10 Rationality Meets the Tribe

Recent Models of Cultural Group Selection[1]

David Hales

INTRODUCTION

Recent agent-based computational simulation models have demonstrated how cooperative interactions can be sustained by simple cultural learning rules that dynamically create simple social structures (Riolo 1997, 2001; Hales 2000, 2006; Hales and Areteconi 2006; Marcozzi and Hales 2008; Traulsen and Nowak 2006). 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 emerging via randomized, or other kinds of, adaptation.

Often agent-based models represent social structures such as groups, firms or networks of friends, as external and *a priori* to the agents. In the models we discuss in this chapter, however, the social structures are endogenous such that agents construct, maintain and adapt them through ongoing behavior. A subset of traits supports the formation and maintenance of simple social structures.

As will be seen, it is the dynamic formation and dissolution of these structures over time that drive, or incentivize, 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-organizing process based on local information and adaptation criteria.

A major advantage of the agent-based approach is that the strict simplifying assumptions of rational action theory can be relaxed because models do not need to be designed with deductive tractability in mind but rather can be explored through computational simulation.

This is particularly useful for exploring complex models in which agents adapt and learn over time without necessarily converging on any equilibrium or where many equilibria are possible but it is not clear which would be selected.

This relaxation of rational action assumptions is possible due to the technical innovation of agent-based modeling and large-scale simulation platforms that allow researchers to empirically experiment with their models by performing many exploratory simulation runs, observing alternative time series (or histories) and changing model parameters (or assumptions). This means that modelers can quickly answer "what if" type questions and assess the impact of broad changes in the behavioral assumptions on which the models are based. The researcher does not have to make an *a priori* commitment to restrictive assumptions. They can be changed (and often are changed) as a result of model exploration (Doran 1998).

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 (Simon 1990, 1997) 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 behaviors in unique situations. The basic idea is "imitate others who appear to be performing well".

The models we present demonstrate that from simple imitation heuristics can emerge social behaviors and structures that display highly altruistic in-group behavior *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 behavior such as free-riding on the group.

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 (A. Smith 1836). The difference is that there is no recognizable market here but rather a dynamic evolution of social structure that can transform egotistical imitative behavior into socially beneficial behavior.

We 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 certain kinds of human societies but rather to indicate the "tribal" nature of all human organizations, 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 behavior.

In these kinds of tribal systems individual behavior cannot be understood from a standpoint of individual rationality without 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.

SITUATING THE MODELS

Diverse models of cultural group selection have been proposed from a wide range of disciplines (Wilson and Sober 1994). More recently attempts to formalize them through mathematical and computer-based modeling have been made.

We wish to situate the models we will discuss in this chapter with reference to the more traditional game theory (Binmore 1994) 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-modeling 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 10.1 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 cannot calculate *a priori* utility outcomes). The benefit that game theory gains by focusing on the bottom left-hand region is analytic tractability by proving equilibrium points such as Nash equilibrium for given games. 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 10.2 shows another two dimensions, learning and utility, along which a broad distinction can be made. Game theory models tend to focus on individual utility maximization 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

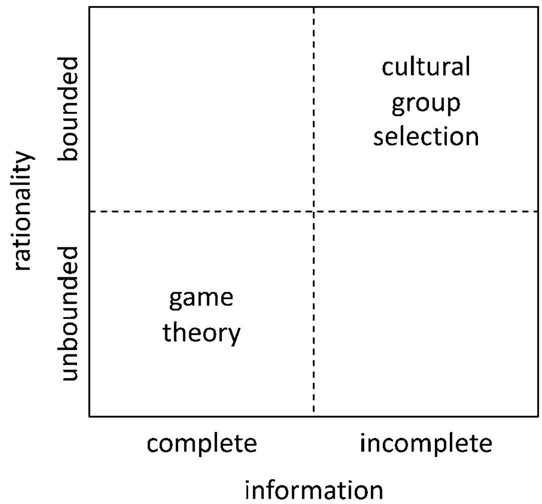

Figure 10.1 Traditionally, game theory models have focused on agents with unbounded rationality and complete information. The cultural group selection models presented here focus on highly bounded rationality and incomplete information.

