

Learning dynamics explains human behaviour in Prisoner's Dilemma on networks

Giulio Cimini and Angel Sánchez

J. R. Soc. Interface 2014 **11**, 20131186, published 19 February 2014

Supplementary data

["Data Supplement"](#)

[http://rsif.royalsocietypublishing.org/content/suppl/2014/02/14/rsif.2013.1186.DC1.htm](http://rsif.royalsocietypublishing.org/content/suppl/2014/02/14/rsif.2013.1186.DC1.html)
|

References

[This article cites 37 articles, 6 of which can be accessed free](#)

<http://rsif.royalsocietypublishing.org/content/11/94/20131186.full.html#ref-list-1>

Subject collections

Articles on similar topics can be found in the following collections

[mathematical physics](#) (45 articles)

Email alerting service

Receive free email alerts when new articles cite this article - sign up in the box at the top right-hand corner of the article or click [here](#)

Research



Cite this article: Cimini G, Sánchez A. 2014 Learning dynamics explains human behaviour in Prisoner's Dilemma on networks. *J. R. Soc. Interface* **11**: 20131186.
<http://dx.doi.org/10.1098/rsif.2013.1186>

Received: 19 December 2013

Accepted: 29 January 2014

Subject Areas:

mathematical physics

Keywords:

evolutionary game theory, Prisoner's Dilemma, social networks, moody conditional cooperation, reinforcement learning

Author for correspondence:

Giulio Cimini

e-mail: g.cimini@math.uc3m.es

Electronic supplementary material is available at <http://dx.doi.org/10.1098/rsif.2013.1186> or via <http://rsif.royalsocietypublishing.org>.

Learning dynamics explains human behaviour in Prisoner's Dilemma on networks

Giulio Cimini¹ and Angel Sánchez^{1,2}

¹Grupo Interdisciplinar de Sistemas Complejos (GISC), Departamento de Matemáticas, Universidad Carlos III de Madrid, 28911 Leganés, Madrid, Spain

²Instituto de Biocomputación y Física de Sistemas Complejos (BIFI), Universidad de Zaragoza, 50018 Zaragoza, Spain

Cooperative behaviour lies at the very basis of human societies, yet its evolutionary origin remains a key unsolved puzzle. Whereas reciprocity or conditional cooperation is one of the most prominent mechanisms proposed to explain the emergence of cooperation in social dilemmas, recent experimental findings on networked Prisoner's Dilemma games suggest that conditional cooperation also depends on the previous action of the player—namely on the 'mood' in which the player is currently in. Roughly, a majority of people behave as conditional cooperators if they cooperated in the past, whereas they ignore the context and free ride with high probability if they did not. However, the ultimate origin of this behaviour represents a conundrum itself. Here, we aim specifically to provide an evolutionary explanation of moody conditional cooperation (MCC). To this end, we perform an extensive analysis of different evolutionary dynamics for players' behavioural traits—ranging from standard processes used in game theory based on pay-off comparison to others that include non-economic or social factors. Our results show that only a dynamic built upon reinforcement learning is able to give rise to evolutionarily stable MCC, and at the end to reproduce the human behaviours observed in the experiments.

1. Introduction

Cooperation and defection are at the heart of every social dilemma [1]. While cooperative individuals contribute to the collective welfare at a personal cost, defectors choose not to. Owing to the lower individual fitness of cooperators arising from that cost of contribution, selection pressure acts in favour of defectors, thus making the emergence of cooperation a difficult puzzle. Evolutionary game theory [2] provides an appropriate theoretical framework to address the issue of cooperation among selfish and unrelated individuals. At the most elementary level, many social dilemmas can be formalized as two-person games where each player can either cooperate (C) or defect (D). The *Prisoner's Dilemma* (PD) game [3] has been widely used to model a situation in which mutual cooperation leads to the best outcome in social terms, but defectors can benefit the most individually. In mathematical terms, this is described by a pay-off matrix (entries correspond to the row player's pay-offs)

	C	D
C	R	S
D	T	P

where mutual cooperation yields the reward R , mutual defection leads to punishment P , and the mixed choice gives the cooperator the sucker's pay-off S and the defector the temptation T . The essence of the dilemma is captured by $T > R > P > S$: both players prefer any outcome in which the opponent cooperates, but the best option for both is to defect. In particular, the temptation to cheat ($T > R$) and the fear of being cheated ($S < P$) can put cooperation at risk, and according to the principles of Darwinian selection, cooperation extinction is inevitable [4].

Despite the conclusion above, cooperation is indeed observed in biological and social systems alike [5]. The evolutionary origin of such cooperation hence remains a key unsolved issue, particularly because the manner in which individuals adapt their behaviour—which is usually referred to as evolutionary dynamics or strategy update—is unknown *a priori*. Traditionally, most of the theoretical studies in this field have built on update rules based on pay-off comparison [6–8].¹ While such rules fit in the framework of biological evolution, where pay-off is understood as fitness or reproductive success, they are also questionable, especially from an economic perspective, as it is often the case that individuals perceive the others' actions but not how much they benefit from them. Indeed, experimental observations [11–13] (with some exceptions [14], but see also the reanalysis of those data in [15]) point out that human subjects playing PD or Public Good games do not seem to take pay-offs into consideration. Instead, they respond to the cooperation that they observe in a reciprocal manner, being more prone to contribute the more their partners do.

Reciprocity [16] has been studied in two-player games through the concept of reactive strategies [17], the most famous of which is *Tit-For-Tat* [18] (given by playing what the opponent played in the previous run). Reactive strategies generalize this idea by considering that players choose their action with probabilities that depend on the opponent's previous action. A further development was to consider memory-one reactive strategies [17], in which the probabilities depend on the previous action of both the focal player and her opponent. In multiplayer games, conditional cooperation, i.e. the dependence of the chosen strategy on the amount of cooperation received, had been reported in related experiments [11] and observed also for the spatial iterated PD [14] (often along with a large percentage of free-riders). The analysis of the two largest scale experiments to date with humans playing an iterated multiplayer PD game on a network [12,13] extended this idea by including the dependence on the focal player's previous action, giving rise to the so-called *moody conditional cooperation* (MCC).

