

ALTRUISM AMONG SIMPLE AND SMART VAMPIRES

GENNARO DI TOSTO

MARIO PAOLUCCI

ROSARIA CONTE

*Institute of Cognitive Science and Technology, CNR
Via San Martino della Battaglia 44, 00185 Rome, ITALY*

In this work we present a simulation study drawn upon a famous ethological example: food-sharing habit of vampire bats in Central America. A norm of reciprocity is introduced when agents are expected to exhibit a cooperative behavior under serious environmental strictures (infrequent but lethal food scarcity). We compare different agents' architectures (with different levels of cognitive complexity) in the evolution of food-sharing habit and we analyse their performances against the presence of cheaters in the population. The experiments are aimed at studying the role of cognitive mediators (i.e. goals) in the spreading of altruism, described as a purposeful action.

Keywords: Agent modelling; cognition; altruism; evolution of cooperation.

1. What is the Problem

That altruism is advantageous from a biological point of view is consolidated scientific knowledge. Sociobiologists *à la* Trivers¹⁵ proposed a biological theory of reciprocal altruism, in terms of its reproductive advantage for individual donors. This theory received support from experimental ethologists,¹⁷ from evolutionary game theorists,¹¹ and from evolutionary psychologists.^{5,9} The biological advantage of altruism is also acknowledged by group selection theorists,¹ although in their view altruism benefits (and is reproduced by) the group as a whole rather than its members.

However, most authors agree that a pure, unconditioned form of altruism has not many chances to spread in presence of cheaters. Rather, what is found to contribute to the reproductive success of individuals (and groups) is a conditioned form of altruism, i.e. reciprocity or acting to the benefit of others known as altruists. Reciprocity, in its turn, raises another fundamental question, namely how altruists select recipients of help. How tell that current recipients will be future donors? Essentially, either direct experience (one's image of the potential recipient) or indirect one (its reputation) assists us in selecting recipients. Apparently, the capacity to apply and reason upon a norm of reciprocity evolved as early as the hominids' habit to settle in large non-kin groups.⁷ The evolutionary psychological view of reciprocity shed a new light on the issue of social norms. Unlike conventions, these are seen as *mental* mechanisms, used for social reasoning and for social action.

The follow-up question then is *what type of mental construct corresponds to the norm of reciprocity* (and possibly, to any other social norm). Whereas this question received attention by evolutionary psychologists, and is addressed in natural experiments with human agents,¹⁰ it is overlooked within the simulation-based study of altruism. There is a consistent number of works in the field using altruism to deal with problems of optimization in MAS¹⁴ or using indirect communication based on stigmergy to observe the emergence of cooperative behavior in ant-like species⁸. Our focus is, however, directed inside the mind of the agents, and our intention is to shed some light on the role played by intelligent systems in the evolution of pro-social behavior.

Acquired results from evolutionary game theory show that cooperative behavior can lead a population to converge on a solution that is utility-maximizing in the long term, even if it is irrational in the short one; i.e. an altruistic behavior (a behavior with a cost for the performing agent, whose benefit are addressed to some other agent) is evolutionary fit only if reciprocated. It is possible to observe a similar phenomenon when the reasoning activity of the agents is driven by internal motivational state instead of an utility function? This paper revolves around this question. Rather than providing an exhaustive answer, we will concentrate our attention on the comparison between two extreme options: a simple, or sub-cognitive, altruistic algorithm and a smart, or quasi-cognitive one. In the former, agents apply routines under given conditions. In the latter, agents execute actions to achieve their goals. By this means, we intended to model motivations and to observe their impact on altruism.

Theories of intelligent or rational action emphasise the role of (limited) knowledge (Simon¹³, Kahneman and Tversky¹⁶, etc.). Analogously, the evolutionary theory of mind stresses the representational side of cognition. In both cases, motivations are usually overlooked. In other fields of science (AI and agent systems), where the purpose is to construct autonomous systems applied to some task, the role of goals as sources and guide of action is explicitly acknowledged. The scarce attention paid to the agents' minds in the (simulation-based) study of altruism may also be based on a folk theorem, according to which altruism and intelligence are essentially incompatible.

The common observation that altruism is frequent in ant-like species, where individual autonomy is not fully developed and behaviour is rigidly programmed by biological evolution, led to the general assumption that the higher the species' evolutionary rank, the less frequent the altruistic behaviour. Does this mean that intelligence hinders or even prevents altruism? What are the effects of mental complexity and individual autonomy on social action? This question becomes crucial when considering that altruism enhances biological fitness.

