

Evolution of Commitment and Level of Participation in Public Goods Games

The Anh Han · Luís Moniz Pereira ·
Tom Lenaerts

Received: date / Accepted: date

Abstract Before engaging in a group venture agents may require commitments from other members in the group, and based on the level of acceptance (participation) they can then decide whether it is worthwhile joining the group effort. Here, we show in the context of Public Goods Games and using stochastic evolutionary game theory modelling, which implies imitation and mutation dynamics, that arranging prior commitments while imposing a minimal participation when interacting in groups induces agents to behave cooperatively. Our analytical and numerical results show that if the cost of arranging the commitment is sufficiently small compared to the cost of cooperation, commitment arranging behavior is frequent, leading to a high level of cooperation in the population. Moreover, an optimal participation level emerges depending both on the dilemma at stake and on the cost of arranging the commitment. Namely, the harsher the common good dilemma is, and the costlier it becomes to arrange the commitment, the more participants should explicitly commit to the agreement to ensure the success of the joint venture. Furthermore, considering that commitment deals may last for more than one encounter, we show that commitment proposers can be lenient in case of short-term agreements, yet should be strict in case of long-term interactions.

T. A. Han

School of Computing, Teesside University, Borough Road, Middlesbrough, TS1 3BA, UK
E-mail: T.Han@tees.ac.uk

L. M. Pereira

NOVA Laboratory for Computer Science and Informatics (NOVA LINCS), Faculdade de Ciências e Tecnologia, Universidade Nova de Lisboa, 2829-516 Caparica, Portugal
E-mail: lmp@fct.unl.pt

Tom Lenaerts

MLG, Département d'Informatique, Université Libre de Bruxelles, Boulevard du Triomphe CP212, 1050 Brussels, Belgium
AI lab, Computer Science Department, Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussels, Belgium
E-mail: Tom.Lenaerts@ulb.ac.be

Keywords Commitment · Cooperation · Participation · Evolutionary Game Theory · Public Goods Games

1 Introduction

Before engaging in a group venture individuals often secure prior commitments from other members of the group, and based on the level of participation (i.e. how many group members commit) they can then decide whether it is worthwhile joining the group effort (Nesse, 2001, Sterelny, 2012, Barrett, 2003). Many group ventures can be launched only when the majority of the participants commit to contribute to a common good (Cherry and McEvoy, 2013, Chen and Komorita, 1994). A cooperative hunting effort (both in animals, such as lions and some birds, and in humans) usually requires a sufficient number of participants “on board” to embark (Stander, 1992, Alvard and Nolin, 2002). While some international agreements require ratification by all parties before entering into force, most (especially global treaties) require a minimum less than the total number of negotiating countries (Barrett, 2003, Cherry and McEvoy, 2013). In group or coalition formation in multi-agent systems, a sufficient number of participants needs to agree on the terms of the agreement for it to be binding (Ray, 2007, Hasan and Raja, 2013). In general, it appears that the required participation level depends on the nature of the problem in place. Here we investigate analytically and numerically whether commitment strategies, in which players propose, initiate and honor a deal, evolve as viable strategies for the evolution of cooperative behavior in the Public Goods Game (PGG), while at the same time analyzing the effect of the participation level and the transition from a single to multiple-rounds version of the game.

In a typical PGG, all players can choose whether to cooperate, contributing an amount c to the public good, or to defect, taking advantage of the public good without contributing to it (Hauert et al., 2007, Sigmund, 2010). The total contribution is multiplied by a constant factor, $r > 1$, and is distributed equally among all players. A group of cooperators hence do better than a group of defectors, but defectors always have a better payoff in their group. Only when r is smaller than the group size (denoted by N), does the PGG represent a social dilemma, i.e. every individual player is better off defecting than cooperating, no matter what the other players do, although in this case it is the worst possible outcome for the group (Hauert et al., 2002). Evolutionary game theory (EGT) models (Hofbauer and Sigmund, 1998, Sigmund, 2010) predict the destruction of cooperation – famously known as ‘the tragedy of commons’ (Hardin, 1968).

In our commitment extension to the PGG, agents have, before playing the PGG, the option to propose other members in the group to commit to contribute, where the proposers pay a personal cost ϵ , to make it credible. If a sufficient number of the members commit (participation level F), the PGG is played. Otherwise, the commitment proposers refuse to play. Those

who committed but then do not contribute have to compensate others at a personal cost, δ . Further details of the game are provided in Section 3.

In the two-player setting (namely, the Prisoner’s Dilemma) the behavior of proposing prior commitments has been shown to promote the evolution of cooperation if the cost of arranging commitment ϵ is sufficiently small compared to the cost of cooperation (Han et al., 2013a, 2015b,c). But when larger groups of actors are involved, decision-making becomes much more complex (Van Segbroeck et al., 2012, Han et al., 2012c, Gokhale and Traulsen, 2010, Duong and Han, 2015, 2016). Instead of a clear, full commitment or, no-commitment, from the co-player as in the two-player game, when moving to the multi-player decision scenario of the PGG, there can be several possible intermediate degrees of participation (as many as the size of the group, i.e. N). It is not clear which minimal participation level would evolve in a the population given the settings of the PGG.

To answer this question, we will distinguish N different participation levels for the one-shot PGG, encoded in terms of commitment-proposing strategies, $COMP_F$ where $F \in \{1, \dots, N\}$. $COMP_F$ contributes c to the public good when there are at least F players in the group (including herself) that agree or commit to contribute; otherwise, the strategy refuses to play. Examples for such a minimum membership requirement can be found in the creation of treaties that address international environmental issues (Barrett, 2003, Cherry and McEvoy, 2013)¹ or the formation of coalitions in multi-agent systems (Shehory and Kraus, 1998, Shehory et al., 1998). These new strategies allow us to investigate how the severity of the game (defined by $r < N$, where lower r values correspond to a tougher PGG) and the parameters of the commitment system (ϵ and δ) influence the required participation level. Second, we examine how strict, in case the PGG is repeated for multiple rounds R , these $COMP_F$ players should be when they notice that among those that committed to contribute, some of them did not honor the deal: should they immediately claim the compensation or might it be worthwhile to be lenient and continue the game? In that case how lenient should an agent be? Again we determine here how the three parameters, r , ϵ and δ , affect the answers to these questions.

The remainder of this article is structured as follows. In Section 2 we highlight the research, both analytical and experimental, that are most closely related to our work. In Section 3 the game structure is discussed in more detail and the Evolutionary Game Theory methods used to obtain the results are explained. The results section (Section 4) is divided into three parts. First, a mathematical analysis is performed to determine under which conditions the $COMP_F$ strategies are evolutionarily viable. Second, results are provided that

¹ For example, the Kyoto Protocol to the United Nations Framework Convention on Climate Change required ratification by at least 55 parties prior to its entry into force in February 2005. Similarly, the Montreal Protocol on Substances that Deplete the Ozone Layer required at least 11 countries to ratify it before it entered into force in 1989. While many treaties require only a subset of affected parties to ratify for entry into force, some require all parties to join. For example, the Convention for the Protection of the Marine Environment of the North-East Atlantic required accession of all negotiating parties (Barrett, 2003).

show what the required participation level in an agreement should be for cooperation to thrive. Third, the results for multiple-round PGGs are presented. Finally, a discussion of all the presented results and conclusions are provided in Section 5.

2 Related Work

The problem of explaining the emergence of collective behavior has been studied extensively from a wide range of research fields, including Anthropology, Sociology, Economics, Evolutionary Biology, Psychology, and, more recently, Artificial Intelligence (AI) and Multi-agent Systems (MAS) (Axelrod and Hamilton, 1981, West et al., 2007, Nowak, 2006b, Sigmund, 2010, Tuyls and Parsons, 2007, Hofmann et al., 2011, Han et al., 2011, 2012b, Lerat et al., 2013, Han et al., 2013b, Hasan and Raja, 2013, Ranjbar-Sahraei et al., 2014, Airiau et al., 2014, Sasaki et al., 2015, Han, 2016, Han and Lenaerts, 2016). The Public Goods Game (PGG) is a standard framework to study this problem, as it captures the tension between the benefit of mutual cooperation and the temptation to exploit the efforts of others in a joint venture (Ostrom, 1990, Fehr and Gächter, 2000, Santos et al., 2008, Hauert et al., 2007, Sigmund et al., 2010, Van Segbroeck et al., 2012). Several mechanisms have been proposed that promote the evolution of cooperation within the context of the PGG (Nowak, 2006b, West et al., 2007, Sigmund, 2010), including (but not limited to) kin and group selection (Traulsen and Nowak, 2006, West et al., 2007), reputation and repeated interactions, networked reciprocity (Szolnoki and Perc, 2010, Santos et al., 2008, Perc et al., 2013), and punishment/reward (Szolnoki and Perc, 2013, 2012, Sigmund et al., 2001, Hauert et al., 2007, Szolnoki and Perc, 2010). In this work, we study whether the behavior of proposing prior agreements regarding posterior compensations can resolve the tension observed within the PGG, without taking into account relatedness, reputation, structured populations or repeated interaction effects. The focus on strategies capable of creating agreements makes this work more closely linked with evolution of minimal cognitive capabilities. We studied such strategies before in the context of pair-wise interactions (Han et al., 2012b, 2013a, 2015c, Han, 2016). Yet, when moving from pair-wise to group interactions, the outcome is more complex since there are more possible participation levels commitment proposers can insist on. These levels play a crucial role regarding the effectiveness of the mechanism, which cannot be seen in the pair-wise interaction setting.

