Getting along with your neighbours – emergent cooperation in networks of adaptive agents

Tania Bransden and David G. Green

Faculty of Information Technology, Monash University, Clayton Victoria 3800 Australia

Email: tania.bransden@infotech.monash.edu.au

Abstract. The ways in which social order emerges from communication in human networks provides a model from which to draw lessons about the design of computing networks. In particular it provides lessons about the likely behaviour of swarms of intelligent agents. In human societies, cooperation and social order are maintained by a combination of laws that govern behaviour (prescription) and peer group pressure to conform. In this study we use simulations of social networks to show that both prescription and peer pressure are needed to achieve conformity. Our results show that dishonesty increases as social interaction decreases, but strongly interacting societies can promote either dishonesty as well as honesty. However, in a strongly interacting society, even a small degree of enforcement ensures almost universal conformity.

Key words: self-organisation, networks, complexity, social order, agents, evolutionary computation

1. Introduction

A major issue in modern computing is how to coordinate the activity of many different "agents". Is it possible for large swarms of independent agents with varied functions and priorities, to contribute towards a common goal? In distributed computing, for instance, many processors, at many different locations, need to communicate and coordinate their processing. Likewise, there is a growing recognition that large-scale Internet functions, such as web services and agent technologies, must be self-organising in order to be effective. Other emerging multi-agent technologies pose similar challenges. They include monitoring via sensor swarms and nanotechnology, where there is a potential need for large populations of nanobots to act in concert to perform useful tasks.

One of the most fruitful ideas in modern computing has been to solve complex problems by mimicking natural processes (T1,2). The field of intelligent and evolutionary systems, for instance, is based around computing analogies for human intelligence (e.g. automatic reasoning, machine learning, and neural networks) and evolution (e.g. genetic algorithms). The general approach is to identify a natural process, such as evolution, and identify ways in which it links with computing systems.

We argue that ways in which order emerges in human societies holds many important lessons for coordinating agent swarms. One of the great achievements of civilization, for instance, has been to create large, stable communities, even though they may consist of many millions of individuals. In this study, we test hypotheses about the emergence of human cooperation in large societies. In particular, we examine ways in which peer networks contribute to large-scale cooperation.

2. Social cooperation

The need for social order is great. Without communication and cooperation, a society quickly breaks apart and degenerates into conflict and self-interest. If everyone breaks a law, then it becomes unenforceable. We see elements of this phenomenon in extreme situations where law and order break down and society descends into anarchy. In 1977, for instance, a blackout in New York City resulted in widespread looting and violence as the normal restraints on behaviour were removed (3). Similar incidents often occur in extreme conditions, such as war and natural disasters. In 2005, hurricane Katrina in the Louisiana led to widespread disruption of social infrastructure, including law enforcement. As a result, the storm's aftermath included theft and violence on a scale never before seen in U.S. history.

One of the concerns in modern society is that modern technology and facilities have adverse side effects, including increases in crime and other anti-social behaviour (2). The reason for this is that they create conditions that reduce the effectiveness of peer groups.

One example of this is the increase in violence attributed to TV (4). One effect is that TV can change opinions by pumping out a message that is contrary to the prevailing beliefs or opinions. Previous simulation studies with our colleagues have shown that such a mass medium is so highly influential that only the most tightly bound peer network can resist (5,6). What happens is that initially the peer network counters the message put out by the mass medium by influencing members to retain the majority opinion. However, when the number of affected individuals reaches a critical level, the peer network actively accelerates the spread of the media's message.

In this study, we investigate another way in which modern technology has contributed to social breakdown. This is by the degeneration of the social network, leading to social isolation. Some examples will serve to illustrate the ways in which social communication networks can dissolve.

First, a side effect of mass media is that people have changed the way they entertain themselves. Instead of going out to social events, such as the theatre, or a party, they are increasingly likely to stay at home and watch TV or video. The result is a decrease in the amount of contact and communication with friends, family and neighbours (7).

A similar effect arises from some kinds of home technology. In Australia, for instance, the introduction of refrigeration made it possible for families to keep perishable foods longer, and reduced the frequency with which they needed to go shopping. At the same time, it also helped supermarkets to flourish, so driving many local stores out of business. Both processes contributed to a decrease in the frequency of random contacts that shoppers have with neighbours and tradespeople.

