

Suicide as an Evolutionarily Stable Strategy

Steven Mascaro, Kevin B. Korb and Ann E. Nicholson

School of Computer Science

and Software Engineering

Monash University

Clayton, Victoria 3800, Australia

E-mail: {stevenm, korb, annn}@csse.monash.edu.au

Abstract

Research into ethics has traditionally been the domain of philosophers and theologians, pursuing their investigations a priori, since social experimentation is not a viable option. After Axelrod's work on the evolution of behaviour in the iterated prisoner's dilemma (Axelrod 1984), artificial life (a-life) methods have been applied to social simulation. Here we use an a-life simulation to pursue experiments with ethics. We use a utilitarian model for assessing what is ethical, as it offers a direct and computationally clear means of measuring the ethical value of actions, based on the utility of outcomes. We investigate the particular action of altruistic suicide fostering the survival of others. Our experimental simulation results demonstrate that suicide can be an evolutionarily stable strategy (ESS).

Keywords: Experimental ethics, altruism, altruistic suicide, artificial life, social simulation, evolutionarily stable strategies.

1. Introduction

Since its inception, research into ethics has been the domain of philosophers and theologians, pursuing their investigations for the most part *a priori*, without benefit of experimental interventions in society. Since experimental manipulations of society just to investigate the ethical value or disvalue of some behaviour would themselves almost certainly be unethical, this is hardly surprising. Nevertheless, there is an important, and ancient, philosophical tradition which would support at least the relevance of such experimentation, namely utilitarianism. The absolutist ethics of most religions dictate that the ethical virtue, if any, of an act inheres in the act itself, and so ethical experimentation would not only be unethical it would also be pointless: the experimental outcomes come after the act in question, and so are irrelevant to determining its ethical value. Utilitarianism, by contrast, asserts that the moral worth of an action is measured by the utility of its outcomes; that is, an act is good if it leads to states of high positive utility and bad if it leads to states of high negative utility. Adopting such a view, we can pursue an experimental program to determine the ethical value or disvalue of such actions as euthanasia, abortion, rape and racial discrimination. Of course, experimentation in human society would be unlikely to be tolerated; however, the recent emergence of artificial life (a-life) simulation methods allow us to pursue such a program ethically.¹

The aim of our project, then, is to explore the consequences of actions typically construed as ethical or unethical under various circumstances in an a-life simulation environment. As our a-life populations *evolve*, our project is equally aimed at exploring the evolvability of ethical and unethical actions, and also at whether various social actions (e.g., altruistic actions) present in the initial populations are lost due to selection pressure or contrariwise maintained in a stable state at some level above the random mutation rate — i.e., are evolutionarily stable strategies (ESS).

In this paper we describe both our simulation environment and our first experiment, the evolutionary and ethical characteristics of suicide fostering the survival of others, i.e., altruistic suicide. We also present results obtained in a modified form

¹Unless, of course, one is inclined to take *very* seriously Chris Langton's claim that artificial life forms are truly alive. But we think such claims are (so far) unjustified.

of this experiment, in which suicide is modelled after the process of apoptosis, or programmed cell death, in yeast. These experiments illustrate environments where suicide is not only clearly altruistic, but also an evolutionarily stable strategy: the population retains altruistic suicide at a rate much larger than can be explained by the (re-) introduction of behaviours by random mutation.

2. Background

The ideas behind this research project stem in part from the pioneering work of Axelrod (Axelrod 1984) on the evolution of behaviour in the iterated prisoner’s dilemma. Since then, a large variety of “social simulation” studies have been pursued. Researchers are using a-life to investigate, for example, social communication networks (Epstein and Axtell 1996), economics (Gilbert and Conte 1995; Holland and Miller 1991) and disease transmission (Epstein and Axtell 1996). The extension of social simulation techniques to the study of specifically ethical questions seems a very natural one. Such an extension does require the introduction of some way of measuring the ethical value of the actions under study, in order to interpret the results. The more orthodox, absolutist ethical systems are not very promising for this purpose; it is at best unclear how to rank the ethical value of actions all of which violate some qualitative principles of behaviour. Fortunately, utilitarianism offers a direct and computationally clear means of measuring the ethical value of actions. Utilitarianism is an important theory in the western philosophical tradition, commonly associated with Bentham (Bentham 1789), but originating at least with Plato’s *Republic*. Such a use of utilitarianism is in principle a direct extension of the approach to artificial intelligence that views its subjects as decision-theoretic agents, as advocated, for example, by (Russell and Norvig 1995).

Given that agents have utilities — numeric values assigned to the situations which the agents are in (the experiences they have) which reflect their preferences² — utilitarianism asserts (in one of its forms) that the ethical action in any situation is the one that maximizes the total utility over all the agents in a population. The value of an action a for an individual following this approach is

$$v(a) = \sum_i u_i(o) \tag{1}$$

where $v(a)$ is the moral value of action a and $u_i(o)$ is the utility of outcome state o (resulting from action a) for individual i .³

Since the moral worth of an action can only be judged properly according to the known and knowable circumstances surrounding the act (i.e., the agent cannot be expected to be prescient of the outcomes), the *actual* outcomes cannot truly be used to assess moral worth when deciding what action to take. Instead, the outcomes reasonably anticipated must serve; in other words, the moral worth of an action is more properly asserted to be the sum of its *expected* utilities over all agents in the environment:

$$v(a) = \sum_j \sum_i u_i(o_j)p(o_j|a) \tag{2}$$

where o_j is an outcome state and $p(o_j|a)$ is its probability given the action a .