In the attempt to shed some light on these questions, and to check the validity of the previous folk theorem, we run a series of simulation experiments inspired by a famous ethological example, i.e. the vampire bats' altruism.¹⁷ This species has puzzled ethologists for decades, while at the same time providing evidence for the

theory of reciprocal altruism. Vampire bats evolved a rather unique form of food sharing, consisting of successful hunters regurgitating a portion of the blood ingested in favour of unlucky fellows. Apparently, the rationale of this behaviour cannot be found in kin selection, since the average rate of relatedness among individuals living in the same roosts is rather low (around 6%). Instead, simulation findings¹⁸ showed that altruism allows the rate of survivors to rise up to 80% of the initial population, as opposed to the bare 20% obtaining in simulated roosts where individuals do not help one another. Sociobiologists⁶ interpreted vampires' food-sharing as supporting the reciprocal altruism theory. Simulations supporting an interpretation based upon group selection theory have recently been run.¹² Whatever the biological rationale of vampire bats, this species offers a good target for modelling altruism because its survival is strongly interdependent with the evolution of altruistic behaviour. However, our study is aimed neither at contributing to the sociobiological debate around altruism nor at promoting ethological simulation, however fascinating such an interdisciplinary enterprise may appear. We intend to explore altruism at an abstract level, and find out its underlying mental ingredients. Far from addressing the question as to what the real vampire mind is like, we explore the relation between some aspect of cognition and altruism, possibly showing their co-evolution. In sum, the paper aims to

- Model altruism at an abstract level, taking inspiration by a real-world example (vampire bats) which provides good evidence for setting parameters to given values in a non-arbitrary way.
- Explore the internal ingredients allowing agents to exhibit altruism; in particular compare two different algorithms, simple and smart, and observe their respective performance.
- Check the validity of the folk theorem that intelligence hinders altruism: should we assume vampires to be stupid enough as to rigidly apply some pre-established routines, or else is their behaviour compatible with a more intelligent and sophisticated mental equipment? And if so, what are the specific advantages of such a higher mental complexity?

2. Simple and Smart Mechanisms

A social norm operates in different ways according to the level of complexity of the agent architecture. With simple agents, essentially endowed with a set of production rules and a repertoire of actions, the output of a production rule will always be an action.

Smart agents are endowed not only with actions, but also with mental states such as beliefs and goals. Therefore, the norm of reciprocity will be a more complex construct taking beliefs as input conditions and giving a set of mental representations, namely beliefs and goals, as outputs. If the norm reads “under a certain condition c , a given behaviour b ought to be executed”, and c is verified, it can be inferred that something (b) ought to be done. The main output of the rule is then a normative

belief: it tells you how you should behave, what is expected of you. Whether you will actually conform or not depends upon a further mental process initialised by the rule, i.e., the formation of a normative goal based upon the normative belief.⁴ Let us call these two options the simple and the smart mechanism for implementing social norms, or, for short, the simple and the smart algorithms. The question is, what is the difference between them? Are they interchangeable (what would obviously render the simpler always preferable), or is there a significant difference of functionality, and if so, which one?

Our hypothesis is that the simple and the smart architectures are not interchangeable, and that they correspond to different stages in the evolution of social intelligence. In particular, we expect that simple agents always apply the norm with reasonably high life expectancies, whereas smart agents may not (since a goal corresponding to the norm is not necessarily formed). In this sense, a simple architecture is preferable to a smart one.

However, other effects of both architectures may be envisaged, pointing to less obvious advantages of the smart one. A smart implementation of the norm may be preferable when environmental conditions are such that existing means and actions for executing the norm are inadequate or insufficient. Far from interchangeable, the two architectures may be expected to correspond to different stages in the evolution of social intelligence. Simple agents don't fail to apply the norm but show adaptive behaviour only by means of genetic recombination. On their side, smart agents are less reliable but more flexible and are expected to observe the norm even with low life expectancies. Which one is better under which condition?

We have designed a simulation experiment to compare the two algorithms and provide initial answers to this question.

3. The Simulation Model

In nature, groups of vampire bats share the same physical sites (caves or trees), called roosts. the roost is the place where they get back to after hunting, in which they reproduce and perform other social activities (nursing, grooming and sharing food). From now on, we will refer to agents who live in the same roost as in-roosts.