In summary, many processes may have contributed to decreasing social contact within modern society. We argue that this decreased contact has led to a similar decrease in the effectiveness of social networks.

Law and order

Why do people obey the law? Laws work only so long as most people obey them. All the laws in the world are meaningless if no one obeys them. Traditionally, the chief method of ensuring that people obey laws is enforcement. This is managed either by social vigilance (detection and punishment by peers), or by a specific group (e.g. a police force) that apprehends and punishes offenders.

Fear of punishment provides widespread motivation for obedience. Societies sometimes combine the two. In communist Russia, for instance, citizens were expected to act as informers against dissidents. When transgressions occurred, not only were the transgressors punished, but also those associated with them. Recent research has shown that this kind of strategy can be extremely effective at maintaining strict obedience (8).

However, if everyone breaks a law, then enforcement becomes impossible. If everyone stole whatever they needed, whenever they needed it, then society would quickly cease to function. Police would be unable to cope and even if they could, you cannot lock up everyone or basic infrastructure would break down. In other words, maintaining a law-abiding society depends on people cooperating. And cooperation depends on people's attitudes.

The aim of rules and laws is to ensure that benefits for all members of a society. However, in many circumstances, individuals can gain greater benefits by breaking laws rather than conforming. This betrayal of social trust is often represented in terms of the game theory model of prisoner's dilemma.

Prisoners Dilemma

One of the most widely studied models of social cooperation is the so-called prisoner's dilemma. In this "game", two prisoners are being held for a crime that they committed. If they cooperate with each other and neither confesses, then they both receive a light sentence. However, if either prisoner defects and confesses to the crime, then he gets off free whilst the other prisoner receives a heavy sentence. Although cast in terms of crime and punishment, the prisoner's dilemma game has been applied to many different social situations, such as profit and loss in business dealings.

Despite its widespread relevance, the prisoner's dilemma game focuses on individuals trying to maximize their advantage. It sees the question of social cooperation versus betrayal as a purely economic issue, based on deliberate individual strategies. In one sense peer group pressure can be seen as a matter of benefits: betray the group and you no longer enjoy the benefits of belonging to the group. However, real life cooperation is not only a matter of costs and benefits. It depends on many things. In practice it also involves strong elements of communication, habit, learning and inter-personal bonding.

Social discipline

Traditionally, discipline in human societies is maintained in two ways: pressure and prescription. In small bands of people, peer pressure is enough to ensure social cohesion. This phenomenon was identified by Dunbar (9), who pointed out that speech enables traditional human social groups to form larger bands (100-150 individuals) than apes and monkeys (40-60 individuals), which maintain social cohesion by mutual grooming. However peer pressure fails in large societies: cohesion breaks down because communication between individuals is not enough to ensure full connectivity (10).

In general, large societies and social groups maintain order by prescription. This takes many forms, such as the customs of a social group, edicts of chiefs and elders, and laws passed by parliament. In modern society, laws are enforced by police and administered by an extensive legal system.

Social networks

Prescription alone is not enough to ensure compliance by members of a population. An important contribution to social obedience comes from peer group pressure. Without pressure from a peer group, social prescriptions fail. Peers reinforce views and attitudes that promote cooperation and conformity within their group.

In social networks, with strong interactions, rebellion can lead to new peer groups. At Yoyogi Park, Harajuku on Sundays, thousands of Tokyo teenagers all rebel as a group. They dress in outrageous punk costumes and play loud music. They try to assert their individual identities against the norms of Japanese society. And yet in a sense they do not rebel at all. They do all this en masse, in effect asserting that they belong to this rebellious teen social group.

3. Methods

3.1 The model

To answer the questions that we posed above, we use a simulation model that is similar to the ones used in previous studies, both by ourselves and our colleagues (e.g. 5, 6).

Individuals

In our simulation model of social networks, we represent a social system as a network of "actors" (agents, individuals), each of which represents an individual (person, computer, software agent).

Each of the actors in the social network has a state that reflects the attitude or opinion of the person concerned. In this model, the state (the actor's "attitude") represents commitment to obeying the law. So their attitude can take one of two values: HONEST or DISHONEST. If they are honest, then they will obey the law, except under the most extreme circumstances. If dishonest, they will not hesitate to break the law. Admittedly this is a simplification of what happens in real life, but we feel that it is a reasonable one.

