



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 12-011

“Cooperation in Large Societies”

by

Francesc Dilmé

<http://ssrn.com/abstract=2034396>

Cooperation in Large Societies*

Francesc Dilmé[†]

March 26, 2012

Abstract

Consider the following situation involving two agents who belong to a large society. One of the agents needs help to avoid a big loss. The other agent may either incur a low cost to help him or do nothing. If agents do not recognize each other, providing incentives for socially optimal behavior (helping) is, in general, very difficult. We use a repeated anonymous random matching setting in a large society to understand how, in the previous situation, help may take place in equilibrium. We find explicit equilibria that, unlike other models proposed in the literature, feature smooth aggregate behavior over time and robustness to many perturbations, such as the presence of behavioral types or trembles. We consider the joint limit of increasing the size of the society and making it more interactive (or patient.) Under this limit, our equilibria resemble the tit-for-tat strategy for the prisoner's dilemma, introducing some small probability of forgiveness. The model is also applied to bilateral trade, where the mechanism used to spread deviations is transmissive instead of contagious. The smooth evolution of the aggregate variables over time makes the model suitable for empirical work.

Keywords: Cooperation, Many Agents, Repeated Game, 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 Pennsylvania. fdilme@sas.upenn.edu

1 Introduction

Helping other people implies, in general, some (opportunity) cost to the helper. In large societies, where most people do not know each other, it is difficult to provide incentives for people to help. Indeed, even when helping is socially optimal, it is myopically suboptimal, and anonymity makes rewarding helpers and punishing deviators in the future difficult. In order to understand why people help each other in our societies, we construct a model in the potentially most hostile environment: a society where a large number of anonymous players uniformly randomly match each period. In such an environment, the main contribution of this paper is to find non-trivial equilibria that have smooth aggregate properties and are robust to many perturbations. Therefore, we provide a game theoretical rationale for some behavior that is potentially difficult to explain in standard economic models.

Non-trivial equilibrium strategies balance two potential deviations: not helping to get a high current period (but “worsening the society”) or forgiving a deviation (and therefore “improving the society”). The first deviation implies that the consequences of not helping have to have a high persistence, since the deviator must be punished in the future. The second deviation requires agents to be willing to perpetuate bad behavior in the society. We will see that, in equilibrium, agents will be indifferent about helping or not along all the path of play.

The smooth evolution of aggregate variables (such as the number of people helping each period) is a feature that makes our equilibria more plausible and attractive from the empirical point of view than other equilibria proposed in the literature. The difficulty on achieving this arises from the fact that the player who does not receive help does not know whether his counterpart is deviating or just responding to the behavior of his matches in the past. In our equilibria exponential contagion will be avoided by keeping, after a defection, the expected aggregate number of defections small but persistent. On the contrary, the proposed equilibria in the literature in environments like ours (see the literature review below) rely on a fast collective change of behavior to punish deviations. Shortly after a deviation almost none of the members of the society help each other. Finally, a public signal is needed to make cooperation possible again. This makes their equilibria very fragile to small perturbations, that would generate oscillations between complete cooperation and complete defection; we show that our equilibria are not.

Our equilibria will be constructed using a particular class of strategies (named affine strategies). We will see that, unlike in other models in the literature, the existence of equilibria generated by these strategies is robust with respect to a number of perturbations. This is an important feature when we model societies using games with a large number of agents, since mistakes, behavioral types, or entry and exit of agents are likely to exist at the aggregate level.

The model can also be used to understand bilateral trade. In environments where a lot of agents trade products with characteristics that are unobservable at the moment the trade takes place, some inefficiencies may arise from the fact that it may be difficult to generate trust in the quality of the product sold. We will also see that in this environment we can find affine equilibria. Nevertheless, our results will only be valid for separable utility, which is a natural assumption of bilateral trade, but non-generic.

We consider the joint limit of increasing the number of players and letting the discount factor approach one. This limit will provide us with simple equations that are easy to interpret and will allow us to easily compute the conditions for the existence of equilibria under each perturbation. Also, it will provide some intuition about the mechanism through which deviations spread to the society, which will be either contagious or transmissive, depending on which model is used.

After this introduction, there is a review of the literature related to our model. Section 2 sets the general model and applies it to the dictator game and the prisoner's dilemma. In Section 3 we analyze the implications of the model for large societies. Robustness to perturbations like entry/exit of players, behavioral types and trembles is studied in Section 4. Section 5 concludes, and the Appendix is used to provide the proofs of the results stated in the previous sections.

1.1 Literature Review

Most of the 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 player knows everything that happens in the game. Okuno-Fujiwara and Postlewaite (1995) assume that players have a social status that is updated over time according to some social norm. In Takahashi (2010) players only observe 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 players are patient enough, all rational payoffs can be achieved.

Kandori (1992), Ellison (1994) and Harrington (1995), instead, assume that there is no information about the opponent. In Kandori (1992) the equilibrium is very fragile to perturbations, since it is a grim trigger. Ellison (1994) introduces a public randomization device that allows everyone to return to cooperation. Harrington (1995) investigates grim-trigger strategies when matching is not uniform. In these models aggregate dynamics are non-smooth when noise is introduced (society oscillates between full cooperation or total defection). Furthermore, the noise should be small enough such that at each period there are almost no mistakes in the whole society.

Our model will make dynamics smooth and will allow mistakes to happen every period.

As in other models in the literature, our equilibria will be belief-free equilibria. 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 Repeated Dictator Game

With the aim of explaining unilateral help, this section is devoted to modeling and finding equilibria of repeated games where, in each period, only one of the players takes an action. The stage game is a generalization of the dictator game, where the actions space and payoff functions are restricted to be finite. We will see that in this environment we can generically find explicit equilibria displaying nontrivial intertemporal incentives under repeated random matching with many agents.

Consider a set $\mathcal{N} \equiv \{1, \dots, N\}$ of players (or agents), N even. Consider the following 2-player stage game Γ :

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 $a \in A$, $|A| < \infty$. The player’s payoff after the action is chosen is $u^H(a)$ for the helper and $u^R(a)$ for the receiver.

Below (on page 8) we will interpret the helper as an agent who can choose to help or not the receiver, incurring a personal cost if he decides to help him. If the receiver is not helped, he incurs some cost, which is saved otherwise. If the cost savings of the receiver are higher than the cost of help from the helper, cooperation is socially optimal. Nevertheless, the myopic incentive for the helper is not to help.

Time is discrete, $t = 0, 1, 2, \dots$. In each period players are uniformly randomly matched in pairs and play Γ . All players have a common discount factor δ . For notational convenience, we define $\Delta(A)$ as the simplex

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

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 assume that players are anonymous; that is, a player is not able to recognize her opponent in each stage game. This is in line with Kandori (1992), Ellison (1994) and Harrington (1995). Anonymity and the lack of aggregate information make punishing deviators potentially very difficult in this environment.

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 $h^i \in \mathcal{H}_{\text{ind}}$, $h_t^i \in \Theta \times A$ specifies what role player i played in period t and the action that was chosen at period t 's stage game (by him if $h_{t,1}^i = H$ or by his opponent if $h_{t,1}^i = R$). 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 functions¹ $\{\sigma_t\}_{t=0}^{T-1}$ such that $h_{t,1}^i \neq h_{t,1}^{\sigma_t(i)}$ (the matched agents have different roles) and $h_{t,2}^i = h_{t,2}^{\sigma_t(i)}$ (the action realized is the same), for all $i \in \mathcal{N}$ and $t \in \{0, \dots, T-1\}$. 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 strategy for player i is a function $\alpha^i : \mathcal{H}_{\text{ind}} \rightarrow \Delta(A)$.

Fix a strategy profile $(\alpha^i)_{i \in \mathcal{N}}$, where, for each $t \geq 0$ and $i \in \mathcal{N}$, $\alpha_t^i \equiv \alpha^i(h^{i,t})$. 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 $h^t \in \mathcal{H}^*$).² We write $\tilde{V}^i(h^t)$ recursively in the following way:

$$\tilde{V}^i(h^t) = \frac{1}{2} \left((1 - \delta) u^H(\alpha_t^i) + \delta \mathbb{E}_t[\tilde{V}^i(h^{t+1})|H] \right) \quad (2.1)$$

$$+ \frac{1}{2} \left((1 - \delta) \sum_{j \neq i} \frac{u^R(\alpha_t^j)}{N-1} + \delta \mathbb{E}_t[\tilde{V}^i(h^{t+1})|R] \right) \quad (2.2)$$

$$= (1 - \delta) u(\alpha_t^i, \bar{\alpha}_t^{-i}) + \delta \mathbb{E}_t[\tilde{V}^i(h^{t+1})], \quad (2.3)$$

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

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

is the average distribution of actions that player i will face after history h^t under uniform random matching.

