Environmental Risk, Cooperation and Communication Complexity

Peter Andras, Gilbert Roberts, John Lazarus
Department of Psychology
University of Newcastle
Newcastle upon Tyne, NE1 7RU, UK
{peter.andras;gilbert.roberts;j.lazarus}@ncl.ac.uk

Abstract

The evolution of cooperation and communication in communities of individuals is a puzzling problem for a wide range of scientific disciplines, ranging from evolutionary theory to the theory and application of multi-agent systems. A key issue is to understand the factors that affect collaboration and communication evolution. To address this problem, here we choose the environmental risk as a compact descriptor of the environment in a model world of simple agents. We analyse the evolution of cooperation and communication as a function of the environmental risk. Our findings show that collaboration is more likely to rise to high levels within the agent society in a world characterised by high risk than in one characterised by low risk. With respect to the evolution of communication, we found that communities of agents with high levels of collaboration are more likely to use less complex communication than those which show lower levels of collaboration. Our results have important implications for understanding the evolution of both cooperation and communication, and the interrelationships between them.

1 Introduction

Understanding how cooperation between unrelated individuals can arise in animal and human societies has puzzled evolutionary theorists. Early solutions to the problem were found in terms of reciprocal altruism; the mutual exchange of benefits between pairs of individuals (Trivers 1971, Axelrod & Hamilton 1981, Axelrod 1984, Roberts & Sherratt 1998). More recently, 'indirect reciprocity', in which individuals who are seen to be more generous receive more help from others, has been proposed as an additional route to cooperation (Alexander 1987, Nowak & Sigmund 1998, Leimar & Hammerstein 2001).

In these models individuals have information about the past behaviour of others on which to base decisions but there is no communication of intentions: individuals simply act; cooperating, defecting or declining to interact. (This is in contrast to the evolutionary modelling of competitive behaviour in which signalling has played a central role: e.g. Maynard Smith 1974, Enquist 1985.) Although there seems to be little theoretical work on intentional signalling in the context of cooperation, arbitrary signals correlating with altruism (the 'green beard effect', Dawkins 1976), tags indicating individual identity (Riolo et al. 2001) and signalling of partner quality (Leimar 1997) have been considered. While existing evolutionary models of cooperation are important in examining the minimal conditions for the evolution of cooperation they are also impoverished - at least for the human case - in excluding the possibility that, intentionally or unintentionally, individuals may communicate their intention to cooperate before interacting. In the development and maintenance of human relationships cooperation is accompanied by signals of short-term intentions and longer term commitment. Honest communication, and consequent trust, are of the greatest importance for the development of a stable collaborative relationship; deceit and mistrust are inimical to it (Boon & Holmes 1991).

In this paper we develop an agent world model to examine the evolution of cooperation when individuals communicate their intentions. This communication helps individuals to decide whether to enter into cooperation with another: it allows partner choice. Many of the earlier models provide no such choice, so that cheats can be avoided only if individuals have information on their past behaviour. When the behaviour of others is to some extent predictable, however, individuals can derive the intentions of others and cooperators can choose to interact with other cooperators while cheats can be ostracized (Roberts 1998).

We ask how cooperation and communication respond to variation in the risk or complexity of the environment. Risk and complexity are important general properties of the environment that impact on an individual's success in ways that can be influenced by cooperating with others. For example, resources may be predictable (low risk) or unpredictable (high risk), and threats from predators or competitors may vary in a similar way. Resource acquisition, the avoidance of predators and success in competition can all be enhanced by collaboration with others.

The second focus of the paper is how the complexity of communication itself evolves in this context and whether it differs between cooperators, cheaters and those who decline to interact at all. This interest in the communication of intentions follows other work in artificial societies (e.g. Schillo et al. 2000), which investigate how collecting information about the intentions of other agents can enhance the development of collaboration.

Our agents communicate by sequences of signals, each of which informs the potential partner about the cooperative intentions of the agent. The reliability of the signals differs between agents, who use the information from communication in present and past interactions to make decisions about whether to share resources with a potential partner. The interaction between agents occurs within a risky environment where risk refers to the variability of the gains that result from cooperative behavior. Our results will be applicable to agent societies, animal societies, and human social systems.

The rest of the paper is structured as follows. First, we describe the agent world model. Next, we present the simulation results. Finally, we discuss the implications of our results.