Peer-peer interaction

Communication between actors in the society occurs as a series of social exchanges or communication "events" involving pairs of agent, reflecting (say) conversations between people. These events occur either at regular intervals or at random, within the social networks that form between friends, colleagues, family and casual acquaintances. In each event, two actors communicate and influence one another. For any interaction between a pair of individuals, A and B say, the possible outcomes will be as follows:

1. If A and B initially agree, then neither changes opinion.
2. If A and B initially disagree, then either:
   - One of the pair (selected randomly) switches to share the other's opinion. This outcome occurs with probability p_{change}.
   - Neither changes opinion. This outcome occurs with probability $1-p_{change}$.

The model processes peer-peer interactions in random asynchronous fashion (11). That is, the model selects a pair at random, and determines the outcome of their interaction, before going on to process another pair.

Social structure

The edges in the network represent social "ties" and define the patterns of communication links between the individuals. The network structure is fixed but may have different patterns of connections, including random, scale-free, small-world and hierarchical (12).

Time

The model simulates successive time periods (e.g. days). During each time period, the individuals undergo a number of interaction events. The number of events within each time period is set by a parameter N_{events}. Also, at the conclusion of each time period, every actor has an opportunity to commit a crime.

Crime and punishment

Independent of the peer network, the model includes both incentives and disincentives to break the law.

The chief incentive to break the law, and perhaps become dishonest, is financial stress. We assume that honest individuals will not break the law, unless they are in financial distress, such as not being able to pay bills, or pay for food and shelter.

Here we represent this issue by introducing an index F, (range 0-1), which indicates (say) the degree of economic stress within the community. The analogy here is with real-life issues that individuals must face, such as cost of living or taxation. In the model, we represent the current financial needs of actors as a number I, and we assume that honest actors do not break the law unless $I > F$, that is (for instance) if they owe more than society can support them.

During each time period, every actor has an opportunity to break the law. In the model we compute which people break the law at the end of each time period. This is done as follows:

1. If attitude(A)= DISHONEST, then A breaks the law.
2. If attitude (A)= HONEST, then either:
   IF I(A) > F, then A breaks the law;
   ELSE A obeys the law.

For each time period, the chief outputs of the model are:

1. The *crime rate* (the number of individuals who broke the law);
2. The number of honest individuals in the society.

Apart from peer influence, the main motivation for being honest is the threat of punishment. Here we assume that a percentage of individuals who break the law get caught. This percentage corresponds to (say) the amount of effort put into law enforcement within the community.

We do not introduce a cost for actors who are caught breaking the law. However, we assume that the attitudes of individuals may change according to whether they get caught when they do break the law.

Individuals who do not commit a crime do not change their attitude. If an individual does commit a crime, then his/her state may change, depending on whether or not he/she is caught. Table 1 summarizes the possible outcomes.

Table 1. Probability of changing attitude after a crime

Initial attitude of the criminal	Result of the crime	
	Caught	Evades capture
DISHONEST	p_{reform}	0
HONEST	0	$p_{corrupt}$

3.2 Virtual experiments

We set up virtual experiments using our model to answer the following two questions for our virtual society.

Experiment 1.

Question. How does discipline change as social contact declines? This is motivated by the hypothesis that peer group pressure is essential to keep people obeying the law. So we are testing whether decreasing social contact leads to an increase in crime and other anti-social behaviour.

Method. To address this question we observe how the crime rate changes as the number of social contacts (peer pressure) changes? One extreme is where there is no peer contact at all.

Table 2. Parameter settings for Experiment 1.