A standard approximation, sanctioned by the Law of Large Numbers, is the substitution of long-term frequencies for probabilities. In an evolutionary setting it is fair to say that the repeated exposure of a type of action to various kinds of environments will produce a large number of tests of the properties of the action — with the result that the tendency of that type of action to induce overall positive or negative utilities may be assessed relative to those environments, while additionally the reproductive fitness of that type of action may also be assessed. By examining the tendency of an action to produce either positive or negative average utilities under various circumstances, an evolutionary experimental ethics may find answers to moral questions which no amount of purely analytic ratiocination may yield.

3. Evolutionary Ethics

“Evolutionary ethics” as a philosophical thesis avers that ethical values are to be found in what promotes or impedes survival and reproduction. Here we introduce quite a different meaning: evolutionary ethics as the empirical exploration of utilitarian value in an evolving population. In any particular environment, some actions will be evolvable without being ethical (e.g., rape in many environments), and others will be ethical without being evolvable (e.g., mass altruistic suicide of a

²See (von Neumann and O.Morgenstern 1947) for the very general conditions in which numeric utility functions may be constructed from preference orderings.

³The outcomes are, of course, future outcomes dependent upon the action selected. Therefore, it would be appropriate to take into account the time discount factor for valuing future events. For simplicity, we ignore that issue here.

| World X | ETHICAL | UNETHICAL |
|----------------|-----------------------|---------------------|
| EVOLVABLE | Altruistic Suicide | Rape |
| UNEVOLVABLE | Mass Suicide | (Widespread) Incest |

Table 1. Actions mapped by their ethical value and evolvability in some possible world X

kin group). Thus, we can imagine a 3-dimensional lattice for classifying behaviours according to their evolvability and ethical value relative to the environment. Table 1 shows a slice of this lattice, with ethical nature on one axis and evolvability on the other, the slice capturing a hypothetical “World X”. For example, according to this table, World X is one where altruistic suicide is both ethical and evolvable, whereas rape could evolve but would be unethical.

By a behaviour being evolvable we mean that the behaviour is an ESS — an *evolutionarily stable strategy*, as defined by John Maynard Smith: “An ESS is a strategy such that, if all the members of a population adopt it, then no mutant strategy could invade the population under the influence of natural selection” (Smith 1982, p. 10). A prototype of an ESS is TIT-FOR-TAT in the iterated prisoner’s dilemma, which resists invasion by either the SUCKER or DEFECTOR strategies (Axelrod 1984). In our experiments, strategies are always mixed (i.e., probabilistic); therefore, the criterion for an ESS we apply is to test whether the behaviour resists disappearing in the presence of alternative behaviours. In particular, we consider here whether the probability of altruistic suicide drops to the background mutation rate (below which it cannot reliably fall) when the option of not committing suicide is available (as it always is).

For many behaviours it will be possible to discover environments (worlds) for which they are variously ethical or unethical, and others again for which they are variously evolvable or unevolvable. What is of interest is not, then, simply inventing worlds where behaviours often thought to be unethical turn out to be ethical; rather, it is investigating boundaries separating those environments in which a behaviour is ethical vs. unethical, and again evolvable vs. unevolvable, and determining what factors explain the differences, that is of real interest. If, then, the demarcating factors correspond to factors in our world, there is the real possibility that evolutionary ethical simulations will shed light on issues of importance to us.

4. The Ethics Simulation Environment


Our ethics simulation environment shares with genetic algorithms the concepts of a population composed of agents who have genotypes which guide their behaviour (Holland 1975). Genetic operators during reproduction mix the genes of reproducers (crossover) and modify the genetic determinants of behaviour randomly in mutations. As is standard in a-life (but not genetic algorithms), the environment in which the agents act include not only externally determined features (such as food), but also the agents themselves and their behaviour. Agents may eat, attempt to mate, move around their two-dimensional world, and in the end will die from poor health, old age, or suicide (see Figure 1).

There are two units of time in the simulation, one for the purposes of simulation and the other for statistics collection. Every agent gets a chance to perform one action during a *cycle*; the order in which agents perform actions is randomly selected for each cycle. The second time unit is the *epoch*, at the end of which assorted statistics are collected; the epoch extends over a fixed number of cycles, specifically five in this study.

The Board. The board consists of a two-dimensional grid of cells that can contain any number of items in the simulation, such as food or agents.

