Choosing Opponents in Games of Cooperation and Coordination

Peter Engseld[*] Andreas Bergh[†]

March 5, 2005

Abstract

We analyze a cooperation game and a coordination game in an evolutionary environment. Agents make noisy observations of opponent's propensity to play dove, called reputation, and form preferences over opponents based on their reputation. A game takes place when two agents agree to play. Socially optimal cooperation is evolutionarily stable when reputation perfectly reflects propensity to cooperate. With some reputation noise, there will be at least some cooperation. Individual concern for reputation results in a seemingly altruistic behavior. The degree of cooperation is decreasing in anonymity. If reputation is noisy enough, there is no cooperation in equilibrium. In the coordination game, the efficient equilibrium is chosen and agents with better skills to observe reputation earn more.
JEL classification: C70; C72

Keywords: Cooperation; Coordination; Conditioned Strategies; Prisoners Dilemma; Signaling; Reputation; Altruism; Evolutionary Equilibrium

1 Introduction

We analyze a cooperation game and a coordination game in an evolutionary environment where the matching of opponents is based on agents' preferences over possible opponents. Agents choose opponents based on a signal (interpreted as reputation) which is a noisy observation of each agent's true propensity to take a certain action. In this setting we find that reputational concerns can explain seemingly altruistic behavior. In the cooperation game, we show that if reputation perfectly reflects each agent's true propensity to cooperate, then the payoff

[*]Dept of Economics, Lund University, P.O. Box 7082, 220 07 Lund, Sweden. Email: peter.engseld@nek.lu.se, Tel. +46-46-222 79 13, Fax: +46-46-222 46 13.
[†]Dept of Economics, Lund University, P.O. Box 7082, 220 07 Lund, Sweden. Email: andreas.bergh@nek.lu.se, Tel. +46-46-222 48 10, Fax: +46-46-222 46 13.

maximizing strategy in evolutionarily stable populations, when mutual cooperation is socially optimal, is to cooperate and choose to play with cooperative agents. When reputation is noisy the degree of cooperation is decreasing in reputation noise. Only when the degree of noise is sufficiently high, all cooperation ceases to exist.

In the coordination game, we show that if reputation sufficiently accurately reflects the actual behavior, the efficient outcome is a unique stable equilibrium. Since the payoff is strictly increasing in the agents ability to observe the true propensity, evolutionary forces will favor prosocial behavior in the long run.

The literature on the problem of cooperation is huge and spans several disciplines, see e.g. Hammerstein (2003). Following Axelrod (1984), the tit for tat strategy became widely known as the best strategy in repeated games. However, as pointed out by Boyd & Lorberbaum (1987), a population of cooperating tit-for-tats can be invaded by nice but less retaliatory strategies, resulting in a population vulnerable to invasion by defecting strategies.

Other important theories regarding the emergence of cooperation are based on genetic relations (kin selection) and group selection. However, these approaches can not explain why cooperation arises even among genetically unrelated agents who meet only once. Theories of group selection must also explain the absence of free riding within the group. Explanations based on punishment of those who violate the cooperative norms have been suggested or discussed by for example Ostrom et al. (1992), and Boyd & Richerson (1992). This type of explanation only raises another question: Why would agents engage in costly punishment of non-cooperators?

It is our belief that despite having many merits, existing theories do not sufficiently explain under what circumstances agents cooperate at a personal cost with unrelated strangers in one-shot interactions. We argue that a crucial part of the strategy in repeated games is the choice of opponent. To substantiate our argument, we also apply reputation based choice to a classical coordination game previously analyzed by e.g. Kandori et al. (1993) and Young (1993). Our model produces the same results for this game.

2 The model

Consider a population I with a large even number N agents who are repeatedly matched to play a symmetric 2×2 game. The figures below show the game of cooperation (left) and the game of coordination (right).

	H	D
H	0, 0	1, β
D	β, 1	α, α

$\Gamma_1(\alpha, \beta)$

The cooperation game

	H	D
H	γ, γ	0, 0
D	0, 0	δ, δ

$\Gamma_2(\gamma, \delta)$

The coordination game

In the cooperation game, $\Gamma_1(\alpha, \beta)$, we assume that $\alpha \in (0, 1)$ and $\beta \in (-\infty, 0)$. This will include the Prisoner's Dilemma as the special case when co-operation is socially optimal $(2\alpha > 1 + \beta)$. In the coordination game, $\Gamma_2(\gamma, \delta)$, we assume that $\delta > \gamma > 0$. The action set is $A \equiv \{H, D\}$, where $a \in A$. Actions are taken in discrete time, $t \in \mathbb{Z}$. We think of the actions H and D as playing Hawk (hard, defect) and Dove (nice, cooperate) respectively.

We use the term propensity as a measure of inclination to D, formally defined as follows:

Definition 1 *The propensity P_i^t of agent $i \in I$ at time t is a recursive function P_i^t, where $P_i^t \equiv \rho \Pr(a_i^t = D \mid \cdot) + (1 - \rho) P_i^{t-1}$ and $\rho \in (0, 1)$.*

The propensity in period t is defined as a weighted average of the probability to play D in period t and the propensity in the previous period.

It is unlikely that agents are able to perfectly observe the propensity of other agents, therefore we assume that every agent $i \in I$ observes the *reputation* r_{-i}^{t} of an opponent $-i$, which is a realization of the stochastic variable $\mathbf{R}_{-i}^{t}$. $\mathbf{R}_{-i}^{t}$ is symmetrically and unimodally distributed around P_{-i}^{t}, see figure below. The value of r_{-i}^{t} is private information for i.

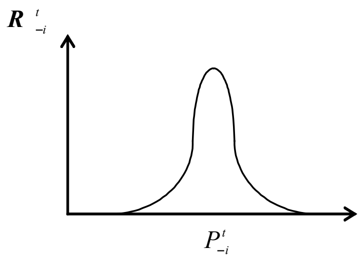

Let $\underline{\mathbf{R}}_{i,I}$ be the stochastic vector of all agents in I ordered according to reputation by agent i. Denote the vector of all agents in I ordered decreasingly according to propensity P_I, where P_I^k denotes the agent with k highest propensity. Let $\underline{R}_{i,I}$ be a realization of $\underline{\mathbf{R}}_{i,I}$. Let the observational skills, denoted $O_i \in \mathbb{R}_+, \forall i \in I$, satisfy:

Assumption 1 $\frac{\partial \operatorname{corr}\left[\underline{\mathbf{R}}_{i,I}, P_I\right]}{\partial O_i} > 0$, *and* $O_i = \infty \Rightarrow \underline{R}_{i,I} = P_I, \forall i \in I$.