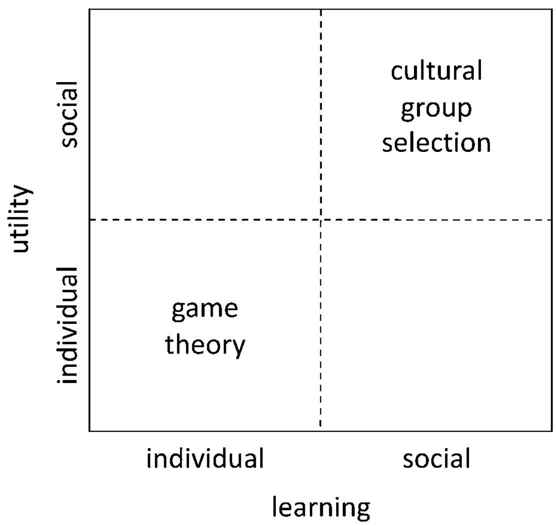

Figure 10.2 Cultural group selection models also differ from the traditional game theory approach in their focus on social learning and (often emergent) social utility over individual utility.

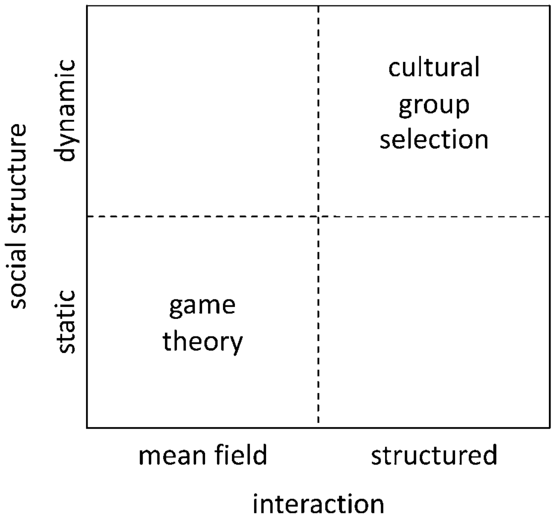

Figure 10.3 The cultural group selection models represent interactions within dynamic social structures whereas game theory has tended towards static "mean field" structures.

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 10.3 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 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".

RECENT CULTURAL GROUP SELECTION 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 naïve non-formalized group selection approaches were exposed as logically incoherent by biologists. However these objections have been challenged due to recent advances in the area as a result of extensive use of computational (often agent-based) modeling and a theoretical shift that accepts that from 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 by Wilson and Sober (1994).

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.

Key aspects that define different forms of group selection are: how group boundaries are formed, 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.

The "success" of any group selection model is judged by how well the system self-organizes towards achieving a collective goal—whatever that may be. Often this will be maximizing the sum of individual utility but could involve other measures such as equality or fairness for example.

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 behavior (or strategy). Either they can 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 (Hardin 1968).

Often this is formalized as a prisoner's 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 utility cost c that a pro-social individual incurs and an associated utility benefit b that individuals within a group gain. A group containing only pro-social individuals will lead each to gain a utility of $b—c$. However, a group containing only selfish individuals will lead each to obtain a utility of zero. But a selfish individual within a group of pro-socials will gain highest utility. In this case the selfish individual will gain b but the rest will gain less than $b—c$. Given that b and c are positive then it is always in an individual's interests (to maximize utility) to behave selfishly. In an evolutionary scenario in which the entire population interacts within a single group then selfish behavior will tend to be selected because this increases utility. This ultimately leads to an entire population of selfish individuals and a suboptimal average population level fitness of zero. This is the Nash equilibrium (Nash 1950) and an evolutionary stable strategy for such a system (J.M. Smith 1982).

There have been various models of cooperation and pro-social behavior based on reciprocity using iterated strategies within the PD (Axelrod 1984; Riolo 1997). 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 (on the order of millions) or cheating behavior that allow individuals (or computer nodes) to fake new identities.