Food. Agents can eat food that is distributed uniformly around the board. The food contains nutrition that the agents can absorb to increase their health, which enables them to survive and reproduce. Food is generated each cycle and has a short lifetime relative to agents. In this study, the food is distributed across cycles according to either a sinusoidal or constant distribution function. Sinusoidal distribution provides for periods of high and low availability of food, mimicking seasonal fluctuations or, perhaps, periodic droughts.

Agents. Each agent has a chromosome (which determines its behavioural tendencies), a location and orientation, an age, a level of health and a randomly selected age of “natural” death (assuming it does not die of “non-natural” causes, such as



Actions

- Move forward
- Turn
- Rest
- Eat
- Mate
- Suicide

Observations

- Number of Agents
- Amount of Food
- Self-Age
- Self-Health

Figure 1. Pictorial representation of part of the simulation environment.

lack of food, in the meantime). The chromosomes have “heads” and “bodies”. The head contains general information for the agent, including mutation rates for other values in the chromosome. During mating, these mutation rates are themselves modified in offspring by global meta- mutation rates.

The chromosome body is made up of a list of production rules (condition-action pairs). The size of the list varies from individual to individual. The rules are examined during each simulation cycle, with the first rule having a condition that matches the agent’s environment observation used to select the agent’s action. The action part of a rule consists of a probability vector across action types; i.e., the action taken is determined probabilistically, in a way that depends upon the environmental conditions which the rule is matching. The values that are stored with the agent are not true probabilities — they are free to exceed 1 and fall below 0 through mutation. Before an action is chosen, these pseudo-probabilities are converted to genuine probabilities such that all negative values are treated as 0, and the remaining values are normalised so as to sum to 1. If no rule is matched during a cycle, a default action probability vector is used; that vector is in the chromosome’s head.

Rules are conditioned on the agent’s observations of the environment. The conditions are patterns that are matched against the environment, allowing for values that are greater than, less than or equal to some value specified in the condition, or possibly accepting any environmental value for a particular measurement. If all conditions in a rule are met, then its action probability vector is used to select the agent’s action. The actions currently available to agents are: moving forward, turning, resting, eating, mating and suiciding. An agent will rest for one cycle when its attempt to perform another action fails, such as attempting to eat when there is no food available.

Health and Utility. An agent’s initial health is given to it by its parents; that is, each parent donates a portion of its own health to give its offspring a start in life, currently 50 health units from each parent. Subsequently, an agent’s health depends upon events. Eating nutritious food results in a direct positive health effect; movement costs a unit of health. Each action also has a utility associated with it, given the circumstances. Thus, successful mating has a positive utility, as does eating nutritious food. Utilities flow on to health, and so are simultaneously indirect health determiners. Utility and health units are kept on the same scale, so that one negative utile (utility unit) decrements health one unit. Hence, the health of agent i at cycle c in the simulation is given by:

$$\text{Health}_{ic} = \text{initial health}_i + \sum_{t=0}^c (\text{utility}_{it} + \text{health effect}_{it}) \quad (3)$$

Observations. An agent makes four observations of its environment and itself. The health and age of the agent make up the internal observations; the external observations consist of the number of agents on the board and the amount of food available.

Eating. An agent may consume any food available within either their own cell or any neighboring cells. The act results in the removal of the food from the board and, on the current implementation, an increase in the agent’s health consisting of a