We recently performed a first analysis of commitments in a group interaction setting (Han et al., 2015b), analyzing commitment strategies that follow two different approaches when facing non-committers, i.e. those who do not agree to contribute. Prior work showed that these defectors can be persistent at higher initiation costs. Because PGG is by definition non-exclusive, it may be costly to prevent the non-committing players from enjoying the public goods without external measures (Maier-Rigaud et al., 2010). The article com-

pared two different strategies: (i) AVOID, which specifies that the player does not participate in the creation of the common good whenever there are non-committers; (ii) RESTRICT, which has the capacity to impose boundaries (at a cost) on the common good so that only those that have committed to make it work have (better) access or that the benefit non-contributors can acquire is reduced. The analysis reveals that RESTRICT, rather than AVOID, leads to more favorable societal outcomes, in terms of contribution levels, especially when the group size and/or the benefit of the PGG increase. The AVOID strategy is equivalent to one of the strategies we investigate here: $COMP_F$ with the participation level F being equal to the size of the group ($F = N$). The present work investigates participation levels smaller than the group size, focussing on how varying F affects the cooperation in the population and which level is evolutionarily preferred. The current work complements the prior publication, describing a novel, alternative approach to cooperation enhancement in the PGG, as costly restriction measures may not always be possible and they may take additional effort and time to implement (Ostrom, 1990). Furthermore, the newly defined participation level factor allows us to extend the definition of commitment from the one-shot PGG to its iterated version. That enables us, for the first time, to show how agents would behave given the length of a commitment deal; for instance, should they be more or less strict in terms of the participation level from the co-players?

The research addressed in this article is directly related to the results produced by Van Segbroeck et al. (Van Segbroeck et al., 2012) where a N-person direct reciprocity mechanism was analyzed in the context of a repeated PGG. The authors considered strategies that cooperate only if the number of group members that cooperated in the previous round reaches a certain threshold. Although these threshold strategies resemble the conditional commitment proposers as defined in our model, they are different in the same manner that our work differs from punishment or reward strategies traditionally studied in the context of the evolution of cooperation: The N-player direct reciprocity strategies are reactive in the sense that they respond immediately to the previous behavior of the opponents (Sigmund et al., 2001, Chen et al., 2015). Differently, our commitment strategies with participation conditions strategically decide on how to behave before the actual game is played. As a consequence, these strategies can be investigated even in the context of the extended one-shot PGG.

Several behavioral economic experiments on commitments in PGGs have been performed, and our results are in close accordance with the outcome of these experiments (Chen and Komorita, 1994, Cherry and McEvoy, 2013). For example, high levels of cooperation were observed in a PGG experiment where a binding agreement, which was enabled through a prior communication stage among the members of the group, could be arranged before the PGG interaction occurred (Chen and Komorita, 1994). The experiment also showed that whenever a commitment deal is not binding or not enforced, corresponding to a low compensation cost in our commitment model, defectors are widespread and the contribution level is low. The main difference with this experimental

setting is that players, once they agreed to commit, still have the possibility to defect (even when they then have to pay a compensation); while in the experiment, any participant that agrees to contribute is enforced to do so. Hence, the experiments do not reveal the effect of varying the compensation cost, which as we also show here, is a major factor in a commitment system. Our work reveals that, whenever the compensation reaches a certain threshold, increasing it does not lead to further improvement (in terms of cooperation levels). It implies that, when designing laws (whether in real life or in a self-organizing MAS), it is not necessary to have an infinitely large compensation or sanction against law breakers; a sufficient, predefined, one is enough for a wide range of situations.

Another commitment experiment (Cherry and McEvoy, 2013) takes the form of a deposit and refund scheme (Sasaki et al., 2015): In this scheme, players that agree to commit have to deposit an amount which will be refunded only if they honor the commitment and contribute to the common good. The main difference with our analysis is that the agreement is set up exogenously by a third party instead of being implemented as a strategic behavior (i.e. it is not an option to propose a commitment in the experiment). Nonetheless, the outcome of the experiment revealed that whenever the deposit amount, corresponding to the compensation cost in our model, is sufficiently high, the contribution level is significant (Cherry and McEvoy, 2013). But again, they considered only one value of the compensation cost, preventing us to compare the effects of varying this essential feature. Furthermore, it is worth noting that in both these experiments (Chen and Komorita, 1994, Cherry and McEvoy, 2013), the cost of setting up the commitment is always set to 0, thereby leading to effortless and effective commitment strategies. But as we will show, this cost is the decisive factor for the viability of commitment strategies as well as the emerging participation level. In short, the current paper suggests that further experiments are required to explore the effects of varying the essential parameters driving commitments in group interactions, including the costs of arranging commitment and of compensation.

Last but not least, there is a large literature on commitments in AI and MAS (Schelling, 1990, Castelfranchi and Falcone, 1998, Wooldridge and Jennings, 1999, Harrenstein et al., 2007, Castelfranchi and Falcone, 2010, Winikoff, 2007). These works focus on how to use commitment for regulating individual and collective behaviors. They study how to formalize commitments and their different aspects, such as norms and conventions, in a MAS. Our work hence provides insights into the design of such self-organizing MAS when dealing with group interactions. For instance, what are the appropriate actual degrees of commitment one should require from group members leading to highest levels of cooperation. In a similar manner, our research may also have important implications for the work on group and coalition formation (Shehory and Kraus, 1998, Shehory et al., 1998), where a decision on such a formation naturally depends on the number of agents agreeing with the terms of the formation (Ray, 2007, Sugawara, 2011, Hasan and Raja, 2013). In turn, as these AI and MAS works focus more on the formalization of more complex decision

making aspects, such as in arranging an agreement, they provide us with more sophisticated extensions to our current commitment-organising mechanisms. For instance, in (Hasan and Raja, 2013) the authors formalize commitment as a device to incentivise other players in a structured population to form coalitions within which they can enjoy mutual cooperation. Agents play an iterated Prisoner’s Dilemma game with their neighbors and offer commitments to their wealthiest neighbors in order to form coalitions. The commitment mechanism is implemented similar to our model (Han et al., 2012b, 2013a). The authors analyze the conditions regarding network structure and payoff configurations under which an optimal coalition is achieved, with highest levels of cooperation.

In (Schelling, 1990, Harrenstein et al., 2007), different forms of complex commitments are described, such as commitments conditional on other players’ commitments and actions. The players do not act or commit simultaneously and it is strategically important to take into account the order of making a commitment. In the current PGG framework, commitments are made simultaneously and agents are not aware of each others’ commitments (at least the non-proposing ones). It might be interesting to relax this requirement in future work. More complex approaches to designing a commitment or agreement deal are also extensively studied in coalition formation (Ray, 2007, Shehory and Kraus, 1998, Shehory et al., 1998), especially when the terms of an agreement must be negotiated by the coalition members before being enacted.

3 Models and Methods

3.1 Commitment in Public Goods Games

As already discussed in the introduction we examine here the evolution of commitment strategies in the context of the PGG. A commitment strategy is defined as follows (see also Figure 1): Before playing the PGG, the agent or player proposes other members in the group to commit to contribute c , paying a personal cost ϵ , corresponding to the cost of establishing the agreement. When a sufficient number of the participants commit, the PGG is played. Otherwise, a commitment proposer refuses to play the PGG. Those who committed but then do not contribute, i.e. they defect, have to compensate others at a personal cost, δ . Yet, how many participants have to commit so that the PGG is played? And how will the parameterisation of the extended PGG influence the level of participation requested by the commitment proposers?

With a fixed group size N , there are N possible participation levels (i.e. how many players in the group agree to contribute, including the proposer), where we identify this participation by $F \in \{1, \dots, N\}$. We hence denote by $COMP_F$ the corresponding commitment proposing strategy that contributes c if there are at least F players (including herself) in the group that agree to contribute; otherwise, the strategy refuses to play, resulting in a zero payoff for all group members. Note also that the results remain equivalent when the game

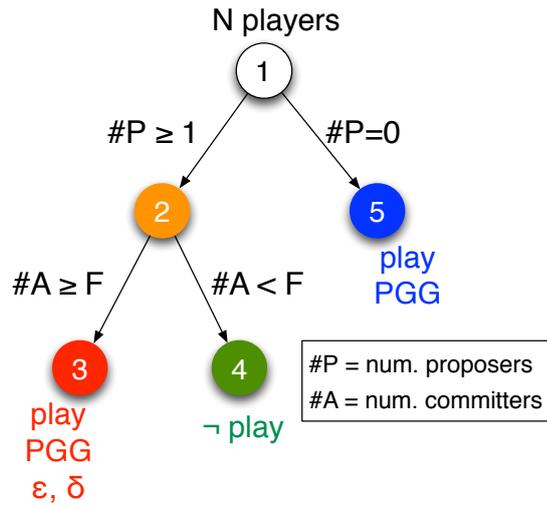


Fig. 1 Structure of the PGG with commitments. When there are no proposers among the N players, the original PGG is played (node 5). When there are proposers $COMP_F$ (node 2), there need to be a sufficient number of acceptors (including the proposers themselves) (node 3), otherwise the game is not played and everyone obtains 0 payoff (node 4). In the former case, the PGG is played but proposers pay the setup cost ϵ and those that defect in the PGG (whether they are proposers or acceptors) will have to pay a compensation δ .

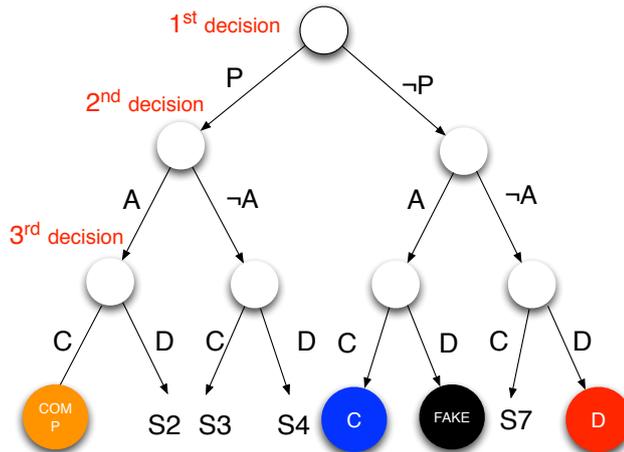


Fig. 2 Strategy taxonomy. There are three different decisions: 1) before the PGG, whether to propose a commitment deal or not; 2) upon receiving a proposal, whether to accept or reject it; and finally, 3) whether to cooperate or defect in the PGG itself. That leads to eight possible strategies, where four of them can be excluded due to being dominated. Hence, four strategies that are unconditional on the presence of a commitments are named: COMP, C, D, FAKE.

is still played and $COMP_F$ defects whenever there are not enough committers, as that would result in the same outcome of zero payoff for everyone in the group.

