate
OtherAgent  $\leftarrow$  test for adjacent neighbour who can reproduce
if PopulationSize < AgentPopulationCap and OtherAgent exists then
  Resources  $\leftarrow$  Resources - ParentalInvestmentCost
  OtherAgent's Resources  $\leftarrow$  OtherAgent's Resources - ParentalInvestmentCost
  NewAgent  $\leftarrow$  one point crossover of Agent and OtherAgent (See Algorithm 4)
  Population  $\leftarrow$  Population + NewAgent
  PopulationSize  $\leftarrow$  PopulationSize + 1
end if
if Age > AgeLimit or Resources < 0 then
  Population  $\leftarrow$  Population - Agent
  PopulationSize  $\leftarrow$  PopulationSize - 1
end if

```

Appendix C

Average Misperception

Probabilities

Appendix C contains the simulation data from the evolutionary variant of Akaishi and Arita's simulation, previously graphed in Figure 6.1 in Chapter 6.

Table C.1 lists the average misperception probabilities produced by the parameter sets. It also indicates which seven parameter sets had a higher average misperception probability, indicating a potential benefit from misperception, and the labels applied to those specific parameter sets.

Parameter sets 1–7 identify instances where the average misperception probability was greater than the mutation range of 0.0392. Parameter sets 8 and 9 are included for completeness, as later tests suggested that those parameter sets also support the hypothesis of beneficial misperception.

Agent Density	Resource Density	Basic Metabolic Rate	Parental Investment Cost	Average Misp Probability	Above Mutation Range?	Label
30%	25%	0.2	25	0.0140		
30%	25%	0.15	100	0.0177		
30%	25%	0.1	50	0.0927	√	1
30%	25%	0.05	500	0.0391		8
30%	10%	0.2	25	0.0153		
30%	10%	0.15	100	0.0194		
30%	10%	0.1	50	0.0147		
30%	10%	0.05	500	0.0302		
30%	5%	0.2	25	0.0213		
30%	5%	0.15	100	0.0290		
30%	5%	0.1	50	0.0193		
30%	5%	0.05	500	0.0437	√	4
25%	15%	0.2	25	0.0122		
25%	15%	0.15	100	0.0180		
25%	15%	0.1	50	0.0119		
25%	15%	0.05	500	0.0264		
20%	10%	0.2	25	0.0149		
20%	10%	0.15	100	0.0199		
20%	10%	0.1	50	0.0136		
20%	10%	0.05	500	0.0297		
15%	10%	0.2	25	0.0160		
15%	10%	0.15	100	0.0197		
15%	10%	0.1	50	0.0173		10
15%	10%	0.05	500	0.0313		
15%	5%	0.2	25	0.0220		
15%	5%	0.15	100	0.0288		
15%	5%	0.1	50	0.0191		
15%	5%	0.05	500	0.0401	√	7
10%	5%	0.2	25	0.0233		
10%	5%	0.15	100	0.0288		
10%	5%	0.1	50	0.0202		
10%	5%	0.05	500	0.0413	√	6
5%	5%	0.2	25	0.0288		9
5%	5%	0.15	100	0.0418	√	5
5%	5%	0.1	50	0.0497	√	3
5%	5%	0.05	500	0.0559	√	2

Table C.1: Average Misperception probabilities for the simulation's parameter sets, with labelling of instances where higher misperception probabilities evolved. Those instances support the hypothesis that misperception can benefit affected entities in an evolutionary environment.

Appendix D

Average Misaction Probabilities

Appendix D contains the simulation data for Misaction-affected foraging, previously graphed in Figure 7.1 in Chapter 7.

Table D.1 lists the average misaction probabilities calculated for each of the parameter sets. It also indicates which parameter sets had an average misaction probability greater than the mutation range of 0.0392, indicating a potential benefit from misaction. The labelled parameter sets from the earlier simulations of Misperception-affected foraging show that increased probabilities of both foraging strategies tend to evolve in similar conditions.

Agent Density	Resource Density	Basic Metabolic Rate	Parental Investment Cost	Average Misaction Probability	Above Mutation Range?	Label
30%	25%	0.2	25	0.0018		
30%	25%	0.15	100	0.0039		
30%	25%	0.1	50	0.0737	√	1
30%	25%	0.05	500	0.0161		8
30%	10%	0.2	25	0.0070		
30%	10%	0.15	100	0.0094		
30%	10%	0.1	50	0.0066		
30%	10%	0.05	500	0.0170		
30%	5%	0.2	25	0.0164		
30%	5%	0.15	100	0.0760	√	11
30%	5%	0.1	50	0.0185		
30%	5%	0.05	500	0.0580	√	4
25%	15%	0.2	25	0.0041		
25%	15%	0.15	100	0.0058		
25%	15%	0.1	50	0.0041		
25%	15%	0.05	500	0.0108		
20%	10%	0.2	25	0.0073		
20%	10%	0.15	100	0.0093		
20%	10%	0.1	50	0.0068		
20%	10%	0.05	500	0.0172		
15%	10%	0.2	25	0.0069		
15%	10%	0.15	100	0.0098		
15%	10%	0.1	50	0.0083		10
15%	10%	0.05	500	0.0168		
15%	5%	0.2	25	0.0152		
15%	5%	0.15	100	0.0431	√	13
15%	5%	0.1	50	0.0186		
15%	5%	0.05	500	0.1316	√	7
10%	5%	0.2	25	0.0101		
10%	5%	0.15	100	0.0498	√	12
10%	5%	0.1	50	0.0231		
10%	5%	0.05	500	0.0731	√	6
5%	5%	0.2	25	0.0186		9
5%	5%	0.15	100	0.0567	√	5
5%	5%	0.1	50	0.0340		3
5%	5%	0.05	500	0.1020	√	2

Table D.1: Average Misaction probabilities for the simulation’s parameter sets, with labelling of instances where higher misperception probabilities were previously observed. Note that six of the nine parameter sets with Average Misaction probabilities greater than the mutation range also had Average Misperception probabilities greater than the mutation range.

Appendix E

Additional Data for Misperception and the Iterated Prisoner's Dilemma

E.1 Probabilities of Noise-Induced State Change

Appendix E.1 discusses the states changes of the Iterated Prisoner's Dilemma game, specifically when both players use the Tit for Tat strategy, exploring the possible transitions between these states due to environmental noise and the probabilities of such events.

The Iterated Prisoner's Dilemma game has four possible states, two of which have symmetrical outcomes for the players involved. This model may be considered to have three different game states — Mutual Cooperation, Mutual Defection and Alternating Cooperation and Defection. Players using the Tit for Tat strategy select strategies that keep the game in the same state. Noise affects this situation by altering a player's move and thereby changing the game's state. In a noisy environment, the likelihood that a game will change states due to noise can be calculated.

Figure E.1 displays the state transition diagram of two Tit for Tat players affected by a single source of noise, which has a 5% chance of affecting each player. Should noise occur, it may affect one player, both players or neither players. The probability of noise affecting neither player is $0.95 \times 0.95 = 0.9025$. The probability of noise affecting one player is $(0.05 \times 0.95) + (0.95 \times 0.05) = 0.095$. The probability of noise affecting both

players is $0.05 \times 0.05 = 0.0025$. Therefore, the probability that noise will initiate a state change is $0.095 + 0.0025 = 0.0975$ or 9.75%.




Figure E.1: States of the Iterated Prisoner's Dilemma game with the probabilities of noise-induced transitions. During an iterated game, Tit for Tat players will keep the game in the same state, while noise changes the state.

The probability that the iterated game will change states between turns can also be calculated if the players are randomly selecting their strategies each turn. When players use the Random strategy there is a probability of 0.5 that they will select the same strategy from the previous game and a probability of 0.5 that they will not. Therefore, the probability that the game remains in the same state is $0.5 \times 0.5 = 0.25$, while the probability that the game changes states is 0.75.

When there is sufficient environmental noise, both the noise-affected Tit for Tat players and the Random players may randomly switch between game states. However, the Random players will change game states much more frequently (75% of the time) than the noise-affected Tit for Tat players (9.75% of the time). Over a long enough time frame, the noise-affected Tit for Tat players will spend an equal number of turns in each of the four game states, with transitions occurring infrequently. The Random players also spend an approximately equal duration in each of the four game states; however, they are much more likely to change states between turns. Therefore both strategies yield similar total payoffs for the iterated game.

E.2 Algorithm for a Trickster-influenced Iterated Prisoner's Dilemma game

Appendix E.2 describes the algorithm used to implement an Iterated Prisoner's Dilemma game with a Trickster randomly influencing the first move of some games.

The noisy Iterated Prisoner's Dilemma game with a Trickster's input differs slightly from a standard IPD game. Algorithm 7 describes the behaviour of an Iterated Prisoner's Dilemma game affected by player Misperception, external noise (implemented as Misaction) and a Trickster. It also demonstrates how the provided simulation parameters affect the games. As noted previously, this game uses that standard dilemma payoffs used by Axelrod (1984) and others.

E.3 Algorithm for an evolutionary Iterated Prisoner's Dilemma game

Algorithms 8 – 11 describe the algorithms of the evolutionary Iterated Prisoner's Dilemma simulation with asymmetric misperception.

Algorithm 8 describes the overall process of the evolutionary Iterated Prisoner's Dilemma played by players affected by Forgiving and Punishing Misperception. The main simulation parameters are *Generations*, *PopulationSize*, *Matches*, *MisactionProbability* and *MutationStdDev*. The default values used for the simulations were *Generations* = 10000, *PopulationSize* = 25, *TournamentLength* = 200, *MisactionProbability* = 0.01 and *MutationStdDev* = 0.01.

Algorithm 9 describes the IPD game played by two players and demonstrates how Misperception and environmental noise (Misaction) both affect the players' implementation of the Tit for Tat strategy. The *RandomProbability* mentioned here obtains random floating-point values between 0.0 and 1.0 from a random number generator that has a uniform distribution. These values are used for testing whether misperceptions or misactions occur.

Algorithm 10 describes the crossover operation that creates new players. *RandomBoolean* is a randomly selected boolean value produced by a random number generator.

Algorithm 7 Iterated Prisoner’s Dilemma game influenced by a Trickster.

Require: *MisperceptionProbability* \leftarrow simulation parameter value
Require: *MisactionProbability* \leftarrow simulation parameter value
Require: *Trickster* \leftarrow simulation parameter value
Require: *Payoffs*[] \leftarrow game payoffs \triangleright Standard payoffs (T=5, R=3, P=1, S=0)
for *Player1* and *Player2* **do**
 Score \leftarrow 0 \triangleright Initialise Score
 History[] \leftarrow *None* \triangleright Initialise History array
 Move \leftarrow C \triangleright Tit for Tat initially cooperates
end for
for 10000 *iterations* **do**
 if *iteration* = 1 **then** \triangleright First move
 if *Trickster* = *Enabled* **then**
 RandomPlayer \leftarrow a random choice of *Player1* and *Player2*
 RandomPlayer’s Move \leftarrow D
 end if
 for *Player1* and *Player2* **do**
 if *RandomProbability* < *MisactionProbability* **then**
 Toggle *Move*
 end if
 Outcome \leftarrow player’s *Move* and opponent’s *Move*
 Score \leftarrow *Score* + *Payoffs*[*Outcome*]
 History[*iteration*] \leftarrow *Move* \triangleright Put player’s move in its history
 end for
 else
 for *Player1* and *Player2* **do**
 ObservedMove \leftarrow opponent’s *History*[*iteration* - 1] \triangleright Observe opponent’s
previous move
 if *RandomProbability* < *MisperceptionProbability* **then**
 Toggle *ObservedMove* \triangleright Misperceive
 end if
 Move \leftarrow *ObservedMove* \triangleright Tit for Tat behaviour
 if *RandomProbability* < *MisactionProbability* **then**
 Toggle *Move*
 end if
 end for
 for *Player1* and *Player2* **do**
 Outcome \leftarrow player’s *Move* and opponent’s *Move*
 Score \leftarrow *Score* + *Payoffs*[*Outcome*]
 History[*iteration*] \leftarrow *Move* \triangleright Append player’s move to its history
 end for
 end if
end for

Its values are used to determine the order in which the players’ chromosomes are considered for the crossover operation.

