Indirect reciprocity in finite populations of explorative agents

Fernando P. Santos
INESC-ID and Instituto Superior Técnico, Universidade de Lisboa
IST-Taguspark, 2780-990 Porto Salvo, Portugal
fernando.pedro@tecnico.ulisboa.pt

Jorge M. Pacheco
CBMA and Departamento de Matemática e Aplicações, Universidade do Minho
Campus de Gualtar, 4710-057 Braga, Portugal
jmpacheco@math.uminho.pt

Francisco C. Santos
INESC-ID and Instituto Superior Técnico, Universidade de Lisboa
IST-Taguspark, 2780-990 Porto Salvo, Portugal
franciscosantos@tecnico.ulisboa.pt

ABSTRACT

Indirect reciprocity (IR) constitutes a powerful mechanism to sustain cooperation between self-regarding agents. With IR, reputations are attributed depending on actions and particular contexts, which implies that altruism today can translate into a social image, prone to be rewarded tomorrow. It is not trivial, however, which social norms (i.e., rules that dictate the reputation to be attributed given actions and contexts) will drive the learning process of agents towards the most cooperative – and socially desirable – states. Here we provide the summary of two recently published works where we employ computational and mathematical methods (inspired in evolutionary game theory) to explore which social norms are able to promote cooperation. We show that the size of a population plays a key role in the cooperation rates sustained by particular social norms. In particular, we show that a single norm systematically leads to the highest cooperative standards in small communities. Additionally, we show that high exploration rates (i.e., the probability that an agent randomly explores the strategy space) may either improve or harm cooperation, depending on the underlying social norm at work.

1. INTRODUCTION

Indirect Reciprocity (IR) is a fundamental mechanism to understand the evolution of cooperation, often said to provide the biological basis of our morality [6]. Under IR one expects a return, not from someone we helped, but from a third party; in this sense, helping the "right" individuals may contribute to a reputation uplift that increases the chance of being helped at a later stage. This reputation shift depends on the socially adopted norms that define what actions (and under which contexts) are reckoned as good (G) or bad (B). In turn, the way reputations are attributed significantly impacts the success metrics associated with each strategy and fundamentally drives the associated learning process of agents; the overall cooperation levels in the system vary concomitantly.

Studying IR allows us to unveil mathematically the fundamental hinges that sustain cooperation in social systems, also providing clues for puzzling aspects of human collective action. On top of that, IR models emerge nowadays as relevant toolkits for artificial intelligence applications: first, there is a growing interest in including morality in artificial decision making [1], an endeavour that can naturally profit from the formal apparatus of IR; second, given the current importance of reputation-based systems in online services and systems supporting sharing economies, IR models can be used as important tools to design reputation schemes that efficiently induce users to cooperate [2, 4, 3].

Most theoretical models employed to date have studied how IR can lead to the emergence and sustainability of cooperation in infinite populations of individuals that always resort to fitness-driven mechanisms of strategy update [7, 6]. However, social interactions often take place in confined groups composed by few dozens of people. In small populations, stochastic finite size effects are not only important, but may even render infinite populations analyses misleading [5]. The stochastic nature of IR interactions is even amplified by the fact that individuals frequently update their behaviour following motives that depart from common fitness-based heuristics. Thus, it remains an open question which norms are particularly efficient promoting cooperation in small sized communities, composed by agents that combine social learning capabilities (i.e., resort to observation and imitation to update strategies) with the spontaneous exploration of the strategy space. With the current extended abstract, we would like to summarise two works recently published in which we provide a new analysis of this problem [9, 8].

We consider a population of individuals who randomly interact in pairs through a donation game: a donor may cooperate (C) and help the recipient at a cost c to herself, conferring a benefit b to the recipient (with $b > c$); otherwise (defection, D) no one pays any costs nor distributes any benefits. Reputations are public and attributed by a bystander who witnesses a pairwise interaction. We adopt a world of binary reputations, Good (G) and Bad (B). This binary reputation scheme, despite its formal simplicity, allows to consider a plethora of moral rules with variable complexity and it is specially amenable to a systematic mathematical treatment, in the framework of population dynamics. To perform an evaluation, the bystander uses a social norm, that is, a rule that converts the combined information stemming from the action of the donor and the reputation of the recipient into a new reputation for the donor. Social norms encoding this type of information are classified as 2nd-order norms [7]. Four of these social norms have been given special attention: Stern-judging (**SJ**, also known as Kandori), which assigns a good reputation to a donor that helps a good recipient or refuses help to a bad one, assigning a bad reputation in the other cases; Simple-Standing (**SS**), similar

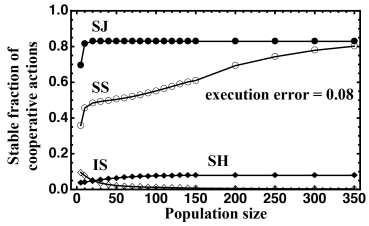

Figure 1: Stern-judging (SJ) is able to foster the highest rates of cooperation, independently of the population size; the efficiency of SS in fostering cooperation strongly depends on the population size and errors see [9].

