Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
http://economics.sas.upenn.edu/pier

PIER Working Paper 14-021

"Cooperation in Large Societies"

Second Version

by

Francesc Dilmé

http://ssrn.com/abstract=2444935

Cooperation in Large Societies*

Francesc Dilmé[†]

June 1, 2014

Abstract

This paper investigates how cooperation can be sustained in large societies even in the presence of agents who never cooperate. In order to do this, we consider a large but finite society where in each period agents are randomly matched in pairs. Nature decides, within each match, which agent needs help in order to avoid some loss, and which agent can help him incurring a cost smaller than the loss. Each agent observes only his own history, and we assume that agents do not recognize each other. We introduce and characterize a class of equilibria, named linear equilibria, in which cooperation takes place. Within this class, efficiency can be achieved with simple one-period strategies, which are close to a tit-for-tat strategy when the society is large, and which generate smooth dynamics of the expected average level of cooperation. Unlike previously suggested equilibria in similar environments, our equilibria are robust to the presence of behavioral agents and other perturbations of the base model. We also apply our model to bilateral trade with many traders, where we find that the mechanism of transmitting defections is transmissive and not contagious as in our base model.

Keywords: Cooperation, Many Agents, Repeated Games, Unilateral Help.

JEL Classification: D82, C73, D64.

*Work in progress. I am very grateful to George Mailath and Andy Postlewaite for their comments and encouragement, and the University of Pennsylvania's Micro Lunch seminar participants for their useful comments.

[†]University of Bonn. fdilme@uni-bonn.de

1 Introduction

The goal of this paper is to rationalize why people help other people who are unknown to them. In order to do this, we construct a model and propose a new class of equilibria that induce help in a large society. Differently from the previously proposed equilibria (see the literature review below), we find that cooperation is robust to some perturbations like, for example, the presence of agents who never cooperate.

We consider a finite but potentially large society. In each period, agents are uniformly randomly matched in pairs. Within each match nature decides who is the (potential) helper and who is the (potential) receiver (of the help). The helper decides between helping, which is costly for him but saves a bigger cost to the receiver, or not helping. We assume anonymity (i.e., agents do not recognize their past opponents) and that agents observe only the outcomes of their own matches. In this environment, where incentivizing cooperative behavior may be difficult, we find equilibria where cooperation takes place even when trembles or behavioral agents are present.

In our model, due to anonymity, individual punishments cannot be used to provide incentives. Also, if, for example, there are behavioral agents who never help, social punishments cannot be too harsh; otherwise, cooperation would be unsustainable and not incentive compatible. Our equilibrium strategies generate non-explosive but persistent punishments that make cooperation incentive compatible and robust.

Our equilibrium strategies balance the incentive of helping (in order to "improve" the continuation play of the rest of the society) and not helping (in order to get a high current individual payoff). In order for an equilibrium strategy to provide incentives to help, when an agent does not help, it must be the case that the rest of the society is "worsened," so then it is more likely that he is going to not be helped in the future. Still, the incentive of an agent to help cannot be too strong, since non-cooperative behavior has to be persistent enough in the society so defectors can be punished. Nevertheless, given that when an agent is not helped he does not know whether his opponent was a deviator himself or he was just perpetuating a deviation in order to punish someone else, the punishers are as punished as the deviators. So, agents could have the incentive to "forgive," that is, to "pretend" that they were helped in order to not be punished themselves. We show that simple one-period-memory strategies balance the incentive to help and not help after each history.

When the society is large, our equilibrium strategies are close to a tit-for-tat strategy, that is, with a very high probability each agent reproduces in the current period the action that was played in his stage game of the previous period. Tit-for-tat-like strategies make the effects of each action in the continuation play of the rest of the society small but persistent enough that each agent's incentives are balanced. As a consequence, a defection does not lead to an exponential contagion of non-cooperative behavior, but instead leads to a contagion that, in expectation, slowly disappears over time. As a result, we show that the expected fraction of agents helping in each period evolves smoothly over time.

We show that the existence of equilibria with cooperation is robust to the presence of perturbations such as trembles, behavioral types, or entry/exit of agents; and it depends only on the "total size" of the perturbation (for example, the number of behavioral types or the total probability of a mistake). The reason is that our mechanism to promote cooperation relies on the persistence of the effects of individual actions on the continuation play of the society. Since the "total size" of the perturbation is a measure of how fast information is lost, it therefore determines the existence of equilibria with cooperation. In contrast, we find that the long-run maximum level of cooperation depends also on the direction of the perturbation. For example, we prove that the existence of equilibria with cooperation depends only on the fraction of behavioral types in the society, while long-run full cooperation is reachable only when they cooperate for sure.

Our results can be extended to model bilateral trade in environments with a lot of traders whose products have characteristics that are unobservable at the moment the trade takes place. In those environments, efficiency is only reached if there traders trust eachother. We find that equilibria with trust exist only when the utility is separable, which is a reasonable assumption in bilateral trade environments. In this case, the mechanism to keep incentives is found to be transmissive, in contrast to the unilateral help case, where it is contagious.

After this introduction, there is a review of the literature related to our model. Section 2 introduces our unilateral help model, we introduce linear equilibria where cooperation takes place and we analyze their properties when the society gets large. Robustness to perturbations like entry/exit of agents, behavioral agents and trembles is studied in Section 3. Section 4 extends our results to a bilateral-trade model and Section 5 concludes. An appendix is used to provide the proofs of the results stated in the previous sections.

1.1 Literature Review

Kandori (1992), Ellison (1994), Harrington (1995) and Deb and González-Díaz (2011) consider repeated games with many agents and random matching and assume that there is no information about the opponent. These papers analyze grim-trigger-like (contagious) equilibria in which each player cooperates until he is defeated. After a defection, there is an exponential contagion, leading to a fully non-cooperative behavior. Using a public randomization device, cooperation can be restored. These equilibria are extremely fragile to small perturbations of the model, such as the introduction of a single non-cooperative behavioral agent. Also, when mistakes are introduced the whole society oscillates between total cooperation or complete defection. Our equilibrium strategies, instead, are close to a tit-for-tat strategy, are robust to some perturbations and feature smooth evolution of the aggregate level of cooperation. They provide us with a different and less fragile mechanism through which cooperation can be sustained, which generates significantly different cooperation dynamics.

Other models in the game theoretical literature on cooperation in large societies assume some degree of knowledge about the opponent in order to sustain cooperation. For example, Kandori (1992) considers the case where every agent knows everything that happens in the game. Okuno-Fujiwara and Postlewaite (1995) assume that agents have a social status that is updated over time according to some social norm. In Takahashi (2010) agents observe only their opponent's past play. Deb (2008) introduces cheap talk before each stage game takes place. In most of these models, folk theorems can be proved; that is, keeping the size of the society constant, if agents are patient enough, all rational payoffs can be achieved. Our model focuses on the potentially most hostile environment for cooperation, where agents have no other information than their own past play. So, even though additional information can potentially improve the efficiency and robustness of equilibrium cooperation, we believe that it is important to know what can be achieved in the absence of it.

Our equilibria are belief-free equilibria, in the sense that the continuation play of an agent is optimal independently of the current total history. These equilibria were introduced by Piccione (2002) and Ely and Välimäki (2002). We will use a version of the concept that Ely, Hörner and Olszewski (2005) adapted to games with incomplete information such as ours. As we will see, freedom from beliefs will allow us to make the strategy depend, in a simple way, on the individual history.

2 The Model

2.1 Unilateral Help

One of the main goals of this paper is to understand myopically suboptimal behavior that takes place in large societies. The simplest model of such behavior is one with "unilateral help." By unilateral help we mean a situation where there is one agent who needs help (the "receiver") and another agent who can decide to help him or not (the "helper"). Helping is costly to the helper, so it is individually suboptimal. Nevertheless, if helping is beneficial enough from the receiver's point of view, it is socially optimal.[1]

In order to model the previous situation, consider the following 2-player stage game Γ (often referred to as the (random) dictator game):

1. First, nature randomly assigns to each player a different role $\theta \in \Theta \equiv \{R, H\}$, where R is called the "receiver" and H is called the "helper."
2. After the role assignment, the helper decides on an action in $a \in A \equiv \{D, C\}$, where D is called "defection" (i.e., not helping) while C is called "cooperation" (i.e., helping).[2]
3. The payoffs after the action $a \in A$ is chosen are u_a^H for the helper and u_a^R for the receiver, which are given by[3]

C	$(u_C^H = -\ell, u_C^R = 0)$
D	$(u_D^H = 0, u_D^R = -g)$

   We assume that $g > \ell > 0$ so, from a social point of view, C is desirable.

