Social Norms of Cooperation in Multiagent Systems

(Doctoral Consortium)

Fernando P. Santos

GAIPS, INESC-ID and Instituto Superior Técnico, Universidade de Lisboa
ATP-Group, Lisboa, Portugal
2780-990 Porto Salvo, Portugal
fernando.pedro@tecnico.ulisboa.pt

ABSTRACT

Explaining cooperation between individuals is a fundamental scientific endeavour. Similarly, devising mechanisms that encourage autonomous agents to cooperate with each other is a central topic in multiagent systems. Indirect reciprocity (IR) emerged as the most elaborated mechanism of cooperation discovered, involving the moral assessment of actions, the spreading of reputations and agents that care about the social image of others. In this context, the way actions are judged and reputations are attributed depends on the socially adopted norms that define what actions (and under which contexts) are reckoned as Good or Bad. It remains an open question which norms are able to maintain high levels of cooperation in a multiagent system, composed by a finite number of agents that may often behave unpredictably. We employ computational and mathematical methods (inspired in evolutionary game theory) to explore which social norms are able to promote cooperation in a multiagent system. We seek to extend the state of art by introducing new models of IR that consider a finite number of agents that are given the opportunity to randomly explore the strategy space and that may commit errors of different nature.

Keywords

Cooperation; Indirect Reciprocity; Reputations; Social Norms; Evolutionary Game Theory; Multiagent System.

1. INTRODUCTION

Mathematically, we can say that an agent – called donor – cooperates (C) when deciding to expend a cost (c) to provide a benefit (b, $b > c$) to a recipient [7]. When defecting (D), no one receives or looses payoff. As cooperation implies the altruistic loss of payoff for the donor, it is not trivial to reason why would a rational and self-regarded agent cooperate with anyone else. When played bilaterally and simultaneously, this cooperation game turns into a Prisoner's Dilemma (PD), which provides a convenient abstraction of a wide range of interactions where the pursuit of self-interest leads to poor collective outcomes. Understanding this dilemma, and each mechanism capable to solve it, is fundamental to understand human evolution [7]. On top of that, under-

standing cooperation is today essential to manage large-scale systems – such as peer-to-peer networks, e-commerce, trading systems, crowdsourcing platforms and sharing economies – where cooperation between agents is paramount [2, 3].

In this context, Indirect Reciprocity (IR) was emphasised as the most elaborated and cognitively demanding mechanism of cooperation discovered so far [5]. Whereas under direct reciprocity one expects to receive help from someone helped before, under IR one expects a return, not from someone we helped, but from someone else; in this sense, helping the "right" individuals may contribute to a reputation uplift that increases the chance of being helped, by someone else, at a later stage. The relation between cooperation and IR has already been addressed within the multiagent systems community. In particular, Peleteiro et al. nicely show that cooperation under IR is boosted by coalitions and agents able to rewire their neighborhoods. In that work, the reputation of agents is uplifted anytime they cooperate [6]. Often, however, the update of reputations also depends on the agents against whom actions are directed to. The notion of social norm is central at this point, as the reputation shifts depend on the adopted social norms that define what actions (and under which contexts) are reckoned as Good/Bad.

Social norms are a wide topic, extensively discussed in multiagent systems [12, 11, 10, 4]. On general, social norms are public and establish an expected pattern of behaviour [12]. When violated, they may lead to "*responses that range from gossip to open censure, ostracism, or dishonour for the transgressor*" [1]. In our work, we follow this notion: we consider that a social norm stands as a rule that dictates the expected behaviour of agents that act as donors, attributing them a new reputation (Good, G or Bad, B), given their action (Cooperate, C or Defect, D) and also the reputation of the recipient (G or B).

Figure 1: Social norms in terms of the new reputation (Bad/Good, inside each matrix) attributed to a donor given its action (C/D, rows) and the reputation (B/G, columns) of the recipient.