By observing the reputation of potential opponents, agents form preferences over opponents. Let $\mathbf{\Psi}$ denote the set of all complete and transitive orderings of the agents in I. The preferences for agent i will be denoted $\succsim_i$.

Assumption 2 $\succsim_i \in \mathbf{\Psi}, \forall i \in I$.

Let $\mathbf{S}$ denote the set of all pure strategies. A pure strategy $s \in \mathbf{S}$ is a mapping from own propensity, and the reputation of the rest of the population onto A and $\mathbf{\Psi}$. More formally we have $s : [0,1]^N \mapsto A \times \mathbf{\Psi}$. Denote the set of strategies where the action is conditioned on the reputation of the opponent $\mathbf{S}^c \subset \mathbf{S}$.

2.1 Matching of the Agents

The games Γ_1 and Γ_2 are played repeatedly by the agents in the population. The matching of the agents is conducted through a procedure defined as follows. Let $C_i \equiv \left(C_i^1, C_i^2, ..., C_i^N\right)$ where C_i^k denotes agent i's k-preferred choice. Thus, C_i^1 denotes i's most preferred opponent, C_i^2 her second best, and so on. Agent i will be matched against j if $C_i^1 = j$ and $C_j^1 = i$.

The procedure makes use of a randomized choosing order, assumed (without loss of generality) to coincide with the numbers 1 to N. First, agent 1 asks her most preferred opponent, who accepts if agent 1 is her most preferred opponent. Then agent 2 asks her most preferred opponent, and when all agents have proposed to their first best choice, the procedure is repeated for second best choices. The procedure continues until all agents are paired. Formally, the matching procedure can be described by the following algorithm, which orders all agents in I into matched pairs in the set of matched pairs $\mathbb{I}$.

Algorithm 1 (Matching procedure) *Let $\mathbb{I}$ be the set of matched pairs.*

Step 0. *Let $\mathbb{I} = \emptyset$, $i = 1$, and $l = 1$.*

Step 1. *If there exists an $m \in [1, l]$ such that if $\left(C_i^l = j\right) \cap \left(C_j^m = i\right) \cap (i, j \notin \mathbb{I})$, then $(i, j) \in \mathbb{I}$.*

Step 2. *Increase i by 1. If $i \leq N$, go to step 1.*

Step 3. *Increase l by 1 and let $i = 1$. If $l \leq N$, go to step 1.*

This procedure is assumed to be repeated an infinite number of times within each period. This ensures that the realized payoff for every agent at each period is equal to the expected payoff. This could be interpreted as inertia in the observations of the opponents' reputation such that the agents react "sluggishly" on any changes in strategies. The probability for agent i to play D within each procedure at each t is equal to $\Pr\left(a_i^t = D \mid \cdot\right)$.

For technical reasons, we allow agents to be matched up with themselves. This can be seen as if all agents appear in pairs or that an agent sooner or later will have an identical offspring. Finally, note that when observational skills are non-existent, this matching procedure is equivalent to random matching.

2.2 Evolutionary Stability

To analyze evolutionary stability in the two games, we assume that every agent in the population in each period with a small probability will change to another mutant strategy. In each time period there is never more than one mutation. The question whether a mutant strategy could invade the current incumbent strategy distribution becomes complicated by the fact that if an agent switches strategy, then the propensity (in almost all cases) will change which in turn could trigger new response actions from its opponents. For this reason, we need to modify the concept of evolutionary stability.

A mixed strategy is denoted σ and is defined as a probability distribution over $\mathbf{S}$. Formally, $\sigma \equiv (\sigma_s)_{s\in\mathbf{S}}$, $\sigma_s \in [0,1], \forall s \in \mathbf{S}$ and $\int_{s\in\mathbf{S}} \sigma_s = 1$. Any mixed strategy can consequently be seen as a vector $\sigma \in \mathbb{R}_+^\infty$, that belongs to the unit simplex Δ_σ, where

$$\Delta_\sigma \equiv \left\{ \sigma \in \mathbb{R}_+^\infty \mid \int_{s\in\mathbf{S}} \sigma_s = 1 \right\}.$$

Let $q_\sigma^t \in [0,1]$ denote the fraction of individuals in I at period t with strategy σ. Thus it is possible to characterize any combination of mixed strategies in the population as a point Q_I in the simplex denoted Δ_Q:

$$\Delta_Q \equiv \left\{ Q_I \in \mathbb{R}_+^\infty \mid Q_I = (q_\sigma)_{\sigma\in\Delta_\sigma}, \int_{\sigma\in\Delta_\sigma} q_\sigma = 1 \right\}.$$

Note that Q_I both can be viewed as a point in $\mathbb{R}_+^\infty$ and as a set of strategies. We use $Q_{I\setminus i}$ to denote the strategy mix in the population $I \setminus i$. In the same manner, let $O_{I\setminus i}$ denote the observational skills of all agents $I \setminus i$. The expected payoff for agent i with strategy σ_i at period t will be $\pi_i^t\left(\sigma_i, O_i; Q_{I\setminus i}, O_{I\setminus i}\right)$. When there is no risk of confusion, let $Q = Q_I$. We use the following evolutionary property:

$$\text{sign}\left(\frac{q_\sigma^{t+1}}{q_\sigma^t} - \frac{q_{\sigma'}^{t+1}}{q_{\sigma'}^t}\right) = \text{sign}\left(\pi^t\left(\sigma, \cdot; \cdot\right) - \pi^t\left(\sigma', \cdot; \cdot\right)\right). \quad (1)$$

This specification means that strategies with higher payoffs will have a higher representation in the population in the next period. The offsprings are assumed to inherit both strategy and propensity from the parent.

Note that the payoff of an agent i depends both on the agent's strategy σ_i and on the opponent's strategy σ_{-i}. Also, the opponent's actions can depend on the agent's reputation $r_{-i,i}$, just as the agent's action can depend on the opponent's reputation r_{-i}. This implies that if an agent changes strategy, her actions and thus her propensity can change, which could trigger different actions from other agents and thereby change their propensity, which in turn might lead to other agents changing their actions ad infinitum.

The following assumption help us avoid such cumbersome dynamic.

Assumption 3 *The adjustment process of the propensity is much faster than the growth/learning process, which in turn is much faster than the process of mutations.*

This will imply that the population on average can be considered stationary in so far as the pair σ_i, P_i is fixed $\forall i \in I$. This also renders the index for time redundant in most cases.

Another consequence of assumption 3 is that there exists a well-defined correspondence between the propensity and the strategies: For all strategy mixes $Q \in \Delta_Q$, there exists a corresponding propensity distribution.

