

The Emergence of Commitments and Cooperation

The Anh Han
Centro de Inteligência Artificial
Departamento de Informática,
Faculdade de Ciências e
Tecnologia
Universidade Nova de Lisboa,
2829-516 Caparical
Portugal
h.anh@campus.fct.unl.pt

Luís Moniz Pereira
Centro de Inteligência Artificial
Departamento de Informática,
Faculdade de Ciências e
Tecnologia
Universidade Nova de Lisboa,
2829-516 Caparical
Portugal
lmp@di.fct.unl.pt

Francisco C. Santos
Centro de Inteligência Artificial
Departamento de Informática,
Faculdade de Ciências e
Tecnologia
Universidade Nova de Lisboa,
2829-516 Caparical
Portugal
fcsantos@fct.unl.pt

ABSTRACT

Agents make commitments towards others in order to influence others in a certain way, often by dismissing more profitable options. Most commitments depend on some incentive that is necessary to ensure that the action is in the agent's interest and thus, may be carried out to avoid eventual penalties. The capacity for using commitment strategies effectively is so important that natural selection may have shaped specialized capacities to make this possible. Evolutionary explanations for commitment, particularly its role in the evolution of cooperation, have been actively sought for and discussed in several fields, including Psychology and Philosophy. In this paper, using the tools of evolutionary game theory, we provide a new model showing that individuals tend to engage in commitments, which leads to the emergence of cooperation even without assuming repeated interactions. The model is characterized by two key parameters: the punishment cost of failing commitment imposed on either side of a commitment, and the cost of managing the commitment deal. Our analytical results and extensive computer simulations show that cooperation can emerge if the punishment cost is large enough compared to the management cost.

Categories and Subject Descriptors

I.2 [Artificial Intelligence]: Multiagent systems

General Terms

Experimentation, Theory

Keywords

Evolution of Commitment, Evolution of Cooperation, Evolutionary Game Theory, Prisoner's Dilemma

1. INTRODUCTION

Over the last few years, several mechanisms have been pointed out to promote the emergence and maintenance of

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), June, 4–8, 2012, Valencia, Spain.

Copyright © 2012, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

cooperation. From group and kin relations, memory and reputation based reciprocity mechanisms, to social diversity and context based reactions, grounded or not on incipient levels of cognition, there has been a large improvement on our capacity to understand the roots of animal and human cooperation [3, 2, 20, 27, 42, 45, 14, 15]. They are certainly hugely important, but are they sufficient? Or, as many have suggested [5, 19, 22], might there be other routes to social behavior that have been neglected? Certainly there are. Commitment, which amounts to expressing an intention rather than having it recognized, may stand as another route to cooperation, even in its simplest form, as we purport to show here.

Agents make commitments towards others when they give up options in order to influence others. Most commitments depend on some incentive that is necessary to ensure that the action is in the agent's interest and thus will be carried out [12], on pain of some heavy penalty. Committers also incur in a small cost when proposing or setting up a commitment so as to make it credible upfront to others, and entice these to accept to commit.

The capacity for using commitment strategies effectively is so important that natural selection may have shaped specialized signaling capacities to make this possible [49, 37, 41, 26, 8, 4]. And it is believed to have an incidence on the emergence of morality [38]. Assuming cooperation to be, at best, just the result of individuals' purely competitive strategies can make it conceptually unstable [31], most especially in non-iterated or history-free interactions. And it seems possible that the spread of simplistic notions, rooted in science, about the evolutionary origins of social relationships could foster a trend to make these relationships more conflicted, and society more brutal. An antidote is an evolutionary approach to behavior that incorporates a capacity for mutual commitment, shown advantageous for all concerned [26], even in non-iterated or memory-free settings.

Hence, our goal in this paper is to examine, through Evolutionary Game Theory (EGT) [20, 45], how the most simple of commitment strategies work, and how they can give rise to the emergence of cooperation. We shall do so in the setting of the non-iterated Prisoner's Dilemma (PD), a well-known game-theoretical framework to study the evolution of cooperation within populations of self-interested agents [2, 20, 27, 45]¹. In an interaction, each player has two options,

¹There are other social dilemmas, such as the Stag Hunt and the Chicken Game. [45], but the Prisoner's Dilemma is

cooperates (C) or defects (D), and defect is the dominant option – it is always better to defect in a one-shot interaction. Both players should choose to defect, while they would be better off by choosing to cooperate instead, thus leading to the destruction of social welfare and individuals’ fitness.

In a nutshell, convincing others of one’s credibility in a commitment proposal amounts to submit to options that change the incentives of the situation. These options, namely commitment cost and penalty for defaulting, can be expressed by the payoffs specified in a game. When opponent players observe meticulously such payoffs, and realize that compliance with a proposed commitment is in the proposing player’s best interests, then, given any opponent player’s open option to commit, these may change their expectations and behavior accordingly, and adopt as a result a strategy which either accepts commitment proposals or ignores them. In general, there are four main reasons to believe a commitment will be fulfilled [26]: i) a commitment can be self-reinforcing if it is secured by incentives intrinsic to the situation; ii) a commitment can be secured by external incentives controlled by third parties; iii) a commitment can be backed by a pledge of reputation; and iv) a commitment can be reinforced by internal emotional motives.

The first two types are secured in much the same way a loan is secured by a collateral. They objectively change the situation so that fulfillment becomes in the individual’s best interests. The latter two types do not change the objective contingencies; they are subjective commitments in that they may involve a continued option of renegeing, according to some or other stance extraneous to the game’s given payoffs matrix.

