

The emergence of reciprocally beneficial cooperation

Sergio Beraldo[♦] and Robert Sugden[♣]

Abstract: We offer a new and robust model of the emergence and persistence of cooperation when interactions are anonymous, the population is well-mixed, and evolution selects strategies according to material payoffs. The model has a Prisoner's Dilemma structure, but with an outside option of non-participation. The payoff to mutual cooperation is stochastic; with positive probability, it exceeds that from cheating against a cooperator. Under mild conditions, mutually beneficial cooperation occurs in equilibrium. This is possible because the non-participation option holds down the equilibrium frequency of cheating. Dynamic properties of the model are investigated theoretically and through simulations based on replicator dynamics.

Keywords: Cooperation; voluntary participation; random payoffs.

JEL Classification: C73.

Date: 9 July 2015

[♦] Department of Economics and Statistics, University of Napoli Federico II, via Cinthia Monte Sant'Angelo, 80124, Napoli, Italy. e-mail address: s.beraldo@unina.it. Tel: +39 081 675179.

[♣] *Corresponding author:* School of Economics and Centre for Behavioural and Experimental Social Science, University of East Anglia, Norwich NR4 7TJ, UK. e-mail address: r.sugden@uea.ac.uk. Tel.: +44 (0) 1603 593423. Fax: +44 (0) 1603 455259

^{*} We thank Pasquale Commendatore for advice on the simulation analysis, Matteo Morini and Francesca Di Iorio for programming and data analysis for the stochastic evolutionary model investigated in an earlier version of this paper, and an associate editor and four anonymous referees for comments. SB acknowledges financial support from the University of Napoli "Federico II" (Programma di breve mobilità) and from the Banco di Napoli Foundation. RS's work was supported by the Economic and Social Research Council of the UK (award no. ES/K002201/1).

1. Introduction

Studies of animal behaviour have found many practices which create collective benefits at some apparent cost or risk to individual participants. Examples include alarm calls, food-sharing, grooming, and participation in inter-group warfare. One of the most fundamental problems in evolutionary biology since Darwin (1859) has been to explain how such forms of cooperation evolve by natural selection. An analogous problem in economics has been to look for explanations of cooperative human practices, such as the fulfilment of market obligations, the provision of public goods through voluntary contributions, and the management of common property resources, that are consistent with the traditional assumption of individual self-interest.¹ Many different theories have been proposed by biologists and economists as possible solutions. Among the mechanisms that have been modelled are direct and indirect reciprocity, kin selection, group selection, and the ‘green beard’ mechanism. (For an overview of these mechanisms, see Nowak, 2006. Tomasello (2014) provides a comprehensive account of why and how cooperation may have evolved among early humans. His hypothesis is that early humans were forced by ecological circumstances into more cooperative modes of life, and this led to the evolution of ways of thinking that were directed towards coordination with others to achieve joint goals.) Some economists have combined biological and economic modes of explanation, hypothesising that human cooperation in the modern world is a product of genetically hard-wired traits that evolved by natural selection to equip *Homo sapiens* for life in hunter-gatherer societies. In some versions of this hypothesis, those traits act as equilibrium selection devices in the modern ‘game of life’ (e.g. Binmore, 1994, 1998); in others, they can generate non-selfish behaviour in modern societies (e.g. Boyd et al. 2005; Bowles and Gintis, 2011).

However, a recent trend in biology has been to question whether such sophisticated explanations are always necessary. Many forms of apparently cooperative behaviour have been found to be forms of mutualism: the ‘cooperating’ individual derives sufficient direct fitness benefit to make the behaviour worthwhile, and any effect on the fitness of others is incidental (e.g. Clutton-Brock, 2002, 2009; Sachs et al., 2004). The Snowdrift game (Sugden, 1986), in which equilibrium involves cooperation by one player and free-riding by the other, is increasingly used in biology as a

¹ In parallel with attempts to explain human cooperation in terms of self-interest, there is a large literature in economics in which cooperation is explained as a product of non-selfish motivations, such as altruism, preferences for equality, fairness or welfare-maximisation, team reasoning, or intrinsic desires to cooperate. Without denying the existence of such motivations, we think it is useful to investigate the explanatory power of the self-interest assumption.

model of such behaviour. In this paper, we present a new model of the evolution of cooperation which fits with this trend of thought.

Our methodological approach treats the biological and economic problems of cooperation as isomorphic to one another. That is, we hypothesise that the emergence and reproduction of human cooperative practices are governed by evolutionary mechanisms that are distinct from, but structurally similar to, those of natural selection. Candidate mechanisms include trial-and-error learning by individuals, imitation of successful neighbours, and cultural selection through inter-group competition. Analyses which use this approach may be both informed by and informative to theoretical biology. For example, Sugden's (1986) analysis of the emergence of social norms was inspired by the earlier work of theoretical biologists, but it developed new models (in particular, the Snowdrift and Mutual Aid games) which have since been widely used in biology (e.g. Leimar and Hammerstein, 2001; Nowak and Sigmund, 2005). The model that we present in this paper can be interpreted as a representation either of natural selection or of trial-and-error human learning.

Our modelling strategy is distinctive in that it uses three assumptions which in combination rule out most of the mechanisms that feature in existing theories of cooperation. Specifically, we assume that interactions are anonymous, that evolution takes place in a large, well-mixed population, and that the evolutionary process selects strategies according to their material payoffs. The assumption of anonymity excludes mechanisms based on reputation, reciprocity or third-party punishment. The assumption of well-mixedness excludes mechanisms of group or kin selection. The assumption that selection is for material payoffs excludes mechanisms which postulate non-selfish preferences as an explanatory primitive. Working within the constraints imposed by these assumptions, we are able to generate a simple and robust model of cooperation.

Our model adapts the familiar framework of a Prisoner's Dilemma that is played recurrently in a large population. We introduce two additional features, which we suggest can be found in many real-world cases of potentially cooperative interaction, both for humans and for other animals.

The first additional feature is that participation in the game is voluntary. One of the restrictive properties of the Prisoner's Dilemma is that, in any given interaction, an individual must act either *pro*-socially (the strategy of cooperation) or *anti*-socially (the strategy of defection or cheating, which allows a cheater to benefit at the expense of a cooperator). There is no opportunity to be simply *asocial*. We add an asocial strategy, that of opting out of the interaction altogether. Of course, if the only difference between anti-social and asocial behaviour was that asocial individuals did not benefit when their co-players chose to cooperate, asociality would be a dominated strategy.

It is an essential part of our model that if both players cheat, both are worse off than if they had opted out of the interaction.²

The second additional feature is that the payoff that each player receives if they both cooperate is subject to random variation. Before choosing his (or her, or its) strategy, each player knows his own cooperative payoff, but not the other player's. With non-zero probability, the payoff from mutual cooperation is greater than that from cheating against a cooperator. Thus, there are circumstances in which it would be profitable for a player to cooperate if he were sufficiently confident that the other player would cooperate too. Crucially, however, it is never *common knowledge* that the payoffs are such that mutual cooperation is a Nash equilibrium. In our model, players receive no information at all about the realisation of the random component of their co-players' cooperative payoff. This is obviously an extreme assumption; we use it only as a modelling simplification. In real interactions, players often have *some* such information. (For example, explanations of animal behaviour in asymmetric contests often depend on the assumption that both contestants recognise some feature of the game which signals which of them is more likely to attach the higher value to the disputed resource [Maynard Smith and Parker, 1976].) But the main qualitative results of our model require only that each player has some private information about his own payoffs such that, with non-zero probability, a player may know that cooperation is a non-dominated strategy for him without knowing whether the same is true for the other player.³

As an illustration of the kind of interaction that our model represents, we offer the following variant of Rousseau's (1755/ 1988, p. 36) story of hunting in a state of nature. Two individuals jointly have the opportunity to invest time and energy to hunt a deer. The hunters can succeed only by acting on a concerted plan out of sight of one another. A hunt begins only if both individuals agree to take part. Each can then cheat by unilaterally pursuing a smaller prey, which the other's deer-hunting tends to flush out and make easier to catch. The anticipated benefit of deer-hunting to an individual, conditional on the other's not cheating, can be different for different individuals and on different occasions. Sometimes, but not always, this benefit is sufficiently low that unilateral cheating pays off.