The space of possible strategies, which do not condition on the presence of a commitment, is determined by three decisions: before the PGG, whether to propose a commitment deal or not; upon receiving a proposal, whether to accept or reject it; and finally, whether to cooperate or defect in the PGG itself. As shown in previous models (Han et al., 2013a, 2015b), some unreasonable strategies can be excluded as they will get eliminated anyway. Namely, the strategies that cooperate in the game but will reject commitment proposals by others (i.e. S3 and S7 strategies in Figure 2) and the strategies that propose an agreement but then defect in the game (i.e. S2 and S4 strategies in Figure 2), independent of whether they accept another player's proposal or not. The former ones are willing to cooperate, and thus should also agree to commit when being asked to since they prefer to cooperate anyway. Additionally, a positive compensation is guaranteed without having to pay the cost of arrangement. The latter ones should not propose commitments while intending to defect, as they would lose the cost of arranging the commitment and moreover have to compensate the other players. Figure 2 summarizes the taxonomy of the possible strategies.

Besides the unconditional strategies, we also include in our model a strategy that behaves conditionally on the presence of a commitment deal (see the FREE strategy below), which appeared to be the main obstruction for commitment strategies to evolve (Han et al., 2013a, 2015b).

Summarizing, we will examine the evolutionary dynamics of five strategies:

- i) The $COMP_F$ strategies with different levels F of participation, with $F \in \{1, \dots, N\}$;
- ii) Traditional unconditional contributors (annotated by C), who always commit when being proposed a commitment deal, contribute whenever the PGG is played, but do not propose commitment;
- iii) Unconditional non-contributors (annotated by D), who do not accept commitment, defect when the PGG is played, and do not propose commitment;
- iv) Fake committers (annotated by FAKE), who accept a commitment proposal yet do not subsequently contribute whenever the PGG is actually played. These players assume that they can exploit the commitment proposing players without suffering the consequences; and
- v) Commitment free-riders (annotated by FREE), who defect unless being proposed a commitment, which they then accept and cooperate subsequently in the PGG. In other words, these players are willing to contribute when a commitment is proposed but are not prepared to pay the cost of setting it up.

The strategies are randomly sampled from a well-mixed, finite population of a constant size Z , which consists of players adhering to one of the five commitment proposing strategies $COMP_F$ ($1 \leq F \leq 5$), C, D, FREE or the

FAKE strategy (i.e. nine strategies in total). In each interaction, N agents are randomly selected from the population for playing the PGG.

Focal strategy i	Opponent strategy j	Payoff Π_{ij}	Payoff constraints
$COMP_{F_1}$	$COMP_{F_2}$	$rc - c - \frac{\epsilon}{N}$	$\forall F_1, F_2 \in \{1, 2, \dots, N\}$ with $F_1 \neq F_2$
$COMP_F$	C $FREE$	$rc - c - \frac{\epsilon}{k}$	$\forall F \in \{1, 2, \dots, N\}$ and k focal strategies with $1 \leq k \leq N$
$COMP_F$	D	$\frac{krc}{N} - c - \frac{\epsilon}{k}$ 0	$\forall F \in \{1, 2, \dots, N\}$ and k focal strategies with $N \geq k \geq F$ otherwise (i.e. $F > k$)
$COMP_F$	$FAKE$	$\frac{krc}{N} - c - \frac{\epsilon}{k}$ $+$ $\frac{(N-k)\delta}{k}$	$\forall F \in \{1, 2, \dots, N\}$ and k focal strategies with $1 \leq k \leq N$
C	$COMP_F$ C	$rc - c$	$\forall F \in \{1, 2, \dots, N\}$ and k focal strategies with $1 \leq k \leq N$
C	D $FAKE$ $FREE$	$\left(\frac{rk}{N} - 1\right)c$	k focal strategies with $1 \leq k \leq N$
$FREE$	$COMP_F$	$rc - c$ 0	$\forall F \in \{1, 2, \dots, N\}$ and k focal strategies with $1 \leq k < N$ otherwise (i.e. $k = N$)
D	$COMP_F$	$\frac{(N-k)rc}{N}$ 0	$\forall F \in \{1, 2, \dots, N\}$ and k focal strategies with $N - k \geq F$ otherwise (i.e. $N - k < F$)
$FAKE$	$COMP_F$	$\frac{(N-k)rc}{N} - \delta$ 0	$\forall F \in \{1, 2, \dots, N\}$ and $\forall k$ with $1 \leq k < N$ otherwise (i.e. $k = N$)
$D,$ $FAKE$ or $FREE$	C	$\frac{rk}{N}c$ 0	k focal strategies with $1 \leq k < N$ otherwise (i.e. $k = N$)
$D,$ $FAKE$ or $FREE$	$D,$ $FAKE$ or $FREE$	0	

Table 1 Specification of the strategy payoffs. F_i or F refers to the number of other players to commit to the PGG, including the strategy itself. N is the group size and k ($N - k$) is the number of focal (opponent) strategy players in the group. The parameter F plays an important role in groups where the non-committing players D are present as the commitment player will only accept to play when at least F in the group agree to contribute, namely, an amount c ; r is the PGG multiplication factor, typically $r < N$; ϵ and δ are respectively the cost of setting up the commitment and the compensation that needs to be paid when a player dishonors the commitment.

Table 1 lists the payoffs each strategy receives when encountering specific other strategies. We denote by $\Pi_{ij}(k)$ and $\Pi_{ji}(k)$ the payoffs of a strategist of

type i and j , respectively, when the random sampling consists of k players of type i and $(N - k)$ players of type j . Next to providing the well-known payoffs for the typical PGG strategies the table specifies also for every $COMP_F$ strategy as well as the *FREE* and *FAKE* strategies how their payoff is determined. The first column captures the first strategy and the second the list of opponent strategies. These payoffs are used in the Evolutionary Game Theory (EGT) methods, which are discussed in detail in the next section.

3.2 Evolutionary dynamics in finite populations

Both the analytical and numerical results obtained here use Evolutionary Game Theory methods for finite as opposed to infinite well-mixed populations (Nowak et al., 2004, Imhof et al., 2005, Sigmund, 2010). Differently from the infinitely large population settings in which the deterministic replicator dynamics equation is adopted to model the population dynamics (Hofbauer and Sigmund, 1998), the finite population approach uses a Moran process (Nowak, 2006a), a type of birth-death process that incorporates random replacement of one individual by the offspring of another individual (selection) with the offspring having the possibility to adopt the behavior of the parent or to pick a random alternative (mutation). The first mechanism eliminates strategies either randomly or in a fitness proportionate manner (exploitation) and the second mechanism provides the means for novel strategies to appear in the population (exploration) (Traulsen et al., 2009). Different from the dynamics in infinitely large populations, which can be characterized by stable equilibrium points where multiple strategies may co-exist, the stochasticity (i.e. the random replacement) in a finite population will, given enough time, drive the population to a homogeneous state, wherein all agents use the same strategy, as long as the state is reachable (Nowak, 2006a). This homogenous state might be escaped only via mutation, as once all agents adopt the same strategy no change to the population strategic composition can be obtained otherwise. Nonetheless, any state is reachable from any other state, since there is always a non-zero probability that an agent imitates another agent's strategy² if the influence of the game payoff on the fitness is not infinitely large (see the Fermi function below). In general, the finite population dynamics defined by the Moran process allow one to characterise the probabilities of moving between the homogeneous population states and how often each homogeneous state is visited in the long run (see details of the Markov process below). For a nice illustration of the difference between the two approaches see some examples in (Imhof et al., 2005) where the approaches were used to analyze strategies in the iterated Prisoner's Dilemma. Finite population dynamics have received increasing attention over the years as it has been shown to be successful in explaining different realistic observations in the study of the evolution of cooperation (Nowak et al., 2004, Imhof et al., 2005, Hauert et al., 2007, Traulsen

² This imitation process is equivalent to the selection in the birth-death process described above.

and Hauert, 2009) as well as having the capacity to elegantly reproduce results from behavioral experiments (Rand et al., 2013, Zisis et al., 2015).

In the current setting, an agent’s payoff represents its *fitness* or social *success*, and the evolutionary dynamics are driven by social learning (Hofbauer and Sigmund, 1998, Sigmund, 2010, Rendell et al., 2010), whereby the more successful agents tend to be imitated more often by the others. In the current work, imitation is modeled through the widely used pair-wise comparison rule (Traulsen et al., 2006), which assumes that an agent A with fitness f_A adopts the strategy of another agent B with fitness f_B with probability given by the Fermi function ³,

$$\left(1 + e^{-\beta(f_B - f_A)}\right)^{-1}.$$

The parameter β represents the ‘imitation strength’ or ‘intensity of selection’, i.e., how strongly the agents base their decision to imitate on the fitness difference ($f_B - f_A$). For $\beta = 0$, we obtain the limit of neutral drift – the imitation decision is random. For large β , imitation becomes increasingly deterministic. It is noteworthy, especially for those who are familiar with other learning literature, that this parameter plays a similar role as the temperature factor in Boltzmann exploration mechanism usually used in Reinforcement Learning to balance between exploitation and exploration (Bloembergen et al., 2015). Indeed, as exploration is introduced below, β balances between greedily mimicking more successful interaction partners and randomly switching to the alternatives available in the population.

As this imitation dynamics only work with the strategies available in the population, novel behaviors can only be explored through mutation (Traulsen et al., 2009). In the evolutionary process, one assumes that, with a certain mutation probability μ , agents switch randomly to one of the potential strategies instead of imitating another agent (which now occurs with the probability $1 - \mu$). Mutation hence seeds the population with new strategies, providing them the possibility to invade (when they are more successful than the resident strategy). As such it provides an additional way of exploring the strategy space.