The MCC strategy can be described as follows [19]: if in the previous round the player defected, she will cooperate with probability $p_D = q$ (approximately independently of the observed cooperation), whereas, if she cooperated, she will cooperate again with a probability $p_C(x) = px + r$ (subject to the constraint $p_C(x) \leq 1$), where x is the fraction of cooperative neighbours in the previous round. There is ample evidence supporting this aggregate behaviour, as it has been observed in at least five independent experiments: the two already quoted [12,13]; another one on multiplayer PD [20]; a lab-in-the-field experiment with people attending a fair in Barcelona, where participants in the age range 17–87 behaved consistently according to the MCC strategy [21], and finally, in [14], as revealed by a recent meta-analysis of those experimental results [15]. On the other hand, it could be argued that MCC behaviour arises from learning processes experienced by the players. In this respect, it is true that when a number of iterations of the PD is regarded as a single 'supergame', repetitions of such supergame show changes in behaviour [22]. This is in agreement with the observations in [12], where two repetitions of the supergame were carried out with the same players (Experiments 1 and 2 in the reference), and it was found that the initial behaviour was indeed different in both. However, analysis that exclude the first few rounds of those experiments shows clear evidence for MCC behaviour which, if anything,

becomes even more marked in the second one. Similar analysis were carried out in all other experiments, precisely to check for the effects of learning, finding in all cases strong evidence in support of the MCC strategy, even in [20], where 100 iterations of the PD were played. Therefore, we are confident that the observation of MCC behaviour is reproducible and correctly interpreted, and we believe it is a good framework to study the problem as we propose here. However, from the viewpoint of ultimate origins and evolutionary stability of this kind of behaviour, conditional cooperation and its moody version are a puzzle themselves. For instance, theoretical results based on replicator dynamics show that the coexistence of moody conditional cooperators with free-riders is not possible beyond very small groups [19]. Additionally, whereas the strategies reported in [12,13] are aggregate behaviours, it is not clear how individual MCC behavioural profiles $\{q, p, r\}$ evolve in time and how many evolutionarily stable profiles can exist among the players.

Here, we aim precisely to address these issues by developing and studying a model for the evolutionary dynamics of MCC behavioural traits. To this end, we perform agent-based simulations of a population consisting of N differently parameterized moody conditional cooperators, either on a well-mixed population or placed on the nodes of a network, who play an iterated PD game with their neighbours (which is the same setting used in recent experiments [12–14]) and whose behavioural parameters $\{q, p, r\}$ are subject to a strategy update process. Specifically, during each round t of the game, each player selects which action to take (C or D) according to her MCC traits, then plays a PD game with their neighbours—the chosen action being the same with all of them—and collects the resulting pay-off π^t . Subsequently, for every τ rounds, players may update their MCC parameters according to a given evolutionary rule.

The key and novel point in this study is that we explore a large set of possible update rules for the MCC parameters, whose details are given in the electronic supplementary material, Material and methods. To begin with, the first set of rules that we consider are of imitative nature, in which players simply copy the parameters from a selected counterpart. Imitation has been related to bounded rationality or to a lack of information that forces players to copy the strategies of others [23]. The rules that we consider here cover different aspects of imitation. Thus, we study the classical imitative dynamics that are based on pay-off comparison: stochastic rules as *Proportional Imitation* [24] (equivalent, for a large and well-mixed population, to the replicator dynamics [6]), the *Fermi rule* [25] (featuring a parameter β that controls the intensity of selection, and that can be understood as the inverse of temperature or noise in the update rule [26,27]) and the *Death–Birth rule* (inspired on Moran dynamics [28]), as well as the deterministic dynamics given by *Unconditional Imitation* (also called 'Imitate the Best') [29]. In all these cases, players decide to copy one of their neighbours with a probability (that may be 1, i.e. with certainty) that depends in a specific manner on the pay-offs that they and their partners obtained in the previous round of the game. To widen the scope of our analysis, we also analyse another imitative mechanism that is not based on pay-off comparison, namely the *Voter model* [30], in which players simply follow the social context without any strategic consideration [31]. Finally, in order to go beyond pure imitation, we also consider another two evolutionary dynamics which are innovative, meaning that they

allow extinct strategies to be reintroduced in the population (whereas imitative dynamics cannot do that). The first one is *Best Response* [26,32], a rule that has received a lot of attention in the literature, especially in economics, and that represents a situation in which each player has enough cognitive abilities to compute an optimum strategy given what her neighbours did in the previous round. The second one is *Reinforcement Learning* [33–35], which instead embodies the condition of a player that uses her experience to choose or avoid certain actions based on their consequences: actions that met or exceeded aspirations in the past tend to be repeated in the future, whereas choices that led to unsatisfactory experiences are avoided. Note that neither of these two last rules relies on the use of information on others' pay-offs.

With the different update schemes that we have summarized above, we have an ample spectrum of update rules representing most of the alternatives that have been proposed to implement evolutionary dynamics. The point of considering such a comprehensive set is directly related to our aim: finding how evolution, in a broad sense, can give rise to situations that are compatible with those seen in the experiments [12,13], in terms of values and stationarity of the MCC parameters, as well as of the final level of cooperation achieved. Additionally, we study different spatial structures determining the interactions among the players: the simple set-up of a well-mixed population (modelled by a random graph of average degree $\bar{k} = m$ rewired after each round of the game), as well as more complex structures—such as Barabási–Albert scale-free network [36] (with degree distribution $P(k) \sim k^{-3}$ and $\bar{k} = m$) and regular lattices with periodic boundary conditions (where each node is connected to its $k = m$ nearest neighbours) as used in the available experimental results. In doing so, we add another goal to our research, namely to check whether evolution can also explain the observed lack of *network reciprocity* [37], which is another important experimental outcome [12,13]. Indeed, experimental results show very clearly that, when it comes to human behaviour, the existence of an underlying network of contacts does not have any influence on the final level of cooperation. Therefore, any evolutionary proposal to explain the way subjects behave in the experiments must also be consistent with this additional observation.

