Repeated Games with Endogenous Discounting*

Asen Kochov   Yangwei Song†

September 20, 2016

The paper studies infinitely repeated games in which the players' rates of time preference may evolve over time, depending on what transpires in the game. A key result is that in any first best equilibrium of the repeated prisoners' dilemma, the players must eventually cooperate. If we assume that the players become more patient as they obtain better outcomes, we show that cooperation prevails from the beginning of the game and is thus the unique outcome of any first best equilibrium. The latter result is suitably extended to all symmetric two player games. A separate contribution of the paper is to propose a framework in which intertemporal trade can emerge as a first best equilibrium of a repeated strategic interaction, generating predictions that differ from those in the standard framework.

KEYWORDS: Repeated games, efficiency, folk theorems, endogenous discounting.

*We are grateful to Paulo Barelli, John Duggan, and Hari Govindan for their helpful comments. Special thanks go to Takuo Sugaya for referring us to the work of Lehrer and Pauzner. The paper also serves as our response to William Thomson who likes to joke that, while fairness sometimes implies efficiency, the converse is never true.

†Department of Economics, University of Rochester, Rochester, NY 14627. E-mail: asen.kochov@rochester.edu, ysong19@ur.rochester.edu

1 Introduction

Strong restrictions on the structure of intertemporal preferences are a common feature in the study of repeated games. In fact, most of the literature assumes that preferences can be represented by a time-additive utility function with a constant rate of time preference. This paper considers a more general class of intertemporal preferences whereby the rate of time preference is endogenous. Specifically, the discounted sum of payoffs takes the recursive form:

$$v_i(a^0, a^1, ...) = g_i(a^0) + \beta_i(a^0) v_i(a^1, a^2, ...). \tag{1}$$

Above, $g_i(a)$ is player i's stage payoff from an action profile a and $\beta_i(a)$ is the player's discount factor as a function of the action profile. We refer to repeated games in which the intertemporal preferences of the players take this form as games with endogenous discounting.[1]

The primary question we ask is: How much of our intuition about repeated interactions relies on the standard preference specification? Consider the classical prisoners' dilemma game. One implication of the folk theorem for repeated games is that cooperation can be sustained in equilibrium provided that the players are sufficiently patient. A more troublesome implication is that the concept of equilibrium loses much of its predictive power. First, there are many equilibria in which the players fail to coordinate on an efficient outcome. There are also many efficient equilibria in which cooperation does not prevail. Such equilibria can generate starkly disparate outcomes. In some, the continuation utility of one of the players is held arbitrarily close to his security level throughout the entire game.

As we step away from the standard preference specification, the following results emerge. First, the paper confirms the familiar conclusion that a wide range of outcomes can arise in equilibrium. In particular, say that a play path $(a^0, a^1, ...)$ is **sequentially individually rational** if the continuation utility of every player exceeds his security level at every point

[1] In a deterministic environment, a special case of the representation in (1) was first introduced by Uzawa [30]. Epstein [9] studied a stochastic framework and showed that the utility function over deterministic outcome streams takes the form in (1) whenever preferences are stationary and random outcome streams are evaluated by the expected utility criterion. The preferences we use have also been employed in the macroeconomic literature and in the study of small open economies in particular. See Backus et al. [3] for a survey. The reader may also benefit from comparing our work with that of Obara and Park [23]. They consider repeated games with time additive but nonstationary preferences, as is the case when the players are hyperbolic discounters. The different preferences they use raise different conceptual problems and lead to different results.

in time. Our folk theorem shows that any such play path can arise in equilibrium if the players are sufficiently patient. The more surprising result is that in every first best equilibrium, the players must eventually cooperate. In fact, for an intuitive specification of the preferences we consider, cooperation prevails from the beginning of the game. Together, the requirements of efficiency and sequential individual rationality can thus be extremely powerful, selecting cooperation as the unique outcome of the repeated prisoners' dilemma.