Algorithm 11 describes the mutation operation upon a newborn player. The floating-point values that store a player’s Forgiving and Punishing Misperception are mutated by

Algorithm 8 An evolutionary IPD tournament with asymmetric misperception.

Generations \leftarrow external simulation parameter
PopulationSize \leftarrow external simulation parameter
TournamentLength \leftarrow external simulation parameter
MisactionProbability \leftarrow external simulation parameter
MutationStdDev \leftarrow external simulation parameter
Population[] \leftarrow None
for *PopulationSize* players **do**
 NewPlayer \leftarrow a new uninitialised player
 NewPlayer's ForgivingMisperceptionProbability \leftarrow random probability (normally distributed, $\mu = 0.05, \sigma = 0.02$)
 if *NewPlayer's ForgivingMisperceptionProbability* < 0.0 **then**
 NewPlayer's ForgivingMisperceptionProbability $\leftarrow 0.0$
 end if
 NewPlayer's PunishingMisperceptionProbability \leftarrow random probability (normally distributed, $\mu = 0.05, \sigma = 0.02$)
 if *NewPlayer's PunishingMisperceptionProbability* < 0.0 **then**
 NewPlayer's PunishingMisperceptionProbability $\leftarrow 0.0$
 end if
 NewPlayer's MisactionProbability \leftarrow *MisactionProbability*
 NewPlayer's TotalScore $\leftarrow 0$
 Population \leftarrow *Population* + *NewPlayer*
end for
for *Generations* iterations **do**
 i $\leftarrow 0$
 for *i* $<$ *PopulationSize* **do**
 j $\leftarrow i$
 for *j* $<$ *PopulationSize* **do**
 if $i \neq j$ **then** \triangleright Prevent player from competing against itself
 Perform IPD game between *Population*[*i*] and *Population*[*j*] (See Algorithm 9)
 end if
 j $\leftarrow j + 1$
 end for
 i $\leftarrow i + 1$
 end for
 Best \leftarrow player with highest *TotalScore* from *Population*
 Random \leftarrow random player from *Population*
 Worst \leftarrow player with lowest *TotalScore* from *Population*
 NewPlayer \leftarrow single-point crossover between *Best* and *Random* (See Algorithm 10)
 Mutate *NewPlayer* (See Algorithm 11)
 Population \leftarrow (*Population* - *Worst* + *NewPlayer*)
 i $\leftarrow 0$
 for *i* $<$ *PopulationSize* **do**
 Population[*i*]'s *TotalScore* $\leftarrow 0$
 i $\leftarrow i + 1$
 end for
end for

Algorithm 9 IPD Game between Tit for Tat players with Forgiving and Punishing Misperception.

```

Players  $\leftarrow$  (PlayerA, PlayerB)
for all Player in Players do
  History[]  $\leftarrow$  None
  GameScore  $\leftarrow$  0
end for
for TournamentLength iterations do
  for all Player in Players do
    Opponent  $\leftarrow$  (Players - Player)
    if iteration = 1 then
      ObservedMove  $\leftarrow$  C ▷ Tit for Tat's initial Cooperation
    else
      ObservedMove  $\leftarrow$  Opponent's History[iteration - 1] ▷ Observe opponent's
      previous move
    end if
    if ObservedMove = C then
      if RandomProbability < Player's PunishingMisperceptionProbability
then
        Toggle ObservedMove ▷ Misperceive (Punishing)
      end if
    end if
    if ObservedMove = D then
      if RandomProbability < Player's ForgivingMisperceptionProbability then
        Toggle ObservedMove ▷ Misperceive (Forgiving)
      end if
    end if
    Move  $\leftarrow$  ObservedMove ▷ Tit for Tat behaviour
    if RandomProbability < Player's MisactionProbability then
      Toggle Move ▷ Misact
    end if
    Player performs Move
  end for
  for all Player in Players do
    Opponent  $\leftarrow$  (Players - Player)
    Outcome  $\leftarrow$  Player's Move and Opponent's Move
    GameScore  $\leftarrow$  GameScore + Payoffs[Outcome]
    History[iteration]  $\leftarrow$  Move ▷ Append player's move to its history
  end for
end for
for all Player in Players do
  TotalScore  $\leftarrow$  TotalScore + GameScore ▷ Update each players' TotalScore
end for

```

adding a small random delta value to each. This delta is randomly generated from a normally distributed random number generator for which $\mu = 0.0$ and $\sigma = \textit{MutationStdDev}$. A player's chromosome contains its Forgiving and Punishing Misperception Probabilities,

Algorithm 10 A single-point crossover operation to produce a new player offspring from two parent players.

```

NewPlayer  $\leftarrow$  new uninitialised player
Chromosome  $\leftarrow$  None
if RandomBoolean = True then
    Head  $\leftarrow$  Best
    Tail  $\leftarrow$  Random
else
    Head  $\leftarrow$  Random
    Tail  $\leftarrow$  Best
end if
CutPoint  $\leftarrow$  random integer between 0 and 2
if CutPoint = 0 then
    Chromosome[ForgMispProb]  $\leftarrow$  Tail's Chromosome[ForgMispProb]
    Chromosome[PunMispProb]  $\leftarrow$  Tail's Chromosome[PunMispProb]
else if CutPoint = 1 then
    Chromosome[ForgMispProb]  $\leftarrow$  Head's Chromosome[ForgMispProb]
    Chromosome[PunMispProb]  $\leftarrow$  Tail's Chromosome[PunMispProb]
else if CutPoint = 2 then
    Chromosome[ForgMispProb]  $\leftarrow$  Head's Chromosome[ForgMispProb]
    Chromosome[PunMispProb]  $\leftarrow$  Head's Chromosome[PunMispProb]
end if
NewPlayer's Chromosome  $\leftarrow$  Chromosome
Return NewPlayer

```

which are referred to here as *Chromosome*[*ForgMispProb*] and *Chromosome*[*PunMispProb*] respectively.

E.4 Comparing Forgiving and Punishing Misperception

Appendix E.4 details the results obtained from the 30 individual simulation runs in Section 8.3. These results are categorised into those simulations that show evidence of a benefit from misperception and those that do not. The population data from each of the 30 individual populations is plotted, distinctly showing both the beneficial and detrimental elements of asymmetric misperception.

The evolutionary simulation that investigated the effects of Forgiving and Punishing Misperception was executed 30 times, with the population data from each of those runs listed here (Figure E.2 to Figure E.31). These plots display the average score obtained by each generation of the population, along with the average Forgiving and Punishing Misperception probabilities of the player population. A higher score indicates that the

Algorithm 11 Mutate a new player’s chromosome by adding small values to each element of the chromosome.

```

fMispDelta ← random mutation delta (Normally-distributed,  $\mu = 0.0, \sigma =$ 
MutationStdDev)
pMispDelta ← random mutation delta (Normally-distributed,  $\mu = 0.0, \sigma =$ 
MutationStdDev)
if Chromosome[ForgMispProb] + fMispDelta > 1.0 then
    Chromosome[ForgMispProb] ← 1.0
else if Chromosome[ForgMispProb] + fMispDelta < 0.0 then
    Chromosome[ForgMispProb] ← 0.0
else
    Chromosome[ForgMispProb] ← Chromosome[ForgMispProb] + fMispDelta
end if
if Chromosome[PunMispProb] + pMispDelta > 1.0 then
    Chromosome[PunMispProb] ← 1.0
else if Chromosome[PunMispProb] + pMispDelta < 0.0 then
    Chromosome[PunMispProb] ← 0.0
else
    Chromosome[PunMispProb] ← Chromosome[PunMispProb] + pMispDelta
end if

```

population is capable of mutual Cooperation despite the noise, while lower scores indicate either random transitions between game states or the development of mutual Defection.

Forgiveness can be considered to benefit these players if the population maintains an average score that is greater than that obtained from random strategy selection (10,800) and close to that of continued mutual Cooperation (14,400). Forgiveness is not beneficial if the population evolves a high Punishing Misperception probability, which leads to average scores less than those from random strategy selection. A high Punishing Misperception probability causes the population to typically defect, resulting in an average score close to that obtained from mutual Defection (4,800).

Based on these criteria the simulations can be divided into two groups. The first case, indicating a benefit from misperception, encapsulates the runs where Forgiving Misperception evolves and helps the players to maintain mutual Cooperation for the majority of the simulation’s execution. Five simulation runs exhibited this behaviour. The second group contains instances where the average player score is approximately equal to or less than the score from random strategy selection for much of the simulation, which indicates that misperception is not beneficial. There were 25 simulation runs that exhibited this behaviour. Table E.1 clarifies how the simulation runs fit into each group.

All the populations quickly evolve Forgiving Misperception, in order to maintain a state of mutual Cooperation in the noisy environment. However, Punishing Misperception

	Count	Run numbers
Beneficial Misperception	5	1, 2, 3, 11, 22
Detrimental Misperception	25	4–10, 12–21, 23–30

Table E.1: Simulation runs where misperception is ultimately beneficial or harmful. In the majority of the simulation runs, misperception did not provide a consistent long-term benefit to the players.

often evolves to exploit this forgiveness and reduces the population’s average score to that received from random strategy selection. In some runs Punishing Misperception evolves to fairly high probabilities (greater than 50%), which greatly reduces the population’s average score. Once the population enters this state, Forgiving Misperception begins to disappear as it is only detrimental at this point. The majority of the simulations show that the populations cannot maintain Forgiving Misperception without Punishing Misperception evolving to exploit it.

In all of the simulated populations there is a direct and inverse relationship between the Average Score and the Average Punishing Misperception Probability, where the score decreases in turn with the increase of the Punishing Misperception.




Figure E.2: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 1: Misperception is beneficial.** The player population quickly evolves forgiving misperception to deal with the noise. The population’s total score declines at two points when Punishing Misperception begins to evolve in the population, before recovering.




Figure E.3: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 2: Misperception is beneficial.** Forgiving Misperception quickly evolves in the player population to counteract the effects of noise. Punishing Misperception later evolves to exploit the forgiving players, before disappearing from the population shortly before the end of the simulation.




Figure E.4: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 3: Misperception is beneficial.** Forgiving Misperception quickly evolves in the player population. Noticeable amounts of Punishing Misperception evolve to exploit the forgiving players in several places, reducing the population's total score.




Figure E.5: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 4: Misperception is detrimental.** Forgiving Misperception quickly evolves in the player population. Punishing Misperception eventually evolves to exploit the forgiving players, reducing the Forgiving Misperception probability and the total score. By the end of the simulation there is no Forgiving Misperception and the score is less than that received from random strategy selection.




Figure E.6: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 5: Misperception is detrimental.** Forgiving Misperception evolves in the player population. Punishing Misperception quickly evolves to exploit the forgiving players and it soon dominates the population. The population stabilises with a high Punishing Misperception probability and a total score less than that received from random strategy selection.




Figure E.7: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 6: Misperception is detrimental.** The population evolves Forgiving Misperception to maintain cooperation. This allows Punishing Misperception to evolve to exploit it, thereby decreasing the population's total score. The average Punishing Misperception probability declines slightly, allowing the Forgiving Misperception probability and the total score to increase slightly.