Using our motivation, we interpret the previous game as a situation where an agent (the receiver) may incur a loss of g. Another agent (the helper) may help the receiver to prevent this loss at cost ℓ. Since $g > \ell > 0$, cooperation is socially optimal. Nevertheless, in the unique Nash equilibrium of the previous game the helper does not help the receiver.

2.2 Repeated Unilateral Help in Societies

As is well known in the repeated-games literature, myopically suboptimal behavior may be sustained when a stage game is (infinitely) repeated. In order for the intertemporal incentives to work, players should value enough the future (so future punishments/rewards

[1]Given that the players are ex-ante identical, by socially optimal we mean that helping is the symmetric strategy profile that maximizes the individual ex-ante payoff of each player (before role assignment) in the stage game.

[2]Even though most of the results are valid for a richer action space, we focus on the 2-action case in order to make our results clear. Still, our notation allows an immediate generalization to the multiple actions case, which is briefly analyzed at the end of this section.

[3]These payoffs are used without lost of generality. Indeed, any four real values u_C^H, u_D^H, u_C^R and u_D^R, satisfying $u_C^H < u_D^H$ (myopic incentive to play D), $u_C^R > u_D^R$ (R prefers to be helped) and $u_C^H + u_C^R > u_D^H + u_D^R$ (helping is socially optimal) generate the same set of equilibria in both the stage game and the repeated game.

counterbalance the myopic incentives) and actions should be observable enough (so the myopically/socially optimal actions can be punished/rewarded).

We now consider our stage (unilateral-help) game Γ played repeatedly in a society composed of a potentially large (but finite) number of players (or agents). We assume anonymity; that is, a player is not able to recognize his previous opponents and also players have no information other than their past stage-game outcomes.[4] In this case, anonymity and the lack of aggregate information play a role similar to non-observability in the standard repeated games models with imperfect monitoring: they make punishing deviators potentially very difficult. Our goal is to find explicit equilibria in this environment in which cooperation is sustained even in the presence of behavioral agents.

Time is discrete, $t = 0, 1, 2, ...$ Consider a set $\mathcal{N} \equiv \{1, ..., N\}$ of players (or agents), N even. In each period players are uniformly randomly matched in pairs and play Γ. All players have a common discount factor $\delta \in (0, 1)$.

Let $\mathcal{H}_{\text{ind}} \equiv \cup_{t=0}^{\infty}(\Theta \times A)^t$ denote the set of individual histories. For a given player $i \in \mathcal{N}$ and history $h^i \in \mathcal{H}_{\text{ind}}$, $h^i_t \in \Theta \times A$ specifies what role was assigned to player i in period t and the action that was chosen in period t's stage game (by him if $h^i_{t,1} = H$ or by his opponent if $h^i_{t,1} = R$). Let's define $\mathcal{H}^* \equiv \cup_{t=0}^{\infty}(\Theta \times A)^{tN}$ as the set of aggregate histories. An aggregate history $h^T \in \mathcal{H}^*$ is *consistent* if there exist T matching functions[5] $\{\sigma_t\}_{t=0}^{T-1}$ such that $h^i_{t,1} \neq h^{\sigma_t(i)}_{t,1}$ (different agents within a macth have different roles) and $h^i_{t,2} = h^{\sigma_t(i)}_{t,2}$ (the action realized is the same), for all $i \in \mathcal{N}$ and $t \in \{0, ..., T-1\}$. Let $\mathcal{H}$ denote the set of all aggregate histories that are consistent. Note that $\mathcal{H}$ is a strict subset of $\mathcal{H}^*$.

A behavior strategy for player i is a function $\alpha^i : \mathcal{H}_{\text{ind}} \to [0, 1]$ indicating the probability that, conditional on being chosen the helper, player i chooses action C (we also use $a^i_t = 1$ to denote that C is chosen, and $a^i_t = 0$ to denote that D is chosen). Fix a strategy profile $(\alpha^i)_{i \in \mathcal{N}}$, where, for each $t \geq 0$ and $i \in \mathcal{N}$, $\alpha^i_t \equiv \alpha^i(h^{i,t})$. Let $\tilde{V}^i(h^t)$ denote the expected future payoff of a player $i \in \mathcal{N}$ at time t (as a function of the aggregate history $h^t \in \mathcal{H}^*$).[6]

[4]This is in line with Kandori (1992), Ellison (1994), Harrington (1995) and Deb and González-Díaz (2011).

[5]A matching function is $\sigma : \mathcal{N} \to \mathcal{N}$ satisfying $\sigma(i) \neq i$ and $\sigma(\sigma(i)) = i$, $\forall i \in \mathcal{N}$.

[6]Note that since each player observes only his individual history, continuation plays are defined in all histories in $\mathcal{H}^*$, so the value function is well defined even for non-consistent histories.

We write $\tilde{V}^i(h^t)$ recursively in the following way:

$$\begin{aligned}\tilde{V}^i(h^t) &= \frac{1}{2}\left((1-\delta)\,u^H(\alpha_t^i) + \delta\,\mathbb{E}_t[\tilde{V}^i(h^{t+1})|H]\right) \\ &\quad +\frac{1}{2}\left((1-\delta)\sum_{j\neq i}\frac{u^R(\alpha_t^j)}{N-1} + \delta\,\mathbb{E}_t[\tilde{V}^i(h^{t+1})|R]\right) \\ &= (1-\delta)\,u(\alpha_t^i,\bar{\alpha}_t^{-i}) + \delta\,\mathbb{E}_t[\tilde{V}^i(h^{t+1})]\ , \end{aligned} \tag{2.1}$$

where $u(\alpha_t^i,\bar{\alpha}_t^{-i}) \equiv \frac{1}{2}\,u^H(\alpha_t^i) + \frac{1}{2}\,u^R(\bar{\alpha}_t^{-i})$, $u^H(\alpha_t^i) \equiv -\alpha_t^i\,\ell$ and $u^R(\bar{\alpha}_t^{-i}) = -(1-\bar{\alpha}_t^{-i})\,g$, $\mathbb{E}_t[\cdot]$ is the expectation over the role assignment and the actions played by players at t conditional on the total history of the game at time t (i.e., conditional on the corresponding aggregate history h^t), $\mathbb{E}_t[\cdot|\theta]$ is further conditioned on i being assigned role $\theta \in \{H,R\}$ at t and, for each history h^t,

$$\bar{\alpha}_t^{-i} \equiv \frac{1}{N-1}\sum_{j\neq i}\alpha_t^j \quad \in \Delta(A)$$

is the average distribution of actions that player i faces at history h^t under uniform random matching.

Linear Strategies

Note that the period-t payoff in equation (2.1) is a linear function of $\bar{\alpha}_t^{-i}$. Therefore, we can apply the expectation operator to the arguments of the payoff in each period:

$$V(h^t) = (1-\delta)\,\mathbb{E}_t\left[\sum_{s=0}^{\infty}\delta^s\,u(\alpha_{t+s}^i,\bar{\alpha}_{t+s}^{-i})\right] = (1-\delta)\sum_{s=0}^{\infty}\delta^s\,u\big(\mathbb{E}_t[\alpha_{t+s}^i],\mathbb{E}_t[\bar{\alpha}_{t+s}^{-i}]\big)\ . \tag{2.2}$$

Notably, the continuation value of an agent i depends on the complete history of the game only (and linearly) through his own strategy and the random variable $\mathbb{E}_t[\bar{\alpha}_{t+s}^{-i}]$. So, in order to find explicit equilibria, we take advantage of the linear structure of the payoffs by looking for strategies where $\mathbb{E}_t[\bar{\alpha}_{t+1}^{-i}]$ evolves linearly in a_t^i and $\bar{\alpha}_t^{-i}$. In this case, the RHS equation (2.1) becomes a linear function of α_t^i and $\bar{\alpha}_t^{-i}$.

