**Evolving Specialisation, Altruism and Group-Level Optimisation Using Tags**

David Hales - April 2002 – Submission for MABS2002
dave@davidhales.com
Centre for Policy Modelling (http://www.cfpm.org)
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). We show that even random (or *dumb*) searching for appropriate partners produces significant levels of specialisation and cooperation. Additionally we demonstrate that non-random (or *smart*) searching supports very high levels of cooperation.

## Introduction

Recent tag 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 their *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 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.

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

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. The individuals within the groups work together symbiotically to produce a more efficient unit to the benefit of all – even though this requires altruistic behaviour between self-interested individuals. We therefore advance the model and results as a minimal demonstration of evolutionary mechanisms that can guide (boundedly rational) individual self-interested utility maximises towards the maximisation of the social (group level) good.

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.

Additionally we demonstrate that populations of agents with *smarter* partner selection strategies support higher levels of specialisation and cooperation. Finally we hypothesize that an evolutionary process would select for smart strategies.

## The Model

The model consists of a population of 100 evolving agents. The tag matching mechanism follows that of Riolo et al (2001). The specialisation process follows Hales (2002a). Here we briefly summarise the model. Each agent has three traits, a tag $\tau$ ε [0,1], a tolerance threshold $1 \leq T \geq 0$, and a skill type S ε {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 from those skills that do not match the receiving agents skill[2]. An agent therefore is never awarded a resource that matches its skill type. Since the agent cannot harvest the resource, it searches for another agent in the population with required skill and tag values.

Donation only occurs if a recipient is found with the required skill type and with 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 here, the benefit b = 1 but the cost c is varied as is the size of the skill set S (see results).

[2] Results obtained from a model in which agents may be awarded resources matching their own skill types produced similar results to those presented in this paper.

In the experiments presented we compare two kinds of agent society: societies composed entirely of agents using *dumb* searching strategies and societies composed entirely of agents using *smart* searching strategies. These strategies indicate the way an agent searches for a potential recipient to donate to. Agents using a dumb strategy select an agent at random from the population, after each resource award, and make a donation if the recipient has a sufficiently similar tag and skill matching the resource skill. Agents using a smart strategy *search the entire population* for a recipient that has a sufficiently similar tag and the required skill. We assume (but do not directly model) that some efficient mechanism exists which allows agents to *find a potential recipient in the population if one exists*[3]. As discussed elsewhere (Hales 2002b) a number of plausible mechanisms can be hypothesised based on spatial and / or cognitive relationships (e.g. "small world" social networks, meeting places, central stores – see the later discussion for more on this).

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 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 $T < 0$ or $T > 1$, it is reset to 0 and 1 respectively. Also with probability 0.1 the offspring is given a new skill type (uniformly randomly selected).

## Results

The first set of results, shown in table 1 and figure 1 below, show the donation rates achieved (as a percentage of total awards made) and the average tolerance values in a 2-skill scenario (for a populations of dumb agents and populations of smart agents). The results are over 30,000 generations with 30 replications. Each replication represents an individual run started with a different pseudo-random number seed. The standard deviations are over the 30 runs executed for each unique P value setting[4]. The column labelled "dumb" shows the results using a *dumb* random recipient search strategy. The "smart" columns show the results of the smart strategy populations. In order to make a "fair" comparison we also consider the results of the smart strategy when the cost, c, is increased from c = 0.1 to c = 0.5 (in all cases the benefit, b, is held at 1).

As can be seen in table 1 and figure 1, the donation rate for the *smart* strategy (when c=0.1) increases dramatically when P=2 (awards) and then increases modestly (though non-monotonically) as P is increased. In comparison with the previous results for the *dumb* searching strategy the donation rate is substantially higher (61% to 2.6% for a single award). Note also that the tolerance values are higher too and the standard

[3] If several suitable recipients exist in the population we assume here that one of them is selected to receive the donation at random.

