

Cooperation through the Endogenous Evolution of Social Structure

David Hales¹ and Shade T. Shuttters²

¹ The Open University, Milton Keynes, UK
dave@davidhales.com
<http://davidhales.com>

² Center for Social Dynamics and Complexity and School of Sustainability,
Arizona State University, Tempe, AZ, USA
sshutte@asu.edu
<http://www.public.asu.edu/~sshutte>

Abstract. A number of recent models demonstrate sustained and high levels of cooperation within evolutionary systems supported by the endogenous evolution of social structure. These dynamic social structures co-evolve, under certain conditions, to support a form of group selection in which highly cooperative groups replace less cooperative groups. A necessary condition is that agents are free to move between groups and can create new groups more quickly than existing groups become invaded by defecting agents who do not cooperate.

Keywords: evolution of cooperative, agents, group selection, prisoner's dilemma, cultural evolution.

1 Introduction

Human society is pervaded by groups. Some are formal, such as corporations, educational institutions and social clubs. Others are informal such as youth tribes, collections of old men in a town square who discuss politics and play chess, and more recently various online forums. Some last for a long time, replenishing their membership over many generations and others are ephemeral and fleeting. Some have distinct and clear boundaries others are more diffuse – formed of overlapping networks of relationships. It would appear hard to make sense of human social behaviour without some reference to groups. Indeed, if individuals are asked to describe themselves then it is highly likely that they will refer to the groups that they hold membership of.

1.1 The Danger of Intuition

Although it seems intuitively clear that humans (and even other species) benefit from organising, coordinating and cooperating within groups, understanding how these might evolve from individual behaviour poses major puzzles for both political economy and evolutionary theory. If we start from individual self-interest or selfish replicators then why would individuals behave for the benefit of the group

if they can “get away” with free riding on the group - extracting the benefits of group membership without making a contribution?

One of the dangers of intuition about group processes is that it can lead one to ascribe agency to a group where none exists. Indeed the idea that if a group has interests - in the sense of something that would benefit all members - then rational or evolutionary individuals will behave in a way that promotes those interests has been debunked by careful analysis. Olson’s famous work clearly describes the folly of this intuition when considering rational agents [23] and biologists have also challenged this idea from an evolutionary perspective [39,14].

Given these results there has been a desire to understand the highly groupish phenomena observed in human societies which appears to be altruistic from the point of view of the the individual. One way to tackle this is to attempt to capture the kinds of cultural evolutionary processes that might support learned behaviour that does not conform to self-interest [4].

In this paper we follow this line by discussing recent evolutionary models that rely on a dynamically evolving population structure - that constrains interaction possibilities - that produce remarkable groupish phenomena. With only minimal assumptions these models support the endogenous formation of groups and high levels of cooperation within the groups even when there are significant incentives for individuals to free ride.

1.2 Recent Models

Recent evolutionary models demonstrate novel forms of group selection based on simple learning rules. They function via the spontaneous formation and dissolution of groupings of selfish agents such that altruistic behaviour evolves within their in-groups.

These social dynamics, offering an alternative to rational action theories, demonstrate several notable features of human systems such as seemingly irrational altruism, highly tribal or “groupish” behaviour, and complex dynamics of social structures over time. We overview several classes of such models some based on evolving network structures and others based on different forms of population structures indicating their key features and potential applications.

Recent agent-based computational simulation models have demonstrated how cooperative interactions can be sustained by simple imitation rules that dynamically create simple social structures [28,27,11,13,12,7,20,37,30] . These classes of models implement agents as adaptive imitators that copy the traits of others and, occasionally, adapt (or mutate) them. Although these models bear close comparison with biologically inspired models – they implement simple forms of evolution – the interpretation can be of a minimal cultural, or social, learning process in which traits spread through the population via imitation and new traits emerge via randomised, or other kinds of, adaption.

Often agent-based models represent social structures such as groups, firms or networks of friends, as external and a priori to the agents – see so-called “network reciprocity” results [19]. In the models we discuss in this paper the social structures are endogenous such that agents construct, maintain and adapt

them through on-going behaviour. A subset of traits supports the formation and maintenance of simple social structures [16].

As will be seen, it is the non-equilibrium processes of dynamic formation and dissolution of these structures over time that drives, or incentivises, the agents to behave cooperatively. Yet, as we will show, it is not necessary for the individual agents to prefer socially beneficial structures or outcomes, rather they emerge through a self-organising process based on local information and adaptation.