² In some simplified models of the Prisoner's Dilemma, each player makes an independent decision about whether to incur a fixed cost to confer a larger fixed benefit on the other player; payoffs in the four possible outcomes are defined as sums of these costs and benefits. Because we want to distinguish between cheating and non-participation, our model does not have this additivity property.

³ The game would be very different (and rather trivial) if all payoffs were common knowledge. In that case, there would be a Nash equilibrium in which there was no cheating and in which each player cooperated if and only if *both* players gained more from mutual cooperation than from cheating against a cooperator.

As a more modern illustration, consider two individuals who make contact through the internet. One of them is offering to sell some good which has to be customised to meet the specific requirements of the buyer; the other is looking to buy such a good. If they agree to trade, each individual invests resources in the transaction (exchanging information, producing and dispatching the good, sending payment). Each may have opportunities to gain by deviating from the terms of the agreement. Sometimes, but not always, the benefit of completing the transaction is sufficiently low that unilateral cheating pays off.

We will show how the interaction of voluntary participation and stochastic payoffs can induce cooperation. Of course, it is well known that voluntary participation can facilitate cooperation when players can distinguish between more and less cooperative opponents. If such distinctions are possible, voluntary participation can allow cooperators to avoid interacting with cheats. This can sustain cooperation without the need for informationally and cognitively more demanding strategies of reciprocity or punishment – an idea that can be traced back to Adam Smith’s (1763/ 1978, pp. 538–539) analysis of trustworthiness among traders in commercial societies. But such mechanisms are ruled out by our anonymity assumption.

In our model, voluntary participation facilitates cooperation by a different route. Because would-be cheats have the alternative option of non-participation, and because non-participation is the best response to cheating, the equilibrium frequency of cheating is subject to an upper limit. If cheating occurs at all, the expected payoff from cheating cannot be less than that from non-participation. Thus, for any given frequency of cooperation, the frequency of cheating is self-limiting. The underlying mechanism is similar to that of the Lotka–Volterra model of interaction between predators and prey: the size of the predator population (the frequency of cheating) is limited by the size of the prey population (the frequency of cooperation).

Clearly, however, this mechanism can support cooperation only if, when the frequency of cheating is sufficiently low, some players choose to cooperate. This could not be the case if, as in the Prisoner’s Dilemma, cooperation was *always* a weakly dominated strategy. In our model, random variation in the payoff from mutual cooperation ensures that players sometimes find it worthwhile to cooperate, despite the risk of meeting a cheat. The players who cooperate are those for whom the benefit of mutual cooperation is sufficient to compensate for this risk. Because cooperators are self-selecting in this way, the average payoff in the game is greater than the payoff to non-participation. In other words, despite the presence of cheats, beneficial cooperation occurs.

As a first step in developing an evolutionary model, we begin (in Section 2) by presenting our variant of the Prisoner’s Dilemma as a one-shot game and identifying its symmetric Bayesian

Nash equilibria. We show that, provided the upper bound of the distribution of cooperative benefit is not too low, the game has at least one such equilibrium in which beneficial cooperation occurs. In Section 3 we investigate some comparative-static properties of equilibria in this game. We show that as the distribution of cooperative benefit becomes more favourable, the maximum frequency of cooperation that is sustainable in equilibrium increases. In Section 4 we examine the dynamics of the model, using simple analytical methods. In Section 5 we supplement this analysis by computer simulations based on replicator dynamics. Our analysis shows that, in the neighbourhood of ‘interior’ equilibria in which some but not all players choose non-participation, the dynamics are similar to those of predator-prey models. Depending on the payoffs of the game, interior equilibria may be locally stable (with evolutionary paths spiralling in from a large zone of attraction) or unstable (with evolutionary paths spiralling out and ending at an equilibrium of non-participation). In Section 6, we discuss the contribution that our model can make to the explanation of cooperative behaviour. We show that, despite sharing some features of existing biological models of mutualism and voluntary participation, our model isolates a distinct causal mechanism.

2. The model: equilibrium properties

In this Section and in Section 3, we present our game in one-shot form and analyse its equilibrium properties. This is a game for two players 1 and 2. For each player $i \in \{1, 2\}$, the benefit x_i that he gains if both players cooperate is an independent realisation of a random variable X whose distribution $f(\cdot)$ is continuous with support $[x_{\min}, x_{\max}]$. Each player knows his own benefit but not that of the other player. Given this knowledge, he plays a game with three pure strategies – to cooperate (C), to cheat (D), or not to participate (N). The payoff matrix is shown in Table 1.

Table 1: Payoff matrix for the game

		player 2		
		N	C	D
player 1	N	0, 0	0, 0	0, 0
	C	0, 0	x_1, x_2	$-b, a$
	D	0, 0	$a, -b$	$-c, -c$

$$x_{\max} > a > x_{\min} \geq 0; b > c > 0.$$

The essential features of the game are contained in the structure of best responses.⁴ The condition $x_{\max} > a > x_{\min}$ implies that either C or D (but not N) may be the better response to C, depending on the relevant player's realisation of X . The condition $b > c$ implies that, as in the Prisoner's Dilemma, D is better than C as a response to D. Given that the payoff to N is normalised to zero, $a > 0$ implies that cheating gives a higher payoff than non-participation if the opponent cooperates; $c > 0$ implies that the opposite is the case if the opponent cheats. No assumption is made about whether $a - b$ (i.e. the net benefit of an interaction in which one player cooperates and the other cheats) is positive, zero or negative. It is easy to imagine real-world applications (such as our example of internet trading) in which any of these possibilities would be plausible.⁵ The condition $x_{\min} \geq 0$ (which is not essential for our main results) implies that players are never worse off from mutual cooperation than from non-participation.

We now consider symmetric Bayesian Nash equilibria (SBNE) of the game. Although our formal analysis in this Section and in Section 3 treats the game as one-shot, our ultimate concern is with SBNEs as possible stationary points in an evolutionary process in a well-mixed population.

We will say that a pure strategy (N, C or D) is *played* in a given equilibrium if and only if the unconditional probability with which it is played is nonzero. Some significant properties of SBNE hold for all parameter values. First, there is a *non-participation equilibrium* in which only N is played; in this equilibrium, players' payoffs are zero and unilateral deviations lead to neither gain nor loss. Second, there is no SBNE in which C is played but D is not. (Against an opponent who might play C but will not play D, i 's best response plays D when $x_i < a$.) Third, there is no SBNE in which D is played but not C. (Against an opponent who might play D but will not play C, N is the unique best response.) Thus, only two types of equilibrium participation are possible. Depending on the parameter values, there may be an *interior equilibrium* in which N, D and C are all played; and there may be a *boundary equilibrium* in which D and C are played but not N.

⁴ Provided that this best response structure is maintained, the main implications of the model are preserved. It is not essential that the payoff from playing C against C is stochastic and that all other payoffs are not; but there must be some random variation in the payoffs, such that the best reply to C is sometimes C and sometimes D.

⁵ In our game, unlike the 'additive' Prisoner's Dilemma games described in footnote 2, the sign of $a - b$ imposes no constraints on whether the payoff from playing C against C is greater or less than that from playing D against D.