Figure E.8: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 7: Misperception is detrimental.** The player population evolves Forgiving Misperception to maintain cooperation. Punishing Misperception then evolves to exploit any forgiveness, reaching a peak and then declining. The population's total score then increases as players use forgiveness to maintain cooperation, before Punishing Misperception evolves again and reduces the total score again.




Figure E.9: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 8: Misperception is detrimental.** The player population quickly evolves Forgiving Misperception to maintain cooperation despite the noise. In two separate instances Punishing Misperception evolves in the population, decreasing the population's total score and the average Forgiving Misperception probability.




Figure E.10: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 9: Misperception is detrimental.** Forgiving Misperception quickly evolves in the player population. It is soon followed by Punishing Misperception, which gradually decreases the population's total score as it increases.




Figure E.11: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 10: Misperception is detrimental.** Forgiving Misperception evolves to enable the players to cooperate in the noisy environment. Once forgiveness is widespread, Punishing Misperception gradually evolves to exploit it. As the Punishing Misperception probability rises, the population's total score gradually declines, until the total score finally stabilises near the score received from mutual Defection.




Figure E.12: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 11: Misperception is beneficial.** The player population quickly evolves Forgiving Misperception to maintain Cooperation. Punishing Misperception evolves in several instances to exploit this forgiveness, however, it only has a small impact on the population's total score in these cases. Punishing Misperception appears to exhibit an irregular cyclic behaviour, which is not present in the other populations.




Figure E.13: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 12: Misperception is detrimental.** Forgiving Misperception is quickly evolved in the population to maintain Cooperation. It is followed by the evolution of Punishing Misperception to exploit it, which quickly increases in the population. This reduces the population’s total score and causes Forgiving Misperception to become nearly extinct, as the high Punishing Misperception probabilities make it a liability.




Figure E.14: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 13: Misperception is detrimental.** The population evolves Forgiving Misperception, which then allows Punishing Misperception to evolve in the population. Punishing Misperception reduces the population’s total score below the score that would be obtained in a noisy environment with no misperception of either type.




Figure E.15: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 14: Misperception is detrimental.** Forgiving Misperception quickly evolved to prevent noise from disrupting Cooperation between players and enabling the population to receive a total score close to that from mutual Cooperation. However, Punishing Misperception evolves to exploit the forgiveness, affecting the population's total score and Forgiving Misperception probability.




Figure E.16: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 15: Misperception is detrimental.** Forgiving Misperception evolves in the population and is soon followed by a low Punishing Misperception that exploits it. This exploitation has a small impact on the population's score, until the Punishing Misperception probability increases to further exploit the forgiveness.




Figure E.17: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 16: Misperception is detrimental.** Forgiving Misperception evolves and allows the population to maintain Cooperation. However, significant Punishing Misperception later evolves in the population, reducing the population's total score and the forgiveness of the population.




Figure E.18: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 17: Misperception is detrimental.** Forgiving Misperception quickly evolved in the population and is then followed by Punishing Misperception to exploit it. The population manages to maintain some Forgiving Misperception, along with Punishing Misperception to exploit it. The population's score at this point is typically slightly less than that received from the random selection of strategies.




Figure E.19: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 18: Misperception is detrimental.** The population quickly evolves Forgiving Misperception to maintain Cooperation. This is soon exploited by Punishing Misperception, which evolves such high probabilities that Forgiveness becomes a significant drawback and dies out in the population. The population evolves an extremely high Punishing Misperception probability, leaving it in a stable state with a total score similar to that of Mutual Defection.




Figure E.20: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 19: Misperception is detrimental.** The population evolves Forgiving Misperception to maintain Mutual Cooperation in the noisy environment. Punishing Misperception evolves to exploit it and eventually becomes significant in the population. Once the Punishing Misperception probability reaches approximately 70%, the population's total score stabilises close to that received for Mutual Defection and Forgiving Misperception disappears, as it has become a penalty.




Figure E.21: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 20: Misperception is detrimental.** Forgiving Misperception is quickly evolved, enabling the population to maintain Cooperation. Punishing Misperception eventually evolves to exploit it, however, preventing Cooperation and decreasing the population's total score. Punishing Misperception then disappears from the population in a short space of time, allowing the forgiving players to maintain Cooperation. Once again, Forgiveness provides a suitable environment for Punishing Misperception, which evolves to exploit the Forgiveness and reduces the population's total score by preventing Cooperation.




Figure E.22: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 21: Misperception is detrimental.** The population quickly evolves Forgiving Misperception to restore mutual Cooperation. Punishing Misperception then evolves to exploit this forgiveness and eventually reduces the population's average score below that obtained from random strategy selection.




Figure E.23: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 22: Misperception is beneficial.** The player population evolves Forgiving Misperception. In this run the population manages to maintain Forgiving Misperception without evolving a high Punishing Misperception probability. In this run Forgiving Misperception is not invaded by exploitative Punishing Misperception.




Figure E.24: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 23: Misperception is detrimental.** Player population initially evolves Forgiving Misperception that is clearly beneficial. However, Punishing Misperception subsequently evolves to exploit the population's forgiveness and reduces the population's average score below that received for random strategy selection.




Figure E.25: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 24: Misperception is detrimental.** The player population quickly evolves Forgiving Misperception to maintain mutual Cooperation. Punishing Misperception eventually evolves to exploit the forgiving players, reducing the population's score. Punishing Misperception then decreases in the population, allowing them to resume mutually Cooperating and thereby increase the average score. However, Punishing Misperception then evolves to exploit the forgiving players and reduces the population's average score. Misperception is sometimes beneficial in this run; however, this benefit is not consistent.




Figure E.26: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 25: Misperception is detrimental.** Forgiving Misperception quickly evolves in the population to allow players to mutually Cooperate. Punishing Misperception then evolves to exploit the forgiveness, before reducing its occurrence and allowing mutual Cooperation to increase the population's score. However, the average Punishing Misperception probability subsequently increases and gradually reduces the population's score close to that of random strategy selection.




Figure E.27: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 26: Misperception is detrimental.** Forgiving Misperception evolves in the population to allow the maintenance of mutual Cooperation. Punishing Misperception evolves to exploit the forgiveness and reduces the population's average score. Towards the end of the simulation, Punishing Misperception begins to rise and is matched exactly by a decrease in the population's average score and Forgiving Misperception probability.




Figure E.28: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 27: Misperception is detrimental.** Forgiving Misperception evolves in the population to maintain Cooperation. However, Punishing Misperception quickly evolves to exploit this forgiveness. The population ends in a final state with almost no Forgiving Misperception and near certain Punishing Misperception. This state is similar to a game where the players use the Always Defect strategy and has a similar average score.




Figure E.29: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 28: Misperception is detrimental.** The population quickly evolves Forgiving Misperception to maintain mutual Cooperation. While Punishing Misperception does evolve to exploit the forgiving players, it disappears from the population and then reappears in several cases. The population receives an average score indicative of mutual Cooperation when the population's Punishing Misperception probability is not too high.




Figure E.30: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 29: Misperception is detrimental.** The population quickly evolves Forgiving Misperception to maintain mutual Cooperation. However, Punishing Misperception invades the population towards the end of the simulation and reduces the average score below that received from random strategy selection.




Figure E.31: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Run 30: Misperception is detrimental.** The population evolves Forgiving Misperception to maintain mutual Cooperation and this state is maintained for much of the simulation. However, Punishing Misperception evolves to exploit the forgiving players and reduces the population's average score below that of random strategy selection. Forgiving Misperception is beneficial due to the Cooperation it maintains; however, it is ultimately detrimental for the population as it allows the invasion of Punishing Misperception.

E.5 Algorithm to limit Forgiving Misperception

In order to implement an upper bound on Forgiving Misperception, the existing simulation was modified to control the maximum Misperception Probability that new offspring could receive. This value was controlled by a new global simulation parameter, referred to here as *MispUpperBound*. Algorithm 12 documents these modifications to the mutation process, which modifies the process originally described in Algorithm 11.

E.6 Results from Restricting the Forgiving Misperception Probability

Appendix E.6 contains additional plots of the simulation data obtained when the Forgiving Misperception probability is restricted by an upper bound. These plots display two-dimensional and wire frame plots to complement and further clarify those in Section 8.4.2, displaying plots of the Forgiving Misperception Probability, Punishing Misperception Probability and the Average Score.

Algorithm 12 Mutate a new player’s chromosome, with an upper bound on its Forgiving Misperception Probability.

```

fMispDelta ← random mutation delta (Normally-distributed,  $\mu = 0.0, \sigma =$ 
MutationStdDev)
pMispDelta ← random mutation delta (Normally-distributed,  $\mu = 0.0, \sigma =$ 
MutationStdDev)
if Chromosome[ForgMispProb] + fMispDelta > MispUpperBound then
    Chromosome[ForgMispProb] ← MispUpperBound
else if Chromosome[ForgMispProb] + fMispDelta < 0.0 then
    Chromosome[ForgMispProb] ← 0.0
else
    Chromosome[ForgMispProb] ← Chromosome[ForgMispProb] + fMispDelta
end if
if Chromosome[PunMispProb] + pMispDelta > 1.0 then
    Chromosome[PunMispProb] ← 1.0
else if Chromosome[PunMispProb] + pMispDelta < 0.0 then
    Chromosome[PunMispProb] ← 0.0
else
    Chromosome[PunMispProb] ← Chromosome[PunMispProb] + pMispDelta
end if

```

E.6.1 Forgiving Misperception Probabilities

Figure E.32a shows a two dimensional plot of the average Forgiving Misperception probabilities, with each Misaction probability (i.e. noise) plotted separately. This perspective clearly shows the steep decline in the population’s average Forgiving Misperception once the upper bound on forgiveness reaches approximately 30% and Punishing Misperception begins to invade the populations.

Figure E.32b shows the expanded fine structure of Figure E.32a, which details the point at which Forgiving Misperception decreases in the populations. This reveals that the point at which the Forgiving Misperception probability begins to decrease occurs at a lower Forgiving Misperception Probability Upper Bound as the Misaction Probability increases.

Figure E.33 shows wire frame versions of the average Forgiving Misperception plots previously shown in Figure 8.11. The wire-frame view clearly displays the smooth increase in the average forgiving misperception probability as the upper bound on this value increases, up until the point at which punishing misperception invades the populations.




Figure E.32: The Average Forging Misperception Probabilities of the Tit for Tat players for the various misaction probabilities, plotted against the Forging Misperception Probability Upper Bound. Figure E.32a shows the forging misperception probability increasing in step with its upper bound until approximately 0.3, where it quickly declines. This is the point at which punishing misperception evolves to exploit the high levels of forgiveness in the population. When the misaction probability is 0%, the absence of environmental noise provides less selective pressure for Forging Misperception, resulting in a lower average Forging Misperception probability. Figure E.32b focuses on the expanded fine structure of the data when the Forging Misperception Probability Upper Bound ranges from 0.25 to 0.45.




Figure E.33: The Average Forgiving Misperception Probabilities plotted against the Forgiving Misperception Probability Upper Bound and the Misaction Probability, shown from two different perspectives. Note the steep decrease in the Forgiving Misperception probability near the Forgiving Misperception Probability Upper Bound of 30%, arising due to exploitation by Punishing Misperception. The misaction probability of 0% produces less selective pressure for Forgiving Misperception, thereby lowering the population's average Forgiving Misperception probability in those simulations.