Tag Model

In Hales (2000) a "tag" model of cooperation was proposed which selected for pro-social groups. It models populations of evolving agents that form groups with other agents who share an initially arbitrary tag or social marker. The tag approach was originally proposed by Holland (1993) and developed by Riolo (1997, 2001). The tag is often interpreted as an observable social label (e.g., style of dress or accent etc.) and can be seen as a group membership marker. It can take any mutable form in an agent-based model (e.g., integer or bit string). The strategies of the agents evolve, as do the tags themselves, through agents imitating others obtaining higher utility than themselves. Interestingly this very simple scheme structures the population into a dynamic set of tag groups and selects for pro-social behavior over a wide range of conditions. Figure 10.4 shows a schematic diagram of tag group evolution and an outline algorithm that generates it.

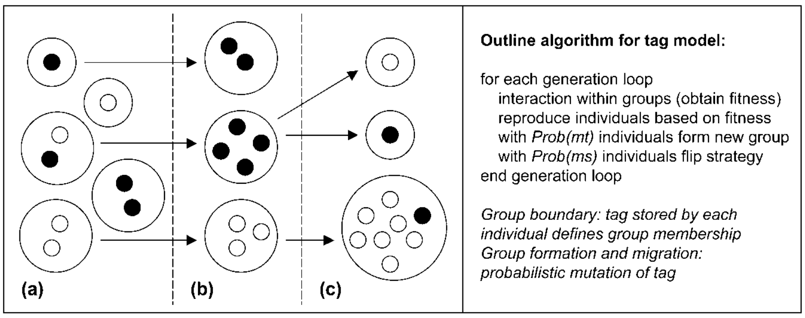

Figure 10.4 Schematic of the evolution of groups in the tag model. Three generations (a–c) are shown. White individuals are pro-social; black are selfish. Individuals sharing the same tag are shown clustered and bounded by large circles. Arrows indicate group linage. Migration between groups is not shown.

In general it was found that pro-social behavior was selected when *b* > *c* and *mt* >> *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. A subsequent tag model (Riolo 2001) produced similar results although it cannot be applied to pro-sociality in general because it does not allow for fully selfish behavior of identically tagged individuals (Roberts and Sherrat 2002).

Network-Rewiring Models

Network-rewiring models for group selection have been proposed with direct application to peer-to-peer (P2P) protocol design and biological systems (Hales 2004, 2006; Santos *et al.* 2006). In these models, which were adapted from the tag model described earlier, 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 neighbors. 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 behavior 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 neighbors and reconnecting to a single randomly chosen neighbor with low probability *mt*. As with the tag model pro-social behavior is selected when *b* > *c* and *mt* >> *ms,* where *ms* is the probability of nodes spontaneously changing strategies. Figure 10.5 shows a schematic of network evolution (groups emerge as cliques within the network) and an outline algorithm that implements it.

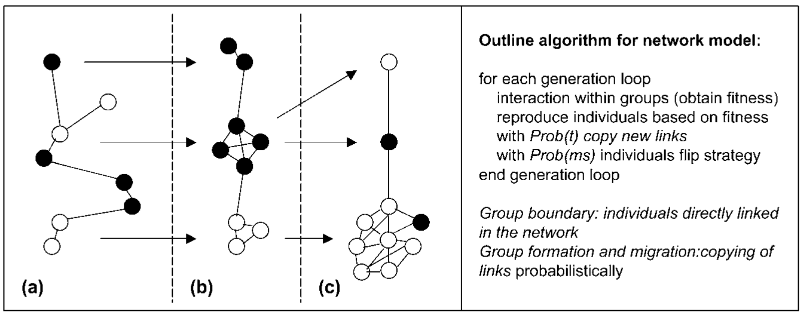

Figure 10.5 Schematic of the evolution of groups (cliques) in the network-rewiring model. Three generations (a–c) are shown. White individuals are pro-social; black are selfish. Arrows indicate group linage. Notice the similarity to the tag model in Figure 10.4.

In this model we see dynamics and properties similar to the tag model but in an evolving graph. This is interesting because social networks can be viewed as graphs. In addition from a computer science perspective graphs can represent P2P networks. In Hales (2006) the same rewiring approach was applied to a scenario requiring nodes to adopt specialized roles or skills within their groups, not just pro-social behavior alone, to maximize social benefit. This indicates that the same kind of group selective process can support the emergence of in-group specialization.