Since the only information on which a player can condition his strategy choice is private to that player, and given the assumed properties of $f(\cdot)$, any SBNE can be described by the values of two variables. The variable $\beta \in [x_{\min}, x_{\max}]$ is defined such that each player i chooses C if and only if $x_i \geq \beta$. The variable $\pi \in [0, 1]$ is the probability that each player i chooses D, conditional on $x_{\min} < x_i < \beta$. In the non-participation equilibrium, $\beta = x_{\max}$ and $\pi = 0$. In an interior equilibrium, $x_{\min} < \beta < x_{\max}$ and $0 < \pi < 1$. In a boundary equilibrium, $x_{\min} < \beta < x_{\max}$ and $\pi = 1$.

We now analyse these equilibria. Consider any player i facing an opponent whose strategy is described by (β, π) , in an interaction in which $x_i = \beta$. Let V_N, V_D, V_C and V_M be the expected payoffs to player i from playing N, D, C and M respectively, where M is the mix of D with probability π and N with probability $(1-\pi)$. Let $g(x) \equiv F(x)/[1-F(x)]$, where $F(\cdot)$ is the cumulative of $f(\cdot)$. It is straightforward to derive the following expressions:

$$V_N = 0 \tag{1}$$

$$V_D = [1-F(\beta)]a - F(\beta)\pi c \tag{2}$$

$$V_C = [1-F(\beta)]\beta - F(\beta)\pi b \tag{3}$$

$$V_M = \pi V_D. \tag{4}$$

In analysing equilibrium, it is convenient to work in a (β, π) space defined by $x_{\min} \leq \beta \leq x_{\max}$ and $\pi \geq 0$. Notice that this space includes points at which $\pi > 1$. Although such points have no interpretation within our model, equations (1)–(4) above define V_N, V_D, V_C , and V_M for *all* values of π . This allows us to define the loci of points in this (β, π) space at which the mathematical equations $V_N = V_D$ and $V_C = V_M$ are satisfied, and then to characterise equilibria in terms of these loci, imposing the inequality $\pi \leq 1$ as an additional constraint. This method of analysis is useful in simplifying the proofs of our results.

First, consider the locus of points in the (β, π) space at which $V_N = V_D$. Any interior equilibrium must be a point on this *ND locus*, with $0 < \pi < 1$; any boundary equilibrium must be a point at which $V_N \leq V_D$ and $\pi = 1$. By (1) and (2), this locus is determined by:

$$V_D \geq (\text{or } <) V_N \Leftrightarrow a/\pi c \geq (\text{or } <) g(\beta). \tag{5}$$

This is a continuous and downward-sloping curve which includes the point $(x_{\max}, 0)$ and is asymptotic to $\beta = x_{\min}$. It divides the (β, π) space into three regions: the set of points *on* the locus, at which $V_N = V_D$; the set of points *inside* the locus (that is, below and to the left), at which $V_N < V_D$; and the set of points *outside* the locus (that is, above and to the right), at which $V_N > V_D$.

Now consider the locus of points at which $V_C = V_M$. Every equilibrium must be a point on this *CM locus*, with either $\pi = 0$ (the non-participation equilibrium), $0 < \pi < 1$ (an interior equilibrium), or $\pi = 1$ (a boundary equilibrium). Combining equations (2)–(4), this locus is determined by:

$$V_C \geq (\text{or } <) V_M \Leftrightarrow (\beta - \pi a) / [\pi(b - \pi c)] \geq (\text{or } <) g(\beta). \quad (6)$$

This is a continuous curve which includes the points $(x_{\min}, x_{\min}/a)$ and $(x_{\max}, 0)$ and has the property that $\pi < 1$ when $\beta \leq a$. It divides the (β, π) space into three regions: the set of points *on* the locus, at which $V_C = V_M$; the set of points *inside* the locus, at which $V_C > V_M$; and the set of points *outside* the locus, at which $V_M > V_C$.

Propositions (5) and (6) together imply the following result about the relative positions of the two loci:

$$\text{if } V_D = V_N \text{ and } \beta < x_{\max} \text{ and } \pi > 0 \text{ then } V_C \geq (\text{or } <) V_M \Leftrightarrow \beta \geq (\text{or } <) ab/c. \quad (7)$$

The loci intersect at the non-participation equilibrium $(x_{\max}, 0)$. If $x_{\max} \leq ab/c$, there is no other intersection and hence no interior equilibrium.⁶ This case is illustrated in Figure 1a. (The loci are shown by the curves ND and CM; N is the non-participation equilibrium. Sections of these loci which pass above $\pi = 1$ are drawn in grey dots to signify that these points have no interpretation within the model. The arrows refer to the dynamic analysis, which will be presented in Section 4.) If instead $x_{\max} > ab/c$, there is exactly one other intersection, at $\beta = ab/c$. There are now three alternative cases.

[Figure 1 near here]

In the first case, illustrated in Figure 1b, this intersection is at $\pi < 1$. This intersection, denoted I, is an interior equilibrium, defined by $\beta = ab/c$, $\pi = a/g(ab/c)$.⁷ These values of β and π imply that the probability with which C is played, conditional on participation in the game (i.e. conditional on N *not* being played) is $c/(a + c)$, ensuring that $V_D = 0$. (Equivalently, the frequencies with which C and D are played are in the ratio $c : a$.) There may also be boundary equilibria; these occur if the CM locus intersects the line $\pi = 1$ to the left of the ND locus.

⁶ In this case, it is possible that the CM locus intersects the line $\pi = 1$, creating boundary equilibria. However, this would require a high proportion of the probability mass of $f(X)$ to be between a and x_{\max} , despite the fact that $x_{\max} \leq ab/c$.

⁷ The equilibrium value of π can be derived from (2) by using the fact that, in an interior equilibrium, $V_D = 0$.

In the second case, the loci intersect at $\pi > 1$. Because the CM locus is continuous, and because $x_{\min}/a < 1$, there must be at least one value of β in the interval $a < \beta < ab/c$ at which the CM locus intersects the line $\pi = 1$. Any such point is a boundary equilibrium. This case is illustrated in Figure 1c; B is a boundary equilibrium. In the third case (not illustrated), the loci intersect exactly at $\pi = 1$. Then this intersection is a boundary equilibrium. In this case, there may be other boundary equilibria.

The foregoing argument establishes:

Result 1. If $x_{\max} > ab/c$, there is at least one (interior or boundary) equilibrium with $0 < \pi \leq 1$ and $x_{\min} < \beta < x_{\max}$.

In other words, provided the upper tail of the distribution of cooperative benefit is not too short, there is at least one equilibrium in which both C and D are played.

We now consider players' payoffs in such equilibria. Let $V^*(\beta, \pi)$ be the ex ante expected payoff to any player i , prior to the realisations of the random variable X , given that i and his opponent play according to β and π . We will call $V^*(\beta, \pi)$ the *value* of the game conditional on (β, π) .

The following results are derived in the Mathematical Appendix:

Result 2. In every interior and boundary equilibrium, the value of the game is strictly positive.

Result 3. Suppose there are two equilibria, (β, π) , (β', π') , such that $\beta < \beta'$. Then $V^*(\beta, \pi) > V^*(\beta', \pi')$.

Result 2 establishes that in every interior and boundary equilibrium, cooperative activity creates positive net benefits relative to the benchmark of non-participation, despite the nonzero probability of cheating. If there are multiple equilibria, one of these is distinguished by its having the lowest value of β . (Since there can be no more than one interior equilibrium, no two equilibria have the same value of β .) Result 3 establishes that this is the equilibrium at which the value of the game is greatest. We will call this the *highest-value equilibrium*.