In our simulation, the world and the food are represented as parameter values, not as objects. The entities explicitly modelled as objects are the agents, the bats. Roosts are modeled as aggregates of bats (around 15 on average). In-roosts are allowed to share food and to groom one another. No other social activity will be modelled.

Each simulation cycle includes one day and one night stage. During the day, the simulated animals perform social activities (grooming and food sharing). At night, they hunt. In our model, hunt is defined as an ecological parameter; in accordance with real-world data, its default value is set to 93%. In substance, each night 93% of the population will find food to survive until the next hunt. The remaining 7% will starve, unless they receive help from some fellow (under the form of regurgitation).

Vampire bats do not accumulate resources (except for credits, see below): hunt is performed only for immediate food consumption. In addition, although the average lifetime of these animals lasts around 14 ys., starvation and death are a constant menace to them, since each good hunt gives them no more than 60 hours' autonomy.

These conditions are exactly reproduced in our simulation model. As to daily activities, grooming has at least two effects in nature: thanks to and during it, animals familiarise and check their respective physical shape. Since satiation causes body volume inflation, a lucky hunter may grow to almost 50% more than its normal size, as can be easily detected by any grooming partner. Likewise, a starving bat is also likely to be recognised. Bluff would immediately be found out.

In our model, animals are immersed in grooming networks, which are randomly activated at the onset of simulation. The network nodes represent potential partners for grooming interaction. Each day, pairs are formed by each animal choosing one partner from its grooming network. As in the real world, also in our model grooming has the effect of increasing the probability of food-sharing among in-roosts: a starving bat will turn to grooming partners for help, and will avoid death if any of them is found to be full (having had a good hunt).

A starving fellow is not denied help by a satiated in-roost: the grooming network creates a familiarity as well as a reciprocity basin. Help-giving allows animals to achieve credits, which will be extinguished if and when help is returned. Lucky hunt may last the short space of one night, and every agent may soon shrink in starvation. Hence, it will be urged to go out for grooming in the hope to meet with a luckier (and fatter) debtor. In less metaphorical terms, the grooming network facilitates re-encounters and therefore the extinguishing of credits.

Consequently, artificial bats are endowed with social knowledge: consisting of a memory of past grooming and food-sharing interactions, and of consequent credits. After each simulation the number of expired agents, the number of altruistic acts performed, and the number of credits turned on or off are recorded.

4. First Study: The Simulation Experiment

Wilkinson¹⁷ showed food-sharing to advantage vampires coming from the same roost, preventing the extinction of in-roosts. Getting back to our main problem, we asked ourselves how is reciprocity dealt with by vampires. Is it a reactive behavior, inevitably leading the animals to behave altruistically in presence of any starving fellow, or is it a more complex rule, which requires a more elaborated mental process? To answer this question, we compared two populations of vampires in terms of mortality and reproduction rates. In one population, the simple group, the norm gives a mere behavioural output, whereas in the other, the smart group, the output is mental, i.e. a subjectively valued motivation.

4.1. *Reciprocity and the Credit Network*

Reciprocity is implemented thanks to a network of credits, by means of which every agent can keep track of its interactions. This tool represents the social memory of the artificial vampires, providing them individual recognition. Every vampire has its own network-like memory, every nodes being a past recipient and every link being an action of food-sharing.

At any donation the network is updated. Either a previous donor is refunded, in which case its credit is extinguished and one link removed; or, a new credit is formed, and a new link is activated between the current donor and its recipient. Whenever donors are reciprocated, their corresponding credits are cancelled. The credit network is investigated any time a request of help is received. In a more restrictive condition, only if no credit link is active with the postulant will the agent give help. Otherwise, help is denied. In a less restrictive condition, help is denied only when the same postulant asks for help more than two times consecutively.

Agents search for potential donors within the grooming network. Only one trial is available. If help is denied, the postulant is bound to die. What is interesting about the credit network is that it can be passed on to one's offspring, which inherit parents' features and credits. Consequently, a given credit can be extinguished during the donor's life or after its death to the benefit of its offspring. Obviously, the more the credits passed on to future generations, the higher the probability of survival of one's offspring.

4.2. *Actions*

In the model, only the following actions can be performed:

Groom: the condition for this action is that two agents are sorted out from the same grooming network. The effect of grooming is that the weight of the link among those partners is reinforced, and the probability that they share food is correspondingly strengthened.