¹A matching function is $\sigma : \mathcal{N} \rightarrow \mathcal{N}$ satisfying $\sigma(i) \neq i$ and $\sigma(\sigma(i)) = i, \forall i \in \mathcal{N}$.

²Note that since each player only observes 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.

In a game with N players, with N potentially large beliefs over aggregate histories may be very complicated. Indeed, individual histories will be very diverse, and even more diverse when behavioral types are introduced or players enter/exit the game, as discussed in Section 4. Then, including players' beliefs about the aggregate histories and making the strategies optimal given these beliefs is in general not feasible. In order to keep the strategies simple, we introduce the concept of a stationary belief-free equilibrium.

Consider a strategy profile $(\alpha^i)_{i \in \mathcal{N}}$. Let $A_*^i \subseteq A$, for each $i \in \mathcal{N}$, denote the set of actions that player i plays with positive probability after some history (on or off the path of playing) under this strategy profile.

Definition 2.1. $(\alpha^i)_{i \in \mathcal{N}}$ is a **stationary belief-free equilibrium** (hereafter, SBFE, or simply equilibrium) if it is a sequential equilibrium and if, at any individual history, for each player $i \in \mathcal{N}$ it is optimal to play any $a^i \in A_*^i$, independently of the aggregate history.

Note that our concept of a stationary belief-free equilibrium differs from the belief-free equilibrium in Ely, Hörner and Olszewski (2005) in that we require players to be indifferent over all actions in A_*^i for all histories. This extra restriction will allow us to find stationary equilibria in large societies where there is constant entry and exit of agents. In our model, this implies that, under a belief-free equilibrium $(\alpha^i)_{i \in \mathcal{N}}$, for all i and $h^t \in \mathcal{H}^*$

$$\tilde{V}_t^i(h^t) \geq (1 - \delta) u(a_t^i, \bar{\alpha}_t^{-i}) + \delta \mathbb{E}_t[\tilde{V}_{t+1}^i(h^{t+1}) | a_t^i] \quad \forall a_t^i \in A \quad (2.4)$$

with equality for all $a_t^i \in A_*^i$. Note that, as we can see in (2.3), for a fixed player i , the only stage-payoff-relevant variables are his own action, a_t^i , and the average expected action of the rest of the players, $\bar{\alpha}_t^{-i}$.

In general, strategies may depend on the individual history in a very complicated way. In order to simplify this dependence, we can use the previous observation, trying to find a strategy profile such that the distribution of $\bar{\alpha}_{t+1}^{-i}$ is only a function of $\bar{\alpha}_t^{-i}$ and the individual action a_t^i . In this case, if we also require freedom from belief, it is easy to recursively write the individual problem in the form of the following Bellman equation:

$$V^i(\bar{\alpha}_t^{-i}) = (1 - \delta) u(a_t^i, \bar{\alpha}_t^{-i}) + \delta \mathbb{E}[V^i(\bar{\alpha}_{t+1}^{-i}) | a_t^i, \bar{\alpha}_t^{-i}] , \quad (2.5)$$

where V^i is the value function and a_t^i is any action in A_*^i .³ In order to find an explicit strategy with the previous properties, we will look for behavior under which updating of $\bar{\alpha}_t^{-i}$ is affine. Doing

³Here, $\mathbb{E}[\cdot | a_t^i, \bar{\alpha}_t^{-i}]$ is interpreted as the expected value if, conditional on $\theta_t^i = H$ player i chooses a_t^i , and the average action of the rest of the players (if chosen helper at t) is $\bar{\alpha}_t^{-i}$.

this, we will be able to find an affine solution for V^i , and therefore we will be able to bring the expectation inside the argument of V^i in the last term of the RHS of (2.5). Since both sides of the equation will be affine in $\bar{\alpha}_t^{-i}$, we will be able to find coefficients of the function V^i to match all terms.

Definition 2.2. We say that a strategy profile $(\alpha^i)_{i \in \mathcal{N}}$ **yields an affine law of motion** if it is symmetric and there exists a jointly affine⁴ function $\Lambda : \Delta(A) \times \Delta(A) \rightarrow \Delta(A)$ such that

$$\mathbb{E}[\bar{\alpha}_{t+1}^{-i} | a_t^i, \bar{\alpha}_t^{-i}] = \Lambda(a_t^i, \bar{\alpha}_t^{-i}) \quad (2.6)$$

for any history $h^t \in \mathcal{H}^*$, $a_t^i \in A$ and $i \in \mathcal{N}$.

Note that we impose the property at all histories, \mathcal{H}^* , not just at consistent ones. The reason is that we want this concept to be robust to a number of perturbations (see Section 4) that expand the set of feasible histories \mathcal{H} to \mathcal{H}^* or even larger sets. For the same reason we impose stationarity, that is, we do not allow Λ to be a function of time. This also allows us to easily model long-lived societies.

The following theorem fully characterizes the individual strategies that yield affine laws of motion:

Theorem 2.1. *A symmetric strategy profile α yields an affine law of motion if and only if $\mathbb{E}_t[\alpha_{t+1}^i | a_t^j] = \Psi(\alpha_t^i, a_t^j)$, where $\Psi(\cdot, \cdot)$ is a jointly affine function of α_t^i and the action received a_t^j . Under these strategies, we have*

$$\Lambda(a_t^i, \bar{\alpha}_t^{-i}) = \frac{\Psi(\bar{\alpha}_t^{-i}, a_t^i)}{N-1} + \frac{N-2}{N-1} \Psi(\bar{\alpha}_t^{-i}, \bar{\alpha}_t^{-i}), \quad (2.7)$$

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 at period t .

Proof. The proof is in the appendix on page 25. □

We call an individual strategy yielding an affine law of motion an **affine strategy**, and a SBFEE yielding an affine law of motion an equilibrium in affine strategies. Although for a given $\Psi(\cdot, \cdot)$ there are many ways to implement these strategies, two of them will be of special relevance for our arguments.⁵ The first is given by

$$\alpha_{t+1}^i = \begin{cases} \Psi^H(a_t) & \text{if } \theta_t^i = H, \\ \Psi^R(a_t) & \text{if } \theta_t^i = R, \end{cases} \quad (2.8)$$

⁴Joint affinity is equivalent to the existence of two matrices M_1 and M_2 such that $\Lambda(\alpha^i, \alpha^j) = M_1 \alpha^i + M_2 \alpha^j$ for all $\alpha^i, \alpha^j \in \Delta(A)$.

⁵The first strategies resemble the one-period memory strategies used in Ely and Välimäki (2002), while the second strategies resemble those that are only dependent on the opponent's actions used in Piccione (2002).

where a_t denotes the pure action that is realized at period t 's stage game. The corresponding function Ψ described in Theorem 2.1 (that is written before the role of the player is determined at t) is $\Psi(a_t^i, a_t^j) = \frac{1}{2} \Psi^H(a_t^i) + \frac{1}{2} \Psi^R(a_t^j)$. Note that this is a one-period memory strategy, since α_{t+1}^i only depends on the actions realized at time t .

Second, we take the expectation out of the left-hand side of the equation $\mathbb{E}_t[\alpha_{t+1}^i | a_t^j] = \Psi(\alpha_t^i, a_t^j)$. In order to keep the strategy well defined (that is, function of the individual history in and out the path of play) α_t^i should be interpreted as the (mixed) action that the player should play given the past history, and not the actual (mixed) action played (they may differ when the player deviates). Therefore, the strategy takes the form:

$$\alpha_{t+1}^i = \begin{cases} \Psi^H(\alpha_t^i) & \text{if } \theta_t^i = H, \\ \Psi^R(\alpha_t^i, a_t^j) & \text{if } \theta_t^i = R. \end{cases} \quad (2.9)$$

Note that in this case α_{t+1}^i is independent of the particular realization of the action of player i at time t , and it depends on the whole past individual history. Under this strategy, at period t , a player's strategy has an individual state variable given by the mixed distribution over the actions she is going to play this period, α_t^i , and it is updated using the information received (that is, the action played by the opponent). For most of our arguments, we will be using this class of strategies.