2. Results and discussion

We have carried out an extensive simulation program on the set of update rules and underlying networks that we have introduced above. In what follows, we separate the discussion of the corresponding results into two main groups: imitative and non-imitative strategies. Additional aspects of our numerical approach are described in the electronic supplementary material, Results.

2.1. Imitative updates

The five topmost sets of plots of figure 1 show the evolution of the level of cooperation c (defined as the percentage of players who cooperate in each round of the game), as well as the stationary probability distribution of the individual MCC parameters among the population, when different evolutionary dynamics are employed to update players' behavioural traits. Note that all the plots refer to the case $\tau = 1$ (meaning that the update takes place after each round). We will show only results for this choice below, because we have observed that

the value of τ basically influences only the convergence rate of the system to its stationary state, but not its characteristic features. As can be seen from the plots, the final level of cooperation here is, generally, highly dependent on the population structure, and often the final outcome is a fully defective state (especially for a well-mixed population).² Then, as expected from non-innovative strategies, the number of profiles $\{q, p, r\}$ that survive at the end of the evolution is always very low and, in general, only one profile is left for every individual realization of the system. Notwithstanding, the surviving profiles are very different among independent realizations (except when the final outcome is full defection, where $q \rightarrow 0$ irrespectively of p and r), indicating the absence of a stationary distribution for MCC parameters, i.e. the lack of evolutionarily stable profiles. The only case in which the parameters q and r tend to concentrate around some stationary non-trivial values is given by games played on lattices and with unconditional imitation updating. Finally, we note that, when the update rule is the voter model, the surviving profile is just picked randomly among the population (as expected from a rule that is not meant to improve pay-offs), and hence the cooperation level remains close on average to the value set by the initial distribution of MCC parameters. A similar behaviour is observed with the Fermi rule for low β , where β is the parameter that controls the intensity of the selection. Whereas for high β (low temperature), errors are unlikely to occur and players always choose the parameters that enhance their pay-offs, resulting in full defection as final outcome, for low β (high temperature) errors are frequent, so that MCC parameters basically change randomly and c remains close to its initial value. It is also worth noting that proportional imitation and the Fermi rule lead to very similar results, except for the parameter q , which makes sense in the view that they are very similar unless β is very small. The fact that both the Fermi rule and the death–birth update lead also to similar outcomes is probably related to those two dynamics being both error-prone, with specific features of either one showing, for instance, in the different results on lattices. Nonetheless, beyond all these peculiarities of each imitative dynamics, the main conclusion of our simulation program is that this type of update scheme is not compatible with the experimental observations.

Note that it is not our goal to explain in detail the effects of a particular updating rule on a given population structure. However, it is possible to gain qualitative insights into the behaviour of the system from rather naive considerations. Take for instance, scale-free networks, which feature hubs (players with high degree) that thus get higher pay-offs than average players do. If the dynamics is of imitative nature, hubs' strategy is stable and tends to spread over the network: there is the possibility for a stable subset of cooperators to form around hubs [38]. This behaviour (which cannot occur in random or regular graphs, where the degree distribution is more homogeneous) is clearly visible when the updating rule is proportional imitation. Notably, the stability of the subset of cooperators is destroyed when mistakes are possible (as with the Fermi rule); on the other hand, it is enhanced when the updating selects preferentially individuals with high pay-offs (as with the death–birth rule or unconditional imitation). In these two latter cases, cooperation becomes sustainable also in lattices, as these structures naturally allow clusters of mutually connected cooperators to emerge. Instead, the independence on the network observed—as we shall see—in the case of reinforcement