Definition 2.1. We say that a strategy profile $(\alpha^i)_{i\in\mathcal{N}}$ *yields a linear law of motion* if it is symmetric and there exists a linear function $\Lambda : \Delta(A)\times\Delta(A)\to\Delta(A)$ such that

$$\mathbb{E}[\bar{\alpha}_{t+1}^{-i}|a_t^i,\bar{\alpha}_t^{-i}] = \Lambda(a_t^i,\bar{\alpha}_t^{-i}) \quad \text{for all } h^t\in\mathcal{H}^*,\ a_t^i\in A \text{ and } i\in\mathcal{N}. \tag{2.3}$$

The following theorem characterizes the strategies that yield linear laws of motion:

Proposition 2.1. *A symmetric strategy profile α yields a linear law of motion iff there exists a linear function $\Psi(\cdot,\cdot)$ such that $\mathbb{E}_t[\alpha^i_{t+1}|a^i_t,\bar{\alpha}^{-i}_t] = \Psi(a^i_t,\bar{\alpha}^{-i}_t)$ for all $h^t \in \mathcal{H}$ and $i \in \mathcal{N}$. Under these strategies, we have*

$$\Lambda(a^i_t,\bar{\alpha}^{-i}_t) = \frac{1}{N-1}\,\Psi(\bar{\alpha}^{-i}_t,a^i_t) + \frac{N-2}{N-1}\,\Psi(\bar{\alpha}^{-i}_t,\bar{\alpha}^{-i}_t)\ , \tag{2.4}$$

and $\mathbb{E}_t[\bar{\alpha}_{t+1}] = \Psi(\bar{\alpha}_t,\bar{\alpha}_t)$, where $\bar{\alpha}_t$ is the average mixed action of all players in period t.

We call an individual strategy yielding a linear law of motion a *linear strategy*, and an equilibrium yielding a linear law of motion an *equilibrium in linear strategies*. Note that $\Psi(\cdot,\cdot) \equiv 0$ (i.e., playing D after all histories) is a linear strategy, and it is easy to see that it leads to an equilibrium in linear strategies. We call equilibria with $\Psi(\cdot,\cdot) \not\equiv 0$ *cooperative equilibria*. Note that if players play a linear strategy, the continuation value (2.2) can be additively separated into a part that depends on the individual actions, and a part that depends only on $\bar{\alpha}^{-i}_t$. So, in any equilibrium, we can write $V(h^t) = V(\bar{\alpha}^{-i}(h^t))$.[7]

For a given linear function $\Psi(\cdot,\cdot)$, there are many strategies that satisfy equation (2.1). One class of particular interest given its simplicity is one-period-memory strategies. It is easy to show that if there is an equilibrium in linear strategies providing a payoff $V(\bar{\alpha}^{-i}_t)$, a one-period-memory equilibrium in linear strategies providing the same payoff also exists. So, we focus on one-period-memory strategies, which allow us to write

$$\alpha^i_{t+1}(h^t,(\theta_t,a_t)) = \Psi^{\theta_t}_{a_t}\ , \tag{2.5}$$

where $\Psi^{\theta}_{a} \in [0,1]$ indicates the probability of playing C if the role in the previous period was θ and the action played was a. The corresponding function Ψ described in Proposition 2.1 can now be written as[8]

$$\Psi(\alpha^i_t,\bar{\alpha}^{-i}_t) = \tfrac{1}{2}\,\Psi^{H}(\alpha^i_t) + \tfrac{1}{2}\,\Psi^{R}(\bar{\alpha}^{-i}_t)\ , \tag{2.6}$$

where $\Psi^{\theta}(\alpha) \equiv \alpha\,\Psi^{\theta}_{C} + (1-\alpha)\,\Psi^{\theta}_{D}$.

A cooperative equilibrium balances two incentives. First, players should be willing to cooperate in each period t, so playing C instead of D should increase the continuation payoff

[7]This separability of the continuation value between the own actions and the average cooperation level implies that equilibria in linear strategies are belief-free. Indeed, the continuation strategy is optimal independently of the current total history of the game. This feature, which is not present in the previous models, simplifies the equilibrium strategies and allows the equilibrium construction to be robust to perturbations, where the set of histories reached in equilibrium is very rich.

[8]Note that $\Psi(\cdot,\cdot)$ is the expected (mix) action played in period t, that is, before the role is assigned.

of an agent, i.e., increase $\mathbb{E}[\bar{\alpha}_{t+1}^{-i}]$. Second, players should be willing to play D after not being helped in the past in order to punish deviators. So, when a player plays D instead of C his continuation payoff cannot decrease too much. We find that these two incentives can be balanced after any history with simple one-period-memory strategies.

The linear nature of the strategies defined in Proposition 2.1 allows us to have neat interpretations of dynamic equations like (2.4). Indeed, given the action of player i (a_t^i), the dynamics of the "average cooperation level" of the rest of the players (given by the function Λ) can be decomposed into two terms. The first term on the RHS corresponds to the direct effect, that is, the expected change in the cooperation level of the opponent of player i. The second term is the expected change of $\bar{\alpha}_t^{-i}$ "by itself," since $N-2$ players are matched and change their cooperative behavior independently of what player i played this period. As one can expect, when N is large the direct effect is small.

Existence and Optimality of Cooperative Equilibria

After defining and characterizing cooperative equilibria, we now establish the conditions under which they exist, and we provide a concept of optimality within this class.

The following proposition establishes the condition for a cooperative equilibrium to exist:

Proposition 2.2. *There exists a cooperative equilibrium if and only if*

$$N \leq \frac{(g-\ell)\,\delta}{2\,\ell\,(1-\delta)} + 1 \ . \tag{2.7}$$

The condition for existence of cooperative equilibria, stated in equation (2.7), is intuitive. As one could expect, the higher the (expected) social gain from cooperation, $\frac{g-\ell}{2}$, compared with the myopic gain from deviating conditional on being the helper, ℓ, the easier it is for cooperative equilibria to exist. Also, a high discount factor relaxes the constraint for the existence of cooperative equilibria. We see that cooperation can be sustained in an arbitrarily large society, as long as its agents are patient enough.[9]

Given that in all equilibria in linear strategies the continuation value of a player i is a function of only $\bar{\alpha}_t^{-i}$, we can compare equilibria by comparing $V^i(\bar{\alpha}_t^{-i})$. We say that a cooperative equilibrium is *optimal* if $\alpha_0^i = 1$ and $V^i(\bar{\alpha}^{-i})$ is maximal for all $\bar{\alpha}^{-i}$ among all equilibria. Note that this is a strong concept of efficiency, since we impose optimality after any history, not just on the path of play. The following proposition establishes the existence and uniqueness of an optimal equilibrium:

[9]Formally, for all N there exists a $\bar{\delta} < 1$ such that if $\delta \geq \bar{\delta}$ a cooperative equilibrium exists.

Proposition 2.3. *If* (2.7) *holds, there exists a unique optimal one-period-memory cooperative equilibrium. In this equilibrium,* $\Psi_D^R = 0$ *and* $\Psi_C^R = \Psi_C^R = 1$*. If* (2.7) *holds with strict inequality, there is asymptotic full cooperation after any history, i.e.,* $\lim_{t\to\infty} \mathbb{E}[\bar{\alpha}_t | h^s] = 1$ *for any* $h^s \in \mathcal{H}^*$*.*

Note that in the optimal cooperative equilibrium a receiver at time t plays in period $t+1$ the same action he received the previous period (if he is chosen to be the helper). Intuitively, in the optimal equilibrium, the average cooperation level in the society $\bar{\alpha}_t$ should quickly approach 1 after any history, since helping is socially optimal. Nevertheless, $\bar{\alpha}_t$ quickly approaching 1 implies that defections are also quickly forgiven. Then, in order for cooperation to be incentive compatible, we need the behavior of the receiver to be fully responsive to the action played by the helper.

In order to get a better intuition of the optimal equilibrium, assume that all players play a tit-for-tat strategy.[10] In this case, it is easy to see that the average cooperation level would be a martingale. Indeed, assume that players i and j are matched at time t. With probability $\frac{1}{2}$, player i is the helper, so next period $a_{t+1}^i = a_{t+1}^j = a_t^i$, and with probability $\frac{1}{2}$ we have $a_{t+1}^i = a_{t+1}^j = a_t^j$. So, in expectation, there is one player in period $t+1$ playing a_t^i and one player playing a_t^j (even though they play the same strategy for sure). Since the expected cooperation level is preserved in each match, it is preserved also at the aggregate level. In this case, the difference in the continuation payoff of playing C or D is

$$-(1-\delta)\,\ell + \sum_{t=1}^{\infty} \frac{1}{N-1}\,\delta^t\,(1-\delta)\,\tfrac{1}{2}\,(g-\ell)\ .$$

If (2.7) holds with equality, then the previous expression is equal to 0, so a tit-for-tat strategy is an equilibrium. If, instead, (2.7) holds with strict inequality, the previous expression is strictly positive. In this case, if the rest of the players play a tit-for-tat, each player has incentives to always play C to increase the probability of being helped in the future. So, in order to make playing D more attractive, a probability that non-cooperative behavior will be forgiven is necessary. Hence, the optimal equilibrium exactly balances the incentives to play C and D by setting a positive probability of forgiveness $\Psi_D^H > 0$.