Ask for help: the condition for the application of this action is that the postulant be starving. The request may be addressed to one agent in the same grooming network, if one is made. Otherwise, other in-roosts may be addressed as well. The effect of a request will be either donation or denial. In the first case, the postulant will ingest some blood and gain some hours of autonomy. In the second, it is bound to die.

Donate: the condition for applying this action is that recipient is starving. The effect is that donor's autonomy is reduced and the recipient's is increased.

Deny help: the condition is that agent received a request for help by someone. The effect is the latter's death.

4.3. *Control Condition: The Simple Vampire*

Simple agents have been modelled according to a rather simple procedure, which considers the following nested sets of conditions:

- If a starving fellow asks for help and you have ingested more than you need, give away surplus. According to Table 1, a postulant will be at 12 hours autonomy (starving point) or less. The donor will give away an amount of blood that corresponds to a loss of 6 hours autonomy for him, and to the recipient gaining 18 hours. This is done without taking too much risk, that is, when donors are at 48 hours autonomy.
- At less than 48 hours autonomy, it is unreasonable to give away blood, because the recipient will gain more than is left to the donor. In such a case, the donor will deny help.

During one day stage of the simulation, every agent has the possibility to play the role of the donor or the recipient. The way food-sharing behavior is designed avoids two agents to groom one another cyclically in order to increase artificially their autonomy: if and only if an agent is provided with at least 48 hours of autonomy, it will take into account the request of a starving fellow (i.e. with less than 12 hours of autonomy). With less than 48 hours of autonomy, together with the impossibility to store food for future needs, blood is a resource too important to share.

4.4. *Experimental Condition: The Smart Vampire*

At first, this has been modelled as a goal-based but rigid system, endowed with two goals, the normative (give help) and the survival goal (stay alive). To keep the two experimental conditions as comparable as possible, we endowed both types of agents with the same repertoire of actions (give blood, deny help). Both goals vary (randomly) according to their motivational force, which is kept constant throughout the simulation. The motivational force of goals is supposed to be set to 0, 1 or 2. Therefore six patterns of relationships between goals emerge:

- (1) 0 for the normative goal (NG) and 1 or 2 for the survival goal (SG);
- (2) 1 or 2 for SG and 0 for NG;
- (3) 0 for NG and 0 for SG;
- (4) 1 or 2 for NG and 1 or 2 for SG ($NG = SG \neq 0$);
- (5) NG higher than SG ($NG \neq SG \neq 0$; $NG > SG$);
- (6) NG lower than SG ($NG \neq SG \neq 0$; $NG < SG$).

By elimination of the third case, in which the animal is apparently affected by a total lack of motivation, five plausible cases are derived. In Table 1 the outputs of the motivational interplay are given together with the names which we will use from now on to refer to each case.

ACTIONS	GOALS (NG, SG)	STRATEGY
Deny	NG=0, SG=1-2	Cheaters
Donate even at 12	NG=1-2, SG=0	Martyrs
Donate 12 at 48, 6 at 24	NG=SG	Fair
Donate 24 at 48, 12 at 24	NG>SG	Generous
Donate 6 at 48, deny at 24	NG<SG	Prudent

Table 1. Outputs of motivations' interplay ACTIONS GOALS (NG, SG)

4.5. Hypotheses

Homogeneous populations of simple vampires, where no cheating occurs, are expected to show higher fitness (lower mortality) than smart ones, which are inherently heterogeneous. Among smart vampires, the higher number of donations produced by the individuals with the normative motivation exceeding the goal of survival should be compensated by the number of those that sacrifice their lives. Complementarily, smart vampires are expected to do better in the long run, when reproduction is turned on. In fact, individuals sacrificing their lives will gain credits transmitted to their offspring. Reproduction is expected to introduce a fundamental asymmetry between strong altruists and strong cheaters. In short, we expected the norm to be observed by all simple vampires if their life expectancies allow its execution. Instead, we expected it to be systematically violated by a consistent minority of smart vampires and conversely applied by another share of this population even when they have short life expectancies. As a consequence, in standard ecological conditions (93% probability of successful hunt) we expected simple vampires to observe the norm more than smart vampires and consequently reach a higher survival rate. The opposite pattern was expected with a lower probability of hunt. As one can see from Figure 1, food-sharing in homogeneous populations of simple vampires leads to the results obtained by Wilkinson: average mortality drops from around 80% to 24%.