The affine nature of the strategies defined in Theorem 2.1 allows us to have neat interpretations of dynamic equations like (2.7). Indeed, given some action of player i (α_t^i), the dynamics of the ‘‘average state’’ 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 state variable of the player faced by player i . The second term is the expected change of $\bar{\alpha}_t^{-i}$ ‘‘by itself,’’ since $N - 2$ players will be matching and updating their states independently of what player i played this period. As one can expect, when N is large the direct effect vanishes.

Unilateral Help

One of the main goals of this paper is to understand helping behaviors that are myopically sub-optimal 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 has to 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 may be socially optimal.

Formally, consider the following 2-player stage dictator game:

1. First nature randomly assigns two roles: the helper and the receiver.
2. After the role assignment, the helper decides on an action in $\{E, S\}$, where E stands for “effort” (in helping) and S for “shirking” (not helping). The players’ payoff after the action is chosen is⁶

$$\begin{array}{c} \emptyset \\ E \quad \boxed{(a = -\ell, b = 0)} \\ S \quad \boxed{(c = 0, d = -g)} \end{array}$$

where the first component in each payoff vector corresponds to the helper’s payoff and the second corresponds to the receiver’s. We assume that $g > \ell$ so, from a social point of view, E is desirable.

The intuition here is a situation where an agent needs help (the receiver), and there is another who can potentially help him (the helper). If the helper helps the receiver, he incurs a cost ℓ . If the receiver is not helped, he loses g . If $g > \ell$, cooperation is socially optimal.⁷ Nevertheless, the myopic incentive for the helper is to not help (play S).

Following our previous notation, now the space of (mixed) actions is the segment $\Delta(\{E, S\}) \equiv \{(p, 1 - p) | p \in [0, 1]\}$ in \mathbb{R}^2 , where the first component corresponds to the probability of playing E . Since each mixed action $\alpha_t^i \in \Delta(\{E, S\})$ is characterized by a unique p_t^i , we will use that to denote it. Then, we can define p_t^{-i} and p_t as the analogous concepts to $\bar{\alpha}_t^{-i}$ and $\bar{\alpha}_t$ in the following way:

$$p_t^{-i} \equiv \frac{1}{N-1} \sum_{j \neq i} p_t^j \quad \text{and} \quad p_t \equiv \frac{1}{N} \sum_{j=1}^N p_t^j .$$

Using Theorem 2.1 we know that the following strategies yield an affine law of motion:⁸

$$p_{t+1}^i = \begin{cases} \eta^H p_t^i + k^H & \text{if chosen helper at } t, \\ \eta^R p_t^i + k_{a_t^j}^R & \text{if chosen receiver and received } a_t^j \text{ at } t, \end{cases} \quad (2.10)$$

⁶These payoffs are used without loss of generality. Indeed, any four real values a, b, c, d satisfying $a < c$ (myopic incentive to play E), $b > d$ (R prefers to be helped) and $a + b > c + d$ (helping is socially optimal) generate the same set of equilibria in both the stage game and the repeated game.

⁷Note that, indeed, the symmetric strategy that maximizes the individual ex-ante payoff (before role assignment) of the stage game is playing E with probability 1.

⁸Note that we are using strategies of the form (2.9). The results of existence and optimality remain unchanged if strategies of the form (2.8) are used, instead. In this second case $\eta^R = 0$ (since, when a player is chosen to be the receiver, the future action only depends on the action received), and $p_t^i = 0$ indicates $a_{t-1}^i = S$ and $p_t^i = 1$ if $a_{t-1}^i = E$.

where

$$(k^H, k^H + \eta^H, k_E^R, k_E^R + \eta^R, k_S^H, k_S^R + \eta^R) \in [0, 1]^6 . \quad (2.11)$$

The restriction on the parameters (2.11) ensures that if $p_t^i \in [0, 1]$, then $p_{t+1}^i \in [0, 1]$. Note that due to the asymmetric information (the helper does not know the action the receiver would have chosen in the case where he had been the helper), a player has two updating functions: one when he is chosen as the helper and the other when he is chosen as the receiver (that is contingent on the action received). The restriction that η^R is independent of the action received is a consequence of the joint affinity imposed by the Theorem 2.1.

The following proposition establishes the condition for SBFE to exist:

Proposition 2.1. *There exists a non-trivial SBFE if and only if*

$$N \leq \frac{(g - \ell) \delta}{2 \ell (1 - \delta)} + 1 . \quad (2.12)$$

Proof. The proof of this proposition is analogous to the proof of Proposition 2.4. Nevertheless, the proof of Proposition 2.4 is algebraically simpler and will provide us with more intuitive insights about the model. Therefore, this proof is left to the reader. \square

Let's now introduce a concept of optimality in order to find the "best" equilibrium in affine strategies (when (2.12) holds). To extend the concept to large and long-lived societies with some perturbations (modeled in Section 4), we will require optimality given any history, in or off the path of play. Let \mathcal{E}^{UH} denote the set of equilibria in affine strategies for the unilateral help, that is

$$\mathcal{E}^{\text{UH}} \equiv \{ (k^H, \eta^H, k_E^R, k_S^H, \eta^R, p_0) \mid (2.10) \text{ is an eq. and } (2.11) \text{ holds} \} . \quad (2.13)$$

Given $e \in \mathcal{E}^{\text{UH}}$, let $V_e^i : [0, 1] \rightarrow \mathbb{R}$ denote the corresponding value function for player $i \in \mathcal{N}$. We say that an equilibrium $e \in \mathcal{E}^{\text{UH}}$ is **optimal** if $p_0^i = 1$ and $V_e^i(p^{-i}) \geq V_{e'}^i(p^{-i})$ for all $p^{-i} \in [0, 1]$, $i \in \mathcal{N}$ and $e' \in \mathcal{E}^{\text{UH}}$. Note that this is a strong concept of efficiency, since we impose optimality given any p , not just on the path of play (we will see below that, on the path of play, $p_t = 1$ for all t). The following proposition establishes the existence and uniqueness of an optimal equilibrium:

Proposition 2.2. *If (2.12) holds, there exists a unique equilibrium in affine strategies that is optimal. In this equilibrium, $k_E^R = 1$, $k_S^R = \eta^R = 0$ and $k^H + \eta^H = 1$. If (2.12) holds with strict inequality, under this equilibrium we have asymptotic efficiency, i.e., $\lim_{t \rightarrow \infty} \mathbb{E}[p_t | h^s] = 1$ for any $h^s \in \mathcal{H}^*$.*

Proof. The proof of this proposition is again completely analogous to the proof of Proposition 2.5, but the proof of Proposition 2.5 is algebraically simpler and will provide us with more intuition. \square

Note that the optimal equilibrium in affine strategies, characterized in the previous proposition, has the highest sensitivity of the individual state to the received action, since $h_t^i = (R, E)$ implies $p_{t+1}^i(h^t) = 1$, and $h_t^i = (R, S)$ implies $p_{t+1}^i(h^t) = 0$. Indeed, conditional on being the receiver, the individual state switches to 0 or to 1, depending on the action received. Intuitively, the best equilibrium is the one where p_t approaches fastest to 1 (everyone helping is socially optimal). Nevertheless, when p_t approaches fast to 1, the change in the aggregate state resulting from a deviation is less persistent. Then, to be incentive compatible, we need the aggregate state to be as sensitive as possible to individual deviations.

We can use the results of the unilateral help 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

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

Lemma 2.1. *A sufficient condition for the existence of equilibrium in affine strategies where the socially optimal a^{SO} is played in every period is*

$$N \leq \frac{\delta}{1 - \delta} \frac{u(a^{SO}, a^{SO}) - u(a^{LE}, a^{LE})}{u^H(a^{LE}) - u^H(a^{SO})} + 1 . \quad (2.14)$$

Proof. The proof is in the appendix on page 26. \square

Intuitively, if the loss in social efficiency of a^{LE} (numerator of the second fraction) is large compared to the short-run incentive to deviate (denominator), then the socially efficient outcome can be sustained for large societies.

2.2 Symmetric Simultaneous Games

In order to compare our results with those obtained in the literature on repeated games with many players, we now adapt the previous section to a symmetric simultaneous stage game. We will see that while Theorem 2.1 is still true, we will have a new necessary condition on the payoffs for the existence of equilibria in affine strategies.

We can use the same notation as for the dictator game. Indeed, equation (2.3) is still valid, but now $u(\cdot, \cdot)$ should be interpreted as the payoff given the (mixed) actions. The concept of strategy profile that yields an affine law of motion also remains the same, so Theorem 2.1 is still valid.