Interestingly it has also been shown recently (Ohtsuki 2006) 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 behavior can be sustained in some limited situations if $b / c > k$, where k is the average number of neighbors over all nodes (the average degree of the graph). This implies that if certain topologies can be imposed then pro-social behavior 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 neighbors.

Group-Splitting Model

In Traulsen and Nowak (2006) a group selection model is presented that sustains pro-social behavior 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 in combination with splitting and extinction processes 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 endogenously 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 10.6 shows a schematic of group-splitting evolution and an outline algorithm that implements it.

APPLICATIONS

We believe that these new models 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

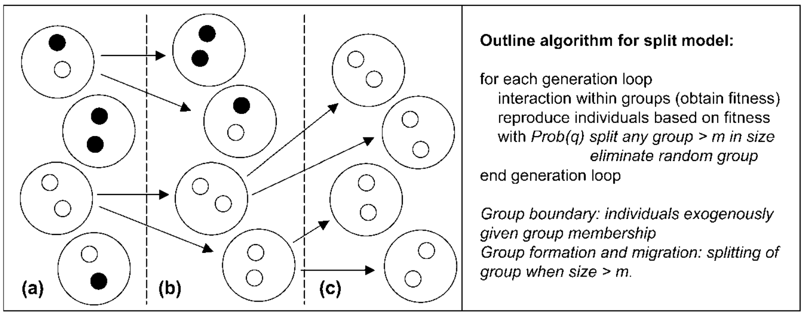

Figure 10.6 Schematic of the evolution of groups in the group-splitting model. Three generations (a–c) are shown. White individuals are pro-social; black are selfish. Individuals along the same group are shown clustered and bounded by large circles. Arrows indicate group linage. Migration between groups is not shown.

tracking and measurement of the dynamics of groups over time (Palla *et al.* 2007). Potentially massive clean data sets can be collected, and the models presented here can be calibrated and validated (or invalidated).

In addition, as has already been implied earlier, P2P systems composed of millions of online nodes could benefit from the application of group selection techniques by applying them directly to the algorithms (or protocols) used by nodes to self-organize productive services for users.

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 (consider online social networking as an example of this).

CONCLUSION

What these models demonstrate is that simple agent heuristics based on imitation directed towards individual improvement of utility can lead to behavior 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, over time, creates conditions that favor pro-social behavior. Agents receive utility by interacting in tribes (simple social structures). Tribes that cannot offer the agent a good "utility deal" 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 behavioral 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 that it is this low cost, and consequent freedom from geographical and organizational constraints, which is a major factor in the recent success of online communities, virtual social networks and other peer-production communities such as Wikipedia (Benkler 2006).

However, this process would not preclude agents with explicit group-level utility preferences—i.e., incorporating "social rationality" functions or the 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 proactively recruit others from the population. However, this introduces issues such as explicit recruitment processes, explicit internal social models and, potentially, transferable utility. This implies the requirement for an effective "store of utility" (i.e., money) that the simple models presented here do not contain. Here we begin to see formation of something that resembles a market. In this context the models we have presented could be seen as "pre-market" exchange structures in which value is not separated from the social structures that produce it because it cannot be easily stored, accumulated, transferred or controlled.

We might argue that where such "pre-market" structures perform well then there will not be any incentive for agents to engage in the additional costs of implementing explicit market structures. 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 incentivize non-cooperation. This is such a serious issue that game theory explicitly addresses it within the emerging area of mechanism design (Dash *et al.* 2003). However, often these models rely on standard rational action assumptions and a high degree of central control that enforce the "rules of the game".

A possible interesting future research direction could be to identify those scenarios in which tribal approaches are appropriate and those in which markets are appropriate. Here perhaps we could forge a third way between markets versus central control.

NOTES

1. This work was partially supported by the Future and Emerging Technologies program FP7-COSI-ICT of the European Commission through project QLectives (Grant no. 231200).