Figure E.34: The Average Punishing Misperception Probabilities of the Tit for Tat players for the various misaction probabilities, plotted against the Forgiving Misperception Probability Upper Bound. Figure E.34a shows rapid increases in the punishing misperception probabilities directly correlates with the step decrease of forgiving misperception in the player population. A misaction probability of 0% provides less selective pressure for the development of Forgiving Misperception, ultimately resulting in a much lower average Punishing Misperception probability. Figure E.34b focuses on the expanded fine structure of the data when the Forgiving Misperception Probability Upper Bound ranges from 0.25 to 0.45.

E.6.2 Punishing Misperception Plots

Figure E.34a shows a two dimensional plot of the average Punishing Misperception probabilities, plotted separately for each of the Misaction probabilities. This perspective clearly shows that for each of the Misaction probabilities, Punishing Misperception is nearly zero in the populations until the Forgiving Misperception upper bound reaches approximately 30%. Once these populations have sufficiently high Forgiving Misperception probabilities, high Punishing Misperception probabilities evolve to take advantage of the forgiveness.

Figure E.34b shows the expanded fine structure of Figure E.34a, focusing on the Forgiving Misperception Probability Upper Bounds where the Punishing Misperception Probability begins to increase in the populations.

Figure E.35 shows a wire-frame plot of the average Punishing Misperception probability, which was previously displayed in Figure 8.12. The two perspectives displayed here show the rapid increase in the population's average Punishing Misperception probability once the Forgiving Misperception upper bound reaches approximately 30%.

E.6.3 Average Final Individual Score Plots

Figure E.36a shows a two dimensional plot of the average final individual scores, plotted separately for each Misaction probability. This perspective highlights the transition between high and low scores that occurs when the Forgiving Misperception upper bound reaches approximately 30%; the same point at which Forgiving Misperception decreases and Punishing Misperception increases. Before the populations evolve significant Punishing Misperception probabilities they are obtaining scores that are close to those available from mutual Cooperation and therefore much higher than those available from the random behaviour that would develop in a noisy environment with no Forgiving Misperception.

Figure E.36b shows the expanded fine structure of Figure E.36a, focusing on the point at which the scores begin to decrease, which occurs at the point when the Average Punishing Misperception probability increases and the Average Forgiving Misperception Probability decreases.

Figure E.37 shows a wire-frame plot of the average final scores, displaying the same data previously shown in Figure 8.13. The two different perspectives of this plot show how the average score steeply declines as higher Punishing Misperception probabilities become widespread among the populations. These plots also show that the average score







Figure E.35: The Average Punishing Misperception Probabilities plotted against the Forgiving Misperception Probability Upper Bound and the Misaction Probability, shown from two different perspectives. The Punishing Misperception probability is very low until the Forgiving Misperception Probability Upper Bound reaches approximately 30%, where it steeply increases. This increase matches the abrupt decrease in the Forgiving Misperception probability, since Punishing Misperception is exploiting Forgiving Misperception. The lower average Punishing Misperception probabilities obtained from a misaction probability of 0% are due to the lower average Forgiving Misperception probabilities in those populations.



(a)



(b)

Figure E.36: The Average Final Individual Score received by the Tit for Tat players for the various misaction probabilities, plotted against the Forgiving Misperception Probability Upper Bound. Figure E.36a shows that the majority of points with a Forgiving Misperception Probability Upper Bound of less than 0.3 have a score above 10,800, indicating that misperception provides a higher payoff in this instance than the payoff received by noisy behaviour. The simulations with a misaction probability of 0% have a higher score since those instances had much lower average Punishing Misperception probabilities than the simulations with higher misaction probabilities. Figure E.36b focuses on the expanded fine structure of the data when the Forgiving Misperception Probability Upper Bound ranges from 0.25 to 0.45.





Figure E.37: The Average Final Individual Score of the Tit for Tat populations plotted against the Forgiving Misperception Probability Upper Bound and the Misaction Probability, shown from two different perspectives. The decrease in the population's score occurs at the same Forgiving Misperception Probability Upper Bound as the increase in Punishing Misperception and the decrease in Forgiving Misperception, showing the effects of Punishing Misperception's exploitation on the population's average score. The simulations with a misaction probability of 0% have a distinctly higher score, since those instances had much lower average Punishing Misperception probabilities than the simulations with higher misaction probabilities.

decreases as the Misaction probability increases, indicating that while forgiveness can counteract the errors produced by misaction, these errors will still have a detrimental impact on the players' scores.

E.7 Algorithm to limit forgiveness

Appendix E.7 describes the modifications made to the original simulation (see Appendix E.3) in order to implement a limit on how often a player may forgive an opponent's Defections. Adding a forgiveness count to limit a player's Forgiving Misperception required the addition of a new variable external simulation parameter, *ForgCount*, and an internal counter to ensure that a player does not forgive more than that number of times during a match (*MyForgCount*). Algorithm 13 describes how the modified Iterated Prisoner's Dilemma game functions and is a modification of Algorithm 9.

E.8 Additional Data from Restricting Forgiveness

Appendix E.8 contains additional plots of the simulation results from player populations where a forgiveness count restricted a player's ability to forgive its opponent's Defections. These plots are intended to complement those in Section 8.5.2.




Figure E.38: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Forgiveness Count = 25, Run 6:** Stable population where Punishing Misperception evolves and disappears at several times, resulting in lower average scores while it exists in the population.

Figure E.38 shows another stable population where Forgiving and Punishing Misperception coexist. The average Forgiving Misperception probability evolves to approximately 30%, while the average Punishing Misperception probability remains near 0%, except for

Algorithm 13 IPD Game between Tit for Tat players with Forgiveness limited by a global count.

```

Players  $\leftarrow$  (PlayerA, PlayerB)
for all Player in Players do
  History[]  $\leftarrow$  None
  GameScore  $\leftarrow$  0
  MyForgCount  $\leftarrow$  0
end for
for TournamentLength iterations do
  for all Player in Players do
    Opponent  $\leftarrow$  (Players - Player)
    if iteration = 1 then
      ObservedMove  $\leftarrow$  C ▷ Tit for Tat's initial Cooperation
    else
      ObservedMove  $\leftarrow$  Opponent's History[iteration - 1] ▷ Observe opponent's
      previous move
    end if
    if ObservedMove = C then
      if RandomProbability < Player's PunishingMisperceptionProbability
then
        Toggle ObservedMove ▷ Misperceive (Punishing)
      end if
    end if
    if ObservedMove = D then
      if RandomProbability < Player's ForgivingMisperceptionProbability then
        if MyForgCount < ForgCount then
          MyForgCount  $\leftarrow$  MyForgCount + 1
          Toggle ObservedMove ▷ Misperceive (Forgiving)
        end if
      end if
    end if
    Move  $\leftarrow$  ObservedMove ▷ Tit for Tat behaviour
    if RandomProbability < Player's MisactionProbability then
      Toggle Move ▷ Misact
    end if
    Player performs Move
  end for
  for all Player in Players do
    Opponent  $\leftarrow$  (Players - Player)
    Outcome  $\leftarrow$  Player's Move and Opponent's Move
    GameScore  $\leftarrow$  GameScore + Payoffs[Outcome]
    History[iteration]  $\leftarrow$  Move ▷ Append player's move to its history
  end for
end for
for all Player in Players do
  TotalScore  $\leftarrow$  TotalScore + GameScore ▷ Update each players' TotalScore
end for

```

several occasions where it rises and then falls. These peaks correspond with decreases

in the average score and, in some cases, with decreases in the average Forgiving Misperception probability. The population receives an average score that correlates with that received from mutual Cooperation, with several durations that show slightly lower scores correlating with increases in the average Punishing Misperception probability.




Figure E.39: Average Player Score, Average Forgiving Misperception Probability and Average Punishing Misperception Probability plotted over time. **Forgiveness Count = 30, Run 6:** Stable population with somewhat cyclic Forgiving and Punishing Misperception and a somewhat regular periodicity to Punishing Misperception.

Figure E.39 shows a stable population where the average Punishing Misperception probability fluctuates significantly. These fluctuations appear to be somewhat cyclic and each causes a corresponding drop in the average score. With the forgiveness count set to 30, Punishing Misperception rises and falls several times in the population, yet it does not threaten to drive out Forgiving Misperception and only has a small effect on the population's average score. The rises in the Punishing Misperception probability all have correlating negative impacts on the population's average score.

E.9 Additional Population Data for Restricting Forgiveness

Appendix E.9 contains additional plots comparing the population data obtained using a forgiveness count to restrict a player's ability to forgive its opponent's Defections. These plots are intended to complement those in Section 8.5.2, displaying plots of the Forgiving Misperception Probability, Punishing Misperception Probability and the Average Score. Included are two-dimensional and wire frame plots of the same data, intending to further clarify the structure of the data within those plots.

E.9.1 Forgiving Misperception Plots




Figure E.40: The Average Forgiving Misperception Probabilities from the various Misaction Probabilities plotted against the Forgiveness Count. Here the Average Forgiving Misperception probability gradually decreases as the Forgiveness Count increases.

Figure E.40 shows a two dimensional plot of the average Forgiving Misperception probabilities, displaying the effects of the forgiveness count on the population's forgiving misperception. Here the Average Forgiving Misperception probability gradually decreases as the forgiveness count is increased.

Figure E.41 shows the average Forgiving Misperception probabilities of the populations when the forgiveness count is used to limit Punishing Misperception, displaying the same data previously shown in Figure 8.18. The average Forgiving Misperception probability gradually decreases as the forgiveness count increases. The increasing forgiveness count allows more opportunities for exploitation, which decreases the adaptivity of Forgiving Misperception and reduces the average Forgiving Misperception probabilities found in the populations.

E.9.2 Punishing Misperception Plots

Figure E.42 shows a two dimensional plot of the Average Punishing Misperception probabilities for the Misaction probabilities. The gradual increase of the Average Punishing





Figure E.41: The Average Forgiving Misperception Probability plotted against the Forgiveness Count and the Misaction Probability, displayed from two different perspectives. The average Forgiving Misperception Probability gradually decreases as the forgiveness count is increased.




Figure E.42: The Average Punishing Misperception Probabilities from the various Mismatch Probabilities plotted against the Forgiveness Count. Here the Average Punishing Misperception probability gradually increases as the forgiveness count increases. This shows that as the population is able to forgive more Defections, Punishing Misperception increases to exploit that forgiveness.

Misperception probabilities indicates that it becomes more advantageous within the populations as there is more forgiveness to exploit.

Figure E.43 shows the average Punishing Misperception probabilities of the simulated populations when the forgiveness count is used to limit Punishing Misperception, displaying the same data previously shown in Figure 8.19. The Punishing Misperception probability increases gradually with the forgiveness count, before rising sharply as the forgiveness count reaches 25. At this point exploiting Forgiving Misperception becomes a better strategy and the average Punishing Misperception probability found in the populations increases.





Figure E.43: The Average Punishing Misperception Probability plotted against the Forgiveness Count and the Misaction Probability, displayed from two different perspectives. There is a gradual increase in the average Punishing Misperception probability as the forgiveness count is increased, which permits more Punishing Misperception to be forgiven.

E.9.3 Average Final Individual Score Plots




Figure E.44: The Average Individual Final Score from the various Misaction Probabilities plotted against the Forgiveness Count. The scores obtained by these players gradually decrease as the forgiveness count is increased.

Figure E.44 shows a two dimensional plot of the Average Final Individual Score for the Misaction probabilities that were investigated. The scores obtained by the players in these populations gradually decline as the forgiveness count increases, due to the correlating increases in the Average Punishing Misperception probabilities.