[4] The standard deviations are not calculated over the percentages given but proportions (i.e. percentages scaled within [0..1] – so 100% would count as 1 and 50% as 0.5 etc.)

deviation of tolerance (over the 30 runs) is much higher (indicating a high heterogeneity over the runs)[5].

| Awards (P) | Donation Rate – Ave % (st.dev. in brackets) | | | Tolerance – Ave (st.dev. in brackets) | | |
|---|---|---|---|---|---|---|
| | Dumb c = 0.1 | Smart C = 0.1 | Smart c = 0.5 | Dumb c = 0.1 | Smart c = 0.1 | Smart c = 0.5 |
| 1 | 2.6 (0.000) | 61.0 (0.079) | 61.0 (0.076) | 0.017 (0.000) | 0.035 (0.157) | 0.038 (0.107) |
| 2 | 2.2 (0.000) | 80.0 (0.011) | 69.3 (0.083) | 0.012 (0.000) | 0.019 (0.050) | 0.055 (0.182) |
| 3 | 2.3 (0.000) | 85.6 (0.048) | 73.9 (0.052) | 0.010 (0.000) | 0.090 (0.216) | 0.044 (0.147) |
| 4 | 6.4 (0.064) | 85.8 (0.031) | 76.8 (0.061) | 0.010 (0.000) | 0.057 (0.147) | 0.062 (0.172) |
| 6 | 30.3 (0.007) | 87.7 (0.046) | 77.9 (0.008) | 0.021 (0.021) | 0.111 (0.217) | 0.013 (0.013) |
| 8 | 32.8 (0.001) | 90.5 (0.062) | 80.6 (0.043) | 0.024 (0.024) | 0.225 (0.290) | 0.049 (0.136) |
| 10 | 33.8 (0.015) | 89.3 (0.056) | 81.1 (0.039) | 0.043 (0.043) | 0.180 (0.279) | 0.040 (0.132) |
| 20 | 35.5 (0.034) | 89.5 (0.057) | 82.0 (0.003) | 0.106 (0.078) | 0.189 (0.274) | 0.012 (0.005) |
| 40 | 36.0 (0.047) | 91.3 (0.047) | 83.1 (0.015) | 0.241 (0.241) | 0.268 (0.305) | 0.025 (0.049) |

*Table 1*

*Donation rates and tolerance levels for different numbers of awards in a 2-skill scenario (i.e. when $S \in \{1,2\}$) for different search strategies and costs. The values in brackets are standard deviations over the 30 replications.*

When the cost, c, is increased to 0.5 the smart strategy does not increase as dramatically when P=2. But still shows substantially higher donation levels than the dumb strategy. The increase in donation rate as P is increased follows a more linear (monotonically increasing) path than when the cost was c=0.1 and the tolerance levels are generally reduced. The variance of the tolerance is reduced by still higher than for the dumb strategy.

In a 5-skill scenario (table 2 and figure 2) a similar (though even more dramatic) pattern is seen for the smart strategy when c = 0.1. Again a big increase in the donation rate takes place when P=2 followed by a much slower (non-monotonic rise) as P is increased. Again the variance over the 30 runs for each P value is much higher than the dumb strategy (evidenced by the standard deviation values).

[5] This result suggests that there may not be a "single story" (i.e. evolutionary trajectory) here. This means that increases in donation and tolerance might be highly contingent on initial conditions and on-going stochasticities (this is discussed later).

