How to cause war and why we may be moving away from it:

Grouping by encounter probability in an FSS society

Els van Dijk
0616419
E.T.vandijk@student.ru.nl

August 23, 2009

Thesis advisors:
dr. W.F.G. Haselager
dr. I.G. Sprinkhuizen-Kuyper

Contents

Abstract

The question of how to obtain high levels of cooperation in a Prisoner's Dilemma-like situation has been the subject of many models and simulations. Few have, however, tackled the problem from as realistic a viewpoint as Nakai and Muto have, with their artificial society in which Friend Selection Strategies are evolved. They try to mimic how human societies show not only periods of 'peace' (high cooperation), but also periods of 'war' (less cooperation). Although results are promising, their model only shows plausible results with regard to human societies when the society is of a certain size; for large societies, the model only shows peaceful behaviour, without periods of war.

In order to improve the plausibility of the model – and thereby also that of its behaviour – for larger societies, we will divide the society into groups with high intra-group encounter probabilities and low inter-group encounter probabilities. Results show that by this addition a war-like state is reintroduced into the model's behavioural repertoire, thus improving the model's behaviour in terms of plausibility.

Introduction

Most people have no problem going out everyday not knocking other people out and taking their money. In fact, most rarely even consider the possibility of taking for themselves what belongs to another. In most human societies social policy dictates that people generally treat one another decently.

However, from a purely rational, self-interested point of view, this makes no sense: why would a person not just take what they want? It may leave someone else deprived of what was previously theirs and perhaps it will even leave them emotionally, or otherwise, scarred. But since there is only benefit in it for the person doing the taking, why should he be bothered by such considerations?

This line of reasoning obviously leads to a dark horizon: if we were all to adhere to this rationale, the society we live in would soon pass into a state of 'war of all against all'. It is plain to us that such a society would not be by far as desirable to its inhabitants as a peaceful one.

This contradiction between what is rational and what is, in the long run, most beneficial to all parties involved is captured by the game theory problem of the Prisoner's Dilemma. Two agents engaging in a prisoner's dilemma will have a rational tendency to 'defect' (take negative action towards the other agent). However, if both defect, both are worse off than if they had both cooperated. To understand how we as rational beings may have overcome this dilemma (to a certain extent), it could therefore be useful to look at how we can get an artificial society of agents who engage in prisoner's dilemma-like tasks to behave as we do: cooperate most of the time.

Many such models have been proposed, tackling the problem from different viewpoints[1]. For most of these models, the aim is to give 'good' results – achieving high levels of cooperation between agents. However, they do not necessarily give plausible results. We all know, whether from history books or personal experience, that peace (or high levels of cooperation) in a society or group, once achieved, does not necessarily last forever. Generally, there will be periods of peace, alternated with periods of 'war' or at least less cooperation. It is exactly this emergence and collapse of peace that Nakai and Muto [4] seek to explain with their model of Friend Selection Strategies, or FSS's.

In this paper, the plausibility of results of the FSS model are examined for societies of different sizes. Analysis of the results suggests that some of the problems that arise may lie in the fact that one of the assumptions made in the model is no longer plausible for large groups with regard to real-world human societies. In order to remove this discrepancy the mechanism of grouping by encounter probability is introduced. The question then is whether this addition improves the model's results in terms of plausibility for large societies.

1 Nakai & Muto: the FSS society

In essence, Nakai and Muto's model is a basic artificial society: a certain number (`n_rounds`) of rounds are played, in which each of the `N` agents in the model meets a certain number (`M`) of other agents (see Table 1 for an overview of the model's most important parameters and their default values, where applicable). In each of these interactions the agent ('performer') who meets the other agent (the 'performed') decides whether to cooperate or defect. Depending on the choice of the performer, both agents obtain a payoff, as per Table 2.

Note that only the performer (agent A) has a say in what happens. The prisoner's dilemma becomes apparent when the payoffs for an encounter where B is the performer and A the performed is included in the payoff table (see Table 3). This form of interaction – 'generalized exchange' – is used, because it is closer to reality than 'restricted exchange', where interaction is mutual: The victim of an assault does not generally have the opportunity to get back at his attacker; only if they happen to meet again in a situation where the victim is in control can punishment be exerted.

[1] See [4] for a more extensive list of related research.