Given the previous components, the population dynamics can be described by a Markov chain, for which the stationary distribution characterizes the average time the population spends in each of monomorphic end state consisting of only one strategic type. For arbitrary mutation probabilities, this stationary distribution is cumbersome to compute analytically as one has to deal with multiple new mutants at the same time (Hauert et al., 2007, Traulsen et al., 2009). In the limit of small mutation rates (Fudenberg and Imhof, 2005, Imhof et al., 2005, Hauert et al., 2007), any newly occurring mutant in a homogeneous population will fixate or become extinct long before the occurrence of the next mutation. Hence, the evolutionary dynamics will proceed with, at most, two strategies in the population, allowing one to describe the behavioral dynamics

³ See (Traulsen and Hauert, 2009) for some alternative approaches to modelling social dynamics in finite population settings.

by a Markov chain, in which each state corresponds to a monomorphic population, whereas the transition probabilities are given by the fixation probability of a single mutant (Fudenberg and Imhof, 2005, Imhof et al., 2005). This approach has been proven useful for explaining different realistic observations in the study of the evolution of cooperation (Nowak et al., 2004, Hauert et al., 2007, Han et al., 2012b), and for generating results close to real experimental data (Rand et al., 2013, Zisis et al., 2015). Below we describe how the fixation probabilities and the stationary distribution are determined analytically.

Payoffs over group samplings

In finite populations, the groups engaging in a PGG are given by multivariate hypergeometric sampling. For transition between two pure states (small mutation), this reduces to sampling (without replacement) from a hypergeometric distribution (Hauert et al., 2007, Sigmund, 2010). Namely, in a population of size Z with x individuals of type i and $Z - x$ individuals of type j , the probability to select k individuals of type i and $N - k$ individuals of type j in N trials is (Hauert et al., 2007)

$$H(k, N, x, Z) = \frac{\binom{x}{k} \binom{Z-x}{N-k}}{\binom{Z}{N}}.$$

Recall that $\Pi_{ij}(k)$ and $\Pi_{ji}(k)$ denote the payoff of a strategist of type i and j , respectively, when the random sampling consists of k players of type i and $N - k$ players of type j (as derived above). Hence, in a population of x i -strategists and $(Z - x)$ j -strategists, the average payoffs to i and j strategists are (Hauert et al., 2007, Sigmund, 2010):

$$\begin{aligned} P_{ij}(x) &= \sum_{k=0}^{N-1} H(k, N-1, x-1, Z-1) \Pi_{ij}(k+1) \\ &= \sum_{k=0}^{N-1} \frac{\binom{x-1}{k} \binom{Z-x}{N-1-k}}{\binom{Z-1}{N-1}} \Pi_{ij}(k+1), \\ P_{ji}(x) &= \sum_{k=0}^{N-1} H(k, N-1, x, Z-1) \Pi_{ji}(k) \\ &= \sum_{k=0}^{N-1} \frac{\binom{x}{k} \binom{Z-1-x}{N-1-k}}{\binom{Z-1}{N-1}} \Pi_{ji}(k). \end{aligned} \tag{1}$$

Now, the probability to change the number k of agents using strategy i by ± 1 in each time step can be written as

$$T^\pm(k) = \frac{Z-k}{Z} \frac{k}{Z} \left[1 + e^{\mp\beta[P_{ij}(k) - P_{ji}(k)]} \right]^{-1}, \quad (2)$$

with T^+ corresponding to an increase from k to $k+1$ and T^- corresponding to the opposite. The fixation probability of a single mutant with a strategy i in a population of $(N-1)$ agents using j is given by (Traulsen et al., 2006, Karlin and Taylor, 1975, Fudenberg and Imhof, 2005, Imhof et al., 2005)

$$\rho_{j,i} = \left(1 + \sum_{i=1}^{Z-1} \prod_{j=1}^i \frac{T^-(j)}{T^+(j)} \right)^{-1}. \quad (3)$$

In the limit of neutral drift (i.e. $\beta = 0$), $\rho_{B,A}$ equals the inverse of population size, $1/Z$.

Considering a set $\{1, \dots, q\}$ of different strategies, these fixation probabilities determine a transition matrix $M = \{T_{ij}\}_{i,j=1}^q$, with $T_{ij, j \neq i} = \rho_{ji}/(q-1)$ and $T_{ii} = 1 - \sum_{j=1, j \neq i}^q T_{ij}$, of a Markov chain. The normalized eigenvector associated with the eigenvalue 1 of the transposed of M provides the stationary distribution described above (Karlin and Taylor, 1975, Fudenberg and Imhof, 2005, Imhof et al., 2005), describing the relative time the population spends in a configuration with only one of the strategies.

Risk-dominance condition

An important analytical criteria to determine the evolutionary viability of a given strategy is whether it is risk-dominant with respect to other strategies (Nowak, 2006b, Gokhale and Traulsen, 2010). Namely, one considers which selection direction is more probable: an i mutant fixating in a homogeneous population of agents playing j or a j mutant fixating in a homogeneous population of agents playing i . When the first is more likely than the latter, i is said to be *risk-dominant* against j (Gokhale and Traulsen, 2010), which holds for any intensity of selection and in the limit of large population size Z when

$$\sum_{k=1}^N \Pi_{ij}(k) \geq \sum_{k=0}^{N-1} \Pi_{ji}(k). \quad (4)$$

4 Results

4.1 Constraints on the evolutionary viability of $COMP_F$

In order to assess whether and when the commitment proposing behavior can be a viable strategy for the evolution of cooperation we first derive conditions for which $COMP_F$ are risk-dominant against other strategies in the population

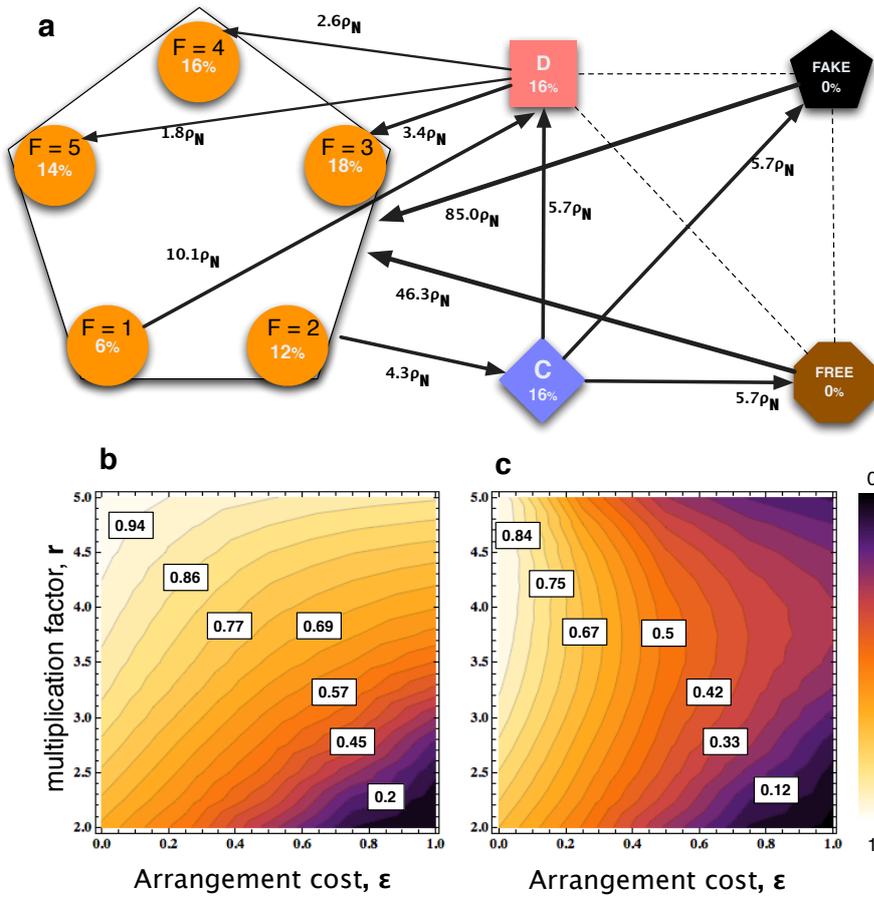


Fig. 3 (a) Stationary distribution and transition probabilities in a population of nine strategies, including five types of $COMP_F$ and the other four strategies. The black arrows identify the transitions that are stronger than neutral, which is $\rho_N = 1/Z$. The dashed lines denote neutral transitions. Also, different types of $COMP_F$, within the pentagon, are neutral among each other. They behave equivalently when facing FAKE, FREE and C players. When playing with D, there are transitions from D to $COMP_3$, $COMP_4$ and $COMP_5$, but it is reversed for $COMP_1$. (b) Cooperation and (c) commitment levels in the population of nine strategies, as a function of ϵ and r . Cooperation level corresponds to the sum of the frequencies of C and all $COMP_F$. Commitment level is the total frequency of all $COMP_F$. In both cases, the smaller the cost of arranging commitment (i.e. the smaller ϵ), the higher level of cooperation and commitment in the population. Furthermore, the cooperation level also increases with r . Parameters: In all panels, $N = 5$, $Z = 100$, $\delta = 2$; $\beta = 0.25$. In panel (a): $r = 4$; $\epsilon = 0.25$.

(Nowak, 2006b, Gokhale and Traulsen, 2010). This derivation requires us to combine the information in Table 1 and Equation (4).

1. $COMP_F$ is risk-dominant against FREE when

$$\sum_{k=1}^N (rc - c - \epsilon/k) \geq (N-1)(rc - c), \quad (5)$$

which can be simplified to

$$\epsilon \leq \frac{c(r-1)}{H_N}, \quad \text{where } H_N = \sum_{k=1}^N 1/k \quad ^4. \quad (6)$$

2. $COMP_F$ is risk-dominant against FAKE if

$$\sum_{k=1}^N \left(\left(\frac{rk}{N} - 1 \right) c + \frac{N\delta - \epsilon}{k} - \delta \right) \geq \sum_{k=1}^{N-1} \left(\frac{r(N-k)}{N} c - \delta \right), \quad (7)$$

which can be simplified to

$$\delta \geq \frac{(N-r)c}{NH_{N-1}} + \frac{H_N}{NH_{N-1}} \epsilon \quad (8)$$

Note that both conditions do not depend on the participation level F .