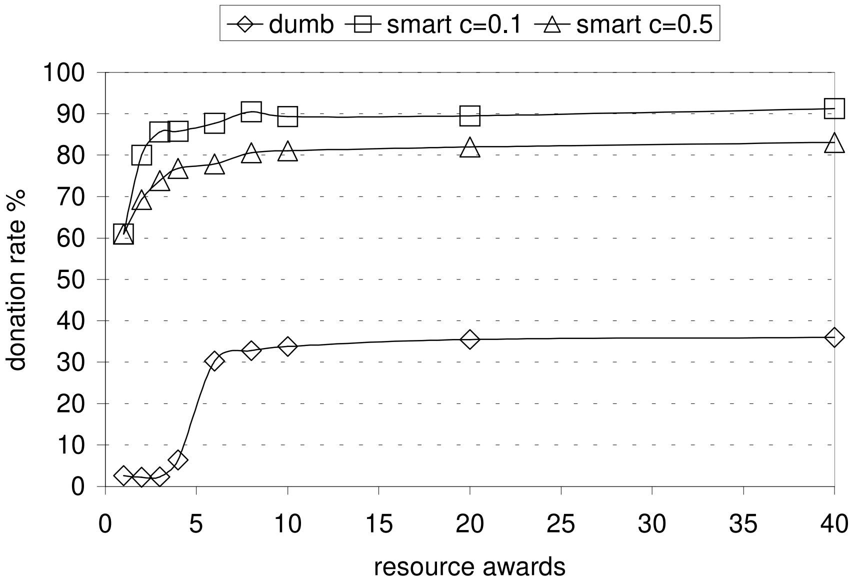


*Figure 1*
*A chart showing the donation rate data given in table 1 (above) for the 3 different population types in the 2-skill scenario. Note that the smart populations (with costs of 0.1 and 0.5) have much higher donation rates that the dumb population runs.*

When c = 0.5 the results are broadly similar to those for the 2-skill scenario. Again notice that for the smart strategy when c = 0.5 the increase when P = 2 is less dramatic, and as P increases the donation rate increases in a more linear (and monotonic) way. We note here however, that tolerance values and the variance of those values are lower than in the 2-skill scenario – almost comparable to the dumb strategy.

**Discussion**

Let us be clear about what is being demonstrated by the model: Agents form groups, based on tag similarity, containing a diversity of skills. Agents donate resources, requiring skills (to harvest) that they do not posses, to other agents within their group *even though this causes them to incur a substantial cost.* This behaviour persists even though the agents are reproduced on the basis of individual utility. Since agents can only pass resources to others (within their group) that posses the required skill, the high-level of donation rate produced indicates that *high levels of skill diversity are being maintained within groups*. This is exactly the kind of group organisation that can best exploit the environmental scenario. So, *agents are forming into very efficient, skill diverse groups*. Since skill diversity means that agents cannot be clones, this evolved structure cannot be the result of a simple form of kin based selection.

| Awards (P) | Donation Rate – Ave % (st.dev. in brackets) | | | Tolerance – Ave (st.dev. in brackets) | | |
|---|---|---|---|---|---|---|
| | Dumb c = 0.1 | Smart c = 0.1 | Smart c = 0.5 | Dumb c = 0.1 | Smart c = 0.1 | Smart c = 0.5 |
| 1 | 1.5 (0.001) | 29.5 (0.081) | 29.5 (0.081) | 0.028 (0.002) | 0.021 (0.118) | 0.021 (0.084) |
| 2 | 1.1 (0.000) | 75.3 (0.016) | 47.9 (0.087) | 0.019 (0.001) | 0.023 (0.034) | 0.030 (0.111) |
| 3 | 1.0 (0.000) | 81.8 (0.037) | 59.9 (0.035) | 0.015 (0.001) | 0.056 (0.110) | 0.017 (0.048) |
| 4 | 0.9 (0.000) | 84.3 (0.060) | 66.3 (0.046) | 0.013 (0.000) | 0.104 (0.222) | 0.028 (0.105) |
| 6 | 0.9 (0.000) | 84.4 (0.051) | 70.9 (0.010) | 0.011 (0.001) | 0.099 (0.210) | 0.011 (0.014) |
| 8 | 0.9 (0.000) | 88.5 (0.077) | 73.1 (0.002) | 0.010 (0.000) | 0.250 (0.319) | 0.009 (0.000) |
| 10 | 2.1 (0.002) | 84.1 (0.049) | 74.5 (0.002) | 0.010 (0.000) | 0.099 (0.208) | 0.009 (0.000) |
| 20 | 12.9 (0.000) | 85.8 (0.070) | 77.3 (0.002) | 0.025 (0.003) | 0.170 (0.261) | 0.010 (0.001) |
| 40 | 13.9 (0.015) | 91.0 (0.015) | 79.3 (0.033) | 0.098 (0.190) | 0.370 (0.341) | 0.038 (0.107) |

*Table 2*

*Donation rates and tolerance levels for different numbers of 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}) for different search strategies and costs. The values in brackets are standard deviations over the 30 replications.*

The results presented here indicate that "smart searching" strategies, i.e. efficient methods of find an appropriate in-group recipient for donation, substantially increase the efficiency of specialisation within groups. Since smart strategies would seem to require a higher cognitive ability within agents, can we conclude that there is strong selection pressure for such cognitive abilities? We *cannot* draw such a conclusion because *the results presented here do not pit dumb strategies against smart ones in the same evolving population*. We hypothesize, however, that smart agents would indeed out compete dumb agents – even when the cost of donating for smart agents was substantially higher. We put this hypothesis to the test in a subsequent paper (Hales 2002c).

