The Development of Cooperation in Evolving Populations through Social Importance

Pedro Sequeira, Samuel Mascarenhas, Francisco S. Melo and Ana Paiva
INESC-ID and Instituto Superior Técnico, University of Lisbon
Av. Prof. Dr. Cavaco Silva, 2744-016 Porto Salvo, Portugal
pedro.sequeira@gaips.inesc-id.pt, samuel.mascarenhas@gaips.inesc-id.pt, fmelo@inesc-id.pt, ana.paiva@inesc-id.pt

***Abstract*—Several agent-based frameworks have been proposed to investigate the possible reasons that lead humans to act in the interest of others while giving up individual gains. In this paper we propose a novel framework for analyzing this phenomenon based on the notions of *social importance* (SI) and *local discrimination*. We analyze such mechanism in the context of a "favors game" where a recipient agent can "claim" a favor to a donor agent, which may in turn "confer" its request at the expense of a certain cost. We perform several agent-based simulations and study both the conditions under which cooperation occurs and the dynamics of the relationships formed within a population. The results of our study indicate that the SI mechanism can promote cooperation in populations where all individuals share a common social predisposition towards the favors game, and also in initially mixed-strategy populations evolving by means of mutation and natural selection. We also show that the framework predicts the emergence of a conservative strategy that makes individuals to be "cautious" when interacting with "acquaintances".**

I. Introduction

Cooperation appears in nature as an organizational mechanism capable of enhancing the reproductive power of individuals that live in groups and help each other. The subject who cooperates, or *donor*, pays a cost so that the fitness of another individual, the *recipient*, is increased [1, 2]. Likewise, a *defector* does not incur in any penalty by helping another individual, nor does he receive the higher benefit of "working together". The success behind cooperation is that the instant cost incurred by an altruistic act can be compensated by long-term benefits if others also help the donor in the future as relationships within the population develop. If every individual in a group cooperates, their average fitness will be higher than if they always defect. In mixed populations, however, defectors will perform better on average than cooperators and natural selection based on competition will decrease the proportion of cooperators until they become extinct [1–4].

In this paper we propose an agent-based framework that follows the principles of a recent model of the general dynamics of human social interaction in the context of virtual agent simulation [5]. The model is based on Kemper's *sociological theory of human motivation* [6], which postulates that we act in favor of those to whom we confer enough *status*. Conversely, the amount of collaboration we can claim from others depends on how much status we assume that we have in their minds. To avoid confusion with other possible meanings of the word status, the concept was modeled as *social importance* (SI) [5]. According to this model, individuals perform *claims* and *conferals* when they believe they are sufficiently "important" to another individual. The SI between individuals develops as a consequence of repeated interactions.

In our framework, the concept of SI proposed in [5] is translated into a scalar denoted by $\sigma(a, b) \in [0; 1]$, which asserts the extent to which individual a will voluntarily comply with the needs or desires of individual b. In other words, $\sigma(a, b)$ represents the SI that a attributes to b. Similarly, an individual can estimate its relevance in the perspective of others as a result of interactions. We denote by $\hat{\sigma}(b, a)$ the SI that individual a believes that b attributes to him. One fundamental aspect of our framework is that of *local discrimination*—each element keeps and updates an individual registry of his relationship with very other member of the population with which he interacts. Such registry is *private*, meaning that no agent in the population has access to the importance that one agent attributes to another, hence the "local" denomination.

In this paper we perform several multiagent simulations to study the conditions under which cooperation occurs, and the impact of the SI mechanism in the emergence of cooperative behavior. We perform two sets of simulations in populations of interacting agents: one in which the agent's predisposition towards claiming/conferring is fixed through time; another where we apply natural selection and mutation to determine the best evolved strategies in a mixed population. Each dyadic interaction affects the payoff received and the change in SI attributed to each agent therein. Overall, the results of our experiments indicate that the SI model is able to promote cooperation even when all the agents in the population are not predisposed to cooperate and the initial SIs are low. More specifically, unexpected benevolent actions are socially rewarded via the SI mechanism, which promotes subsequent cooperative behaviors. Under the influence of evolutionary pressures in mixed populations, the results of our experiments predict the emergence of conservative strategies where agents likely cooperate with individuals with a high SI, defect with non-related ones and have equal probability when interacting with agents with a mid-level SI.

II. The "Favors Game"

For the purposes of our study, we use a recently proposed two-stage, two-player game referred to as the "favors game", detailed in [7], which we now summarize.