Contrary to the dictator game, $u(\cdot, \cdot)$ needs not to be jointly affine 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 belief-free equilibria in affine strategies:

Proposition 2.3. *A necessary condition for the existence of a non-trivial equilibrium in affine strategies is that $u(\cdot, \cdot)$ is jointly affine. This is equivalent to requiring that the utility be separable:*

$$u(a^i, a^j) = u_1(a^i) - u_2(a^j) \tag{2.15}$$

for all $a^i, a^j \in A$. In this case V^i is an affine function.

Proof. The proof is in the appendix on page 26. □

The previous proposition says that the marginal (instantaneous) payoff from playing a^i instead of a^j does not depend on a^j . This is the consequence of the fact that both Ψ and Λ are jointly affine under affine strategies and therefore additively separable. So, the marginal change in $\bar{\alpha}_{t+1}^{-i}$ (and then also the marginal change in future payoff) when player i plays a^i instead of a^j is independent of $\bar{\alpha}_t^{-i}$. Since, by freedom from belief, V^i is independent of the action chosen in A_*^i , the marginal change in the current payoff has to also be independent of the action received.

Prisoner's Dilemma

The dictator game described in the previous section provides a framework where the necessary condition for affine strategies to exist (see Proposition 2.3) holds generically. Still, most of the literature focuses on the prisoner's dilemma. In order to compare our results with those in the literature and to model bilateral trade, in this section we apply our model to the prisoner's dilemma.

There are situations where the payoff condition (2.15) can be reasonably assumed. The most commonly used example that satisfies this condition is bilateral trade. Indeed, in most of the models of bilateral trade, individual gains depend on the sum of the payoffs of one's action (profits from cheating or not on the quality of the product sold, for example) and the trading counterpart's

action (monetary costs of being cheated or not by the trading partner).⁹ Then, consider the following symmetric prisoner's dilemma:

$$\begin{array}{cc|cc}
 & & E & S \\
 E & (g - \ell, g - \ell) & (-\ell, g) \\
 S & (g, -\ell) & (0, 0)
 \end{array} \tag{2.16}$$

The interpretation of the previous payoff matrix is the following. We interpret g as the gain from trade if the counterpart does not cheat (the quality of the good is high), and ℓ as the cost of not cheating (the cost of making a high-quality good). It is important to recall that the utility is separable, so the marginal gain from playing S instead of E is independent of the action selected by the opponent. This allows the existence of equilibria in affine strategies.

Note that now affine strategies have less free parameters than in the unilateral help case, since there is no role differentiation. Indeed, under no role differentiation, Theorem 2.1 establishes that affine strategies have the form

$$p_{t+1}^i \equiv \Psi(p_t^i, a_t^j) = \eta p_t^i + k_{a_t^j} , \tag{2.17}$$

where p_t^i is, as in the dictator game, the probability of playing E at t , a_t^j is the action played by the opponent of player i at t and

$$(k_E, k_S, k_E + \eta, k_S + \eta) \in [0, 1]^4 . \tag{2.18}$$

Recall that, as a consequence of Theorem 2.1, η needs to be independent of the action received.

The following proposition is analogous to Proposition 2.1, establishing the conditions for the existence of affine equilibria for the prisoner's dilemma:

Proposition 2.4. *An equilibrium in affine strategies exists for the two-action game if and only if*

$$N \leq \frac{(g - \ell) \delta}{\ell(1 - \delta)} + 1 . \tag{2.19}$$

Proof. The proof is in the appendix on page 27. □

⁹Consider the following model for the exchange of goods. There are two possible values for the quality of a good (low or high), which is not observable. The value of acquiring a high-quality good is g , while the value of a low-quality good is 0. The cost of producing a good is ℓ if its quality is high and 0 if its quality is low. After exchange occurs, the payoff of each player is the valuation of the acquired good (g or 0) minus the cost of the good given to the counterpart (ℓ or 0).

The condition (2.19) for the prisoner’s dilemma is identical to (2.12) 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. Apart from this, as we said at the beginning of this section, the results are very similar to those of the dictator game.

Let’s now state a result about the characterization of the optimal equilibrium in affine strategies, when it exists. The intuition is the same as for Proposition 2.2:

Proposition 2.5. *If (2.19) holds, there exists a unique optimal equilibrium in affine strategies. In the optimal equilibrium $\eta = 0$ and $k_E = 1$. If (2.19) holds with strict inequality, the optimal equilibrium features asymptotic efficiency, i.e., $\lim_{t \rightarrow \infty} \mathbb{E}[p_t | h^s] = 1$ for all $h^s \in \mathcal{H}^*$.*

Proof. The proof is in the appendix on page 28. □

3 Large Societies

As mentioned before, one of the main goals of our model is to explain cooperation in large societies, that is, when the number of players is large. This section introduces a technique that allows us to directly obtain asymptotic conditions for the existence of equilibria and the aggregate variables dynamics without having to deal with the otherwise tedious finite expressions. This will be especially useful to solve the extensions of the model in Section 4, obtaining simpler expressions that are easier to interpret.

3.1 Asymptotic Properties

In the previous section we obtained conditions for the existence of equilibria ((2.12) and (2.19)). Notice that in both expressions, as δ approaches 1, the upper bound on N increases. In this section we will consider the case where δ is very close to 1 and N is very large, but keeping $(1 - \delta)N$ constant.

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 (2010) consider $N = \infty$ (a continuum of agents) and discuss 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 $\delta \rightarrow 1$ to find folk-theorem-like results. The joint limit can be taken in our model and equilibria in affine strategies will still generically exist.

Let Δ denote the length of the time period, and let t now denote the physical time, $t = 0, \Delta, 2\Delta, \dots$. Let's define $\gamma \equiv \frac{1}{N\Delta}$ and $\rho \equiv \frac{1-\delta}{\Delta}$. 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 $t = \frac{1}{\rho}$. The interpretation of γ is the rate at which a player meets a fraction of the population. Indeed, the expected time at which a player meets a fraction e^{-1} of the players is $t = \frac{1}{\gamma}$.¹⁰ So, the limit is taken such that there is enough interaction in the society, meaning that the discounting time scale is similar to the interaction time scale.

For simplicity we will consider the prisoner's dilemma, although the results are very similar for the dictator game. Fix ρ and γ , and assume that $\Delta > 0$ is such that $\frac{1}{\gamma\Delta} \in 2\mathbb{N}$ and $1 - \rho\Delta > 0$. Let $\mathcal{E}^{\text{PD}}(\Delta)$ be the corresponding set of equilibria in affine strategies. Then:

Lemma 3.1. *Consider a sequence $(\Delta_n)_{n \in \mathbb{N}}$ such that, for all $n \in \mathbb{N}$, $\Delta_n > 0$, $\frac{1}{\gamma\Delta_n} \in 2\mathbb{N}$ and $1 - \rho\Delta_n > 0$, and satisfying $\lim_{n \rightarrow \infty} \Delta_n = 0$. Then,*

$$\exists n_* \text{ such that } \mathcal{E}^{\text{PD}}(\Delta_n) \neq \emptyset \quad \forall n > n_* \quad \Leftrightarrow \quad \begin{cases} \gamma > \gamma_* \text{ or} \\ \gamma = \gamma_* \text{ and } \gamma_* \geq \rho. \end{cases}$$

where $\gamma_* \equiv \frac{\ell\rho}{g-\ell}$.