In the literature on endogenous discounting, it is common to assume that discount factors are a strictly monotonic function of the underlying payoffs. While this assumption is not necessary for our folk theorem, it is the key requirement under which we investigate the interplay between efficiency and sequential individual rationality. There are two cases to consider. The first is when the marginal impatience of each player i, $1 - \beta_i(a)$, decreases the more desirable he finds the constant path $(a, a, ...)$. Informally, one can think of this assumption as capturing the often heard claim that 'time flies faster when people are happy'. The polar case of increasing marginal impatience is defined analogously. One can think of it as capturing a situation in which the promise of high present rewards tempts the person and 'blinds' him to what the future may bring. It should be pointed out that the merits of each case have been debated going back to the classical works of Fisher et al. [11, p.72] and Friedman [12, p.30]. See also Lucas and Stokey [22], Backus et al. [3], and Epstein [9, 10]. In this paper, we do not take sides in the ongoing debate whether increasing or decreasing marginal impatience is a more appropriate restriction on behavior. The two cases are investigated in turn, obtaining qualitatively different results. If marginal impatience is decreasing, we show that cooperation is the outcome of any first best equilibrium. If marginal impatience is increasing, the uniqueness result is less stark: for every first best equilibrium there is some date T after which cooperation prevails. As we explain below however, an interesting aspect of this case is that 'cooperation' can take a novel form.

1.1 Decreasing Marginal Impatience

To gain some intuition for the uniqueness result obtained under decreasing marginal impatience, it is helpful to first consider a dynamic general equilibrium model in which preferences are standard but individuals differ in their rates of time preference. An old conjecture by Ramsey [27] asserts that the most patient individual would eventually possess all the capital. The conjecture was confirmed by Becker [5]. His arguments rely on the fact that competitive markets are efficient. In particular, efficiency mandates that im-patient individuals be rewarded in the early stages of the interaction, whereas their more patient counterparts be rewarded in the latter stages. Lehrer and Pauzner [21] have recently considered how Ramsey's conjecture may play out in a strategic setting. Looking at a repeated prisoners' dilemma game, their work points to a stark tension between the requirements of efficiency and sequential individual rationality. Consistent with Ramsey's conjecture, efficiency once again requires that the utility of the more patient player is eventually maximized. But in a repeated prisoners' dilemma game, this is possible only if the other player is pushed below his security level, violating sequential individual rationality. One reaches the important conclusion that first best outcomes cannot be sustained in an equilibrium of the game, no matter how patient the players. In contrast to Lehrer and Pauzner [21], suppose now that the players are a priori symmetric but marginal impatience is decreasing rather than fixed. If, in the course of such a game, one of the players is promised a high payoff, he will exhibit a greater degree of patience. Efficiency then requires that the same player be rewarded relatively more in the future, which in turn insures that he will sustain the higher level of patience as the game progresses. A self-enforcing dynamic is therefore created that once again pushes one of the players to his most preferred point on the Pareto frontier, while the other player is pushed to 'immiseration'. What makes the games we study different from the games in Lehrer and Pauzner [21], however, is that this dynamic is not inevitable. Suppose the players, who are a priori symmetric, coordinate on a play path along which both of them attain symmetric outcomes. Then, their rates of time preference remain identical and the 'Ramsey dynamic' is never triggered. In the prisoners' dilemma game, this means that every first best equilibrium induces a unique play path whereby the players cooperate in every period.

While the preceding discussion has focused on the prisoners' dilemma, we should emphasize that a suitable generalization is obtained for all symmetric two player games. See Theorem 8.1. When marginal impatience is increasing, the case we preview next, we have been unable to obtain a similar generalization and we limit attention to the prisoners' dilemma.

1.2 Increasing Marginal Impatience

It was observed already that when marginal impatience is increasing, the uniqueness result we obtain is less stark. Rather than beginning immediately, cooperation prevails after some period T. What is interesting about this case, however, is that 'cooperation' may take a different form: Depending on the preference parameters, the players may cooperate by alternating between their most preferred outcomes in the stage game. In the prisoners' dilemma game, this means that the players take turns defecting. We refer to this outcome as one of **intertemporal cooperation**. There are two (related) reasons why this result may be of interest. Observe first that in the benchmark model with time additive utilities and identical discount factors, intertemporal cooperation is never an efficient outcome.[2] Thus, to the extent that one is willing to restrict attention to first best equilibria, the preferences we consider generate different dynamics and lead to new testable implications.