3. $COMP_F$ is risk-dominant against D if

$$\epsilon \leq \left(\frac{r+F-N-1}{H_N - H_{F-1}} \right) c \quad (9)$$

Combining these three conditions, one can see that $COMP_F$ is risk-dominant against all three defecting strategies when δ satisfies Equation (8) and ϵ is bounded as follows

$$\epsilon \leq \min \left\{ \frac{r+F-N-1}{H_N - H_{F-1}}, \frac{r-1}{H_N} \right\} \cdot c. \quad (10)$$

These conditions can be understood intuitively. For a successful commitment, the cost of arranging the commitment needs to be justified with respect to the benefit of mutual cooperation, i.e. $(r-1)c$, as well as a sufficient compensation is to be arranged. Moreover, the necessary condition for the cost, specified in Equation (10), to hold is $r+F \geq N+1$. Since $N > r$ it implies that $F \geq 2$. That said, $COMP_I$ can never be risk-dominant or viable against defectors. Commitment proposers should initiate the PGG only when there is at least one other player, apart herself, agreeing to commit.

This analytical observation is corroborated by the numerical computation of the transition probabilities and stationary distribution (see Methods in the

⁴ It is known that $0.577215 < H_n - \log n \leq 1$ (Knuth, 2011). Using these inequalities we provide, in the Appendix, some further simplifications of the analytical formulas that include H_N which are presented throughout this results section.

previous section) in Figure 3a. Note the transitions between different types of $COMP_F$ players and D. There are transitions from D to $COMP_3$, $COMP_4$ and $COMP_5$, but it is reversed for $COMP_1$. No transitions between $COMP_2$ and D are shown since they are both weaker than neutral, though the one from $COMP_2$ to D is stronger than the opposite one. All $COMP_F$ behave equivalently to each other when facing FAKE, FREE and C players, with the same transition probabilities from and to the three latter strategies. There are cycles from C to defecting strategies and back to certain commitment proposing strategies, namely, $COMP_3$, $COMP_4$ and $COMP_5$.

Note that the risk-dominance conditions are derived in the limit of large population sizes Z . It is therefore interesting (and possible) to analyze these conditions when the group size $N \rightarrow \infty$ (see again Equation 4). In Appendix we provide detailed analysis for this. Apart from conclusions that are obtained from direct application of taking the limit of the right-hand sides in the inequalities (8) and (10), we obtained necessary asymptotic conditions for the multiplication factor $r(N)$, the cost of arranging commitment $\epsilon(N)$, and the compensation $\delta(N)$, all as a function of N , for which $COMP_F$ can be risk-dominant against all the defective strategies, as follows

1. $r(N)$ satisfies that $r(N) = \Omega(\log N)$, i.e.

$$\lim_{N \rightarrow \infty} \frac{r(N)}{\log N} > 0; \text{ and}$$

2. $\epsilon(N)$ grows at most as fast as $\frac{r(N)}{\log N}$ when $N \rightarrow \infty$; and
3. $\delta(N)$ grows at least as fast as $\frac{\epsilon(N)}{N}$ when $N \rightarrow \infty$.

4.2 Emergence of cooperation and sufficient participation levels

To provide further understanding on the viability of $COMP_F$ in dealing with defectors and free-riders, we compute the stationary distributions for a range of values of the commitment cost ϵ and the multiplication factor r , see Figures 3b and 3c. The results show that, when the cost of arranging the commitment is sufficiently small compared to the cost of cooperation, commitment arranging behavior is frequent leading to a high level of cooperation in the population. We plot the total levels of cooperation and commitment in this population (where cooperation level is the sum of the frequencies of C and of all the $COMP_F$ strategies; and commitment level is the sum of frequency of each $COMP_F$). We can see that in both cases, the smaller the cost of arranging commitment (i.e. the smaller ϵ), the higher level of cooperation and commitment in the population. Furthermore, the cooperation level also increases when the dilemma becomes less harsh, i.e. with increasing r . The contour lines of $\epsilon \in [0.25, 0.3]$ mark the transition between more and less than 50% of cooperation and commitment in the PGG. Our additional numerical results (see Figure 7 in Appendix) demonstrate that these results are robust for different values of compensation cost and group size.

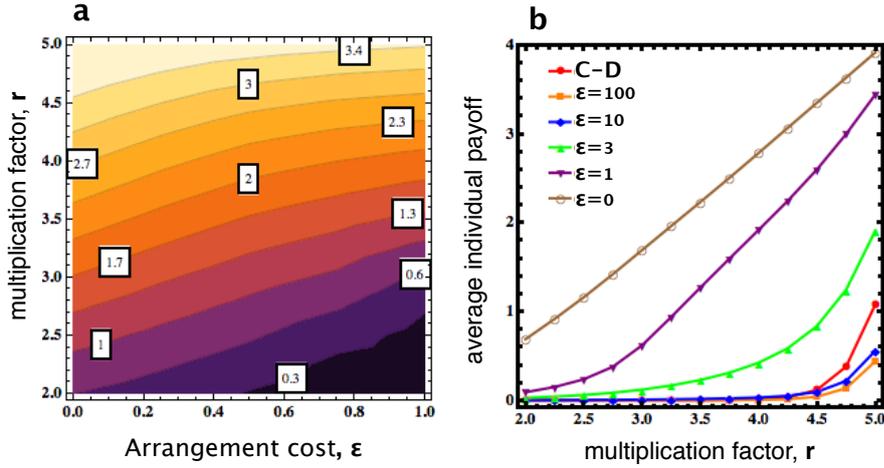


Fig. 4 (a) Average individual payoff in the population of nine strategies, as a function of ϵ and r . Similar to the cooperation level, the smaller the cost of arranging commitment (i.e. the smaller ϵ), the higher the average individual payoff in the population. It also increases with r . (b) Average individual payoff in the population of nine strategies in comparison with that in the population of two strategies C and D, for varying r . For a reasonable ϵ , the average individual payoff (hence also the population total payoff or welfare) in the former population is significantly greater than that in the latter one. For very high values of ϵ , it is reverse when r is high. Parameters: In both panels, $N = 5$, $Z = 100$, $\delta = 2$; $\beta = 0.25$.

One can easily determine the effect of commitment on the average payoff and compare that to the case where commitment is absent. In the latter case, as commitment is absent there is only one decision the players need to make, which is whether to cooperate or defect in the PGG. Hence, we have only two possible strategies C (always cooperate) and D (always defect). In Figure 4a, we plot the average individual payoff when commitment is present, as a function of ϵ and r . This average is calculated by weighting the payoff in each monomorphic state of the Markov chain by its frequency in the stationary distribution. Namely, the average individual payoff in each $COMP_F$ state is $(rc - c - \epsilon/N)$, in C state is $(rc - c)$, and in each defective strategy state (i.e. FREE, D and FAKE) is 0. One can observe in the figure, on one hand, that, similarly to the cooperation level, the smaller the cost of arranging commitment (i.e. the smaller ϵ) and the higher the multiplication factor of the PGG (i.e. r), the higher the average individual payoff in the population. On the other hand, for a reasonable ϵ , the average individual payoff when commitment is present is significantly larger than when it is absent (see Figure 4b), demonstrating that the introduction of commitment not only improves the level of cooperation but, as a consequence, the overall population welfare.

As was shown in Figure 3a, the participation level plays a determining role in the evolutionary viability of the commitment proposing strategies, in relation to their ambition to induce cooperation in the population. We determine now in Figure 5 the appropriate or minimal participation level that

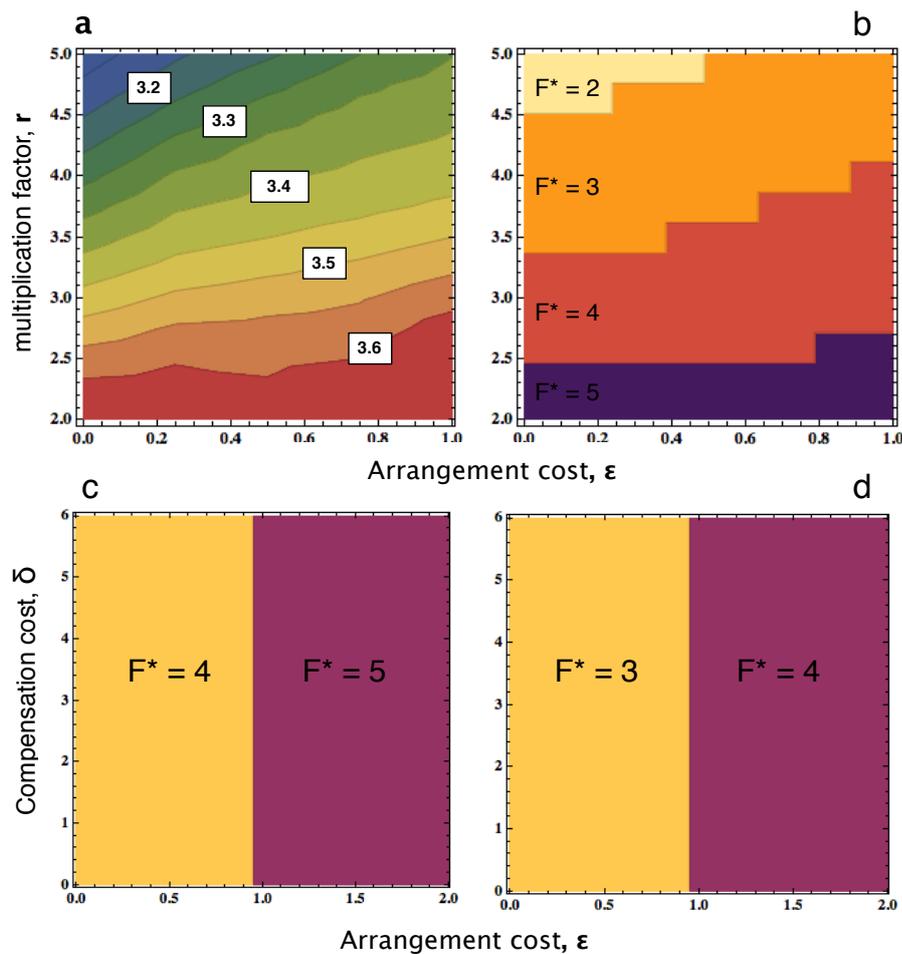


Fig. 5 (a) Average level of commitment among commitment proposing strategies and (b) optimal commitment strategy (denoted by F^*), as a function of r and ϵ , in a population of the nine strategies. There is an intermediate value of F for which commitment proposers are most frequent. The larger the cost of arranging commitment, ϵ , and the harsher the PGG dilemma (i.e. the lower r), the higher F^* is and a higher average number of committers are required to play the game (i.e. the stricter $COMP_F$ should be). Optimal threshold F^* as a function of ϵ and δ , for (c) $r = 2.5$ and (d) $r = 4.0$. There is an intermediate value F^* where the highest frequency is achieved, and the higher ϵ , the higher F^* becomes. Moreover, given the cost of arranging commitment, the optimal commitment threshold does not depend on the compensation δ . Parameters: In all panels: $N = 5$, $Z = 100$, $\beta = 0.25$. In panels (a) and (b): $\delta = 6$.