Parameter	Meaning	Default value
N	the number of agents in the society	–
M	the number of agents each agent interacts with in a round	N-1
R	the reflection ratio: the ratio of agents that adopt the best-performer's FSS at the end of a round	0.1
μ_S	the strategy mutation rate: the probability that an agent's FSS is mutated into a random FSS at the end of a round	0.003
μ_P	the perception error rate: the probability that an agent's social perception is flipped at the end of a round	0.05

Table 1: The most important parameters in Nakai & Muto's model, with default values where applicable.

	payoff A	**payoff B**
A cooperates	0	0
A defects	0.5	-1

Table 2: Payoff table for the game played by two agents in each encounter.

1.1 Friend Selection Strategies

What distinguishes this model from others is the mechanism that determines which action is chosen by the performer in an interaction. In some models, only personal gain in the form of payoff is taken into account by an agent when choosing an action. Humans' behaviour towards others is, however, to a large extent based on how the other is perceived. This idea is also reflected in models that use 'discriminator strategies' ([5], [3], [6]). These strategies allow an agent to determine whether another agent is considered a 'good' or a 'bad' person. Based on this perception, the agent cooperates or defects respectively when encountering that agent.

The perception of an agent as a 'good' or a 'bad' person in discriminator strategies is based on all actions taken by that agent, towards all other agents. This is a strange assumtion in two ways: firstly, all agents' all actions need to be visible to each agent in the society, which is an implausible assumption, especially for large societies. Also, in human societies, people will not only not know about certain actions, they are likely to also not care. Generally, a person will only take into account actions taken towards himself or those who are important to him: his friends.

Nakai and Muto's FSS's are based on this assumption. An agent's FSS allows the agent to determine a social perception of each other agent: is this person a friend or an enemy? If an agent is considered a friend, he will be cooperated with when encountered, if an agent is considered an enemy, he will be defected on when encountered. The social perception is based, depending on the FSS the agent has, on the other agent's behaviour towards either the agent himself, or his friends. See Table 4 for a list of the FSS's used in the model and a brief explanation of each of these. For example, if an agent has the FSS 'US-TFT', if this agent meets another agent, he asks all agents he considers friends whether the other agent attacked them when last they met. If most say yes, the other agent is considered an enemy and is, as such, defected on. If most say no, the agent is considered a friend and is cooperated with.

	B cooperates	B defects
A cooperates	A: 0 B: 0	A: -1 B: 0.5
A defects	A: 0.5 B: -1	A: -0.5 B: -0.5

Table 3: Summed payoff table for two encounters between two agents, where A is performer in one encounter and B is performer in the other; payoffs for A are in the upper left corner of cells, payoffs for B in the lower right corner.

Strategy	Description
ALL-D	any agent met is considered an enemy.
ALL-C	any agent met is considered a friend.
me-TFT	tit for tat: an agent who defected on me is considered an enemy, an agent who cooperated with me is considered a friend.
me-CWD	coward: an agent who defected on me is considered a friend, an agent who cooperated with me is considered an enemy.
us-TFT	tit for tat for friends: an agent who defected on my friends is considered an enemy, an agent who cooperated with my friends is considered a friend.
us-CWD	coward for friends: an agent who defected on my friends is considered a friend, an agent who cooperated with my friends is considered an enemy.

Table 4: The six Friend Selection Strategies.

1.2 Evolving behaviour

Nakai and Muto's model uses evolutionary principles to allow the behaviour of agents to evolve across rounds. Since agents' behaviour is caused by their FSS's, it is to the FSS's that the evolutionary principles are applied.

At the end of each round, every agent's payoffs in that round are summed, yielding a total payoff for each agent. These total payoffs form the fitness scores by means of which agents' performances are compared.

'Survival of the fittest' is now implemented by some ratio (`R`) of agents who performed worst adopting the FSS of the best-performing agent at the end of each round. In order to ensure that a certain amount of variety in FSS's is kept within the society, mutation is added to the mix: at the end of a round each agent has a certain small probability (μ_S) of having its FSS set randomly to one of the six FSS's.

1.3 Results

As mentioned in the introduction, Nakai and Muto try not only to show how peace can emerge, but also how it can collapse. The results they show in their paper are promising; see Figure 4(a) for a graph much like the ones Nakai and Muto show to build a case for their model. This graph shows how the ratio of friends evolves across rounds, in a society where `N`=20 and `M`=19. The friend ratio is the sum of the numbers of agents each agent considers a friend, divided by the maximum possible number of friends in the society; if the friend ratio is 1, every agent considers every other agent a friend and if the friend ratio is zero, none of the agents have any friends.