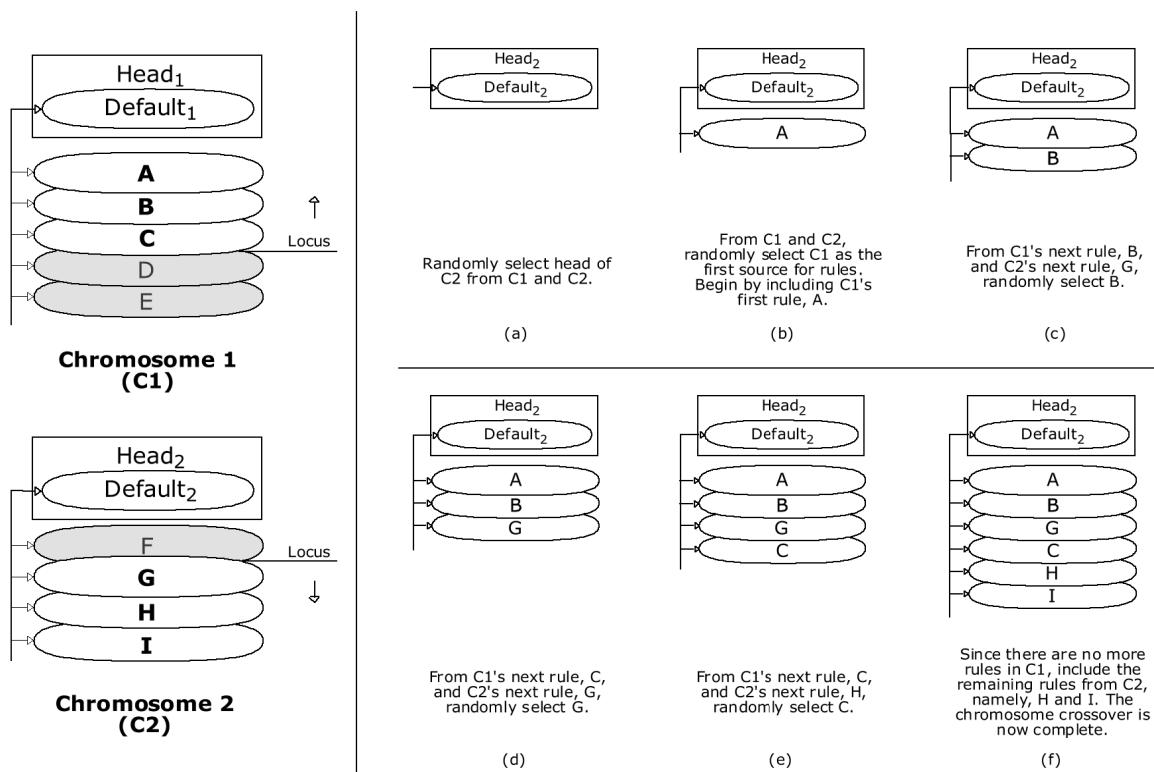


Figure 2. Diagram showing the steps involved in one possible reproduction

positive utility and a large positive health effect (70 units of health).

Movement. Agents can rotate in their cell, turning to face any of the 8 neighboring cells. As a separate action, they may also choose to move forward 1 cell. If they encounter the edge of the board, the action fails, forcing the agent to rest for that turn.

Reproduction. Since parents must donate 50 units of health to their children, there are two preconditions to successful mating: (1) there must be a mate available in the current or a neighbouring cell; and (2) both potential parents must have sufficient health available. This required health level is actually set above 100, at 200 units, implying some “maturation” period for infants to reach reproductive health and not allowing weak individuals to reproduce. If mating is successful, then an offspring is produced whose chromosome is a combination of its parents’ chromosomes.

The offspring’s chromosome head is taken randomly from one parent, with the mutation rates in the head mutated according to the meta-mutation rates. The chromosomal body is a uniform crossover of the parent bodies, subject to mutation. Since the parents’ body lengths can be different, this uniform crossover is somewhat different from the usual genetic algorithm crossover. A locus, as if for a single point crossover, is uniformly randomly selected from each parent chromosome. The offspring chromosome contains as many rules as does the first parent from its first rule to its locus plus as many rules as the second parent from its locus plus one to the end of its chromosome. This allows the complexity (chromosomal length) of offspring to itself evolve. The actual rules composing the child chromosome are taken with equal probability from either parent at the corresponding site, unless one parent’s chromosome is too short, in which case the remaining rules are taken from the longer parent.

Figure 2 illustrates the uniform crossover of reproduction for two parents’ possible chromosomes, $C1$ and $C2$, omitting mutations. Initially, a chromosome head is randomly selected from one of the parents (Figure 2a), which is followed by the body crossover. Loci are selected on both chromosome bodies so that only a portion of the rules from each parent are included. For $C1$, the locus lies between rules C and D , the child’s chromosome receiving rules A through C from this parent; from $C2$, the child’s chromosome receives rules G through I — altogether, a total of 6 rules.




Figure 3. (a) Proportion of suicides to all actions in simulations with sinusoidal (solid) and constant (dashed) food and (b) Average populations for sinusoidal (solid) and constant (dashed) food simulations.

The crossover proceeds by randomly alternating between the two parents’ chromosomes, rules taken in sequence with respect to their original ordering in the parent. In Figure 2b, $C1$ is randomly selected to contribute A , its first available rule; not having chosen G , the first rule in $C2$, it remains for the next random selection. Rules B and G , being the next rules in $C1$ and $C2$ respectively, will vie for the next position. B is randomly selected (Figure 2c), advancing the next available rule in $C1$ to C . In a similar fashion, the next randomly selected rule for inclusion becomes G (chosen from C and G), which is followed by C (from C and H). At this point, the rules in $C1$ have been exhausted and the residual rules in $C2$ are appended to the offspring chromosome, bringing the crossover to completion (Figure 2f).

Suicide. An agent may choose to suicide. Suicide causes an agent’s health to drop to 0.

5. Experimental Results

The behaviour we have chosen to investigate in the first instance is altruistic suicide. In these ethical simulations, an altruistic act may be considered to be any act which reduces the expected fitness of the individual, but is beneficial to the expected fitness of the remaining population. An alternative interpretation, justified on utilitarian grounds, is that an altruistic act is one which reduces one’s own expected future utility while enhancing the global expected future utility of the remaining population. Here we employ the latter interpretation. We assess the altruism of an action by running the simulation multiple times, recording the population’s utility with and without the action being available. In this experiment, the expected future utility is always non-negative. Hence, suicide is (almost) always detrimental to the agent. It is altruistic, therefore, whenever it enhances the utility of the population remaining.