[10] We call the strategy $\Psi_a^\theta = a$ a *tit-for-tat* (strategy). Note that it differs from the standard definition of a tit-for-tat strategy because we have an ex-post asymmetric game, i.e., a game where the nature assigns different roles to different (ex-ante identical) players.

Many Actions

Let's briefly discuss how our results apply when the helper can choose more than two actions, i.e. $|A| > 2$. For notational convenience, we define $\Delta(A)$ as the simplex

$$\Delta(A) \equiv \left\{ v \in \mathbb{R}_+^{|A|} \;\middle|\; \textstyle\sum_{a\in A} v_a = 1 \right\}$$

and we will use $v \in \Delta(A)$ such that $v_a = 1$ to denote $a \in A$. Players are expected utility maximizers. We use $u^H(\alpha_t^i)$ (resp. $u^R(\alpha_t^j)$) to denote the stage game payoff of being the helper (resp. the receiver) and exerting (resp. receiving) a mixed action $\alpha_t^i \in \Delta(A)$.

We can use the results of the unilateral help with two actions to get a sufficient condition for the existence of equilibria in the general dictator game. Let a^{SO} be a "least-myopic" socially optimal action and a^{LE} a "least-efficient" Nash equilibrium of the stage game, that is,

$$\begin{aligned} a^{SO} &\in \operatorname*{argmin}_{\alpha\in\Delta(A)} u^H(\alpha) \quad \text{ s.t. } \alpha \in \operatorname*{argmax}_{\hat\alpha\in\Delta(A)} u(\hat\alpha, \hat\alpha) \ , \\ a^{LE} &\in \operatorname*{argmin}_{\alpha\in\Delta(A)} u(\alpha, \alpha) \quad \text{ s.t. } \alpha \in \operatorname*{argmax}_{\hat\alpha\in\Delta(A)} u^H(\hat\alpha) \ , \end{aligned}$$

where, as before, $u(a, a') \equiv \frac{1}{2}\, u^H(a) + \frac{1}{2}\, u^R(a')$.

Corollary 2.1. *A sufficient condition for the existence of cooperative equilibria where the socially optimal a^{SO} is played in every period is*

$$N \leq \frac{\delta}{1-\delta}\, \frac{u(a^{SO}, a^{SO}) - u(a^{LE}, a^{LE})}{u^H(a^{LE}) - u^H(a^{SO})} + 1 \ . \tag{2.8}$$

Intuitively, if the payoff loss of everyone playing a^{LE} instead of a^{SO} (numerator of the second fraction) is large compared to the short-run incentive to deviate (denominator), then the socially efficient outcome can be sustained in a large society.

2.3 Large Societies

As mentioned before, one of our main goals is to study cooperation in large societies, that is, societies with a large number of agents. We devote this section to analyze the joint limit $N \to \infty$ and $\delta \to 1$. This limit will not only be useful to analyze the dynamics of the cooperation level of the society, but also to find simpler and more intuitive expressions that will make our analysis of the perturbations more clear.

In the pervious section we obtained the condition for the existence of cooperative equilibria given in (2.7). In this expression, as δ approaches 1, the upper bound on N increases.

In this section we consider the case where δ is very close to 1 and N is very large, but keeping $(1-\delta)\,N$ constant.[11] In this limit, the problems that the increase in the number of agents poses to cooperation are alleviated by a higher interaction frequency among agents in a large society.

Since we want to increase the number of agents in the society, in this section we should consider a sequence of models instead of a single fixed model. Note that the parameters in our base model are $\Gamma \equiv (\ell, g, \delta, N)$. Let's fix a strictly decreasing sequence $(\Delta_n)_n$ converging to 0, and for each n let's define the n-th model as a model with parameters $\Gamma_n \equiv (\ell, g, \delta_n, N_n)$. In order to keep our notation simple, we assume that $N_n = \frac{1+\Delta_n\,\gamma}{\Delta_n\,\gamma}$ and $\delta_n = \frac{1}{1+\Delta_n\,\rho}$ for some constants $\gamma, \rho > 0$, and that $N_n \in 2\mathbb{N}$ for all n.

For each $n \in \mathbb{N}$, we interpret Δ_n as the length of the time period. We interpret ρ as the usual discount rate associated with the discount factor δ. Intuitively, the time at which a player discounts his payoff by a factor e^{-1} is $\frac{1}{\rho} + O(\Delta_n)$. The interpretation of γ is the rate at which a player meets a fraction of the population. Indeed, the expected real time at which a player meets a fraction e^{-1} of the players is $\frac{1}{\gamma} + O(\Delta_n)$. So, the limit involves increasing the population and the interaction level in the society at the same rate, meaning that the discounting time scale is similar to the interaction time scale.

Corollary 2.2. *Let $\gamma_* \equiv \frac{2\,\ell\,\rho}{g-\ell}$. Then, a cooperative equilibrium exists for the n-th model iff $\gamma \geq \gamma_*$, for all $n \in \mathbb{N}$.*

The previous corollary is proved by just plugging the values of N_n and δ_n into the existence condition (2.2). From now on in this section we assume that $\gamma \geq \gamma_*$, that is, cooperative equilibria exist. The next proposition sets an important property of the equilibrium strategies for any cooperative equilibrium when the society is large.

Proposition 2.4. *Fix any sequence of cooperative equilibria, one in each n-th model, with strategies denoted $\Psi_n(\cdot,\cdot)$. Then, $|a - \Psi_n^\theta(a)| = O(\Delta_n)$.*

The intuition of Proposition 2.4 is the following. As we argued before, in a tit-for-tat strategy each action has a permanent effect on the expected equilibrium level of cooperation. Also, in order to provide incentives to cooperate and punish in a large society, actions have to have an effect on the society's level of cooperation that is persistent enough. So, when

[11] To our knowledge, the joint limit where N is large and δ is close to 1 is not studied in the related literature. Models like Takahashi's (2010) consider $N = \infty$ (a continuum of agents) and discuss the properties of their equilibria in terms of δ (allowing information about the opponent). Conversely, Kandori (1992) and Ellison (1994) fix N (finite) and then let $\delta \to 1$ to find folk-theorem-like results.

N is large, tit-for-tat-like strategies are needed to ensure that the effects of each action in the continuation play of the rest of the players are persistent enough.

Let's finally focus on the dynamics of the average cooperation level in the society. In this case, for the sake of clarity, we focus now on the sequence of one-period-memory optimal equilibria (established in Proposition 2.3). The next corollary establishes that, when a society is large, the expected cooperation level moves smoothly toward 1 in the optimal equilibrium. Its proof is just standard algebra, so we skip it.

Corollary 2.3. *Fix two positive real time lengths $\tau, \tau' \in \mathbb{R}_+$, and assume $\gamma > \gamma_*$. Then, for the optimal equilibrium,*[12]

$$\mathbb{E}_t[\bar{\alpha}_{\lceil(\tau+\tau')/\Delta_n\rceil}] = 1 - (1 - \bar{\alpha}_{\lceil\tau/\Delta_n\rceil})\, e^{-\rho\left(\frac{\gamma}{\gamma^*}-1\right)\tau'} + O(\Delta_n)\ .$$

We see that when γ is big (that is, N is small), the speed of convergence of $\mathbb{E}_t[\bar{\alpha}_{\lceil(\tau+\tau')/\Delta_n\rceil}]$ to 1 is high. Indeed, as mentioned before, the optimal equilibrium maximizes the speed at which the level of cooperation increases while keeping the incentive to cooperate. So, when the number of players decreases, punishments are more effective, and therefore the speed of convergence can be increased.

3 Robustness

The goal of this section is to show that the existence of cooperative equilibria is robust to some perturbations. We think that robustness is important, since human societies are subject to many "imperfections" that may make cooperation (even) more difficult to sustain. So, in order to improve our understanding of how cooperation is sustained, we introduce some perturbations to the original model and show that the existence of cooperative equilibria remains generic to these perturbations.