Looking at Figure 4(a), we can see that the society starts in a 'war of all against all'-like state,

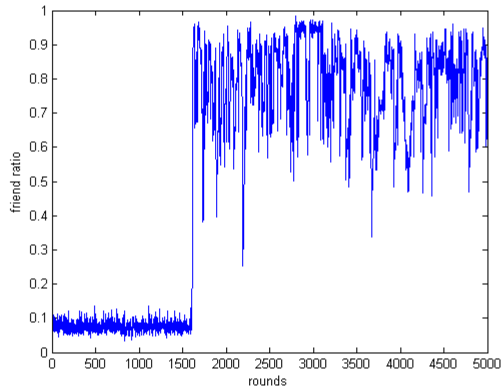

Figure 1: Development of friend ratio across rounds; `n_rounds`=5000, `N`=20, `M`=19.

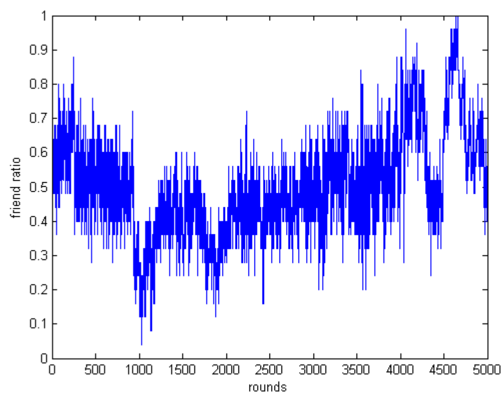

Figure 2: Development of friend ratio across rounds; `n_rounds`=5000, `N`=5, `M`=4.

with friend ratios below 0.1. Then, around the 1600th round, as if by divine intervention[2], 'peace' seems to emerge, with much higher friend ratios. This peaceful state is much less stable than the war-like state, with many small negative peaks that almost look like 'crime waves': they emerge quickly and have quite a negative influence on friendly feelings in the society, but the society soon recovers from them.

2 Problems of the FSS society

Although the effects that can be found in Figure 4(a) can be plausibly coupled to effects we see in our own society, the model's behaviour is not as plausible overall as that particular graph would have us believe. Subsequent runs do confirm the model consistently shows the effects described above for the settings used to obtain Figure 4(a). However, with different settings for `N` and `M`, at least some of the effects disappear.

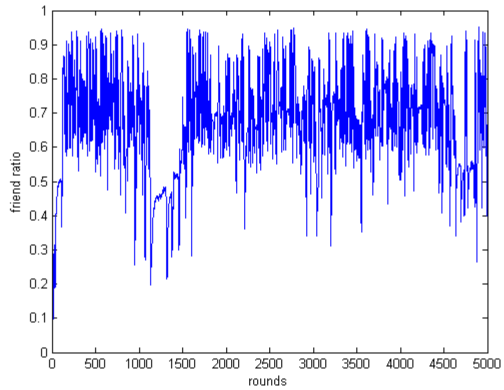

Figure 3: Development of friend ratio across rounds; `n_rounds`=5000, `N`=120, `M`=119.

2.1 Low `N`: too unstable

In a small society (Figure 2), it seems like neither the peaceful nor the war-like state we would hope to find are present. The friend ratio shows what look like random fluctuations; there does not seem to be any pattern there that would suggest the model's behaviour is converging to either a peaceful or a war-like state. The behaviour of this small society is, in short, much less stable than that of the larger society of 20 agents.

The opposite would have made more sense, if we compare to the human societies we are trying to mimic. Liebrand [1] shows that, when faced with a Prisoner's Dilemma-like situation, at least some subjects were more prone to defecting on a larger scale when interacting with a larger group of others. This effect may be explained by the fact that in a small society, each person's influence on the welfare of the entire society – which in turn affects one's own welfare – is relatively large, so cooperation is more crucial to one's own survival.

2.2 High `N`: too stable

In a large society (Figure 3), the peaceful state that was also present in the smaller society of 20 agents is still there. However, the war-like state that made the model's behaviour so uniquely realistic before is not by far as convincing as it was in the smaller society. There are negative peaks, but they do not reach to such low friend ratios as in the smaller society and they also last less long.