As stated previously, in some of the results, very high levels of tolerance were produced. This tended to occur when P (awards) was high and a smart strategy was employed (though not when costs were high). In those cases where the tolerance was high, the variance (over the 30 replication) of tolerance was also high. This indicated heterogeneity of evolutionary trajectories taking place over replications.

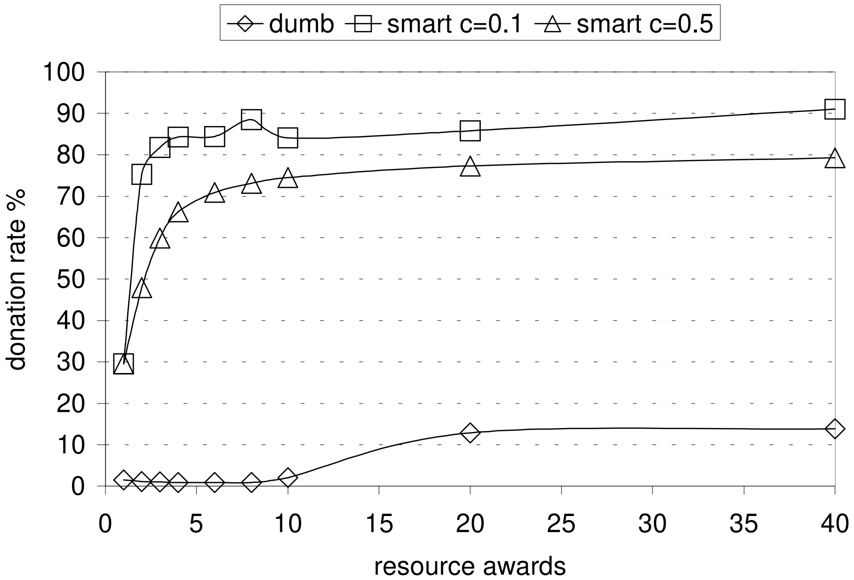


*Figure 2*
*A chart showing the donation rate data given in table 2 (above) for the 3 different population types in the 5-skill scenario. This can be compared to figure 1 – notice that with more skills less donation occurs – particularly for the dumb population type.*

To investigate this we examine the average tolerances and donation rates of the individual replication runs when P=40 in the 5-skill scenario for the smart strategy when cost c=0.1. Table 3 shows the results of the first 15 individual runs. Note that the runs fall into two categories: a) runs with low average tolerances and b) runs with very high average tolerances. The b-type runs produce almost 100% donation rates (the results are rounded to 4 significant figures in the table - the actual results are slightly less than 100%). The a-type runs have lower (yet still high) donation rates. The column labelled "Tag Clone Donations" shows the proportion of awards that resulted in a donation to a recipient agent with an *identical tag* value to the donor. As would be expected, runs that average high tolerances result in high donation to non tag-clones, even though donation rates are almost at maximum. Such results indicate that all agents donate to all others and maintain the required skill diversity to exploit each resource type.

But how is this happening? Intuitively, it would appear that, in the b-type runs the population should be invaded by "cheating" agents with a low tolerance that would restrict donation to agents with tag values much closer to their own while benefiting from donations from more tolerant agents.