Proof. The proof is in the appendix on page 28. □

In the proof of Proposition 2.4 (equation (6.6)) we see that the necessary and sufficient condition for an equilibrium in affine strategies to exist is the existence constants k_S, k_E, η satisfying (2.18) and

$$k_E - k_S = \frac{\frac{1-\delta}{\delta(1-\eta)} + 1}{\frac{g-\ell}{\ell(N-1)} + 1} = \frac{\frac{\rho\Delta}{(1-\rho\Delta)(1-\eta)} + 1}{\frac{\gamma\Delta(g-\ell)}{\ell(1-\gamma\Delta)} + 1}. \quad (3.1)$$

The previous expression indicates that, as Δ becomes small, equilibrium k_S and η should be small as well, and $k_E = 1 + O(\Delta)$ should be close to 1. So, we say that $(\tilde{k}_S, \tilde{k}_E, \tilde{\eta}, p_0) \in \mathbb{R}_+^2 \times \mathbb{R} \times [0, 1]$

¹⁰Indeed, let z_t denote the fraction of players that a given player still has not met before period t . Then, the expected variation of z_t is given by

$$\mathbb{E}_t \left[\frac{z_{t+\Delta} - z_t}{\Delta} \right] = \frac{(1 - z_t)z_t + z_t(z_t - \gamma\Delta) - z_t}{\Delta} = -\gamma z_t.$$

Then, as $\Delta \rightarrow 0$, z_t solves $\dot{z} = -\gamma z$, so $z(t) = e^{-\gamma t}$.

is an **asymptotic equilibrium** if there exists a sequence $(\Delta_n)_n$ satisfying the same conditions of Lemma 3.1 and a sequence of equilibria in affine strategies $(k_{S,n}, k_{E,n}, \eta_n, p_{0,n}) \in \mathcal{E}^{\text{PD}}(\Delta_n)$ such that¹¹

$$\begin{aligned} k_{S,n} &= \tilde{k}_S \Delta_n + O(\Delta_n^2) , \\ k_{E,n} &= 1 - \tilde{k}_E \Delta_n + O(\Delta_n^2) \text{ and} \\ \eta_n &= \tilde{\eta} \Delta_n + O(\Delta_n^2) . \end{aligned}$$

It is easy to show that if $\gamma > \frac{g\rho}{l-g}$ then $(\tilde{k}_S, \tilde{k}_E, \tilde{\eta}, p_0)$ is an asymptotic equilibrium if and only if $\tilde{\eta} \in [-\tilde{k}_S, \tilde{k}_E]$ and

$$\tilde{k}_S + \tilde{k}_E - \tilde{\eta} = \rho \left(\frac{\gamma}{\gamma^*} - 1 \right) . \quad (3.2)$$

Fix an asymptotic equilibrium $(\tilde{k}_S, \tilde{k}_E, \tilde{\eta}, p_0)$. Then, for any sequence $(\Delta_n)_n$ supporting this equilibrium, we have¹²

$$\mathbb{E}_t \left[\frac{p_{t+\Delta_n}^{-i} - p_t^{-i}}{\Delta_n} \middle| p_t^i \right] = \underbrace{\mathbb{E}_t \left[\frac{\Psi(p_t^{-i}, p_t^{-i}) - p_t^{-i}}{\Delta_n} \right]}_{\equiv \dot{p}^{-i}(p_t^i, p_t^{-i})} + O(\Delta_n) , \quad (3.3)$$

The first three terms in the RHS of the previous expression correspond to the first term of the RHS of the expression (2.7), that is, the evolution of p_t^{-i} “by itself” (using the law of motion given by $\Psi(\cdot, \cdot)$). The fourth term in the RHS corresponds to the rate at which p_t^{-i} is changed if p_t^i is played, depends on the size of the population (parameterized in γ) and the expected change in the state of the opponent.¹³ We can use the previous expression to obtain the limiting expression for Bellman equation (2.5) (written in (6.4) for the prisoner’s dilemma):

$$\rho V^i(p^{-i}) = \max_{p^i} (\rho u(p^i, p^{-i}) + V^i(p^{-i}) \dot{p}^{-i}(p^i, p^{-i})) + O(\Delta_n) . \quad (3.4)$$

Note that V^i is affine by Proposition 2.3, so $V^i(p)$ is well defined. Since affine strategies require indifference, we should impose the condition that the argument of the max operator is independent of p^i . So, solving (3.4) we again find condition (3.2).

¹¹We use the usual notation $O(\Delta)$ to mean “terms that go to 0 at least as Δ .” For example, $a = b + O(\Delta)$ means that $\limsup_{\Delta \rightarrow 0} |\frac{a-b}{\Delta}| < \infty$.

¹²Expressions like (3.3) should be interpreted as being true for any sequence of equilibria that generates the corresponding asymptotic equilibrium.

¹³For example, assume $p_t^i = 1$ (E is played with probability 1). Then, since $k_E = 1 + O(\Delta_n)$, $k_S = O(\Delta_n)$ and $\eta = O(\Delta_n)$, we have that $\mathbb{E}[p_{t+\Delta_n}^j - p_t^j] = 1 - p_t^{-i} + O(\Delta_n)$, where j is the opponent of player i at t .

Hence, if we want to obtain the existence condition and dynamics of p with an error $O(\Delta_n)$ we can directly solve (3.4) using (3.3) and imposing indifference (disregarding $O(\Delta_n)$ terms). These are simple equations, especially taking into account the fact that $V^i(\cdot)$ is affine. We will see that this technique is useful to get simple analytical formulae when we introduce perturbations in the model in Section 4.

The optimal equilibrium for each Δ_n (established in Proposition 2.5) is such that $\tilde{k}_E = 0$, $\eta = 0$ and $\tilde{k}_S = \rho(\frac{\gamma}{\gamma^*} - 1) + O(\Delta_n)$. Therefore, we have that

$$\begin{aligned} \dot{p}_t &= \rho \left(\frac{\gamma}{\gamma^*} - 1 \right) (1 - p_t) + O(\Delta_n) \\ \Rightarrow p_t &= 1 - (1 - p_0) e^{-\rho \left(\frac{\gamma}{\gamma^*} - 1 \right) t} + O(\Delta_n). \end{aligned}$$

We see that when γ is big (that is, N is small compared to $1/(1 - \delta)$) the speed of convergence to $p = 1$ is high.

The same exercise of considering the limit $\Delta_n \rightarrow 0$ can be done for the dictator game instead of the prisoner's dilemma (recall that in this case we have a game with role differentiation). It is easy to see that when Δ_n is small, k^H , $1 - \eta^H$, $1 - k_E^R$, k_S^H , η^R are $O(\Delta_n)$. The results and formulae are very similar, so we do not repeat the calculations.

3.2 Comparison

The previous results for the dictator game and the prisoner's dilemma are very similar. Indeed, except for a 2 factor, even for $\Delta = 1$, the corresponding bounds on N , (2.12) and (2.19), are identical. Nevertheless, there is an important qualitative difference in the equilibria in both games, which is given by the way 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 at period t everyone's individual state is 1 (play E for sure) except for player $i \in \mathcal{N}$, whose state is 0 (playing S for sure).¹⁴ Let j be his opponent at period t . Then, we have:

- In the dictator game, two things can happen this period. If i is the receiver, he will be helped with probability 1 and his state (and the state of everyone else) on period $t + \Delta$ will be 1 again. If i is the helper, he will play S and update his state to $p_{t+\Delta}^j = k^H = O(\Delta)$, in which case the number of players that will have a state "close to 0" at $t + \Delta$ is 2. Therefore, although on average the number of "infected" players is less than (and close to) one, there may be two "infected" players.

¹⁴Note that in both games, in the optimal equilibrium, if there is no deviation, $p_{t'}^i = 1$ for all $t' > t$ and $i \in \mathcal{N}$.

- In the prisoner’s dilemma, at period $t + \Delta$ player i has state $p_{t+\Delta}^i = k_E = 1$, player j has state $p_{t+\Delta}^j = k_S = O(\Delta)$, and the rest of the players have state 1. In this case we see that the number of players with a state “close to 0” is 1 for sure. This state will be transmitted until someone, with a probability $O(\Delta)$, “forgives” the deviation.

Note that both optimal equilibria are “quasi-tit-for-tat,” in the sense that a player plays with very high probability the action that his last opponent played. This is clear for the prisoner’s dilemma, since $\eta = 0$, $k_E = 1$ and $k_S = O(\Delta)$. For unilateral help, recall that since $\eta_E^R = 1$ and $k_S^R = \eta^R = 0$, if a player was the receiver last period this period he will be playing the exact same action he received. If the last period in which he was the receiver was T periods ago and he received action E , he will now be playing E with probability 1 (since $p_* = 1$). If, instead, he received action S , he will now be playing E with probability

$$\eta^H \sum_{t=2}^{T-2} (1 - \eta^H)^t = \eta^H \frac{1 - (1 - \eta^H)^{T-1}}{\eta^H} = 1 - (1 - \eta^H)^{T-1} .$$

Note first that since the probability of being the helper each period is $\frac{1}{2}$, the expected value for T is 2. Also, since $\eta^H = 1 - O(\Delta)$, $\mathbb{E}[1 - (1 - \eta^H)] = O(\Delta)$. Therefore, the player will be playing S with a very large probability.