1.3 When in Rome

In the models we present here agents are assumed to have incomplete information and bounded processing abilities (bounded rationality). Given these relaxed assumptions agents use social learning heuristics (imitation) rather than purely individual learning or calculation. It has been argued by Herbert Simon [32,33] that complex social worlds will often lead to social imitation (or “docility” in Simon’s terminology) because agents do not have the information or cognitive ability to select appropriate behaviours in unique situations. The basic idea is “imitate others who appear to be performing well” which might also be captured by the famous quote of Saint Ambrose: “when in Rome, do as the Romans do”.

The models we present demonstrate that simple imitation heuristics can emerge social behaviours and structures that display highly altruistic in-group behaviour even though this is not part of the individual goals of the agents and, moreover, may appear irrational from the point of view of the individual agents. Agents simply wish to improve their own individual condition (or utility) relative to others and have no explicit conception of in- or out-group. Yet a side effect of their social learning is to sustain group structures that constrain the spread of highly non-social (selfish) or cheating behaviour such as freeloading on the group. They can be compared to models of group selection [40,26].

We could replace the term “side effect” with the term “invisible hand” or “emergent property”. We can draw a loose analogy with Adam Smith’s thoughts on the market [34]. The difference is that there is no recognisable market here but rather a dynamic evolution of social structure that can transform egotistical imitative behaviour into socially beneficial behaviour.

1.4 Tribal Systems

One might term these kinds of models “tribal systems” to indicate the grouping effects and tendency for intra-group homogeneity because individuals joining a group often join this group via the imitation of others who are already a member of the group. We do not use the term “tribal” to signify any relationship between these kinds of models and particular kinds of human societies but rather to indicate the tribal nature of all human organisations i.e. that individuals almost always form cliques, gangs or other groupings that may appear arbitrary and may be highly changeable and ephemeral yet have important effects on inter-agent dynamics and behaviour.

In these kinds of tribal systems individual behaviour cannot be understood from a standpoint of individual rationality or equilibrium analysis but rather only

with reference to the interaction history and group dynamics of the system as a whole. The way an individual behaves depends on their history and relationship to the groups or tribes that they form collectively.

2 Situating the Models

Diverse models of cultural group selection have been proposed from a wide range of disciplines [40]. More recently attempts to formalise them through mathematical and computer based modelling have been made.

We wish to situate the models we will discuss in this chapter with reference to the more traditional game theory [3] approach that assumes agents are rational, in the *homo economicus* sense, and have perfect information, common knowledge and no social structures to constrain interactions.

Our aim in this section is to give the non-modelling expert a sense of the relation between rational action approaches (game theory) and the more bio- and socially- inspired approaches of cultural group selection by presenting a number of broad dimensions over which they differ. It is of course the case that the boundaries between approaches is never as clean or distinct as simple categories suggest, however, to the extent that a caricature can concisely communicate what we consider to be key points that distinguish approaches it can be of value.

Figure 1(a) shows two dimensions along which game theory and cultural group selection approaches may be contrasted. Traditionally game theory models have focused on agents with unbounded rationality (i.e. no limit on computational ability) and complete information (i.e. utility outcomes can be calculated for all given actions). The cultural group selection models presented here focus on highly bounded rationality (agents just copy those with higher utility) and highly limited information (agents can not calculate a priori utility outcomes). The benefit that game theory gains by focusing on the bottom left-hand region are analytic tractability by proving equilibrium points such as Nash Equilibrium for given games [21]. Given incomplete information and bounded rationality it generally becomes more difficult to find tractable solutions and hence (agent-based) computer simulation is often used.

Figure 1(b) shows another two dimensions, learning and utility, along which a broad distinction can be made. Game theory models tend to focus on individual utility maximisation and action or strategy selection (a kind of learning) at the individual level via deduction (bottom-left). Cultural group selection focuses on social learning based on imitation in combination with rare innovation events (comparable to mutation in biological models). The emergent result is increase in social utility even though the agents themselves use a heuristic based on trying to improve their own individual utility. Hence cultural group selection could also be placed in the bottom-right quadrant.

Figure 1(c) shows another two dimensions, interaction and social structure, that distinguish the cultural group selection models and game theory. The cultural group selection models presented here represent interactions within dynamic social structures whereas game theory has tended towards static mean

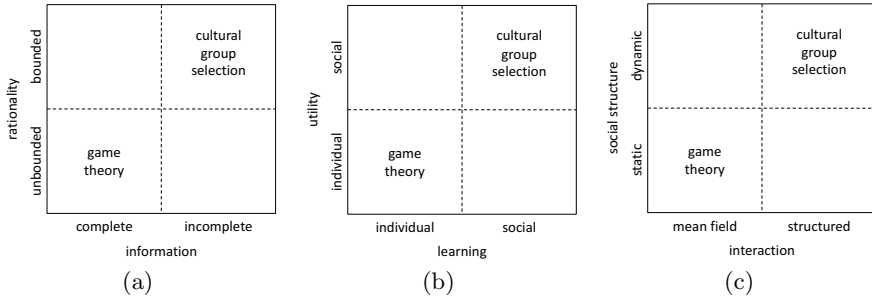


Fig. 1. Six qualitative dimensions distinguishing traditional game theory models and many cultural group selection models

field structures, by which we mean that game interactions are often assumed to occur stochastically, with equal probability, between agents over time. In the cultural group selection models (as will be seen later) a key aspect that drives the evolution of cooperation and increases in social utility is the dynamic formation of in-groups of agents that interact together exclusively, excluding interactions with the out-group.