A. Game Payoffs

The game involves the exchange of work favors, where an individual playing the role of *claimer* chooses whether to ask another individual, playing the role of *conferrer*, to help him perform some work that needs to be done. Working together will incur in some cost, denoted by c_t, in which case the claimer receives the benefit of the work, denoted by b_t. If the conferrer decides not to confer to the claim, the claimer performs the work alone at the expense of some cost, denoted by c_a, and receives a benefit denoted by b_a. The act of claiming involves an extra cost to the claimer, denoted by c_c. If the claimer decides not to ask the favor, the conferrer may still decide to help by performing all the work for him, the latter thereby receiving all the work benefit b_a and the conferrer all the work cost c_a.

The logic of the game and respective payoffs are described in matrix form in Tables I(a)-I(b), where we set $c_c = -0.5$, $c_a = -1$, $b_a = 1$, $c_t = -0.5$, $b_t = 2$. By setting these values we force mutual defection to be the *dominant* action—both in the case of the claimer and the conferrer, defecting (*i.e.*, choosing $\overline{CL}$ or $\overline{CO}$) is always better than cooperating (*i.e.*, choosing CL or CO) *independently* of what the other player chooses—, making it a good testbed for our study of the dynamics behind the SI mechanism.

B. Behavioral Probabilities

In [7], we define the probabilities of two individuals asking and conceding favors to one another according to the level of SI that they attribute to and perceive from each other, namely:

$$\begin{aligned} \mathbb{P}_{cl}(a,b) &= \mathcal{B}(\hat{\sigma}(b,a), \theta_{cl}(a)), && (1)\\ \mathbb{P}_{ce}(a,b) &= \mathcal{B}(\sigma(a,b), \theta_{ce}(a)), && (2)\\ \mathbb{P}_{ci}(a,b) &= \mathcal{B}(\sigma(a,b), \theta_{ci}(a)), && (3) \end{aligned}$$

where $\mathbb{P}_{cl}(a,b)$ denotes the probability of individual a asking individual b a favor, and $\mathbb{P}_{ce}(a,b)$ and $\mathbb{P}_{ci}(a,b)$ are respectively the probabilities of a conferring a claim when b asked a favor—corresponding to an *explicit conferral*, and when b did not ask a for help—an *implicit conferral*. Parameters $\theta_{cl}(a)$, $\theta_{ce}(a)$ and $\theta_{ci}(a)$ are respectively the *claim*, *explicit conferral* and *implicit conferral* thresholds, taking values in $[0,1]$ that represent the *predisposition* of individual a in asking for and conceding others favors. $\mathcal{B}(x, x_0) = (1+e^{-\beta(x-x_0)})^{-1}$ corresponds to the Boltzmann function centered around x_0. As can be seen from (1), the higher the thresholds the harder it will be for a to perform claims and conferrals to b. Also, the greater the (estimated) SI between the individuals the more likely it will be for a to ask and concede favors to b.[1]

III. Social Importance Dynamics

SI is a dynamic mechanism that changes through time as a consequence of repeated interactions—the SI that individuals believe they hold at some other agent's eyes is shown in what they ask of that agent. Similarly, one's response to some claim will translate the SI attributed to that person. To reflect that aspect of the SI model, in this paper we devise an *action assessement mechanism* that functions like a theory-of-mind (ToM) procedure—it asserts how "expected" a player's action seems to have been in the perspective of the other player. Tables I(c)-I(d) depict the *expectancy-assessment rules* associated with each possible joint action for both players in the "favors game", where:

$$\begin{aligned} \hat{\mathbb{P}}_{cl}(a,b) &= \mathcal{B}(\sigma(a,b), \theta_{cl}(a)),\\ \hat{\mathbb{P}}_{ce}(a,b) &= \mathcal{B}(\hat{\sigma}(b,a), \theta_{ce}(a)),\\ \hat{\mathbb{P}}_{ci}(a,b) &= \mathcal{B}(\hat{\sigma}(b,a), \theta_{ci}(a)), \end{aligned}$$

are the *estimated* action probabilities when an individual a plays the game with another individual b. The idea behind these rules is for an individual to *evaluate*, in the form of a "social payoff", the actions of others according to the mismatch between the action he expected another subject to play and the one he in fact played, according to his estimate of their relationship. Given the rules for the role of claimer, expressed in Table I(c), we see that if a conferrer did indeed confer to the favor, the higher the estimated probability the lower the mismatch will be, thus resulting in a lower evaluation of the conferrer's action. If, however, the conferrer did not concede to perform the favor, the associated penalty will be as high as the estimated probability. The rules for a conferrer player are defined in Table I(d). In this case, the more expected a claim is, the lower the penalty will be when a claimer asks a favor. If, however, the claimer does not ask the favor, the more expected this action the more positive the conferrer will judge it, as a sign of recognition towards the claimer.