2 The agent world

In this section we describe the world of our agents. We start with the description of the environment, followed by the description of the agents, communication processes, resource management, the offspring generation rules, and we close with the description of the evolution of the agent society. We also present some argumentation to support our choices with respect to the implemented principles, rules and methods.

2.1 The environment

The environment of our agent world is characterized by a given risk. The environmental risk represents in a compact way the complexity of the surrounding environment. The environmental risk is implemented as the variance of the agent's resource regeneration process.

2.2 The agents

Our agents dispose over some generic resources that they use to maintain themselves and to reproduce.

Each agent speaks the same communication language consisting of the symbols: '0','s','i','y','n','h' and 't'. The meaning of the communication symbols are as follows: '0' – no intention of communication, 's' – start of communication, 'i' – maintaining the communication, 'y' – indication of the willingness to engage into resource sharing, 'n' – indication of no further interest in communication, 'h' – effective sharing of the resources, 't' – not sharing the resources after an indication of willingness to engage into sharing. The last two symbols, 'h' and 't' actually mean the resource-sharing or no-resource-sharing actions. The first four symbols are ranked according to their positive contribution towards engagement in sharing (the least positive is the '0' and the most positive is 'y').

The agents generate communication units (i.e., one of the above symbols) when they engage in communication with another agent. Each agent has its own realization of the language. This language is represented in the form of a two-input probabilistic automaton (i.e., it is equivalent of a probabilistic push-down automaton). The language units are production rules of the form

$$
\begin{aligned}
L:\ & U_{current}, U'_{current} \rightarrow_{p1} U_{new,1} \\
& U_{current}, U'_{current} \rightarrow_{p2} U_{new,2} \\
& \dots \\
& U_{current}, U'_{current} \rightarrow_{pk} U_{new,k}
\end{aligned}
$$

where $U_{current}$ is the last communication unit produced by the agent, $U_{current}$' is last communication unit produced by the partner of the agent, $U_{new,1}$, $U_{new,2}$,… $U_{new,k}$ are the new communication units that can be produced by the agent, and p1,p2,…,pk are the probabilities of production of these communication units, p1+p2+…+pk=1 (an example of a such rule is: L: i,i' $\rightarrow_{0.4}$ y, i,i' $\rightarrow_{0.5}$ i, i,i' $\rightarrow_{0.1}$ n that means that after producing the symbol 'i', and receiving a symbol 'i' from the communication partner, the agent will produce the symbol 'y' with probability 0.4, the symbol 'i' with probability 0.5, and the symbol 'n' with probability 0.1).

The language units obey intention consistency rules, i.e., if U0,U1,U2, and U3 are communication units, and U2 is equally or more positive than U1, and U3 is equally or more positive than U2, and L1 is a language unit that produces U3 after U1 and receiving U0 , and L2 is a language unit that produces U3 after U2 and receiving U0, then the probability of producing U3 using L2 is equal or higher than the probability of producing U3 using L1. Similarly, if L1 is a language unit that produces U3 after U0 and receiving U1 , and L2 is a language unit that produces U3 after U0 and receiving U2, then the probability of producing U3 using L2 is equal or higher than the probability of producing U3 using L1. In other words the intention consistency rules mean that more positive inputs are more likely to lead to positive outputs than are less positive inputs. This choice of the intention consistency rules is in agreement with human and animal behavior, where the expression of friendly signals is more likely to be followed by further friendly signals by the same individual than non-friendly signals.

Each agent has a characteristic intention, which indicates the extent to which it is willing to share resources with other agents. This sharing intention determines the probability of the y,y' → h production rule.

The agents are equipped with a memory. The memory of the agents can store the experiences of collaboration with the last M different partners (M=10 in our implementation). The memory of the agents also fades with time, and if they don't meet an old partner for long time they forget their memories about this partner. For each memorized partner the agent keeps the score of the successful and unsuccessful meeting (i.e., successful means a meeting that led to getting shared resources from the partner).

The agents are located on a two dimensional plane, and they may change their location. The location of an agent determines the neighbourhood of the agent that consists of the N (N=10) closest agents.

The agents live for T (T=60) time units. In each time unit they try to find a collaboration partner in their neighbourhood. At the end of their life time the agents produce their offspring.