Parameter	*Value*
No. of people in society	250
Initial percent of people who are honest	100
Peer-peer interactions per time period	10,000
Peer influence	0.5
Probability of corrupting honest people	0.1
Probability of criminals reforming	0.0
Economic stress	0.1

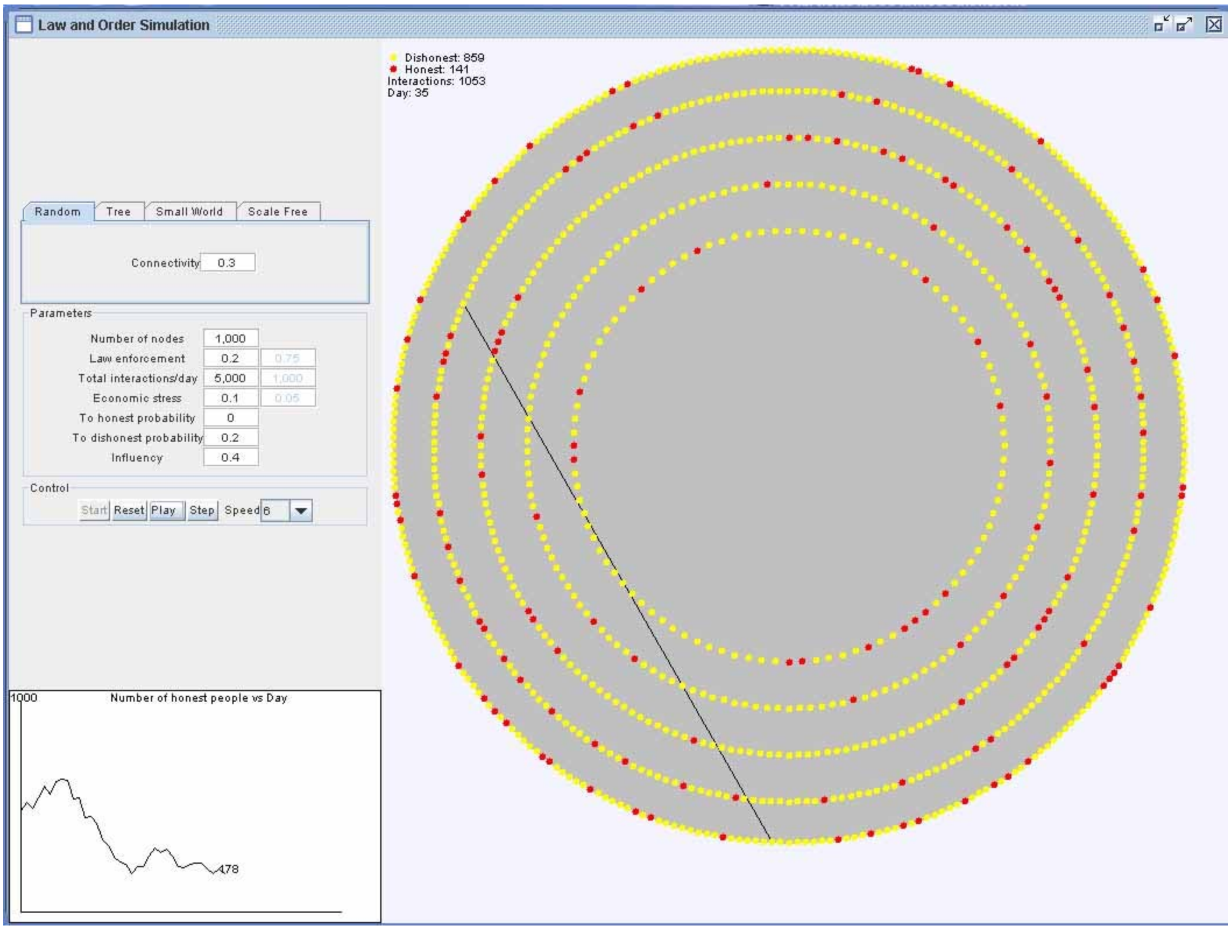

Figure 1. User interface to the visual display version of our "law and order" simulation model. The dots represent individuals in the society. They are arranged in circles for efficient display (a population of 1000 is displayed here). The colour of each dot indicate the person's attitude (dark = honest, light = dishonest). The black line joining a pair of dots represents an interaction that is in progress between two individuals. At the top left of the screen are control buttons and boxes allowing the user to select the type of network and the values for key parameters. At the bottom left, the graph plots changes in the number of honest people in the society with time. In this case, the grey background results from a dense network of ties between pairs of individuals in the network.

Experiment 2

Question. How does discipline change under extreme conditions? This question is motivated by the observation, discussed earlier, that law and order often break down in conditions of extreme hardship.

Method. To address this question we observe how the crime rate changes when the degree of economic stress on individuals approaches extreme values. In our model, this parameter could take values anywhere in the range 0-1.

