Cooperation and Specialisation without Kin Selection using Tags

CPM Working Paper 02-88 (version 2)
David Hales - February 2002

dave@davidhales.com
Centre for Policy Modelling
The Business School
Manchester Metropolitan University
Manchester, UK.

Abstract

We present a model that demonstrates the evolution of groups composed of cooperative individuals performing specialised functions. Specialisation and cooperation results from an evolutionary process in which selection and reproduction is based on individual fitness. Specialists come to help (through the donation of resources) their *non-kin* group members, optimising their behaviour as a team and producing a fitter group. The mechanism that promotes this benevolent, cooperative group behaviour is based on the concept of a "tag". Tags are observable markings, cues or displays. Individuals can observe the tags of others and take alternative actions based on those observations (e.g. to altruistically help or not).

Introduction

Recent tag based models (Hales 2000, Hales 2001, Riolo et al 2001) have shown how benevolent behaviour can be evolved between individuals in one-time interactions. However, in these models the altruistic behaviour of individuals may be interpreted as a form of kin selection (Sigmund & Nowak 2001). This interpretation is possible because all the agents within a cooperative group are identical[1] when cooperation is high.

In this paper, however, we demonstrate tag processes that are sufficient to produce sustained altruistic behaviour towards others who are *not kin related*. Moreover, we show that this non-kin based altruism is a basis for the evolution of groups of heterogeneous (specialised) individuals who, although not kin related, cooperate and work to benefit the *group as a whole*.

The tag processes presented in the model can therefore be *interpreted* as selection at a supra-individual level. Groups of specialised individuals appear to cooperate and evolve to increase group level fitness. Specifically, we note that the model we present is constructed such that individuals cannot help kin directly and hence the cooperative behaviour *cannot* be the result of kin selection.

We advance the model as an example of how, even simple organisms (or artificial computational agents), can evolve to form cooperative heterogeneous groups or teams composed of individuals performing specialised tasks.

---

[1] In both models when *identical* agents are paired they *must* act altruistically.

The model can be interpreted culturally, under the assumption that individuals copy the behaviours of successful individuals, or genetically, where the fittest individuals have a higher reproductive success.

Finally we note that although the model sustains group directed non-kin altruism, agents with *smarter* partner selection strategies could *outperform* the dumb agents presented here. We outline possible smart strategies that would increase altruism and sustain larger groups supporting more highly specialised agent teams.

The Model

The model consists of a population of evolving agents. The tag matching mechanism follows that of Riolo et al (2001). Here we outline his model but extend it with the representation of different agent "skills". Each agent has three traits, a tag $\tau \in [0,1]$, a tolerance threshold $1 \leq T \geq 0$, and a skill type $S \in \{1,2\}$. Initially tags, thresholds and skills are allocated uniformly randomly. In each generation, each agent is awarded some number P of resources. Each resource is assigned a required skill type. Resources can only be "harvested" by agents possessing the required skill type. The skill type assigned to a resource is randomly assigned. If an agent receives a resource award that matches its skill type then it can harvest the resource and gain a benefit[2], b, from it. If the agent cannot harvest the resource it may search for another agent in the population with the required skill and donate the resource. Here the simplest possible search method is used to locate a potential donor. An agent is chosen at random from the population.

Donation only occurs to the randomly chosen potential recipient if it has the required skill type and a sufficiently similar tag value. A recipient tag is considered to be sufficiently similar if it is within the tolerance of the donating agent. Specifically, given a potential donor agent D and a potential recipient R a donation will only be made when $| \tau_D - \tau_R | \leq T_D$. This means that an agent with a high T value may donate to agents over a large range of tag values. A low value for T restricts donation to agents with very similar tag values to the donor. In all cases donation can only occur when the skill type of the receiving agent matches the skill type associated with the resource. If a donation is made the donating agent incurs a cost, c, and the recipient gains a benefit, b (since it can harvest the resource). In all experiments given in this paper, the population is composed of 100 agents, the cost $c = 0.1$ and the benefit $b = 1$. The number of resource awards P and the skill set S is varied (2 skill and 5 skill scenarios are compared) - see results below.

After all agents have been awarded P resources and made any possible donations the entire population is reproduced. Reproduction is accomplished in the following manner – each agent is selected from the population in turn, its score is compared to another randomly chosen agent, and the one with the highest score is reproduced. Mutation is

[2] However, in the experiments presented here, resource awards are never made that match the skill of the receiving agent. This means that *every* resource award may lead to a donation action. This was implemented to aid analysis (eliminating the effects of stochastic donation possibilities). However, the results presented here are not substantially different from those obtained without this simplification.

applied to each trait of each offspring. With probability 0.1 the offspring receives a new tag (uniformly randomly selected). With the same probability, gaussian noise is added to the tolerance value (mean 0, standard deviation 0.01). When $T < 0$, it is reset to 0. Also with probability 0.1 the offspring is given a new skill type (uniformly randomly selected).

Results

The first set of results, in Table 1 below, show the donation rates achieved as a percentage of total resource awards made (over all generations and runs). Also the average tolerance value is shown. Both values are averages over 30,000 generations with 30 replications. Specifically then, the donation rate gives the percentage of awards that resulted in an *actual* donation (for all generations and all 30 runs). The standard deviations are over the 30 individual replications (runs), which were started with different pseudo-random number seeds.

Awards (P)	Donation Rate (Ave. %)	Donation (St. dev.)	Tolerance (Ave.)	Tolerance (St. dev.)
1	2.6	0.000	0.017	0.000
2	2.2	0.000	0.012	0.000
3	2.3	0.000	0.010	0.000
4	6.4	0.064	0.010	0.000
6	30.3	0.007	0.021	0.021
8	32.8	0.001	0.024	0.024
10	33.8	0.015	0.043	0.043
20	35.5	0.034	0.106	0.078
40	36.0	0.047	0.241	0.241

Table 1
Donation rates and tolerance levels for different numbers of resource awards in a 2-skill scenario.