3 Three Kinds of Models

Historically group selection has been seen as controversial within both biological and social sciences due to the difficulty in advancing a plausible theory and the inability of identifying such processes empirically in the field. Also certain kinds of non-formalised group selection approaches were exposed as naive by biologists. However these objections have been challenged due to recent advances in the area due to extensive use of computational (often agent-based) modelling and a theoretical shift that accepts that selection operating at the individual level can, under broad conditions, emerge group level selection at a higher level. The historical debate from a group selectionist perspective is well covered elsewhere [40].

We will not retread old ground here but will concentrate on presenting a specific class of group selection models that have recently emerged in the literature. These models may be interpreted as cultural evolutionary models in which imitation allows traits to move horizontally. We do not concern ourselves here with the biological interpretation of such models but rather the cultural interpretation.

Group selection relies on the dynamic formation and dissolution of groups. Over time individual entities may change groups by moving to those that offer better individual performance. Interaction between entities that determine performance is mainly restricted to those sharing the same group. Essentially then, in a nutshell, groups that support high performance for the individuals that comprise them grow and prosper whereas exploitative or dysfunctional groups

dissolve as individuals move away. Hence functional groups, in terms of satisfying individual goals, are selected over time. It should be noted, as will be seen in the network rewiring models discussed below, that clearly defined group boundaries are not required so long as interactions between agents are sufficiently localised through an emergent structure [41].

Key aspects that define different forms of group selection are: How group boundaries are formed (or interaction is localised); the nature of the interactions between entities within each group; the way that each entity calculates individual performance (or utility) and how entities migrate between groups.

3.1 Pro-social or Selfish Behaviour

In almost all proposed social and biological models of group selection, in order to test if group selection is stronger than individual selection, populations are composed of individuals that can take one of two kinds of social behaviour (or strategy). They can either act pro-socially, for the good of their group, or they can act selfishly for their own individual benefit at the expense of the group. This captures a form of commons tragedy [15].

Often this is formalised as a Prisoners Dilemma (PD) or a donation game in which individuals receive fitness payoffs based on the composition of their group. In either case there is a fitness cost c that a pro-social individual incurs and an associated fitness benefit b that individuals within a group gain. A group containing only pro-social individuals will lead each to gain a fitness of $b - c$ which will be positive assuming $b - c > 0$. However, a group containing only selfish individuals will lead each to obtain no additional fitness benefit or cost. But a selfish individual within a group of pro-socials will gain the highest fitness benefit. In this case the selfish individual will gain b but the rest will gain less than $b - c$. Hence it is always in an individual's interests (to maximise fitness) to behave selfishly. In an evolutionary scenario in which the entire population interacts within a single group then selfish behaviour will tend to be selected because this increases individual fitness. This ultimately leads to an entire population of selfish individuals and a suboptimal situation in which no individual gains any fitness. This is the Nash Equilibrium [21] and an Evolutionary Stable Strategy for such a system [35].

There have been various models of cooperation and pro-social behaviour based on reciprocity using iterated strategies within the PD [1]. However, we are interested in models which do not require reciprocity since these are more generally applicable. In many situations, such as large-scale human systems or distributed computer systems, repeated interactions may be rare or hard to implement due to large population sizes (of the order of millions) or cheating behaviour that allow individuals (or computer nodes) to fake new identities.

3.2 Tag Models

In [11] a tag model of cooperation was proposed which selected for pro-social groups. It models populations of evolving agents that form groups with other

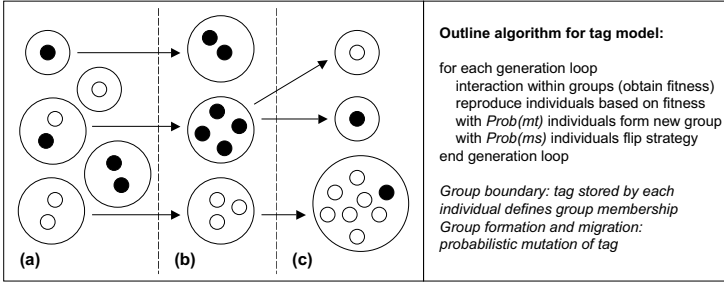


Fig. 2. Schematic of the evolution of groups in the tag model. Three generations (a-c) are shown. White individuals are pro-social, black are selfish. Individuals sharing the same tag are shown clustered and bounded by large circles. Arrows indicate group lineage. Migration between groups is not shown.