Definition 2 $\overrightarrow{QQ'}$ *denotes a sequence* $\left\{Q^j\right\}_{j=1}^M$ *of adjacent points in* Δ_Q *between* $Q = Q^1$ *and* $Q' = Q^M$, *where*

- Q^j *can evolve to* Q^{j+1} *through growth, or*

- $\exists \sigma_i^{'}$ *such that* $\pi_i\left(\sigma_i, \cdot; Q_{I\backslash i}, \cdot\right) \leq \pi_i\left(\sigma_i^{'}, \cdot; Q_{I\backslash i}, \cdot\right)$ *for some* $i \in I$, *where* $Q_{I\backslash i} \cup \sigma_i = Q^j$ *and* $Q_{I\backslash i} \cup \sigma_i^{\prime} = Q^{j+1}, \forall j \in (1...M-1)$.

$\overrightarrow{QQ'}$ implies that there exists a path where the strategy mix in the population evolves from Q to Q' either through growth or through that agents, one at the time, change strategy to another with at least as high payoff.

Let us now define a *Mutation Proof Attraction Set* (*MAS*), which basically is a modified *Absorbing set* (see e.g. Samuelson (1998)), where the set is closed under the growth mechanism and mutations.

Definition 3 (*MAS*) $\mathbf{Q}^{MAS}(\Gamma)$ *is a set of strategy mixes* $Q \in \mathbf{Q}^{MAS}(\Gamma)$ *such that*

- $\exists\ \overrightarrow{QQ'},\ \forall Q' \in \mathbf{Q}^{MAS}(\Gamma)$,
- $\exists\ \overrightarrow{Q'Q},\ \forall Q' \in \mathbf{Q}^{MAS}(\Gamma)$, *and*
- $\nexists\ \overrightarrow{QQ''}$ *for any* $Q'' \notin \mathbf{Q}^{MAS}(\Gamma)$.

Let $\Delta^{MAS}(\Gamma) \equiv \bigcup \mathbf{Q}^{MAS}(\Gamma)$.

Property 1 $\Delta^{MAS}(\Gamma) \neq \emptyset, \forall \Gamma$.

A population I belongs to a *MAS*, precisely if the strategy mix Q in the population belongs to an attraction set $\mathbf{Q}^{MAS}(\Gamma)$ such that $\exists\ \overrightarrow{QQ'},\ \forall Q' \in \mathbf{Q}^{MAS}(\Gamma)$ and $\exists\ \overrightarrow{Q'Q},\ \forall Q' \in \mathbf{Q}^{MAS}(\Gamma)$. That is, each combination of strategies in the population that belongs to the attraction set $\mathbf{Q}^{MAS}(\Gamma)$ must be able to evolve to any other point in the attraction set through growth and/or individual changes to strategies that yields at least as high payoff, i.e. through drift. Moreover, there must not exist any feasible path such that the population could evolve to a point $Q'' \notin \mathbf{Q}^{MAS}(\Gamma)$. Note that *MAS* yields identical equilibria on unconditioned strategies as *NSS* (see Maynard Smith (1982)).

An alternative way to view MAS is through a phase diagram. Each point in Δ_Q is associated with a directional vector whereto the strategy mix Q can evolve through growth as expressed in equation 1, and/or through mutations (drift). An absorbing set is a vector field where there does not exist any directional vectors from the the growth mechanism pointing outside the set. A MAS is consequently an absorbing set with a hull such that the combined set is closed under both the growth mechanism and mutations. That is, MAS is an absorbing set that will prevail minor perpetrations such as mutations.

3 Evaluating the Cooperation Game

Let us first consider $\Gamma_1(\alpha, \beta)$. The payoff is strictly increasing in the propensity of the opponent.

Property 2 $\frac{\partial \pi_i(\Gamma_1)}{\partial P_{-i}} > 0$.

Regardless of whether an agents intends to play H or D, she will earn more if the opponent is more likely to play D. From the definition of reputation we know that the expected value of the reputation equals to the propensity.

Property 3 $\frac{\partial \pi_i(\Gamma_1)}{\partial r_{i,-i}} > 0$.

Denote by $\mathcal{D}$ the set of preferences such that the agent prefers opponents with higher reputation:

Definition 4 $\mathcal{D} \equiv \{\succsim_i \mid r_{i,j} \geq r_{i,k} \Leftrightarrow j \succsim_i k, \forall j,k \in I\}$.

$\mathcal{D}$ implies that agent i prefers agent j as an opponent over agent k if and only if i perceives the reputation of j to be higher than that of k.

3.1 Perfect Observational Skills

Let us begin by analyzing the model when reputation perfectly reflects propensity: $O_i = \infty, \forall i \in I$. This implies that the reputation is identical to the propensity. From property 2 we know that payoff is strictly increasing in the opponent's propensity. Since all agents will be able to avoid opponents with lower propensity, agents will only be matched up against opponents with identical propensity.

Definition 5 *Let* $z_i \equiv \Pr(a_i = D)$.

Since all agents have identical propensity, $z_i = z_{-i} \equiv z$. The payoff for an arbitrary agent i is: $\pi_i(\cdot) = \alpha z^2 + z(1-z)(1+\beta)$, which is maximized for

$$z = \frac{1}{2} + \frac{1}{2}\frac{\alpha}{1+\beta-\alpha}. \quad (2)$$

When cooperation is socially optimal, $2\alpha \geq 1+\beta$, this implies $z = 1$. If $2\alpha < 1+\beta$, the payoff is maximized for mixed strategies where the probability of playing D is equal to $\frac{1}{2} + \frac{1}{2}\frac{\alpha}{1+\beta-\alpha} > \frac{1}{2}$. Thus, the probability of playing action D is strictly increasing in α, and always higher than 50 percent.

Let P^* denote the propensity associated with the payoff maximizing strategies when all agents have identical propensity. Thus, when $2\alpha \geq 1+\beta$ we have that $P^* = 1$. When $2\alpha < 1+\beta$, P^* is associated with strategies such that $z = \frac{1}{2} + \frac{1}{2}\frac{\alpha}{1+\beta-\alpha} > \frac{1}{2}$.

However, when all agents have identical propensity, the population will be vulnerable to a neutral invasion of agents with the same propensity as the incumbents, but with $\succsim_i \neq \mathcal{D}$.

Lemma 1 $P_i = P_j, \forall i,j \in I \Rightarrow \pi_i(\mathcal{D}) = \pi_i(\succsim_i \neq \mathcal{D}), \forall i \in I$.

In other words, the population will through mutations drift away from all agents preferring opponents with higher reputation.