We answered the questions by using sensitivity analysis to see how the crime rate changes. That is we systematically altered values of the stress parameter to see how the system responded.

3.3 Implementation

We implemented the above model in the java language, using many of the same routines that we developed for previous models of social networks (13). We created two versions of the model: (1) An interactive model with visual display (Fig. 1), allowing users to follow the detailed behaviour of the system, and (2) A high performance version aimed at carrying out sensitivity studies and scenario testing involving multiple runs that summarise the results of systematic changes in parameter values.

4. Results

4.1 The degenerate society

In the extreme case where there are no peer-to-peer interactions, our model reduces to a markov process. We can obtain the transition matrix for changes in the attitudes of individuals by combining the rates in Table 1, with the rate at which honest actors commit crimes. The resulting transition matrix of the process is:

$$
\begin{array}{cc} & \begin{array}{cc} H & D \end{array} \\ \begin{array}{c} H \\ D \end{array} & \begin{pmatrix} 1- fp_{corrupt} & fp_{corrupt} \\ p_{reform} & 1- p_{reform} \end{pmatrix} \end{array}
$$

This process becomes absorbing if any of the three parameters is zero. So if either f or p_{corrupt} is zero, then H (HONEST) becomes an absorbing state: that is, the entire population eventually winds up HONEST (H). If p_{reform} is zero, then D is an absorbing state and the entire population ends up DISHONEST. If both terms are zero, then the system never deviates from its initial state. If neither is zero, then the process is ergodic, which implies that the system eventually settles into an equilibrium in which the population contains a proportion of both honest and dishonest people.

4.2 Experiment 1 - Influence of a peer network

The peer network alters the above picture and the extent of the difference depends on how much interaction occurs between peers (Fig. 2). When the number of interactions is small in our simulations, the number of honest people is always low. However, as the number of interactions per period increases, the society tends to polarize, becoming either all honest, or all dishonest.

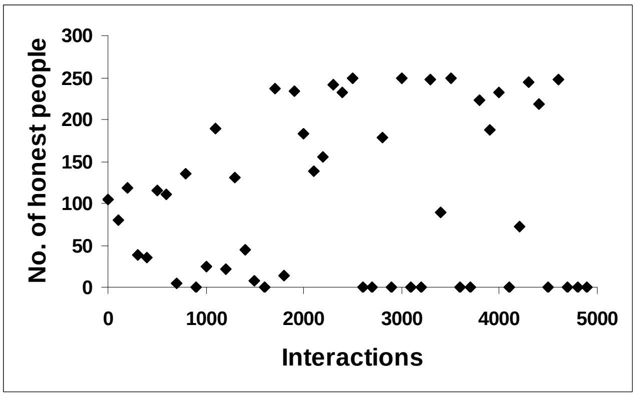

Figure 2. Changes in the number of honest individuals in a society as the number of interactions per time period increases. The values of parameters used in the model are given in Table 2.

4.3 Variations on the initial model

The results in figure 2 assumed that no reform occurred. When the experiment is repeated with even a low percentage (10%) of criminals reforming, the result is dramatically different (Fig. 3). The number of honest people in the network is always high (usually >80%).

In the above model, the crimes that people committed had no repercussions. However, inter-personal crimes (e.g. robbery, assault). always have victims who suffer as a result. To test whether inter-personal crimes led to different behaviour by the society, we implemented a version of our model that differed from the initial model in two ways.

(1) Individuals committed crimes against randomly selected individuals in the society. However, who commits crimes was decided in the same way as before.

(2) We assumed that victims of crimes reacted by immediately becoming HONEST, irrespective of their previous attitude.

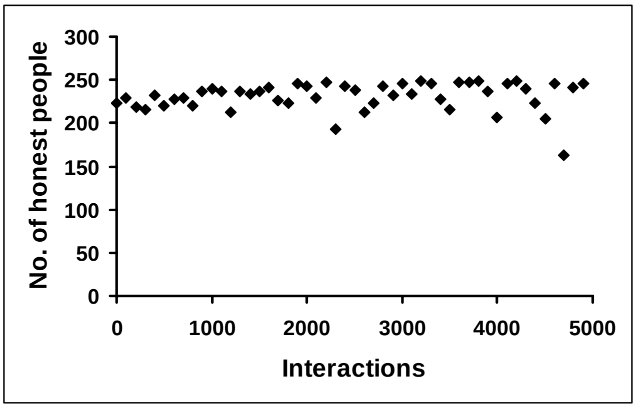

Figure 3. Changes in the number of honest individuals in a society as the number of interactions per time period increases. The values of parameters used in the model are all as in Figure 2, except for reform, which is set at 0.1.