But a moment's reflection indicates that agents may not benefit by reducing tolerance in the small decrements produced by the gaussian mutation method (described above) when the cost of donation is low and the initial tolerance is high. This is because a small decrement in the tolerance would mean that agents with higher tolerances could still be recipients of donations. Only if an agent was produced (via mutation) that had a very uncommon tag and a very low tolerance could a substantial increase in fitness be produced. It should be noted that when the cost is increased to c=0.5 the b-type runs disappear. This could be explained by the additional evolutionary advantage that the extra cost gives to "cheaters". But do these reflections explain the results? In tables 1 and 2, the b-type runs only occur in resource rich (i.e. when awards (P) are high) environments. We currently do not have an explanation for this. Consequently the b-type runs are currently under closer investigation.

| Run | Donation Proportion | Tolerance Ave | Tag Clone Donations |
|---|---|---|---|
| 1 | 0.811 | 0.020 | 0.768 |
| 2 | 1.000 | 0.779 | 0.013 |
| 3 | 1.000 | 0.673 | 0.015 |
| 4 | 0.881 | 0.188 | 0.429 |
| 5 | 0.998 | 0.600 | 0.018 |
| 6 | 1.000 | 0.669 | 0.015 |
| 7 | 0.861 | 0.109 | 0.499 |
| 8 | 0.811 | 0.021 | 0.750 |
| 9 | 1.000 | 0.659 | 0.015 |
| 10 | 0.829 | 0.045 | 0.699 |
| 11 | 1.000 | 0.768 | 0.013 |
| 12 | 1.000 | 0.901 | 0.012 |
| 13 | 0.917 | 0.418 | 0.297 |
| 14 | 0.834 | 0.046 | 0.641 |
| 15 | 0.813 | 0.022 | 0.760 |

*Table 3*

*The first 15 individual runs for the smart strategy (c=0.1 and P=40) in the 5-skill scenario. Each run represents a single execution of the model to 30,000 generations starting with a different pseudo-random number seed. The Donation Rate shows the proportion of awards that resulted in a successful donation. The Tolerance Average shows the average tolerance over all agents over each generation. The Tag Clone Donations shows the proportion of successful donations made to an agent with an identical tag value.*

The forms of interaction and specialisation possible in the model presented in this paper are limited. The assumption that donation only takes place between two individuals (i.e. donation results from dyadic pairing of a donor and recipient) precludes the possibility of an agent donating to several agents at once. The environmental scenario, in which agents are rewarded individually for diverse skills, does not allow agents to become *full-time* specialists in other roles – such as internal group organisation roles (e.g. redistribution agents, that collect and redistribute resources to the in-group or policing agents that punish potential free-riders within the group). Intuitively, it would seem, agents would need to have *smarter redistribution* (or donation) strategies than those presented in the present model and richer interaction abilities. A forthcoming paper will explore the evolution of smarter redistribution strategies.

One of the aims of this paper is to demonstrate empirically from an evolutionary algorithm (since we currently lack a deductive theory) that altruism, group formation and specialisation are intimately linked in the formation of integrated higher-level social entities which are better adapted to their environment than atomised individuals. Specifically we are trying to show that the underlying processes that bring about such organisations can be driven by essentially self-interested learning behaviour. In this sense we question the intuition that evolution leads to selfish replicators. We argue that even in our simple scenario, this intuition would lead to confusion and a failure to understand the behaviour of individual agents observed over time.

This final observation seems particularly poignant when applied to the concept of "cultural evolution". In modern societies we appear to observe large numbers of institutions and social groups composed of many individuals with highly specialised roles and seemingly altruistic behaviours within and even between groups. We argue that such complexity and seemingly "unselfish" behaviour does not necessarily undermine the project of attempting to explain such phenomena from an evolutionary perspective[6].

**Acknowledgements**

Work presented here has greatly benefited from discussions with, and ideas from, Bruce Edmunds, Centre for Policy Modelling, Manchester Metropolitan University, UK. The results presented here are based on two previous internal CPM discussion papers – Hales (2001a), Hales (2001b).

**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. (2002a), Cooperation and Specialisation without Kin Selection using Tags. *CPM Working Paper 02-88*. The Centre for Policy Modelling, Manchester Metropolitan University, Manchester, UK (available at: http://www.cpm.mmu.ac.uk/cpmreps.html).

Hales, D. (2002b), Smart Agents Don't Need Kin – Evolving Specialisation and Cooperation with 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).

[6] In the context of cultural evolution the fundamental units have been hypothesised as memes (behaviours, beliefs and traits that may be copied between individuals).

Hales, D. (2002c), Wise-Up! - Smart Tag Pairing Evolves and Persists. *CPM Working Paper 02-90*. 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.