In Figure 2, the performance of homogeneous simple vampires with food-sharing is compared with smart vampires. Findings indicate a clear advantage of the former. In Figure 3, the same situation is shown, but with agent reproduction activated.

4.6. Preliminary Discussion

So far, simulation findings indicate that a simple algorithm performs rather well. One could ask how such a good strategy evolved. One interesting answer¹² points to the effects of roosting. Simulation findings indicated that vampires probably evolved thanks to the co-evolution of food-sharing and roosting. Food-sharing allowed them to survive in roosts of homogenous altruistic populations, whereas heterogeneous roosts extinguish. Therefore, reciprocity emerged as an effect of inter-roost competition. Since only homogeneous altruistic roosts survived, reciprocity and altruism




Fig. 1. The effect of food-sharing. Values are number of agents per strategy (y), and ticks (x). 10 roosts of 15 agents each populate the world.




Fig. 2. Simple and smart vampires. Comparison of simulations with 10 roosts of 15 agents each at startup. Values as for Figure 1.

reproduced. Those findings then seemed to show that a simple algorithm for altruism evolved thanks to roost competition. The present ones, in turn, show that such

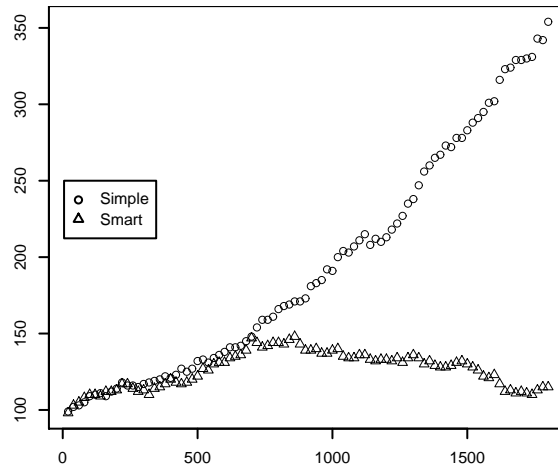


Fig. 3. Reproduction among simple and smart vampires. Comparison of simulations with 100 agents at startup. Values as for Figure 1.

a simple algorithm is rather efficient and powerful, doing even better than a more complex one. Why bother with higher mental complexity when a simpler algorithm accounts for satisfactory results? Although performed by stupid agents, the simple rule for food-sharing is probably not the result of one single mutation, but rather the effect of a process of trials and errors that has led the species to select and mould the most convenient algorithm under given ecological pressures. During this process, vampires “learned” when, how and to what extent to be altruistic. Can we imagine an alternative way to achieve the same result? One would say that smart systems learn during their experience to do something at least comparable with what is learned by the species during its evolution. Is this true?

5. Second Study: Dynamic Smart Systems

In the second phase of our study, we set to model a dynamic variant of smart agents. Goal-dynamics is an essential aspect of cognition.^{2,3} In a cognitive architecture, a goal is a highly dynamic mental construct, which, thanks to beliefs, may be generated, abandoned, worked out, suspended, interrupted, achieved, compromised, etc. We introduced in our smart algorithm a small number of rules modifying the values of agents’ goals, in particular of the normative goal. This was supposed to increase or decrease as an effect of others’ impact on one’s own conditions. If one receives help, the force of the altruistic motivation increases, whereas it decreases if one is denied help. This effect is even stronger, if one is not reciprocated and its credits are therefore not extinguished. What happens under these circumstances? What are the effects of goal-dynamics on altruism, and more specifically on the fitness of the smart population?

5.1. *The Experiment*

In the second battery of experiments, we have studied the behaviour of a population of dynamic smart agents. The motivational value for survival (SG) has been fixed to 0, while the normative motivation is left free to change. The normative value is incremented by altruistic acts: the experience of an effective application of the norm (donation and, to a greater amount, reciprocation) raise the value of NG, while the increase of unreciprocated donations lowers it. As a consequence, the strategy applied by an agent is dictated by the value of NG:

Normative Goal	Strategy
$NG < -2$	Cheater
$-2 \leq NG < 0$	Prudent
$0 \leq NG \leq 1$	Fair
$1 < NG \leq 4$	Generous
$NG > 4$	Martyr

5.2. *Findings*

Simulations showed that different strategies — corresponding to different patterns of relationships among agents' goals — emerge in the experimental condition (see Figure 4, left), and their difference increases when strategies are inherited by the agents' offspring. In this condition, the dominant strategy appears to be martyrdom. This should not come as a surprise since this strategy is self-reinforcing: as soon as cheaters start to decline, donations will increase and reciprocity will emerge. Consequently, the normative goal will grow almost ad libitum.

