

Coevolution of cognition and cooperation in structured populations under reinforcement learning

Rossana Mastrandrea^{a,*}, Leonardo Boncinelli^b, Ennio Bilancini^a

^a IMT Alti Studi Lucca, Piazza S. Francesco 19, 55100 Lucca, Italy

^b University of Florence, Italy

ARTICLE INFO

Keywords:

Dual process cooperation
Cognition
Prisoner's Dilemma
Network
Agent based model
Structured populations

ABSTRACT

We study the evolution of behavior under reinforcement learning in a Prisoner's Dilemma where agents interact in a regular network and can learn about whether they play one-shot or repeatedly by incurring a cost of deliberation. With respect to other behavioral rules used in the literature, (i) we confirm the existence of a threshold value of the probability of repeated interaction, switching the emergent behavior from intuitive defector to dual-process cooperator; (ii) we find a different role of the node degree, with smaller degrees reducing the evolutionary success of dual-process cooperators; (iii) we observe a higher frequency of deliberation.

1. Introduction

The evolution of cooperation has been investigated intensely in various disciplines, such as biology, economics, computer science, physics and psychology. There are two important dimensions, among many [1–3], that have been shown to affect the evolution of cooperation: the interaction structure, i.e., who interacts with whom [4], and the mode of cognition, i.e., the extent of deliberation as opposed to intuition [5]. While for the interaction structure there is a substantial consensus that sparse and heavily clustered networks help the spread of cooperation [1,6], for the mode of cognition results are more articulated and depend on specific features of the social dilemma [7,8] and of the cost of deliberation [9].

An important aspect in evolutionary models is the behavioral rule adopted by agents, which heavily contributes to determining the trajectories of the dynamic adjustment. The literature on evolutionary dynamics of cooperation and interaction has extensively considered behavioral rules encompassing best reply [10], imitation [11] and processes of the type death–birth or birth–death [6]. More recently, reinforcement learning has also been considered [12–15]. Reinforcement learning is indeed a prominent behavioral rule originated in behavioral sciences [16,17] and recently become extremely popular in computer and social sciences, with many different applications [18,19]. Reinforcement learning, which is a particular type of Q-learning [20], is known for its capacity to solve complex problems, at least those with a fixed environment [21], in a very parsimonious way in terms

of informational requirements. However, reinforcement learning performs poorly to attain cooperation in a system with many interacting agents [12], unless the interaction structure is endogenous [22].

We contribute to the literature by addressing the effects of inserting reinforcement learning into the model of Mosleh and Rand [23], which is basically the same model of Bear and Rand [7] with the addition of an interaction structure in the form of a k -regular lattice. In their model, agents accumulate over time payoffs across games played with all of their neighbors, and the agents who are randomly selected to update behavior copy the strategy of a neighbor with a probability that is increasing in the neighbor's accumulated payoff. In particular, agents are involved in either a one-shot prisoner dilemma or in a repeated prisoner dilemma; the latter interaction is represented in a reduced form as a coordination game, with an efficient Nash equilibrium where cooperation is sustained through conditional strategies and an inefficient (but possibly risk-dominant) Nash equilibrium where defection occurs ([24,25]; for a survey, see [26]). In our model, when behavior is updated an agent increases the probability of choices that have best performed in the past, regarding the cognitive mode and the action in the Prisoner's Dilemma conditional on the information acquired.

We provide three main results. First, we confirm previous findings, obtained with other behavioral rules [7,23], that for low probabilities of repeated interaction the intuitive defector behavior (i.e., always defect, never deliberate) is favored by evolution, while for higher probabilities it is the dual-process cooperator behavior (i.e., cooperate if not deliberating, deliberate to switch to defection when payoff-maximizing) to be favored. The intuition is similar to the one of the

* Corresponding author.

E-mail address: rossana.mastrandrea@imtlucca.it (R. Mastrandrea).

previous literature [7–9,23]: when the probability of repeated interaction is very low, it is not worth incurring the cost of deliberating to learn if the current interaction is actually repeated or one-shot, with the result that agents always expect the interaction to be mostly one-shot and systematic defection becomes very attractive. In this respect, reinforcement learning does not change the quality of results.