In our EGT setting however, we will simply assume that a game’s payoff matrix, concerning a set of strategies, summarily ingrains and expresses in its structure the impingement of all such contingencies. For instance, often a capacity for commitment allows individuals to act in ways that reap the benefits of image scoring through maintaining a reputation, or the access of others to a social history of prior interactions. In this study, for simplicity but also for exhibiting the purity and power of the commitment mechanism, we ignore the effect of repeated interactions [52], and of any reputation [29, 32] associated with particular individuals. We aim to show that the simplest of core commitment mechanisms can improve cooperation, and leave any other complications for the future, most promisingly how commitment can be combined with and reinforce other known mechanisms of cooperation. And perhaps surprisingly we can do so. Thus, no credibility of commitment will be taken into account [6] beyond that which is expressed in a game’s payoff matrix. No reputation appraisal of the commitment proposer is made by its co-player, and no historical or social data is even available to do so. Each pairwise interaction is purely based on fixed individual strategies that might involve commitment or the lack thereof. Also, no “cheater or deceit detection” or “intention recognition” is in place [14, 15]. Nevertheless, systematic unconditional bluffing on the part of a player is a possible fixed feature of its strategy, in the sense that, from the start, the player does not intend to fulfill commitments.

It will be seen in our commitment model that players defaulting on their commitments, be they the proposing or the known to represent one of the most difficult or fierce environments for cooperation to emerge.

accepting party, will be subject to evolutionary disadvantage for a wide range of parameters.

We show that more elaborate commitment strategies are not strictly necessary for commitment to become evolutionarily advantageous. Neither an aptitude for higher cognition, nor for empathy, nor for mind reading are needed. These aptitudes would only be required for more sophisticated forms of commitment, scaffolded atop the core one. We will explain the evolution, in a population, of the capacity for a simple form of commitment as the result of otherwise being excluded from a group of committed promise abiding cooperators, in the sense that this strategy tends to invade the game playing population under rather general conditions.

The remainder of this paper is organized as follows. In Section 2, we discuss the relevant literature. In Section 3, our EGT commitment model and its methods are defined and explained. Forthwith, in Section 4, we proffer results obtained with the model, both analytic and via numeric and computer simulations. We conclude the paper with a discussion section on commitment and its EGT modeling.

2. RELATED WORK

Evolution of cooperation has been a central research topic of many fields, including Biology, Economics, Artificial Intelligence, Political Science and Psychology [2, 20, 27, 42, 45, 15, 21]. Several mechanisms responsible for promoting cooperative behavior have been recently identified (see surveys in [27, 45]). In these contexts, several aspects have been shown to play an important role in the emergence of cooperation. Differently, our model does not require any of those aspects, namely it does not assume kinship or in-group relatedness of agents, nor repeated interactions or reputation consideration, nor concrete structures of population distribution. However, we envisage that the mechanism of commitment could reinforce the existing mechanisms of cooperations, e.g., easing the conditions for the emergence of cooperation therein. This will be the subject of the future work.

Evolutionary explanations of commitment, particularly its role in the evolution of cooperation, have been actively sought for and discussed in several fields, including Psychology and Philosophy [26, 12, 18, 6, 8, 4, 38]. But there are only a few computational models that show the evolutionary advantages of commitment in problems where cooperative acts are beneficial [49, 37, 41]. In addition, often models rely on repeated interactions or long-term relationships [8, 4], alike the conditions where Triver’s direct reciprocity [52] may play a role. Here we provide an analytic model in the framework of evolutionary game theory showing that, with the availability of the mechanism of commitment, cooperation can emerge even without assuming repeated interactions.

Last but not least, it is undoubtedly important to mention the extensive literature of AI and Multi-agent System research on commitment, e.g., [43, 54, 18, 6, 53, 16, 7]. The main concern therein is how to formalize different aspects of commitment and how a commitment mechanism can be implemented in multi-agent interactions to enhance them (e.g. for improved collaborative problem solving [54]), especially in the context of game theory. In contradistinction, our concern is in the nature of an evolutionary explanation of commitment, particularly how it can promote the emergence of cooperation.

3. MODELS AND METHODS

3.1 Model

Let us consider a commitment variant of the Prisoner's Dilemma game in which a new type of cooperator (denoted by COM_C) that, before each interaction, asks the co-player whether it commits to cooperate. If the co-player does not so commit, there is no interaction. Both players get 0. Otherwise, if the co-player commits, they then go on to play with each other in the present interaction. If the co-player keeps to its commitment, both players obtain the reward payoff, R^2 . Otherwise (if the co-player fails its commitment), the proposing or focal player obtains the sucker payoff, S , and its co-player obtains the temptation payoff, T . However, the one that fails the commitment will suffer a penalty cost, and its non-defaulting co-player gains a compensation for the potential loss due to its default of fulfilling the commitment. For simplicity, we assume that these two amounts (penalty and compensation) are equal, being denoted by δ . The penalty cost can be a real monetary one, e.g., in the form of prior debit (e.g., in the case of accommodation rental) or of a subsequent punishment cost (e.g., commitment was performed in terms of a legal contract, and one who fails commitment must pay a cost to compensate for the other), or an imaginary abstract value, e.g., public spread of good/bad reputation (bad reputation for the one that fails, and sympathy for the other), or even an emotional suffering [26, 12, 18, 38]. How this cost is set up depends on the types of commitment at work, or the reason for which the commitment is believed to be fulfilled (see Introduction), which topic is beyond the scope of this paper. However, various techniques can be seen in [43, 18].