This section is particularly relevant given that equilibria proposed in the literature when no information is available, such as the ones discussed in Kandori (1992), Ellison (1994) and Harrington (1995), are fragile to these perturbations. Those equilibria are grim-trigger-like equilibria: everyone cooperates as long as their opponents cooperate. Any non-cooperative behavior generates a quick non-cooperative contagion that leads to a total failure of cooperation. Eventually, using public signals, everyone may simultaneously return to full cooperation. In these grim-trigger-like equilibria, cooperation is possible as long as, in each period, everyone believes that with a very high probability there is no defection contagion

[12] $\lceil x \rceil = \min\{n \in \mathbb{Z} | n \geq x\}$ denotes the smallest integer no lower than x.

taking place. So, if agents' mistakes are introduced into the model, there may be an equilibrium in this class of strategies only if the probability of a mistake happening is very low *even at the aggregate level.* As a consequence, the expected (real) time when a mistake happens for one particular player must be very high when the society is large. We will see that in our equilibria, instead, cooperative equilibria may exist even if the expected real time of individual mistakes is not large. More strikingly, if there is only one behavioral agent who defects in every period while the rest of the $N-1$ players are normal, no grim-trigger-like equilibrium with cooperation can be constructed, independently of the value of δ. Our equilibria will survive in the presence of a small fraction of behavioral agents.

3.1 Trembles

Consider first the case where players may tremble (make mistakes) when deciding on their action. We assume that if a player plays C, his opponent perceives D with probability $\varepsilon_1 \Delta_n$, and if he plays D, his opponent perceives C with probability $\varepsilon_2 \Delta_n$.[13] For simplicity, we assume that none of the players observe whether a tremble happens when they act or receive their (perturbed) opponent's action.

Fix a player $i \in \mathcal{N}$. For a given mixed action α_t^i let $f(\alpha_t^i)$ be the expected probability that his opponent perceives the action as C. This is given by

$$f(\alpha_t^i) = (1 - \varepsilon_1 \Delta_n)\, \alpha_t^i + (1 - \alpha_t^i)\, \varepsilon_2 \Delta_n \ .$$

Note that $f(\cdot)$ is linear. In this case, for a given linear strategy $\Psi(\cdot,\cdot)$, it is easy to show that we can write the evolution of the average level of cooperation as

$$\mathbb{E}[\bar{\alpha}_{t+1}^{-i} | \alpha_t^i, \bar{\alpha}_t^{-i}] = \frac{\Psi(\bar{\alpha}_t^{-i}, f(\alpha_t^i))}{N-1} + \frac{N-2}{N-1}\, \Psi(\bar{\alpha}_t^{-i}, f(\bar{\alpha}_t^{-i})) \ .$$

Intuitively, we change the formula for $\mathbb{E}[\bar{\alpha}_{t+1}^{-i} | \alpha_t^i, \bar{\alpha}_t^{-i}]$ in the unperturbed model (given in (2.4)) to incorporate mistakes in the action received. Again, we have that $\mathbb{E}[\bar{\alpha}_{t+1}^{-i} | \alpha_t^i, \bar{\alpha}_t^{-i}]$ is linear in α_t^i and $\bar{\alpha}_t^{-i}$. The following proposition establishes the conditions under which cooperative equilibria exist for this case:

Proposition 3.1. *A cooperative equilibrium exists for n large enough if*

$$\gamma > \gamma'_* \equiv \frac{\ell\,(2\,\rho + \varepsilon_1 + \varepsilon_2)}{g - \ell} \ . \qquad (3.1)$$

If, instead, $\gamma < \gamma'_*$, then there is some n^* such that no cooperative equilibrium exists for all $n > n^*$. When $\gamma > \gamma'_*$ and n is large enough, there is a unique one-period-memory optimal

[13]In order to keep our expressions simple and intuitive, we write our results in the large society limit.

equilibrium, with $\Psi^H_C = \Psi^R_C = 1$ and $\Psi^R_D = 0$. The asymptotic expected fraction of players playing C under the optimal equilibrium, $\bar{\alpha}^*_\infty$, is given by

$$\bar{\alpha}^*_\infty = 1 - \frac{\varepsilon_1\,\ell}{(g-\ell)\,\gamma - 2\,\ell\,\rho} + O(\Delta_n)\ . \tag{3.2}$$

3.2 Behavioral Types

Assume now that, in each n-th model, a fraction $\phi_n\,\Delta_n$ of the players are behavioral (or action) types (so $\phi_n\,\Delta_n\,N_n \in \mathbb{N}$), and that $(\phi_n)_n$ converges to some $\phi > 0$. We assume that behavioral types play C with probability $\pi \in [0,1]$ independently of their history.[14] Let $N^\phi_n \equiv N_n\,\phi_n\,\Delta_n$ be the number of behavioral players. The evolution of $\bar{\alpha}^{-i}_t$ is now given by

$$\mathbb{E}[\bar{\alpha}^{-i}_{t+1}|\alpha^i_t,\bar{\alpha}^{-i}_t] = \Lambda(\alpha^i_t,\bar{\alpha}^{-i}_t) + \frac{N^\phi_n}{N_n-1}\Big(\pi - \Psi\big(\pi, \tfrac{(N_n-1)\,\bar{\alpha}^{-i}_t + \alpha^i_t - \pi}{N_n-1}\big)\Big)\ , \tag{3.3}$$

where Λ is defined in (2.4) and Ψ in (2.3).[15] The intuition behind the previous formula is the following. If there were no behavioral players, the term in the RHS of (3.3) would be just $\Lambda(\alpha^i_t,\bar{\alpha}^{-i}_t)$ (using equation (2.3)). Nevertheless, since behavioral types do not update the action they play, an adjustment term has to be included. This term, whose weight is $\frac{N^\phi_n}{N_n-1}$, undoes the update in the action specified by the strategy of the normal types and sets it to π again.

The following proposition establishes the conditions under which cooperative equilibria exist for this case:

Proposition 3.2. *The same result as in Proposition 3.1 holds for behavioral types, replacing $\varepsilon_1+\varepsilon_2$ by ϕ and $\frac{\varepsilon_2}{\varepsilon_1+\varepsilon_2}$ by π.*

> *Remark* 3.1. *(forgiveness delegation)* The strategies that we find in our model require agents to "forgive" (that is, play C after they played D in the previous period) with a very small probability ($O(\Delta_n)$). In the real world it may be difficult to individually generate these small probabilities. Instead, societies may reach the efficient equilibria by "generating" a small fraction of forgivers (behavioral types) that play C with a very high probability independently of the action they received. The rest

[14] One could assume that there is a mass $\phi\,\Delta_n$ of behavioral types whose mixing probability ν is distributed according to some distribution F. Nevertheless, by linearity it is easy to see that this is equivalent to all of them having the same mixing probability equal to $\pi \equiv \mathbb{E}[\nu]$.

[15] Here $\Lambda(\alpha^i_t,\bar{\alpha}^{-i}_t)$ can be interpreted as the unperturbed $\mathbb{E}[\bar{\alpha}^{-i}_{t+1}|\alpha^i_t,\bar{\alpha}^{-i}_t]$, so we can separate the effect of the perturbation.

of the players may then play a tit-for-tat strategy, a strategy that is much simpler from the individual point of view. Therefore, priests or charities could provide efficiency in a world with bounded rationality.

3.3 Entry/Exit

We finally assume that in each period after playing the stage game, $N_n^\phi \equiv \phi_n \, \Delta_n \, N_n$ players exit the game (die), and they are replaced by N_n^ϕ new players (are born), where (ϕ_n) satisfies the same conditions as in Section 3.2.[16] The entrants are aware of the equilibrium that is played, but they do not have any previous history. Note that (one-period-memory) linear strategies are a function of the previous history only through the previous play. Therefore, we assume that new players draw an initial mixed action ν drawn from a distribution F. As for the behavioral types case above, linearity is equivalent to all entering agents playing a mixed action equal to $\pi \equiv \mathbb{E}[\nu]$.

For practical purposes we assume the following timing within each period. First, the agents are randomly matched and play. Second, a fraction $\phi_n \, \Delta_n$ of them die. Finally, new agents are born with initial mixed action π, and the overall population remains constant. This timing allows us to write the following expression for the evolution of $\bar{\alpha}_t^{-i}$:

$$\mathbb{E}[\bar{\alpha}_{t+1}^{-i} | \alpha_t^i, \bar{\alpha}_t^{-i}] = f(\Lambda(\alpha_t^i, \bar{\alpha}_t^{-i})) \tag{3.4}$$

where

$$f(\bar{\alpha}_t^{-i}) \equiv \frac{N_n - 1 - N_n^\phi}{N_n - 1} \, \bar{\alpha}_t^{-i} + \frac{N_n^\phi}{N_n - 1} \, \pi \ ,$$

and where Λ is defined in (2.4) using (2.3). Intuitively, conditional on player i surviving period t, $\mathbb{E}_t[\bar{\alpha}_{t+1}^{-i}]$ is a transformation of the basic model's prediction to account for entry and exit of players.