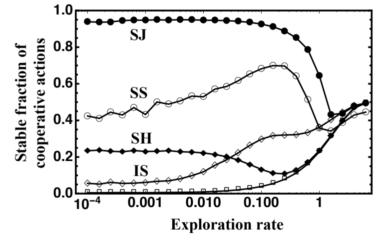

Figure 2: Cooperation under Simple-standing (SS) and Image-score (IS) profits from high exploration rates

to SJ, but more "benevolent" by assigning a good reputation to any donor that cooperates; Shunning (**SH**), similar to SJ but less "benevolent", by assigning a bad reputation to any donor that defects; and Image Score (**IS**, actually a 1st order norm) where all that matters is the action of the donor, who acquires a good reputation if playing C and a bad reputation if playing D [6].

In the space of 2nd-order norms that we consider, a duple p suffices to unambiguously define a strategy, by specifying the action directed at a G or B recipient. This leads to the following 4 possible strategies: unconditional Defection (AllD, $p = (D, D)$), unconditional Cooperation (AllC, $p = (C, C)$), Discriminator strategy (Disc, $p = (C, D)$), that is, cooperate with those in good reputation, and defect otherwise), and paradoxical Discriminator strategy (pDisc, $p = (D, C)$, the opposite of Disc). Unlike previous studies, in [9, 8] we investigate the evolutionary dynamics of these 4 strategies within finite populations by means of a stochastic birth-death process, both analytically and through large-scale computer simulations. We assume that agents resort to their strategy to interact with different opponents over time, accumulating payoff and having their own reputation updated after each encounter. In an asynchronous fashion, agents are giving the opportunity to adapt and acquire a different strategy: with a probability μ (also called, exploration or mutation rate), an agent A will randomly adopt any of the 4 strategies; with probability $1 - \mu$, agent A will pick a random opponent B and copy her strategy with a probability that grows with their fitness (i.e, average payoff) difference. This way, different social norms will judge strategies differently, which will impact the adoption of behaviours over time and lead to the distinctive cooperation levels depicted in Fig. 1 and Fig. 2.

2. SUMMARY OF RESULTS

As detailed in Fig. 1, we show that SJ clearly stands out for small population sizes, dominating with SS for large population sizes. Indeed, it can be shown that only SJ and SS are able to combine a high prevalence of an ALL-Disc configuration with the incidence of Good reputations in this configuration, efficiently fostering high levels of cooperation (Fig. 1). Yet, in small-scale societies, SJ significantly promotes more cooperation than SS, as the latter fails to prevent the invasion of unconditional defectors (AllD) in small populations. On the other hand, SJ fosters an ideal coordination between strategy and prevailing reputations, assuring the stability of configurations where individuals cooperate in the donation game. These results remain valid for a wide interval of reputation assignment time-scales, errors of execution and reputation-assignment inaccuracies.

In [8], and as depicted in Fig. 2, we show that considering highly explorative agents nurtures an environment particularly favourable for increasing cooperation levels under IS and SS. Notwithstanding, SJ remains the leading social norm in promoting cooperation.

3. ACKNOWLEDGEMENTS

This research was supported by Fundação para a Ciência e Tecnologia (FCT) through grants SFRH/BD/94736/2013, PTDC/EEI-SII/5081/2014, PTDC/MAT/STA/3358/2014 and by multi-annual funding of CBMA and INESC-ID (under the projects UID/BIA/04050/2013 and UID/CEC/50021/2013 provided by FCT).

4. REFERENCES

[1] V. Conitzer, W. Sinnott-Armstrong, J. S. Borg, Y. Deng, and M. Kramer. Moral decision making frameworks for artificial intelligence. *AAAI-17*, 2017.

[2] C. Dellarocas. Reputation mechanism design in online trading environments with pure moral hazard. *Inf. Syst. Res.*, 16(2):209–230, 2005.

[3] H. Franks and N. Griffiths. Robust reputation in decentralized markets. *Comput. Intell.*, 31(4):569–592, 2015.

[4] C.-J. Ho, Y. Zhang, J. Vaughan, and M. Van Der Schaar. Towards social norm design for crowdsourcing markets. In *AAAI Workshops*, 2012.

[5] L. A. Imhof, D. Fudenberg, and M. A. Nowak. Evolutionary cycles of cooperation and defection. *Proc. Natl. Acad. Sci. USA*, 102(31):10797–10800, 2005.

[6] M. A. Nowak and K. Sigmund. Evolution of indirect reciprocity. *Nature*, 437(7063):1291–1298, 2005.

[7] H. Ohtsuki and Y. Iwasa. The leading eight: social norms that can maintain cooperation by indirect reciprocity. *J. Theor. Biol.*, 239(4):435–444, 2006.

[8] F. P. Santos, J. M. Pacheco, and F. C. Santos. Evolution of cooperation under indirect reciprocity and arbitrary exploration rates. *Sci. Rep.*, 6(37517), 2016.

[9] F. P. Santos, F. C. Santos, and J. M. Pacheco. Social norms of cooperation in small-scale societies. *PLoS Comput. Biol.*, 12(1):e1004709, 2016.