Figure E.45 shows the average final score obtained by the simulated populations when the forgiveness count is used to limit Punishing Misperception, displaying the same data previously shown in Figure 8.20. The score increases up towards the maximum obtainable from mutual cooperation, before gradually decreasing once there is enough forgiveness to be exploited by Punishing Misperception. This decrease is much more gradual than that observed from the Forgiving Misperception probability upper bound simulations.





Figure E.45: The Average Final Individual Score plotted against the Forgiveness Count and the Misaction Probability, viewed from two different perspectives. As the forgiveness count increases, the average score increases up to the theoretical limit available for mutual cooperation, before gradually decreasing once Punishing Misperception evolves to exploit the forgiving players.

E.10 Additional Lotka-Volterra Model Fit

Appendix E.10 details another attempt to fit the predator-prey model to recorded simulation data that exhibits a cyclical relationship between Forgiving Misperception and Punishing Misperception. This comparison was performed in the same manner as that found in Section 8.6 and examines a cyclical relationship identified in Appendix E.8 (Figure E.39), in a population with a misaction probability of 0.05 and a forgiveness count of 30.




Figure E.46: Predator-prey data (from coefficients $a = 0.000620$, $b = 0.003067$ and $d = 0.016103$) plotted against the quasi-periodic average Punishing and Forgiving Misperception probabilities. The simulation results were obtained with a forgiveness count of 25. While the recorded misperception probabilities do exhibit the offset increases and decreases in Forgiving and Punishing Misperception that suggest a predator-prey relationship, the lack of periodicity in the oscillations of the two misperception probabilities produces a poor fit against the Lotka-Volterra model.

The search identified the set of coefficients of $a = 0.000620$, $b = 0.003067$ and $d = 0.016103$ as producing a predator-prey model that best fit the simulation data. This model is plotted against the simulation data in Figure E.46. While there does appear to be some correlation between the simulation data and the predator-prey model, the fit is again rather poor. A statistically significant test result shows that the predator-prey model almost certainly did not generate the recorded simulation data. The noisy Punishing Misperception probability data does not fit well with the smooth results from the Lotka-Volterra equations. A better argument against this population exhibiting a

predator-prey relationship is the lack of clear periodicity observed in either the Forgiving or Punishing Misperception. The Forgiving Misperception data fails to exhibit the periodicity of a predator-prey relationship, while the Punishing Misperception data does demonstrate distinct quasi-periodic peaks and troughs; however, again without the clear periodicity required for a conventional predator-prey relationship. As such, the population cannot be argued to operate as a predator-prey system, although it is possible that the simulation data is too noisy to allow the development or maintenance of clearly distinguishable predator-prey relationships. Such a question is better left for future research, and a deeper quantitative analysis of these results.

Vita

Publications arising from this thesis include:

- L. Brumley, C. Kopp, and K. B. Korb (2005)** Misperception, Self-Deception and Information Warfare, *in* G. Pye and M. Warren (eds), *Proceedings of the 6th Australian Information Warfare & Security Conference 2005*, School of Information Systems, Deakin University, Geelong VIC Australia, pp. 71–79.

- L. Brumley, C. Kopp, and K. B. Korb (2006)** The Orientation step of the OODA loop and Information Warfare, *in* C. Vailli and A. Woodward (eds), *Proceedings of the 7th Australian Information Warfare & Security Conference 2006*, School of Computer and Information Science, Edith Cowan University, Perth WA Australia, pp. 18–25.

- L. Brumley, C. Kopp, and K. B. Korb, (2006)** Causes and Effects of Perception Errors, *Journal of Information Warfare*, **5**(3), School of Computer and Information Science, Edith Cowan University, Perth WA Australia, pp. 41–53.

- L. Brumley, K. B. Korb, and C. Kopp, (2007)** An Evolutionary Benefit from Misperception in Foraging Behaviour, *Lecture Notes in Artificial Intelligence*, vol 4828, Springer, Germany, pp. 96–106.

- L. Brumley, K. B. Korb, and C. Kopp, (2009)** Using Misperception to Counteract Noise in the Iterated Prisoner’s Dilemma, *Lecture Notes in Artificial Intelligence*, vol 5865, Springer, Germany, pp. 53–62.

- L. Brumley, C. Kopp, and K. B. Korb, (2012)** Cutting Through the Tangled Web: An Information-Theoretical Perspective on Information Warfare, *Air Power Australia Analyses*, 9, no. 2, 2012

For these publications I completed the experimental work, simulation programming, testing, statistical analysis and initial written drafts, under the supervision of Kevin Korb and Carlo Kopp.

Permanent Address: Clayton School of Information Technology
Monash University
Australia

This thesis was typeset with $\text{\LaTeX} 2_{\epsilon}$ ¹ by the author.

¹ $\text{\LaTeX} 2_{\epsilon}$ is an extension of \LaTeX . \LaTeX is a collection of macros for \TeX . \TeX is a trademark of the American Mathematical Society. The macros used in formatting this thesis were written by Glenn Maughan and modified by Dean Thompson and David Squire of Monash University.

References

- Abramson, D., Sosic, R., Giddy, J. and Hall, B. (1995). Nimrod: A Tool for Performing Parameterised Simulations Using Distributed Workstations, *4th IEEE Symposium on High Performance Distributed Computing*, pp. 112–121.
- <http://www.csse.monash.edu.au/~davida/papers/nimrod.ps.Z> [Online; accessed 5-June-2007].
- Adami, C. (1998). *Introduction to Artificial Life*, Springer, New York.
- Akaishi, J. (2004). Unpublished Email Correspondence.
- Akaishi, J. and Arita, T. (2002a). Misperception, Communication and Diversity, in R. K. Standish, M. A. Bedau and H. A. Abbass (eds), *Proceedings of the 8th International Conference on Artificial Life*, pp. 350–357.
- Akaishi, J. and Arita, T. (2002b). Multi-agent Simulation Showing Adaptive Property of Misperception, *FIRA Robot World Congress*, pp. 74–79.
- Allan and Pease, B. (2006). *The Definitive Book of Body Language*, Bantam, New York.
- Andrews, C. (2004). Belief Systems, Information Warfare, and Counter Terrorism, *Proceedings of the 5th Australian Information Warfare & Security Conference 2004 (IWAR 2004)*, Perth, WA, pp. 92–99.
- Axceleon, Inc (2003). *EnFuzion 8.0 User Manual*.
- <http://www.csse.monash.edu.au/cluster/enfman80/book1.htm> [Online; accessed 7-March-2006].
- Axelrod, R. (1984). *Evolution of Cooperation*, Basic Books, New York.
- Axelrod, R. (1997). *The Complexity of Cooperation*, Princeton University Press, New Jersey.

- Axelrod, R. and Dion, D. (1988). The Further Evolution of Cooperation, *Science* **242**: 1385–1389.
- Bateson, G. (1972). Form, Substance and Difference, *Steps to an Ecology of Mind*, Chandler Publishing Co., New York. <http://www.rawpaint.com/library/bateson/formsubstanceanddifference.html> [Online; accessed 23-July-2009].
- Bazovsky, I. (2004). *Reliability Theory and Practice*, Dover Publications, Mineola, New York.
- Bennett, D. A. and Aggarwal, N. T. (2004). Alzheimer's Disease and Other Dementias, in W. J. Weiner and C. G. Goetz (eds), *Neurology for the Non-Neurologist*, Lippincott Williams and Wilkins, pp. 271–286.
- Bennett, P. G. (1977). Towards a theory of Hypergames, *Omega* **5**: 749–751.
- Bennett, P. G. (1980). Hypergames : Developing a Model of Conflict, *Futures* **12**: 489–507.
- Bennett, P. G. and Dando, M. R. (1979). Complex Strategic Analysis: A Hypergame Study of the Fall of France, *Journal of the Operational Research Society* **30**(1): 23–32.
- Berkowitz, B. and Hahn, R. W. (2003). Cybersecurity: Who's Watching the Store?, *Issues in Science and Technology* .
<http://www.issues.org/19.3/berkowitz.htm> [Online; accessed 14-September-2009].
- Berlekamp, E. R., Conway, J. H. and Guy, R. K. (1982). *Winning Ways, for your mathematical players*, Vol. 2, Academic Press, London.
- Berne, E. (1964). *Games People Play: The psychology of human behaviour*, Ballantine Books, New York.
- Berne, E. (1973). *What Do You Say after You Say Hello?: the psychology of human destiny*, Bantam Books, New York.
- Boisot, M. H. (1998). *Knowledge Assets : Securing Competitive Advantage in the Information Economy*, Oxford University Press, Oxford.
- Bonabeau, E. W. and Theraulaz, G. (1995). Why Do We Need Artificial Life?, in C. Langton (ed.), *Artificial Life: An Overview*, The MIT Press, Cambridge, Massachusetts, pp. 303–325.

- Borden, A. (1999). What is Information Warfare?, *Aerospace Power Chronicles, United States Air Force, Air University, Maxwell AFB* .
<http://www.airpower.maxwell.af.mil/airchronicles/cc/borden.html> [Online, accessed 17-August-2005].
- Borden, A. (2001). Unpublished correspondence between the Author and C. Kopp.
- Bose, P. (2003). *Alexander the Great's Art of Strategy : Business Lessons from the great Empire builder*, Allen and Unwin, Crows Nest, Australia.
- Boyd, J. R. (1986). Patterns of Conflict, Slideshow.
http://www.dnipogo.org/boyd/patterns_ppt.pdf [Online; accessed 22-October-2011].
- Boyd, J. R. (1987). Strategic Game of ? and ?, Slideshow.
<http://www.dnipogo.org/boyd/pdf/strategy.pdf> [Online; accessed 22-October-2011].
- Boyd, J. R. (1992). Conceptual Spiral, Slideshow.
http://pogoarchives.org/m/dni/john_boyd_compendium/conceptual-spiral-20110900.pdf [Online; accessed 22-October-2011].
- Boyd, J. R. (1996). The Essence of Winning and Losing, Slideshow.
http://pogoarchives.org/m/dni/john_boyd_compendium/essence_of_winning_losing.pdf [Online; accessed 22-October-2011].
- Brain, C. (2000). *Advanced subsidiary psychology: approaches and methods*, Nelson Thornes, Cheltenham, United Kingdom.
- Brams, S. J. (1977). Deception in 2×2 Games, *Journal of Peace Science* **2**: 171–203.
- Brodie, E. D. (1993). Differential Avoidance of Coral Snake Banded Patterns by Free-Ranging Avian Predators in Costa Rica, *Evolution* **47**(1): 227–235.
- Broom, D. (2001). Evolution of pain, *Vlaams Diergeneeskundig Tijdschrift* **70**(1): 17–21.
- Brumley, L. (2004). *HYPANT: A Hypergame Analysis Tool*, Honours Thesis, Monash University, Clayton Campus.

- Byrne, C. C. and Kurland, J. A. (2001). Self-Deception in an Evolutionary Game, *Journal of Theoretical Biology* **212**(4): 457–480.
- Caesar, J. and Handford, S. (1951). *The Conquest of Gaul*, Penguin Books Limited.
- Cammack, R., Attwood, T. K., Campbell, P. N., Parish, J. H., Smith, A. D., Stirling, J. L. and Vella, F. (2006). *Oxford Dictionary of Biochemistry and Molecular Biology*, 2nd edn, Oxford University Press.
- Carlson, E. A. (2004). *Mendel's Legacy: The Origin of Classical Genetics*, Cold Spring Harbor Laboratory Press, Cold Spring Harbor, New York.