Proposition 1 *If* $Q \in \Delta^{MAS}(\Gamma_1)$ *and* $O_i = \infty, \forall i \in I$ *then*

- $2\alpha \geq 1+\beta \Rightarrow \lim_{t\to\infty} \Pr(a_i = D) = 1$, *and*

- $2\alpha < 1 + \beta \Rightarrow \lim_{t\to\infty} \Pr(a_i = D) = \frac{1}{2} + \frac{1}{2}\frac{\alpha}{1+\beta-\alpha}, \forall i \in I.$

The intuition is as follows: Since agents with $\succsim_i \neq \mathcal{D}$ only can make a neutral invasion, they will only represent a small fraction of I. This fraction can be exploited by strategies more inclined to play H, with preferences $\mathcal{D}$. The later category will initially yield more than all other strategies in the population. However, strategies with $\succsim_i \neq \mathcal{D}$ will yield less than all other strategies and thus also grow slower. This implies that agents with low propensity strategies to a higher degree will become matched up themselves, and thus earn a lower payoff. This process will eventually stabilize when the expected payoff for agents with $\succsim_i \neq \mathcal{D}$ equals that of agents with low propensity strategies.

3.2 Imperfect Observational Skills

Assume now that reputation is noisy and only imperfectly reflects propensity.

Lemma 2 *$O_i < \infty, \forall i \in I$, and $\exists P_i \neq P_j$, for some $i, j \in I \Rightarrow \pi(\sigma \mid \mathcal{D}) > \pi(\sigma \mid \succsim \neq \mathcal{D}), \forall \sigma \in Q$.*

When the population contains agents with different propensity, preferences $\mathcal{D}$ will yield payoff.

Consider an incumbent strategy σ and a mutant strategy σ' with corresponding propensities P and P', such that $P' < P \leq P^*$. Let $d_{P,P'} \equiv |P - P'|$. Let $x_{P,P'}$ be the probability that an arbitrary agent with observational skills O perceives a potential opponent with propensity P to have a higher reputation than a potential opponent with propensity P', formally $x_{P,P'} \equiv \Pr(r \geq r' : P \geq P', O)$. From the definition of reputation we have the following general property:

Property 4 *$\frac{\partial x_{P,P'}}{\partial d_{P,P'}} > 0$, and $\frac{\partial x_{P,P'}}{\partial O} > 0$.*

Now consider the payoffs $\pi(\sigma)$ and $\pi(\sigma')$. For any population size N and for each pair P, P' there exists a minimum $x_{P,P'}$ for which σ has a higher payoff than σ'. Denote this value $x^{\min}_{P,P'}$. When $d_{P,P'}$ is smaller, O must be higher to ensure that the incumbent strategy σ yields more:

Lemma 3 $\lim_{P' \to P} x^{\min}_{P,P'} = 1, \forall P \in (0,1]$.

Corollary 1 *For any given level of observational skills, there exists a population size N^* such that all populations larger than N^* can be invaded.*

In other words, unless observational skills are perfect, a sufficiently large population can always be invaded by agents with $P' < P$. Hence, from lemma 2 it follows that:

Corollary 2 $O_i < \infty, \forall i \in I,$ *and* $Q \in \Delta^{MAS}(\Gamma_1) \Rightarrow \succsim_i = \mathcal{D}, \forall i \in I.$

Less than perfect observational skills are sufficent for incumbents to earn more than mutants with $P' = 0$, given that incumbents have sufficiently high propensity, $P \geq \widehat{P}$:

Lemma 4 $x^{\min}_{P,0} < 1, \forall P \in \left(\widehat{P}, 1\right]$.

From Lemma 3 and 4, we can draw a qualitative figure over $x^{\min}_{1,P}$ as in the left hand figure below. Assume that $2\alpha > 1 + \beta$, which implies that $P^* = 1$. The exact shape of $x^{\min}_{1,P}$ depends on the agents' observational skills, and the payoffs α and β. If the propensities of two agents are arbitrarily close, it will be impossible to discriminate them. If reputation symmetrically and unimodally distributed around P, we expect the probability $x_{P_i,P_{-i}}$ to be as depicted by the dotted line in the left-hand figure below. Since we know from lemma 3 that $\lim_{P' \to P} x^{\min}_{P,P'} = 1$, it follows that there exists an interval immediately below P^* where invasions are successful. Denote this interval $\widetilde{P}^{-}_{P^0}$ and let $P^0 \equiv P^*$.

Consider the case when $\left|\widetilde{P}^{-}_{P}\right| < P$, i.e. when there exists an interval where invasions are not possible. Assume there is a complete invasion at $P^1 \in \widetilde{P}^{-}_{P^0}$. This implies that the entire population will have the propensity P^1. Now there exists an interval $\widetilde{P}^{-}_{P1}$ where invasions will be successful. Envision a complete invasion at $P^2 \in \widetilde{P}^{-}_{P^1}$. If $P^2 < P^0 - \left|\widetilde{P}^{-}_{P^0}\right|$ then, in addition to the interval with lower propensity $\widetilde{P}^{-}_{P^2}$, there exists an interval $\widetilde{P}^{+}_{P^2}$, where agents with higher propensity $P > P^2 + \left|\widetilde{P}^{-}_{P^0}\right|$ can invade. These agents' probability of matching with each other is high enough to ensure a higher average payoff.

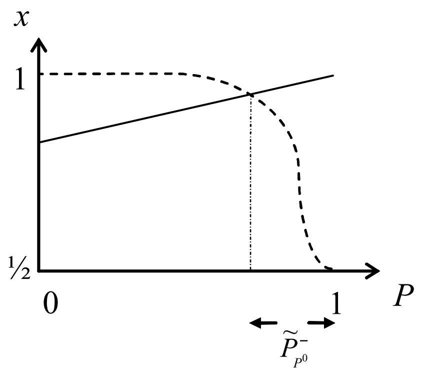

Qualitative figure of $x^{\min}_{1,P}$ and $x_{P_i,P_{-i}}$.

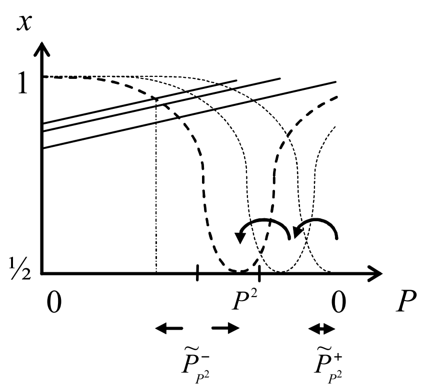

The emergence of the upper interval $\widetilde{P}^{+}_{P2}$.

Thus, there exists a positive probability that a complete invasion will take place at P^0, in which case the process starts over. Finally, note that if invasions are incomplete, all populations will contain agents with different propensities.

Property 5 *If* $\left|\widetilde{P}^{+}_{P=0}\right| > 0$ *and invasions are complete, then there exists a* $P' \neq P$ *where invasions will be successful. If* $\left|\widetilde{P}^{+}_{P=0}\right| > 0$ *and invasions are incomplete, then* $q_\sigma, q_{\sigma'} > 0$ *where* σ *and* σ' *have different corresponding propensities,* $P \neq P'$.