Second, in contrast with specific [23] and general [1,6] findings of the previous literature that fewer connections help the spread of cooperation, we find that a lower degree in the regular network hampers the adoption of dual-process cooperation, making intuitive defection more likely to be adopted. Therefore, reinforcement learning drastically changes the effects of a higher number of connections, which here becomes a factor promoting cooperation.

Third, we find that reinforcement learning increases the observed frequency of deliberation, for every probability of repeated interaction, comparing our results with those of Mosleh and Rand [23]. It can be interesting to notice that a cognitively cheap behavioral rule, as reinforcement learning is, is associated with a higher reliance of the more demanding cognitive mode of deliberation. In turn, this greater reliance on deliberation has an impact on cooperation rates, depending on whether deliberation is more cooperative than intuition (which happens for low probabilities of repeated interaction) or less cooperative (which happens for high probabilities of repeated interaction).

Overall, reinforcement learning does not change the conclusion that there exists a threshold value of the probability of repeated interaction switching the emergent behavior from intuitive defector to dual-process cooperator. At the same time, we can conclude that the role of the number of connections for the spread of cooperation is moderated by the behavioral rule, with a higher network degree favoring cooperation under reinforcement learning. Finally, we observe that reinforcement learning increases the observed frequency of deliberation.

2. Methods

In this section we introduce the model, the dynamics and simulations setup, highlighting similarities and differences with [23].

2.1. The model

Our results are based on agent based simulations. Here, we describe the model setup.

Population. The representative population consists of $N = 100$ agents simultaneously playing a modified version of the Prisoner's Dilemma (PD). At each iteration, the two interacting agents can play Tit-for-Tat (TFT) or always defect (AllD) in a one-shot or repeated game. In the one-shot game, *TFT* corresponds to cooperation, while *AllD* stands for defection. From now on, we will use for simplicity *C* or *D* to indicate the two choices in both scenarios.

Stage game. With probability p_G the stage game is a repeated PD which, given the restriction to strategies TFT and AllD, can be represented as a one-shot symmetric coordination game with the following payoff matrix (for the row player)

	<i>TFT</i>	<i>AllD</i>
<i>TFT</i>	$b - c$	0
<i>AllD</i>	0	0

With probability $1 - p_G$ the stage game agents is a one-shot symmetric PD with the following payoff matrix (for the row player)

	<i>TFT</i>	<i>AllD</i>
<i>TFT</i>	$b - c$	$-c$
<i>AllD</i>	b	0

The parameter b represents the benefits of cooperation generated for the other agent while c is the cost of cooperation incurred by the cooperating agent. In our simulations we set $b = 4$ and $c = 1$.

Deliberation. At each round, a deliberation cost d^* is randomly drawn from a uniform distribution in $[0, 1]$. Each agent has her own individual threshold cost for deliberation, d_i , $i \in \{1, \dots, N\}$: if $d_i \leq d^*$, the agent will deliberate acquiring information about the game type (one-shot/repeated); if $d_i > d^*$, she will act under intuition ignoring the type of game.

Interaction network. All agents are placed on a regular lattice with fixed number of neighbors for each agent, i.e. fixed *node degree* k . We consider different degree values: $k \in \{2, 4, 8, 20, 40\}$. The graph is mathematically represented by an adjacency matrix $A \equiv (a_{ij})_{1 \leq i, j \leq N}$, where $a_{ij} = 1$ means that agent i is linked to agent j , while $a_{ij} = 0$ represents the absence of agents connection.

Each agent is characterized by a vector of parameters, completely defining her strategy:

$$(p_{i,int}, p_{i,del}^{1s}, p_{i,del}^{rep}, d_i) \quad i \in \{1, \dots, N\} \quad (1)$$

where $p_{i,int}$ is the probability to play *C* under intuition (i.e., independently on the type of game); $p_{i,del}^{1s}$ is the probability to play *C* under deliberation in a one-shot PD; $p_{i,del}^{rep}$ is the probability to play *C* under deliberation in a repeated PD; d_i is the agent deliberation cost threshold.