- CERN (2006). Colt Open Source Libraries for High Performance Scientific and Technical Computing.
<http://acs.lbl.gov/software/colt/> [Online; accessed 13-January-2007].
- Chambers, J. (1988). *The Devil's Horsemen: The Mongol Invasion of Europe*, Cassell Publishers Ltd, London.
- Chartrand, G. (1977). *Introductory Graph Theory*, Dover Publications, New York.
- Chauvel, G. (2008). Dispersion in Optical Fibers, Report, Anritsu Corporation.
www.ausoptic.com/Alltopic/Download/Disp_in_Opt_Fibers_PMD_CD.pdf [Online; accessed 1-July-2012].
- Conley, H. P. (1988). A History of Camouflage: Concealment and Deception, *Technical Report A216 593*, Air War College.
- Daniels, C. B. (1974). Self-deception and interpersonal deception, *The Personalist* **55**: 244–252.
- Darwin, C. (1859). *On the Origin of Species*, John Murray, London.
- Darwin, C. (1871). *The Descent of Man, and Selection in Relation to Sex*, Penguin, London. Reprinted 2004.
- Dawkins, R. (1976). *The Selfish Gene*, Oxford University Press, Oxford.
- Dawkins, R. (1986). *The Blind Watchmaker*, Longman Scientific & Technical, Essex.

- de Jong, P. T. V. M. (2006). Age-Related Macular Degeneration, *The New England Journal of Medicine* **355**(14): 1474–1485.
- de Zavala, A. G., Cichocka, A., Eidelson, R. and Jayawickreme, N. (2009). Collective Narcissism and its Social Consequences, *Journal of Personality and Social Psychology* **97**(6): 1074–1096.
- Demos, R. (1960). Lying to Oneself, *The Journal of Philosophy* **57**(18): 588–595.
- Denning, D. E. (1999). *Information Warfare and Security*, Addison-Wesley, Boston.
- Department of Communications (1986). *Radio frequency interference handbook*, Australian Government Publishing Service, Canberra.
- DeRoos, J. B. (2002). Introduction to the Predator Prey Problem.
<http://home.messiah.edu/~deroos/CSC171/PredPrey/PPIntro.htm> [Online; accessed 6-June-2011].
- Deutsch, D. (1980). The Octave Illusion and Auditory Perceptual Integration, in J. V. Tobias and E. D. Schubert (eds), *Hearing Research and Theory*, Vol. 1, Academic Press, New York, pp. 99–142.
- Don, J. (2013). Tons of Released Drugs Taint US Water, AP news article.
<http://www.usnews.com/science/articles/2009/04/19/tons-of-released-drugs-taint-us-water> [Online; accessed 8-June-2013].
- Doran, J. (1998). Simulating Collective Misbelief, *Journal of Artificial Societies and Social Simulation* **1**(1).
<http://jasss.soc.surrey.ac.uk/1/1/3.html> [Online; accessed 1-December-2006].
- Doty, R. L., Yousem, D. M., Pham, L. T., Kreshak, A. A., Geckle, R. and Lee, W. W. (1997). Olfactory Dysfunction in Patients With Head Trauma, *Archives of Neurology* **54**(9): 1131–1140.
- Edelstein-Keshet, L. (1988). *Mathematical Models in Biology*, Random House, New York.
- Edwards, W. (1968). Conservatism in human information processing, in B. Kleinmütz (ed.), *Formal Representation of Human Judgment*, John Wiley, New York, pp. 17–52.

- Engen, R. (2008). Killing for their Country: A new look at “Killology”, *Canadian Military Journal* **9**(2): 120–128. Book Review Essay
<http://www.journal.dnd.ca/vo9/no2/doc/16-engen-eng.pdf> [Online; accessed 23-May-2012].
- Everson, T. (2007). *The Gene: A Historical Perspective*, Greenwood Guides to Great Ideas in Science, Greenwood Press, Westport, Connecticut.
- Ewen, R. B. (2003). *An Introduction to the Theories of Personality*, 6th edn, Lawrence Erlbaum Associates, Mahwah, New Jersey.
- Fang, L., Hipel, K. W. and Kilgour, D. M. (1989). Conflict models in graph form: Solution concepts and their interrelationships, *European Journal of Operations Research* **41**: 86–100.
- Fang, L., Hipel, K. W. and Kilgour, D. M. (1993). *Interactive Decision Making*, J. Wiley, New York.
- Festinger, L. (1957). *A Theory of Cognitive Dissonance*, Stanford University Press, Stanford, California.
- Fingarette, H. (1969). *Self-deception*, Routledge and Kegan Paul, London.
- Fischer, B. B. (1997). *A Cold War Conundrum: The 1983 Soviet War Scare*, Center for the Study of Intelligence, Central Intelligence Agency.
<https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/books-and-monographs/a-cold-war-conundrum/source.htm> [Online; accessed 4-April-2008].
- Fodor, J. A. (1983). *The Modularity of Mind*, The MIT Press, Cambridge, Massachusetts.
- Frankfurt, H. G. (2005). *On Bullshit*, Princeton University Press, Princeton, New Jersey.
- Fraser, N. M. and Hipel, K. W. (1984). *Conflict Analysis, Models and Resolutions*, Elsevier Science Publishing Co. Inc., New York.
- Gaddis, J. L. (2005). *The Cold War*, The Penguin Press, New York.
- Gigerenzer, G. (2002). *Reckoning with Risk: Learning to live with uncertainty*, Penguin Books, London.

- Gladwell, M. (2005). *Blink: The power of thinking without thinking*, Little, Brown and Company, New York.
- Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, Reading, Massachusetts.
- Goldstein, D. G. and Gigerenzer, G. (2002). Models of Ecological Rationality: The Recognition Heuristic, *Psychological Review* **109**(1): 75–90.
- Goldstein, J. (1999). Emergence as a Construct: History and Issues, *Emergence* **11**, pp. 49–72.
- Goldstein, J. L. and Krasner, S. D. (1984). Unfair Trade Practices: The Case for a Differential Response, *The American Economic Review* **74**(2): 282–287.
- Gould, S. J. and Lewontin, R. (1978). The Spandrels of San Marco and the Panglossian Paradigm, *Proceedings of the Royal Society of London* **205**: 581–598.
- Griffith, H. W. (1995). *Complete Guide to Symptoms, Illness and Surgery*, 3rd edn, The Berkley Publishing Group, New York.
- Grossman, D. (1995). *On Killing: The Psychological Cost of Learning to Kill in War and Society*, Little, Brown and Company, Boston.
- Grossman, P. Z. (1994). The Dilemma of Prisoners, *Journal of Conflict Resolution* **38**(1): 43–55.
- Haight, M. R. (1980). *A study of self-deception*, Harvester Press, Sussex.
- Hamilton, W. D. (1963). The Evolution of Altruistic Behavior, *The American Naturalist* **97**(896): 354–356.
- Hamilton, W. D. (1964). The Genetical Evolution of Social Behaviour, *Journal of Theoretical Biology* **7**: 1–16, 17–52.
- Hamilton, W. D. (1970). Selfish and Spiteful Behaviour in an Evolutionary Model, *Nature* **228**(5277): 1218–1220.
- Hanley, J. E., Orbell, J. and Morikawa, T. (2002). The Cost of Misperception if Deadly Conflicts: Hawk-Dove Games and Suicidal Terrorism, *Politics and the Life Sciences* **21**(1): 11–15.

- Hanlon, R. T. and Messenger, J. B. (1996). *Cephalopod Behaviour*, Cambridge University Press.
- Hanson, N. R. (1958). *Patterns of Discovery: An Inquiry into the Conceptual Foundations of Science*, Cambridge University Press, Cambridge.
- Hanson, N. R. (1970). *Perception and Discovery: An Introduction to Scientific Inquiry*, Wadsworth.
- Hardin, G. (1968). The Tragedy of the Commons, *Science* **162**: 1243–1248.
- Haswell, J. (1979). *The Intelligence and Deception of the D-Day Landings*, Batsford, London.
- Haswell, J. (1985). *The Tangled Web: The Art of Tactical and Strategic Deception*, John Goodchild Publishers, Wendover.
- Hecht, J. (2006). *Understanding Fiber Optics*, 5th edn, Pearson, Upper Saddle River, New Jersey.
- Held, G. (1997). *Understanding Data Communications*, 2nd edn, John Wiley and Sons, New York.
- Heuer, Jr, R. J. (1999). *Psychology of Intelligence Analysis*, Center for the Study of Intelligence, Central Intelligence Agency, Langley, Virginia, USA. <https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/books-and-monographs/psychology-of-intelligence-analysis/PsychofIntelNew.pdf> [Online; accessed 12-March-2007].
- Hodges, J. R. (1994). *Cognitive Assessment for Clinicians*, Oxford University Press, Oxford.
- Hofstadter, D. R. (1986). *Metamagical themas : questing for the essence of mind and pattern*, Bantam Books, New York.
- Holden, G. (1994). *Russia after the Cold War : History and the Nation in Post-Soviet Security Politics*, Westview Press, Boulder, Colorado.
- Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor.

- Holt, T. (2004). *The Deceivers : Allied military deception in the Second World War*, Scribner, New York.
- Howard, N. (1971). *Paradoxes of Rationality : Theory of Metagames and Political Behaviour*, M.I.T. Press.
- Howard, N. (1987). The Present and Future of Metagame Analysis, *European Journal of Operational Research* **32**: 1–25.
- Hughes, H. C. (2001). *Sensory Exotica*, The MIT Press, Cambridge, Massachusetts.
- Hughes-Wilson, J. (2006). *A Brief History of the Cold War*, Robinson, London.
- Hutchinson, W. and Warren, M. (2001). Principles of Information Warfare, *Journal of Information Warfare* **1**(1): 1–6.
- Islam, M. M., Pose, R. and Kopp, C. (2005). Suburban Ad-Hoc Networks in Information Warfare, in G. Pye and M. Warren (eds), *Proceedings of the 6th Australian Information Warfare & Security Conference 2005 (IWAR 2005)*, School of Information Systems, Deakin University, Geelong, Victoria, pp. 71–79.
- Janis, I. L. (1982). *Groupthink: Psychological studies of policy decisions and fiascoes*, Houghton Mifflin, Boston.
- Jauss, H. R. (1982). *Toward an Aesthetic of Reception*, University of Minnesota Press, Minneapolis. Translated from German by Timothy Bahti.
- Jenik, A. (2009). Cyberwar in Estonia and the Middle East, *Network Security* **2009**(4): 4–6.
- Jormakka, J. and Mölsä, J. V. E. (2005). Modelling Information Warfare as a Game, *Journal of Information Warfare* **4**(2): 12–25.
- Kaminski, M. M. (2004). *Games Prisoners Play: The Tragicomic Worlds of Polish Prison*, Princeton University Press, Princeton.
- Kearns, M., Littman, M. L. and Singh, S. (2001). Graphical Models for Game Theory, in J. S. Breese and D. Koller (eds), *17th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 253–260.

- Kerr, I. (2007). To Observe and Protect? How Digital Rights Management Systems Threaten Privacy and What Policymakers Should Do about It, in P. K. Yu (ed.), *Intellectual Property and Information Wealth: Issues and Practices in the Digital Age*, Vol. 1, Praeger Publishers, Westport, Connecticut, pp. 321–343.
- Kilgour, D. M., Hipel, K. W. and Fang, L. (1987). The Graph Model for Conflicts, *Automatica* **23**(1): 41–55.
- Kim, H. (1969). Ideology and Indoctrination in the Development of North Korean Education, *Asian Survey* **9**(11): 831–841.
- King, D. (1997). *The Commissar Vanishes*, Metropolitan Books, New York.