2.3 Communication processes

After selecting a collaborating partner the agents may engage in a communication process. The communication process starts properly after both agents communicated the 's' symbol. We set a limit (L_1) for the preliminary communication. If the two agents do not reach the proper start of the communication in a communication of length L_1 we consider that they stop their communication at this moment.

During the communication process the agents use their own realization of the common language to produce communication units. The communication process ends either with the communication of an 'n' symbol (i.e., signalling no further interest), or with the communication of the 'y' symbol by both partners. After this each agent decides whether to share or not to share their resources with the other agent by producing the action symbol 'h' or 't'. We impose a communication length limit (L_2) on the proper collaboration oriented communication. If the agents do not reach the stage of communicating the 'y' symbols in L_2 communication steps, we consider that they stop their communication.

At the start of each communication process, the agents update their language unit probabilities according to their memories of the agent they are currently interacting with (note that if there had been no previous interaction with this agent there is no update). The updated version of their language applies only to the present communication process. In the case of more positive experiences (those that led to sharing) the probabilities leading to more positive symbols are increased, while in the case of negative experiences these are decreased. The probabilities for each language unit are normalized after effecting the above changes (i.e., the probability of producing all the allowed new symbols is always one for each language unit).

During each communication process, as an agent produces equally or more positive symbols their willingness to share increases. We note that although this increase happens in all agents, those who have very low intention to collaborate will increase an originally low probability, which means that they will not necessarily share at the end of the communication process. We adopted this collaboration willingness increase principle in conformity with human and animal behavior, where a sequence of expression of friendly signals increases the likelihood of the friendly ending of the interaction, even if the original intentions were less friendly.

2.4 Resource management

The agent dispose over their own resources that they use to maintain themselves and reproduce. In each turn the agents use their available resources to produce new resources. If they manage to find a partner who is willing to share its resources they can use the combined resources to generate the new resources for the next turn.

The mean resource generating function is a squashing function of the form:

$$\overline{R}_{new} = a \cdot \frac{1}{1+e^{-R+R_0}}$$

where R is the amount of available resources, and R_0 and a are parameters. Operating at the convex half of the squashing function (i.e., $R < R_0$) means that using more added resources is more advantageous than using the resources separately.

Resource generation happens in a probabilistic manner. The environmental risk specifies the variance of the resource regeneration process. The amount of new resources is found by taking a sample from a uniform distribution that has the calculated mean and the variance specified by the environmental risk. We use the notation N(R) for the amount of new resources generated by using R amount of available resources.

The variance of the resource regeneration increases with the time spent in negotiation about resource sharing. This risk increase principle is in agreement with how environmental risk changes in the real world. To exemplify it we consider an example from business. If two companies start negotiations about a joint business, lengthy negotiations may proceed while a competitor enters in the market, and the final gain of the two collaborating companies will be reduced. At the same time lengthy negotiations may lead to a well-designed contract that makes possible to avoid future impasses, making the collaboration more profitable. If the negotiations are short, the deal is made quickly, and the companies may start gaining some new market share. At the same time they may run into some unregulated disputes that may slow down their cooperation and the increase of their market share. As we can see from this example, if the communication process is short the variance of the expected benefits is likely to be smaller than in the case when the communication process becomes lengthy.

When two agents meet, having resources R1 and R2, and they both decide to share their resources they may receive extra new resources. The extra resources for both partners are calculated as the half of the difference

$$N(R1+R2) – (N(R1) + N(R2))$$

Such agents are called collaborators.

If an agent is engaged in a communication process, convinces its partner to share, but then withholds its own resources from sharing, it is called a cheater. The gain of a cheater is the whole amount of the difference

$$N(R1+R2) – (N(R1) + N(R2))$$

In such case the one that is cheated generates only N(R2) new resources for itself, i.e., it does not benefit from the sharing.

If two agents select each other as communicating partners, but they do not manage to decide about the sharing of their resources (i.e., their communication end with an 'n' symbol) we call them non-collaborators.

If an agent does not reproduce enough resources to maintain itself (i.e., the maintenance costs are higher than the amount of own resources) the agent reaches the zero resource level and dies.

2.5 Offspring generation

When the agents reach the end of their lifetime they generate their offspring. The number of the offspring depends on the available resources of the agent.

If the agent has R resources, and the mean amount of the resources in the agent society at that moment is R_m, and the standard deviation of the resources is R_s, then the number of offspring of the agent is calculated as

$$n = \alpha \cdot \frac{R - (R_m - \beta \cdot R_s)}{R_s} + n_0$$

where α, β, n_0 are parameters.