agents who share an initially arbitrary tag or social marker. The tag approach was discussed by Holland [16] and developed by Riolo [28,27]. The tag is often interpreted as an observable social label (e.g. style of dress, accent etc.) and can be seen as a group membership marker. It can take any mutable form in a model (e.g. integer or bitstring). The strategies of the agents evolve, as do the tags themselves. Interestingly this very simple scheme structures the population into a dynamic set of tag-groups and selects for pro-social behaviour over a wide range of conditions. Figure 2 shows a schematic diagram of tag-group evolution and an outline algorithm that generates it.

In general it was found that pro-social behaviour was selected when $b > c$ and $mt \gg ms$, where mt is the mutation rate applied to the tag and ms is the mutation rate applied to the strategy. In this model groups emerge from the evolution of the tags. Group splitting is a side effect of mutation applied to a tag during reproduction. Figure 3 shows typical output from a tag simulation model visualising a portion of the tag space. Each line on the vertical axis represents a unique tag value (i.e. a possible group). Groups composed of all cooperative agents are shown in light grey, mixed groups of cooperators and defectors (mixed) are dark grey and groups composed of all defectors are black. The size of each group is not shown. Notice that over time new groups form, persist for some period and then become invaded by some defectors and quickly die. This cyclical process continues persistently. Since mixed groups die quickly the number of defectors in the entire population at any time instant are small. Hence at any given time the highest scoring agents are those who are defecting but this is not a sustainable strategy because they destroy the groups they are situated in.

A subsequent tag model [27] produced similar results with a different tag and cooperation structure. Tags are defined as single real values $[0..1]$ and cooperation is achieved through a donation game in which a related tolerance value $[0..1]$ specifies a range around the tag value to which agents will donate if they find other agents with tags within this range. The tolerance and tag are traits that evolve based on fitness (payoff). This provides the potential for overlapping group

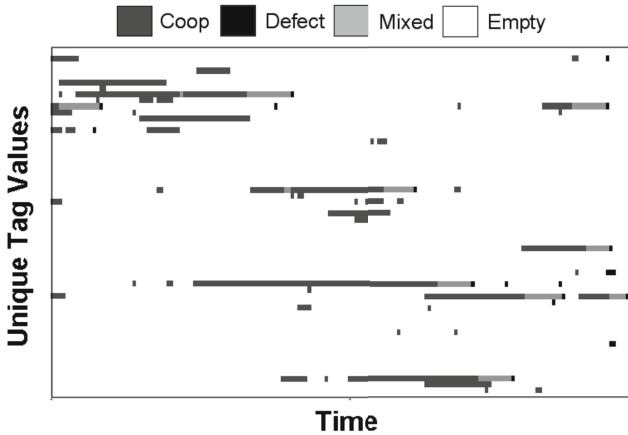


Fig. 3. Evolution of tag groups over time. Here 200 cycles from a single typical simulation run are shown for a portion of the tag space. Cooperative groups come into existence, persist for some time, and then die.

boundaries. Also the model works through random sampling of the population for game partners rather than only sampling within a group. However It has been argued that this approach cannot be applied to pro-sociality in general because it does not allow for fully selfish behaviour between identically tagged individuals [29,6]. Put simply, those agents which share the same tag must cooperate and can not evolve a way to defect since there is no separate trait defining the game strategy.

Recent work by Traulsen and Nowak examined the tag process both analytically and in simulation deriving the necessary conditions under which high cooperation can be sustained [38]. They derived an elegant mathematical description of the stochastic evolutionary dynamics of tag-based cooperation in populations of finite size. They found that in a cultural model with equal mutation rates between all possible phenotypes (tags and strategies), the crucial condition for high cooperation is $b/c > (K + 1)/(K - 1)$, where K is the number of tags. A larger number of tags requires a smaller benefit-to-cost ratio. In the limit of many different tags, the condition for cooperators to have a higher average abundance than defectors becomes $b > c$. Hence this result indicates that it is not necessary to have higher mutation rates applied to tags per se but rather to have enough tag space relative to the benefit-to-cost ratio.

Finally, more recent work rigorously replicates and compares several existing tag models from the literature and introduces the notion of weak and strong cheating [31]. It was found that the way that agents implement their cheating mechanism can have dramatic effects of the results obtained from some previous tag models. This has implications for both cultural and biological interpretations of the models.