Two players that defect in an interaction obtain the punishment payoff, P^3 . As usual, for the Prisoner's Dilemma, the payoff entries satisfy the ordering, $T > R > P > S$, whereas the four possible outcomes can be written down as a payoff matrix

$$\begin{array}{cc} & \begin{array}{cc} C & D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \begin{pmatrix} R, R & S, T \\ T, S & P, P \end{pmatrix} \end{array}$$

For setting up a commitment, the proposer must pay a small management cost, ϵ . The cost of proposing and setting up the commitment might be high, but it is reasonable to assume that this cost is quite small compared to the mutual benefit of a cooperation strategy guaranteeing commitment, $\epsilon \ll R$.

Given the nature of a situation expressed in terms of payoff entries, one can naturally expect that if a proposed punishment cost, δ , is high enough compared to the cost of managing the commitment, ϵ – to convince and guarantee that cooperation is in the proposer's interest and also drive away potential exploiters – cooperation can emerge, even in the fierce environment of the Prisoner's Dilemma. This penalty

²Note that here we do not yet take into account execution noise (see, e.g., [45, 32]), i.e. the agents might misimplement their intended choice, from cooperate to defect or vice versa. Thus, COM_C will never misimplement the intended commitment choice.

³For the sake of a clear representation, in our analysis we adopt $P = 0$ [40, 45] (as in the Donation game), even if the more general case can be analyzed in the same manner and portray similar results to the ones presented below.

and management relation is subject to detailed study below, both analytically and by means of computer simulations.

We consider a finite population of a constant size, consisting of four strategies: COM_C (as described above), C (always cooperates, without proposing to commit), D (always defects, and does not commit when being asked to), and D_COM (always defects, though commits when being asked to). Here, we assume that cooperators, including COM_C and C players, always commit whenever being asked to since they are better off to do so, as cooperation is their default choice, and reasonable commitment deals only are proposed. Hence, for the sake of exposition, the two other (unreasonable) strategies, those of cooperators that refuse to commit and of defectors that propose commitment, are omitted here (they would become eliminated anyway). The former is dominated by the pure cooperator strategy, C, while the latter is by the pure defector strategy, D.

In each round, two random players are chosen from the population for an interaction. For the row player, the (average) payoff matrix reads

$$\begin{array}{cccc} & \begin{array}{cccc} COMC & C & D & DCOM \end{array} \\ \begin{array}{c} COMC \\ C \\ D \\ DCOM \end{array} & \begin{pmatrix} R - \epsilon/2 & R - \epsilon & -\epsilon & S + \delta - \epsilon \\ R & R & S & S \\ 0 & T & P & P \\ T - \delta & T & P & P \end{pmatrix} \end{array} \quad (1)$$

Note that when a COM_C interacts with another COM_C, only one of them pays the cost of having proposed commitment, ϵ (e.g., the arbitrary one that proposes). Therefore, the average payoff of a COM_C in playing with another COM_C is, $R - \epsilon/2$.

3.2 Methods

Our analysis is based on evolutionary game theory methods for finite populations [28, 23]. In the context of evolutionary game theory, the individuals' or agents' payoff represents their *fitness* or social *success*. The dynamics of strategy change in a population is governed by social learning, that is, the most successful agents will tend to be imitated by the others. There are many ways to model social learning [20, 45, 36]. Adopting one of the most frequently used ones, we shall consider the so-called pairwise comparison rule [51], which assumes that an agent A with fitness f_A adopts the strategy of another agent B with fitness f_B with probability given by

$$\frac{1}{1 + e^{-\beta(f_B - f_A)}}$$

where β controls the 'imitation strength', i.e., how strongly the agents are basing the decision to imitate on fitness comparisons. For $\beta = 0$, we obtain the limit of neutral drift – the imitation decision is random. For large β , imitation becomes increasingly deterministic.

In the absence of mutations, the end states of evolution are inevitably monomorphic: once such a state is reached, imitation cannot produce change. We thus further assume that, with a certain mutation probability $\mu > 0$ (also dubbed the exploration rate [50]), an agent switches randomly to a different strategy without imitating another agent. The resulting Markov Chain has a stationary distribution, which characterizes the average time the population spends in each of these monomorphic end states. Yet, for arbitrary exploration rates and number of strategies, stationary distribu-

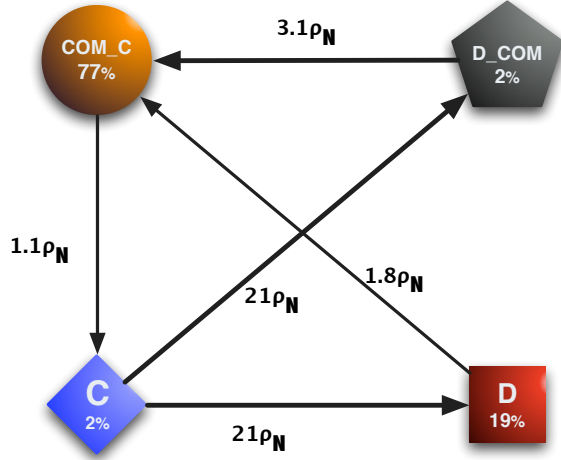


Figure 1: Stationary distribution and fixation probabilities. The population spends most of the time in the homogenous state of COM_C. The black arrows stand for the transitions that are rather stronger than neutral. The strongest transitions are from C to D and D_COM, and the slowest one is from COM_C to C. There are rather strong transitions from D and D_COM to COM_C. Parameters: $T = 2$, $R = 1$, $P = 0$, $S = -1$; $\delta = 4$; $\epsilon = 0.05$; imitation strength, $\beta = 1$; population size, $N = 100$; $\mu_N = 1/N$ denotes the neutral fixation probability.

tions are often cumbersome to compute [17, 46, 39].