If n is negative or R=0 we consider that the agent produces no offspring. If $n > n_{max}$, where n_{max} is the allowed upper limit of offspring, we cut back n to n_{max}.

In order to avoid strong generational effects the newly generated offspring have random ages between 1 and A_0.

The offspring of an agent inherit from their parent its language and collaboration intention with small random modifications. They also inherit the resources of their parent equally divided between the offspring.

2.6 The evolution of the agent society

At the beginning we start with randomly initialised agents, i.e., the transition probabilities of their language units, their collaboration intentions, initial resources and initial positions are set randomly.

The agent society evolves through the interaction and reproduction of the agents. The agents search for collaborating partners. They try to share their resources, or to cheat the collaborating partner, or they may not manage to make the decision about sharing. In each turn each agent may choose one partner from its neighbours. After each turn the agents make a random move, changing their position, and possibly finding a new neighbourhood.

The agents regenerate their resources alone or in collaboration with another agent in each turn, and they pay a part of their resources to maintain themselves. At the end of their lifetime the agents generate their offspring if they have enough resources.

3 Simulation results

This section presents our simulation results. The objective of these simulations were twofold. First, to determine how environmental risk affects both the level of collaboration, and the complexity of communication. Second, to examine how the various strategists (collaborators, cheaters, non-collaborators) differ in complexity of their communications in an evolving society. The results are presented in this order, after examining some general effects of risk level on the agent society.

We selected five levels of risk in the range of 0.1 – 0.9 (the risk levels were 0.1, 0.2, 0.5, 0.7, and 0.9). We measured communication complexity by measuring the average length of the communication processes.

For each risk level we run 20 simulations to obtain valid estimates of average values and variances of the measured variables. The number of agents in each simulation was the same at the beginning (1500). We run each simulation for 400 time units, or until the agent population died out or reached the maximum allowed level of number of individuals (5000).

3.1 General effects of the environmental risk

To see the general effects of environmental risk on the agent society we looked at how the number of agents and the resource level varied with time. We considered the average amount of resources separately for collaborators, cheaters and non-collaborators.

Figure 1 shows the average number of agents in the agent societies for the five risk levels. Figures 2 – 4 show the change over time of the average amount of resources of collaborators, cheaters and non-collaborators.

The first segment of dropping in the graphs represents the period when the randomly initialised population selects those who are able to survive. This segment corresponds to one generation (i.e., around 60 time units).

Following the initial drop the societies start to grow in number and in average amount of resources in all cases. The graphs show that this growth happens much faster in agent societies living in low risk environment than in those which live in high risk environments.

These results confirm the standard expectation that the average level of populations and their available resources is lower in high risk environments than in low risk environments.

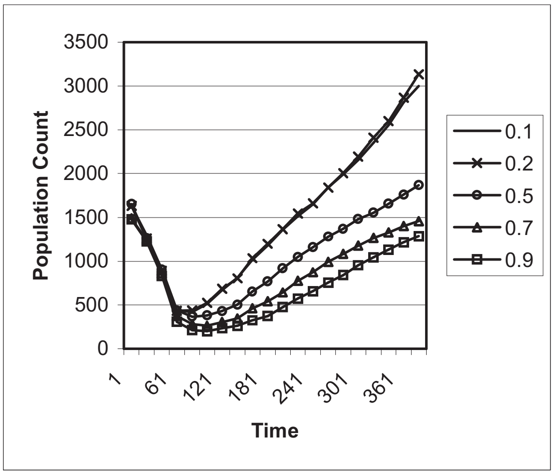

Figure 1: The evolution of the number of agents in agent societies living in environments with different risk levels.

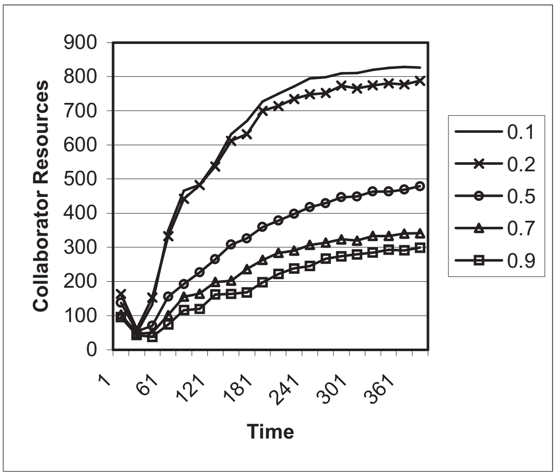

Figure 2: The evolution of the average amount of resources of collaborators in agent societies living in environments with different risk levels.