3. The model: comparative statics

The frequency of cooperative behaviour that can be sustained in equilibrium depends on the distribution of cooperative benefit X . To keep the exposition simple, we analyse the effect of a

rightward shift from one distribution F to an unambiguously superior distribution G when there is no change in the support $[x_{\min}, x_{\max}]$. That is, for all $x_{\min} < z < x_{\max}$, $G(z) < F(z)$. The values of all other parameters are held constant.

Using (5) it can be shown that if some point (β, π) is on the ND locus for the distribution F , it is inside the corresponding locus for G . Similarly, using (6), if some point (β, π) is on the CM locus for the distribution F , it is inside the corresponding locus for G . Thus, an improvement in the distribution of cooperative benefit moves both loci outwards. Figure 2 illustrates the effects of a shift in the distribution from F (inducing the loci $\text{ND}(F)$ and $\text{CM}(F)$) to G (inducing the loci $\text{ND}(G)$ and $\text{CM}(G)$).

As this diagram shows, if the game has interior equilibria for both distributions, those equilibria have the same value of β , namely ab/c , but the G equilibrium has a higher value of π . Since $G(ab/c) < F(ab/c)$, and since the frequencies with which C and D are played are in the fixed ratio $c : a$, both C and D are played with higher frequency in the G equilibrium than in the F equilibrium. More intuitively, the relationship between cooperation and cheating is analogous to that between prey and predator. If the distribution of cooperative benefit becomes more favourable, a higher frequency of cooperation is induced; but the more cooperation there is, the more cheating can be sustained.

If the game has boundary equilibria for both distributions, the highest-value G equilibrium must be to the left of the highest-value F equilibrium. (This can be seen by considering the effect of an outward shift of the CM locus in Figure 1c.) Thus, the former equilibrium induces a higher frequency of cooperation than the latter.

The following general result is proved in the Appendix:

Result 4: Suppose $x_{\max} > ab/c$ and let F, G be two distributions of X such that G is rightward of F . Then in the highest-value G equilibrium, the frequency of cooperation and the value of the game are both strictly greater than in the highest-value F equilibrium.

Thus, as the distribution of cooperative benefit becomes progressively more favourable, the maximum sustainable frequency of cooperation increases.⁸ Increases in cooperation are associated with increases in cheating until the frequency of non-participation falls to zero.

⁸ This comparative-static property is compatible with evidence that in both human and non-human interaction, the level of cooperation is greater, the higher the payoffs to cooperation (Clutton-Brock, 2002; Capraro et al., 2014).

4. The model: dynamics

We now embed our game in an evolutionary process. We consider a finite population of potential players, sufficiently large to legitimate the use of the law of large numbers. In each of a long series of periods, individuals from this population are randomly and anonymously matched to play the game. Since we are presenting an evolutionary analysis, we do not define any concept of ‘lifetime’ utility that players maximise. Instead, we assume that, at the population level, behaviour gravitates towards whatever pattern of play is *currently* payoff-maximising for individuals, given the behaviour of the population as a whole.

In this Section, we present a simple dynamic analysis that can be represented in the (β, π) space of Figure 1. For the purposes of this analysis, β and π are interpreted as descriptions of the mix of strategies played in the population at any given time: $\beta \in [x_{\min}, x_{\max}]$ is the critical value of X such that C is chosen by any player i if and only if $x_i \geq \beta$, and $\pi \in [0, 1]$ is the relative frequency of D choices among players for whom $x_i < \beta$. We assume that β and π evolve independently. (In a biological application, this is equivalent to assuming that β and π are determined by distinct sets of genes.) Thus, the direction of change of β depends on the relative values of V_C and V_M : β tends to increase (respectively: decrease) if $V_M > V_C$ ($V_M < V_C$). The direction of change of π depends on the relative values of V_D and V_N : π tends to increase (decrease) if $V_D > V_N$ ($V_D < V_N$). This gives the dynamics shown in phase-diagram form in Figure 1.

In the case shown in Figure 1a, the dynamics in the neighbourhood of the non-participation equilibrium (N) are cyclical. At the level of generality at which we are working, it is not possible to determine whether this equilibrium is locally stable, but it has a non-empty zone of attraction, including at least all points at which $\beta = x_{\max}$. In the cases shown in Figures 1b and 1c, the non-participation equilibrium is locally unstable. (In these cases, all paths from points close to N but below the ND locus lead away from N, and must eventually pass through or to the left of the interior equilibrium I. As in the case shown in Figure 1a, N has a non-empty zone of attraction.) However, in states in which almost all players choose non-participation, selection pressure is weak, and so the dynamics shown in the diagrams might work very slowly in the region close to N.

It is clear from Figure 1b that, in the neighbourhood of an interior equilibrium, the dynamics exhibit cyclical or spiralling paths. Described in terms of the evolution of the frequencies of the three strategies N, C and D, these paths are similar to those of the Rock–Paper–Scissors game. (The frequency of cooperation is greatest towards the left of the diagram, where the value of β is low.

From there, evolutionary paths lead towards the top right, where the values of β and π are both high, and the frequency of cheating is greatest. From there, paths lead towards the bottom right, where β is high and π is low, and the frequency of non-participation is greatest. And from there, paths lead back towards the left. If paths spiral outwards, they may lead into the zone of attraction of the non-participation equilibrium.) These paths resemble predator–prey interactions, cheats acting as predators and cooperators as prey.

If the CM locus cuts the line $\pi = 1$ at a point where $\beta < ab/c$, this point is a boundary equilibrium. If, as in the case shown in Figure 1c, points to the left of this equilibrium are outside the locus, the equilibrium is locally stable. Not all boundary equilibria have this property, but whenever the ND and CM loci intersect at $\pi > 1$, there must be at least one locally stable boundary equilibrium.

5. Simulations

In this Section we briefly illustrate some basic features of our theoretical model by means of computer simulations. Using a simple deterministic replicator dynamics, we analyse the evolution of β and π over time, and the associated relative frequencies p_N , p_C and p_D with which strategies N, C and D respectively are played. Further details are provided in the Mathematical Appendix.⁹

In applying replicator dynamics to our game, we cannot treat N-players, C-players and D-players as distinct sub-populations which replicate separately. This is because, in our model, players' decisions about whether or not to cooperate in any given interaction are conditioned on the relevant realisation of the random variable X . Our method is to assume that at any given time t , all players are characterised by the same (β, π) pair. In replicator dynamics, the growth rate of the population fraction using a given strategy is proportional to the difference between the current payoff of that strategy and the weighted average of the current payoffs of all strategies, each strategy being weighted by the relative frequency with which it is played (Taylor and Jonker, 1978). For the purposes of our analysis, we define the current payoff to each strategy as the expected payoff to any player i from choosing that strategy, conditional on the current values of p_N , p_C and p_D and conditional on x_i being equal to the current value of β . At each time t , the rates of change of p_N , p_C and p_D are determined by the replicator equations; these changes are then implemented through

⁹ In a preliminary version of this paper, we presented simulations based on a stochastic evolutionary process in which selection acted directly on β and π . However, as a referee suggested, the main evolutionary properties of our model can be conveyed more simply through replicator dynamics.

changes in β and π . This method allows mathematically tractable simulations while conserving the essential features of the dynamics described theoretically in Section 4.¹⁰

Our simulations use the parameter values $a = 3$, $b = 4$, $x_{\min} = 0$ and $x_{\max} = 9$; the distribution of X over the interval $[x_{\min}, x_{\max}]$ is assumed to be uniform. We investigate the dynamics at different values of c , with particular emphasis on values in the interval $[1.33, 2.4]$ in which interior equilibria occur. (If $c \leq 1.33$, $x_{\max} \leq ab/c$, and so the CM and ND loci intersect only at the non-participation equilibrium; if $c \geq 2.4$, the loci intersect at $\pi \geq 1$.) Intuitively, increases in c (that is, increases in the cost incurred by each player when both cheat) reduce the rewards from cheating and so favour cooperation.