Fortunately, in the case of small exploration or mutation rates, analytical computation of this stationary distribution can be conveniently computed [11, 23, 17, 41]. The small exploration rates guarantee that, any newly occurred mutant in a homogeneous population will fixate or become extinct long before the occurrence of another mutation. Hence, the population will always consist of at most two strategies. This allows one to describe the evolutionary dynamics of our population in terms of a reduced Markov Chain, whose size is equal to the number of strategies being considered (which is 4 in our case), and each state represents a possible monomorphic end state of the population associated with a one of the strategies. The transitions between states are defined by the fixation probabilities of a single mutant of one strategy in a homogeneous population of individuals adopting another strategy (see Figure 1 for better understanding).

$$T^\pm(k) = \frac{N-k}{N} \frac{k}{N} \frac{1}{1 + e^{\mp\beta[\Pi_A(k) - \Pi_B(k)]}} \quad (2)$$

More precisely, let N be the size of the population. Suppose there are at most two strategies in the population, say, k agents using strategy A ($0 \leq k \leq N$) and $(N-k)$ agents using strategies B. Thus, the (average) payoff of the agent that uses A and B can be written as follows, respectively,

$$\begin{aligned} \Pi_A(k) &= \frac{(k-1)\pi_{A,A} + (N-k)\pi_{A,B}}{N-1} \\ \Pi_B(k) &= \frac{k\pi_{B,A} + (N-k-1)\pi_{B,B}}{N-1} \end{aligned} \quad (3)$$

where $\pi_{X,Y}$ stands for the payoff an agent using strategy X obtained in an interaction with another agent using strategy Y , given in the payoff matrix (1).

Now, the probability to change the number k of agents using strategy A by \pm one in each time step can be written as

The fixation probability of a single mutant with a strategy A in a population of $(N-1)$ agents using B is given by [51, 25, 11, 23, 17]

$$\rho_{B,A} = \frac{1}{1 + \sum_{i=1}^{N-1} \prod_{j=1}^i \frac{T^-(j)}{T^+(j)}} \quad (4)$$

In the limit of neutral selection ($\beta = 0$), $T^-(j) = T^+(j) \forall j$. Thus, $\rho_{B,A} = 1/N$. Considering a set $\{1, \dots, q\}$ of different strategies, these fixation probabilities determine a transition matrix $M = \{T_{ij}\}_{i,j=1}^q$, with $T_{ij, j \neq i} = \rho_{ji}/(q-1)$ and $T_{ii} = 1 - \sum_{j=1, j \neq i}^q T_{ij}$, of a Markov Chain. The normalized eigenvector associated with the eigenvalue 1 of the transposed of M provides the stationary distribution described above [25, 11, 23, 17], describing the relative time the population spends adopting each of the strategies.

Now let us recall some important analytic measures which will be used in our analytical study. In a pair-wise comparison of strategy A with strategy B, we say that A is advantageous (against B) if an A mutant has a fixation probability in a population of agents using B greater than that of the neutral selection (which equals the inverse of population size, $1/N$) [28, 27, 45]. Interestingly, it was shown that this condition holds if

$$(N-2)\pi_{A,A} + (2N-1)\pi_{A,B} > (N+1)\pi_{B,A} + (2N-4)\pi_{B,B} \quad (5)$$

which, in the limit of large N , is simplified to

$$\pi_{A,A} + 2\pi_{A,B} > \pi_{B,A} + 2\pi_{B,B} \quad (6)$$

Another important measure to compare the two strategies A and B is which direction the transition is stronger or more probable, an A mutant fixating in a population of agents using B or a B mutant fixating in the population of agents using A. It can be shown that the former is stronger if [24, 45]

$$(N-2)\pi_{A,A} + N\pi_{A,B} > (N-2)\pi_{B,A} + N\pi_{B,B} \quad (7)$$

which, in the limit of large N , is simplified to

$$\pi_{A,A} + \pi_{A,B} > \pi_{B,A} + \pi_{B,B} \quad (8)$$

4. RESULTS

We compute the fixation probabilities and stationary distribution numerically for small mutation or exploration rates (see Methods). The population spends most of the time in the homogeneous state where all individuals utilize the commitment strategy (Figure 1).

In general, amongst the monomorphic states of the population, the strongest transitions are from C to D and C to D_COM. The difference of a small cost of proposing commitment, ϵ , between COM_C and C, leads to a near-neutral transition from COM_C to C. The more intricate transitions are between COM_C and D or D_COM, which are the central part of our analysis.