BIBLIOGRAPHY

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

Benkler, Y. (2006) *The Wealth of Networks: how social production transforms markets and freedom,* New Haven, CT: Yale University Press.

Binmore, K. (1994) *Game Theory and the Social Contract,* vol. 1: 'Playing fair', Cambridge, MA: MIT Press.

Dash, R., Jennings, N. and Parkes, D. (2003) 'Computational-mechanism design: a call to arms', *IEEE Intelligent Systems,* November, 40–7 (Special Issue on Agents and Markets).

Doran, J. (1998) 'Simulating collective misbelief', *Journal of Artificial Societies and Social Simulation,* 1 (1).

Hales, D. (2000) 'Cooperation without space or memory: tags, groups and the prisoner's dilemma', in S. Moss and P. Davidsson (eds) *Multi-Agent-Based Simulation,* LNAI 1979, Berlin: Springer, 157–66.

Hales, D. (2004) '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* (p2p2004), Washington, DC: IEEE Computer Society Press.

Hales, D. (2006) 'Emergent group-level selection in a peer-to-peer network', *Complexus,* 3: 108–18 (DOI: 10.1159/000094193).

Hales, D. and Arteconi, S. (2006) 'SLACER: a self-organizing protocol for coordination in P2P networks', *IEEE Intelligent Systems,* 21 (2): 29–35.

Hardin, G. (1968) 'The tragedy of the commons', *Scienc* 162: 1243–48.

Holland, J. (1993) 'The effect of labels (tags) on social interactions', Santa Fe Institute working paper 93–10–064, Santa Fe, NM.

Lindenberg, S. (2001) 'Social rationality as a unified model of man (including bounded rationality)', *Journal of Management & Governance,* 5 (3–4): 239–51.

Marcozzi, A. and Hales, D. (2008) 'Emergent social rationality in a peer-to-peer system', *Advances in Complex Systems,* 11 (4): 581–95.

Nash, J. (1950) 'Equilibrium points in *n*-person games', *Proceedings of the National Academy of Sciences,* 36 (1): 48–9.

Ohtsuki, H. *et al.* (2006) 'A simple rule for the evolution of cooperation on graphs and social networks', *Nature,* 441 (25): 502–5.

Palla, G., Barabasi, A.L. and Vicsek, T. (2007) 'Quantifying social group evolution', *Nature,* 446: 664–7.

Richerson, P. and Boyd, R. (2001) 'The biology of commitment to groups: a tribal instincts hypothesis', in R.M. Nesse (ed.) *Evolution and the Capacity for Commitment,* New York: Russell Sage Press.

Riolo, R. (1997) 'The effects of tag-mediated selection of partners in Evolving populations playing the iterated prisoner's dilemma', Santa Fe Institute working paper 97–02–016, Santa Fe, NM.

Riolo, R., Cohen, M.D. and Axelrod, R. (2001) 'Cooperation without reciprocity', *Nature,* 414: 441–3.

Roberts, G. and Sherratt, N.T. (2002) 'Similarity does not breed cooperation', *Nature,* 418: 449–500.

Santos, F.C., Pacheco J.M. and Lenaerts T. (2006) 'Cooperation prevails when individuals adjust their social ties', *PLoS Computational Biology,* 2 (10): 140.

Simon, H.A. (1990) 'A mechanism for social selection and successful altruism', *Science,* 250: 1665–8.

Simon, H.A. (1997) *Models of Bounded Rationality,* vol. 3: 'Empirically grounded economic reason', Cambridge MA: MIT Press.

Smith, A. (1836, first edn 1776) *An Inquiry into the Nature and Causes of the Wealth of Nations*, England: William Clowes and Sons.
Smith, J.M. (1982) *Evolution and the Theory of Games,* Cambridge: Cambridge University Press.
Traulsen, A. and Nowak, M.A. (2006) 'Evolution of cooperation by multilevel selection', *Proceedings of the National Academy of Sciences,* 130 (29): 10952–5.
Wilson, D.S. and Sober, E. (1994) 'Re-introducing group selection to the human behavioural sciences', *Behavioural and Brain Sciences,* 17 (4): 585–654.