That is, if observational skills allow successful invasions of a population with propensity $P = 0$, then there will exist periods with agents of different propensity in the population. Regardless of whether invasions are complete or not, we find that:

Proposition 2 $\left|\widetilde{P}^{+}_{P=0}\right| > 0$ *and* $Q \in \Delta^{MAS}(\Gamma_1) \Rightarrow \frac{\partial \pi_i}{\partial O_i} > 0, \forall i \in I.$

The intuition for Proposition 2 is that when observational skills are good enough to ensure some degree of cooperation, agents with better observational skills will earn more because they will be more successful in avoiding agents with lower reputation and more successful in being matched against opponents with higher reputation. For the same reason, evolutionary pressure will continuously improve observational skills. Since mutations are rare, all agents will have identical observational skills, improving over time:

Corollary 3 $\left|\widetilde{P}^{+}_{P=0}\right| > 0$ *and* $Q \in \Delta^{MAS}(\Gamma_1) \Rightarrow O_i = O, \forall i \in I$ *and* $\lim_{t\to\infty} O = \infty$.

Since the interval $\left|\widetilde{P}^{-}_{P}\right|$ is decreasing in observational skills, it follows that there exists a crucial observational skill which we call O^* such that if $O_i \geq O^*$ for at least some agent, then there will be at least some cooperation in the population. Moreover, in the long run observational skills will improve until the population eventually reaches the case of perfect observational skills.

Denote a strategy mix where all choose action D and H respectively:

$$Q^D \equiv \{Q \mid z_i = 1, \forall i \in I\}, \text{ and } Q^H \equiv \{Q \mid z_i = 0, \forall i \in I\}.$$

Now consider the case when $\left|\widetilde{P}^{-}_{P}\right| \geq P$, i.e. when invasions are possible anywhere below P. In this case, the observational skills are so poor that more hawkish behavior is always beneficial. Thus, if observational skills are poor enough, $O_i < O^*, \forall i \in I$ such that $\left|\widetilde{P}^{-}_{P}\right| \geq P, \forall P \in [0,1]$, we have the following result:

Proposition 3 $\left|\widetilde{P}^{-}_{P}\right| \geq P, \forall P \in [0,1] \Rightarrow \nexists \left|\widetilde{P}^{+}_{P'}\right| > 0$ *for any* $P' \in [0,1]$.

Corollary 4 $O_i < O^*, \forall i \in I \Rightarrow Q^{\mathrm{H}} = \Delta^{MAS}(\Gamma_1)$.

Having described the two cases when $O_i = \infty$ and $O_i < O^*$ $\forall i \in I$, we now turn to the intermediate case when $O^* < O_i < \infty$, $\forall i \in I$.

Lemma 5 $\frac{\partial \Pr((i,j)\in\mathbb{I}|P_i=P_j)}{\partial O} > 0, \forall i,j \in I.$

The probability that agents with identical propensity are matched up is strictly increasing in observational skills. Better observational skills also mean that agents with higher propensity will be more successful in avoiding agents with lower propensity. This leads to the following conclusion:

Proposition 4 *The degree of cooperation in the population is strictly increasing in observational skills when* $O_i \geq O^*, \forall i \in I$.

To summarize the analysis of the cooperation game, if observational skills are poor enough, such that $\nexists\ P_i \neq P_j$, for any $i,j \in I$, then $Q^H = \Delta^{MAS}(\Gamma_1)$. If on the other hand the population for a given observational skill is in a *MAS*, such that $\exists\ P_i \neq P_j$, for some $i,j \in I$, then evolutionary pressure will lead to increasing observational skills. If agents are able to evolve better observational skills, then $\Delta^{MAS}(\Gamma_1)$ will eventually converge to Q^D, where all agents are playing *D*.

4 Evaluating the Coordination Game

Let us now consider the coordination game $\Gamma_2(\alpha,\beta)$. First define $\mathcal{H}$ as the set of preferences such that the agent prefers opponents with lower reputation:

Definition 6 $\mathcal{H} \equiv \{\succsim_i | \ r_{i,j} \leq r_{i,k} \Leftrightarrow j \succsim_i k, \forall j,k \in I\}$.

4.1 Perfect Observational Skills

Definition 7 *Let* σ^{D} *be the the set of strategies such that* $\sigma \in \sigma^{\mathrm{D}} \Leftrightarrow \sigma \in Q \in Q^{\mathrm{D}} \subset \Delta_Q$*, and* σ^{H} *such that* $\sigma \in \sigma^{\mathrm{H}} \Leftrightarrow \sigma \in Q \in Q^{\mathrm{H}} \subset \Delta_Q$.

First, we analyze the case when reputation perfectly reflects propensity: $O_i = \infty, \forall i \in I$. In a *MAS*, all agents utilize σ^D.

Proposition 5 $O_i = \infty, \forall i \in I \Rightarrow Q^{\mathrm{D}} = \Delta^{MAS}(\Gamma_2)$.

As in the cooperation game above, the population is vulnerable to a neutral invasion by σ^D where $\succsim \neq \mathcal{D}$. However, when a mutation with a propensity corresponding to the preferences of the drifting strategies occurs, both strategies will yield a lower payoff than the incumbent strategy.

We now turn to the more interesting case of imperfect observational skills.

4.2 Imperfect Observational Skills

Let us first focus on strategies such that $\sigma_s = 0, \forall s \in \mathbf{S}^c$, i.e. where the probability to play D is independent of the opponent.

Lemma 6 *$O_i < \infty, \forall i \in I$, and $\exists P_i \neq P_j$, for some $i, j \in I \Rightarrow \forall Q \in \Delta_Q$ we have that $\pi_i(\mathcal{D}) > \pi_i(\succsim_i \neq \mathcal{D})$ or $\pi_i(\mathcal{H}) > \pi_i(\succsim_i \neq \mathcal{H})$, for σ where $\sigma_s = 0, \forall s \in \mathbf{S}^c$.*

Less formally, if the observational skills are imperfect and there exist agents with different propensities, then for all unconditioned strategies the payoff is maximized by either $\mathcal{D}$ or $\mathcal{H}$.

Using lemma 6 it can be shown that for any $Q \in \Delta_Q$ to be a *MAS*, either all agents will play D or all agents will play H.

Proposition 6 $O_i < \infty, \forall i \in I \Rightarrow Q^{\mathrm{H}} \subseteq \Delta^{MAS}(\Gamma_2)$ *and* $Q^{\mathrm{D}} \subseteq \Delta^{MAS}(\Gamma_2)$.