For relationships (and its associated SI values) to change over time as a consequence of the players actions, we defined the following SI *attribution rule*:

$$\sigma(a,b) = \sigma(a,b) + \alpha(\mathcal{W}(a,b) + \mathcal{S}(a,b)), \tag{4}$$

updating the SI attributed by individual a to individual b after some round in the "favors game", where α corresponds to a learning rate. $\mathcal{W}(a,b)$ is the *work-related payoff* for player a upon performing some joint action with player b. Payoffs are taken from the game's payoff matrices in Tables I(a)-I(b) according to the players' roles and chosen actions. Similarly, $\mathcal{S}(a,b)$ corresponds to individual a's evaluation of b's action according to the *social payoffs* in Tables I(c)-I(d). The decisions made by each individual in the game will therefore affect both the payoff it receives and its relationships. Additionally, the SI update in (4) takes information from the subjects' *external* and *social* environments. We also devised an update rule for the player a's estimate $\hat{\sigma}(b,a)$,

$$\hat{\sigma}(b,a) = \hat{\sigma}(b,a) + \alpha(\mathcal{W}(b,a) + \mathcal{S}(b,a)). \tag{5}$$

In this case, player a applies ToM reasoning and updates the SI estimate as if players a and b had swapped roles and actions. As will become apparent during our study, because each individual may have unique predispositions (thresholds)

[1]Note that for conferrals the true value $\sigma(a,b)$ is used.

TABLE I: (a)-(b) Payoffs for each player within the "favors game", where CL and $\overline{CL}$ refer to the actions of claiming and not claiming, respectively, and CO and $\overline{CO}$ refer to the actions of conferring and not conferring, respectively. (c)-(d) Social payoffs evaluating the actions of the other player, where a is the claimer and b the conferrer. See the text for more details.

Game payoff, $\mathcal{W}(a,b)$, …

(a) …for the claimer.

$a \setminus b$	CO	$\overline{CO}$
CL	1	−0.5
$\overline{CL}$	1	0

(b) …for the conferrer.

$b \setminus a$	CL	$\overline{CL}$
CO	−0.5	−1
$\overline{CO}$	0	0

Social payoff, $\mathcal{S}(a,b)$, …

(c) …attributed by the claimer.

$a \setminus b$	CO	$\overline{CO}$
CL	$1-\hat{\mathbb{P}}_{ce}(a,b)$	$-\hat{\mathbb{P}}_{ce}(a,b)$
$\overline{CL}$	$1-\hat{\mathbb{P}}_{ci}(a,b)$	$-\hat{\mathbb{P}}_{ci}(a,b)$

(d) …attributed by the conferrer.

$b \setminus a$	CL	$\overline{CL}$
CO / $\overline{CO}$	$\hat{\mathbb{P}}_{cl}(b,a)-1$	$\hat{\mathbb{P}}_{cl}(b,a)$

for each possible action, SI estimation errors may occur, and therefore the SI value one estimates, $\hat{\sigma}(b,a)$, may be quite different from its real value, $\sigma(b,a)$.

IV. Experiments and Results

In this section we describe our empirical study of the SI mechanism detailed in the previous section. In our study, we model the "favors game" as a two-stage game: claims occur in the first stage of the game, conferrals in the second.[2] For each experiment, we detail the procedure and parameters used and discuss the obtained results. In all experiments, we model a population of K agents interacting and playing the above-detailed "favors game" for T time-steps. Each agent is modeled as an independent individual capable of playing the "favors game" in both the role of claimer and conferrer, the choice of which is made randomly at each time-step. The agents keep and update a register of the SI and estimated SI for all other agents in the population according to the result of each interaction. For the purposes of our study we set an initial value of $\sigma(a,b) = \hat{\sigma}(b,a) = 0.2$ for each pair of individuals (a,b) in the population.[3] From the description of our framework it should be apparent that the general game strategy of an individual playing the "favors game" is determined by the values of the three thresholds. Formally, we denote the strategy of a player as the vector $\Theta = [\theta_{cl}, \theta_{ce}, \theta_{ci}]$. The initial strategies for each agent and the choice of interacting agents at each time-step is experiment-dependent. In all experiments we simulate a population of $K = 100$ interacting agents and use a Boltzmann function smoothness of $\gamma = 20$ and a learning rate $\alpha = 0.05$.