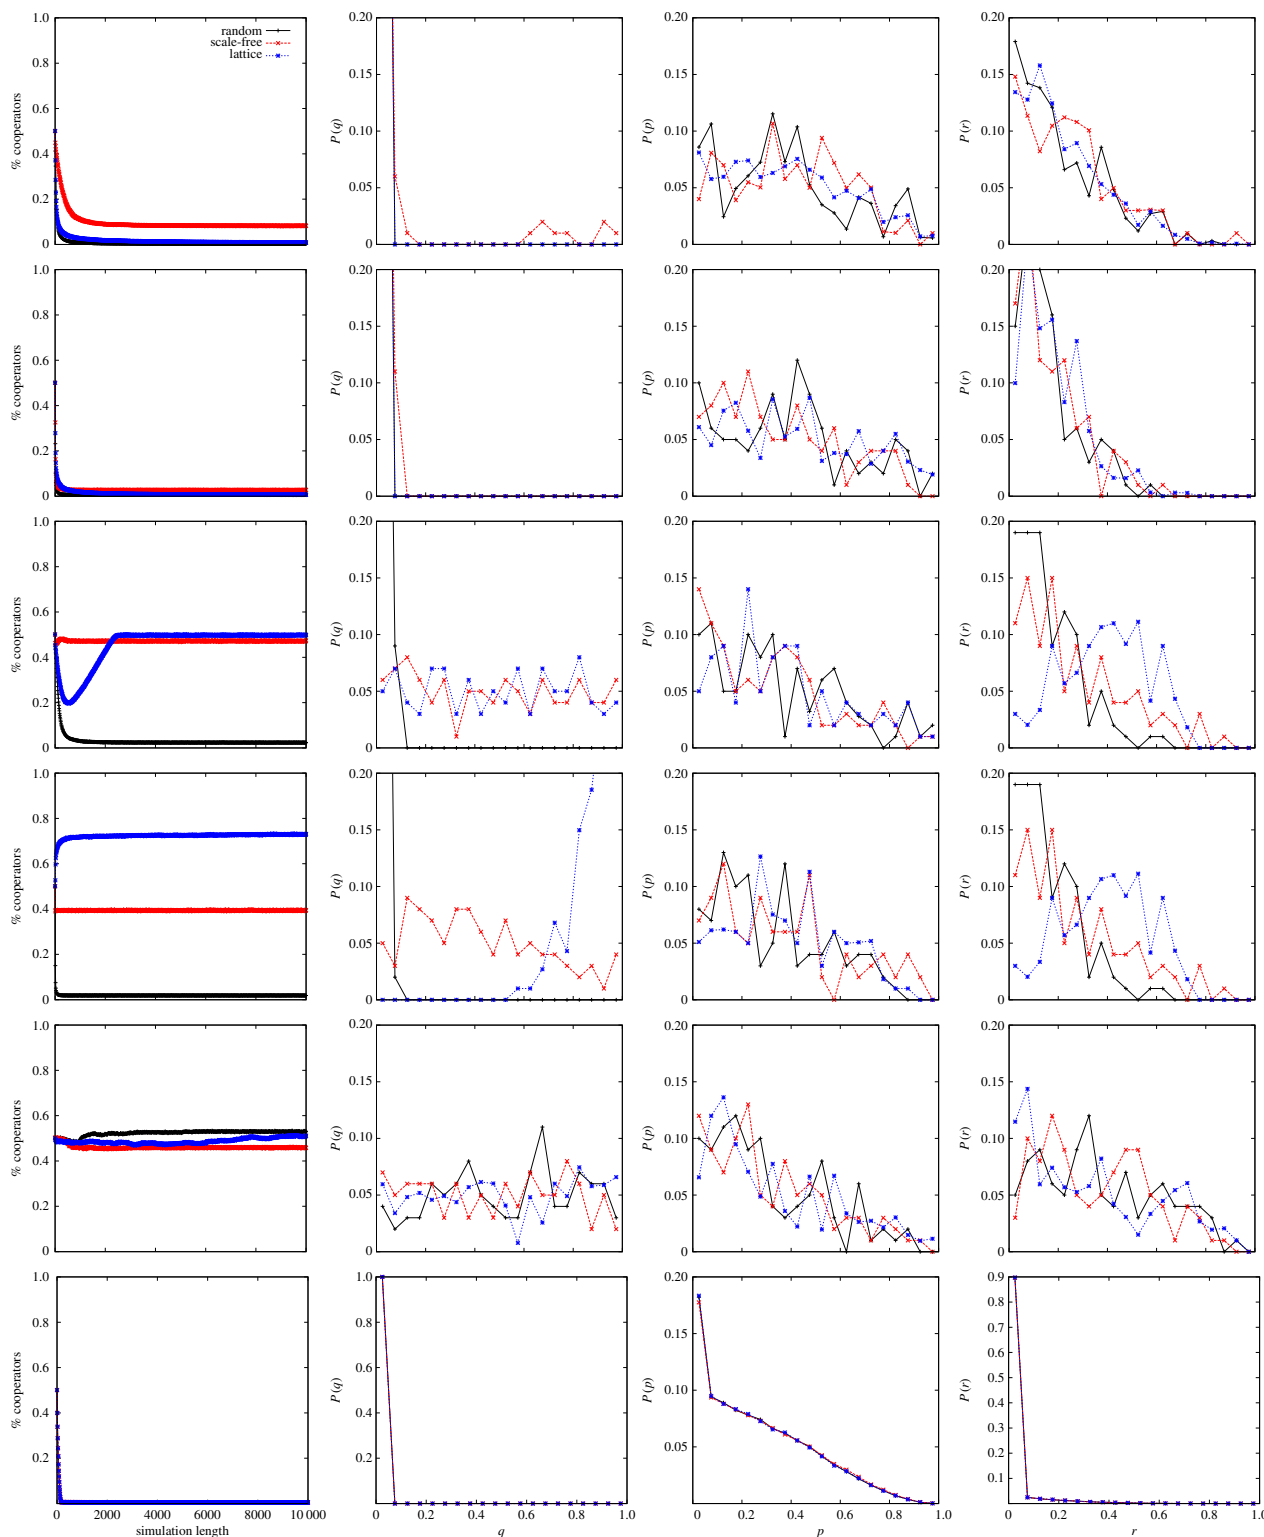


Figure 1. Evolution of the level of cooperation (left column) and stationary distributions of MCC parameters (from left to right: q , p and r) when the evolutionary dynamics is (from top to bottom): proportional imitation, Fermi rule with $\beta = 1/2$, death–birth rule, unconditional imitation, voter model and best response with $\delta = 10^{-2}$. Results are averaged over 100 independent realizations. (Online version in colour.)

learning is easily explained by players not even looking at each other, which makes the actual population structure irrelevant.

2.2. Non-imitative updates

A first general finding about this type of evolutionary rule is that, because of their own nature, they allow for a very large number of surviving MCC profiles ($\sim N$), even when the parameters tend to concentrate around specific values. The bottom set of plots of figure 1 summarizes our results for

the best response dynamics, which is the most ‘rational’ of the ones that we are studying here. For this choice, the system always ends up in a fully defective state, irrespectively of the network’s structure, which is the outcome that would be obtained by global maximization of the individual pay-offs. In this sense, the amount δ by which parameters are shifted at each update influences only the convergence rate of the system: higher δ arrives faster to full defection ($q = r = 0$). We then see that evolution by best response fails completely to explain any of the main experimental results.