At each round, all pairs of connected agents are selected to play the game. The random cost $d^* \sim U[0, 1]$ is drawn at the beginning of each round and kept fixed for all playing agents.

2.2. Dynamics

When all agents have played, the *reinforcement learning* rule is applied to update the strategy of each agent, both in terms of actions taken and deliberation cost threshold. At the end of the game, each agent compares the average payoffs of each choice and game-type and decide to increase/decrease the related probabilities accordingly. We want to stress that each agent can play different strategies (cooperate/defect or decide intuitively/deliberately) with different neighbors in the same round, but according to the same strategy vector (1) kept fixed. Indeed, the update of the probabilities/cost is done only when all agents have played with all neighbors.

Deliberation cost. At each round and for each agent, we compute the experiences payoffs of playing under intuition and under deliberation at that time, averaged over the two possible choices *C* or *D*: $\pi_{i,int}$, $\pi_{i,del}$. Then, the individual deliberation cost is updated according to the following rule:

$$d_i = \begin{cases} d_i + \lambda & \text{if } \pi_{i,int} < \pi_{i,del} - d_i \\ d_i - \lambda & \text{if } \pi_{i,int} > \pi_{i,del} - d_i \end{cases} \quad (2)$$

with $\lambda = 0.1$.

In other terms, the agent increases (decreases) her deliberation cost of a small amount, λ , if the average payoff obtained under deliberation – net of the deliberation cost – is greater (smaller) than the average payoff obtained acting under intuition. If $\pi_{i,int} = \pi_{i,del} - d_i$ there is no update of the threshold cost. We set 0 and 1 respectively as the maximum and minimum values for the deliberation cost threshold, therefore if $d_i < 0$ we approximate it with 0 and if $d_i > 1$ we approximate it with 1. This choice avoids meaningless outcomes, such as a negative or extremely high deliberation cost threshold, without loss of generality. Indeed, the effect of having a deliberation threshold cost equal to 0 (1) is exactly the same of having a negative (greater than 1) one, but speeding up the process avoiding to reach very small (high) threshold values. Finally, it is worth noting that the update is done only when both payoffs are observed (i.e., the agent acts both under deliberation and under intuition at least once in the same round).

Action strategy. At the end of each round, the following average payoffs are computed: (i) the payoffs resulting from cooperation, $\pi_{i,int}^C$, and defection, $\pi_{i,int}^D$, under intuition; (ii) the payoffs resulting from cooperation, $\pi_{i,del,1s}^C$, and defection, $\pi_{i,del,1s}^D$, under deliberation when the game is one-shot; (iii) the payoffs resulting from cooperation, $\pi_{i,del,rep}^C$, and defection, $\pi_{i,del,rep}^D$, under deliberation when the game is repeated.

Then, the strategy played by every agent i is updated according to the following rules:

$$p_{i,int} = \begin{cases} p_{i,int} + \epsilon & \text{if } \pi_{i,int}^C > \pi_{i,int}^D \\ p_{i,int} - \epsilon & \text{if } \pi_{i,int}^C < \pi_{i,int}^D \end{cases} \quad (3)$$

$$p_{i,del,1s} = \begin{cases} p_{i,del,1s} + \epsilon & \text{if } \pi_{i,del,1s}^C > \pi_{i,del,1s}^D \\ p_{i,del,1s} - \epsilon & \text{if } \pi_{i,del,1s}^C < \pi_{i,del,1s}^D \end{cases} \quad (4)$$

$$p_{i,del,rep} = \begin{cases} p_{i,del,rep} + \epsilon & \text{if } \pi_{i,del,rep}^C > \pi_{i,del,rep}^D \\ p_{i,del,rep} - \epsilon & \text{if } \pi_{i,del,rep}^C < \pi_{i,del,rep}^D \end{cases} \quad (5)$$

with $\epsilon = 0.1$.