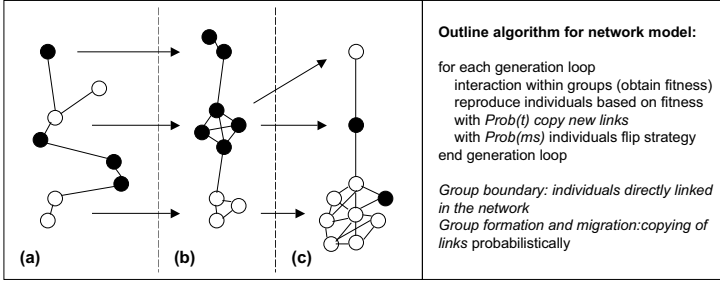


Fig. 4. Schematic of the evolution of groups (cliques) in the network-rewiring model. Three generations (a-c) are shown. White individuals are pro-social, black are selfish. Arrows indicate group lineage. Notice the similarity to the tag model in figure 2.

3.3 Network Rewiring Models

Network rewiring models for group selection have been proposed with direct application to peer-to-peer (P2P) protocol design [7,8]. In these models, which were adapted from the tag model described above, individuals are represented as nodes on a graph. Group membership is defined by the topology of the graph. Nodes directly connected are considered to be within the same group. Each node stores the links that define its neighbours. Nodes evolve by copying both the strategies and links (with probability t) of other nodes in the population with higher utility than themselves. Using this simple learning rule the topology and strategies evolve promoting pro-social behaviour and structuring the population into dynamic arrangements of disconnected clusters (where $t = 1$) or small-world topologies (where $0.5 < t < 1$). Group splitting involves nodes disconnecting from all their current neighbours and reconnecting to a single randomly chosen neighbour with low probability mt . As with the tag model pro-social behaviour is selected when $b > c$ and $mt \gg ms$, where ms is the probability of nodes spontaneously changing strategies. Figure 4 shows a schematic of network evolution (groups emerge as cliques within the network) and an outline algorithm that implements it. Figure 5 shows output from a typical simulation run where $t = 1$. Four snapshots of the network structure are shown at different key time cycles. As can be seen the system self-organises into a dynamic ecology of cooperative components. The evolution of the network is therefore the outcome of a coevolutionary process between nodes which effects both the strategy and the network structure based on local node information. This can be compared to global forms of network optimisation, that interestingly, also implement localised mutation and rewiring [18].

In this model we have translated dynamics and properties similar to the tag model into a graph. In [8,7] the rewiring approach was suggested as a possible protocol that could be applied in an online peer-to-peer system. In [12] the same fundamental rewiring protocol was applied to a scenario requiring nodes to adopt specialised roles or skills within their groups, not just pro-social behaviour alone,

to maximise social benefit. Also another network rewire model shows similar properties [30].

Interestingly it has also been shown recently [22] in a similar graph model tested over *fixed topologies* (e.g. small-world, random, lattice, scale-free) that under a simple evolutionary learning rule pro-social behaviour can be sustained in some limited situations if $b/c > k$, where k is the average number neighbours over all nodes (the average degree of the graph). This implies that if certain topologies can be imposed then pro-social behaviour can be sustained without rewiring of the topology dynamically. Although analysis of this model is at an early stage it would appear that groups form via clusters of pro-social strategies forming and migrating over the graph via nodes learning from neighbours. This reinforces the insight that localisation of interaction rather than strict group boundaries is sufficient to produce cooperation [41].

3.4 Group-Splitting Model

In [37] a group selection model is given that sustains pro-social behaviour if the population is partitioned into m groups of maximum size n so long as $b/c > 1 + n/m$. In this model group structure, splitting and extinction is assumed *a priori* and mediated by exogenous parameters. Splitting is accomplished by explicitly limiting group size to n , when a group grows through reproduction beyond n it is split with (high) probability q into two groups by probabilistically reallocating each individual to one of the new groups. By exogenously controlling n and m a detailed analysis of the model was derived such that the cost / benefit condition is shown to be *necessary* rather than just sufficient. The model also allows for some migration of individuals between groups outside of the splitting process. Significantly, the group splitting model can potentially be applied recursively to give multilevel selection – groups of groups etc. However, this requires explicit splitting and reallocation mechanisms at each higher level. Figure 6 shows a schematic of group-splitting evolution and an outline algorithm that implements it. We include this example here for comparison but it should be noted that it is not a fully endogenous group formation model.

4 Discussion

What do the kinds of models presented in this paper actually tell us? They do not aim to capture any particular phenomena but rather to demonstrate certain kinds of possible process under given assumptions. Hence they should be seen as an aid to intuition. Models such as these deduce logically, through a computer program, outcomes of given assumptions. As we have discussed previously unchecked intuition can be unreliable when applied to complex systems such as these. They can be viewed, then, as thought experiments with the aid of a computer.