Assume that agents get some utility $\bar{U}$ when they leave (die), independent of α_t^i and $\bar{\alpha}_t^{-i}$. So, each player's equivalent discount factor is $(1 - \Delta_n \, \phi_n) \, \delta$, to account for the possibility of dying.

Proposition 3.3. *The same result as in Proposition 3.2 holds for entry and exit of players, replacing all discount factors ρ by the equivalent discount factor $\rho + \phi$.*

---

[16]For simplicity we assume that the number of players that exit is certain. It is easy to see that, given the linearity in our strategies and payoffs, our results are the same if, instead, we assume that each player has an independent probability of leaving equal to $\phi_n \, \Delta_n$. This could be more realistic in some interpretations.

3.4 Comparison

Now let's see why the three perturbations discussed in this section lead to similar results and let's analyze them. The intuition will be provided using the optimal equilibrium.

The relationship between the perturbations with entry/exit of players and with behavioral types is obvious when the optimal equilibrium is considered. Indeed, in the entry/exit case, a fraction $\phi_n \Delta_n$ of the agents die and are reborn with initial mixing probability π in each period. Since the action played in the period after is only a function of the action played/received in this period, it is equivalent to pick up these players randomly (entry/exit case) or letting them always be the same (behavioral types).

For the trembles case, a fraction $\varepsilon_2 \Delta_n$ of players who are supposed to receive D end up receiving C. In the entry/exit case, instead, a fraction $\phi_n \Delta_n$ of the players who received D die and are replaced by players with initial mixing probability π. It is easy to see that given the linearity of our equilibrium strategies, changing the mixing probability of a fraction $\varepsilon_2 \Delta_n$ of the players from 0 to 1 is equivalent to changing the mixing probability of a fraction $\phi_n \Delta_n$ of players from 0 to $\pi = \frac{\varepsilon_2}{\phi_n} + O(\Delta_n)$. Doing the same for the players who receive D instead of C, we get the transformation rule of Propositions 3.1 and 3.2.

Notice that the condition for cooperative equilibria to exist is more restrictive in the perturbed model than in the not-perturbed model. We see in equation (3.1) that only the "size" of the perturbation ($\varepsilon_1 + \varepsilon_2$ in the trembles case, ϕ in the behavioral types and entry/exit cases) matters for the existence of equilibria. In particular, for example, the existence of equilibria in the presence of behavioral types depends only on the number of them, not on the actions they perform. The reason is that when there is a source of exogenous "deviations" (coming from entry/exit, behavioral types or trembling) it is more difficult but still possible to punish deviators. Intuitively, the punishment for non-cooperative behavior cannot be too persistent, since it takes place with positive probability in each period. So, it is more difficult to provide incentives to help. Less intuitively, if behavioral players are "too cooperative," deviations may be forgiven too fast, making defection too attractive. Hence, the "size" of the perturbation is a measure of how much information about the past play is lost in each period. Given that in cooperative equilibria incentives are provided through the persistence of the effects of the actions in the continuation play of the society, the size of the loss of information due to perturbations determines the existence of such equilibria.

Equation (3.2) shows that, conditional on the existence of cooperative equilibria, asymptotic efficiency is reachable only when the perturbation is in the "right" direction. Indeed, if $\varepsilon_1 = 0$ in the trembles case (i.e., when the helper plays C, the receiver receives C for sure) or if $\pi = 1$ in the entry/exit and the behavioral types cases (i.e., behavioral types and new entrants play C with probability 1), we have that efficiency can still be achieved.

4 Bilateral Trade

In this section we extend our results in order to model how trust can be sustained in large markets with bilateral trade of goods that have unobservable characteristics. We show that when payoffs are linearly separable, cooperative equilibria may sustain cooperation as in our unilateral-help model, but with a different mechanism to spread defections.

Consider a market with a lot of traders. In each period, traders are uniformly and randomly matched in pairs. Before trade takes place, each trader decides between producing a high-quality good (action denoted C) or producing a low-quality good (action denoted D). Let $u(a^i, a^j)$ be the utility derived from a trader who produces a good of quality a^i and exchanges it for a good of quality a^j.

Contrary to the dictator game, $u(\cdot,\cdot)$ (defined immediately after equation (2.1)) need not be linear in general in a simultaneous symmetric game. Indeed, $u(\alpha_t^i, \bar{\alpha}_t^{-i})$ may contain terms where components of α_t^i and $\bar{\alpha}_t^{-i}$ are multiplying each other. The following proposition establishes a necessary condition for the existence of cooperative equilibria:

Proposition 4.1. *A necessary condition for the existence of cooperative equilibria is that $u(\cdot,\cdot)$ be linear. This is equivalent to requiring the utility function to be additively separable:*

$$u(a^i, a^j) = u_1(a^i) - u_2(a^j) \quad \textit{for all } a^i, a^j \in A = \{D, C\}. \tag{4.1}$$

The previous proposition implies that cooperative equilibria exist only when the marginal (instantaneous) payoff from playing a'^i instead of a^i does not depend on a^j. This is the consequence of the fact that both Ψ and Λ need to be linear under linear strategies and therefore additively separable.

Still, the restriction in the payoff function (4.1) can be reasonably assumed when we model bilateral trade. Indeed, assume that the cost of producing a good is ℓ if its quality is high and 0 if its quality is low. Also, assume that the value of acquiring a high-quality good is g, while the value of acquiring a low-quality good is 0. In this case, the payoffs matrix of the corresponding game is the following prisoner's dilemma:

	C	D
C	$(g-\ell, g-\ell)$	$(-\ell, g)$
D	$(g, -\ell)$	$(0,0)$

(4.2)

Note that the payoff function (given by $u(a^i, a^j) = -\ell\, a^i + g\, a^j$) satisfies the condition (4.1). As in the unilateral help model, we assume that playing C (now interpreted as exchanging high-quality goods) is socially optimal but myopically suboptimal, that is, $g > \ell > 0$.

The following proposition is analogous to Propositions 2.2 and 2.3, establishing the conditions for the existence cooperative equilibria and characterizing the optimal one-period-memory equilibrium for the bilateral trade case (its proof is also analogous, so we omit it):

Proposition 4.2. *A cooperative equilibrium exists for the bilateral trade game iff*

$$N \leq \frac{(g-\ell)\,\delta}{\ell\,(1-\delta)} + 1\ . \tag{4.3}$$

In this case there is a unique optimal one-period-memory cooperative equilibrium, in which each player plays C at t if the action of his opponent was C at $t-1$, and if the opponent's action was D at $t-1$, he plays C with some fixed probability independent of the history.

The condition (4.3) for the prisoner's dilemma is identical to (2.7) for the dictator game except for a factor 2. This is driven by the fact that in the dictator game each player exerts an action, on average, every 2 periods, while in the prisoner's dilemma an action is exerted each period. Equivalently, the expected gain per period under full cooperation is $\frac{1}{2}(g-\ell)$ for unilateral help and $g-\ell$ for bilateral trade. This makes it easier to provide the incentives in the prisoner's dilemma.

Comparison

Most of the results we obtained for the unilateral help case can be almost identically reproduced for the bilateral trade case. Nevertheless, there is an important qualitative difference in the equilibria in both games, which is given by the way in which a deviation spreads. The dictator game's equilibrium is contagious, while the prisoner's dilemma's equilibrium is transmissive.

To see the difference, consider the optimal equilibrium in each game. Let's assume that in period t all players are supposed to play C for sure except for some player $i \in \mathcal{N}$, who is supposed to play D for sure. Let j be his opponent in period t. Then, we have:

- In the unilateral help case, two things may happen in period t. If i is the receiver, he will be helped with probability 1 and everyone will play C thereafter. If i is the helper, instead, he will play D. In this case, his opponent will play D in period $t+1$ (if is chosen to be the receiver) while he will play D with some positive probability (close to 1 if the society is large). Therefore, when the society is large, even though on average the number of "infected" players is less than (and close to) one, there may be two "infected" players. After T periods, the number of players playing D ranges from 0 to 2^T, but the average number of them converges to 0.

- In the bilateral trade case, both players i and j play an action. So, in period $t+1$ player i plays C for sure (since $a_t^j = C$), while player j will play D (since $a_t^j = D$) with some positive probability (close to 1 if the society is large), and the rest of the players will play C for sure. In this case we see that the number of players who play D is either 0 or 1. After T periods, the number of players playing D will still be 0 or 1, and the probability of total cooperation tends asymptotically to one.