The probability of cooperating in the three cases (under intuition, under deliberation when the game is one-shot, under deliberation when the game is repeated) is increased (decreased) by a small amount, ϵ , if the related payoff averaged over all games played by the agent in that round (i.e., its number of neighbors or node degree) is greater (smaller) than the average defection payoff obtained respectively in the same three scenarios. No update takes place in the above thresholds in case of equality or if one of the payoffs to be compared is absent.

Perturbations. We assume that with probability $\mu = 0.05$ a mutation occurs: the agent updates her action strategy or her deliberation cost inverting the inequalities in (2), (3), (4), (5). This choice introduces the possibility of a random mistake in the update process at each decision level.

2.3. Simulations

We considered a discrete range of variability for the probability to have a repeated game with step equal to 0.1: $p_G \in \{0, 0.1, 0.2, \dots, 1\}$. At the beginning of each simulation run, all strategy vectors are independently initialized from a uniform distribution. Each simulation run continued over a number of generations until no more than one agent updates its strategy for 10^2 generations. All results are averaged over 1000 initializations.

3. Results

In this work the same setting proposed by [23] has been employed, except for the behavioral rule used by agents for updating their strategy. Hence, the comparison between our and their outcomes allows us to focus on the effect of reinforcement learning. In the following sections, we highlight three main results standing out from this comparison.

Emergence of dual-process cooperation. Fig. 1(a) shows that as the probability of repeated interaction, p_G , increases, the observed frequency of cooperation under intuition gets larger and larger, starting from values close to 0.1 and reaching almost 1. These S-shaped curves correspond to the switch from situations mainly characterized by intuitive defection (for low values of p_G) to situations where dual-process cooperation prevails (for high values of p_G). This is qualitatively the same result as in Mosleh and Rand [23].

However, a difference emerges for small values of p_G : while the observed frequency of cooperation under intuition in Mosleh and Rand [23] is low for large values of k and high for small values of k , we find a low frequency of cooperation irrespective of k . This outcome

suggests that when a reinforcement learning update rule is employed, the network structure/number of neighbors becomes negligible for the emergence of cooperation under intuition when the game is mainly one-shot, remaining its value around 0.1 independently on the number of games played by each agent.

Dual-process cooperation and node degree. In contrast with Mosleh and Rand, we find that the curves associated with a smaller k are located further down, as shown in Fig. 1(a). Our results suggest that a higher number of neighbors favors the adoption of dual-process cooperation behavior, while the opposite is true for Mosleh and Rand [23]. This outcome highlights the role played by the reinforcement learning update rule: by increasing the number of games played by each agent in each round, the probability to move from intuitive defection to dual-process cooperation rapidly increases; on the contrary for $k = 2$ this passage arrives at high value of p_G (> 0.5). Fig. 1(b) summarizes this effect, showing a negative slope curve representing the relation between the node degree and the probability p_G (the analogous figure in Mosleh and Rand [23] exhibits instead a positive slope).

Furthermore, by looking at Figs. 1(c) and 1(d), we recognize that the frequency of cooperation under deliberation when the game is repeated rises rapidly as p_G increases, at least when the node degree is rather large, and they never go down. In our case, hence, we do not observe the frequency get pulled back towards 0.5 as in Mosleh and Rand [23], figure 3.

Frequency of deliberation. Fig. 1(e) shows the maximum threshold cost for deliberation averaged over all agents and for all values of the probability p_G . Even if the curves shapes (inverted V) are quite similar to the analogous graph of Mosleh and Rand, in our setting two main differences occur: (i) for all values of k , except for $k = 2$, we observe very close trends of the maximum deliberating cost with respect to the probability p_G with a slight leftward shift of the peak as k increases (exactly the opposite tendency exhibiting in the analogous figure of Mosleh's paper Mosleh and Rand [23]); (ii) the maximum threshold costs of deliberation double that shown in the paper by Mosleh and Rand Mosleh and Rand [23] for almost all values of p_G . Both outcomes, summarized in Fig. 1(f), suggest a higher frequency of deliberation when a reinforcement learning update rule is used.