We determine the *fitness* of some population P according to the average cumulative payoff of all agents, *i.e.*, :

$$\mathcal{F}(\mathrm{P}) = \frac{1}{K}\sum_{k=1}^{K} f(k,T) = \frac{1}{K}\sum_{i=1}^{K}\sum_{t}\mathcal{W}_t(k,\cdot), \qquad (6)$$

where $f(k,D)$ is agent k's *individual fitness* expressed as the game-related payoff accumulated as a result of his interactions with all other agents during the last D time-steps. For the purposes of this experiment, it is also important to assess the "quality" of the relationships that each agent has with others. In that respect, we say that two agents a and b are *strongly-related* if all following relations are cumulatively verified,

$$\begin{array}{ll} \hat{\sigma}(b,a) \geq \theta_{cl}(a) & \sigma(a,b) \geq \theta_{ci}(a) \\ \hat{\sigma}(a,b) \geq \theta_{cl}(b) & \sigma(b,a) \geq \theta_{ci}(b), \end{array} \qquad (7)$$

implying that the probability of both asking a favor to the other and conceding one without the other having to ask is very high. In addition to measuring the population fitness, we also measure the number of strong-relationships each agent has, and the number of *strongly-related groups* in which it is inserted, corresponding to sub-groups of the population where all agents are strongly-related.[4]

A. Social Importance Promotes Cooperation

In this first experiment we test the ability of the SI mechanism in promoting cooperation under different initial conditions. To achieve that, we performed a sensitivity analysis by varying the action thresholds of all the agents in the population. Each player's predisposition is a slight variation of a common profile, *i.e.*, some vector $\Theta = [\theta_{cl}, \theta_{ce}, \theta_{ci}]$, indicating the general tendency of claiming and conferring within that population. In a certain manner, a profile Θ relates to the way a given culture affects the behavior of a group of individuals interacting with each other.

1) Experimental Setup: The algorithm for this experiment proceeds as follows. At each time-step $t = 1, \ldots, T$ we randomly select $K/2$ pairs (a,b) of agents to play a round of the "favors game", where agent a plays the role of claimer and b the role of conferrer. In each interaction, the claimer chooses its action (CL or $\overline{CL}$) according to the action probability $\mathbb{P}_{cl}(a,b)$ in (1). The conferrer then chooses his action (CO or $\overline{CO}$) according to either $\mathbb{P}_{ce}(b,a)$ or $\mathbb{P}_{ci}(b,a)$, depending on player a's action. After playing the game, each player receives the corresponding payoff as defined in the game matrices in Tables I(a)-I(b), and also updates the SI and estimated SI based on the assessment matrices in Tables I(c)-I(d) and the update rules in (4) and (5). This process repeats for T time-steps.

Each agent's strategy is initially sampled from the interval defined by $\Theta \pm [0.05, 0.05, 0.05]$ to introduce some disturbance in the population's strategies and is held fixed throughout time. We tested 8 different predisposition profiles resulting from combining values for each threshold $\theta_i \in \Theta$ from the

[2] This means that a conferrer knows if the claimer asked a favor but a claimer takes its decisions based solely on his estimated SIs $\hat{\sigma}(a,b)$.

[3] This value represents a baseline for describing initial relationships according to the SI model [5], *i.e.*, we can think of a relationship of this kind as one in which the individuals are "strangers".

[4] We can think of strongly-related individuals as "friends", and strongly-related groups as "circles of friends". Agents can have many friends and belong to many different groups.

TABLE II: Results for the first experiment where different populations of agents play some strategy Θ in the "favors game". Population labels (A-J) were added for convenience of explanation.