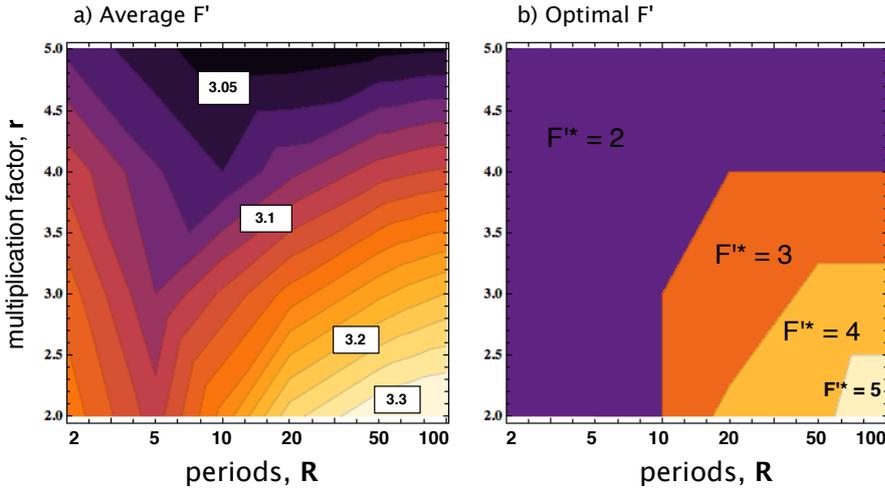


Fig. 6 (a) Average F' level in the population among commitment proposing strategies and (b) optimal F' level (with the highest frequency), as a function of r and R , in a population of all types of $COMP_{F, F'}$ and the other four strategies. Parameters: $N = 5$, $Z = 100$; $\epsilon = 0.25$, $\delta = 20$; $\beta = 0.25$.

ensures sufficient frequencies of the commitment proposer and cooperation, for different values of the most relevant parameters. The results show that the participation level depends both on the dilemma at stake and the cost of arranging the commitment. Namely, the harsher the dilemma (i.e. small r) and the costlier the commitment arrangement (i.e. bigger ϵ), the more agents need to accept to contribute in the PGG in order for commitment proposers to invest in the common good themselves (see Figures 5a and 5b). Figure 5b also reveals that for a given situation there is an optimal threshold (F^*) that varies most strongly with the severity of the dilemma. Moreover, given the cost of arranging commitments, this optimal commitment threshold is (almost) independent of the compensation δ (see Figures 5c and 5d). This observation is robust for a predefined r . These results also imply that ϵ is the essential parameter in a group commitment system, determining which participation level is required to ensure high levels of cooperation. Indeed, as shown in Figure 7 in Appendix—where we plot the total frequency of the commitment proposing strategies as a function of ϵ and δ —when δ reaches a certain threshold, increasing it does not lead to notable improvement.

4.3 Lenience in long-term commitments

Suppose now that commitments may last more than 1 round (denoted by $R > 1$)⁵. When facing FAKE players who commit but then do not contribute,

⁵ “Delayed-return systems in all their variety (for almost all human societies are of this type) have basic implications for social relationships and social groupings: they depend

$COMP_F$ can choose to take immediately the compensation as stated in the agreement thereby ceasing the group interaction for the rest of the commitment time ($R - 1$). Yet, the commitment player may see that although the expected number F was not attained, there is still sufficient participation to make it worthwhile to continue for the remaining rounds. The model discussed so far can be easily extended to incorporate this kind of behavior. The question then becomes for which parameters the $COMP_F$ players can be lenient and for which conditions they need to be strict.

The model is extended to incorporate such behavior by adding another threshold F' ($1 \leq F' \leq N$) saying that, as long as the number of contributors in the group is at least F' , $COMP_{F,F'}$ will not demand the compensation (thereby ceasing the interactions) and continue to interact in the current group.

We compare all types of $COMP_{F,F'}$ with different values of F and F' (i.e. N^2 of them) within a population of these strategies with the other four non-proposing commitment strategies (i.e. 29 strategies in total). To understand the evolutionary viability of $COMP_{F,F'}$ and to calculate the stationary distributions we need to consider how this increase of complexity to the strategies affects the payoff values listed in Table 1 (note that we now calculate payoffs averaged over all rounds):

- The payoffs between two different types of $COMP_{F,F'}$ are the same as the payoffs between two different types of $COMP_F$ in the one-shot game (i.e. $rc - c - \frac{c}{N}$) because they all commit and cooperate in all the rounds.
- The interactions between $COMP_{F,F'}$ on one hand, and C, D and FREE, on the other hand, do not depend on the parameter F' as they all behave the same in all the R rounds (C and FREE commit and cooperate in all the rounds; while, in case of D, the game is either not formed in the first place when there are not sufficient committers in the group or when it is otherwise formed, it stays the same as D players do not commit in the first place – hence no compensation can be enforced). Hence, the payoffs for all these pairs are identical to those in Table 1.
- Only the payoffs for each $COMP_{F,F'}$ strategy when interacting with FAKE remain to be determined. If $k \geq F'$, $COMP_{F,F'}$ will not demand the compensation even when FAKE committed and does not contribute. Hence, $\Pi_{COMP_{F,F'},FAKE}(k) = \Pi_{COMP_F,FAKE}(k)$ and $\Pi_{FAKE,COMP_{F,F'}}(k) = \Pi_{FAKE,COMP_F}(k)$. Otherwise, i.e. if $k < F'$, only the first round takes place and then the commitment is broken as $COMP_{F,F'}$ will demand immediate compensation from FAKE players. Hence, $\Pi_{COMP_{F,F'},FAKE}(k) = \frac{1}{R}\Pi_{COMP_F,FAKE}(k)$ and $\Pi_{FAKE,COMP_{F,F'}}(k) = \frac{1}{R}\Pi_{FAKE,COMP_F}(k)$.

The results as visualized in Figure 6, show that for increasing R (i.e. the longer the commitment lasts), the stricter $COMP_{F,F'}$ should be when there is an agreement in place, requiring more contributors in the group to decide to continue to play. Concretely, take for instance the results in Figure 6 for

for their effective operation on a set of ordered, differentiated, jurally-defined relationships through which crucial goods and services are transmitted. They imply binding commitments and dependencies between people.” Quoted from (Woodburn, 1982)

$\epsilon = 0.25$: To set up the agreement, different starting participation levels F are necessary depending on the severity of the game (see Figures 5a-b). For instance, for $r = 3$, a commitment proposer should play the PGG if only three other players accept to commit (or 2.5 additional players on average). Taking this result, we can see in Figure 6 that even when starting out more strict, the proposer can be lenient towards the actual level of participation as long as the number of rounds is low (in this case $R < 10$): Only when the number of rounds becomes bigger than 40 will she need to be as strict as at the start of the game (i.e. $F = F' = 4$). Notice that a rather high δ (namely, equal 20) is used in Figure 6 to ensure a sufficient average compensation cost for varying R , though our additional analysis shows the results are qualitatively the same when varying δ . Furthermore, we see the same tendency when computing the average frequency of F' in the population (as can be seen in Figure 6a).

5 Conclusions

We have provided a new evolutionary game theory model, which shows that arranging prior commitments in multiagent group interactions, not just pair-wise ones, provides a pathway towards the evolution of cooperation in the typical Public Goods Game (PGG). Our analytical and numerical results clearly exhibit that if the cost of arranging commitment is sufficiently small compared to cost of cooperation, then commitment arranging behavior becomes frequent, leading henceforth to high levels of cooperation in a population sporting a representative variety of playing strategies. Furthermore, an optimal prior commitment participation level emerges, dependent both on the common goods dilemma and on the cost of arranging commitment. In particular, the harsher the dilemma and the costlier the commitment, the higher the required commitment participation level to ensure the success of the joint venture. Additionally, as a commitment deal may last for more than one round, we evince that longer-lasting commitments require a greater strictness upon fake committers than short ones.

The results we obtain are in close accordance with experimental economic outcomes obtained by others (Cherry and McEvoy, 2013, Chen and Komorita, 1994). But the present work further reveals that, whenever the compensation that needs to be paid by fake committers reaches a certain threshold, increasing it does not lead to improvement in terms of cooperation levels. It implies that, when designing norms, whether in real life or a self-organizing MAS, it is not necessary to have an infinitely large compensation or sanction against law breakers, for a sufficient one is enough for a wide range of situations. Moreover, the current paper suggests the need for further behavioral experiments to explore the effects of varying the essential parameters that drive commitments.

Finally, our work provides specific insights into the design of self-organizing MAS when dealing with group interactions. For instance, in finding the effective degrees of commitment one should require from group members which lead to highest levels of cooperation. Another instance concerns the use of

commitment with a view to cooperation in lieu of intention recognition (Han et al., 2011, 2012a, 2015a), when the latter is not reliable enough. Commitment can be seen as a cogent form of intention manifestation, and achieving joint intentions is an exclusive characteristic of the human species (Tomasello, 2014). Notwithstanding, as groups increase in size, intention recognition may need to be supplemented by explicit intention commitment to ensure joint success (Tomasello et al., 2005).