The heat-maps in Fig. 2 provide a graphical summary of our main results showing the variation of the probability of cooperating under intuition, in Fig. 2(a), and the maximum threshold cost, in Fig. 2(b), with respect to the node degree, k , and the probability to have a repeated game, p_G . This is particularly useful to highlight the differences with the setup proposed by Mosleh [23], hence, to evaluate the effect of a reinforcement learning update rule.

Finally, we have repeated the same analysis with different network structures: Erdős-Rényi random graph, Watt-Strogatz and Barabasi-Albert model with varying node degree ($k \in \{2, 6, 8, 20, 40\}$). These findings show high similarity to the case of agents placed on regular lattices, (i) confirming that the observed outcomes are related to the network density rather than to other specific network features as already pointed out by Mosleh and Rand; (ii) stressing that the differences observed with respect to their setup are due to the introduction of a reinforcement learning update rule (figures available upon request).

4. Discussion

The debate in the literature whether people are more cooperative under intuition or deliberation has not focused much so far on the specific rule adopted by agents when updating behavior. In this paper we have shown that the main qualitative conclusion obtained in Mosleh and Rand [23] is robust to agents following a reinforcement learning rule: for a low probability of repeated interaction the prevailing behavior is intuitive defection, while dual-process cooperation spreads as such probability increases until it becomes prevalent for high probabilities. This pattern holds for every number of interacting partners. In