As described in the previous section, it would make sense for agents in a smaller society to be more friendly amongst themselves than they would be in a larger society. This effect of higher `N` causing more stable and higher levels of cooperation, therefore, is far from plausible with regard to human societies.

2.3 The likely culprit: `M`

There may be a common cause for both the instability of the friend ratio in small societies and the extreme stability of the friend ratio in large societies: `M`.

If `N` is small, `M` – being `N-1` – is also small. This means that each agent only has a few encounters with other agents in each round. As the summed payoffs across encounters in a round are used as the fitness function, any chance encounter can now have a very large influence on an agent's fitness score. It is therefore hardly guaranteed that the agent with the highest fitness

[2] see [4] for an elaborate explanation of the mechanisms that cause these sudden jumps from war to peace and vice versa.

score at the end of a round is also the agent that would have done best in the long run. Any agent, regardless of its potential, can, relatively easily, seem the best-performing agent and can, as such, influence the behaviour of the society as a whole. For small `M`, the fitness function loses its ability to identify the actual best-performing agent, so the evolution of agents' behaviour is based on chance, more than anything else, causing the society's behaviour to fluctuate wildly rather than converge.

For large `N` and `M`, the opposite is the case: in each round, agents have every opportunity to 'show their stuff' and since a chance encounter is only one out of many, it can only have very small influence on an agent's fitness score. The agent with the highest fitness score at the end of a round is therefore likely to be one with a high actual fitness, making it very easy for the evolutionary mechanism to do its work and push the society towards the optimal behavioural pattern of cooperation.

The alternation between war and peace we see in a 'medium-sized' society of 20 can now be explained as an effect of `M` being at a threshold value, where the model's behaviour exhibits a combination of the effects of small and large `M` described above. `M` is large enough that the two behavioural patterns of peace (optimal) and war (suboptimal) can emerge. However, `M` is too small for these behaviours to be maintained once they emerge, so the model jumps between the optimal and suboptimal behaviours.

3 Grouping by encounter probability

The fact that the FSS-mechanism introduced by Nakai and Muto only gives realistic results for medium-sized societies would suggest the quasi-stability in smaller and larger human societies are caused by other, or additional, mechanisms.

If we look at how large groups interact in Nakai and Muto's model, we see that the results are not realistic, but neither are the settings: for a society of 120, it is assumed that all agents in that society meet every other agent in the society equally often. In human societies, this is simply not the case: those who are close to us, be it geographically or socially, we meet often. Those who are far away we hardly ever, or even never, meet. For instance, one might meet one's next-door neighbour each day, but only see that lady who lives down the street at the annual barbeque. There is more to the story: the people we meet more often are not just random people – they are part of our family, soccer team, group of co-workers, etc. The whole of a society is divided into smaller sub-societies of people who interact regularly, whereas interactions between people from different subsocieties are more scarce.

This grouping by proximity-dependent interaction probability is not reflected in Nakai and Muto's model. Adding this mechanism may help the model behave more realistically for large groups.

The concept of adding proximity is not new. Venkat and Wakeland [7] examined the influence on proximity-dependent cost of transactions on an agent-based model of a simple artificial economy. This model however, deals with economic interactions rather than social ones and, as such, makes use of nor exhibits any social mechanisms.

Macy and Skvoretz's research [2] is a little more similar to that described in this paper: they divide an artificial society into neighbourhoods with high interaction densities within, and low interaction densities between neighbourhoods, to examine how a strategy for dealing with strangers in one-shot Prisoner's Dilemma's may evolve. However, these strategies in turn have the goal of boosting cooperation between agents to as high a level as possible, where the Friend Selection Strategy society – and with that this extended version of it with grouping by proximity-dependent interaction probability – aims not only to mimic and explain the emergence of peace, but also the collapse thereof.

In the remainder of this paper we will discuss the introduction of proximity between agents to the model and the effects this merits. The problem of small societies is mostly left for another time; some thoughts on the subject will be touched upon in Section 6.

3.1 Implementation

As established in the previous section, human societies are divided into smaller sub-societies and in order to let the model reflect this, the artificial society, too, will be divided into a number of subsocieties, or groups (`n_groups`).