Irrespective of the value of c , the non-participation equilibrium is unstable under replicator dynamics. (Since C is the best response to every non-degenerate mix of C and N, any path starting at a point at which $\pi = 0$ and $x_{\max} > \beta > x_{\min}$ must move away from the non-participation equilibrium along the line $\pi = 0$.) However, there are also paths that converge to the non-participation equilibrium. (Since N is the best response to every non-degenerate mix of N and D, any path starting at a point at which $\beta = x_{\min}$ and $1 > \pi > 0$ must have this property.) The stability properties of interior equilibria are less immediately obvious. It turns out that, as c varies, there are qualitative changes in the dynamics associated with interior equilibria.

Table 2 reports the type and stability properties of the interior equilibrium at different values of c . This equilibrium is unstable for $1.33 < c < 2.3094$ but asymptotically stable for $2.3094 < c < 2.4$. At $c = 2.3094$ there is a *Hopf bifurcation* (see the Mathematical Appendix for details).

Table 2: Type and stability properties of interior equilibria

$c \leq 1.33$	No interior equilibrium
$1.33 < c < 1.66$	Improper node, unstable
$c = 1.66$	Proper node, unstable
$1.66 < c < 2.3094$	Spiral node, unstable
$c = 2.3094$	Hopf bifurcation occurs
$2.3094 < c < 2.4$	Spiral node, asymptotically stable
$c \geq 2.4$	No interior equilibrium

¹⁰ The main difference between the two ways of modelling the dynamics relates to states in which at least one of p_N , p_C and p_D is equal to zero (i.e. where $\beta = x_{\min}$, $\beta = x_{\max}$, $\pi = 0$ or $\pi = 1$). In replicator dynamics, a strategy that is not played in any one period is *never* played, irrespective of its payoffs. Since our main concern in this Section is with interior equilibria, this difference is not particularly significant.

Figure 3 shows the dynamics at different values of c . In Figure 3a, $c = 1$ and the CM and ND loci intersect only at the non-participation equilibrium. Paths that start close to the non-participation equilibrium but below the CM locus initially move away from that equilibrium, but then approach it from above ND locus; all paths converge asymptotically to the $\beta = x_{\max}$ line. As c increases (Figures 3b and 3c) the general picture does not change until the Hopf bifurcation occurs: paths that start close to the interior equilibrium eventually approach the non-participation equilibrium. But after this bifurcation, paths that start below the ND locus spiral in to the interior equilibrium (Figure 3d). The implication is that there is a range of parameter values for which a stable interior equilibrium exists and has a large zone of attraction.

[Figure 3 near here]

6. Discussion

We do not intend to claim that our model represents *the* mechanism that underlies human and animal cooperation. There is no good reason to suppose that cooperation is a single phenomenon with a unified causal explanation. We find it more plausible to view cooperation as a family of loosely-related phenomena which may have multiple causes. We offer our model as a stylised representation of *one* mechanism by which cooperation might emerge and persist.

Our model is unusually robust in that it assumes only materially self-interested motivations and applies to anonymous, well-mixed populations. In claiming this as a merit of the model, we do not deny that individuals sometimes act on pro-social motivations. It has long been known that experimental subjects often cooperate in non-repeated and anonymous Prisoner's Dilemmas (Sally, 1995). Nor do we deny that many recurrent cooperative interactions are between individuals who are known to one another, or that populations of potential cooperators are often structured into clusters of individuals who interact mainly with their neighbours. Each of these factors can contribute to the explanation of cooperation in particular environments. Nevertheless, theories that depend on non-anonymity, or on population structures taking particular forms, have restricted domains of application. And since self-interest is a very common and reliable motivation, models which assume only self-interest can be expected to be particularly robust.

As an illustration of how theories with less robust assumptions can be restricted in their application, we consider the currently much-discussed hypothesis of altruistic punishment (Fehr and Gächter, 2000; Gintis et al., 2005). The key insight is that multilateral cooperation can be sustained

in equilibrium if individuals have low-cost options of punishing one another, and if even a relatively small proportion of individuals have relatively weak preferences for punishing non-cooperators. However, the general effectiveness of this mechanism depends on the cost of punishing being low relative to the harm inflicted, and on the absence of opportunities for punishees to retaliate (Herrmann et al., 2008; Nikiforakis, 2008); and it requires that at least some individuals have non-selfish preferences for punishing. Such preferences might be sustained by *cultural* group selection in hunter-gatherer economies, where groups are small and inter-group warfare is frequent, but these conditions are not typical of the modern world; even among hunter-gatherers, *biological* group selection of altruistic punishment would be frustrated by inter-group gene flow (Boyd et al., 2005). Altruistic punishment should be understood as a mechanism that can sustain cooperation in specific types of environment, not as *the* solution to the problem of explaining cooperation. We claim no more than this for our own model.

We have said that our model is in the same spirit as some recent work by biologists, which finds apparently cooperative behaviour to be directly beneficial to the individual cooperator (see Section 1 above). But, as we now explain, the explanatory principles used by these biologists are not the same as those exhibited in our model.

One of the fundamental features of our model is that the cooperative behaviour it describes is *reciprocally beneficial*. By this, we mean the following. Such cooperation is not simply a unilateral action by one individual which, intentionally or unintentionally, confers benefits on another; it is the *composition* of cooperative actions by two or more individuals, the combined effect of which is to benefit each of them. In other words, each cooperator benefits from his action *only if* this action is reciprocated by one or more other individuals. In the absence of enforceable promises, reciprocally beneficial cooperation requires at least one individual to choose a cooperative action without assurance that others will reciprocate. In our model, any player who chooses to cooperate incurs a risk of loss, which is realised if his opponent cheats. One might think (as we are inclined to do) that reciprocal benefit is a hallmark of genuine, as opposed to apparent, cooperation (see also Sachs et al., 2004; West et al., 2007). In biological models of mutualism, cooperation is not reciprocally beneficial, in the sense we have defined.

In the Snowdrift game, which is often used to model apparently cooperative animal behaviour, cooperation and cheating are best responses to one another. In the original story, two drivers are stuck in the same snowdrift. Both drivers have shovels, and so each can choose whether or not to dig. If either driver digs a way out for his own car, the other can drive out too. Each would rather be the only one to dig than remain stuck. This defines a game with Chicken payoffs;

in a pure-strategy Nash equilibrium, one driver digs and the other free-rides (Sugden, 1986). Such an equilibrium is not a case of reciprocally beneficial behaviour.

Clutton-Brock (2009) offers the Soldier's Dilemma as a model of mutualism in biology. In this game, a patrol of soldiers is ambushed by the enemy. Soldiers who fire back attract incoming fire and increase their chance of being killed. By firing back, however, each individual reduces the probability that the patrol will be overrun. The gain from this may be such that from an individual's perspective there is no dilemma at all: firing back may give the best chance of individual survival, irrespective of what the others do. A biological equivalent to this game (or perhaps to Snowdrift) can be found in the behaviour of certain birds and mammals, such as Arabian babblers and meerkat, which feed in predator-rich environments. Individuals of these species go on sentinel duty once they have fed for long enough to be close to satiation (Clutton Brock et al., 1999). In these games, cooperation is chosen either as a dominant strategy or as a best response to other players' non-cooperation; it is not reciprocally beneficial.