5.1. The Evolutionary Stability of Suicide

Figure 3a shows the proportion of suicides to all actions performed each epoch over an average of 39 simulations when a constant amount of food is generated per cycle (dashed line) and again when food is generated over time according to a sine wave with an extreme amplitude (solid line) over an average of 44 simulations. Figure 3b shows the average population sizes in both of these simulation types. In all simulations the proportion of suicides performed by the agents is very high in the beginning, because the initial agents are generated randomly; suicide begins as approximately equiprobable amongst all of the actions available. In both simulations, these proportions both drop off sharply to begin with, as would be expected. Soon, however, the suicides in the sinusoidal simulations become clearly cyclical. Recalling that the graph shows suicide as a proportion of all actions, the presence of this cyclicity indicates that the agents have evolved this behaviour in response to the seasonal nature of the environment. Contrasting this to the constant food simulation, we find no cyclicity. We confirmed this by performing a significance test on the difference between the peaks and troughs of the suicides of the sinusoidal simulations ($t(22) = 18.5, p < 0.01$).

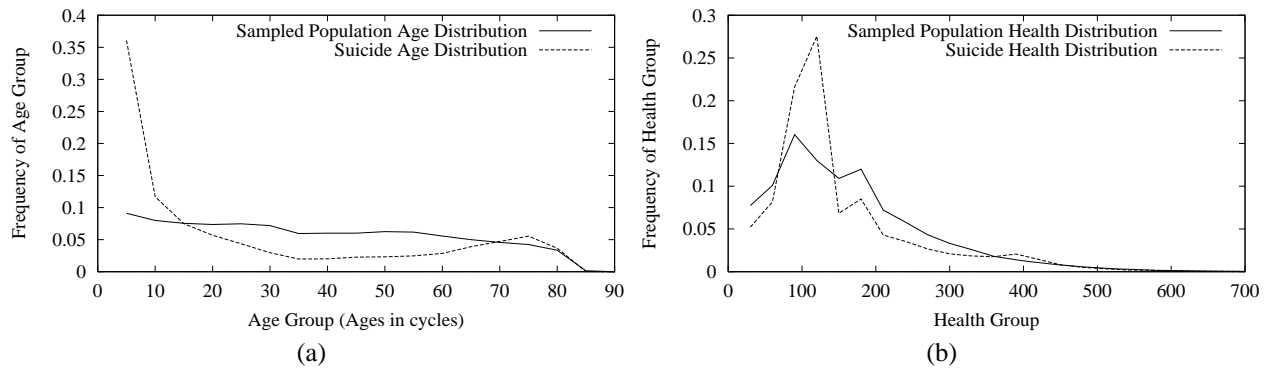


Figure 4. (a) The distribution of ages and (b) the distribution of health amongst a sample of 2000 agents in the general population and amongst 2000 of those that suicide over roughly the last 100 epochs.

The graphs in Figure 3 demonstrate that suicide, given cyclical food distribution, is evolutionarily stable. It is clearly above the background mutation rate. We can establish this even though we cannot measure the background mutation rate directly (since the mutations of the action probabilities are governed by multiple mutation rates), by comparison of the suicide rate with and without cyclical food distribution (which itself cannot be *less* than the background mutation rate).⁴

5.2. Explaining Suicide

The evolution of altruistic behaviour has long been of interest to biologists and sociobiologists (cf. Wilson’s *Sociobiology* 1975). One accepted explanation of the possibility of such evolution is that it occurs via kin selection: the altruistic act reduces the probability of the actor reproducing, but enhances the probability of its kin reproducing, thus tending to preserve or increase the presence of the gene(s) supporting the altruistic action (Hamilton 1964). Actions studied by biologists have been either less dramatic than suicide (e.g., donating food) or have been directed at close relatives (e.g., sacrificing one’s life for one’s children). Nevertheless, kin selection seems to be the only possible explanation of our demonstration, under appropriate conditions, of suicide being an ESS.

Who is suiciding? The graphs in Figure 4 contrast the distributions of age and health of the general population with those of the agents that perform suicide. There is a marked difference between these, with age presenting the greatest contrast. Somewhat surprisingly, younger agents are performing the bulk of suicides; more intuitive is the smaller peak in suicides amongst old agents. We expected that agents with lower health would have occasion to suicide altruistically, and Figure 4b supports that. We speculated, therefore, that youth suicide in the simulation is a side-effect of the correlation between youth and low health.

The benefits of suicide when the agent is either elderly or low in health are quite evident. By performing the sacrifice at a late age, the consequences to an agent’s fitness will not be great, since they will have few reproductive opportunities remaining. An elderly agent may benefit its own genetic endowment (through its children and other kin) more by withdrawing from the environment and releasing resources than by attempting to reproduce. Such freed resources include food and also mating opportunities. By disengaging in competition for these resources, nearby relatives will gain some advantage in reproduction, which in turn will promote genes that favour such altruistic behaviour.

The case for suicide amongst agents with low health is similar. Such an agent would not be able to mate prior to rebuilding its health. If resources are scarce, or about to become scarce, altruistic suicide favouring nearby kin would again be selected for. Of course, the negative effects of suicide would be at a minimum when an agent has both low health and old age.