- Kolb, D. A., Rubin, I. M. and McIntyre, J. M. (1971). *Organizational Psychology: An Experiential Approach*, Prentice Hall, Englewood Cliffs, New Jersey.
- Komov, S., Korotkov, S. and Dylevski, I. (2007). Military aspects of ensuring international information security in the context of elaborating universally acknowledged principles of international law, *Disarmament Forum* (3): 35–43.
<http://www.unidir.org/files/publications/pdfs/icts-and-international-security-en-332.pdf> [Online; Accessed 1-November-2010].
- Kopp, C. (2000a). Information Warfare: A Fundamental Paradigm of Infowar, *Systems: Enterprise Computing Monthly* pp. 46–55.
<http://www.ausairpower.net/OSR-0200.html> [Online; accessed 17-August-2005].
- Kopp, C. (2000b). Moore's Law and its Implications for Information Warfare, in R. Sabilia (ed.), *Proceedings of the 3rd International Association of Old Crows (AOC) Electronic Warfare Conference, Zurich, 2000*, Association of Old Crows, Alexandria, Virginia.
- Kopp, C. (2003). Shannon, Hypergames and Information Warfare, *Journal of Information Warfare* **2**(2): 108–118.
- Kopp, C. (2005a). Classical Deception Techniques and Perception Management vs. the Four Strategies of Information Warfare, in G. Pye and M. Warren (eds), *Proceedings of the 6th Australian Information Warfare & Security Conference 2005 (IWAR 2005)*, School of Information Systems, Deakin University, Geelong, Victoria, pp. 81–89.

- Kopp, C. (2005b). The Analysis of Compound Information Warfare Strategies, in G. Pye and M. Warren (eds), *Proceedings of the 6th Australian Information Warfare & Security Conference 2005 (IWAR 2005)*, School of Information Systems, Deakin University, Geelong, Victoria, pp. 90–97.
- Kopp, C. (2006a). Considerations on Deception Techniques used in Political and Product Marketing, in C. Valli and A. Woodward (eds), *Proceedings of the 7th Australian Information Warfare & Security Conference 2006 (IWAR 2006)*, School of Computer and Information Science, Edith Cowan University, Perth, Western Australia, pp. 62–71.
- Kopp, C. (2006b). CSE 468 Information Conflict, Lecture Slides.
<http://www.csse.monash.edu.au/courseware/cse468/2006/Lectures/CSE-468-04.pdf> [Online; accessed 24-August-2006].
- Kopp, C. (2009). *NCW101: An Introduction to Network Centric Warfare*, 1st edn, Air Power Australia, Melbourne, Australia.
- Kopp, C. (2010). The Four Strategies of Information Warfare and their Applications, *IO Journal* **1**(4): 28–33.
- Kopp, C. (2011). Battlefield ISR – needs and goals, *Defence Today* **9**(2): 10–13.
- Kopp, C. and Mills, B. (2002). Information Warfare and Evolution, in W. Hutchinson (ed.), *Proceedings of the 3rd Australian Information Warfare & Security Conference 2002 (IWAR 2002)*, Edith Cowan University, Perth, Western Australia, pp. 352–360.
- Korotkov, S. (2008). Legal Aspects of Informational Operations, Conference Presentation – United Nations Institute for Disarmament Research: Information & Communication Technologies and International Security.
- Kruger, J. and Dunning, D. (1999). Unskilled and Unaware of It: How Difficulties in Recognizing One’s Own Incompetence Lead to Inflated Self-Assessments, *Journal of Personality and Social Psychology* **77**(6): 1121–1134.
- Kuehl, D. T. (1999). *Strategic Information Warfare: A Concept*, number 332 in *SDSC Working Papers Series*, Strategic and Defence Studies Centre, Australian National University, Canberra, Australia.

- Langton, C. G. (1995). *Artificial Life: An Overview*, The MIT Press, Cambridge, Massachusetts.
- Lewis, R. S. (1988). *Challenger : The Final Voyage*, Columbia University Press, New York.
- Libicki, M. C. (1995). *What is Information Warfare?*, United States Government Printing, Washington DC.
- Libicki, M. C. (2007). *Conquest in Cyberspace: National Security and Information Warfare*, Cambridge University Press, New York.
- Lombardo, M. P. (1985). Mutual Restraint in Tree Swallows: A Test of the TIT for TAT Model of Reciprocity, *Science* **227**(4692): 1363–1365.
- Lotka, A. J. (1920). Analytical Note on Certain Rhythmic Relations in Organic Systems, *Proceedings of the National Academy of Sciences of the United States of America* **6**(7): 410–415.
<http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1084562/> [Online; accessed 27-June-2011].
- Luce, R. D. and Raiffa, H. (1957). *Games and Decisions: introduction and critical survey*, Wiley, New York.
- Lynch, A. (1996). *Thought Contagion: How Belief Spreads Through Society*, Basic Books, New York.
- Majerus, M. E. N. (1998). *Melanism: Evolution in Action*, Oxford University Press, Oxford.
- Majeski, S. J. (1994). Arms Races as Iterated Prisoner's Dilemma Games, *Mathematical Social Sciences* **7**(3): 253–266.
- Martin, M. W. (1986). *Self-Deception*, University Press of Kansas, Lawrence.
- Maynard Smith, J. (1964). Group Selection and Kin Selection, *Nature* **201**(4924): 1145–1147.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*, Cambridge University Press, New York.

- Maynard Smith, J. and Price, G. R. (1973). The Logic of Animal Conflict, *Nature* **246**: 15–18.
- McCauley, M. (2003). *Stalin and Stalinism*, Pearson/Longman, Harlow.
- McGarty, C., Yzerbyt, V. Y. and Spears, R. (2002). Social, cultural and cognitive factors in stereotype formation, in C. McGarty, V. Y. Yzerbyt and R. Spears (eds), *Stereotypes as Explanations: The formation of meaningful beliefs about social groups*, Cambridge University Press, Cambridge, pp. 1–15.
- Mele, A. R. (1987). *Irrationality: An essay on Akrasia, Self-Deception and Self-Control*, Oxford University Press, New York.
- Mendel, G. (1865). Experiments in Plant Hybridization.
<http://www.esp.org/foundations/genetics/classical/gm-65.pdf> [Online; accessed 22-February-2006].
- Metcalf, R. M. (1995). Metcalfe's law: a network becomes more valuable as it reaches more users., *InfoWorld* **17**(40): 53.
- Meyer, B., Beekman, M. and Dussutour, A. (2008). Noise-induced Adaptive Decision-Making in Ant-Foraging, in M. A. et al. (ed.), *Proceedings of the The Tenth International Conference on the Simulation of Adaptive Behaviour*, Springer-Verlag, pp. 415–425.
- Milinski, M. (1987). Tit for Tat in sticklebacks and the evolution of cooperation, *Nature* **325**: 433–435.
- Miller, G. (2001). *The Mating Mind: How Sexual Choice Shaped the Evolution of Human Nature*, Anchor, New York.
- Molander, P. (1985). The Optimal Level of Generosity in a Selfish, Uncertain Environment, *The Journal of Conflict Resolution* **29**(4).
- Morgenstern, O. and von Neumann, J. (1947). *Theory of Games and Economic Behaviour*, Princeton University Press, Princeton.
- Nash, J. F. (1950). Equilibrium Points in n-Person Games, *Proceedings of the National Academy of Sciences of the United States of America* **36**(1): 48–49.

- National Commission on Terrorist Attacks Upon the United States (2004). *The 9/11 Commission Report: Final Report of the National Commission on Terrorist Attacks Upon the United States*, W. W. Norton & Company, Washington DC.
<http://www.gpo.gov/fdsys/pkg/GPO-911REPORT/pdf/GPO-911REPORT.pdf> [Online; accessed 31-May-2012].
- National Transport Safety Bureau (1982). Aircraft Accident Report : Air Florida, Inc., Boeing 737-222, N62AF, Collision with 14th Street Bridge near Washington National Airport.
- Neisser, U. (1976). *Cognition and Reality*, W. H. Freeman, San Francisco.
- Norman, D. A. (1990). *The Design of Everyday Things*, Doubleday, New York.
- Northcutt, W. (2002). *The Darwin Awards: Evolution in Action*, Penguin.
- Nyquist, H. (1928). Certain topics in telegraph transmission theory, *Transactions of the American Institute of Electrical Engineers* **47**: 617–644.
- Office of Fair Trading (2006). Research on impact of mass marketed scams: A summary of research into the impact of scams on UK consumers.
http://www.offt.gov.uk/shared_offt/reports/consumer_protection/oft883.pdf [Online; accessed 1-September-2012].
- Orwell, G. (1949). *Nineteen Eighty-Four*, Secker and Warburg, London.
- Patterson, C. (1999). *Evolution*, 2nd edn, Cornell University Press, Ithaca, New York.
- Peterson, L. L. and Davie, B. S. (2007). *Computer networks: a systems approach*, 4th edn, Morgan Kaufman.
- Piattelli-Palmarini, M. (1994). *Inevitable Illusions: How Mistakes of Reason Rule our Minds*, John Wiley and Sons, New York.
- Pitt, B. (ed.) (1994). *The Military History of World War II*, Chancellor Press.
- Poundstone, W. (1992). *Prisoner's Dilemma: John von Neumann, Game Theory, and the Puzzle of the Bomb*, Oxford University Press, Oxford.
- Price, G. R. (1970). Selection and Covariance, *Nature* **227**(5257): 520–521.

- Raiffa, H. (1968). *Decision Analysis: Introductory Lectures on Choices under Uncertainty*, Random House, New York.
- Ramachandran, V. S. (1995). Anosognosia in Parietal Lobe Syndrome, *Consciousness and Cognition* **4**: 22–51.
- Ramachandran, V. S. (1996). The Evolutionary Biology of Self-Deception, Laughter, Dreaming and Depression: Some Clues from Anosognosia, *Medical Hypotheses* **47**: 347–362.
- Ramachandran, V. S. and Hirstein, W. (1998). The Perception of Phantom Limbs. The D. O. Hebb Lecture, *Brain* **121**(9): 1603–1630.
- Ramsey, F. (1926). Truth and Probability, in R. B. Braithwaite (ed.), *The Foundations of Mathematics and other Logical Essays*, Kegan, Paul, Trench, Trubner & Co., London, pp. 156–198. 1999 Electronic Edition
<http://fitelson.org/probability/ramsey.pdf> [Online; accessed 5-October-2011].
- Rapoport, A. and Guyer, M. (1966). A Taxonomy of 2×2 Games, *General Systems* **11**.
- Rattray, G. (2001). *Strategic Warfare in Cyberspace*, The MIT Press, Cambridge, Massachusetts.
- Ray, T. S. (1992). Evolution, ecology and optimization of digital organisms, Santa Fe Institute working paper 92-08-042.
<http://www.santafe.edu/media/workingpapers/92-08-042.pdf> [Online; accessed 2-February-2008].
- Repast Development Team (2006). Repast Suite.
<http://repast.sourceforge.net/> [Online; accessed 22-April-2007].
- Reynolds, C. W. (1987). Flocks, Herds, and Schools: A Distributed Behavioral Model, *ACM SIGGRAPH Computer Graphics*, Vol. 21, pp. 25–34.
- Ridley, M. (2004). *Evolution*, 3rd edn, Blackwell Publishing, Massachusetts.
- Rios, M. J., Magalhães, S. T., Santos, L. and Jahankhani, H. (2009). The Georgia's Cyberwar, in H. Jahankhani, A. G. Hessami and F. Hsu (eds), *Global Security, Safety,*

- and Sustainability*, Vol. 45 of *Communications in Computer and Information Science*, Springer Berlin Heidelberg, pp. 35–42. http://dx.doi.org/10.1007/978-3-642-04062-7_5 [Online; accessed 19-October-2010].