In order to ensure an agent will have a high probability of meeting another agent within his group and a low probability of meeting an agent from a different group, choosing an interaction partner for an agent will be done in two steps: first, it is decided whether the interaction partner will be from the agent's own group or from another group. There is a 0.95 probability of the interaction partner being from one's own group and a 0.05 probability of the partner being from another group. One agent is then chosen randomly from the thus defined group of eligible interaction partners.

The above probabilities may seem unfairly skewed in favour of encounters with one's own group's members. However, most people spend most of their time either at home or in a working or school environment where interacting with strangers is rare. As such, having only one in twenty encounters be with a stranger does not seem unrealistic.

In the original model, `M` is `N-1` and no double encounters with the same agent within a round are allowed. If both these constraints were kept in the altered model, nothing would change, since each agent would be obliged to meet every other agent in each round. For reasons described in Section 2.3, changing `M` is not really an option. Because of this, and also because it is the more realistic option, we will allow agents to meet the same agent more than once in a round.

4 Results

See Figure 4 for graphs showing the behaviour of the original model (left), compared to the extended version with grouping (right), for societies of 60, 90 and 120. Note that `n_groups` was chosen such that the resulting size of groups in each case was around 6. The reason for this setting is that with a groups of approximately six agents the intended results are most visible (for a comparison of results for different group sizes, see Appendix A).

Two effects stand out in the rightmost graphs: firstly, a war-like state has been reintroduced into the model's behavioural repertoire. Secondly, the war-like state we see in these graphs differs from that in Figure 4(a), in the sense that, although the friend ratio is relatively low and stable, it is not by far as low as it is in war-like periods in a society of 20 without grouping.

5 Conclusions

The difference between the war-like states in large, grouped societies and medium-sized, ungrouped societies may be explained by these war-like states each being like a different kind of war in human societies. For a society of 20 in the original model, the war-like state looks like a 'war of all against all', with friend ratios below 0.1. The graphs for large, grouped societies show a war-like state with friend ratios around 0.3, which is likely to be more like civil war than war of all against all, with groups fighting amongst each other, but less amongst themselves. This, at least for large societies, is a much more plausible state than that of war of all against all.

The fact that a group size of six agents is optimal in terms of showing the desired effect of emergence and collapse of peace also seems plausible: even within a group of eight or ten, people have a tendency to interact more with some than with others. Six seems like a plausible maximum size for a group all members of which interact equally often.

As expected, adding an extra touch of reality to Nakai and Muto's model in the form of grouping by interaction probability merits more realistic results. Where the original model performs unrealistically well in terms of levels of cooperation for large societies, the extended model shows signs of not only emergence of peace, but also collapse thereof into a civil-war-like state, providing a much more plausible set of results for high `N`.

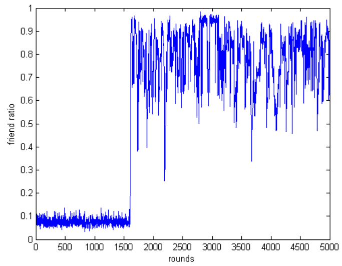

(a) N=20, M=19, original model.

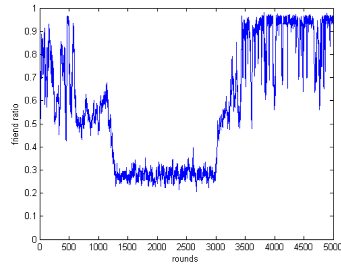

(b) N=20, M=19, n_groups=4.

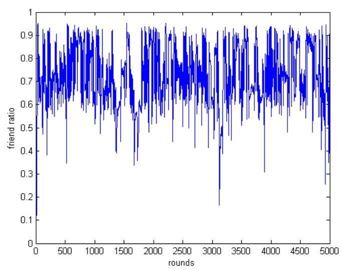

(c) N=60, M=59, original model.

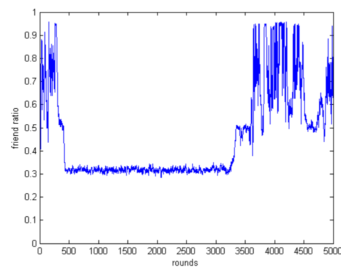

(d) N=60, M=59, n_groups=10.

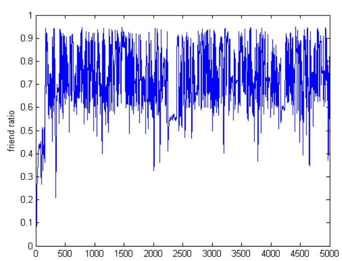

(e) N=90, M=89, original model.