Social norms encoding this type of information are classified as 2nd-order norms [5]. Four of these social norms

have been given special attention (Fig.1): 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 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 [5]. In the space of 2nd-order norms that we consider, a duple p suffices to unambiguously define a strategy of an agent, by specifying the action directed at a G or B recipient. This leads to the following 4 strategies to be used by agents: 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, we seek to investigate the evolutionary dynamics of these 4 strategies within a multiagent system, by means of a stochastic birth-death process, both analytically [9] and through large-scale computer simulations [8]. We fix a social norm and let the system evolve by allowing agents to adopt strategies through social learning (i.e., the observed strategies leading to higher payoffs are adopted with higher probability) or through random exploration of the strategy space. The following questions are posed: **1) What social norms lead to higher levels of cooperation? 2) What social norms are robust to the inclusion of explorative agents, i.e., agents that behave non-strategically by randomly exploring the strategy space? 3) What social norms are robust if agents are allowed to error, both in attributing/retrieving wrong reputations or failing to act as their strategy conveys?**

2. RESULTS
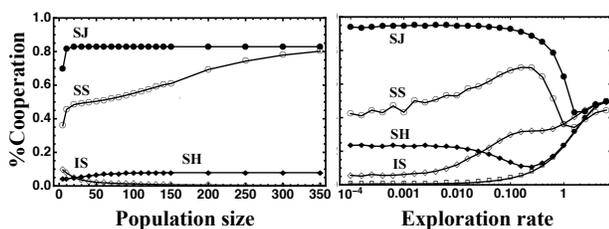
Figure 2: Left: The effect of population size on the emergent cooperation levels for each social norm; a single norm known as Stern-Judging (SJ) stands out for small population sizes. Right: Cooperation under Simple-standing (SS) and Image-score (IS) profits from high exploration rates.

Through analytical tools [9] and large-scale computational simulations [8], we show that population size strongly influences the merits of each social norm (Fig.2). While more than one social norm is able to promote high levels of cooperation in large populations (SJ and SS), in small populations only a single social norm proves to be efficient (SJ), a

property that can be comprehended by assessing the corresponding collective dynamics of strategy adoption [9].

Additionally, we show that high exploration rates (i.e., allowing agents to explore the strategy space with high probability) particularly leverage cooperation in systems where the governing social norm is either SS or IS (Fig.2). Again, we provide a simulation toolkit that allows comprehending these results, by keeping a detailed track of the dynamics resulting from the process of strategy adoption by agents [8]. Future work shall emphasise the characteristics of agents that allow an effective adherence to social norms as well as how different social norms may compete with each other.

3. ACKNOWLEDGEMENTS

REFERENCES

[1] C. Bicchieri. *The grammar of society: The nature and dynamics of social norms*. Cambridge University Press, 2005.

[2] N. Griffiths and M. Luck. Changing neighbours: improving tag-based cooperation. In *Proceedings of the 9th AAMAS*, pages 249–256. IFAAMAS, 2010.

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

[4] J. Morales, M. Lopez-Sanchez, J. A. Rodriguez-Aguilar, M. Wooldridge, and W. Vasconcelos. Automated synthesis of normative systems. In *Proceedings of the 2013 AAMAS*, pages 483–490. IFAAMAS, 2013.

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

[6] A. Peleteiro, J. C. Burguillo, and S. Y. Chong. Exploring indirect reciprocity in complex networks using coalitions and rewiring. In *Proceedings of the 2014 AAMAS*, pages 669–676. IFAAMAS, 2014.

[7] D. G. Rand and M. A. Nowak. Human cooperation. *Trends in cognitive sciences*, 17(8):413–425, 2013.

[8] F. P. Santos, J. M. Pacheco, and F. C. Santos. Evolution of cooperation under indirect reciprocity and arbitrary exploration rates. *Scientific Reports*, 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.

[10] B. T. R. Savarimuthu and S. Cranefield. Norm creation, spreading and emergence: A survey of simulation models of norms in multi-agent systems. *Multiagent and Grid Systems*, 7(1):21–54, 2011.

[11] S. Sen and S. Airiau. Emergence of norms through social learning. In *IJCAI*, volume 1507, page 1512, 2007.

[12] M. Wooldridge. *An introduction to multiagent systems*. John Wiley & Sons, 2009.