In the story of the Soldier's Dilemma, it would be natural to assume that cooperation would be a dominant strategy only if the number of soldiers in the patrol was relatively small, so that each of them received a significant share of the total benefit created by his own cooperative action. Hauert et al. (2002) present a model which can be understood as a version of the Soldier's Dilemma in which the size of the patrol is endogenous. This is an n -player model of voluntary contributions to a public good, but with an outside option of non-participation. A player who takes the outside option receives a small positive payoff σ with certainty, but forgoes any share in the benefits of the public good. Players who participate can either cooperate (contribute to the public good) or cheat (not contribute). Each cooperator incurs a cost of 1 and creates a benefit of r (where $1 < r < n$ and $r > \sigma + 1$), which is divided equally between all participants. This game has no pure-strategy Nash equilibrium. (If all one's opponents take the outside option, the best response is to cooperate; if they all cooperate, the best response is to cheat; if they all cheat, the best response is the outside option.) There is a unique symmetrical mixed-strategy Nash equilibrium in which the expected payoff to all three strategies is σ . More intuitively, in equilibrium the expected number of participants in each game is sufficiently small that cooperation and cheating are equally profitable. Replicator dynamics have the Rock–Paper–Scissors cyclical pattern.

There are some similarities between Hauert et al.'s model and ours: both models include a non-participation option, and both induce mixed-strategy equilibria with predator–prey characteristics. However, Hauert et al.'s model differs from ours in two significant ways. First, the mechanism that induces cooperation works through variation in the number of participants in the

cooperative activity. For this reason, the model cannot represent cooperative activities which require a fixed number of participants. In particular, it cannot represent activities which inherently involve just two individuals – as, for example, most forms of market exchange do. Second, because the costs and benefits of contributing to the public good are non-stochastic, the expected payoffs to cooperation, cheating and non-participation are equal in equilibrium. Thus, although some cooperative activity takes place in equilibrium, this activity generates no net benefit relative to non-participation: it is not reciprocally beneficial.

We suggest that our analysis provides a stylised but essentially realistic account of a mechanism by which reciprocally beneficial cooperation can emerge and persist in anonymous, well-mixed populations in which strategies are selected according to their material payoffs. Using two simple components – voluntary participation and stochastic payoffs – that have not previously been put together, we have constructed a robust general-purpose model of cooperation.

We are conscious that, for some theoretically-oriented economists, the mechanism we have described may seem rather prosaic. For decades, the Prisoner's Dilemma has been used as the paradigm model of cooperation problems, and the problem of explaining cooperation in that game has been treated as a supreme theoretical challenge. Viewed in that perspective, a modelling strategy which relaxes the assumption that cooperation is always a dominated strategy may seem too easy. But we share the view of Worden and Levin (2007) that many real-world cooperation problems are less intractable than the Prisoner's Dilemma. Neglecting these cases results in an incomplete body of theory and fosters unwarranted pessimism about the possibility of spontaneous cooperation.

Mathematical appendix

Proof of Result 2

Let (β, π) be any interior or boundary equilibrium, and consider any player i . With probability $F(\beta)$, $x_i < \beta$ and i plays N or D. In an interior equilibrium, $V_D = V_N = 0$. In a boundary equilibrium, $V_D \geq V_N = 0$ and N is not played. In either case, i 's expected payoff is equal to V_D and is non-negative. With probability $1 - F(\beta)$, $x_i \geq \beta$ and i plays C. If $x_i = \beta$, i is indifferent between C and D and the expected payoff is again V_D . If $x_i > \beta$, i plays C; his expected payoff (conditional on $x_i > \beta$) exceeds that in the $x_i = \beta$ case by $[1 - F(\beta)](x_i - \beta)$; here $1 - F(\beta)$ represents the probability that i 's opponent plays C. Hence:

$$V^*(\beta, \pi) = V_D + [1 - F(\beta)] E[\max(x_i - \beta, 0)], \quad (\text{A1})$$

where E is the expectation operator. Since $V_D \geq 0$ and $\beta < x_{\max}$, the value of $V^*(\beta, \pi)$ is strictly positive.

Proof of Result 3

If (β, π) and (β', π') are both interior and/or boundary equilibria, Result 3 can be derived from (A1) using the fact that V_D is decreasing in β (an implication of (2)). If (β, π) is the non-participation equilibrium, $V^*(\beta, \pi) = 0$ and so Result 3 follows trivially from Result 2.

Proof of Result 4

Suppose $x_{\max} > ab/c$. Let (β, π) be the highest-value F equilibrium and let (β', π') be the highest-value G equilibrium. There are three possibilities. *Case 1:* (β, π) and (β', π') are both interior equilibria. Then $\beta' = \beta = ab/c$ and $\pi' > \pi$. (This case is illustrated in Figure 2.) Since $G(\beta') < F(\beta)$, the frequency of cooperation is higher in the G equilibrium. Using (A1) and the fact that $V_D = 0$ in every interior equilibrium, it can be shown that the value of the game is strictly greater in the G equilibrium. *Case 2:* (β, π) and (β', π') are both boundary equilibria. Then (because the CM locus for G lies outside the CM locus for F) $\beta' < \beta$ and $\pi' = \pi = 1$. Since $G(\beta') < F(\beta)$, the frequency of cooperation is higher in the G equilibrium. Using (2), it can be shown that V_D is strictly greater in the G equilibrium. Then, using (A1) in relation to the distributions F and G , it can be shown that the value of the game is strictly greater in the G equilibrium. *Case 3:* (β, π) is an interior equilibrium and (β', π') is a boundary equilibrium. Then $\beta' \leq \beta$ and $1 = \pi' > \pi$. Since $G(\beta') < F(\beta)$, the frequency of cooperation is higher in the G equilibrium. In the interior equilibrium, $V_D = 0$. In

the boundary equilibrium, $V_D \geq 0$. Then, using (A1), it can be shown that the value of the game is strictly greater in the G equilibrium.

Mathematical details of replicator-dynamic analysis

As explained in Section 5, our simulations assume that, at any given time t , all players are characterized by the same (β, π) pair. Under this assumption, and given our assumptions about the distribution of X , the frequency of cooperation and cheating in the population can be expressed as

$$p_C(t) = 1 - F(\beta(t)) = \frac{x_{\max} - \beta(t)}{x_{\max}} \quad (\text{A2})$$

and

$$p_D(t) = (1 - p_C(t))\pi(t) = \left(\frac{\beta(t)}{x_{\max}}\right) \cdot \pi(t) \quad (\text{A3})$$

respectively. Let

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & \beta(t) & -b \\ 0 & a & -c \end{bmatrix}$$

be the payoff matrix of the game and let $\mathbf{p}(t) = [p_N(t), p_C(t), p_D(t)]'$ indicate the frequencies of the three strategies at t . As explained in Section 5, the payoff to each strategy at t is defined as the expected payoff from that strategy for a player i for whom $x_i = \beta(t)$. Given this definition, the payoff vector and the average payoff are given by

$$(\mathbf{A} \cdot \mathbf{p})' = [0, \beta(t)p_C(t) - bp_D(t), ap_C(t) - cp_D(t)]$$

and

$$\mathbf{p}'\mathbf{A}\mathbf{p} = p_C(\beta(t)p_C(t) - bp_D(t)) + p_D(ap_C(t) - cp_D(t))$$

respectively. In what follows, we drop t to simplify notation.