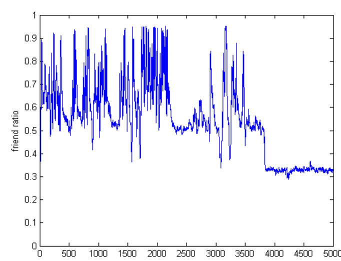

(f) N=90, M=89, n_groups=15.

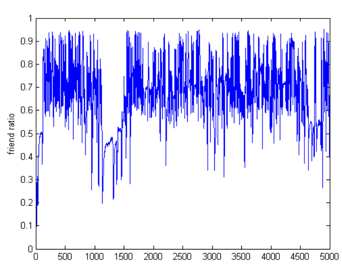

(g) N=120, M=119, original model.

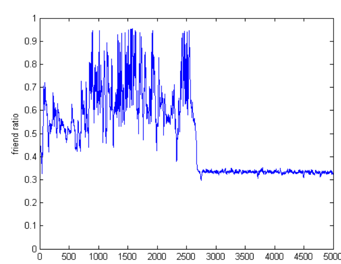

(h) N=120, M=119, n_groups=20.

Figure 4: Comparison of the development of the friend ratio across 5000 rounds for: N=20 and M=19, original model (a) and with grouping (b); N=60 and M=59, original model (c) and with grouping (d); N=90 and M=89, original model (e) and with grouping (f); N=120 and M=119, original model (g) and with grouping (h).

6 Discussion

With the addition of grouping to Nakai and Muto's model, the stage has been set for the addition of more intricate social constructs. In order for the grouping mechanism to still be applicable when dealing with even larger societies, extending the current model with the option of grouping within groups may be useful. Grouping in human societies is usually hierarchical in nature, with the probability of an encounter with some person being proportional to the number of group memberships shared with that person.

However, such extensions should not be attempted until the loose ends of the current model are tied up. Although the results presented in this paper are promising, much work needs to be done. There are clear similarities between the results of these simulations and phenomena known from human societies, but the real question should of course not be how we can mimic reality, but how we can explain it. In depth analysis of mechanisms in the extended model that cause emergence and endurance of the civil-war-like state described in Section 4 will allow comparison of the model's inner workings to reality to see if those, too, are plausible with respect to human societies.

Another issue that needs to be addressed is the awkward results the model gives for small societies. These results suggest that, as was the case with large societies, a mechanism is missing in the current model that is essential to human' success in achieving cooperation in small groups.

Which leaves the question: what can we say about the implications of the model at this point? In the extended version of the model by Nakai and Muto, the existence of groups within the society is simply assumed, since this grouping by interaction probability is present in the reality we seek to mimic and understand. However, in human societies, the tendency of people to form small groups was not a given; it must have evolved at some point.

If we combine this idea with the effect of grouping on cooperation rates as found in the simulations described in this paper, there seems to be a paradox. Large societies exhibit more stable high levels of cooperation without grouping than they do with the addition of grouping. Since higher levels of cooperation mean higher payoffs for the agents, it is clear that from an evolutionary standpoint the division of a society into groups is a bad idea. Yet, grouping did evolve in human societies at some point.

We as modern people have no particular reason not to live together with hundreds, thousands or even millions of others and our modern modes of travel make it almost as easy to encounter our neighbours as to meet someone half-way across the world. Why then are we stuck with this suboptimal tendency to form smaller groups within a large society?

The paradox is only apparent: the evolution of grouping occurred under circumstances very different from current ones. Evolution would have been constrained by many considerations that are no longer relevant today. Small groups may have been more desirable due the fact that a group of only a handful of people could travel faster and was less visible a prey to predators. Also, geographical obstacles like rivers and mountain ridges would reduce the probability of an encounter between persons on opposite sides to almost zero, automatically grouping people by location.

Looking at modern times with this insight in mind, we can see that since grouping in human societies is a legacy from constraints that are no longer relevant, it makes sense that borders between groups, tribes and even nations that were absolute before have in recent history started to fade. And if it is indeed the case, as the simulations in this paper imply, that less grouping makes for more peace, it might be a hopeful suggestion that with the increasing globalization, modern society is becoming increasingly stable in terms of peaceful cooperation.