Under what conditions is suicide more likely? Figure 5a shows how suicide probabilities have evolved in the chromosomes of agents given a cyclical food pattern — that is, the genetic basis of suicide in the simulation. All the possible extreme

⁴We also introduced a “behaviour” (birth defect) of death-upon-birth in order to directly measure the background mutation rates with a trait that could have no evolutionary advantage, finding a rate slightly below the suicide rate in the constant food simulation.

environments were examined to find the probabilities for suicide the agents have in such conditions, making sixteen different environments. These environments are listed along the horizontal axis, while the vertical axis indicates the normalized probabilities of suicide under those conditions.




Figure 5. Suicide with sinusoidal food distribution. (a) Average chromosomal probabilities of suicide for the last 50 epochs. On the horizontal axis, A=Age, P=Population, H=Health, F=Food; 0 indicates a low observation (the low observations are Health: 0 health units; Age: 0 cycles; Population: 0 agents; Food: 0 pieces of food), while 1 indicates a high observation (the high observations are Health: 600 health units; Age: 80 cycles; Population: 1600 agents; Food: 300 pieces of food). (b) Mean differences of suicide probabilities given observation extremes (e.g., high vs. low population).

There is a clear correlation between age and the probability of suicide. This supports our view that young age suicides are prominent because of the strong correlation between youth and low health in our simulation — that is, the chromosomal probabilities above show that it is not young age as such that is causing suicides. Also, high populations lead to greater suicide rates, as indicated in Figure 5b.

We also investigated the chromosomes from another angle, shown in the graphs of Figure 6 (averaged over 10 simulations). The graphs on the left show the frequencies of ages, health units and population densities in the simulation, and those on the right show the probabilities that are contained in the rule triggered by these observed values. These graphs help us understand the genetic background to suicide in this simulation, although they do ignore more complex conditions which trigger behaviour dependent upon more than one condition.

Figure 6a and b show that high youth suicide is resulting from a combination of higher youth population and higher propensity for the youth to suicide. Figure 6b makes plainer the increased tendency of the old to commit suicide.

Figure 7 breaks up suicide probabilities according to whether they are contained in rules that have a *less than* or *greater than* in the age condition. In the former case (Figure 7a), we can see that, for rules containing *less than* age conditions, ages less than 55 cycles trigger probabilities that decrease with greater age. After this point, probabilities flatten. Figure 7b tells a complementary story about suicide probabilities in the elderly.

Considering now health, we can find two peaks in the frequency of suicide in Figure 6c, one at 100 health units and another, weaker peak at 170 health units. Recalling that agents are born with a health of 100 units and that eating food increases health by around 80 units, these peaks seem clearly to be associated with youth. The peak in chromosomal suicide probability at 300 health units (Figure 6d) certainly looks out of place, but with the low frequency of agents at this health level it hard to to assess its significance. After health of 300, suicide probabilities become erratic and quite large at points. Given the extremely low frequencies of agents that have reached such heights of health, it would appear that natural selection has not been given enough opportunity to bring these probabilities to a consistent level.

Finally, we look at the relation between population size and suicide. Figure 6e shows the frequencies that different population levels are attained in the simulations. The suicide probabilities of Figure 6f are more erratic than in the other graphs, but one feature seems worth noting, namely that the suicide probabilities below 700 agents are low, which seems to support our idea of when agents are suiciding.