	Profile Θ θ_{cl}	θ_{ce}	θ_{ci}	**Mean Fitness** $\mathcal{F}(\Theta)$	**Mean SI** σ	**Mutual** Coop.	Def.
A	1.0	1.0	1.0	0.0 ± 0	0.20 ± 0.2	0.00	1.00
B	0.8	0.8	0.8	-0.1 ± 1	0.20 ± 0.2	0.00	1.00
C	0.8	0.8	0.2	$4,216.6 \pm 312$	0.51 ± 0.7	0.40	0.40
D	0.8	0.2	0.8	14.1 ± 4	0.20 ± 0.2	0.00	1.00
E	0.8	0.2	0.2	$10,612.5 \pm 52$	0.94 ± 0.9	0.92	0.02
F	0.2	0.8	0.8	-271.0 ± 68	0.33 ± 0.2	0.19	0.60
G	0.2	0.8	0.2	$-1,045.2 \pm 87$	0.35 ± 0.3	0.19	0.44
H	0.2	0.2	0.8	$4,168.0 \pm 137$	0.48 ± 0.5	0.46	0.49
I	0.2	0.2	0.2	$9,493.6 \pm 164$	0.87 ± 0.9	0.91	0.08
J	0.0	0.0	0.0	$12,438.2 \pm 3$	0.96 ± 0.9	1.00	0.00

set $\{0.2, 0.8\}$. For a baseline comparison, we also tested the threshold vector $\Theta = [0, 0, 0]$, corresponding to a population of "sucker" agents that almost always claim and confer, and vector $\Theta = [1, 1, 1]$, corresponding to a population of "loner" agents that have the lowest probability of claiming and conferring. For each population, we ran 100 simulations of $T = 5 \times 10^4$ time-steps each.

2) Results: Table II depicts the results of this experiment, where we present the mean population fitness, mean final SI and the percentage of the population playing mutual cooperation, *i.e.*, (CL,CO) and mutual defection, *i.e.*, ($\overline{CL}$,$\overline{CO}$). Fig. 1(a) depicts the percentage of the population playing the four possible joint actions at around 0.03% of the total time of the simulation, *i.e.*, between $t = 1,500$ and $t = 2,000$.[5] Figs. 1(b)-1(c) show the progress of the mean number of strong relationships and number of groups that agent participates, respectively. We omitted the results of populations A and B as they were redundant (near 0) throughout time for both metrics.

The first thing we can observe is that high SI levels in the population, *e.g.*, populations I and J, are associated with high population fitness and high levels of mutual cooperation. The inverse is also true, low levels of SI, *e.g.*, populations A and B, lead to mutual defection and therefore to generalized low levels of fitness. By looking at populations with different thresholds in their strategies, we see that the way the SI mechanism promotes cooperation is related to the appearance of unexpected charitable acts by the conferrers—having a high claim threshold of $\theta_{cl} = 0.8$ means that it is difficult for agents to work together and thus mutually cooperate; however, having a low implicit conferral threshold of $\theta_{ci} = 0.2$, as is the case of populations C and E, allows agents to start gaining SI by perform implicit conferrals even to "strangers", *i.e.*, peers with a low-value of SI, thus strengthening their relationships. We can see this effect in Fig. 1(a) in which a high percentage of players in these populations are playing the joint action ($\overline{CL}$,CO). Despite being an inefficient strategy because players are "wasting" their resources by individually performing the work for others, as time progresses this loss is compensated—the gain in SI is reflected in more agents claiming for favors that are conceded. By the end of the simulations, these populations show high levels of average fitness, as can be seen in Table II. Without the opportunity to confer implicitly, the SI mechanism cannot promote altruistic actions and inherently non-cooperative populations, such as B and E, will inevitably end up in generalized mutual defection.

The expectancy mechanism behind the SI framework can also be detrimental for a population's fitness—if the claim threshold is low when compared to the explicit conferral threshold, *i.e.*, when $\theta_{cl} = 0.2$ and $\theta_{ce} = 0.8$, claims start not being conceded and claimers end up doing the work all by themselves and also spending the cost associated with asking the favor. We can observe this phenomenon by looking at the percentage of agents playing this joint action (CL,$\overline{CO}$) in populations F and G in Fig. 1(a). By having this false "hope" of being corresponded, claimers play a strategy that is actually worse than mutual defection. This fact together with the lack of opportunity to perform charitable actions makes the SI level to not improve and lead to the populations' to have a negative accumulated fitness at the end, as depicted by the results in Table II. In this particular dispute between populations F and G, having a high implicit conferral threshold of $\theta_{ci} = 0.8$ makes population F to shift to the more payoff efficient strategy of mutual defection which explains the difference in fitness between the populations.