In any *MAS*, agents with better observational skills will earn more:

Proposition 7 $O_i < \infty, \forall i \in I$, *and* $Q \in \Delta^{MAS}(\Gamma_2) \Rightarrow \frac{\partial \pi_i}{\partial O_i} > 0$.

The intuition for this is that agents with better observational skills will be more successful in avoiding mutant strategies.

Proposition 8 *$\exists O^* < \infty$ such that if $\exists i \in I$ where $O_i \geq O^*$ then $Q^{\mathrm{D}} = \Delta^{MAS}(\Gamma_2)$.*

Corollary 5 $O_i < O^*, \forall i \in I \Rightarrow Q^{\mathrm{H}} \cup Q^{\mathrm{D}} = \Delta^{MAS}(\Gamma_2)$.

To summarize, evolutionary pressure will lead to increasing observational skills. If agents are able to evolve better observational skills, then $\Delta^{MAS}(\Gamma_2)$ will eventually converge to Q^D, where all agents are playing D.

5 Conclusions and Remarks

We have shown that reputation based choice of opponents in games of cooperation and coordination is a mechanism that can explain the emergence and stability of cooperation and prosocial behavior. If the observational skills are perfect, there exists a *MAS* where almost the entire population plays D when cooperation is socially optimal. For given sufficiently accurate observational skills, the degree of cooperation in the population will vary in an unpredictable manner. However, in this case the payoff is strictly increasing in observational skills. The degree of cooperation is increasing in observational skills and decreasing in population size. Total defection is only possible when observational skills are sufficiently poor.

The cooperative strategy presented bears some resemblance to the tit for tat strategy, in that defections trigger retaliatory actions. In our model, the retaliatory action is choosing to play with someone else, and it is triggered by reputation, rather than the latest action of the opponent.

In the coordination game, we show that evolution works to increase observational skills. When observational skills becomes good enough, the efficient outcome is a unique equilibrium, just as predicted by *SSS*, see Young (1993) and Kandori et al. (1993), and the risk dominance criteria, see Harsanyi & Selten (1988).

Reputation based choice provides a possible explanation for the big impact of the degree of anonymity on behavior. When reputation does not perfectly reflect behavior, there are situations cooperation games where the payoff associated with playing hawkish will outweigh the reputational costs. For experimental evidence, see for example Cherry et al. (2002) or Hoffman et al. (1996). In Bergh & Engseld (2005) we demonstrate that existing empirical and experimental evidence support the idea that individuals care about their reputation and that reputational concerns affect behavior. As for the effect of reputation based choice of opponent, there is less experimental evidence available. McCabe et al. (2003) pair participants in a trust game are based on their degree of trust and trustworthiness, which allows cooperation to emerge and protects cooperation from being invaded by defecting players.

This supports our idea that an important key to understanding cooperation in repeated games is the matching procedure. Random/tournament matching here represents one extreme, whereas reputation based choice as analyzed in this paper represents another. In practice, people encounter some situations where they are able to choose their opponent in strategic interactions and some situations where they are forced to play games of cooperation against random agents in the population. The implications of such mixed matching procedures deserve to be examined closely. The results are likely to be positive for cooperation: As long as there is at least some degree of free opponent choice, agents must take into consideration the reputational consequences of their actions also when they play against randomly assigned opponents.

A Appendix

Proof of Lemma 1. $P_i = P_j, \forall i, j \in I \Rightarrow z_i = z_j, \forall i, j \in I$. Consequently, $\forall \succsim_i \in \mathbf{\Psi}$, all opponents will yield the same payoff. ■

Proof of Proposition 1. From Property 2 we know that $\frac{\partial \pi_i(\Gamma_1)}{\partial P_{-i}} > 0$. When observational skills are perfect, all agents can avoid opponents with lower propensity. Each agent i maximizes her payoff when $z_i = 1$ if $2\alpha \geq 1 + \beta$, and $z_i = \frac{1}{2} + \frac{1}{2}\frac{\alpha}{1+\beta-\alpha}$ if $2\alpha < 1 + \beta$. This results in a constant $P_i = K, \forall i \in I$. Denote the incumbent strategies σ. From Lemma 1 we know that the incumbent population can be neutrally invaded by σ' with $P' = K$ but with $\succsim \neq \mathcal{D}$. This in turn makes the population vulnerable to invasion by σ^* with $P^* < K$ and with $\mathcal{D}$. This will result in $\pi(\sigma') < \pi(\sigma)$, implying that σ' will grow slower than σ. Initially, $\pi(\sigma^*) \geq \pi(\sigma)$ because σ^* will only become matched up with σ' and σ^*. Since $\pi(\sigma^*) \geq \pi(\sigma) > \pi(\sigma')$, σ^* will, due to the relatively smaller $q_{\sigma'}$, gradually to a higher degree become matched up with σ^* and consequently earn a lower payoff. Thus, q_{σ^*} decreases until $\pi(\sigma^*) = \pi(\sigma')$, and $\pi(\sigma) > \pi(\sigma^*)$. This in turn implies that both $q_{\sigma'}$ and q_{σ^*} will decrease. No more mutations are possible as long as $q_{\sigma^*} > 0$. In a *MAS* the population will converge to:

$$\begin{aligned} \lim_{t\to\infty} \Pr(a_i = D) &= 1, \forall i \in I \text{ if } 2\alpha \geq 1 + \beta, \text{ and} \\ \lim_{t\to\infty} \Pr(a_i = D) &= \frac{1}{2} + \frac{1}{2}\frac{\alpha}{1+\beta-\alpha}, \forall i \in I \text{ if } 2\alpha < 1 + \beta. \end{aligned}$$

■

Proof of Lemma 2. From Property 2 and 3 we know that $\frac{\partial \pi_i(\Gamma_1)}{\partial P_{-i}} > 0$ and $\frac{\partial \pi_i(\Gamma_1)}{\partial r_{-i}} > 0$. If $\exists P_i \neq P_j$, for some $i, j \in I$, then $\mathcal{D}$ will result in a higher probability of avoiding opponents with a lower propensity and requesting opponents with a higher propensity than any $\succsim \neq \mathcal{D}$. ■

Proof of Lemma 3. Consider strategies σ and σ' with associated propensities P and P', such that $P' < P \leq P^*$. Imagine a population of incumbents with propensity P. Let the probability that incumbents are matched up with each other be denoted ρ. Envision a single mutant agent with propensity P', and let ρ' be the probability that this agent is matched up with herself.

Consider the additional times when the mutant plays H and the incumbent plays D. In these cases incumbents earn $\pi = \alpha\rho + (1-\rho)\beta$ whereas mutants earn $\pi' = 1 - \rho'$. We know that both ρ and ρ' are increasing in $x_{P,P'}$, and that ρ' is decreasing in N. Thus, for any $x_{P,P'} < 1$, a sufficiently large population will ensure that $\pi' > \pi$. Thus $\lim_{P' \to P} x^{\min}_{P,P'} = 1, \forall P \in (0,1]$. ■

Proof of Lemma 4. Let $\pi(\sigma : \sigma')$ denote the payoff for σ when matched up against σ'.