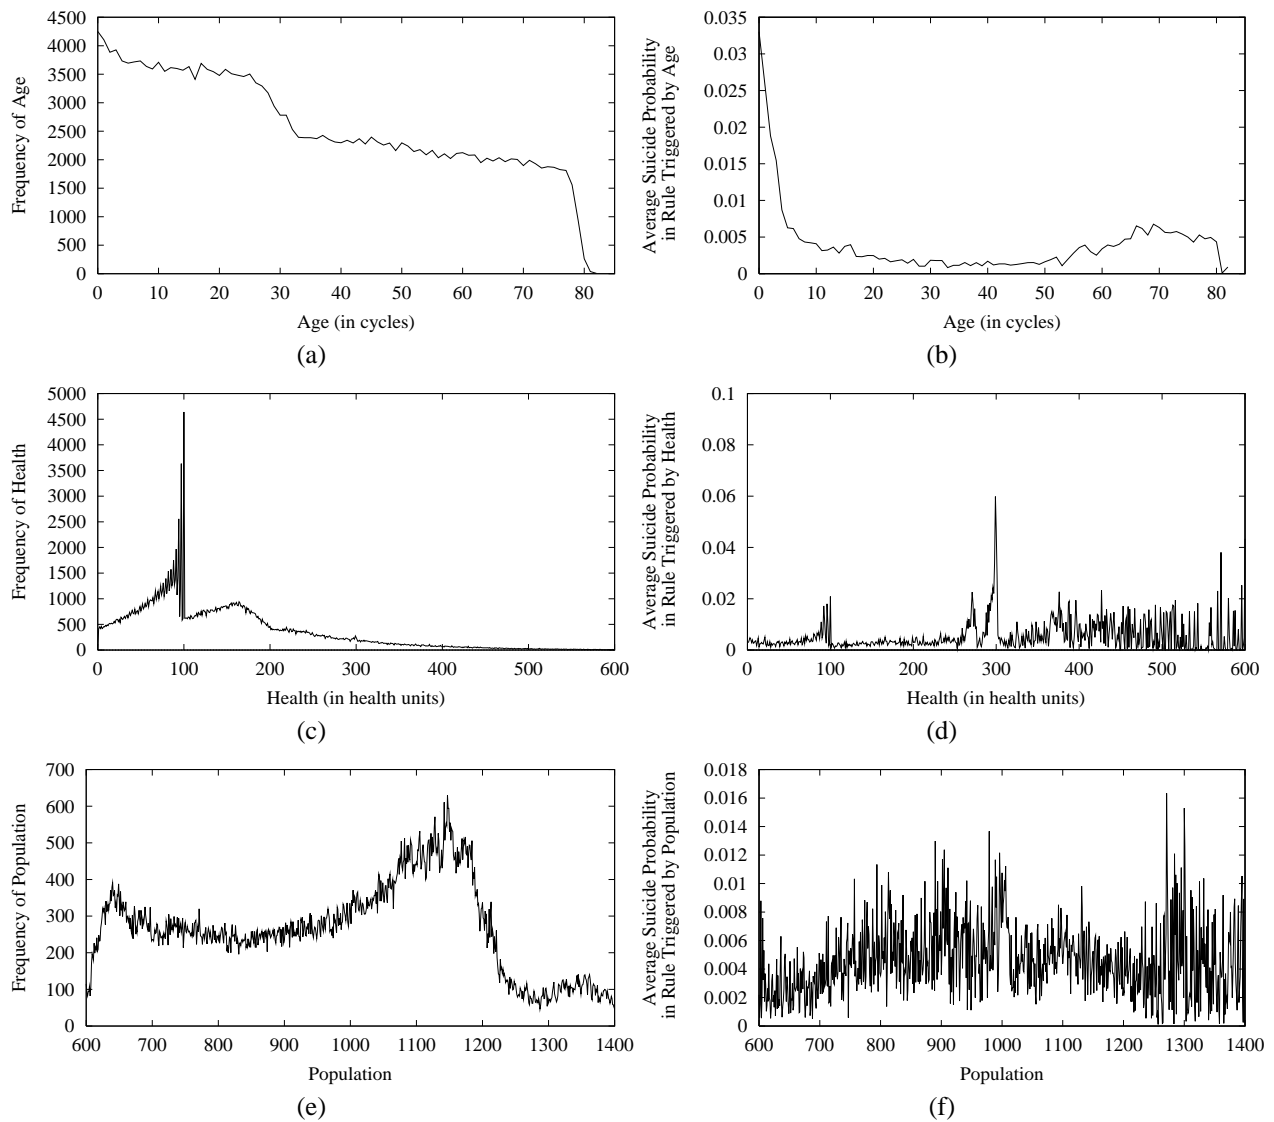


Figure 6. On the left hand side, frequencies of (a) age, (c) health and (e) populations encountered by agents during the simulation. On the right, probabilities in the rules of the chromosome triggered by (b) age, (d) health and (f) populations. Recall also that for these simulations, 5 cycles is equivalent to 1 epoch, and 1 health unit to 1 utile.




Figure 7. The distribution of probabilities in the rules of the chromosome for different ages given (a) a 'less than' and (b) a 'greater than' in the age condition of a rule.




Figure 8. (a) The mean correlation between genetic diversity and Manhattan distance over time, with confidence intervals. (b) Mean genetic diversity over time. 30 pairs of agents are sampled per epoch over periods of 20 epochs.

5.3. Kin Selection

One would expect the genetic relatedness between agents to be correlated to the distance between them. We measure the genetic diversity (the inverse of genetic relatedness) for two agents a and b stochastically as:

$$D(a, b) = \sqrt{\sum_{i=0}^n \frac{\|apv(a, obs_i) - apv(b, obs_i)\|^2}{n}} \quad (4)$$

where n is the number of environments (observation vectors) randomly sampled, obs_i is that vector of observations, apv is a function that takes an agent and an observation vector and returns the agent’s action probability vector. The difference between action vectors for the two agents is calculated for n observation vectors, and the root-mean-square of these differences gives the final value. This is a measure of the average difference between the probability vectors of the two agents under various environments; the environments chosen are weighted according to the environments actually encountered by the agents during the simulation.

We investigated whether genetic diversity is related to the Manhattan distance between agents. The Manhattan distance is simply a sum of the magnitudes of the x-translation and the y-translation needed to move from one agent to the other.⁵ We found that genetic diversity is low at zero distances (i.e., in the agent’s own square), but greater at other distances, allowing kin selection to operate. Although there is only a small linear correlation of 0.039 between distance and genetic diversity (measured during the early epochs 0 – 20, when the population genetics is undergoing the greatest change), this was statistically significant ($p < 0.001$).