As we see, the mechanism of transmission of punishments differs depending on the amount of information conveyed in the stage game. In the unilateral help case only one action is played, so the helper does not obtain information about the action that would have been played by the receiver. As a consequence, the mechanism used for keeping deviations on the society is contagious: with some probability it grows exponentially. This is necessary since after some player defects, his opponent is helped with probability $\frac{1}{2}$ at $t+1$ and therefore "forgives" the past defection. Therefore, in order to guarantee that deviators are punished, the threat of a big contagion (happening with small probability) is needed. In the bilateral trade model, instead, two actions are played, so there is a higher transmission of information. In this case, the mechanism is transmissive: the punishment gets transmitted from player to player for a long time until some player "forgives" it. In this case, a deviator knows that, even though the effect of the defection in the society remains small for sure, it is also (very) persistent.

The expected fraction of cooperators in both cases has very similar dynamics, but the dynamics of the actual fraction of cooperators are more volatile in the unilateral help case.

5 Conclusions

This paper shows that cooperative behavior can be robustly sustained in large societies. By constructing explicit equilibria, we show that even when non-cooperative behavior takes place in every period (for example, due to behavioral players), it is possible to sustain cooperation in equilibrium. This contrasts with previously proposed equilibria, which were fragile to many perturbations of the unperturbed game.

We find that sustaining cooperative behavior in a robust manner requires balancing the incentives to defect and to forgive defections. Linear strategies are able to balance these incentives history by history. These can be implemented by simple one-period strategies, which are close to a tit-for-tat strategy when the society is large. This feature makes social punishments small on a per-period basis, but persistent enough to guarantee that players want to follow the equilibrium strategy. Also, we show that the existence of cooperative equilibria is robust to the introduction of trembles or behavioral types.

Overall, the simplicity of our equilibrium strategies (one-period memory), their robustness to perturbations and the smooth dynamics of aggregate variables provide us with neat insights into how cooperation can be sustained in large societies. Future research shall be devoted to extending our results to other games, and to empirically estimate the strength of our proposed mechanism for sustaining cooperative behavior in large groups of people.

References

Deb, J., 2008, "Cooperation and Community Responsibility: A Folk Theorem for Repeated Matching Games with Names," Working Papers 08-24, New York University, Leonard N. Stern School of Business, Department of Economics.

Deb, J., and J. González-Díaz, 2011, "Community Enforcement beyond the Prisoner's Dilemma," Working paper.

Ellison, G., 1994, "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching," *Review of Economic Studies*, 61(3), 567–88.

Ely, J. C., J. Hörner, and W. Olszewski, 2005, "Belief-Free Equilibria in Repeated Games," *Econometrica*, 73(2), 377–415.

Ely, J. C., and J. Välimäki, 2002, "A Robust Folk Theorem for the Prisoner's Dilemma," *Journal of Economic Theory*, 102(1), 84–105.

Harrington Jr., J. E., 1995, "Cooperation in a one-shot Prisoners' Dilemma," *Games and Economic Behavior*, 8(2), 364 – 377.

Kandori, M., 1992, "Social Norms and Community Enforcement," *Review of Economic Studies*, 59(1), 63–80.

Okuno-Fujiwara, M., and A. Postlewaite, 1995, "Social Norms and Random Matching Games," *Games and Economic Behavior*, 9(1), 79–109.

Piccione, M., 2002, "The Repeated Prisoner's Dilemma with Imperfect Private Monitoring," *Journal of Economic Theory*, 102(1), 70–83.

Takahashi, S., 2010, "Community Enforcement When Players Observe Partners' Past Play," *Journal of Economic Theory*, 145(1), 42–62.

A Omitted Proofs

Proof of Proposition 2.1 (page 8)

Proof. We will prove that a symmetric strategy profile that yields a linear law of motion satisfies $\mathbb{E}_t[\alpha^i_{t+1}|a^j_t] = \Psi(\alpha^i_t, a^j_t)$, for some $\Psi(\cdot,\cdot)$ linear (the reverse implication is trivial), and that this implies (2.4). We prove it for a general (finite) set of actions A, as it is presented at the end of Section 2.2.

We first show that in any strategy profile that yields a linear law of motion, for any history h^t where player j plays α^j_t, the expected action in the next period (i.e. $\mathbb{E}_t[\alpha^j_{t+1}]$) will only be a function of α^j_t and the action played by his opponent at t, denoted by a^k_t. Note that for any player j and strategy α^j, there exists a unique (history dependent) matrix M^j_t such that $\mathbb{E}_t[\alpha^j_{t+1}|a^k_t] = M^j_t\, a^k_t$. So, by linearity, $\mathbb{E}_t[\alpha^j_{t+1}] = M^j_t\, \bar{\alpha}^{-j}_t$. Therefore:

$$\mathbb{E}_t[\bar{\alpha}^{-i}_{t+1}|a^i_t] = \frac{\sum_{j\neq i} M^j_t\, \bar{\alpha}^{-i}_t}{N-1} + \frac{\sum_{j\neq i} M^j_t\,(a^i_t - \alpha^j_t)}{(N-1)^2}\ .$$

Each of the terms in the sum of the first term on the RHS in the previous expression encodes the change in the action of player j if he faced the exact same population as player i. Nevertheless, he faces player i instead of himself, so the second term on the RHS of the previous expression corrects this (note that $\bar{\alpha}^{-i}_t + a^i_t/(N-1) = \bar{\alpha}^{-j}_t + \alpha^j_t/(N-1)$).

Assume $(\alpha^i)_{i\in\mathcal{N}}$ yields a linear law of motion and is such that there exist two individual histories $h^{j,t}, h^{j,t'} \in \mathcal{H}_{\text{ind}}$ such that $\alpha^j(h^{j,t}) = \alpha^j(h^{j,t'})$ and $\mathbb{E}_t[\alpha^j(h^{j,t})|h^{j,t}, a^k_t] \neq \mathbb{E}_t[\alpha^j_{t'}|h^{j,t'}, a^k_{t'}]$ for all j (recall that the strategy profile is symmetric) for some $a^k_t = a^k_{t'}$. This implies that $M^j_t(h^{j,t}) \neq M^j_{t'}(h^{j,t'})$. Let $h^t \equiv (h^{j,t})^N$ and $h^{t'} \equiv (h^{j,t'})^N$ be the aggregate histories where the corresponding individual histories are $h^{j,t}$ and $h^{j,t'}$ for all j, respectively. Then, note that

$$\mathbb{E}[\bar{\alpha}^{-i}_{t+1}|h^t, a^i_t] = \frac{1}{N-1}\, M^j_t(h^{j,t})\, a^i_t + \frac{N-2}{N-1}\, M^j_t(h^{j,t})\, \bar{\alpha}^{-i}_t(h^{j,t})\ ,$$
$$\mathbb{E}[\bar{\alpha}^{-i}_{t+1}|h^{t'}, a^i_{t'}] = \frac{1}{N-1}\, M^j_t(h^{j,t'})\, a^i_{t'} + \frac{N-2}{N-1}\, M^j_t(h^{j,t'})\, \bar{\alpha}^{-i}_t(h^{j,t'})\ .$$

Since $M_t^j(h^{j,t})$ and $M_t^j(h^{j,t'})$ are independent of a_t^i and $a_{t'}^i$, and since they are different, there is no Λ such that equation (2.3) is satisfied. This is a contradiction.

Let $\Psi(\alpha_t^i, a_t^j) \equiv \mathbb{E}_t[\alpha_{t+1}^i | \alpha_t^i, a_t^j]$ followed by the players, where a_t^j is the action received by player i at t and $\Psi : \Delta(A) \times A \to \Delta(A)$. As usual, we extend Ψ to $\Delta(A) \times \Delta(A)$ using linearity. Then

$$\Lambda(\alpha^i, \alpha^{-i}) = \sum_{j \neq i} \left(\frac{\Psi(\alpha^j, \alpha^i)}{(N-1)^2} + \sum_{k \neq i,j} \frac{\Psi(\alpha^j, \alpha^k)}{(N-1)^2} \right) = \sum_{j \neq i} \frac{\Psi(\alpha^j, \alpha^{-j})}{N-1} \, . \tag{A.1}$$

Fix $j \neq i$. Note that, by assumption, the LHS of the previous expression is linear in α^j (since Λ is linear and α^{-i} is linear in α^j). Furthermore, all terms in the sum of the RHS are linear in α^j except, maybe, $\Psi(\alpha^j, \alpha^{-j})$. Nevertheless, since the equality holds for all $\alpha^j \in \Delta(A)$, $\Psi(\cdot, \cdot)$ is linear in the first argument.