Consider a mutant strategy H with a corresponding $P = 0$ and incumbents with a strategy σ with a corresponding propensity $P > \widehat{P}$, where $\widehat{P}$ is the corresponding propensity of strategy $\widehat{\sigma}$ such that $\pi(\widehat{\sigma} : \widehat{\sigma}) = 0$. This implies:

$$\underbrace{\pi(H : \sigma)}_{=X \leq 1} > \underbrace{\pi(\sigma : \sigma)}_{=Y \lesseqgtr \alpha} > \underbrace{\pi(H : H)}_{=0} > \underbrace{\pi(\sigma : H)}_{=Z \geq \beta},$$

As before, let ρ denote the probability that an agent with strategy σ is matched against an agent with the same strategy. Analogously, ρ_H denotes the probability that an agent with strategy H is matched up against an agent with the same strategy. If $\pi(\sigma) > \pi(H)$, we have:

$$\begin{aligned} \rho\pi(\sigma : \sigma) + (1-\rho)\pi(\sigma : H) &> \rho_H \pi(H : H) + (1-\rho_H)\pi(H : \sigma) \\ \rho Y + (1-\rho) Z &> (1-\rho_H) X \\ \rho &> \frac{(1-\rho_H) X - Z}{Y - Z}. \end{aligned}$$

Note that $\lim_{\rho_H \to 1} \frac{(1-\rho_H)X - Z}{Y-Z} = \frac{-Z}{Y-Z} < 1 \Rightarrow \exists \rho_H < 1$ where $\frac{(1-\rho_H)X-Z}{Y-Z} < 1$. Thus $\exists \rho_H < 1$ such that $1 > \rho > \frac{(1-\rho_H)X-Z}{Y-Z}$. We know that $x_{P,0} < 1 \Leftrightarrow \rho, \rho_H < 1$. Then it follows directly that $x^{\min}_{P,0} < 1, \forall P \in \left(\widehat{P}, 1\right]$. ■

Proof of Proposition 2. $\left|\widetilde{P}^{+}_{P=0}\right| > 0 \Rightarrow \exists P_i \neq P_j$, for some $i, j \in I$. From Lemma 2 we know that $\exists P_i \neq P_j$, for some $i, j \in I \Rightarrow \pi(\mathcal{D}) > \pi(\succsim_i \neq \mathcal{D})$, which implies that $Q \in \Delta^{MAS}(\Gamma_1) \Rightarrow \succsim_i = \mathcal{D}, \forall i \in I$.

If the observational skills O_i of a single agent i improve and $\succsim_i = \mathcal{D}$, then from Assumption 1 and the matching procedure, we know that she is more likely to become matched up with opponents with higher propensity P_{-i}, i.e. $\frac{\partial P_{-i}}{\partial O_i} > 0$. Hence, $\frac{\partial \pi_i(\Gamma_1)}{\partial P_{-i}} > 0$, and $\frac{\partial P_{-i}}{\partial O_i} > 0 \Rightarrow \frac{\partial \pi_i(\Gamma_1)}{\partial O_i} > 0$. ■

Proof of Proposition 3. The upper interval for an arbitrary P where invasions are successful is given by $\left|\widetilde{P}_P^+\right| = P^* - P - \widetilde{P}_{P'}^- \Leftrightarrow P' = P + \left|\widetilde{P}_{P'}^-\right|$ for some $P' \in (P, P^*]$.

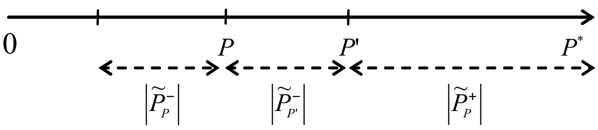

However, since $\left|\widetilde{P}_{P''}^-\right| = P'', \forall P'' \in [0,1]$, it follows that $\nexists \left|\widetilde{P}_P^+\right| > 0$, for any $P \in [0,1]$. ■

Proof of Lemma 5. From Assumption 1, we know that $\frac{\partial \operatorname{corr}\left[\underline{\mathbf{R}}_{i,I}, P_I\right]}{\partial O_i} > 0, \forall i \in I$. From Corollary 2 we know that $\succsim_i = \mathcal{D}, \forall i \in I$, which implies that $\frac{\partial \Pr\left(C_i^k = P_I^k\right)}{\partial O_i} > 0, \forall i, k \in I$. Assuming that all agents have identical observational skills O, from the matching procedure it follows that $\frac{\partial \Pr((i,j) \in \mathbb{I} | P_i = P_j)}{\partial O} > 0, \forall i, j \in I$. ■

Proof of Proposition 4. From Lemma 5 we know that when observational skills improve, the agents will be more likely to be matched up with opponents of identical propensity. Since $\frac{\partial \pi_i(\Gamma_1)}{\partial P_{-i}} > 0$, the agents with higher propensity will receive a higher payoff when observational skills improve and through growth become more frequent in the population. ■

Proof of Proposition 5. Let σ^d denote strategies that always play D. Note that $\sigma^d \subset \sigma^D$. Since the agents maximize their payoff if they play D and meet an agent that plays D, there are no other strategies in Δ_σ yielding a higher payoff than σ^d with $\mathcal{D}$, denoted $\sigma^{d:\mathcal{D}}$. As soon as $Q = Q^D$ we have $\pi\left(\sigma^D\right) = \pi\left(\sigma^{D:\mathcal{D}}\right) > \pi\left(\sigma\right), \forall \sigma \in \Delta_Q \setminus \sigma^D$, where $\sigma^{D:\mathcal{D}}$ denotes strategies σ^D

with $\mathcal{D}$. Finally, $\forall \sigma \in \Delta_Q \setminus \sigma^D$ and $q_\sigma, q_{\sigma^D} > 0$ we have that $\pi\left(\sigma^{d:\mathcal{D}}\right) \geq \pi\left(\sigma^D\right)$, and $\pi\left(\sigma^{d:\mathcal{D}}\right) > \pi(\sigma)$. ■

Proof of Lemma 6. The payoff for each agent $i \in I$ is $\pi_i(\sigma) = z_i z_{-i}\delta + (1-z_i)(1-z_{-i})\gamma$ where $z_i = \Pr(a_i = D \mid \succsim_i, Q)$, and where z_{-i} denotes the opponents $\Pr(a_{-i} = D \mid P_i)$. Note that P_i is independent of $\succsim_i$.