The result of this model was similar to the initial model with reform, but the numbers of honest people was always close to 100% (Fig. 4).

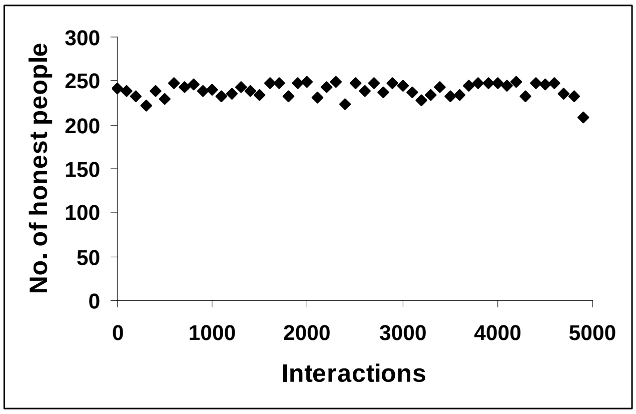

Figure 4. Effect of inter-personal crimes. The graph shows the changes in the number of honest individuals in a society as the number of interactions per time period increases. The model settings are all as in Figure 3, except that crimes are inter-personal (see text for explanation).

4.4 Experiment 2

In the experiment to test the effects of increasing economic stress, we used the same settings as in Figure 3, but varied social stress, instead of interactions. We fixed the number of peer-peer interactions per time period at a constant 5000. We varied economic stress from 0 to 1.0.

The results (Fig. 5) showed that crime rate increased steadily with economic stress. However, the number of honest people remained surprisingly high except at extremely high levels of economic stress.

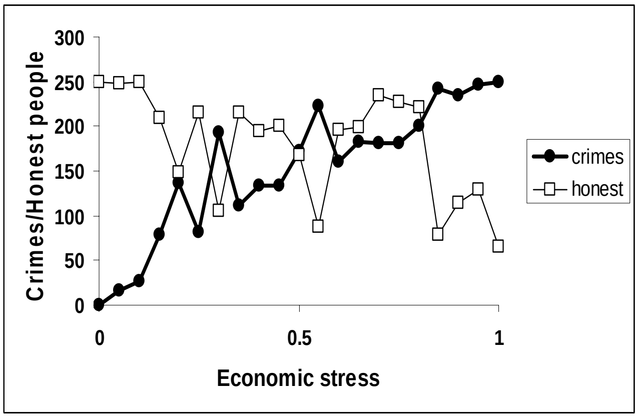

Figure 5. Change in crime rate and honesty within a social network as economic stress increases.

5. Discussion

The behaviour of our model social simulation, as the number of interactions increases, is consistent with our previous studies of consensus in social networks (10). In that model the entire society ended up in the same state, provided the network connectivity was sufficiently great. In those experiments, no account was taken of *which* state the society members all ended up in. However, in the present model, the end state does matter. What happens here is that in the majority of cases, the entire society ends up either honest or dishonest, with no one dissenting.

The results support the claim that decreasing social interaction leads to increases in crime rate.

The surprising implication of our results is that a peer network can reinforce either attitude (HONEST or DISHONEST), but it may not be possible to predict which one will win. It is important to ask whether this result is valid: can this happen in real human societies? At first sight it may seem odd to have a society in which everyone is a criminal. The key here is that in our model, the crimes do not have any impact, either on other individuals, or on global features of the society as a whole. Obviously this does not happen for (say) robbery, in which victims lose money, or (say) assault, in which victims are injured.

Nonetheless, there are crimes of the kind represented in our model. A good example is drivers speeding on the road. In the vast majority of cases, exceeding the speed limit has no effect at all. Compared to the number of times that drivers exceed the speed limit, accidents are relatively rare. In many cities and countries around the world, exceeding the speed limit is so common as to be almost the norm. In other words, virtually the entire society can become dishonest in respect to speeding.