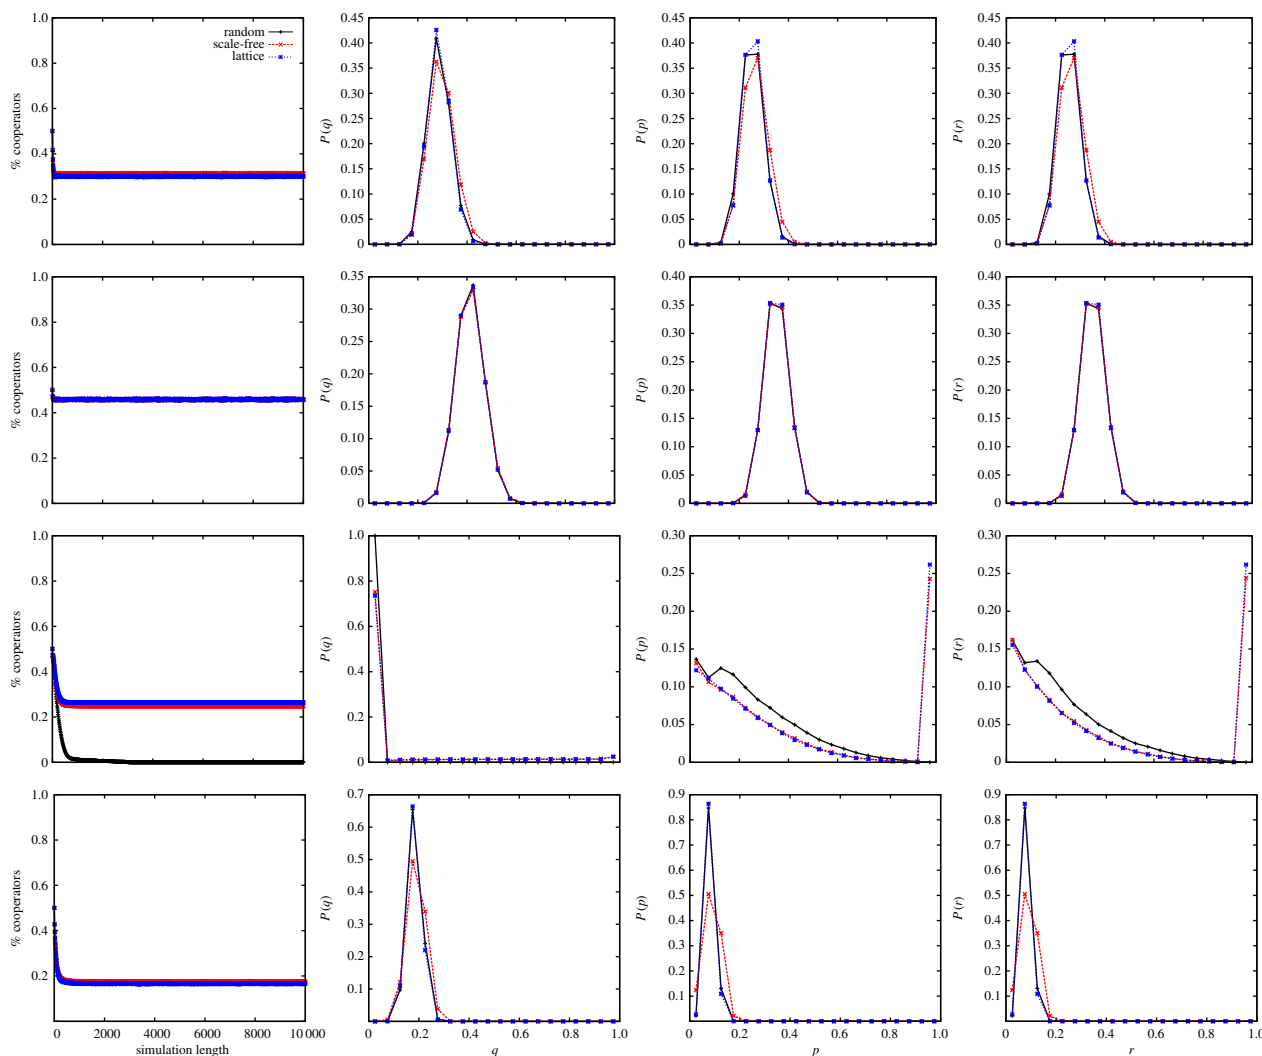


Figure 2. Evolution of the level of cooperation (left column) and stationary distributions of MCC parameters (from left to right: q , p and r) when the evolutionary dynamics is reinforcement learning with learning rate $\lambda = 10^{-1}$. From top to bottom: $A = 1/2$, $A = 5/4$, $A = -1/4$ and adaptive A with $h = 0.2$ and $A^0 = 1/2$. Results are averaged over 100 independent realizations. (Online version in colour.)

Our other rule of choice in this type is reinforcement learning. We will begin by assuming that aspiration levels A remain fixed in time. Our results regarding this rule are presented in figure 2. When A is midway between the punishment and reward pay-offs ($P < A < R$), we observe a stationary, non-vanishing level of cooperation around 30% that does not depend on the population structure. This behaviour, that is robust with respect to the learning rate λ , is in good qualitative agreement with the experimental observations [12,13]. However, the most remarkable outcome of this dynamic is that, contrary to all other update procedures that we have discussed so far, the values of the MCC parameters $\{q, p, r\}$ concentrate around some stationary, non-trivial values which are independent of the population structure and of the initial conditions of the system. Indeed, we have checked that the stationary values of $\{q, p, r\}$ do not depend on the initial form of their distributions, and also that fixing one of these three parameters does not influence the stationary distributions of the others. More importantly, these values are compatible with the ones obtained by linear fits of the aggregate MCC behaviour extracted from the experiments [12,13]. Reinforcement learning thus represents the only mechanism (among those considered here) which is able to give rise to evolutionarily stable moody conditional cooperators, while at the same time reproducing the cooperation level and the lack of network reciprocity (note that, as we already said, the type of network on which

the population sits does not affect the cooperation level). It is worth mentioning two other features of this dynamics. First, we have checked that the value of λ influences only the convergence rate of the system; however, if players learn too rapidly ($\lambda \sim 1$) then the parameters change too quickly and too much to reach stationary values—a phenomenon typical of this kind of learning algorithms. Second, if we introduce in the system a fraction d of players who always defect (recall that full defectors coexist with moody conditional cooperators in the experiments), what happens is that the final cooperation level changes—it drops to 25% for $d = 0.2$ and to 20% for $d = 0.4$ —but the stationary distributions of MCC parameters are not affected. This means that reinforcement learning is able to account for the heterogeneity of the behaviours observed in the experimental populations, which is consistent with the fact that this update rule does not take into account either the pay-offs or the actions of the rest of the players.