Without inheritance, instead, strategies will not differentiate, and the population gradually and slowly decreases (Figure 4, right). Apparently, findings are qualitatively indifferent to the initial strategy (see Figure 6). If no upper limit is introduced to the increasing value of the normative goal, martyrdom proves to be the fittest strategy.

By introducing an upper limit to the increase of the normative goal's value, strategies appear to be more wrapped up. None of them far exceeds all others, but the population varies periodically (Figure 5).

No selective mechanism of retaliation is actually implemented in our model. To simulate retaliation, we can make agents deny help to their debtors if these did not extinguish their debts first. This mechanism is not really selective, since even generous but unlucky agents can be blacklisted. The effect of this kind of retaliation is dramatic: the more severe it is, the faster the extinction of altruistic strategies and of the whole population (Figure 7).

In short, altruism emerges and spreads in populations of dynamic goal-based systems, at least in a multi-roost world with a relatively small number of agents per




Fig. 4. Emergence of altruistic strategy, with (left) and without (right) inheritance. Values are number of agents per strategy (y), and ticks(x). 10 roosts of 15 agents each populate the world. Every agent is Prudent at the beginning (reproducing the Simple behaviour). Note the difference in scale between the vertical axes.

roost (around 15 individuals) — which is not far away from the actual distribution in nature — and mild retaliation. This is so even when the initial population includes a high proportion of cheaters (up to 70%). In the control condition, instead, simple altruists can put up with a minority of cheaters (up to 20%: with higher percentages, the roosts extinguish).

6. Final Discussion

The findings discussed so far show a remarkable impact of intelligent, goal-based systems on the spreading of altruism, provided these systems are highly dynamic. Whereas rigid goal-based systems are worse-off than agents applying a given routine, dynamic ones provide a much more stable effect on altruism, neutralising even a majority of cheaters.

However, these results must be analysed carefully and accepted with some caution. They are obtained under given conditions, from which some lessons can be drawn.

First, the fundamental mechanism that is probably at the origin of altruism evolution is grouping, or roosting. Vampires living in a uni-roost world have no chance to survive, whether they share food or not.

Secondly, and somewhat analogously, the number of agents per roost plays a role as well. The larger the roost, the lower the possibility that a credit network emerges.




Fig. 5. Upper limit, for long timescale (40000 ticks). None of the strategies seem to prevail. Values and initial agents as for Figure 4. Initial behavior is equally distributed among all five strategies.

But without such a network, reciprocity has no chance to emerge, and cheaters will exploit altruists to death. Consequently the entire population will collapse, since in the long run cheating is self-defeating.

A third factor to be considered is retaliation. No selective mechanism of retaliation is actually implemented in our model. However, agents will deny help to their debtors if these did not extinguish their debts first. This rule, if rigidly applied, reduces exploitation but penalises more generally all unlucky agents, regardless of whether they are altruists or cheaters. The milder variant of the rule performs much better (such that no more than two requests in a row by the same debtor are accepted). It is less efficient with regard to cheaters, but it is less unfair with regard to unsuccessful altruists.

Under all of these conditions, which are not unrealistic by the way, dynamic smart agents perform exceptionally well. Even when cheaters represent a vast majority of the population, a handful altruists are able to spread their memes to the whole population (altruistic strategies proselyte horizontally), and their genes to future generations (altruistic strategies reproduce more than cheaters). Indeed, one of the interesting result obtained is the convergence between memetic and genetic fitness. In our simulations, success in terms of proselytes overlaps with success in terms of fitness. However, the rationale for this correlation, which should be further investigated, does not reside in the strategies' perceived success or popularity, but rather on their good effects: people donate if they get help.