Relating the structure of the population, we observe that, in general, the fitness of the population is proportional to the average number of strong relationships, as can be seen by comparing the results in Table II with Fig. 1(b). The exceptions are populations F and G, in which the average number of strong relationships per agent is approximately 20 by the end of the simulation. By looking at the final distribution of joint actions for these populations in Table II, last columns, we see that this result is due to a great dispersion in strategies being played in the population—there is a great portion of the population, especially when compared to the other populations, that is neither mutually cooperating nor mutually defecting. This fact, together with the high average number of groups in which each agent is inserted, as depicted in Fig. 1(c), informs us that the structure of these populations is not very cohesive—it is easy for claimers to ask for favors, but they will only be conceded to strongly-related individuals. The analysis of the number of groups and distribution of mutual strategies for the other populations seem to reinforce the idea that more cohesive populations attain greater degrees of fitness.

B. Conservative Strategies in Mixed Populations

Having established that the SI mechanism promotes cooperation under some conditions in inherently competitive settings, as is the case of the "favors-game", in this second experiment we are interested in determining which kind of strategy emerges from a mixed-population playing the game. Therefore, in this experiment each agent k plays according to its own profile, denoted by $\Theta(k)$, initially selected at random.

[5] An informal analysis on the evolution of joint strategies showed that during this initial period of the simulation the SI mechanism starts to make effect on the dynamics of the played strategies. Usually, after this point strategies start to converge to their final values.

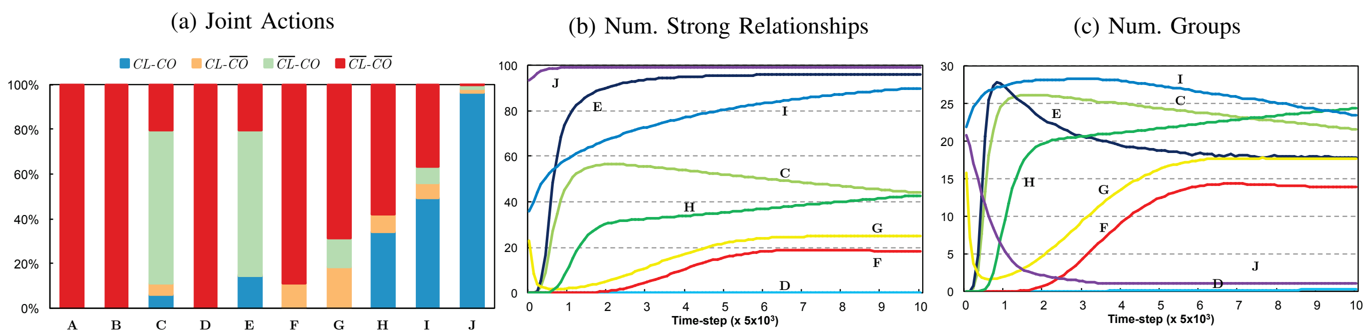

Fig. 1: Comparative results for the first experiment: (a) Mean distribution of the joint actions played by pairs of agents in the "favors game" for all tested populations, taken between $t = 1,500$ and $t = 2,000$; (b) Mean number of strong relationships per agent; (c) Mean number of strongly-related groups per agent. Results are averages over 100 independent Monte-Carlo trials.

The objective is to analyze, based on the thresholds of the emerged strategy, the kind of behavior exhibited and also the relationships the individuals create throughout time.

1) Experimental Setup: We devised an evolutionary procedure based on the mechanisms of natural selection and mutation. We modified the algorithm used in the previous experiment to support evolution as follows. At each time-step $t = 1 \dots T$ the agents play the "favors game" as before. At intervals of G time-steps, a new *generation* of agents replaces the current population. The selection mechanism randomly selects for each agent a in the population another player b which threshold vector he imitates with a probability of:

$$\mathbb{P}_{imit}(a, b) = \mathcal{B}(f(b, G), f(a, G)).$$

This means that the higher the difference in the payoff accumulated during the last generation by the agent being imitated, *i.e.*, b, in relation to that of agent a, the higher the imitation probability $\mathbb{P}_{imit}(a, b)$ will be. If indeed this occurs, agent a's threshold vector "approximates" that of b by:

$$\Theta(a) = \Theta(a) + \mathbb{P}_{imit}(a, b)(\Theta(b) - \Theta(a)),$$

where $\mathbb{P}_{imit}(a, b)$ is here used as an approximation factor. After imitation, a certain percentage $\mathcal{M}$ of the population's weakest players, as determined by their individual fitness $f(k, G)$, are selected for mutation according to:

$$\mathcal{M}(t) = \gamma(1 - 10T^{-1})^t,$$

where γ is a mutation factor. The profile Θ of mutant agents is randomly generated. Strategies evolve and mutate in a process interleaved with that of playing the game, until T is reached. In this experiment we used a mutation factor of $\gamma = 0.9$ and performed selection at intervals of $G = 1,000$ time-steps. We ran 100 simulations for an evolving population as described above for $T = 1.5 \times 10^5$ time-steps each.