3.2 Environmental risk and the level of collaboration

To analyse the effect of the environmental risk on the level of collaboration we looked at the percentage of collaborators, cheaters and non-collaborators within the society (note that the percentage of those who were cheated is the same as the percentage of cheaters). The change of these percentages over time is shown in Figures 5 – 7.

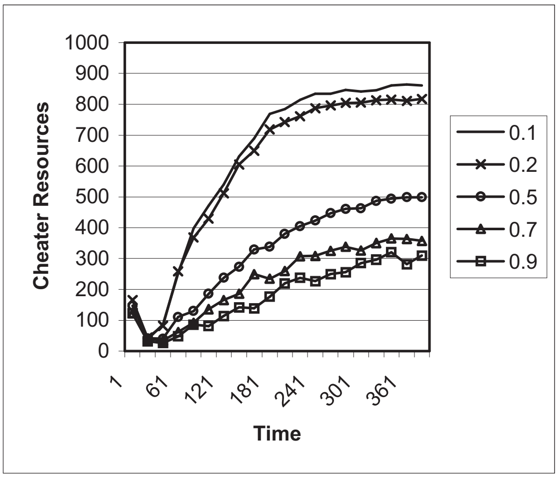

Figure 3: The evolution of the average amount of resources of cheaters in agent societies living in environments with different risk levels.

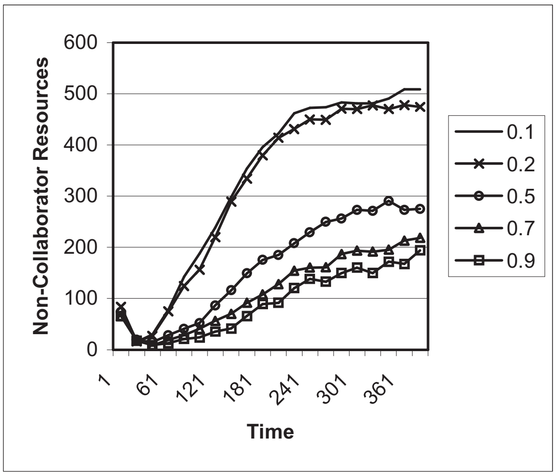

Figure 4: The evolution of the average amount of resources of non-collaborators in agent societies living in environments with different risk levels.

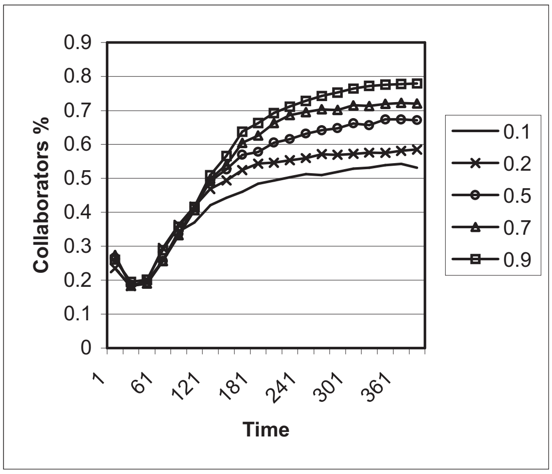

Figure 5: The evolution of the average percentage of collaborators within the agent societies living in environments with different risk levels.

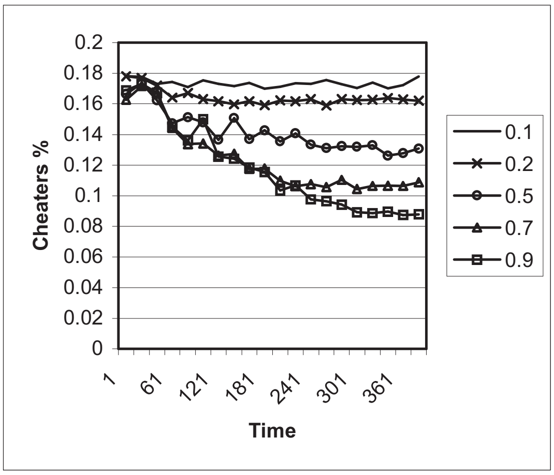

Figure 6: The evolution of the average percentage of cheaters within the agent societies living in environments with different risk levels.