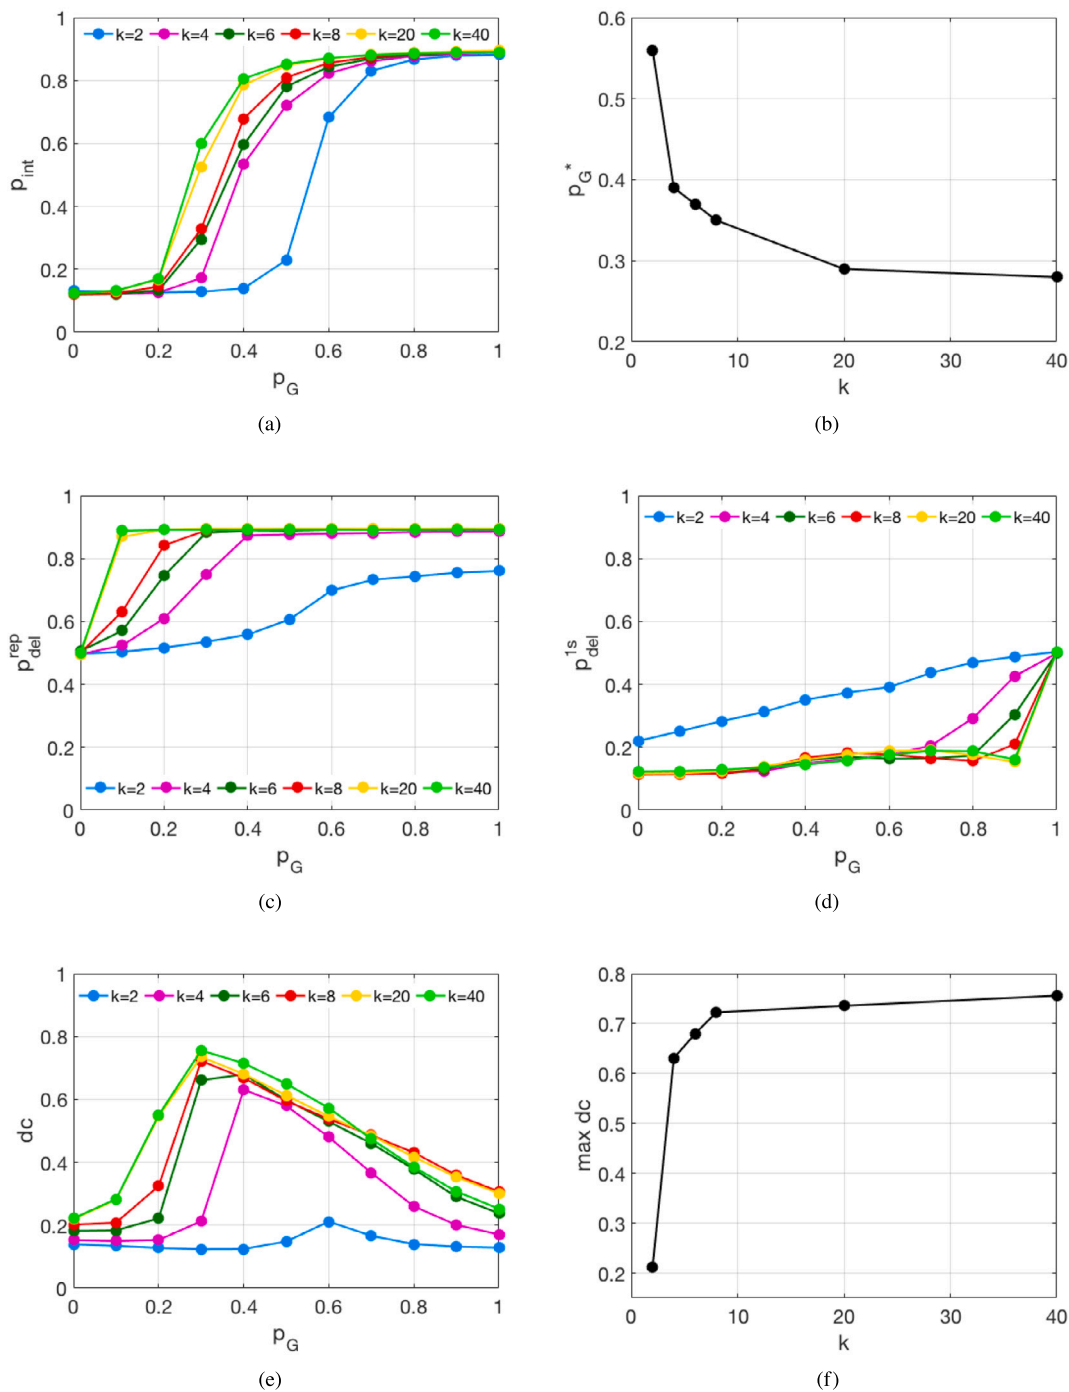


Fig. 1. Strategy evolution and network structure. (a) Probability of cooperating under intuition, p_{int} , as function of the probability to have a repeated game, p_G . (b) Critical value of the probability that the interaction is repeated, p_G , for which $p_{int} = 0.5$ across the different values of $k \in \{2, 4, 6, 8, 20, 40\}$. (c) Probability of cooperating under deliberation when the game is repeated, p_{del}^{rep} , as function of the probability that the game is repeated, p_G . (d) Probability of cooperating under deliberation if the game is one-shot, p_{del}^{1s} , as function of the probability that the game is repeated, p_G . (e) Maximum threshold cost of deliberation, d^* , as function of the probability that the game is repeated, p_G . (f) Maximum value of the deliberation cost across the whole discrete range of variability of the probability to have a repeated game, $p_G \in \{0, 0.1, \dots, 1\}$. In all panels the number of each node's neighbors is fixed, $k \in \{2, 4, 6, 8, 20, 40\}$.

this respect we differ from Mosleh and Rand [23], where dual-process cooperation is quite common also for low probabilities of repeated interaction when the number of interacting partners is small. Fig. 2 summarizes our main findings.

At the same time, there are important qualitative differences brought about by reinforcement learning. The most important one concerns the widespread belief in literature that a smaller number of interacting partners promotes the spread of cooperation. With reinforcement learning, we find the opposite: the smaller the number of partners, the more

difficult it is for the dual-process cooperation to prevail. The reasons for this result could be fruitfully investigated in future research.

Another interesting observation concerns the observed frequency of deliberation. The use of reinforcement learning, which is a behavioral rule that consumes few cognitive resources, increases the frequency of deliberation, which is more cognitively costly than intuition.