Given an understanding of these possible processes it may then be possible to identify real world phenomena that may be evidencing them. This is a much




Fig. 5. Evolution of network structure in a network rewiring model. From an initially random topology (not shown) composed of all nodes playing the defect strategy (dark shaded nodes), components quickly evolve, still containing all defect nodes (a). Then a large cooperative component emerges in which all nodes cooperate (b). Subsequently the large component begins to break apart as defect nodes invade the large cooperative component and make it less desirable for cooperative nodes (c). Finally an ecology of cooperative components dynamically persists as new components form and old components die (d). Note: the cooperative status of a node is indicated by a light shade.

harder task requiring extensive empirical work. However, the models can suggest the kinds of data that would be required and the kinds of hypotheses that can be tested against that data.

4.1 Institutional Economics

The idea of an evolutionary and cyclical group competition process goes at least as far back as Ibn Khaldun writing in the 14th Century [17]. Khaldun, who is considered one of the founding fathers of sociology and economics, developed a

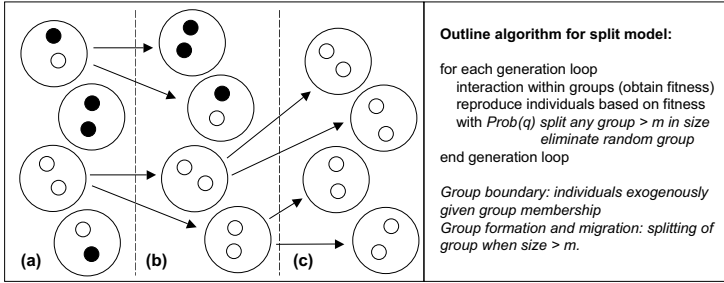


Fig. 6. Schematic of the evolution of groups in the group-splitting model. Three generations (a-c) are shown. White individuals are pro-social, black are selfish. Individuals sharing the same group are shown clustered and bounded by large circles. Arrows indicate group lineage. Migration between groups is not shown.

concept of “asabiyyah” (roughly meaning social cohesion) to explain how groups arise, become powerful and ultimately are conquered by other groups. We do not claim that the models presented here capture the sophistication of Kaldun’s model - indeed he discusses groups with highly sophisticated internal power structures with roles for a ruling elite, artisans and the like - however, we believe, that some of the essential intuitive dynamics are captured in them in a minimal and formal way.

Essentially we have a model of perpetual social conflict – an arms race – in which groups die from within after going through several stages in which social cohesion is successively weakened.

More generally we argue that kinds of models presented here can be viewed as an initial (and minimal) way to establish a link between a line of work called “institutional economics” and the more individualistic evolutionary approaches which have been heavily influenced by biology. The aim is to find the kinds of simple mechanisms that can support institutional-like dynamics. We might view the groups in our models as highly abstract and simple “photo-institutions”. We have shown elsewhere that similar models can support limited forms of specialisation within groups – where different agents perform different functions for the good of the group [12].

Ostrom’s famous work describes in detail how groups with known boundaries – and other “design principles” – can self-organise their own solutions to Common Pool Resource dilemmas without market mechanisms or central control from government [24]. We believe that the kinds of mechanism evidenced within the models presented could be sufficient to support the evolution of more sophisticated cooperation mechanisms rather than simply cooperate or defect. i.e. policing and punishment strategies.

4.2 Applications

The models presented here could potentially have applications in both understanding real existing social systems and engineering new tools that support new kinds of social systems – particularly in online communities. Increasingly online Web2.0 and other communities allow for the tracking and measurement of the dynamics of groups overtime [25]. Massive clean datasets can now be collected and it might be possible for (versions of) the models presented here to be calibrated and validated (or invalidated). However, to our knowledge, this has so far not been achieved for full-blown cultural group selection, however recent highly detailed work shows great promise in this area [36].

In addition, as has already been discussed above, peer-to-peer (P2P) systems composed of millions of online nodes (running client software on users machines) could benefit from the application of group selection techniques by applying them directly to the algorithms (or protocols) used by nodes to self-organise productive services for users. There has already been some deployed on-going experiments with a P2P file-sharing client called “Tribler”¹ which applies a so-called “group selection design pattern” to self-organise cooperative communities for media content sharing [10,9].

These two kinds of application of the models are not independent because by increasing our understanding of productive human social processes we can automate aspects of those processes into computer algorithms to increase their speed and reach - one might consider the success of online social networking as an example of this.