To prove the linearity of Ψ fix $j \neq i$. Note that, by assumption, the LHS of (A.1) is linear in α^i and α^j (keeping constant α^k for $k \neq j$). In the sum of the middle term, all terms in $\Psi(\alpha^k, \alpha^i)$ and $\Psi(\alpha^j, \alpha^k)$, for $k \neq i, j$, are clearly linear in (α^i, α^j). Therefore, $\Psi(\alpha^j, \alpha^i)$ is equal to a sum of terms linear in (α^i, α^j), and therefore it is itself linear. Therefore the first part of the theorem holds.

To prove (2.4), note that the terms of the middle equality in (A.1) can now be simplified using the fact that, as we have just proven, Ψ is linear. It is easy to see that the first term coincides with $\frac{\Psi(\bar{\alpha}^{-i}, \alpha^i)}{N-1}$. The second term can be expressed as follows

$$\begin{aligned} \sum_{j \neq i} \sum_{k \neq i,j} \frac{\Psi(\alpha^j, \alpha^k)}{(N-1)^2} &= \frac{\Psi(\sum_{j \neq i} \alpha^j, \sum_{k \neq i} \alpha^k)}{(N-1)^2} - \sum_{j \neq i} \frac{\Psi(\alpha^j, \alpha^j)}{(N-1)^2} \\ &= \Psi(\bar{\alpha}^{-i}, \bar{\alpha}^{-i}) - \frac{1}{N-1} \Psi(\bar{\alpha}^{-i}, \bar{\alpha}^{-i}) \end{aligned}$$

This shows the result. □

Proof of Proposition 2.2 (page 9)

Proof. Let's fix t and let p^{-i} be the fraction of players other than player i who play C in period t. This differs from $\bar{\alpha}_t^{-i}$ in that p^{-i} is the realization of the randomizations of the players. Then, we compute the incentives for each realization, and the linearity on the payoff and continuation value functions make the use of p^{-i} instead of $\bar{\alpha}_t^{-i}$ equivalent, but more convienient. In this case, the indifference condition imposes

$$-(1-\delta)\,\ell + \delta\, V\big(\tfrac{1}{N-1}\,\Psi_C^R + \tfrac{N-2}{N-1}\,\Psi(p^{-i}, p^{-i})\big) = \delta\, V\big(\tfrac{1}{N-1}\,\Psi_D^R + \tfrac{N-2}{N-1}\,\Psi(p^{-i}, p^{-i})\big) \, . \tag{A.2}$$

We also have the standard Bellman equation for the continuation value:

$$\begin{aligned} V(p^{-i}) &= \frac{1}{2}\Big[-(1-\delta)\,\ell + \delta\, V\big(\tfrac{1}{N-1}\,\Psi_C^R + \tfrac{N-2}{N-1}\,\Psi(p^{-i},p^{-i})\big)\Big] \\ &+ \frac{1}{2}\Big[-(1-\delta)\,p^{-i}\,g + \delta\, V\big(\tfrac{1}{N-1}\,(\Psi^H(p^{-i}) + \tfrac{N-2}{N-1}\,\Psi(p^{-i},p^{-i})\big)\Big]\ . \end{aligned} \tag{A.3}$$

Since as we know $V(\cdot)$ is linear, the previous equations are satisfied only if:

$$N = \frac{g-\ell}{2\,\ell}\,\frac{(\Psi_C^R - \Psi_D^R)\,\delta}{2 - (\Psi_C^H - \Psi_D^H + \Psi_C^R - \Psi_D^R)\,\delta} + 1\ . \tag{A.4}$$

The RHS of this equation is increasing in Ψ_C^H and Ψ_C^R, and decreasing in Ψ_D^H and Ψ_D^R. Given that they belong to $[0,1]$, the maximum N that can be part of a cooperative equilibrium is given by plugging $\Psi_C^H = \Psi_C^R = 1$ and $\Psi_D^H = \Psi_D^R = 0$, which leads to (2.7). □

Proof of Proposition 2.3 (page 10)

Proof. We proceed as in the proof of Proposition 2.2. We isolate Ψ_C^R from the equilibrium condition (A.4). We plug it in the indifference condition (A.2) and the Belleman equation (A.3), and we get the following expression for the value function:

$$V(p^{-i}) = \frac{\ell\,(N-1)}{2\,\delta\,(\Psi_C^R - \Psi_D^R)}\left(\tfrac{1}{2}\,(\Psi_D^H + \Psi_D^R)\,\delta + (1-\delta)\,p^{-i}\right) \tag{A.5}$$

Note that from equation (A.4) we know that $\Psi_C^R > \Psi_D^R$. Also, if $\Psi_C^R < 1$ we can increase Ψ_C^R and Ψ_D^R in the same amount leaving (A.4) the same but increasing the continuation value. Therefore, $\Psi_C^R = 1$. Note that, since the value function is increasing in Ψ_D^R but independent of Ψ_C^H, we can increase both still satisfying (A.4) until $\Psi_C^H = 1$. Finally, it is easy to verify that the value function is decreasing in Ψ_D^R when we change Ψ_D^R and Ψ_D^H so that equation (A.4) holds, so $\Psi_D^R = 0$. □

Proof of Corollary 2.1 (page 11)

The proof relies on constructing a strategy analogous to the 2-actions strategy. Playing C is identified with playing a^{SO} and playing D with playing a^{LE}. Therefore, α_t^i should be interpreted as the probability of playing a^{SO} and $1-\alpha_t^i$ as the probability of playing a^{LE}. Also, in order to construct the equilibrium, we specify that the continuation strategy treats past actions different from a^{SO} in the same way. In this case, since a^{LE} provides the highest myopic payoff, it is easy to see that there are no incentives to deviate to an action different from a^{LE} or a^{SO}, since the continuation play is independent of the action played. In this case, equation (2.8) is then just a generalization of (2.7) to arbitrary payoffs.

Proof of Proposition 2.4 (page 12)

The result is clear using equation (A.4), where we have that $N = O(\Delta_n^{-1})$ and $\delta = 1+O(\Delta_n)$.

Proof of Proposition 3.1 (page 14)

Proof. We proceed similarly to the proof of Propositions 2.2 and 2.3. Let's define, for the first part of this proof, $p_C \equiv \mu_C\,\Delta_n$ and $p_C \equiv \mu_C\,\Delta_n$. Now, we can write the strategy of a player after the stage-action profile (a^i, a^j) as

$$\Psi(a^i, a^j) = \tfrac{1}{2}\,(a^i\,\Psi_C^H + (1-a^i)\,\Psi_D^H) + \tfrac{1}{2}\left(f(a^j)\,\Psi_C^R + (1-f(a^j)\,\Psi_D^R\right)\ .$$

Note that it $\Psi(a^i, a^j)$ is linear. We can now proceed as in the proof of Proposition 2.2, the equation analogous to (A.4) is now

$$N = \frac{g\,(1-p_C-p_D)-\ell}{2\,\ell}\,\frac{(1-p_C-p_D)\,(\Psi_C^R-\Psi_D^R)\,\delta}{2-(\Psi_C^H-\Psi_D^H+(1-p_C-p_D)\,(\Psi_C^R-\Psi_D^R))\,\delta}+1\ .$$

Proceeding similarly as before, we find that the RHS in the previous equality is reached again when $\Psi_C^H = \Psi_C^R = 1$ and $\Psi_D^H = \Psi_D^R = 0$. Using the expressions $\delta = \frac{1}{1+\rho\,\Delta_n}$ and $N = \frac{1+\gamma\,\Delta_n}{\gamma\,\Delta_n}$ we get that the condition for existence of cooperative equilibria is given by equation (3.1).

The asymptotic cooperation level is found using $\Psi(p_\infty^*, p_\infty^*) = p_\infty^*$ and some algebra. □

Proof of Proposition 4.1 (page 18)

Proof. In a cooperative equilibrium each agent i has to be indifferent on playing any of the actions. Assume a pure strategy $\bar{\alpha}^{-i} \in A$. If player i pays $a^i \in A$, there is no uncertainty about the distribution in the next period (given by $\Lambda(a^i, \bar{\alpha}^{-i})$), so we have

$$V^i(\bar{\alpha}^{-i}) = (1-\delta)\,u(a^i, \bar{\alpha}^{-i}) + \delta\,V^i(\Lambda(a^i, \bar{\alpha}^{-i}))\ .$$

The previous equation must be valid for the two actions a^i (and all linear combinations of them). Since LHS does not depend on a^i and the stage-game payoff (and $\Lambda(a^i, \bar{\alpha}^{-i})$) is linear in a^i, then V^i must be also linear linear. So, since both V^i and Λ are linear, the stage-game payoff must be also linear, and therefore the proposition holds. □