Figure 8a shows the correlations between genetic diversity and the Manhattan distance (averaged over 44 simulations). The graph shows the correlation diminishing from an initial value of 0.039 to zero, after about 150 epochs; after that they appear to drift randomly. This is what we would expect given an initial period of rapid evolution, in which conditional suiciding behaviour is taken up by the wider population due to selection pressure, followed by an equilibrium state in which the population is homogeneous. Figure 8b shows, in support of this theory, that genetic diversity itself declines rapidly in the early stages of the simulation, with genetic homogeneity achieved around epoch 150.

5.4. The Ethical Value of Suicide

To judge the ethical value of suicide we can compare the total cumulative utilities of simulations which are identical except in allowing or disallowing suicide. Figure 9a is the outcome for cumulative total utility given sinusoidal food supply⁶. (The ripples show the deleterious effect of drought on the utilities of the agents.) Whereas the total utility in the simulations

⁵We used this for simplicity, despite the fact that it is not the optimal distance measure, since the agents may move diagonally.

⁶The averages of 44 simulations with suicide and of 31 simulations without suicide were used.




Figure 9. Comparison of (a) cumulative total utilities and (b) difference in average utility (suicide minus no suicide).

with suicide is initially lower, the situation soon reverses, with the cumulative total utilities of the two simulation types subsequently diverging. Figure 9b shows the difference in average utilities of agents in these same two simulation types. The average utility of agents when suicide is an option is greater at every epoch, hence the lower total utility in the early epochs of the suiciding simulations is entirely due to the smaller population sizes. Despite smaller early populations, it does not take long for the ethical virtue of assisting others to survive droughts to assert itself in higher cumulative utility. It is worth noting what is sometimes forgotten in discussions of utilitarian ethics, that the proper assessment of the ethical value of the action requires us to consider not just the immediate but also the future utilities of agents: the cumulative graphs of both average and total utilities settle into a roughly linear form only after about 80 epochs. After about 150 epochs, total utility tells the ethical story: in this world (that is, under the circumstances of this simulation), altruistic suicide is quite often the ethical option.

5.5. A Variation on Suicide: Altruistic Yeast

Recent research has uncovered altruistic behaviour in yeast via apoptosis, or programmed cell death (Frohlich and Madoo 2000). Apoptosis is a common form of cell regulation in multicellular animals, where cell death can be triggered by various chemical signals. However, its presence in monocellular yeast is suggestive of some altruistic function. Suiciding yeast decompose in such a way as not to release damaging enzymes and end up as effective nutrition for neighbouring yeast cells. Since the neighbour cells are most likely clones of the suiciding yeast cell, the selective advantage to kin is clear.

Our simulation is easily adapted to investigate the evolution of this type of behavior. We modified it so that suiciding agents leave behind a bundle of food of nutritive value equal to the health of the agent prior to the act of suicide. Some pertinent results of this adapted simulation are shown in Figure 10 (an average of 19 simulations). In comparison to the previous simulations, suicide rates are somewhat higher for the average, which would be expected given its additional benefit. The troughs in particular are higher ($t(22) = 8.44, p < 0.001$). What is quite striking is Figure 10b, showing how a low population is a very strongly preferred occasion for suicide in comparison with the prior simulation. Presumably, when the population is low the provision of a new supply of food has a substantial effect on the few agents left in the agent's neighbourhood.

6. Conclusion

We have demonstrated that suicide can be both an evolutionarily stable strategy and altruistic. More generally and importantly, we have illustrated the use of a-life simulations for testing applied ethical propositions of general interest. We believe this creates a new opportunity to discover empirical facts about important ethical problems computationally. We plan to use this a-life simulation framework to investigate other such actions, including euthanasia, abortion, rape and racial discrimination.

References

Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.




Figure 10. (a) Proportion of yeast-like suicide (suicide and turn into food) over all possible actions and (b) average chromosomal probabilities of suicide (the order of the environments is kept the same as Figure 5 for comparison).

- Bentham, J. (1789). *Introduction to the principles of morals and legislation*.
- Epstein, J. and R. Axtell (1996). *Growing artificial societies*. MIT Press.
- Frohlich, K.-U. and F. Madeo (2000). Apoptosis in yeast - a monocellular organism exhibits altruistic behaviour. *FEBS letters* 473, 6–9.
- Gilbert, N. and R. Conte (1995). *Artificial societies: The computer simulation of social life*. London: UCL.
- Hamilton, W. (1964). The genetical evolution of social behavior. *Journal of Theoretical Biology* 7, Papers I&II, 1–16 & 17–52.
- Holland, J. (1975). *Adaptation in natural and artificial systems*. University of Michigan Press.
- Holland, J. and J. Miller (1991). Artificial adaptive agents in economic theory. *American Economic Review* 81, 365–370.
- Russell, S. and P. Norvig (1995). *Artificial intelligence: A modern approach*. Prentice-Hall.
- Smith, J. M. (1982). *Evolution and the theory of games*. Cambridge University Press.
- von Neumann, J. and O. Morgenstern (1947). *Theory of games and economic behavior*. Princeton University Press.
- Wilson, E. (1975). *Sociobiology: The new synthesis*. Harvard.