Further evidence for the robustness of the reinforcement learning evolutionary dynamics arises from extending our study to other aspiration levels, including dynamic ones. In general, what we observe is that the higher A , the higher the final level of cooperation achieved. When $R < A < T$, players are not satisfied with the reward of mutual cooperation; however, an outcome of mutual defection leads to a great stimulus towards cooperation in the next round. This is why players'

parameters tend to concentrate around values that allow for a strategy which alternates cooperation and defection, and that brings to stationary cooperation levels around 50%. Instead, if $S < A < P$, then defection–defection is a satisfactory outcome for each pair of players. In this case, cooperation may thrive only on stationary networks (where clusters of cooperator may form). However, for a well-mixed population, the final state is necessarily fully defective ($q \rightarrow 0$). Hence, we observe in this case, a dependence on the network structure which is not observed in the experiments; nonetheless, setting an aspiration level below punishment is at least questionable. Therefore, unless players make very strange decisions on their expectations from the game, we find behaviours that agree qualitatively with the experiments. Finally, we consider the case in which players adapt their aspiration level after each round: $A^{t+1} \leftarrow (1-h)A^t + h\pi^t/k$, where h is the adaptation (or habituation) rate and $P < A^0 < R$. What we observe now is that the stationary level of cooperation lies around 20%, the absence of network reciprocity is recovered and players' average aspiration levels remain in the range $P < \bar{A} < R$. Thus, this case is again compatible with experimental observations, and the fact that aspiration levels of an intermediate character are selected (corresponding to the case that better describes them) provides a clear rationale for this choice in the preceding paragraph.

A final important validation of reinforcement learning comes from studying the experience-weighted attraction (EWA) updating [39], an evolutionary dynamics that combines aspects of belief learning models (to which best response belongs) and of reinforcement learning. Results for this choice of the updating scheme (which are reported in the electronic supplementary material, EWA) confirm in fact that reinforcement learning is the determinant contribution which achieves situations matching with empirical outcomes.

3. Conclusion

Understanding cooperation is crucial because all major transitions in evolution involve the spreading of some sort of cooperative behaviour [5]. In addition, the archetypical tensions that generate social dilemmas are present in fundamental problems of the modern world: resource depletion, pollution, overpopulation and climate change. This work, inspired by experiments [12,13], with the aim of finding an evolutionary framework capable of modelling and justifying real people behaviour in an important class of social dilemmas—namely PD games. To this end, we have studied the evolution of a population of differently parametrized MCC whose parameters can evolve. We have considered several rules for parameters' changes—both of imitative nature and innovative mechanisms, as well as rules based on pay-off comparison and others based on non-economic or social factors. Our research shows that reinforcement learning with a wide range of learning rates is the only mechanism able to explain the evolutionary stability of MCC, leading to situations that are in agreement with the experimental observations in terms of the stationary level of cooperation achieved, average values and stationary distributions of the MCC parameters and absence of network reciprocity. Note that we have considered only PD games; however, given that in our set-up, players have to play the same action with all their neighbours, it is clear that our results should be related to Public Goods experiments (where conditional cooperation was first observed [11]). Our

findings thus suggest that MCC can also arise and be explained through reinforcement learning dynamics in repeated Public Goods games.

We stress that this is a very relevant result, as for the first time to our knowledge we are providing a self-consistent picture of how people behave in PD games on networks. Indeed, starting from the observation that players do not take others' pay-offs into account, we find that if this behaviour is to be explained in an evolutionary manner, it has to be because people learn from what they experience, and not from the information they may gather on their neighbours. Such a learning process is in turn very sensible in the heavily social framework in which we as humans are embedded and compatible with the knowledge that we have on the effects of our choices on others. On the other hand, the evolutionary dynamics that our work eliminates as possibly responsible for how we behave are, in fact, difficult to justify in the same social context, either because they would require a larger cognitive effort (best response) or, on the contrary, because they assume a very limited rationality that only allows imitation without reflecting on how we have been affected by our choices. Our work thus provides independent evidence that, at least in the context of human subjects interacting in PD, the observed behaviours arise mostly from learning. Of course, this does not mean that other ways to update one's strategy are not possible: indeed, a large fraction of people have been observed to be full defectors, a choice they may have arrived at by considering the PD game from a purely rational viewpoint. In addition, specific individuals may behave in idiosyncratic ways that are not described within our framework here. Still, as we have seen, our main result, namely that reinforcement learning explains the behaviour of a majority of people and its macro-consequences (level of cooperation, lack of network reciprocity) would still hold true in the presence of these other people.

Although a generalization of our results to other classes of social dilemma beyond PD and Public Goods is not straightforward, our conclusions here should guide further research on games on networks. We believe that the experimental results, to which this work provides a firm theoretical support, allow the conclusion that many of the evolutionary dynamics used in theory and in simulations simply do not apply to the behaviour of human subjects and, therefore, their use should be avoided. As a matter of fact, much of the research published in the last decade by using all these update schemes is only adding confusion to an already very complicated problem. Even so, our findings do not exclude the plausibility of other strategy updating in different contexts. For instance, analytical results with imitative dynamics [40] display an agreement with experimental outcomes on dynamical networks [41], where it was also shown that selection intensity (which can be thought as a measure of players' rationality) can dramatically alter the evolutionary outcome [42]. It is also important to stress that our findings here relate to human behaviour, and other species could behave differently; for instance, it has been recently reported that bacteria improve their cooperation on a spatial structure [43] and this could arise because of more imitative 'strategies'. Finally, a promising line of research could be to compare the distribution of values for the MCC parameters that we have obtained here with the observations on single individuals, thus going beyond the check against aggregate data to address the issue of reproducing whole histograms. Unfortunately, the data that we currently have are not good

in terms of individual behaviour, as observations are noisy and statistics are insufficient to assign significant values of the parameters to single participants. In this respect, experiments specifically designed to overcome this difficulty could be a very relevant contribution to further verifying our claims.