Acknowledgments

TAH acknowledges the support of the Teesside university URF funding (URF-11200174). TL acknowledges the support of *Fondation de la Recherche Scientifique* - FNRS through the grant FRFC nr. 2.4614.12 and the *Fonds voor Wetenschappelijk Onderzoek* - FWO. through the grant nr. G.0391.13N. LMP acknowledges the support from FCT/MEC NOVA LINCS PEst UID/CEC/04516/2013.

References

- Airiau, S., Sen, S., and Villatoro, D. (2014). Emergence of conventions through social learning. *Autonomous Agents and Multi-Agent Systems*, 28(5):779–804.
- Alvard, M. S. and Nolin, D. A. (2002). Rousseau’s whale hunt? *Current Anthropology*, 43(4):533–559.
- Axelrod, R. and Hamilton, W. (1981). The evolution of cooperation. *Science*, 211:1390–1396.
- Barrett, S. (2003). *Environment and Statecraft: The Strategy of Environmental Treaty-Making: The Strategy of Environmental Treaty-Making*. Oxford University Press.
- Bloembergen, D., Tuyls, K., Hennes, D., and Kaisers, M. (2015). Evolutionary dynamics of multi-agent learning: a survey. *Journal of Artificial Intelligence Research*, pages 659–697.
- Castelfranchi, C. and Falcone, R. (1998). Principles of trust for MAS: Cognitive anatomy, social importance, and quantification. In *Proceedings of International Conference on Multi Agent Systems*, pages 72–79. IEEE.
- Castelfranchi, C. and Falcone, R. (2010). *Trust Theory: A Socio-Cognitive and Computational Model (Wiley Series in Agent Technology)*. Wiley.
- Chen, X., Sasaki, T., Brännström, Å., and Dieckmann, U. (2015). First carrot, then stick: how the adaptive hybridization of incentives promotes cooperation. *Journal of The Royal Society Interface*, 12(102):20140935.
- Chen, X.-P. and Komorita, S. S. (1994). The effects of communication and commitment in a public goods social dilemma. *Organizational Behavior and Human Decision Processes*, 60(3):367–386.
- Cherry, T. L. and McEvoy, D. M. (2013). Enforcing compliance with environmental agreements in the absence of strong institutions: An experimental analysis. *Environmental and Resource Economics*, 54(1):63–77.
- Duong, M. H. and Han, T. A. (2015). On the expected number of equilibria in a multi-player multi-strategy evolutionary game. *Dynamic Games and Applications*, pages 1–23.
- Duong, M. H. and Han, T. A. (2016). Analysis of the expected density of internal equilibria in random evolutionary multi-player multi-strategy games. *Journal of Mathematical Biology*, pages 1–34. In Press.
- Fehr, E. and Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review*, 90(4):980–994.
- Fudenberg, D. and Imhof, L. A. (2005). Imitation processes with small mutations. *Journal of Economic Theory*, 131:251–262.
- Gokhale, C. S. and Traulsen, A. (2010). Evolutionary games in the multiverse. *Proceedings of the National Academy of Sciences of the United States of America*, 107(12):5500–5504.
- Han, T. A. (2016). Emergence of social punishment and cooperation through prior commitments. In *Proceedings of the Conference of the American Association of Artificial Intelligence (AAAI’2016)*, pages 2494–2500, Phoenix, Arizona, USA.

- Han, T. A. and Lenaerts, T. (2016). A synergy of costly punishment and commitment in cooperation dilemmas. *Adaptive Behavior*. In Press.
- Han, T. A., Lenaerts, T., Santos, F. C., and Pereira, L. M. (2015a). Emergence of cooperation via intention recognition, commitment and apology—a research summary. *AI Communications*, pages 1–7.
- Han, T. A., Moniz Pereira, L., and Lenaerts, T. (2015b). Avoiding or Restricting Defectors in Public Goods Games? *Journal of the Royal Society Interface*, 12(103):20141203.
- Han, T. A., Pereira, L. M., and Santos, F. C. (2011). The role of intention recognition in the evolution of cooperative behavior. In Walsh, T., editor, *Proceedings of the 22nd international joint conference on Artificial intelligence (IJCAI'2011)*, pages 1684–1689. AAAI.
- Han, T. A., Pereira, L. M., and Santos, F. C. (2012a). Corpus-based intention recognition in cooperation dilemmas. *Artificial Life journal*, 18(4):365–383.
- Han, T. A., Pereira, L. M., and Santos, F. C. (2012b). The emergence of commitments and cooperation. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'2012)*, pages 559–566. ACM.
- Han, T. A., Pereira, L. M., Santos, F. C., and Lenaerts, T. (2013a). Good agreements make good friends. *Scientific reports*, 3(2695).
- Han, T. A., Pereira, L. M., Santos, F. C., and Lenaerts, T. (2013b). Why Is It So Hard to Say Sorry: The Evolution of Apology with Commitments in the Iterated Prisoner’s Dilemma. In *Proceedings of the 23rd international joint conference on Artificial intelligence (IJCAI'2013)*. AAAI Press.
- Han, T. A., Santos, F. C., Lenaerts, T., and Pereira, L. M. (2015c). Synergy between intention recognition and commitments in cooperation dilemmas. *Scientific reports*, 5(9312).
- Han, T. A., Traulsen, A., and Gokhale, C. S. (2012c). On equilibrium properties of evolutionary multiplayer games with random payoff matrices. *Theoretical Population Biology*, 81(4):264–272.
- Hardin, G. (1968). The tragedy of the commons. *Science*, 162:1243–1248.
- Harrenstein, P., Brandt, F., and Fischer, F. (2007). Commitment and extortion. In *Proceedings of the 6th international joint conference on Autonomous agents and MultiAgent systems, AAMAS '07*, New York, NY, USA. ACM.
- Hasan, M. R. and Raja, A. (2013). Emergence of cooperation using commitments and complex network dynamics. In *Web Intelligence (WI) and Intelligent Agent Technologies (IAT), 2013 IEEE/WIC/ACM International Joint Conferences on*, volume 2, pages 345–352. IEEE.
- Hauert, C., Monte, S. D., Hofbauer, J., and Sigmund, K. (2002). Volunteering as red queen mechanism for cooperation in public goods games. *Science*, 296(5570):1129. [DOI: 10.1126/science.1070582].
- Hauert, C., Traulsen, A., Brandt, H., Nowak, M. A., and Sigmund, K. (2007). Via freedom to coercion: The emergence of costly punishment. *Science*, 316:1905–1907.
- Hofbauer, J. and Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press.

- Hofmann, L.-M., Chakraborty, N., and Sycara, K. (2011). The evolution of cooperation in self-interested agent societies: a critical study. In *The 10th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '11*, pages 685–692.
- Imhof, L. A., Fudenberg, D., and Nowak, M. A. (2005). Evolutionary cycles of cooperation and defection. *Proceedings of the National Academy of Sciences of the United States of America*, 102:10797–10800.
- Karlin, S. and Taylor, H. E. (1975). *A First Course in Stochastic Processes*. Academic Press, New York.
- Knuth, D. E. (2011). *The Art of Computer Programming, Volume 4A: Combinatorial Algorithms, Part 1*. Pearson Education India.
- Lerat, J.-S., Han, T. A., and Lenaerts, T. (2013). Evolution of common-pool resources and social welfare in structured populations. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 2848–2854. AAAI Press.
- Maier-Rigaud, F. P., Martinsson, P., and Staffiero, G. (2010). Ostracism and the provision of a public good: experimental evidence. *Journal of Economic Behavior & Organization*, 73(3):387–395.
- Nesse, R. M. (2001). *Evolution and the capacity for commitment*. Russell Sage Foundation series on trust. Russell Sage.
- Nowak, M. A. (2006a). *Evolutionary Dynamics: Exploring the Equations of Life*. Harvard University Press, Cambridge, MA.
- Nowak, M. A. (2006b). Five rules for the evolution of cooperation. *Science*, 314(5805):1560. DOI: 10.1126/science.1133755.
- Nowak, M. A., Sasaki, A., Taylor, C., and Fudenberg, D. (2004). Emergence of cooperation and evolutionary stability in finite populations. *Nature*, 428:646–650.
- Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge university press.
- Perc, M., Gómez-Gardeñes, J., Szolnoki, A., Floría, L. M., and Moreno, Y. (2013). Evolutionary dynamics of group interactions on structured populations: a review. *Journal of The Royal Society Interface*, 10(80):20120997.
- Rand, D. G., Tarnita, C. E., Ohtsuki, H., and Nowak, M. A. (2013). Evolution of fairness in the one-shot anonymous ultimatum game. *Proc. Natl. Acad. Sci. USA*, 110:2581–2586.
- Ranjbar-Sahraei, B., Bou Ammar, H., Bloembergen, D., Tuyls, K., and Weiss, G. (2014). Evolution of cooperation in arbitrary complex networks. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 677–684. International Foundation for Autonomous Agents and Multiagent Systems.
- Ray, D. (2007). *A game-theoretic perspective on coalition formation*. Oxford University Press.
- Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M. W., Fogarty, L., Ghirlanda, S., Lillicrap, T., and Laland, K. N. (2010). Why copy others? insights from the social learning strategies tournament. *Science*, 328(5975):208–213.