To explore potential avenues for future research, it is important to recognize that the dual-process approach has been applied in the literature beyond cooperation to a wide array of social decisions as reviewed




Fig. 2. Dual-process cooperation and network structure. Heat-maps of the (a) probability of cooperating under intuition, p_{int} , and (b) the maximum threshold cost of deliberation, dc^* , as a function of the number of neighbors (i.e., the number of games played by each agent in each round), k , and the probability to have a repeated game, p_G .

in a recent meta-analysis by Capraro [27]. This observations opens up a fertile ground for computational research aimed at unraveling the evolution of social behaviors from a dual-process standpoint. While there have been studies focusing on the evolution of trust, trustworthiness, and deceit, as highlighted by Kumar et al. [28] and Capraro et al. [29], a gap remains in applying a dual-process approach to these areas of investigation. The only contribution in this regard, to best of our knowledge, comes from the attempt by Zucchiatti [30] to investigate the coevolution of cognition and behaviors in the Ultimatum Game. Addressing this gap holds the promise of enriching our understanding of the dynamics underpinning these fundamental social behaviors.

In general, the results of our study call for further investigation into the role played by behavioral rules in explaining the prevalence of cooperation and intuition within models of dual-process theory. Along this research line, it would be interesting to explore cases where the behavioral rules differs between deliberation and intuition, while in this paper, and in related studies in the literature, only the available information changes between one cognitive mode and the other.

Finally, we aim to study how the network of interactions evolves over time in response to the payoffs earned. This would require a model where cognition and cooperation coevolve over a dynamic network, which we leave for further research.

CRediT authorship contribution statement

Rossana Mastrandrea: Writing – review & editing, Writing – original draft, Supervision, Methodology, Investigation, Formal analysis, Conceptualization. **Leonardo Boncinelli:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Conceptualization. **Ennio Bilancini:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Ennio Bilancini reports financial support was provided by Italian Ministry of Education, University and Research (MIUR), PRIN grant n. 20178293XT. Leonardo Boncinelli reports financial support was provided by Italian Ministry of Education, University and Research (MIUR), PRIN grant n. 20178293XT. Leonardo Boncinelli reports financial support was provided by IMT School for Advanced Studies Lucca, PAI Project PRO.CO.PE. Rossana Mastrandrea reports financial support was provided by IMT School for Advanced Studies Lucca, PAI Project PRO.CO.PE. Ennio Bilancini reports financial support was provided by IMT School for Advanced Studies Lucca, PAI Project PRO.CO.PE.

Rossana Mastrandrea reports financial support was provided by Horizon 2020 European, Grant Agreement n.871042. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgments

EB and LB acknowledge financial support from the Italian Ministry of Education, University and Research (MIUR) through the PRIN project Co.S.Mo.Pro.Be. “Cognition, Social Motives and Prosocial Behavior ” (grant n. 20178293XT); RM acknowledges support from the European Union, Horizon 2020 Program under the scheme “INFRAIA-01-2018-2019, Integrating Activities for Advanced Communities”, Grant Agreement n.871042, “SoBigData++: European Integrated Infrastructure for Social Mining and Big Data Analytics” (<http://www.sobigdata.eu>); all authors acknowledge support from the IMT School for Advanced Studies Lucca through the PAI project Pro.Co.P.E. “Prosociality, Cognition, and Peer Effects”. RM is member of the INDAM group *Gruppo Nazionale per l'Analisi Matematica, la Probabilità e le loro Applicazioni (GNAMPA)*”.

References

- [1] Nowak MA. Five rules for the evolution of cooperation. *Science* 2006;314(5805):1560–3.
- [2] Lehmann L, Keller L. The evolution of cooperation and altruism—a general framework and a classification of models. *J Evol Biol* 2006;19(5):1365–76.
- [3] Bowles S, Gintis H. *A cooperative species: Human reciprocity and its evolution*. Princeton University Press; 2011.
- [4] Santos FC, Pacheco JM, Lenaerts T. Cooperation prevails when individuals adjust their social ties. *PLoS Comput Biol* 2006;2(10):e140.
- [5] Capraro V. The dual-process approach to human sociality: A review. 2019, Available at SSRN 3409146.
- [6] Ohtsuki H, Hauert C, Lieberman E, Nowak MA. A simple rule for the evolution of cooperation on graphs and social networks. *Nature* 2006;441(7092):502–5.
- [7] Bear A, Rand DG. Intuition, deliberation, and the evolution of cooperation. *Proc Natl Acad Sci* 2016;113(4):936–41.
- [8] Bear A, Kagan A, Rand DG. Co-evolution of cooperation and cognition: the impact of imperfect deliberation and context-sensitive intuition. *Proc R Soc B* 2017;284(1851):20162326.
- [9] Jagau S, van Veelen M. A general evolutionary framework for the role of intuition and deliberation in cooperation. *Nat Hum Behav* 2017;1(8):1–6.
- [10] Bilancini E, Boncinelli L. The co-evolution of cooperation and defection under local interaction and endogenous network formation. *J Econ Behav Organ* 2009;70(1–2):186–95.