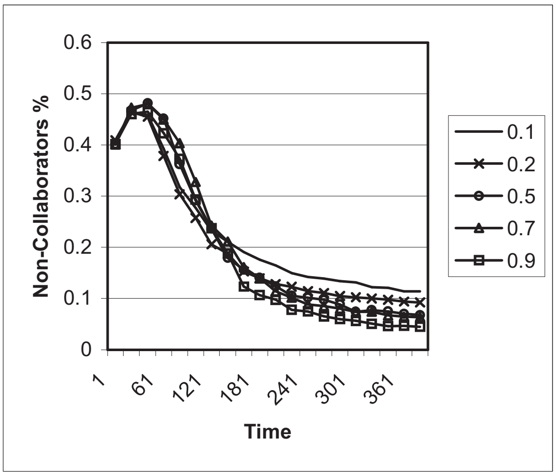

Figure 7: The evolution of the average percentage of non-collaborators within the agent societies living in environments with different risk levels.

The figures show that the level of collaboration increases in all conditions. After the first generation (i.e., around 60 time units) the level of collaborators increases steadily until it stabilizes (above 50%). In the case of cheaters and non-collaborators there is a corresponding decline to stabilization at below 18% for cheaters and below 12% for non-collaborators.

The figures show that the stable level of collaborators is lower in low risk conditions than in high risk conditions, and that levels of cheaters and non-collaborators are higher in low risk conditions than in high risk conditions. This indicates that the agent societies living in a high risk environment are more likely to achieve high level of collaboration than those which live in low risk environments. We note also than in the high risk environments it is more likely that the population dies out than in low risk environments.

3.3 The complexity of communications

We analysed the complexity of communications by measuring the average length of communication processes within the whole society.

Figures 8 shows the evolution over time of the communication complexity in the whole society.

The figure shows that there is no clear ordering of the stable levels of communication complexity as a function of the level of environmental risk. The same is true when each of the three agent strategies is examined independently.

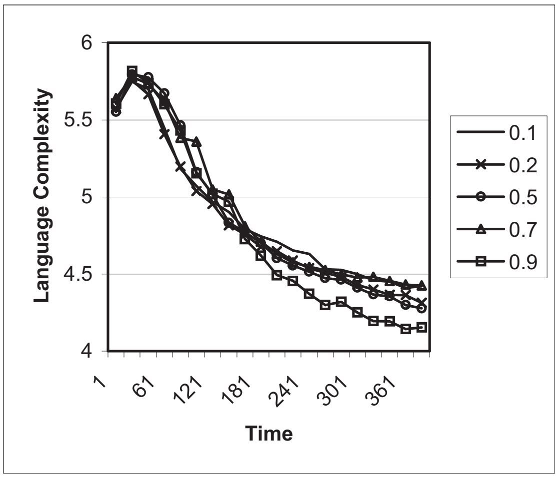

Figure 8: The evolution of the average communication complexity within the whole agent societies living in environments with different risk levels.

3.4 Collaboration and communication complexity

First we analysed the correlation between the levels of collaborators, cheaters and non-collaborators and the average complexity of communications within the society. These correlations are shown in Tables 1 – 3.

Risk	0.1	0.2	0.5	0.7	0.9
Correlation	-0.94	-0.98	-0.99	-0.99	-0.99

Table 1: The correlation between the average percentage of collaborators and the average complexity of communications within societies living in environments with various risk levels.

Risk	0.1	0.2	0.5	0.7	0.9
Correlation	0.15	0.74	0.93	0.95	0.97

Table 2: The correlation between the average percentage of cheaters and the average complexity of communications within societies living in environments with various risk levels.

Risk	0.1	0.2	0.5	0.7	0.9
Correlation	0.97	0.98	0.99	0.99	0.99

Table 3: The correlation between the average percentage of non-collaborators and the average complexity of communications within societies living in environments with various risk levels.

These results indicate that the proportion of collaborators is strongly negatively correlated with the complexity of communications at all risk levels. In the case of cheaters we see that their proportion is moderately positively correlated with the average complexity of communications at low risk levels, and that the correlation gets much stronger for high risk levels. In the case of non-collaborators their percentage is strongly positively correlated with the average communication complexity at all risk levels. These results together suggest that those who collaborate tend to communicate in shorter sequences, while those who cheat or do not collaborate are likely to use longer communication sequences.