A standard replicator equation with the game matrix \mathbf{A} induces the following vector field on the 2-simplex

$$\begin{bmatrix} \dot{p}_N = p_N[-(p_C(\beta p_C - bp_D) + p_D(ap_C - cp_D))] \\ \dot{p}_C = p_C[(\beta p_C - bp_D) - (p_C(\beta p_C - bp_D) + p_D(ap_C - cp_D))] \\ \dot{p}_D = p_D[(ap_C - cp_D) - (p_C(\beta p_C - bp_D) + p_D(ap_C - cp_D))] \end{bmatrix}$$

As $p_N = 1 - p_C - p_D$, we will consider the simplified dynamical system

$$\begin{bmatrix} \dot{p}_C \\ \dot{p}_D \end{bmatrix} = \begin{bmatrix} F(p_C, p_D) \\ G(p_C, p_D) \end{bmatrix} = \begin{bmatrix} p_C[(\beta p_C - bp_D) - (p_C(\beta p_C - bp_D) + p_D(ap_C - cp_D))] \\ p_D[(ap_C - cp_D) - (p_C(\beta p_C - bp_D) + p_D(ap_C - cp_D))] \end{bmatrix}.$$

This system has both interior and boundary fixed points, depending on the parameter values. Here we consider only the interior equilibrium. By (2) and (3) in the main text, this equilibrium is characterized by $\beta^* = \frac{ab}{c}$ and $\pi^* = \frac{c \cdot x_{\max} - a \cdot b}{b \cdot c}$.

By (A2) and (A3), the equilibrium frequencies of C and D are given by $p_C^* = \frac{x_{\max} - \beta^*}{x_{\max}}$ and $p_D^* = \left(\frac{\beta^*}{x_{\max}}\right) \cdot \pi^*$.

The Jacobian, evaluated at the interior equilibrium, is

$$\mathbf{J} = \begin{bmatrix} \frac{-63c^2 - 216c + 176}{3c^3} & \frac{27c^2 - 120c + 112}{9c^2} \\ \frac{36c^3 - 183c^2 + 312c - 176}{c^4} & -\frac{9c^3 - 75c^2 + 168c - 112}{3c^3} \end{bmatrix}$$

with eigenvalues

$$e_{1,2} = \frac{-3c^2 + 16}{6c^3} \pm \frac{(3c - 4)(\sqrt{9c^4 + 180c^3 - 768c^2 + 576c + 256})}{6c^3}$$

and trace

$$\text{TR} = -\frac{63c^2 - 216c + 176}{3c^3} - \frac{9c^3 - 75c^2 + 168c - 112}{3c^3}.$$

At $c = 2.3094$, the trace of the Jacobian matrix \mathbf{J} equals zero and a pair of imaginary eigenvalues, $e_{1,2}(I) = \pm i0.2402$, cross the imaginary axis. This is a degenerate Hopf bifurcation, i.e. the resulting limit cycle is unstable, as is the case for any bifurcation of this type in a two-player 3×3 symmetric game under (continuous-time) replicator dynamics (Zeeman, 1980; on this, see also Ochea, 2010, 2013).

References

- Binmore, K., 1994. *Game Theory and the Social Contract, Volume 1: Playing Fair*. MIT Press, Cambridge, MA.
- Binmore, K., 1998. *Game Theory and the Social Contract, Volume 2: Just Playing*. MIT Press, Cambridge, MA.
- Bowles, S. and Gintis, H., 2011. *A Cooperative Species: Human Reciprocity and its Evolution*. Princeton University Press, Princeton, NJ.
- Boyd, R., Gintis, H., Bowles, S. and Richerson, P.J., 2005. The evolution of altruistic punishment. In Gintis et al., 2005, pp. 215-227.
- Capraro, V., Jordan, J.J., and Rand, D.G., 2014. Heuristics guide the implementation of social preferences in one-shot Prisoner's Dilemma experiments. *Scientific Reports* 4, article no. 6790, 1–5.
- Clutton-Brock, T.H., 2002. Breeding together: kin selection and mutualism in cooperative vertebrates. *Science* 296, 69-72.
- Clutton-Brock, T.H., 2009. Cooperation between non-kin in animal societies. *Nature* 462, 51-57.
- Clutton-Brock, T. H., O'Rian, M.J., Brotherton, P.N.M., Gaynor, D., Kansky, R., Griffin, A.S., Manser, M., 1999. Selfish sentinels in cooperative mammals. *Science* 284, 1640-1644.
- Darwin, C., 1859. *On the Origin of Species by means of Natural Selection*. John Murray, London.
- Fehr, E., Gächter, S., 2000. Cooperation and punishment. *American Economic Review* 90(4), 980-994.
- Gintis, H., Bowles, S., Boyd, R., Fehr, E. (eds), 2005. *Moral Sentiments and Material Interests*. MIT Press, Cambridge, MA.
- Hauert, C., De Monte, S., Hofbauer, J., Sigmund, K., 2002. Volunteering as Red Queen mechanism for cooperation in public goods game. *Science* 296, 1129-1132.
- Herrmann, B., Thöni, C., Gächter, S., 2008. Antisocial punishment across societies. *Science* 319, 1362-1367.
- Jost, C., Arditi, R., 2000. Identifying predator-prey processes from time-series. *Theoretical Population Biology* 57, 325-337.
- Leimar, O., Hammerstein, P., 2001. Evolution of cooperation through indirect reciprocity. *Proceedings of the Royal Society* 268, 745-753.

- Maynard Smith, J., Parker, G., 1976. The logic of asymmetric contests. *Animal Behaviour* 24: 159-175.
- Nikiforakis, N., 2008. Punishment and counter-punishment in public good games: can we really govern ourselves? *Journal of Public Economics* 92, 91-112.
- Nowak, M.A., 2006. Five rules for the evolution of cooperation. *Science* 314, 1560-1563.
- Nowak, M.A., Sigmund, K., 2005. Evolution of indirect reciprocity. *Nature* 437, 1291-1298.
- Ochea M.I., 2010. *Essays on Nonlinear Evolutionary Game Dynamics*. PhD thesis. Amsterdam School of Economics Research Institute (ASE-RI).
- Ochea M.I., 2013. Evolution of repeated prisoner's dilemma play under logit dynamics. *Journal of Economic Dynamics & Control* 37, 2483-2499.
- Rousseau, J.-J., 1755/ 1998. Discourse on the origin and foundations of inequality among men. In Ritter, A., Bondanella, J.C. (eds), 1998, *Rousseau's Political Writings*. Norton, New York, 3-57.
- Sachs, J.L., Mueller, U.G., Wilcox, T.P., Bull, J.J., 2004. The evolution of cooperation. *Quarterly Review of Biology* 79 (2), 135-160.
- Sally, D., 1995. Conversation and cooperation in social dilemmas: a meta-analysis of experiments from 1958 to 1992. *Rationality and Society* 7, 58-92.
- Smith, A., 1763; 1978. *Lectures on Jurisprudence*. Oxford University Press, Oxford.
- Sugden, R., 1986. *The Economics of Rights, Cooperation and Welfare*. Blackwell, Oxford. (Second edition 2004, Palgrave-Macmillan, Basingstoke.)
- Taylor, P.D. and Jonker, L.B., 1978. Evolutionarily stable strategies and game dynamics. *Mathematical Biosciences* 40, 145-156.
- Tomasello, M., 2014. *A Natural History of Human Thinking*. Harvard University Press, Cambridge, MA.
- Worden, L., Levin, S.A., 2007. Evolutionary escape from the prisoner's dilemma. *Journal of Theoretical Biology* 245, 411-422.
- Zeeman, E.C., 1980. Population dynamics from game theory. In Nitecki, Z. and Robinson C. (eds), *Global Theory of Dynamical Systems*. New York: Springer, 471-497.