5 Conclusion

The models presented here show how simple agent heuristics based on imitation directed towards individual improvement of utility can lead to behaviour in which agents behave “as if” there is a motivating force which is higher than self-interest: the interests of the group or tribe. This higher force does not need to be built-in to agents but rather emerges through time and interactions - a historical process. The formation of social structures, overtime, creates conditions that favour pro-social behaviour. Agents receive utility by interacting in tribes (simple social structures that localise game interactions). Tribes that cannot offer the agent a good “utility deal” relative to other tribes will disband as agents “vote with their feet” by joining other better tribes based on their individual utility assessment. Of course movement between tribes, here, is not interpreted as a physical relocation but rather a social and mental one. By copying the traits of others who have higher utility the appropriate social structures emerge. Increasingly in electronic and virtual communities the cost of such movement is converging towards zero or very low individual cost. It could be conjectured it is this low cost, and consequent freedom from geographical and organisational constraints,

¹ See: <http://tribler.org>

which is a major factor in the recent success of online communities, virtual social networks and other peer-production communities such as Wikipedia [2].

However, this process would not preclude agents with explicit group level utility preferences i.e. incorporating “social rationality” functions or such like. Such agents could potentially improve general welfare through a modicum of explicit planning and encouragement of pro-social group formation. The models presented here rely on random trial and error to find cooperative pro-social “seeds” which then are selected and grow via the evolutionary process as other agents join the seed. We speculate that an agent with a correctly aligned internal model of what would make a successful seed could pro-actively recruit others from the population. However, this introduces issues such as explicit recruitment processes, explicit internal social models and, potentially, transferable utility. Here we begin to see the formation of something that resembles a market or an institutional strategy. In this context the models we have presented could be seen as “pre-market” mechanisms or “proto-institutions” in which value is not separated from the social structures that produce it because it cannot be stored, accumulated, transferred or controlled.

The models we have presented mainly focus on social dilemma scenarios - situations in which individuals can improve their own utility at the expense of the group or tribe they interact with. Often the application of the market in these situations does not resolve the dilemma in a socially equitable way (i.e. does not lead to cooperation) but rather can incentivise non-cooperation. This is such a serious issue that game theory explicitly addresses it within the emerging area of Mechanism Design [5]. However, often these models rely on standard Rational Action assumptions and a high degree of central control that enforce the rules of the game.

As previously discussed, the kinds of models presented here might be viewed as the first steps on a bridge between simple evolutionary – often biologically inspired – cultural models and some of the puzzles and finding of institutional economics. Rather than starting from the assumption of existing and complex institutions we attempt to grow them from individuals following simple copying heuristics.

References

1. Axelrod, R.: The evolution of cooperation. Basic Books, NY (1984)
2. Benkler, Y.: The Wealth of Networks: How Social Production Transforms Markets and Freedom. Yale University Press (2006)
3. Binmore, K.: Game Theory and the Social Contract. Playing Fair, vol. 1. The MIT Press, Cambridge (1994)
4. Boyd, R., Richerson, P.: Culture and cooperation. In: Beyond Self-Interest, Chicago University Press, Chicago (1990)
5. Dash, R., Jennings, N., Parkes, D.: Computational-Mechanism Design: A Call to Arms. IEEE Intelligent Systems, 40–47 (November 2003)
6. Edmonds, B., Hales, D.: Replication, Replication and Replication - Some Hard Lessons from Model Alignment. Special Issue on Model-2-Model Comparison, Journal of Artificial Societies and Social Simulation 6(4) (2003)