Between D and COM_C, for $\epsilon \ll R$, COM_C is advantageous. Namely, by a pairwise comparisons of COM_C and D

[28, 27] that condition always holds if (in the limit of large N , see Eq. (6))

$$\epsilon < \frac{2R}{5} \quad (9)$$

This inequality also guarantees that, for a population of size $N > 4$, the more probable transition is from D to COM_C, i.e., satisfying that [24, 45] (see Eq. 8)

$$(N-2)(R - \frac{\epsilon}{2}) - N\epsilon > 0 \quad (10)$$

Similarly, for big enough δ , COM_C is advantageous against D_COM; namely, if

$$\delta > \frac{T-R-2S}{3} + \frac{5\epsilon}{6} \quad (11)$$

It guarantees that the transition of D_COM to COM_C is more probable than the opposite if

$$(N-2)(R - \frac{\epsilon}{2}) + N(S + \delta - \epsilon) > (N-2)(T - \delta) \quad (12)$$

which holds if

$$\delta > \frac{N}{2N-2}(T-R-S + \frac{3\epsilon}{2}) \quad (13)$$

For large enough N , it is simplified to

$$\delta > \frac{T-R-S}{2} + \frac{3\epsilon}{4} \quad (14)$$

Hence, for

$$\delta > \max\{\frac{T-R-S}{2} + \frac{3\epsilon}{4}, \frac{T-R-2S}{3} + \frac{5\epsilon}{6}\} \quad (15)$$

the transition of D_COM to COM_C is the more probable one, as well as greater than neutral.

Taking together with the fact that the transition of COM_C to C is near neutral, one can expect that if the two parameters δ and ϵ satisfy the inequalities (9) and (15), COM_C will prevail – the population will spend most of the time in its homogenous state. This expectation is supported by the numerical results in Figures 2 and 3. For a given payoff matrix of the PD, for strong enough punishment cost of failing commitment, δ , and small enough cost of setting up the commitment, ϵ , the population spends most of the time in the homogeneous state of COM_C (Figure 2). In addition, this result is also flexible with respect to the payoff values of the PD (Figure 3). For the sake of a clear representation of the result, we use in this numerical experiment the Donation game [46] – a special case of PD – where $T = b$, $R = b - c$, $P = 0$, $S = -c$, satisfying that $b > c > 0$; b and c stand for “benefit” and “cost” of cooperation, respectively. It shows that, for given δ and ϵ , for large enough b/c , the population spends most of the time in the homogeneous state of COM_C.

So far, our analytic and numerical results were obtained in the limit of small mutation rates. Next, by extensive computer simulations, we show that this remarkable performance of the commitment strategy COM_C is flexible with respect to mutation rates (Figure 4). Namely, for all the mutation rates up to 0.1, the population always spends most of the time in the homogenous state of COM_C. It also noteworthy, that our analytic results for small imitation strengths and under the extremes of low and high mutation or exploration rates — based on the methods described in [1] — comply with this simulation results.




Figure 2: Contour plot of the frequency of COM_C as a function of ϵ and δ . In a population of COM_C, COM_D, C, and D individuals, for a wide range of ϵ and δ , the population spends most of the time in the homogeneous state of COM_C. The smaller the cost of proposing commitment, ϵ , and the greater the punishment cost of failing commitment, δ , the greater the frequency of COM_C. The payoffs being used are, $T = 2$, $R = 1$, $P = 0$, $S = -1$; imitation strength, $\beta = 1$; population size, $N = 100$.

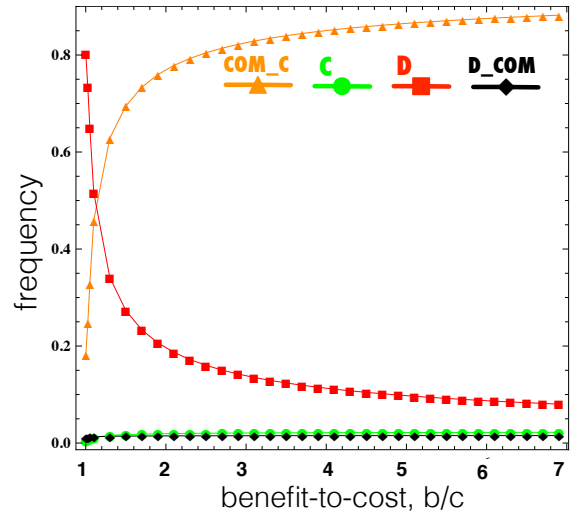


Figure 3: Frequency of each strategy as a function of benefit-to-cost ratio, b/c , for Donation game ($T = b$, $R = b - c$, $P = 0$, $S = -c$, with $b \geq c$). In a population of COM_C, COM_D, C, and D individuals, for a large enough benefit-to-cost ratio, the population spends most of the time in the homogeneous state of COM_C, while D prevails when this ratio is very small. Parameters: $\delta = 4$; $\epsilon = 0.05$; imitation strength, $\beta = 1$; population size, $N = 100$.

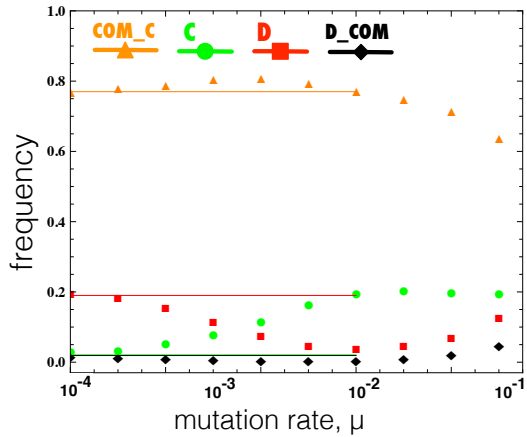


Figure 4: Frequency as a function of mutation rates. Symbols indicate results from computer simulations (averaged over 10^9 update steps), and solid lines show the numerical approximation results for small mutation rates. The population consists of COM_C, COM_D, C, and D individuals. COM_C always dominates for all the mutation rates up to 0.1. Parameter values (the same as in Figure 1): $T = 2, R = 1, P = 0, S = -1; \delta = 4; \epsilon = 0.05$; imitation strength, $\beta = 1.0$; population size, $N = 100$. The simulation results are obtained by averaging 40 runs, and the initial population is equally likely to be in one of the homogenous states.