Figure 1. Equilibria and dynamics

Figure 1a: Non-participation the only equilibrium

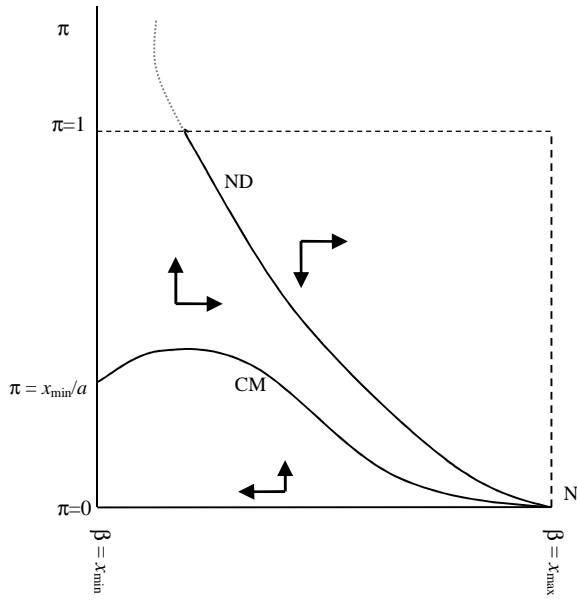


Figure 1b: An interior equilibrium

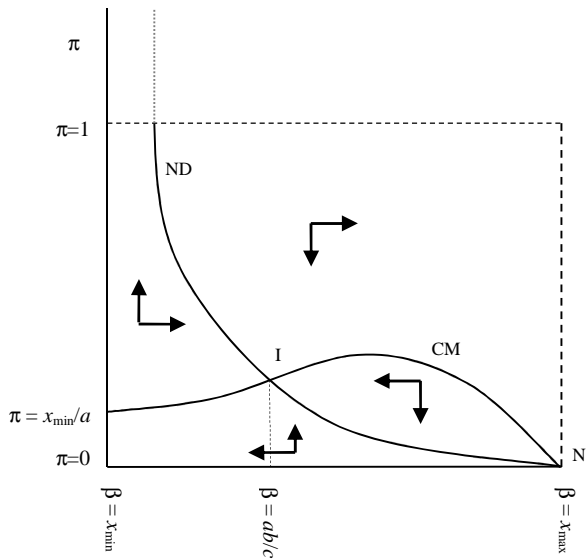


Figure 1c: A boundary equilibrium

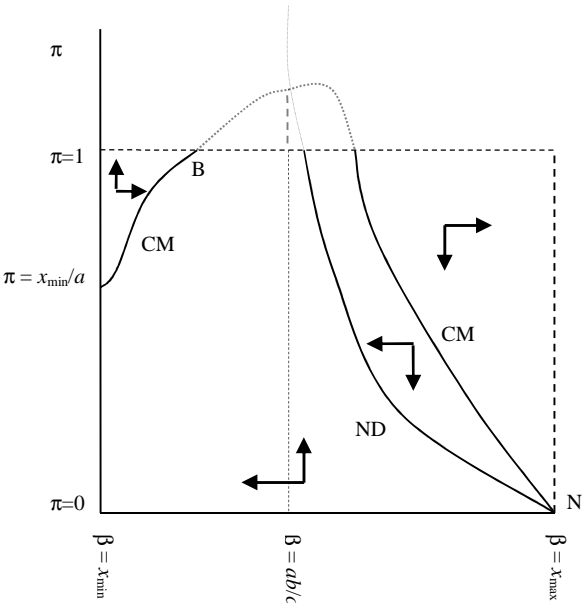


Figure 2: Effects of a shift in the distribution of cooperative benefit

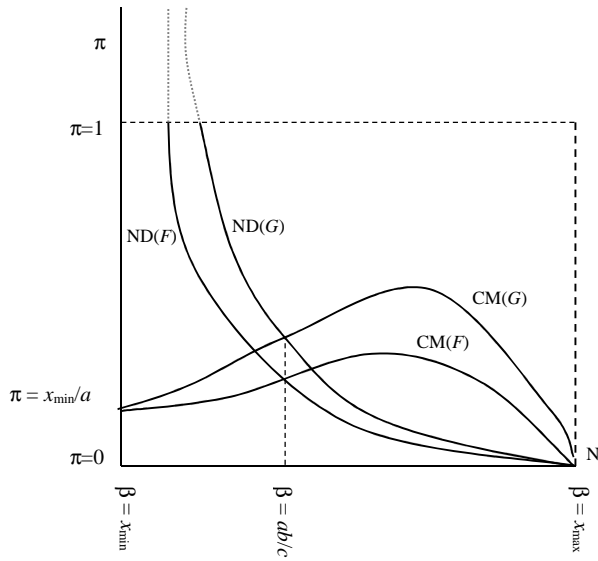


Figure 3: Replicator dynamics for different values of c

Figure 3a. Initial point: $\beta=8.9, \pi=0.01, c=1$.

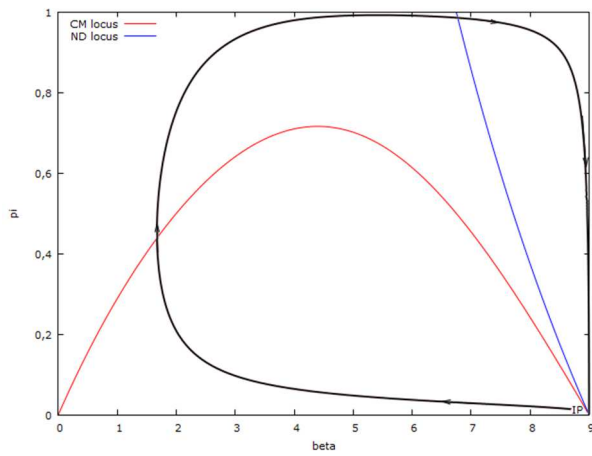


Figure 3b. Initial point: $\beta=6, \pi=0.74, c=1.6$.

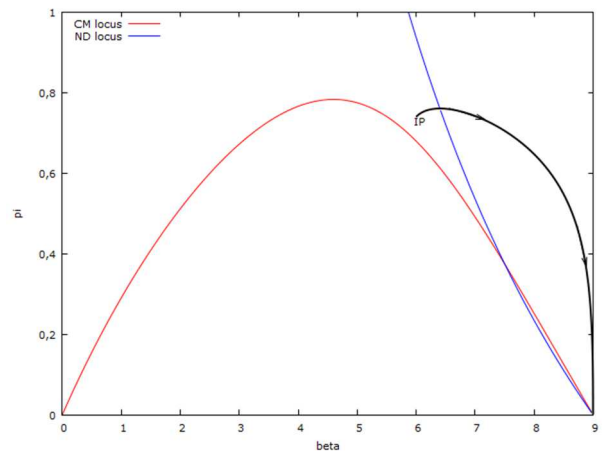


Figure 3c. Initial point: $\beta=6, \pi=0.74, c=2$.

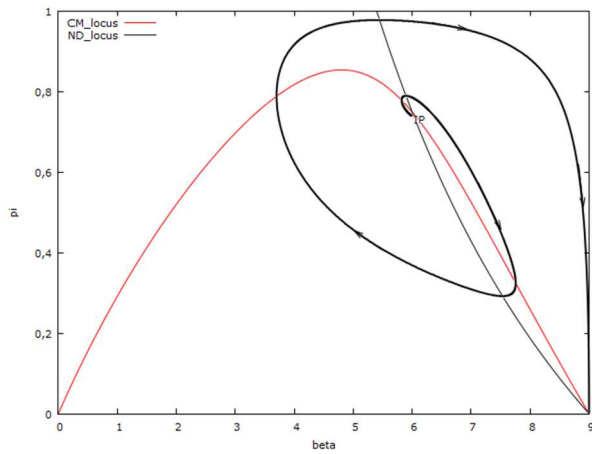


Figure 3d. Initial point: $\beta=8.9, \pi=0.01, c=2.35$.