- [11] Levine DK, Pesendorfer W. The evolution of cooperation through imitation. *Games Econom Behav* 2007;58(2):293–315.
- [12] Tanabe S, Masuda N. Evolution of cooperation facilitated by reinforcement learning with adaptive aspiration levels. *J Theoret Biol* 2012;293:151–60.
- [13] Ezaki T, Horita Y, Takezawa M, Masuda N. Reinforcement learning explains conditional cooperation and its moody cousin. *PLoS Comput Biol* 2016;12(7):e1005034.
- [14] Horita Y, Takezawa M, Inukai K, Kita T, Masuda N. Reinforcement learning accounts for moody conditional cooperation behavior: experimental results. *Sci Rep* 2017;7(1):39275.
- [15] Jia D, Guo H, Song Z, Shi L, Deng X, Perc M, et al. Local and global stimuli in reinforcement learning. *New J Phys* 2021;23(8):083020.
- [16] Skinner BF. *The behavior of organisms: An experimental analysis*. Appleton-Century-Crofts; 1938.
- [17] Skinner BF. *Science and human behavior*. Macmillan; 1938.
- [18] Nian R, Liu J, Huang B. A review on reinforcement learning: Introduction and applications in industrial process control. *Comput Chem Eng* 2020;139:106886.
- [19] Perc M, Gómez-Gardenes J, Szolnoki A, Floría LM, Moreno Y. Evolutionary dynamics of group interactions on structured populations: a review. *J R Soc Interface* 2013;10(80):20120997.
- [20] Wang L, Jia D, Zhang L, Zhu P, Perc M, Shi L, Wang Z. Lévy noise promotes cooperation in the prisoner's dilemma game with reinforcement learning. *Nonlinear Dynam* 2022;108(2):1837–45.
- [21] Watkins CJ, Dayan P. Q-learning. *Mach Learn* 1992;8:279–92.
- [22] Song Z, Guo H, Jia D, Perc M, Li X, Wang Z. Reinforcement learning facilitates an optimal interaction intensity for cooperation. *Neurocomputing* 2022;513:104–13.
- [23] Mosleh M, Rand DG. Population structure promotes the evolution of intuitive cooperation and inhibits deliberation. *Sci Rep* 2018;8(1):1–8.
- [24] Blonski M, Ockenfels P, Spagnolo G. Equilibrium selection in the repeated prisoner's dilemma: Axiomatic approach and experimental evidence. *Am Econ J Microecon* 2011;3(3):164–92.
- [25] Dal Bó P, Fréchette GR. The evolution of cooperation in infinitely repeated games: Experimental evidence. *Amer Econ Rev* 2011;101(1):411–29.
- [26] Newton J. Evolutionary game theory: A renaissance. *Games* 2018;9(2):31.
- [27] Capraro V. The dual-process approach to human sociality: Meta-analytic evidence for a theory of internalized heuristics for self-preservation. *J Personal Soc Psychol* 2024.
- [28] Kumar A, Capraro V, Perc M. The evolution of trust and trustworthiness. *J R Soc Interface* 2020;17(169):20200491.
- [29] Capraro V, Perc M, Vilone D. The evolution of lying in well-mixed populations. *J R Soc Interface* 2019;16(156):20190211.
- [30] Zucchiatti S. *Essays on the evolution of prosocial behaviors* [Ph.D. thesis], IMT School for Advanced Studies Lucca; 2023.