Another important suggestion arising from our research is the relevance of theoretical concepts derived within reinforcement learning to the study of games on networks. In this respect, it is very interesting to recall that a theoretical line of work based on reinforcement learning models for two-player repeated games has received quite some attention recently [44,45]. In this context, a generalized equilibrium concept has been introduced in order to explain the findings in simulations of two-player PD [34,35], called self-correcting equilibrium: it obtains when the expected change of parameters is zero but there is a positive probability to incur into a negative as well as positive stimulus. The extension of the reinforcement learning dynamics to multiplayer PD that we have presented here points to the explanatory power of such equilibrium concepts in the framework of network games, as the level of cooperation observed in experiments is in close agreement with the predicted equilibrium. Importantly, it has recently been shown that behavioural rules with intermediate aspiration levels, as the ones we find here to be relevant, are the most successful ones among all possible reactive strategies in a wide range of two-player games [46]. This suggests that this type of evolutionary dynamics may indeed be relevant in general. It would therefore be important to study whether or not the associated equilibrium concept is also the most important one when other types of games are played on an underlying network. If that is the case, we would have a very powerful tool to understand and predict human behaviour in those situations.

References

- Dawes RM. 1980 Social dilemmas. *Annu. Rev. Psychol.* **31**, 169–193. (doi:10.1146/annurev.ps.31.020180.001125)
- Maynard Smith J, Price GR. 1973 The logic of animal conflict. *Nature* **246**, 15–18. (doi:10.1038/246015a0)
- Axelrod R. 1984 *The evolution of cooperation*. New York, NY: Basic Books.
- Hofbauer J, Sigmund K. 1998 *Evolutionary games and population dynamics*. Cambridge, UK: Cambridge University Press.
- Maynard Smith J, Szathmari E. 1995 *The major transitions in evolution*. Oxford, UK: Freeman.
- Hofbauer J, Sigmund K. 2003 Evolutionary game dynamics. *Bull. Am. Math. Soc.* **40**, 479–519. (doi:10.1090/S0273-0979-03-00988-1)
- Szabó G, Fáth G. 2007 Evolutionary games on graphs. *Phys. Rep.* **446**, 97–216. (doi:10.1016/j.physrep.2007.04.004)
- Roca CP, Cuesta J, Sánchez A. 2009 Evolutionary game theory: temporal and spatial effects beyond replicator dynamics. *Phys. Life Rev.* **6**, 208–249. (doi:10.1016/j.plrev.2009.08.001)
- Lomborg B. 1996 Nucleus and shield: the evolution of social structure in the iterated Prisoner's Dilemma. *Am. Sociol. Rev.* **61**, 278–307. (doi:10.2307/2096335)
- van Veelen M, García J, Rand DG, Nowak MA. 2012 Direct reciprocity in structured populations. *Proc. Natl Acad. Sci. USA* **109**, 9929–9934. (doi:10.1073/pnas.1206694109)
- Fischbacher U, Gächter S, Fehr E. 2001 Are people conditionally cooperative? Evidence from a public goods experiment. *Econ. Lett.* **71**, 397–404. (doi:10.1016/S0165-1765(01)00394-9)
- Grujić J, Fosco C, Araujo L, Cuesta JA, Sánchez A. 2010 Social experiments in the mesoscale: humans playing a spatial Prisoner's Dilemma. *PLoS ONE* **5**, e13749. (doi:10.1371/journal.pone.0013749)
- Gracia-Lázaro C, Ferrer A, Ruiz G, Cuesta JA, Sánchez A, Moreno Y. 2012 Heterogeneous networks do not promote cooperation when humans play a Prisoner's Dilemma. *Proc. Natl Acad. Sci. USA* **109**, 10 922–20 926. (doi:10.1073/pnas.1206681109)
- Traulsen A, Semmann D, Sommerfeld RD, Krambeck HJ, Milinski M. 2010 Human strategy updating in evolutionary games. *Proc. Natl Acad. Sci. USA* **107**, 2962–2966. (doi:10.1073/pnas.0912515107)
- Grujić J, Gracia-Lázaro C, Milinski M, Semmann D, Traulsen A, Cuesta JA, Moreno Y, Sánchez A. Submitted. A meta-analysis of spatial Prisoner's Dilemma experiments: conditional cooperation and payoff irrelevance. *EPJ Data Sci.*
- Trivers RL. 1971 The evolution of reciprocal altruism. *Q. Rev. Biol.* **46**, 35–57. (doi:10.1086/406755)
- Sigmund K. 2010 *The calculus of selfishness*. Princeton, NJ: Princeton University Press.
- Axelrod R, Hamilton WD. 1981 The evolution of cooperation. *Science* **211**, 1390–1396. (doi:10.1126/science.7466396)
- Grujić J, Cuesta JA, Sánchez A. 2012 On the coexistence of cooperators, defectors and conditional cooperators in the multiplayer iterated Prisoner's Dilemma. *J. Theor. Biol.* **300**, 299–308. (doi:10.1016/j.jtbi.2012.02.003)
- Grujić J, Eke B, Cabrales A, Cuesta JA, Sánchez A. 2012 Three is a crowd in iterated Prisoner's Dilemmas: experimental evidence on reciprocal behavior. *Sci. Rep.* **2**, 638. (doi:10.1038/srep00638)
- Gutiérrez-Roig M, Gracia-Lázaro C, Moreno Y, Perelló J, Sánchez A. Submitted. From selfish youth to cooperative old age: transition to prosocial behavior in social dilemmas. *Nat. Commun.*
- Dal Bó P, Frechette G. 2011 The evolution of cooperation in infinitely repeated games: experimental evidence. *Am. Econ. Rev.* **101**, 411–429. (doi:10.1257/aer.101.1.411)

4. Material and methods

Agent-based simulations of the model were carried out using the following parameters: $c_0 = 0.5$ (initial fraction of cooperators),³ $R = 1$, $P = 0$, $S = -1/2$, $T = 3/2$ (entries of the PD's pay-off matrix, such that $T > R$, $S < P$ and $2R > T + S$),⁴ $N = 1000$ and $m = 10$ (network parameters).⁵ The MCC behavioural parameters $\{q, p, r\}$ are all drawn for each player before the first round of the game from a uniform distribution $\mathcal{U}[0,1]$, with the additional constraint $p + r \leq 1$ to have $0 \leq p_C(x) \leq 1$. Note that the particular form of the initial distribution as well as the presence of the constraint does not influence the outcome of our experiments.