To analyse this suggestion directly, we examined how communication complexity evolves in the three different groups at given risk levels. Figures 9 and 10 show two examples for risk levels 0.2 and 0.9.

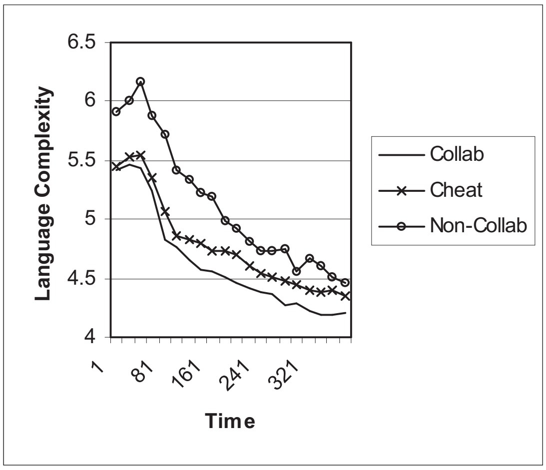

Figure 9: The evolution of the average communication complexity in the groups of collaborators, cheaters and non-collaborators living an environment characterized by risk level r = 0.2.

These figures confirm the suggestion that those who collaborate are likely to use less complex communication between themselves, and those who do not collaborate use more lengthy communication processes.

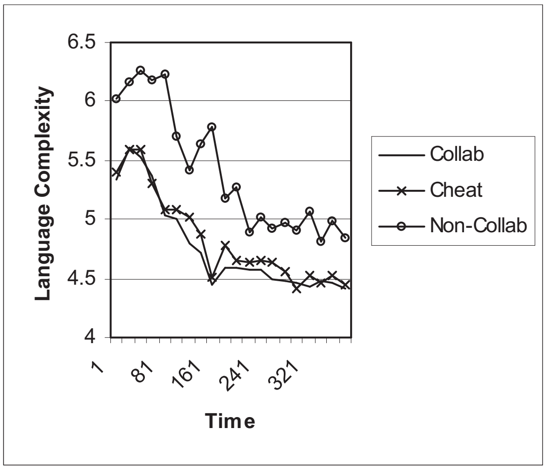

Figure 10: The evolution of the average communication complexity in the groups of collaborators, cheaters and non-collaborators living an environment characterized by risk level r = 0.9.

4 Discussion

First, we interpret our results from the more general point of view of evolution of collaboration and communication in societies of individuals. Second, we discuss the implications of the presented results for agent worlds and multi-agent systems.

4.1 Evolution of collaboration and communication in societies of individuals

Under the assumptions of our model society, cooperation can thrive and its frequency increases with environmental risk, while both cheating and non-cooperation decline. The increase in cooperation with environmental risk probably comes about because cooperation can be crucial for survival when resources are very low and/or provides particularly large rewards (compared to cheating) when resources are very high. This is because while cheating is profitable in the short term (for a few interactions), in the longer term cheaters fail to find other agents who will interact with them. The fact that population size and average resource level decline as risk increases supports the conclusion that cooperation becomes increasingly advantageous in difficult or harsh environments, as measured by risk.

The prediction that cooperation is more likely in risky environments can be tested in animal societies, in human experimental groups and in the real world of human social and economic behaviour, at the level of both individuals and of groups such as firms and nations. For example, this prediction might help to explain the phenomenon of increased feelings of community during wartime. It also suggests that cooperation might be enhanced by increasing the risk or complexity of the problem at hand. Although there may also be costs associated with increasing risk, if perceived risk increased while objective risk remained unchanged then cooperation might be enhanced without cost. However, as perceived risk increases the population of those willing to participate is likely to decline. A possible application area here is communication on the Internet, although there would be ethical issues involved in deceiving users about the risk or complexity of the Internet environment.

In contrast to its effect on cooperation, environmental risk in the model had no clear effect on the length of communication. This may have been because the model language was too simple, varying between only 4 and 6 elements at the outset. For a richer language we predict that communications will be shorter as risk increases since: (1) there is a positive correlation between collaboration level and risk, and negative correlations between cheating and non-collaboration levels and risk (Figures 5-7), and (2) there is a negative correlation between collaboration level and language complexity, and positive correlations between cheating and non-collaboration levels and language complexity (Tables 1-3).