Finally, note that unlike in the other models in the literature, the evolution of p_t (the fraction of people helping each period) is smooth on average in both the dictator game and the prisoner’s dilemma. Indeed, $\mathbb{E}_t[p_{t+\Delta}]$ differs from p_t only in a term $O(\Delta)$. Furthermore, it is easy to see that the variance of $p_{t+\Delta} - p_t$ is asymptotically $O(\Delta^2)$ when $\Delta \rightarrow 0$. This is a feature that seems more plausible from an empirical point of view than global coordination of random public signals to return to cooperation.

4 Robustness

Large societies are usually subject to a variety of distortions that may make it difficult for equilibria to exist. Among others, the entry/exit of agents, mistakes or preference heterogeneity may play a crucial role in preventing the mechanisms for punishing deviations to act properly.

Equilibria proposed in the literature when no information is available, such as the ones discussed in Kandori (1992), Ellison (1994) and Harrington (1995), are very fragile to these perturbations. Their equilibria are such that everyone cooperates until the moment when there is a (single) deviation. At this point, there is a quick contagion and the society ends up with everyone defeating

each other. Eventually, using public signals, everyone may simultaneously return to cooperation. So, deviations need to be rare even at the aggregate level. If the number of players is large, even individually small probabilities of a deviation happening (trembles, for example) at each period may lead to a high aggregate probability, and this may be 1 if there are behavioral types. In this case, the equilibria mentioned would fail. We will see that our equilibria exist even when deviations happen for sure every period.

Note that equilibria in affine strategies generate smooth evolution of the aggregate variables. As we saw in the last section, even when there is a large portion of the society playing S in a given period, $p_{t+\Delta}$ is, on expectation, different from p_t only in a $O(\Delta)$ term. Also, independently of the initial state p_0 of the economy, we have convergence to a long run steady state. Therefore, our equilibrium is more likely to be robust to small perturbations.

In the next sections we discuss briefly three perturbations of the standard game and establish some new conditions for the existence of equilibria in affine strategies in the form of propositions. Since, as we will see, the results are similar, at the end of the section we discuss them and provide some interpretation. Again, for simplicity we will consider the prisoner's dilemma, although the results are very similar for the dictator game.

4.1 Behavioral Types

Many models in the literature of games with a large number of agents and no information are very sensitive to the existence of behavioral types. Indeed, even if a single agent in the economy defects at every period, equilibria like those proposed in Kandori (1992), Ellison (1994) and Harrington (1995) fail. The problem is that since agents can masquerade as behavioral, deviations have to be punished. Nevertheless, the existence of behavioral types implies that too harsh punishments destroy all cooperation in the society. We will see that our model is robust to this perturbation.

Assume that a fraction $\phi \Delta$ of the players have behavioral (or action) types. We assume that behavioral types play E with probability $\pi \in [0, 1]$ independently of their history.¹⁵ Let $n_\phi \equiv N \phi \Delta$ be the number of players that have behavioral types. The evolution of p_t^{-i} is now given by

$$\mathbb{E}[p_{t+\Delta}^{-i} | p_t^i, p_t^{-i}] = \Lambda(p_t^i, p_t^{-i}) + \frac{n_\phi}{N-1} \left(\pi - \Psi\left(\pi, \frac{(N-1)p_t^{-i} + p_t^i - \pi}{N-1}\right) \right), \quad (4.1)$$

¹⁵We could assume that there is a mass $\phi \Delta$ of behavioral types whose mixing probability ν is distributed according to some F . Nevertheless, affinity implies that this is equivalent to all of them having the same mixing probability equal to $\pi \equiv \mathbb{E}[\nu]$.

where Λ is defined in (2.7) and Ψ in (2.17). The intuition behind the previous formula is the following. If there were no behavioral players the term in the RHS of (4.1) would be just $\Lambda(p_t^i, p_t^{-i})$ (using equation (2.6)). Nevertheless, since behavioral types do not update their state, an adjustment term has to be included. This term, whose weight is $\frac{n_\phi}{N-1}$, removes the adjustment that the base model implies for their individual states and sets them to π again.

The following proposition establishes the conditions under which equilibria in affine strategies exist for this case:

Proposition 4.1. *A non-trivial equilibrium in affine strategies exists if*

$$\gamma > \gamma_*(\phi) \equiv \frac{\ell(\phi + \rho)}{g - \ell}. \quad (4.2)$$

and does not exist if $\gamma < \gamma_*(\phi)$.¹⁶ In this case, the optimal equilibrium has $\eta = 0$ (no memory). The asymptotic expected fraction of players playing E under the optimal equilibrium, p_∞^* , is given by

$$p_\infty^* = 1 - \frac{(1 - \pi)\phi\ell}{(g - \ell)\gamma - \ell\rho} + O(\Delta). \quad (4.3)$$

Proof. The proof is in the appendix on page 29. □

Remark 4.1. (forgiveness delegation) The strategies that we find in our model require agents to “forgive” (that is, play E when they receive S) with a very small probability (since $k_S = O(\Delta)$). In the real world it may be difficult to individually find devices to generate these small probabilities. Instead, societies may reach the efficient equilibria by “generating” a small fraction of forgivers (behavioral types) that play E with a very high probability independently of the action they received. The rest of the players may then play a pure tit-for-tat, 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.

¹⁶As in Lemma 3.1 when $\gamma = \gamma_*(\phi)$ there is an additional condition for existence, that can be found if the model is solved for finite Δ and then the limit $\Delta \rightarrow 0$ is taken. Since this is a non-generic case that does not have any economic relevance, we do not extend on this.

4.2 Entry/Exit

Societies are usually not completely closed, so new agents/players enter (are born) and exit (die) constantly. Therefore, in order to explain cooperation in societies, it is important to know the effects of inflows/outflows of agents in our model.

We now assume that in each period after playing the stage game, $n_\phi \equiv \phi \Delta N$ players exit the game (die), and they are replaced by n_ϕ new players (are born).¹⁷ The entrants are aware of the equilibrium that is played, but they do not have any previous history. Note that affine strategies are a function of the previous history only through an individual state, p_t^i . Therefore, we assume that new players are given an initial state ν drawn from a distribution F . As for the behavioral types case above, affinity is equivalent to all entering agents having an individual state 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 $\phi \Delta$ of them die. Finally, new agents are born with individual state π , and the overall population remains constant. This timing allows us to write the following expression for the evolution of p_t^{-i} :

$$\mathbb{E}[p_{t+\Delta}^{-i} | p_t^i, p_t^{-i}] = f(\Lambda(p_t^i, p_t^{-i})) \quad (4.4)$$

where

$$f(p_t^{-i}) \equiv \frac{N-1-n_\phi}{N-1} p_t^{-i} + \frac{n_\phi}{N-1} \pi,$$

and where Λ is defined in (2.7) using (2.17). Intuitively, conditional on player i surviving period t , $\mathbb{E}_t[p_{t+\Delta}^{-i}]$ 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 p_t^i and p_t^{-i} . This implies that each player has an equivalent discount factor of $(1 - \Delta \phi) \delta$, to account for the possibility of dying. Therefore, we have:

Proposition 4.2. *The same result as in Proposition 4.1 holds for entry and exit of players, replacing all discount factors ρ by the equivalent discount factor $\rho + \phi$.*

Proof. The proof is analogous to the proof of Proposition 4.1. □

¹⁷For simplicity we assume that the number of players that exit is certain. It is easy to see that, by affinity, our results are the same if, instead, we assume that each player has an independent probability of leaving equal to $\phi \Delta$. This could be more realistic in some interpretations.

4.3 Trembles

Consider finally the case where players may tremble when deciding on their action. That is, if a player plays E , his opponent perceives S with probability $\varepsilon_1 \Delta$, and if he plays S , his opponent perceives E with probability $\varepsilon_2 \Delta$. For simplicity, we assume that none of the players know if a tremble has happened when they act and observe their opponent's action.