Funding statement. This work was supported by the Swiss Natural Science Foundation through grant PBFPR2_145872, by Ministerio de Economía y Competitividad (Spain) through grant PRODIEVO, by the ERA-Net on Complexity through grant RESINEE, and by Comunidad de Madrid (Spain) through grant MODELICO-CM.

Endnotes

¹There are indeed more sophisticated approaches, e.g. modelling strategies as evolving automata or bitstrings that can cover a large strategy space [9,10].

²Note that the general MCC behaviour includes full defectors and full cooperators as special cases: if for a given player $q = 0$, then as soon as she defects she will keep defecting until q changes; vice versa, if $r = 1$, then as soon as she cooperates she will keep cooperating until r changes.

³This value is close to the initial level of cooperation observed in the experiments [12,13] and otherwise represents our ignorance about the initial mood of the players.

⁴These values are also close to the ones set in the experiments [12,13].

⁵We checked that our results are robust with respect to system size and link density.

23. Schlag KH. 1998 Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *J. Econ. Theory* **78**, 130–156. (doi:10.1006/jeth.1997.2347)
24. Helbing D. 1992 Interrelations between stochastic equations for systems with pair interactions. *Physica A* **181**, 29–52. (doi:10.1016/0378-4371(92)90195-V)
25. Szabó G, Toke C. 1998 Evolutionary Prisoner's Dilemma game on a square lattice. *Phys. Rev. E* **58**, 69–73. (doi:10.1103/PhysRevE.58.69)
26. Blume LE. 1993 The statistical mechanics of strategic interaction. *Games Econ. Behav.* **5**, 387–424. (doi:10.1006/game.1993.1023)
27. Traulsen A, Pacheco JM, Imhof LA. 2006 Stochasticity and evolutionary stability. *Phys. Rev. E* **74**, 021905. (doi:10.1103/PhysRevE.74.021905)
28. Moran PAP. 1962 *The statistical processes of evolutionary theory*. Oxford, UK: Clarendon Press.
29. Nowak MA, May RM. 1992 Evolutionary games and spatial chaos. *Nature* **359**, 826–829. (doi:10.1038/359826a0)
30. Holley R, Liggett TM. 1975 Ergodic theorems for weakly interacting infinite systems and the voter model. *Ann. Probab.* **3**, 643–663. (doi:10.1214/aop/1176996306)
31. Fehr E, Gächter S. 2000 Fairness and retaliation: the economics of reciprocity. *J. Econ. Perspect.* **14**, 159–181. (doi:10.1257/jep.14.3.159)
32. Matsui A. 1992 Best response dynamics and socially stable strategies. *J. Econ. Theory* **57**, 343–362. (doi:10.1016/0022-0531(92)90040-0)
33. Bush R, Mosteller F. 1955 *Stochastic models of learning*. New York, NY: John Wiley and Son.
34. Macy MW, Flache A. 2002 Learning dynamics in social dilemmas. *Proc. Natl Acad. Sci. USA* **99**, 7229–7236. (doi:10.1073/pnas.092080099)
35. Izquierdo SS, Izquierdo LR, Gotts NM. 2008 Reinforcement learning dynamics in social dilemmas. *J. Artif. Soc. Soc. Simulat.* **11**, 1.
36. Barabási AL, Albert R. 1999 Emergence of scaling in random networks. *Science* **286**, 509–512. (doi:10.1126/science.286.5439.509)
37. Nowak MA. 2006 Five rules for the evolution of cooperation. *Science* **314**, 1560–1563. (doi:10.1126/science.1133755)
38. Gómez-Gardeñes J, Campillo M, Floría LM, Moreno Y. 2007 Dynamical organization of cooperation in complex networks. *Phys. Rev. Lett.* **98**, 034101. (doi:10.1103/PhysRevLett.98.034101)
39. Camerer C, Ho TH. 1999 Experience-weighted attraction learning in normal form games. *Econometrica* **67**, 827–874. (doi:10.1111/1468-0262.00054)
40. Wu B, Zhou D, Fu F, Luo Q, Wang L, Traulsen A. 2010 Evolution of cooperation on stochastic dynamical networks. *PLoS ONE* **5**, e11187. (doi:10.1371/journal.pone.0011187)
41. FehI K, van der Post DJ, Semmann D. 2011 Co-evolution of behaviour and social network structure promotes human cooperation. *Ecol. Lett.* **14**, 546–551. (doi:10.1111/j.1461-0248.2011.01615.x)
42. Van Segbroeck S, Santos FC, Lenaerts T, Pacheco JM. 2011 Selection pressure transforms the nature of social dilemmas in adaptive networks. *N. J. Phys.* **13**, 013007. (doi:10.1088/1367-2630/13/1/013007)
43. Hol FJH, Galajda P, Nagy K, Woolthuis RG, Dekker C, Keymer JE. 2013 Spatial structure facilitates cooperation in a social dilemma: empirical evidence from a bacterial community. *PLoS ONE* **8**, e77402. (doi:10.1371/journal.pone.0077402)
44. Erev I, Roth AE. 2001 Simple reinforcement learning models and reciprocity in the Prisoner's Dilemma Game. In *Bounded rationality: the adaptive toolbox*. (eds G Gigerenzer, R Selten), pp. 215–231. Cambridge, MA: MIT Press.
45. Bendor J, Mookherjee D, Ray D. 2001 Reinforcement learning in repeated interaction games. *Adv. Theor. Econ.* **1**, 3. (doi:10.2202/1534-5963.1008)
46. Martínez Vaquero LA, Cuesta JA, Sánchez A. 2012 Generosity pays in the presence of direct reciprocity: a comprehensive study of 2×2 repeated games. *PLoS ONE* **7**, e35135. (doi:10.1371/journal.pone.0035135)