Acknowledgements

Thanks to Y. Nakai and M. Muto for graciously allowing their code to be used for comparative purposes.

Also, many thanks to my devoted thesis advisors, for their guidance and kindness above and beyond the call of duty.

References

[1] Wim B. G. Liebrand. The effect of social motives, communication and group size on behaviour in an n-person multi-stage mixed-motive game. *European Journal of Social Psychology*, 14:239–264, 1984.

[2] Michael W. Macy and John Skvoretz. The evolution of trust and cooperation between strangers: A computational model. *American Sociological Review*, 63(5):638–660, 1998.

[3] Manfred Milinski, Dirk Semmann, Theo C. M. Bakker, and Hans-Jürgen Krambeck. Cooperation through indirect reciprocity: image scoring or standing strategy? *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 268(1484):2495–2501, December 2001.

[4] Yutaka Nakai and Masayoshi Muto. Emergence and collapse of peace with friend selection strategies. *Journal of Artificial Societies and Social Simulation*, 11(3):6, 2008.

[5] Martin A. Nowak and Karl Sigmund. Evolution of indirect reciprocity by image scoring. *Nature*, 393:573–577, June 1998.

[6] Nobuyuki Takahashi and Rie Mashima. The emergence of indirect reciprocity: Is the standing strategy the answer? Paper presented at the annual meeting of the American Sociological Association, Atlanta Hilton Hotel, Atlanta, August 2003.

[7] Kumar Venkat and Wayne Wakeland. Emergence of networks in distance-constrained trade. Paper presented at the 6th International Conference on Complex Systems (ICCS), Boston, June 2006.

Appendix A: additional graphs

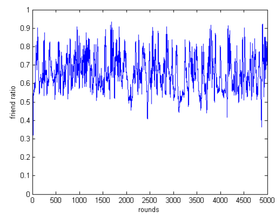

(a) N=60, M=59,n_groups=5

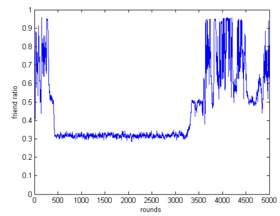

(b) N=60,M=59,n_groups=10

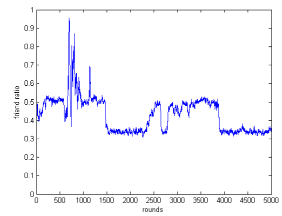

(c) N=60,M=59,n_groups=20

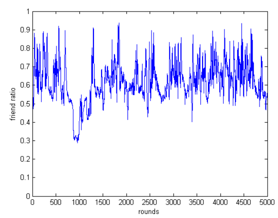

(d) N=90,M=89,n_groups=9

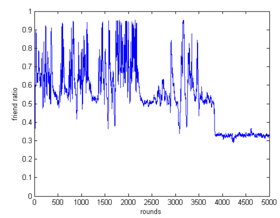

(e) N=90,M=89,n_groups=15

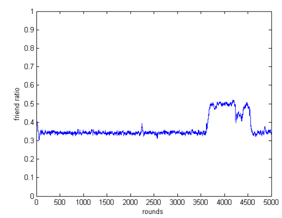

(f) N=90,M=89,n_groups=30

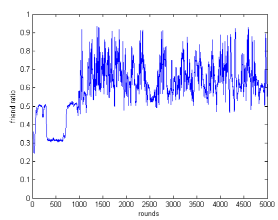

(g) N=120,M=119,n_groups=10

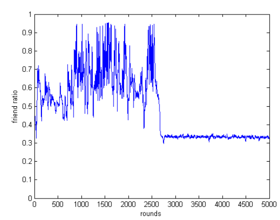

(h) N=120,M=119,n_groups=20

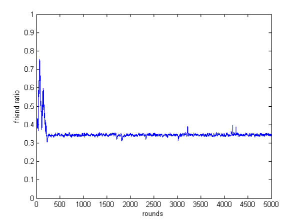

(i) N=120,M=119,n_groups=40

Figure 5: Comparison of the development of the friend ratio across 5000 rounds for: N=60 and M=59, with n_groups=5 (a), n_groups=10 (b) and n_groups=20 (c); N=90 and M=89, with n_groups=9 (d), n_groups=15 (e) and n_groups=30 (f); N=120 and M=119, with n_groups=10 (g), n_groups=20 (h) and n_groups=40 (i);