It is likely that a social system would behave differently for different types of crimes. Crimes against individuals, for instance, directly impact upon the social network itself. In future research, we plan to examine the effects of different kinds of crimes, and interactions.

In most of the above experiments, we assumed that dishonest people never reformed through punishment. Nevertheless, peer pressure did lead to reform of many criminals, and kept most honest people honest. Variations on the original model show that when supported by peer pressure, law enforcement can succeed in keeping almost everyone in a population honest. These results were even more marked for inter-personal crimes, assuming that the experience of being the victim of a crime turns people honest.

The second experiment showed that even though increasing social stress may lead to an increase in crime, it does not necessarily lead to a prevalence of dishonesty amongst the population.

The results reported here are far from definitive. Further research needs to link our findings to changes observed in real societies. The relevance of the results is limited by the assumptions that underlie the models. In real societies, many other factors can come into play. Because of this complexity, real societies are likely to behave in more plastic fashion than our models.

Finally, our results have implications for computing. For instance, the individuals could be agents, and their "attitudes" might be their priority (COOPERATE or IGNORE) with respect to some goal of the network. Suppose that all the agents need to cooperate at a particular time in order to achieve the goal. Then our results make clear predictions about the ability of a network to achieve a goal. If the agents do not communicate frequently enough, then no consensus will emerge. If they do communicate frequently, then is likely that the entire network will either work towards a common goal, or else all ignore it.

As with our previous research on social networks (5, 6, 10), this study shows that simulation of networks of agents provides a potentially powerful took for examining deep social questions, as well as providing useful lessons about computers and computer networks.

Acknowledgments

This work was supported by the Australian Centre for Complex Systems Science, which is funded by the Australian Research Council. We are grateful to our colleagues for their help, especially Dr Suzanne Sadedin, who provided useful comments on the manuscript, and Alex Tee Neng Heng, who implemented the java versions of the model.

References

1 Kirley, M., Li., X., and Green, D. G. (1998). "Investigation of a cellular genetic algorithm that mimics landscape ecology", In X Yao et. al. (eds) *Simulated Evolution and Learning, Lecture Notes in Artificial Intelligence 1585*, Springer, pp. 90-97.

2 Green, D.G. (2004). *The Serendipity Machine-A Voyage of Discovery through the Unexpected World of Computers*. Allen and Unwin, Sydney.

3 Anonymous (2005). *The Blackout History Project*. http://blackout.gmu.edu/

4 Williams, T.M. (1986). *The Impact of Television - a Natural Experiment in Three Communities*. Academic Press. New York.

5 Stocker, R., Cornforth, D. and Green, D.G. (2002). The impact of television on cohesion in social networks - a simulation study. In Namatame, A., Green, D., Aruka, Y. &Sato, H. (eds) *Complex Systems 2002*. Chuo University, Tokyo. pp. 222-228.
6 Stocker, R., Cornforth, D. and Green, D.G. (2003). A simulation of the impact of media on social cohesion. *Advances in Complex Systems* 6(3), 349-359.
7 Putnam, R.D. (2001). *Bowling Alone: the Collapse and Revival of American Community*. Simon and Schuster, New York.
8 Newth, D. (2005). "Altruistic Punishment, Social Structure and the Enforcement of Social Norms." KES 2005, LNAI 3683, pp. 806–812. Springer-Verlag Berlin Heidelberg
9 Dunbar, R. (1996). *Grooming, Gossip and the Evolution of Language*. London: Faber and Faber.
10 Stocker, R. Green, D. G., and Newth, D. (2001). "Consensus and cohesion in simulated social networks", *Journal of Artificial Societies and Social Simulatio*n 4(4), http://www.soc.surrey.ac.uk/JASSS/4/4/5.html
11 Cornforth, D., Green, D.G., Newth, D. and Kirley, M.R. (2004). Ordered asynchronous processes and modularity in complex systems. *Physica D Nonlinear Phenomena* 204, 70-82.
12 Green, D.G. and Bransden, T. (2006). "Complexity theory". In *McGraw-Hill Encyclopedia of Science and Technology*. McGraw-Hill, New York, pp. 507-511.
13 Heng, T.N. and Green, D.G. (2004). VLAB – The Artificial Life Laboratory Online. http://www.complexity.org.au/vlab/