- Santos, F. C., Santos, M. D., and Pacheco, J. M. (2008). Social diversity promotes the emergence of cooperation in public goods games. *Nature*, 454:214–216.
- Sasaki, T., Okada, I., Uchida, S., and Chen, X. (2015). Commitment to cooperation and peer punishment: Its evolution. *Games*, 6(4):574–587.
- Schelling, T. C. (1990). *The strategy of conflict*. London: Oxford University Press.
- Shehory, O. and Kraus, S. (1998). Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1):165–200.
- Shehory, O. M., Sycara, K., and Jha, S. (1998). Multi-agent coordination through coalition formation. In *Intelligent Agents IV Agent Theories, Architectures, and Languages*, pages 143–154. Springer.
- Sigmund, K. (2010). *The Calculus of Selfishness*. Princeton University Press.
- Sigmund, K., Hauert, C., and Nowak, M. (2001). Reward and punishment. *P Natl Acad Sci USA*, 98(19):10757–10762.
- Sigmund, K., Silva, H. D., Traulsen, A., and Hauert, C. (2010). Social learning promotes institutions for governing the commons. *Nature*, 466:7308.
- Stander, P. E. (1992). Cooperative hunting in lions: the role of the individual. *Behavioral ecology and sociobiology*, 29(6):445–454.
- Sterelny, K. (2012). *The evolved apprentice*. MIT Press.
- Sugawara, T. (2011). Emergence and stability of social conventions in conflict situations. In *IJCAI*, pages 371–378.
- Szolnoki, A. and Perc, M. (2010). Impact of critical mass on the evolution of cooperation in spatial public goods games. *Physical Review E*, 81(5):057101.
- Szolnoki, A. and Perc, M. (2012). Evolutionary advantages of adaptive rewarding. *New Journal of Physics*, 14(9):093016.
- Szolnoki, A. and Perc, M. (2013). Correlation of positive and negative reciprocity fails to confer an evolutionary advantage: Phase transitions to elementary strategies. *Physical Review X*, 3(4):041021.
- Tomasello, M. (2014). *Natural History of Human Thinking*. Harvard University Press.
- Tomasello, M., Carpenter, M., Call, J., Behne, T., and Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and brain sciences*, 28(05):675–691.
- Traulsen, A. and Hauert, C. (2009). Stochastic evolutionary game dynamics. *Reviews of nonlinear dynamics and complexity*, 2:25–61.
- Traulsen, A., Hauert, C., De Silva, H., Nowak, M. A., and Sigmund, K. (2009). Exploration dynamics in evolutionary games. *Proc. Natl. Acad. Sci. USA*, 106(3):709–712.
- Traulsen, A. and Nowak, M. A. (2006). Evolution of cooperation by multilevel selection. *Proceedings of the National Academy of Sciences*, 103(29):10952–10955.
- Traulsen, A., Nowak, M. A., and Pacheco, J. M. (2006). Stochastic dynamics of invasion and fixation. *Phys. Rev. E*, 74:11909.
- Tuyls, K. and Parsons, S. (2007). What evolutionary game theory tells us about multiagent learning. *Artificial Intelligence*, 171(7):406–416.

- Van Segbroeck, S., Pacheco, J. M., Lenaerts, T., and Santos, F. C. (2012). Emergence of fairness in repeated group interactions. *Physical review letters*, 108(15):158104.
- West, S., Griffin, A., and Gardner, A. (2007). Evolutionary explanations for cooperation. *Current Biology*, 17:R661–R672.
- Winikoff, M. (2007). Implementing commitment-based interactions. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, AAMAS '07, pages 868–875, New York, NY, USA. ACM.
- Woodburn, J. (1982). Egalitarian societies. *Man, New Series*, 17(3):431–451.
- Wooldridge, M. and Jennings, N. R. (1999). The cooperative problem-solving process. In *Journal of Logic and Computation*, pages 403–417.
- Zisis, I., Guida, S. D., Han, T. A., Kirchsteiger, G., and Lenaerts, T. (2015). Generosity motivated by acceptance - evolutionary analysis of an anticipation games. *Scientific reports*, 5(18076).

A Some simplifications of the obtained analytical results

Here, using the well-known inequalities (Knuth, 2011)

$$\log N + \gamma < F_N = \sum_{k=1}^N \frac{1}{k} \leq \log N + 1$$

where $\gamma = 0.577215$, we provide some simplifications of the conditions obtained in the main text. First of all, regarding the conditions for risk-dominance of $COMP_F$ against D, FREE and FAKE:

$$\begin{aligned} \epsilon &\leq \min \left\{ \frac{r + F - N - 1}{H_N - H_{F-1}}, \frac{r - 1}{H_N} \right\} \times c, \\ \delta &\geq \frac{N - r}{NF_{N-1}}c + \frac{F_N}{NF_{N-1}}\epsilon. \end{aligned} \quad (11)$$

They can be simplified to

$$\begin{aligned} \epsilon &\leq \min \left\{ \frac{r + F - N - 1}{\log \frac{N}{F-1} + \gamma - 1}, \frac{r - 1}{\log N + \gamma} \right\} \times c, \\ \delta &\geq \frac{(N^2 - rN)c + \epsilon}{N^2 (\log(N - 1) + 1)} + \frac{\epsilon}{N}. \end{aligned} \quad (12)$$

B Risk-dominance conditions when $N \rightarrow \infty$

Note that the risk-dominance condition described in the Method section is valid for the large population size limit ($Z \rightarrow \infty$). We now analyze these conditions for $N \rightarrow \infty$. We assume that the multiplication factor r is a function of N , i.e. $r(N)$.

From (12), if $\log N$ grows much faster than $r(N)$ (or using little-o notation, $r(N) = o(\log N)$), i.e. if

$$\lim_{N \rightarrow \infty} \frac{r(N)}{\log N} = 0,$$

then the right-hand side of the first inequality in (12) becomes 0. It implies that when $N \rightarrow \infty$ the cost of arranging the commitment must be infinitely small so that COMP can be risk-dominant against defective strategies (namely, in this case, the FREE strategy). Hence, provided that arranging commitment is costly ($\epsilon > 0$), it must hold that $r(N) \neq o(\log N)$, or, using the big omega notation,

$$r(N) = \Omega(\log N) \text{ (as } N \rightarrow \infty \text{)}. \quad (13)$$

In this case, we have

$$\epsilon \leq \min \left\{ \lim_{N \rightarrow \infty} \frac{r(N) + F - N - 1}{\log \frac{N}{F-1}}, \lim_{N \rightarrow \infty} \frac{r(N)}{\log N} \right\} c \quad (14)$$

Hence, the necessary condition for this inequality to hold is that $r(N) = \Omega(\log N)$ and ϵ grows at most as fast as $\frac{r(N)}{\log N}$.

Consider now the second risk-dominance inequality

$$\delta \geq \frac{(N^2 - rN)c + \epsilon}{N^2 (\log(N - 1) + 1)} + \frac{\epsilon}{N}. \quad (15)$$

If ϵ is a constant, clearly the right hand side becomes 0 as $N \rightarrow \infty$, i.e. δ can be indefinitely small while still ensuring the condition is satisfied (i.e. COMP is risk-dominant against FAKE). Let's assume that ϵ is a function of N . Then, the same conclusion applies if

$$\epsilon(N) = o(N).$$

Otherwise, i.e. if $\epsilon(N) = \Omega(N)$, then the the inequality can be rewritten as

$$\delta \geq \lim_{N \rightarrow \infty} \left(\frac{\epsilon(N)}{N^2 (\log(N-1) + 1)} + \frac{\epsilon(N)}{N} \right). \quad (16)$$

That is, roughly, δ needs to grow at least as fast as $\frac{\epsilon(N)}{N}$ when $N \rightarrow \infty$ to guarantee risk-dominance of COMP against FAKE. We also observe that the contribution cost c does not play a role in determining the lower boundary of δ .

In short, as the group size $N \rightarrow \infty$, the risk-dominance conditions are defined in (14) and (16). For these to hold, it is necessary that

1. the multiplication factor $r(N)$ satisfies that $r(N) = \Omega(\log N)$, i.e.

$$\lim_{N \rightarrow \infty} \frac{r(N)}{\log N} > 0$$

2. the cost of arranging commitment $\epsilon(N)$ grows at most as fast as $\frac{r(N)}{\log N}$ as $N \rightarrow \infty$
3. the compensation cost, $\delta(N)$ grows at least as fast as $\frac{\epsilon(N)}{N}$ when $N \rightarrow \infty$

C Results for varying group sizes

In Figure 7, we plot the total frequency of commitment strategies for varying ϵ and δ , and for increasing group sizes. The results show that the observations seen in the main text are robust for varying group sizes N as well as compensation cost δ . Furthermore, ϵ is the essential parameter because as soon as δ reaches a certain threshold, increasing it does not lead to notable improvement.

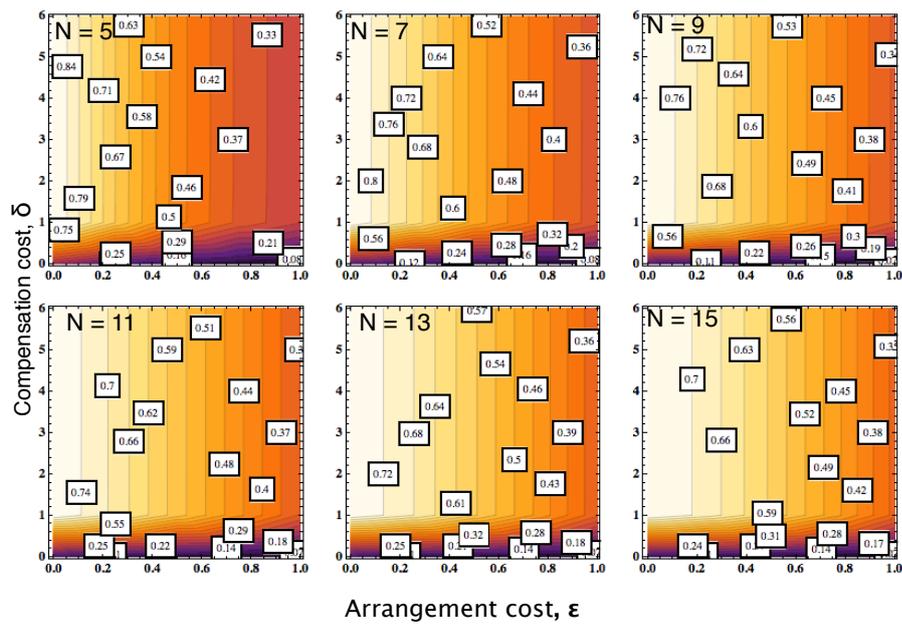


Fig. 7 Results in the main text are robust for varying group size N . We plot the total frequency of the commitment proposing strategies as a function of ϵ and δ . In general, when ϵ is small enough and δ is sufficiently high, commitment proposing strategies are frequent, leading to high levels of cooperation in the PGG. Parameters: $Z = 100$; $r = 4$; $\beta = 0.25$.