Those who collaborated had shorter communication strings than those who cheated or failed to collaborate. This is because if a collaborator meets an agent for whom it has a memory biased towards collaboration then it has higher probabilities of production for positive communication symbols and therefore moves more quickly (i.e., with fewer communication steps) into an interaction that is likely to be collaborative. Thus collaborating agents, by positive feedback, build an increasingly cooperative relationship with each other, in a manner analogous to that described by Roberts & Sherratt (1998). The direct complement of this process is that meeting a past cheater for which an agent has a memory increases that agent's likelihood of cheating in the present interaction after a longer series of communication symbols (see sub-section 2.3 on the communication process).

Collaboration thus brings with it the bonus of a saving on communication effort. Such effort may be trivial, but it may also be considerable, as in some forms of human negotiation. The prediction that collaboration simplifies the communication process (compared to both cheating and avoiding interaction) can be tested in the scenarios already described for examining the relationship between cooperation and risk.

4.2 Collaboration and communication in agent worlds

We see two directions of implications of our work in the context of agent worlds and multi-agent systems.

First, our results indicate that appropriate setting of the environmental risk factors of an agent world can determine to a significant extent the level of collaboration within the agent world. This may have applications in the design of multi-agent systems where the developers wish to achieve some desired mix of collaborative/non-collaborative behavioural patterns that fits to the objective of the system. It is important to note that pure collaborative behavior in an open agent world may pose significant risks to the proper working of that world, as malignant agents may appear, and may abuse the default benevolent behavior of other agents. This means that some level of non-cooperative behavior should be allowed in an open agent system (Sherratt and Roberts 2001).

Second, our results suggest that it is possible to predict the expectable level of collaboration and communication complexity in an agent world, if enough information is available about the environmental risk factors characterizing this world. Such predictions can form the basis for checks of the validity of risk factor assumptions, and for corrective actions aimed to keep the agent world within the desired range of macro parameters.

References

R.D. Alexander. *The Biology of Moral Systems*. Aldine de Gruyter , New York, 1987.

R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.

R. Axelrod, & W.D. Hamilton. The evolution of cooperation. *Science,* 211:1390-1396, 1981.

S.D. Boon and J.G. Holmes. The dynamics of interpersonal trust: resolving uncertainty in the face of risk. In: R.A. Hinde and J. Groebel (eds.) *Cooperation and Prosocial Behaviour*, pp. 190-211. Cambridge University Press, Cambridge, 1991.

R. Dawkins. *The Selfish Gene*. Oxford University Press, Oxford, 1976.

M. Enquist. Communication during aggressive interactions with particular reference to variation in choice of behaviour. *Animal Behaviour*, 33:1152-1161, 1985.

O. Leimar. Reciprocity and communication of partner quality. *Proceedings of the Royal Society of London Series B: Biological Sciences,*264:1209-1215, 1997.

O. Leimar and P. Hammerstein. Evolution of cooperation through indirect reciprocity. *Proceedings of the Royal Society of London Series B: Biological Sciences,* 268:745-753, 2001.

J. Maynard Smith 1974 The theory of games and the evolution of animal conflicts. *Journal of Theoretical Biology*, 47:209-221.

M.A. Nowak, and K. Sigmund. Evolution of indirect reciprocity by image scoring. *Nature,* 393:573-577, 1998.

R.L. Riolo, M.D. Cohen, and R. Axelrod. Evolution of cooperation without reciprocity. *Nature,* 414:441-443, 2001.

G. Roberts. Competitive altruism: From reciprocity to the handicap principle. *Proceedings of the Royal Society Series B: Biological Sciences*, 265:427-431, 1998.

G. Roberts and T.N. Sherratt. Development of cooperative relationships through increasing investment. *Nature*, 394:175-179, 1998.

M. Schillo, P. Funk and M. Rovatsos.Using trust for detecting deceitful agents in artificial societies. *Applied Artificial Intelligence*, 14:825-848, 2000.

T.N. Sherratt and G. Roberts. The role of phenotypic defectors in stabilizing reciprocal altruism. *Behavioural Ecology*, 12:313-317, 2001.

R.L. Trivers. The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46:35-57, 1971.