- Rogers Commission (1986). Report of the Presidential Commission on the Space Shuttle Challenger Accident, *Technical report*.
<http://science.ksc.nasa.gov/shuttle/missions/51-1/docs/rogers-commission/table-of-contents.html> [Online; accessed 25-September-2006].
- Russell, S. J. and Norvig, P. (2009). *Artificial Intelligence: A modern approach*, 3rd edn, Prentice Hall.
- Ruxton, G. D. (2002). The possible fitness benefits of striped coat coloration for zebra, *Mammal Review* **32**: 237–244.
- Sartre, J. P. (1957). *Being and Nothingness*, Methuen and Co. Ltd, London.
- Schermerhorn, P. and Scheutz, M. (2009). The Impact of Communication and Memory in Hive-based Foraging Agents, *2009 IEEE Symposium on Artificial Life (ALIFE 2009) Proceedings*, pp. 29–36.
- Scheutz, M. and Schermerhorn, P. (2008). The Limited Utility of Communication in Simple Organisms, in S. Bullock, J. Noble, R. Watson and M. A. Bedau (eds), *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, MIT Press, Cambridge, MA, pp. 521–528.
- Schneier, B. (2012). Lance Armstrong and the Prisoners’ Dilemma of Doping in Professional Sports, *Wired*. <http://www.wired.com/opinion/2012/10/lance-armstrong-and-the-prisoners-dilemma-of-doping-in-professional-sports/>.
- Schwartz, W. (1994). *Information Warfare: Chaos on the Electronic Superhighway*, Thunder’s Mouth Press, New York.
- Shanghai Cooperation Organization (2009). Agreement between the Governments of the Member States of the Shanghai Cooperation Organization on Cooperation in the Field of International Information Security, Translated Accord Agreement.
http://media.npr.org/assets/news/2010/09/23/cyber_treaty.pdf [Online; Accessed 9-November-2010].

- Shannon, C. (1948). A Mathematical Theory of Communication, *Bell System Technical Journal* **27**: 379–423, 623–656.
- Shannon, C. E. (1949). Communication in the presence of noise, *Proceedings of the Institute of Radio Engineers* **37**(1): 10–21.
- Shannon, C. E. and Weaver, W. (1949). *The Mathematical Theory of Communication*, The University of Illinois Press, Urbana.
- Shubik, M. (1971). The Dollar Auction Game: a paradox in noncooperative behaviour and escalation, *Journal of Conflict Resolution* **15**: 109–111.
- Simon, H. A. (1957). *Models of Man*, J. Wiley, New York.
- Simpson, J. and Weiner, E. (eds) (1989a). *Oxford English Dictionary*, 2nd edn, Oxford University Press, New York. ‘Misperception’.
- Simpson, J. and Weiner, E. (eds) (1989b). *Oxford English Dictionary*, 2nd edn, Oxford University Press, New York. ‘Self-deception’.
- Sims, K. (1994a). Evolving 3D Morphology and Behaviour by Competition, in R. Brooks and P. Maes (eds), *Artificial Life IV Proceedings*, pp. 28–39.
- Sims, K. (1994b). Evolving Virtual Creatures, *Computer Graphics, Annual Conference Series, (SIGGRAPH '94 Proceedings)*, pp. 15–22.
- Skolnik, M. I. (1962). Introduction to radar, *Radar Handbook* p. 1990.
- Skyrms, B. (2004). *The stag hunt and the evolution of social structure*, Cambridge University Press, Cambridge, New York.
- Smillie, G. (1999). *Analogue and Digital Communication Techniques*, Butterworth-Heinemann, London.
- Sober, E. (2001). The Two Faces of Fitness, in R. Singh, D. Paul, C. Krimbas and J. Beatty (eds), *Thinking about Evolution: Historical, Philosophical, and Political Perspectives*, Cambridge University Press, pp. 309–321.
- <http://sober.philosophy.wisc.edu/selected-papers/ET-2001-TwoFacesOfFitness.pdf> [Online; accessed 12-March-2007].

- Spinney, C. (2013). Iraq Invasion Anniversary: Inside the Decider's Head, Battleland Blog, Time Magazine.
<http://nation.time.com/2013/03/22/iraq-invasion-anniversary-inside-the-decidere-head/> [Online; accessed 17-May-2013].
- Stallings, W. (2007). *Data and Computer Communications*, 8th edn, Pearson, Upper Saddle River, New Jersey.
- Starbuck, W. H. and Milliken, F. J. (1988). Challenger: Fine-Tuning the Odds until something breaks, *Journal of Management Studies* **25**(4): 319–340.
- Stauber, J. and Rampton, S. (1995). *Toxic Sludge is Good for You: Lies, Damn Lies and the Public Relations Industry*, Common Courage Press, Monroe, Maine.
- Sun Tzu (1963). *The Art of War*, Clarendon Press, Oxford. Translated and with an introduction by Samuel B. Griffith ; with a foreword by B.H. Liddell Hart.
- Sun Tzu (1981). *The Art of War*, Hodder and Stoughton, London. Translated by James Clavell.
- Sun Tzu (1993). *Sun Tzu: The new Translation*, Quill, New York. Translated by J. H. Huang.
- Sveiby, K. E. (1994). What is Information?
<http://web.archive.org/web/20080103161714/http://www.sveiby.com/Portals/0/articles/Information.html> [Online; accessed 15-December-2007].
- Symantec Corporation (2009). Symantec Report on Rogue Security Software: July 08 – June 09, *Technical report*, Symantec Corporation.
http://eval.symantec.com/mktginfo/enterprise/white_papers/b-symc_report_on_rogue_security_software_exec_summary_20326021.en-us.pdf [Online; accessed 25-July-2011].
- Takahashi, M. A., Fraser, N. M. and Hipel, K. W. (1984). A procedure for analyzing hypergames, *European Journal of Operational Research* **18**: 111–122.
- Thompson, F. (1995). Business Strategy and the Boyd Cycle, *Journal of Contingencies and Crisis Management* **3**(2): 81–90.

- Titus, R. M. (1999). *The Victimology of Fraud, Restoration for Victims of Crime Conference*, Australian Institute of Criminology, Melbourne.
www.aic.gov.au/media_library/conferences/rvc/titus.pdf [Online; accessed 21-April-2008].
- Tooby, J. and Cosmides, L. (2005). Conceptual Foundations of Evolutionary Psychology, in D. M. Buss (ed.), *The Handbook of Evolutionary Psychology*, Wiley, Hoboken, New Jersey, pp. 5–67.
- Trivers, R. (1976). Preface In: *The Selfish Gene* (Dawkins, R., ed.) , Oxford University Press, Oxford, pp. v–vii.
- Trivers, R. (2000). The Elements of a Scientific Theory of Self-Deception, *Annals of the New York Academy of Sciences* **907**: 114–131.
- Trivers, R. (2011). *Deceit and Self-Deception: Fooling Yourself the Better to Fool Others*, Penguin Books, London.
- Trivers, R. L. (2002). *Natural selection and social theory : selected papers of Robert L. Trivers*, Oxford University Press, New York.
- Trivers, R. L. and Newton, H. P. (1982). The Crash of Flight 90: Doomed by Self-deception?, *Science Digest* pp. 66,67,111.
- Tucker, A. W. (1950). A two-person dilemma, Mimeograph. Reprint, On jargon: The prisoner's dilemma, UMAP Journal 1 [1980]: 101.
- Tuckett, A. (2003). *Truth-telling in aged care : a qualitative study*, PhD thesis, Queensland University of Technology.
- Tversky, A. and Kahneman, D. (1974). Judgement under uncertainty: Heuristics and biases, *Science* **185**(4157): 1124–1131.
- Tversky, A. and Kahneman, D. (1988). Rational Choice and the Framing of Decisions, in D. E. Bell, H. Raiffa and A. Tversky (eds), *Decision Making*, Cambridge University Press, New York, pp. 167–192.
- Tye, L. (1998). *The Father of Spin: Edward L. Bernays and the Birth of Public Relations*, Crown Publishers, New York.

- Van Brunt, L. B. (1978). *Applied ECM*, 1st edn, EW Engineering, Dunn Loring.
- Van Evera, S. (2002). Why States Believe Foolish Ideas: Non-Self-Evaluation By States and Societies.
<http://dspace.mit.edu/handle/1721.1/5533> [Online; accessed 15-August-2005].
- Van Leeuwen, D. S. N. (2007). The Spandrels of Self-Deception: Prospects for a Biological Theory of a Mental Phenomenon, *Philosophical Psychology* **20**(3): 329–348.
- Vane, R. R. (2006). Advances in Hypergame Theory, in S. Parsons and P. Gmytrasiewicz (eds), *8th Workshop on Decision Theoretic and Game Theoretic Agents*, Future University, Hakodate, Japan.
- Volterra, V. (1928). Variations and Fluctuations of the Number of Individuals in Animal Species Living Together, *Conseil International pour l'Exploration de la Mer* **3**(1): 3–51. Translated by Mary Evelyn Wells.
- von Clausewitz, C. (1969). *On War*, Vol. 1, Routledge and Kegan Paul, London. Translated by Colonel J. J. Graham.
- von Neumann, J. (1951). The General and Logical Theory of Automata, in A. H. Taub (ed.), *John von Neumann: Collected Works*, Pergamon Press, New York, pp. 288–328.
- Wade, N. (1982). *The Art and Science of Visual Illusions*, Routledge and Kegan Paul, London.
- Wagner, C., Cheung, K. S. K., Ip, R. K. F. and Lee, F. S. L. (2005). Deceptive Communication in Virtual Communities, *Proceedings of the 38th Annual Hawaii International Conference on System Sciences (HICSS'05) - Track 1*, p. 21c.
- Walker, M. (1993). *The Cold War: And the Making of the Modern World*, Fourth Estate, London.
- Wang, M., Hipel, K. W. and Fraser, N. M. (1989). Solution Concepts in Hypergames, *Applied Mathematics And Computation* **34**: 147–171.
- Weik, M. H. (1995). *Communications Standard Dictionary*, 3rd edn, Birkhäuser.
- Wertsch, J. V. (1999). Revising Russian History, *Written Communication* **16**(3): 267–295.

- Widnall, S. E. and Fogelman, R. R. (1997). Cornerstones of Information Warfare, Doctrine/Policy Document, United States Air Force.
http://web.archive.org/web/20050305025521re_/www.af.mil/lib/corner.html
[Online; accessed 3-April-2006].
- Wiener, N. (1961). *Cybernetics: or control and communication in the animal and the machine*, 2nd edn, MIT Press, Cambridge, Massachusetts.
- Wilson, D. S. and Sober, E. (1994). Re-introducing Group Selection to the Human Behavioral Sciences, *Behavioral and Brain Sciences* **17**(4): 585–654.
<http://web.archive.org/web/20071102061057/http://www.bbsonline.org/documents/a/00/00/04/60/bbs00000460-00/bbs.wilson.html> [Online; accessed 5-June-2011].
- Witkop, Jr, C. J., Quevedo, Jr, W. C., Fitzpatrick, T. B. and King, R. A. (1989). *The metabolic basis of inherited disease*, Vol. 2, McGraw-Hill, chapter 119 – Albinism, pp. 2905–2947.
- Wu, J. and Axelrod, R. (1995). How to Cope with Noise in the Iterated Prisoner's Dilemma, *The Journal of Conflict Resolution* **39**(1): 183–189.
- Wynne-Edwards, V. C. (1962). *Animal dispersion in relation to social behaviour*, Oliver and Boyd, Edinburgh.
- Zahavi, A. (1975). Mate Selection – A selection for a Handicap, *Journal of Theoretical Biology* **53**: 205–214.
- Zumdahl, S. S. (2007). *Chemical Principles*, Cengage Learning.