As can be seen in Table 1, the donation rate increases dramatically at 6 awards. As the number of awards is increased from 6 upwards the donation rate increases slightly but not dramatically.

In a 5 skill scenario (Table 2) a similar pattern is seen but more awards are required to produce significant levels of donation. The transition occurs between 10 and 20 awards but the increase is less spectacular (a jump from approximately 2% to 13% donation).

Additional results (not shown here) confirmed a hypothesis that as the number of skills is increased, the donation rate is reduced and more awards are required to produce a significant transition.

Awards (P)	Donation Rate (Ave. %)	Donation (St. dev.)	Tolerance (Ave.)	Tolerance (St. dev.)
1	1.5	0.001	0.028	0.002
2	1.1	0.000	0.019	0.001
3	1.0	0.000	0.015	0.001
4	0.9	0.000	0.013	0.000
6	0.9	0.000	0.011	0.001
8	0.9	0.000	0.010	0.000
10	2.1	0.002	0.010	0.000
20	12.9	0.000	0.025	0.003
40	13.9	0.015	0.098	0.190

Table 2
Donation rates and tolerance levels for different numbers of resource awards in a 5-skill scenario (i.e. there are 5 skill types, such that each agent has a skill S ε{1,2,3,4,5}).

Discussion

Although the donation rates here do not appear high, they indicate that the tag processes applied in the model are *sufficient* to support *significant levels* of donation between non-kin specialists. The fact that donation is occurring at all indicates that in-group diversity of skills is being maintained. In the 2 skill scenario when the number of awards are high, over one third of awards are donated to non-kin group members and in the 5 skill scenario over 10% donated.

Let us take stock of what a run of the model is demonstrating. An observer of such an artificial population would see agents receiving awards and passing those awards to others within their group with the appropriate skill *even though this causes the passing agent to incur a cost*. Such behaviour would appear to be a *non-kin form of group directed altruism supporting in-group specialisation*.

However, it would appear that if this is a form of group selection, then the selective pressure appears to be low – since only moderate (though significant) levels of donation are selected for. Why is this? How could the model be altered to produce much higher levels of altruistic in-group behaviour, supporting high degrees of specialisation in the form of more skills? What might be sufficient conditions to improve this group directed non-kin altruistic process?

In the model discussed so far, agents attempt to find partners by randomly selecting from the entire population. We might think of this as the simplest possible strategy for finding a partner – a *dumb* search strategy. Intuitively, it would seem that the probability of finding an appropriate partner (an in-group agent with the appropriate skill) would become lower as the population size increases and the number of skills increases. If agents cannot find their other group members then they cannot make a donation. For small population sizes with a small number of skills a dumb search strategy would appear

to be sufficient. But the dumb strategy won't scale for larger and more skill diverse populations. Would smarter partner selection strategies support high levels of in-group altruism?

What kinds of smart selection strategies could agents use? If we endow agents with high cognitive capabilities, the tag could be interpreted as a physical location or "central store" through which agents pass resources. In this context, each group (set of agents sharing similar enough tags) would maintain a central store or "clearing house" through which resources were passed. Alternatively, agents could utilise knowledge of the social networks within which they were embedded to pass resources to appropriate in-group members. In the model presented so far, agents either pass a resource directly to another agent with the appropriate skill or do not pass at all. But the model could be extended to allow agents to pass to any other in-group agent in the hope that that agent could find an appropriately skilled agent who could harvest the resource. In this latter case, a kind of "supply chain" could emerge based on social networks.

In all such cases, smart partner selection would require smart agents with social knowledge and rudimentary planning abilities. However, if smart selection strategies can outperform dumb strategies (such as random selection from the population) then this indicates that there is selection pressure for the development of such smart strategies. From this view we might hypothesize that the higher the cognitive ability of agents, the better able they are to sustain large and specialised cooperative social groups.

However, to add weight to the claim that there *is* significant pressure for smart partner selection strategies, it should also be shown that smart strategies could support additional costs and still outperform dumb ones. Cognition costs – sophisticated passing behaviours and storage of social knowledge would require more effort than dumb random selection.

In a following paper (Hales 2002) experiments with smart strategies are described. It is show that such strategies *do indeed* support higher levels of altruism and specialisation (even when costs are high).

Acknowledgements

Thanks to Bruce Edmunds (Centre for Policy Modelling, Manchester Metropolitan University, Manchester, UK) for comments and suggestions on previous drafts of this paper.

References

Hales, D. (2000), "Cooperation without Space or Memory: Tags, Groups and the Prisoner's Dilemma". In Moss, S., Davidsson, P. (Eds.) *Multi-Agent-Based Simulation. Lecture Notes in Artificial Intelligence 1979*. Berlin: Springer-Verlag (available at http://www.davidhales.com).

Hales, D. (2001), *Tag Based Cooperation in Artificial Societies*. Unpublished Ph.D. Thesis, Department of Computer Science, University of Essex (available at: http://www.davidhales.com/thesis).

Hales, D. (2002), Smart Agents Don't Need Kin – Evolving Specialisation and Cooperation Using Tags. *CPM Working Paper 02-89.* The Centre for Policy Modelling, Manchester Metropolitan University, Manchester, UK (available at: http://www.cpm.mmu.ac.uk/cpmreps.html).

Riolo, R., Cohen, M. D. & Axelrod, R. (2001), Cooperation without Reciprocity. *Nature* 414, 441-443.

Sigmund & Nowak (2001), Tides of tolerance. *Nature* 414, 403-405.