2) Results: The results of the second experiment are presented in Figs. 2(a)-2(b) and in Table III. Fig. 2(a) depicts the evolution of the mean thresholds in the population. To get a sense of the distribution of strategies within the population we also present the standard deviations. As we can observe, a threshold vector close to $\Theta = [0.56, 0.48, 0.53]$ ends up being adopted by the generality of the population

TABLE III: Results of the experiments using the evolutionary procedure in the "favors game" according to several metrics.

Metric	**Population Mean**		
Profile (Θ)	$[0.56, 0.48, 0.53]$		
Fitness ($\mathcal{F}(\Theta)$)	$15,087.67$	$\pm$	$1,057.6$
SI (σ)	0.47	$\pm$	0.0
Mutual Cooperation	0.46	$\pm$	0.0
Mutual Defection	0.54	$\pm$	0.0
Strong Relationships	45.42	$\pm$	3.3
Num. Groups	28.84	$\pm$	2.3

by means of evolution based on natural selection. We can also see the impact of mutation in the early stage of the evolutionary process. The departing mixed population evolves into a "standardized" one—as time progresses, the deviation in the population decreases meaning that the thresholds become widely adopted as a result of imitation. The emerged vector makes the agents engage in a conservative strategy when playing the "favors game". By further analyzing the results in Table III we see that such behavior is justified by the mid-level around, $\theta_i = 0.5$, of all thresholds together with the average SI within the population, $\sigma = 0.47$. What these results mean is that by the end of the simulation, the agents claim and confer to individuals having a high SI level and defect with those having a low SI level. Individuals that are at the level of "acquaintances", *i.e.*, with an SI of 0.5, have about 50% chance of being asked and conceded favors, as dictated by the action probabilities defined in (1)-(3). Such cautious strategy enabled the agents to start creating strong relationships in the beginning, converging to a number of 45.42 "friends" at the end, *i.e.*, creating strong relationships with approximately half of the members of the population. This analysis is reinforced by looking at the evolution of average strategies in Fig. 2(b). We clearly see that by the end of the simulation, agents only mutually cooperate with peers within their strongly-related groups, *i.e.*, the probability of being matched with a "friend" is about 0.45, which is related to the percentage of mutual cooperations being performed, 0.46 (see Table III).

V. Conclusions

In this paper we proposed a novel framework inspired by a theory that simulates the general dynamics of human

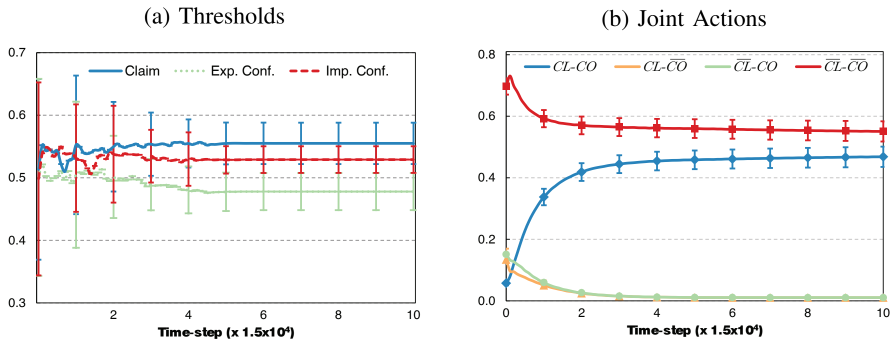

Fig. 2: Progress of results of the experiments using the evolutionary procedure: (a) Evolution of the mean threshold profile Θ, where "Claim" corresponds to θ_{cl}, "Exp. Conf." corresponds to θ_{ce} and "Imp. Conf." corresponds to θ_{ci}; (b): Evolution of the mean percentage of joint actions being played in the "favors game" within the population.

social interaction through the concept of social importance. The goal was to provide a more rich and flexible social framework that relies on a continuous and differentiating evaluation of an individual's interactions with other members of the population. We analyzed our framework in the context of an inherently competitive "favors game" where agents "claim" and "confer" favors. The results of our empirical study show that this mechanism is able to promote cooperation by means of unexpected altruistic actions. By applying an evolutionary procedure in mixed-populations, we observed the emergence of a conservative strategy that provides agents a flexible mechanism to interact with others.