Fix a player $i \in \mathcal{N}$. Let $f(p_t^i)$ be, for a given mixed stage action p_t^i , the expected probability that his opponent will perceive the action as E . This is given by

$$f(p_t^i) = (1 - \varepsilon_1 \Delta) p_t^i + (1 - p_t^i) \varepsilon_2 \Delta .$$

Note that $f(\cdot)$ is affine. Therefore, we can write

$$\mathbb{E}[p_{t+\Delta}^{-i} | p_t^i, p_t^{-i}] = \frac{\Psi(p_t^{-i}, f(p_t^i))}{N-1} + \frac{N-2}{N-1} \Psi(p_t^{-i}, f(p_t^{-i})) ,$$

where Ψ is defined in (2.17). Intuitively, we change the formula for Λ (given in (2.7)) to incorporate mistakes in the action received. Again, we have that $\mathbb{E}[p_{t+\Delta}^{-i} | p_t^i, p_t^{-i}]$ is jointly affine in p_t^i and p_t^{-i} . This allows us to proceed similarly to the previous cases, so we have:

Proposition 4.3. *The same result as in Proposition 4.1 holds for trembles, replacing ϕ by $\varepsilon_1 + \varepsilon_2$ and π by $\frac{\varepsilon_2}{\varepsilon_1 + \varepsilon_2}$.*

Proof. The proof is analogous to the proof of Proposition 4.1. □

4.4 Comparison

Now let's see why the three extensions of the model discussed in this section lead to similar results and analyze them. The intuition will be provided using the optimal equilibrium, since it is the one that leads to the bound on γ and the expression for p_∞ . This equilibrium is especially helpful to get intuitive results because it is a one-period memory equilibrium, where the action played in the next period is only a function of the action received in this period.

The relationship between the perturbations of entry/exit and behavioral types is obvious when the optimal equilibrium is considered. Indeed, in the entry/exit case, in each period a fraction $\phi \Delta$ of the agents die and are reborn with state π . Since the state at the next period is only a function of the action received in this period, this is equivalent to picking these players randomly (entry/exit case) or letting them always be the same (behavioral types).

For the trembles case, note that a fraction $\varepsilon_2 \Delta$ of players who are supposed to receive S (and have their state be almost 0) end up receiving E (moving their state to 1). In the entry/exit case, instead, a fraction $\phi \Delta$ of the players who received S (and have their state to be almost 0) die and are replaced by players with state π . Affinity implies that changing the state of a fraction $\varepsilon_2 \Delta$ from 0 to 1 is equivalent to changing the state of a fraction $\phi \Delta$ to $\pi = \frac{\varepsilon_2}{\phi}$. Doing the same for the other players (players who are supposed to receive E), we get the transformation rule of Proposition 4.3.

Hence, we see that affine strategies can easily be adapted to perturbations when these are included in the model. Notice that when perturbations are included the condition for equilibria to exist is more restrictive than in the not-perturbed model but still holds generically. Indeed, if 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. This requires a fine balance between punishing and forgiving since, in our model, punishment is performed by keeping deviations in the game for a time long enough that the deviator will be worse off. Therefore, if punishments are too harsh, punishing mistakes would destroy the equilibria, while too soft punishments would generate incentives to shirk, so that the deviation masquerades as a mistake.

We also see that asymptotic efficiency is in general no longer reachable. This is true except if the perturbation is in the “correct” direction. Indeed, if $\pi = 1$ in the entry/exit and the behavioral types extensions, we have that we can still achieve efficiency. It is also the case when $\varepsilon_1 = 0$ in the trembles case.¹⁸

5 Conclusions

Sustaining cooperation in large societies requires balancing incentives to defect and to forgive defections. By introducing affine strategies, we are able to balance these incentives history by history. Therefore, by constructing explicit equilibria, we found that cooperation may be both smoothly evolving over time and robust to many perturbations. Hence, our model is a step forward toward understanding how anonymous help may take place in our societies.

The key feature of our equilibria is that affine strategies make aggregate dynamics affine. This allows us to find belief-free equilibria, simplifying the players’ incentive constraints. We found that

¹⁸When $\pi = 1$ or $\varepsilon_1 = 0$ the threat to the existence of equilibria comes from the excessive forgiveness of the society, not from the inability to distinguish deviations from perturbations. Indeed, if a (normal) player plays S when he is supposed to play E , under $\pi = 1$ or $\varepsilon_1 = 0$ there will be a higher probability that this deviation is “forgiven” before it influences his future payoff.

even when the game is perturbed, there is an optimal equilibrium that maximizes the payoff for all priors. It is extremely simple, since it features a one-period memory for the prisoner's dilemma case. In the dictator game, the optimal equilibrium also has one-period memory when the receiver role is played.

If we introduce perturbations in the model such as trembles or behavioral types, the existence of equilibria remains generic. We develop a technique to easily get formulae for very large societies under these perturbations. Indeed, in the joint limit where the number of players increases and the length of the time interval goes to 0, equilibria are characterized by a set of differential equations that can be easily solved. Equilibrium strategies are close to a tit-for-tat in this limit.

Future research will be devoted to extending our results to other games, such as the (sequential game) ultimatum game. Also, we will study the effect of incorporating more players in the stage game, what leads to non-affine dynamics.

References

- Deb, J., 2008, "Cooperation and Community Responsibility: A Folk Theorem for Repeated Matching Games with Names," 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, J. E., and Jr., 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.

6 Appendix

Proof of Theorem 2.1 (page 7)

Proof. We will prove that a symmetric strategy profile that yields an affine law of motion satisfies $\mathbb{E}_t[\alpha_{t+1}^i | a_t^j] = \Psi(\alpha_t^i, a_t^j)$, for some $\Psi(\cdot, \cdot)$ jointly affine (the reverse implication is trivial), and that this implies (2.7).

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

$$\mathbb{E}_t[\bar{\alpha}_{t+1}^{-i} | a_t^i] = \frac{\sum_{j \neq i} M_t^j \bar{\alpha}_t^{-i}}{N-1} + \frac{\sum_{j \neq i} M_t^j (a_t^i - \alpha_t^j)}{(N-1)^2}.$$

Assume $(\alpha^i)_{i \in \mathcal{N}}$ yields an affine 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 $\mathbb{E}_t[\alpha^j(h^{j,t}) | h^{j,t}, a_t^k] \neq \mathbb{E}_t[\alpha^j(h^{j,t'}) | h^{j,t'}, a_{t'}^k]$ for all j (recall that the strategy profile is symmetric) for some $a_t^k = a_{t'}^k$. This implies that $M_t^j(h^{j,t}) \neq M_{t'}^j(h^{j,t'})$. 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

$$\begin{aligned} \mathbb{E}[\bar{\alpha}_{t+1}^{-i} | h^t, a_t^i] &= \frac{1}{N-1} M_t^j(h^{j,t}) a_t^i + \frac{N-2}{N-1} M_t^j(h^{j,t}) \bar{\alpha}_t^{-i}(h^{j,t}), \\ \mathbb{E}[\bar{\alpha}_{t+1}^{-i} | h^{t'}, a_{t'}^i] &= \frac{1}{N-1} M_{t'}^j(h^{j,t'}) a_{t'}^i + \frac{N-2}{N-1} M_{t'}^j(h^{j,t'}) \bar{\alpha}_t^{-i}(h^{j,t'}). \end{aligned}$$

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 (2.6) 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 α_t^j is the action received by player i at t and $\Psi : \Delta(A) \times A \rightarrow \Delta(A)$. As usual, we affinely extend Ψ to $\Delta(A) \times \Delta(A)$. Then

$$\Lambda(\alpha^i, \alpha^{-i}) = \sum_{j \neq i} \left(\frac{\Psi(\alpha^j, \alpha^i)}{(N-1)^2} + \sum_{k \neq i, j} \frac{\Psi(\alpha^j, \alpha^k)}{(N-1)^2} \right) = \sum_{j \neq i} \frac{\Psi(\alpha^j, \alpha^{-j})}{N-1}. \quad (6.1)$$

Fix $j \neq i$. Note that, by assumption, the LHS of the previous expression is affine in α^j (since Λ is affine and α^{-i} is affine in α^j). Furthermore, all terms in the sum of the RHS are affine 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 affine in the first argument.

To prove the joint affinity of Ψ fix $j \neq i$. Note that, by assumption, the LHS of (6.1) is jointly affine 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 affine in (α^i, α^j) . Therefore, $\Psi(\alpha^j, \alpha^i)$ is equal to a sum of terms affine in (α^i, α^j) , and therefore it is itself affine. Therefore the first part of the theorem holds.

To prove (2.7), note that the terms of the middle equality in (6.1) can now be simplified using the fact that, as we have just proven, Ψ is jointly affine. It is easy to see that the first term coincides with $\frac{\Psi(\bar{\alpha}^{-i}, \alpha^j)}{N-1}$. 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 Lemma 2.1 (page 11)

The proof relies on constructing an strategy analogous to the 2-actions strategy (2.10). Playing E is identified with playing a^{SO} and playing S with playing a^{LE} . Therefore, p_t^i should be interpreted as the probability of playing a^{SO} and $1 - p_t^i$ as the probability of playing a^{LE} . Also, in (2.10), $k_{a_t^j}$ is the same for all actions different from a^{SO} . In this case, since a^{EL} provides the highest myopic payoff, it is easy to see that there is no incentives to deviate to any other action, since the continuation play is independent of the action played. Equation (2.14) is then just a generalization of (2.12) to arbitrary payoffs.