In our simulations, martyrdom appears as both highly fit and very good at




Fig. 6. Relative impact of initial strategies. On the left, the initial population is composed by martyrs and cheaters only. On the right, initial population is equally distributed among all five strategies. Values and initial agents as for Figure 4.

making proselytes. Should we consider such a result in the light of corresponding historical examples (see the origins of catholic religion)? Probably not, since this effect is mainly due to an unlimited growth of the normative goal's value. When an upper limit is set, this strategy behaves more or less like all others. Differentiation among strategies lowers, but the whole population still grows. To what extent? Is such a growth unlimited? Again, probably not. The fate of the population is linked to the chances that, by means of learning, cheaters reappear every now and then. When they grow, population declines, until cheaters collapse again. Thereafter, the population starts to increase anew, like in Figure 5. In a follow-up study, the width and frequency of these waves ought to be investigated. A further interesting aspect of the dynamic algorithm explored in this paper resides in its learning mechanism. Whereas in the literature on adaptive systems, learning is usually obtained by means of imitation of either the most frequent or the fittest behaviours, our dynamic agents are more likely to be generous when helped by others, and less when their requests are denied. In real societies, agents cannot perceive the average fitness or frequency of social behaviours, but can easily observe their effects on themselves. Whether dynamic intelligence allows agents to accomplish in their lives the same learning trajectory that the species followed during its evolution, is something we are not entitled to say on the grounds of these findings. Analogously, we are not entitled to infer anything as regard the actual mental constructs underlying vampires' food-sharing. All we can say is that, far from hindering it, intelligence may contribute




Fig. 7. Severity of retaliation. Rapid extinction for initial population of prudent (left) and of martyrs and cheaters (right). Values and initial agents as for Figure 4.

to altruism, at least when it is flexible, by keeping low and inoffensive the cheating sub-population. The question is, why it does not always do so?

References

1. Bowles, S. and Gintis, H. (2001), The Evolution of Strong Reciprocity, Santa Fe Institute Discussion Paper.
2. Castelfranchi C. (1997). Principles of Individual Social Action. In Tuomela R., Hini-tikka G. (Ed.), Contemporary Action Theory (pp. 24-29). Kluwer, Dordrecht.
3. Conte R. (2000). Memes through (Social) Minds. In Aunger R. (Ed.), Darwinizing Culture. The Status of Memetics as a Science. The Oxford University Press.
4. Conte, R. and Castelfranchi, C. (1995). Cognitive and Social Action. Univ. College of London Press.
5. Cosmides, L. and J. Tooby (1994), Better than rational: Evolutionary psychology and the invisible hand, American Economic Review - Papers and Proceedings, 84 (2):327-332.
6. Dawkins, R. (1976) The Selfish Gene, Oxford Univ. Press.
7. Dean, A. (2000) Complex Life: Non Modernity and the Emergence of Cognition and Culture. Aldershot: Ashagate.
8. Dorigo, M.; Birattari, M.; Blum, C.; Gambardella L.M.; Mondada, F.; Sttzle, T. (2004) Ant Colony Optimization and Swarm Intelligence. Proceeding of the 4th International Workshop, ANTS 2004, Brussels, Belgium, September 5-8 2004, Springer.
9. Kruger, D.J. (2002). What is Evolutionary Psychology? Ann Arbor, MI: Altralological Press. Available On-line: <http://www-personal.umich.edu/~kruger>.
10. Kruger, D.J. (2003). Evolution and altruism: Combining psychological mediators with naturally selected tendencies. Evolution and Human Behavior, 24, 118-125.

11. Nowak, M.A. and Sigmund, K. (1998) Evolution of Indirect Reciprocity by Image Scoring. *The Dynamics of Indirect Reciprocity*, *Nature*, 393, 11 June.
12. Paolucci, M., Di Tosto, G., Conte, R.. (2003) Reciprocal Vs Group Altruism Among Vampires, NAACSOS Methods, Toolkits, and Techniques, The University of Chicago, October 3-4, 2003.
13. Simon, Herbert A. (1981) *The sciences of the artificial*. Cambridge Mass., London, MIT Press.
14. Simonin, O. and Gechter, F. (2006) An environment-based principle to design reactive multi-agent systems for problem solving. Second International Workshop on Environments for Multiagent Systems, E4MAS'05.
15. Trivers, R. (1972). The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46, 35-37.
16. Tversky, A. and Kahneman, D. (1974) Judgment under uncertainty: Heuristics and biases. *Science*, 185, 3-20.
17. Wilkinson, G.S. (1984) Reciprocal food sharing in the vampire bat. *Nature* 308, 181-184.
18. Wilkinson, G.S. (1990) Food sharing in vampire bats. *Scientific American* 2, 64-70.