Several mechanisms have been proposed that help explain the evolution and stabilization of cooperation in competitive settings (see [2] for a recent formalization of several such mechanisms). Many mechanisms involve the notion of *relatedness* to balance the decision of cooperating or not with others [4, 8]. Others assume the ability to retaliate after uncooperative acts by means of *reciprocation* and *punishment* [1, 3]. Mechanisms based on *indirect reciprocity* [9–11] assume the existence of a public record determining the *reputation* of individuals according to evaluations made by an observer. In general, very simple rules, strategies and cognitive capabilities are required to explain the behavior of individuals when cooperation is studied under standard frameworks of evolutionary game dynamics [2].

Our framework departs from mechanisms that promote cooperation based on indirect reciprocity as it relies on local interactions to influence interpersonal relationships. Our results point towards the idea that if individuals do not interact as often with all the members in the population, it may be more efficient to have preferences with which they cooperate and to discriminate those outside of their "friend circle". We believe the proposed SI mechanism to be a new form of direct reciprocity in which agents perform charitable acts in the hope of future mutual cooperation. The mechanism is still flexible enough so that agents stop conceding claimed favors after perceiving "unkind" acts from others. By treating interactions as private, our agents develop different attitudes towards individuals inside and outside of their groups. In addition, our matching mechanism based on the SI levels resembles that of network reciprocity [12], in which the matching of interacting players is based on a spatial structure of the population. The SI framework follows the idea that punishment is not a necessary vehicle to achieve cooperation [2]—by lowering SI levels as a consequence of "unworthy" acts, subjects decrease the confidence on such peers which will in turn influence future interactions between them.

In the future we want to use our SI framework to endow our agents with different cultural aspects influencing their decisions and analyze how distinct cultures may arise and cohabit within the same population.

Acknowledgment

This work was supported by national funds through Fundação para a Ciência e a Tecnologia (FCT) with reference UID/CEC/50021/2013. The first author acknowledges a BDP grant CMUP-ERI/HCI/0051/2013 from the INSIDE project funded by the FCT. The second author acknowledges a PhD grant SFRH/BD/62174/2009 from the FCT.

References

[1] R. M. Axelrod, *The evolution of cooperation*. New York, USA: Basic Books, Inc., Publishers, 1984.
[2] M. A. Nowak, "Five rules for the evolution of cooperation." *Science*, vol. 314, no. 5805, pp. 1560–1563, 2006.
[3] R. Trivers, "The evolution of reciprocal altruism," *The Quart. Review of Biology*, vol. 46, no. 1, pp. 35–57, 1971.
[4] R. Dawkins, *The Selfish Gene*, 30th ed. New York, USA: Oxford University Press, 2006.
[5] S. Mascarenhas, R. Prada, A. Paiva, and G. J. Hofstede, "Social importance dynamics: A model for culturally-adaptive agents," in *Intelligent Virtual Agents*. Springer, 2013, no. 8108, pp. 325–338.
[6] T. D. Kemper, *Status, Power and Ritual Interaction*. Ashgate, 2011.
[7] P. Sequeira, S. Mascarenhas, F. S. Melo, and A. Paiva, "The "favors game": A framework to study the emergence of cooperation through social importance (extended abstract)," in *Proc. 14th Int. Conf. on Autonomous Agents and Multiagent Syst. (AAMAS)*, 2015.
[8] W. D. Hamilton, "The genetical evolution of social behaviour." *J. Theoret. Biol.*, vol. 7, no. 1, pp. 1–16, 1964.
[9] M. A. Nowak and K. Sigmund, "The dynamics of indirect reciprocity," *J. of Theoret. Biol.*, pp. 561–574, 1998.
[10] I. Seinen and A. Schram, "Social status and group norms: Indirect reciprocity in a repeated helping experiment," *Europ. Econ. Review*, vol. 50, no. 3, pp. 581–602, 2006.
[11] 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 comp. biology*, vol. 2, no. 12, pp. 1634–1638, 2006.
[12] M. A. Nowak and R. M. May, "Evolutionary games and spatial chaos," *Nature*, vol. 359, no. 6398, pp. 826–829, 1992.