Proof of Proposition 2.3 (page 12)

Proof. Assume freedom from belief and non-triviality. This implies that there are at least two actions where agent i is indifferent to play. Assume a pure strategy $\bar{\alpha}^{-i} \in A$. Then, if player i plays $a^i \in A$, there is no uncertainty about the distribution in the next period (given by $\Lambda(a^i, \bar{\alpha}^{-i})$),

so we have

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

The previous equation must be valid for at least two values of a^i and all affine combinations of them. Since LHS does not depend on a^i and u is affine, then V^i must be affine. Furthermore, since V^i is affine and Λ is jointly affine, u must be jointly affine, and therefore the proposition holds. \square

Proof of Proposition 2.4 (page 13)

Proof. Since $V(\cdot)$ is affine by Proposition 2.3, let's denote $V(x) \equiv V_0 + V_1 x$. Since players are indifferent between playing E and S for all p_t^{-i} , it must be the case that:

$$(1 - \delta) (-\ell + g p_t^{-i}) + \delta V(\Lambda(p_t^{-i}, 1)) = (1 - \delta) (g p_t^{-i}) + \delta V(\Lambda(p_t^{-i}, 0)) \quad (6.2)$$

for all p_t^{-i} . Using (2.7) and the form of the strategies (2.17), it is easy to see that the coefficient on p_t^{-i} on each side of (6.2) is the same. So, we have that

$$V_1 = \frac{\ell(N - 1)(1 - \delta)}{(k_E - k_S)\delta} . \quad (6.3)$$

The second equation that must be satisfied in equilibrium is the recursive expression of the value function V (Bellman equation, (2.5)):

$$V(p_t^{-i}) = (1 - \delta) (-\ell a + g p_t^{-i}) + \delta V(\Lambda(p_t^{-i}, a)) , \quad (6.4)$$

for both $a = 0, 1$. Condition (6.2) ensures that the RHS is independent of a . Then, since both V and Λ are affine, each side is affine in p_t^{-i} , so (6.4) is a system of two equations (one for the intercept and the other for the slope in p_t^{-i}). Equating the constant terms on each side we obtain

$$V_0 = \frac{\ell(N - 1)k_S}{k_E - k_S} . \quad (6.5)$$

Equating now the terms on p_t^{-i} of each side and using (6.3) we have

$$N = 1 + \frac{(g - \ell)\delta(k_E - k_S)}{\ell(1 - \delta(k_E + \eta - k_S))} . \quad (6.6)$$

Given our parametric restrictions, it is clear that the right-hand side reaches its maximum when $k_E = 1$ and $k_S = \eta = 0$. Using these values, we find the condition (2.19). \square

Proof of Proposition 2.5 (page 14)

Proof. Given the proof of Proposition 2.4, the set of parameters for an equilibrium for the prisoners' dilemma is given by

$$\mathcal{E}^{PD} \equiv \{ (\eta, k_E, k_S) \mid (k_E, k_S, k_E + \eta, k_S + \eta) \in [0, 1]^4 \text{ and (2.18) holds} \} .$$

(this analogous to the definition for the unilateral help case at (2.13)).

We will first find the subset of equilibrium parameters that maximize V_0 . Assume that $(\eta, k_E, k_S) \in \mathcal{E}^{PD}$ is such that the corresponding V_0 (defined in (6.5)) is maximized among all elements of \mathcal{E}^{PD} . From equation (6.6) we know that $k_E > k_S$. Furthermore, note that if $k_E < 1$ and $k_E + \eta < 1$ then we can increase both k_E and k_S in the same (small) amount, leaving η constant and increasing V_0 . Therefore, either $k_E = 1$ or $k_E + \eta = 1$. Assume first that $\eta < 0$, so $k_E = 1$. Keeping $k_E = 1$ constant, it is easy to see that both η and k_S can be increased in (different) small amounts such that (6.6) still holds and V_0 increases. Assume that $\eta > 0$ so $k_E + \eta = 1$. Now it is easy to see that we can decrease η and increase k_E in the same amount and increase k_S such that (6.6) still holds and V_0 increases. So, the maximum value of V_0 is achieved when $\eta = 0$ and $k_E = 1$.

Note that $V(1) \equiv V_0 + V_1$ is maximized when $\eta + k_E = 1$. Indeed, in this case, if $p_t^{-i} = 1$ and $p_t^i = 1$, they both remain constant forever (if no one deviates). This provides the highest possible equilibrium payoff. Since both $V(0) \equiv V_0$ and V_1 are maximized if and only if $\eta = 0$ and $k_E = 1$, $V(p)$ is also only maximized under these parameter values for all p . This proves the proposition. \square

Proof of Lemma 3.1 (page 15)

Proof. We can easily prove this result using the condition (2.19) for the existence of affine equilibria. Indeed, this expression can be rewritten for Δ_n as

$$\frac{\gamma - \gamma_*}{\gamma} \geq \Delta_n (\rho - \gamma_*) .$$

It is clear that if $\gamma > \gamma_*$, there exists a $\bar{\Delta} > 0$ such that all $0 < \Delta \leq \bar{\Delta}$ satisfy the previous condition. In the case $\gamma = \gamma_*$ the previous condition can only hold for some $\Delta_n > 0$ if $\rho \leq \gamma_*$, in which case it is satisfied for all $\Delta_n > 0$. Finally, if $\gamma < \gamma_*$, there is always some $\bar{\Delta} > 0$ such that for all $0 < \Delta \leq \bar{\Delta}$ the previous condition is not satisfied. That is exactly the condition stated in the proposition. \square

Proof of Proposition 4.1 (page 20)

Proof. It is not difficult but very long and tedious to get expressions for finite Δ like (2.19). Furthermore, when there is entry and exit of players the expressions are long and difficult to interpret. Instead, we will use the technique developed in Section 3 to easily get the limiting first-order expressions when Δ is small. It is left to the reader to get the expressions for finite Δ and explicitly verify that (4.2) and (4.3) are true for small Δ .

The Bellman equation (3.4) remains the same. The main difference now is the expression of the dynamics of p_t , which now, instead of (3.3), is given by

$$\dot{p}_{BT}^{-i}(p_t^i, p_t^{-i}) \equiv \lim_{\Delta \rightarrow 0} \frac{\mathbb{E}[p_{t+\Delta}^{-i} | p_t^i, p_t^{-i}] - p_t^{-i}}{\Delta} = \dot{p}(p_t^i, p_t^{-i}) + (\pi - p_t^{-i}) \phi, \quad (6.7)$$

where $\dot{p}(p_t^i, p_t^{-i})$ is the law of motion for the basic model, defined in (3.3). We see that we have an extra term that pushes p_t^{-i} towards π with strength ϕ . Note that the \dot{p}_{BT}^{-i} is still affine in (p_t^i, p_t^{-i}) . Using this expression in (3.4) and imposing indifference we find

$$\tilde{k}_E + \tilde{k}_S - \tilde{\eta} = \frac{g\gamma - \ell(\gamma + \rho + \phi)}{\ell}. \quad (6.8)$$

Now, imposing $\eta \in [-\tilde{k}_S, \tilde{k}_E]$ and $\tilde{k}_E, \tilde{k}_S \in \mathbb{R}_+$, it is easy to see that the expression (4.2) is necessary and sufficient for the existence of equilibria.

From (3.4) and (6.7) we can find $V^i(p)$, which is given by

$$V^i(p) = \frac{\ell\rho}{\gamma} p + \frac{g\gamma - \ell(\gamma + \rho + (1 - \pi)\phi + \tilde{k}_E - \tilde{\eta})}{\gamma}.$$

Note that $V^i(p)$ is maximized (for all p) when $\tilde{k}_E - \tilde{\eta}$ is minimized. Given the restriction that $k_E \in \mathbb{R}_+$ and $\tilde{\eta} \in [-\tilde{k}_S, \tilde{k}_E]$, we have $\tilde{k}_E - \tilde{\eta} \geq 0$. So, setting $\tilde{k}_E = \tilde{\eta}$ (and \tilde{k}_S solving (6.8)) maximizes $V^i(p)$. The corresponding value for p_∞^* is found by imposing $\dot{p}(p_\infty^*, p_\infty^*) = 0$, and we get (4.3). \square