All in all, our study exhibits that, in spite of the absence of repeated interactions, reputation effect, network reciprocity, as well as group and kin selection, the strategy of commitment proposal may enable the emergence of cooperation. By imposing a high cost for failing a commitment, when compared to the cost of setting up or managing the commitment deal, the commitment cooperative agents COM_C can get rid of the fake committers (D.COM) as well as avoid being exploited by the pure defectors (D), while playing approximately equally well against the pure cooperators (C). The results of this study suggest that our specialized capacity for commitment, which might have been shaped by natural selection [26] consists in a capacity for managing to impose a high cost of punishment, whether it is monetary or of abstract emotional or reputation value, with a relatively small cost.

We note that there is a significant difference between our commitment model and the works on costly punishment [30, 35, 9, 17, 10]. A commitment deal must be agreed by both sides of it in advance, thereby giving credibility and justification to punish any defaulting player. In addition, the prior agreement gives rise to compensation—the amount of which, in some cases, is agreed explicitly in advance—to the non-defaulting player. This compensation for the non-defaulting player is the significant difference that makes successful those players using the commitment strategy, while those using the costly punishment strategy have only a narrow margin of efficiency [30]; does not stand out as a winning strategy [9]; nor does it promote cooperation at all when taking into account antisocial punishment [35]. The compensation might bring benefit to the commitment strategists once

an appropriate deal would be arranged.

This suggests that although costly punishment, whether it is social or antisocial, might not promote the evolution of cooperation, what we call ‘justified’ punishment, which is warranted by an appropriate commitment deal, does. This kind of punishment might not be costly at all, and can even bring net benefit to its upholder, hence leading to the emergence of cooperation.

5. DISCUSSIONS

Within the general game theory concept of commitment, several distinctions can help separate different subtypes. In particular, some commitments are upfront promises of a next move that can help, while others are upfront threats of a subsequent move that can harm. Commitments can be conditional or unconditional. Threats are usually attempts to influence another person’s next move by stating a conditional subsequent move, and that’s how we may envisage them. Promises are more likely to be unconditional, and that’s how we may conceive of them, though more generally they can be conditional on the other fulfilling a matching promise. Concerning this, we note a difference between a commitment and a convention. A convention is a means for monitoring a commitment: it specifies under what circumstances a commitment can be abandoned and how an agent should behave both locally and towards others when one of these conditions arises [54]. Commitments can also be just towards oneself, taking into account the evolution of possible futures afforded by actions and events, and the individual’s prior and post preferences, in what might be classically seen as a game against nature.

In [34, 33], three different types of individual commitment – hard, revocable, and momentary – are studied in such an evolution context. Let us recall that commitment, in the context of game theory, is a device or mechanism to decide the outcome with the other party [43]. Schelling distinguishes between commitment pure and simple and commitment that takes the form of a threat. What he calls ‘ordinary’ commitment corresponds, in game theory, to the making of an opening announcement in a sequential play, which we dub preemptive, just before both players make their actual move. To constitute a preemption, a player’s announcement action must be irrevocable, that is a promise that is assuredly kept. Preemptive commitment is not necessarily profitable, because it hinges on the opponent’s actual move. Schelling however does not assume the other type of commitment as a ‘threat’, which pertains to the a player’s move in reaction to the opponent’s move. Threats, being conditional, may be of the ‘if-then-else’ form, and can thus combine a threat and a promise, the latter albeit implicit whenever there are just two possible moves. We prefer instead to label ‘reactive’ such so-called threat commitments. In the game context, these occur when the player with the last move irrevocably pledges to respond, in a specified but contingent way, to the opponent’s prior choice [19].

In a nutshell, some players can be ‘preemptive’ committers – those that always propose and always accept proposed commitments–, others may be ‘reactive’ committers – those that always make a ‘reactive’ statement and comply with the implicit requests in such statements–, while other players, though accepting to commit nevertheless default on their commitment, and even others simply omit and ignore preemptive or reactive commitments in their strategies –

they might for instance be persistent defectors or persistent cooperators as we have seen, or, for that matter, follow any other strategy ignorant of commitment. Moreover, in iterated games, commitments can concern future rounds and not just the present one.

We purport to have shown that a simple commitment abiding cooperative strategy can be evolutionarily advantageous even in a non-iterated game setting. But much remains to be explored. In the more general setting and to avoid confusion, it can be helpful to distinguish, even if only conceptually, between “execution moves” and “pre-play moves” [19]. The terms first move and last move then always refer exclusively to execution moves – the choices that actually generate the payoffs. In contrast, commitments come earlier with respect to execution moves: they are pre-play moves. A preemptive commitment is a pre-play move that allows the player making it to take the first execution move. A reactive commitment, although also a pre-play move, can be made only by the player who has the last execution move. In either case, by giving up on his or her choice through committing, the commitment player leaves the opponent with “the last clear chance to decide the outcome” [43].