7. Hales, D., Arteconi, S.: SLACER: A Self-Organizing Protocol for Coordination in P2P Networks. *IEEE Intelligent Systems* 21(2), 29–35 (2006)
8. Hales, D., Edmonds, B.: Applying a socially-inspired technique (tags) to improve cooperation in P2P Networks. *IEEE Transactions in Systems, Man and Cybernetics - Part A: Systems and Humans* 35(3), 385–395 (2005)
9. Hales, D., et al.: QLECTIVES Report D2.1.1: Candidate theory models for cooperation algorithms (2010), <http://davidhales.com/qlectives/D2.1.1.pdf>
10. Hales, D., Arteconi, S., Marcozzi, A., Chao, I.: Towards a group selection design pattern. In: Meyer, F. (ed.) *Proceedings of the Final Workshop on The European Integrated Project “Dynamically Evolving, Large Scale Information Systems (DELIS)”*, Barcelona, February 27–28, Heinz Nixdorf Institute, University of Paderborn, Band 222 (2008)
11. Hales, D.: Cooperation without Memory or Space: Tags, Groups and the Prisoner’s Dilemma. In: Moss, S., Davidsson, P. (eds.) *MABS 2000. LNCS (LNAI)*, vol. 1979, pp. 157–166. Springer, Heidelberg (2001)
12. Hales, D.: Emergent Group-Level Selection in a Peer-to-Peer Network. *Complexus* 3(1), 108–118 (2006)
13. Hales, D.: From Selfish Nodes to Cooperative Networks Emergent Link-based Incentives in Peer-to-Peer Networks. In: *Proceedings of the Fourth IEEE International Conference on Peer-to-Peer Computing (p2p 2004)*. IEEE Computer Society Press (2004)
14. Hamilton, D.W.: Geometry for the selfish herd. *Journal of Theoretical Biology* 31, 295–311 (1971)
15. Hardin, G.: The tragedy of the commons. *Science* 162, 1243–1248 (1968)
16. Holland, J.: The Effect of Labels (Tags) on Social Interactions. Santa Fe Institute Working Paper 93-10-064. Santa Fe, NM (1993)
17. Khaldun, I.: *The Muqaddimah – An Introduction to History*, 9th edn. Princeton University Press (1989)
18. Kaluza, P., Vingron, M., Mikhailov, A.S.: Self-correcting networks: Function, robustness, and motif distributions in biological signal processing. *Chaos* 18, 1–17 (2008)
19. Konno, T.: A condition for cooperation in a game on complex networks. *Journal of Theoretical Biology* 269(1), 224–233 (2011)
20. Marcozzi, A., Hales, D.: Emergent Social Rationality in a Peer-to-Peer System. *Advances in Complex Systems (ACS)* 11(4), 581–595 (2008)
21. Nash, J.F.: Equilibrium Points in N-Person Games. *Proc. Natl. Acad. Sci. USA* 36, 48–49 (1950)
22. Ohtsuki, H., et al.: A simple rule for the evolution of cooperation on graphs and social networks. *Nature* 441(25), 502–505 (2006)
23. Olson, M.: *The Logic of Collective Action*. Harvard University Press, Cambridge (1970)
24. Ostrom, E.: *Governing the Commons – The evolution of institutions for collective action*. Cambridge University Press, UK (1990)
25. Palla, G., Barabasi, A.-L., Vicsek, T.: Quantifying social group evolution. *Nature* 446, 664–667 (2007)
26. Richerson, P., Boyd, R.: *The Biology of Commitment to Groups: A Tribal Instincts Hypothesis in Evolution and the Capacity for Commitment*. Russell Sage Press, New York (2001)
27. Riolo, R., Cohen, M.D., Axelrod, R.: Cooperation without Reciprocity. *Nature* 414, 441–443 (2001)

28. Riolo, R.: The Effects of Tag-Mediated Selection of Partners in Evolving Populations Playing the Iterated Prisoners Dilemma. Santa Fe Institute Working Paper 97-02-016. Santa Fe, NM (1997)
29. Roberts, G., Sherratt, N.T.: Similarity does not breed cooperation. *Nature* 418, 449–500 (2002)
30. Santos, F.C., Pacheco, J.M., Lenaerts, T.: Cooperation prevails when individuals adjust their social ties. *PLoS Comput Biol.* 2(10), e140 (2006)
31. Shuttters, S.T., Hales, D.: Tag-mediated altruism is dependent on how cheaters are defined. *Journal of Artificial Societies and Social Simulation* (expected pub. date November 2012) (in press)
32. Simon, H.A.: A Mechanism for Social Selection and Successful Altruism. *Science* 250 (1990)
33. Simon, H.A.: *Models of Bounded Rationality, Empirically Grounded Economic Reason*, vol. III. MIT Press, Cambridge (1997)
34. Smith, A.: *An inquiry into the nature and causes of the wealth of nations*, 1st edn. William Clowes and Sons, England (1776, 1836)
35. Smith, J.M.: *Evolution and the Theory of Games* (1982)
36. Szell, M., Thurner: Measuring social dynamics in a massive multiplayer online game. *Soc. Networks* 32, 313–329 (2010)
37. Traulsen, A., Nowak, M.A.: Evolution of cooperation by multilevel selection. *Proceedings of the National Academy of Sciences* 130(29), 10952–10955 (2006)
38. Traulsen, A., Nowak, M.A.: Chromodynamics of Cooperation in Finite Populations. *Plos One* 2(3), e270 (2007)
39. Williams, G.C.: *Adaptation and Natural Selection: A Critique of Some Current Evolutionary Thought*. Princeton University Press, Princeton (1966)
40. Wilson, D.S., Sober, E.: Re-introducing group selection to the human behavioural sciences. *Behavioural and Brain Sciences* 17(4), 585–654 (1994)
41. Wilson, D.S., Wilson, E.O.: Rethinking the theoretical foundation of sociobiology. *Q. Rev. Biol.* 82, 327–348 (2007)