Consider the case when $z_i z_{-i}\delta > (1-z_i)(1-z_{-i})\gamma$, the proof is analogue for the case $z_i z_{-i}\delta < (1-z_i)(1-z_{-i})\gamma$. The expected propensity of agent i's opponent, P_{-i}, is higher when $\succsim_i = \mathcal{D}$ than if $\succsim_i \neq \mathcal{D}$. Hence, $z_{-i} : \succsim_i = \mathcal{D} > z_{-i} : \succsim_i \neq \mathcal{D}$, which implies that $\pi_i(\mathcal{D}) > \pi_i(\succsim_i \neq \mathcal{D})$. ■

Proof of Proposition 6. From Lemma 6 we know that $\pi_i(\mathcal{D}) > \pi_i(\succsim_i \neq \mathcal{D})$ or $\pi_i(\mathcal{H}) > \pi_i(\succsim_i \neq \mathcal{H})$, when $\sigma_s = 0, \forall s \in \mathbf{S}^c$, and $\forall Q \in \Delta_Q$. If $\pi_i(\mathcal{D}) > \pi_i(\succsim_i \neq \mathcal{D})$ then $\frac{\partial \pi_i}{\partial P_i} > 0$. Analogously, $\pi_i(\mathcal{H}) > \pi_i(\succsim_i \neq \mathcal{H})$ implies that $\frac{\partial \pi_i}{\partial P_i} < 0$. Thus, in a *MAS* all agents with $\mathcal{D}$ or $\mathcal{H}$ will have strategies $\sigma^{D:\mathcal{D}}$ or $\sigma^{H:\mathcal{H}}$.

Now consider any strategy mix Q where $\pi\left(\sigma^{d:\mathcal{D}}\right) > \pi\left(\sigma^{h:\mathcal{H}}\right)$. The proof is analogous when $\pi\left(\sigma^{d:\mathcal{D}}\right) < \pi\left(\sigma^{h:\mathcal{H}}\right)$. Since all matchings yield the same payoff for both the agent and the opponent, it follows that any strategy $\sigma' \in \Delta_Q \setminus \sigma^D$ yields $\pi(\sigma') \in \left(\pi\left(\sigma^{d:\mathcal{D}}\right), 0\right]$, since σ' will have a corresponding $P' < 1$. If $\sigma' \in \sigma^D \setminus \sigma^{d:\mathcal{D}}$ then $\pi\left(\sigma^{d:\mathcal{D}}\right) \geq \pi(\sigma')$.

Hence, in a *MAS* we have a Q^D where $\sigma_i^D, \forall i \in I$, or a Q^H where $\sigma_i^H, \forall i \in I$. ■

Proof of Proposition 7. From Proposition 6 it follows that if $Q \in \Delta^{MAS}$ then $\sigma_i^D, \forall i \in I$ or $\sigma_i^H, \forall i \in I$. In any *MAS*, better observational skills imply that agents with $\sigma^{d:\mathcal{D}}$ and $\sigma^{h:\mathcal{H}}$ will more often successfully identify mutants and thereby avoid coordination failures, earning them a higher payoff. ■

Proof of Proposition 8. From Proposition 6 we know that $Q \in \Delta^{MAS} \Rightarrow$

$\sigma_i^D, \forall i \in I$ or $\sigma_i^H, \forall i \in I$. Consider the case where $\sigma_j^{h:\mathcal{H}}, \forall j \in I \setminus i$. Imagine an invading strategy $\sigma_i^{d:\mathcal{D}}$ with a corresponding $P = 1$. Let $z_H = \Pr\left(\sigma^{h:\mathcal{H}} : \sigma^H\right)$ denote the probability that an agent with $\sigma^{h:\mathcal{H}}$ is matched up against an agent with identical propensity. Analogously, let $z_D = \Pr\left(\sigma^{d:\mathcal{D}} : \sigma^D\right)$. Note that $\frac{\partial z_H}{\partial O_j} > 0, \forall j \in I \setminus i$, and $\frac{\partial z_D}{\partial O_i} > 0$. The payoffs are $\pi\left(\sigma^{h:\mathcal{H}}\right) = z_H \gamma$ and $\pi\left(\sigma^{d:\mathcal{D}}\right) = z_D \delta$. Since $\delta > \gamma$, $\exists z_D < 1$ such that $\pi\left(\sigma^{d:\mathcal{D}}\right) > \pi\left(\sigma^{h:\mathcal{H}}\right)$. Consequently, $\exists O^* < \infty$ such that $O_i \geq O^*$ for some $i \in I \Rightarrow \pi_i\left(\sigma_i^{d:\mathcal{D}}\right) > \pi_j\left(\sigma_j^{h:\mathcal{H}}\right), \forall j \in I \setminus i$. ■

References

Axelrod, R. (1984), *The Evolution of Cooperation*, Basic Books, New York.

Bergh, A. & Engseld, P. (2005), 'The problem of cooperation and reputation based choise', *Lund University, Department of Economics Working Paper Series* **2005**.

Boyd, R. & Lorberbaum, J. (1987), 'No pure strategy is evolutionarily stable in the repeated prisoner's dilemma game', *Nature* **327**, 58–59.

Boyd, R. & Richerson, P. J. (1992), 'Punishment allows the evolution of cooperation (or anything else) in sizable groups.', *Ethology and Sociobiology* **13**, 171–195.

Cherry, T. L., Frykblom, P. & Shogren, J. F. (2002), 'Hardnose the dictator', *American Economic Review* **92**, 1218–1221.

Hammerstein, P., ed. (2003), *Genetic and Cultural Evolution of Cooperation*, Dahlem workshop reports, Dahlem University Press, USA.

Harsanyi, J. C. & Selten, R. (1988), *A General Theory of Equilibrium Selection in Games*, MIT Press, Cambridge.

Hoffman, E., McCabe, K. & Smith, V. (1996), 'Social distance and other-regarding behavior in dictator games', *American Economic Review* **86**(3), 653–660.

Kandori, M., Mailath, G. J. & Rob, R. (1993), 'Learning, mutation, and long run equilibria in games', *Econometrica* **61**(1), 29–56.

Maynard Smith, J. (1982), *Evolution and the Theory of Games*, Cambridge University Press, Cambridge, Massachusetts.

McCabe, K., Rigdon, M. L. & Smith, V. L. (2003), Sustaining cooperation in trust games, Mimeo, Harvard University.

Ostrom, E., Walker, J. & Gardner, R. (1992), 'Convenants with and without a sword: Self-governance is possible', *American Political Science Review* **86**, 404–417.

Samuelson, L. (1998), *Evolutionary Games and Equilibrium Selection*, MIT Press, Cambridge, Massachusetts.

Young, H. P. (1993), 'The evolution of conventions', *Econometrica* **61**(1), 57–84.