In our present game setting, however, there was no need to make the distinction between the first and the second to play, because each possible player strategy move is exhibited and fixed from the start, as expressed and ingrained in the payoff matrix. By so introducing the several committed unconditional move strategies – though the payoff is of course conditional on the opponent’s move–, we can emulate what would happen in a round if a move sequence actually existed. Put briefly, our commitment model is of the simplest kind and, moreover, it is brought to bear solely on the very next move fold of a pair of players, with no history available on prior commitments. Nevertheless, it captures core features of commitment, namely the high cost of defaulting to discourage false commitment, and thus make it plausible, and a comparatively small but non-zero cost of commitment proposal to lend it initial credibility. On top of this core model more elaborate models affording commitment can subsequently be rooted, including those involving delayed deceit..

What’s more, commitment (or intention manifestation) and intention recognition, are but two sides of a coin really, and their future joint study in the EGT setting is all but unavoidable. It has become increasingly obvious that maximizing reproductive success often requires keeping promises and fulfilling threats, even when that requires in turn sacrifices regarding individual short-term interests. That natural selection has shaped special mental capacities to make this possible seems likely, including a capacity for commitment [26] and for intention recognition [14, 15]. The commitment stance goes yet further, and many aspects of human groups seem shaped by effects of commitments and intention recognition, namely group boundaries, initiation rituals, ideologies, and signals of loyalty to the group [47, 48, 49]. Conversely, many aspects of groups seem to exist largely to facilitate commitment to cooperate and to limit the utility of coercive threats.

The generalized ability for commitment to support cooperative interaction is an important aspect of plasticity in human behavior, and humans support their deal-making in lots of ways. The law is full of instances of people using techniques of commitment to establish the honesty of their

intentions, namely through a variety of contracts [13]. Institutions themselves are supported on committal contracts, and the law of the land proffers methods for constituting and of accountability of social institutions [44].

We believe that studies of commitment will benefit greatly from rigorous models that allow for their analytical study and computer simulation, and in particular within the fold of EGT for the better to examine the emergence of complex social behavior.

6. ACKNOWLEDGMENT

HTA and FCS acknowledge the support from FCT-Portugal (grant SFRH/BD/62373/2009 and R&D project PTDC/FIS/101248/2008, respectively).

7. REFERENCES

- [1] T. Antal, A. Traulsen, H. Ohtsuki, C. E. Tarnita, and M. A. Nowak. Mutation-selection equilibrium in games with multiple strategies. *J. Theor. Biol.*, 258:614–622, 2009.
- [2] R. Axelrod. *The Evolution of Cooperation*. Basic Books, ISBN 0-465-02122-2, 1984.
- [3] R. Axelrod and W. Hamilton. The evolution of cooperation. *Science*, 211:1390–1396, 1981.
- [4] I. Back and A. Flache. The Adaptive Rationality of Interpersonal Commitment. *Rationality and Society*, 20(1):65–83, 2008.
- [5] C. Boehm. The natural selection of altruistic traits. *Human Nature*, 10(3):205–252, 1999.
- [6] C. Castelfranchi and R. Falcone. *Trust Theory: A Socio-Cognitive and Computational Model (Wiley Series in Agent Technology)*. Wiley, 2010.
- [7] A. K. Chopra and M. P. Singh. Multiagent commitment alignment. In *Proceedings of the 8th International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 937–944, 2009.
- [8] de Vos, R. Smaniotto, and D. Elsas. Reciprocal altruism under conditions of partner selection. *Rationality and Society*, 13(2):139–183, 2001.
- [9] A. Dreber, D. G. Rand, D. Fudenberg, and M. A. Nowak. Winners don’t punish. *Nature*, 452(7185):348–351, 2008.
- [10] E. Fehr and S. Gächter. Altruistic punishment in humans. *Nature*, 415:137–140, 2002.
- [11] D. Fudenberg and L. A. Imhof. Imitation processes with small mutations. *Journal of Economic Theory*, 131:251–262, 2005.
- [12] H. Gintis. Beyond selfishness in modeling human behavior. In R. M. Nesse, editor, *Evolution and the capacity for commitment*. New York: Russell Sage, 2001.
- [13] O. R. Goodenough. Law and the biology of commitment. In R. M. Nesse, editor, *Evolution and the capacity for commitment*, pages 262–291. New York: Russell Sage, 2001.
- [14] T. A. Han, L. M. Pereira, and F. C. Santos. Intention recognition promotes the emergence of cooperation. *Adaptive Behavior*, 19(3):264–279, 2011.
- [15] T. A. Han, L. M. Pereira, and F. C. Santos. The role of intention recognition in the evolution of cooperative

- behavior. In *Proceedings of the 22nd international joint conference on Artificial intelligence (IJCAI'2011)*, pages 1684–1689, 2011.
- [16] P. Harrenstein, F. Brandt, and F. Fischer. Commitment and extortion. In *The 6th international joint conference on Autonomous agents and MultiAgent systems*, AAMAS '07. ACM, 2007.
- [17] C. Hauert, A. Traulsen, H. Brandt, M. A. Nowak, and K. Sigmund. Via freedom to coercion: The emergence of costly punishment. *Science*, 316:1905–1907, 2007.
- [18] J. Hirshleifer. Game-theoretic interpretations of commitment. In R. M. Nesse, editor, *Evolution and the capacity for commitment*, pages 77–93. New York: Russell Sage, 2001.
- [19] J. Hirshleifer. There are many evolutionary pathways to cooperation. *Journal of Bioeconomics*, (1):73–93, 1999.
- [20] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge U. P., 1998.
- [21] L.-M. Hofmann, N. Chakraborty, and K. Sycara. The evolution of cooperation in self-interested agent societies: a critical study. In *The 10th International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '11, pages 685–692, 2011.
- [22] N. K. Humphrey. Varieties of altruism. *Social Research*, (64):199–209, 1999.
- [23] L. A. Imhof, D. Fudenberg, and M. A. Nowak. Evolutionary cycles of cooperation and defection. *Proc. Natl. Acad. Sci. USA*, 102:10797–10800, 2005.
- [24] M. Kandori, G. J. Mailath, and R. Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56, 1993.
- [25] S. Karlin and H. E. Taylor. *A First Course in Stochastic Processes*. Academic Press, New York, 1975.
- [26] R. M. Nesse. Natural selection and the capacity for subjective commitment. In R. M. Nesse, editor, *Evolution and the capacity for commitment*, pages 1–44. New York: Russell Sage, 2001.
- [27] M. A. Nowak. Five rules for the evolution of cooperation. *Science*, 314(5805):1560, 2006. DOI: 10.1126/science.1133755.
- [28] M. A. Nowak, A. Sasaki, C. Taylor, and D. Fudenberg. Emergence of cooperation and evolutionary stability in finite populations. *Nature*, 428:646–650, 2004.
- [29] M. A. Nowak and K. Sigmund. Evolution of indirect reciprocity. *Nature.*, 437(7063):1291–1298, 2005.
- [30] H. Ohtsuki, Y. Iwasa, and M. A. Nowak. Indirect reciprocity provides only a narrow margin of efficiency for costly punishment. *Nature*, 457(7601):79–82, 2009.
- [31] S. Oyama. *Evolution's Eye: A Systems View of the Biology-Culture Divide*. Durham, N.C.: Duke University Press., 2000.
- [32] J. M. Pacheco, F. C. Santos, and F. A. C. C. Chalub. Stern-judging: A simple, successful norm which promotes cooperation under indirect reciprocity. *PLoS Comput Biol*, 2:12:e178, 2006.
- [33] L. M. Pereira and T. A. Han. Evolution prospection. In *Proceedings of International Symposium on Intelligent Decision Technologies (KES-IDT'09)*, pages 51–63. Springer Studies in Computational Intelligence 199, 2009.
- [34] L. M. Pereira and T. A. Han. Evolution prospection in decision making. *Intelligent Decision Technologies*, 3(3):157–171, 2009.
- [35] D. G. Rand and M. A. Nowak. The evolution of antisocial punishment in optional public goods games. *Nature Communications*, 2:434, 2011.
- [36] L. Rendell, R. Boyd, D. Cownden, M. Enquist, K. Eriksson, M. W. Feldman, L. Fogarty, S. Ghirlanda, T. Lillicrap, and K. N. Laland. Why copy others? insights from the social learning strategies tournament. *Science*, 328(5975):208–213, 2010.
- [37] A. Robson. Efficiency in evolutionary games: Darwin, nash, and the secret handshake. *J Theo Biol*, 144:379–396, 1990.
- [38] M. Ruse. Morality and commitment. In R. M. Nesse, editor, *Evolution and the capacity for commitment*, pages 221–236. New York: Russell Sage, 2001.
- [39] F. C. Santos and J. M. Pacheco. Risk of collective failure provides an escape from the tragedy of the commons. *Proc Natl Acad Sci USA*, 108:10421–5, 2011.
- [40] F. C. Santos, J. M. Pacheco, and T. Lenaerts. Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proc. Natl. Acad. Sci. USA*, 103:3490–3494, 2006.
- [41] F. C. Santos, J. M. Pacheco, and B. Skyrms. Co-evolution of pre-play signaling and cooperation. *J Theo Biol*, 274:30–35, 2011.
- [42] F. C. Santos, M. D. Santos, and J. M. Pacheco. Social diversity promotes the emergence of cooperation in public goods games. *Nature*, 454:214–216, 2008.
- [43] T. C. Schelling. *The strategy of conflict*. London: Oxford University Press, 1990.
- [44] J. R. Searle. *Making the Social World: The Structure of Human Civilization*. Oxford University Press, 2010.
- [45] K. Sigmund. *The Calculus of Selfishness*. Princeton U. Press, 2010.
- [46] K. Sigmund, H. D. Silva, A. Traulsen, and C. Hauert. Social learning promotes institutions for governing the commons. *Nature*, 466:7308, 2010.
- [47] B. Skyrms. *Evolution of the Social Contract*. Cambridge University Press, 1996.
- [48] B. Skyrms. *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press, 2003.
- [49] B. Skyrms. *Signals: Evolution, Learning, and Information*. Oxford University Press, 2010.
- [50] A. Traulsen, C. Hauert, H. De Silva, M. A. Nowak, and K. Sigmund. Exploration dynamics in evolutionary games. *Proc. Natl. Acad. Sci. USA*, 106(3):709–712, 2009.
- [51] A. Traulsen, M. A. Nowak, and J. M. Pacheco. Stochastic dynamics of invasion and fixation. *Phys. Rev. E*, 74:11909, 2006.
- [52] R. L. Trivers. The evolution of reciprocal altruism. *The Quarterly Review of Biology*, 46:35–57, 1971.
- [53] M. Winikoff. Implementing commitment-based interactions. In *The 6th international joint conference on Autonomous agents and multiagent systems*, AAMAS '07, pages 868–875. ACM, 2007.
- [54] M. Wooldridge and N. R. Jennings. The cooperative problem-solving process. In *Journal of Logic and Computation*, pages 403–417, 1999.