Intertemporal cooperation is an instance of what may be more broadly termed **intertemporal trade**: a mutually advantageous exchange that is driven by differences in the rate of time preference. This observation, which we clarify in Section 7, takes us to the second reason why we believe the results under increasing marginal impatience may be of interest. There is little doubt that intertemporal trade is a feature of many real world situations. Heterogeneity in observed rates of time preference, whether exogenous or endogenous, is well documented. See, for instance, Becker and Mulligan [4], Lawrance [20], and Harrison et al. [18]. While not all of the evidence supports the assumption of increasing marginal impatience, the evidence does suggest that opportunities for intertemporal trade exist. Lehrer and Pauzner [21] provide one way of incorporating such opportunities within a standard strategic situation. This paper provides another, with qualitatively different results.[3]

2 The Model

Time is discrete and varies over an infinite horizon: $t \in \{0, 1, ...\} =: \mathcal{T}$. There is a finite set of players: $I := \{1, 2, ..., n\}$. In each period t, player i can choose a pure action a_i in a finite set A_i. Let $A := \times_{i \in I} A_i$. Player i's mixed actions are denoted by $\alpha_i \in \Delta(A_i)$. We permit public randomization: at each stage the players can condition their actions on an exogenous random variable. A complete history up to period $t \in \mathcal{T}$ consists of all the past mixed actions and public signals. We assume perfect monitoring: each player can

[2] Throughout, we focus on prisoners' dilemma games such that if preferences are standard, cooperation (of the usual form) is the unique symmetric first best outcome. See equation (9) in Section 7 and the discussion in Section C .3.

[3] Building on Lehrer and Pauzner [21], Obara and Zincenko [24] and Opp and Zhu [25] show how familiar results in the study of oligopoly pricing and moral hazard may change once intertemporal trade is allowed for. Our results differ from those in Lehrer and Pauzner [21] and may thus prove useful in applications as well.

condition his action at time t on the entire history.[4] Let Σ_i denote the corresponding set of behavioral strategies for player $i \in I$ and let $\Sigma := \times_{i\in I}\Sigma_i$. A generic strategy profile is denoted as $\sigma := (\sigma_i)_i \in \Sigma$. A play path $\mathbf{a} = (a^0, a^1, ...) \in A^\infty$ is a sequence of pure action profiles. Given a play path $\mathbf{a} = (a^0, a^1, ...) \in A^\infty$ and a time period $t \in \mathcal{T}$, ${}_t\mathbf{a}$ denotes the continuation path $(a^t, a^{t+1}, ...)$ starting from period t. To describe player i's preferences, focus first on the ranking of pure play paths. Define a utility function $v_i : A^\infty \to \mathbb{R}$ as follows:

$$v_i(\mathbf{a}) := g_i(a^0) + \beta_i(a^0)g_i(a^1) + \beta_i(a^0)\beta_i(a^1)g_i(a^2) + ... = g_i(a^0) + \beta_i(a^0)v_i({}_1\mathbf{a}), \tag{2}$$

where $g_i : A \to \mathbb{R}$ is player i's stage payoff and $\beta_i : A \to (0,1)$ is his discount factor. Given (2), preferences are extended to random strategy profiles in the usual manner. In particular, note that each strategy profile $\sigma \in \Sigma$ induces a probability distribution on A^∞. Abusing notation, we denote the induced measure by σ as well. Player i's expected payoff from a strategy profile σ is then $v_i(\sigma) := \mathbb{E}_\sigma v_i(\mathbf{a})$. Note that, if each $\beta_i : A \to (0,1)$ is a constant function, one obtains the standard time additive model with a constant rate of time preference.

A repeated game with endogenous discounting is a tuple $(A, (g_i, \beta_i)_{i\in I})$. By **equilibrium**, we always mean a subgame perfect equilibrium that induces a pure play path. An equilibrium is **efficient** or **first best** if it induces a play path that is not Pareto dominated by any other, potentially random play path. For every $i \in I$, player i's **security level** or **minmax payoff** is the number $\underline{v}_i := \inf_{\sigma_{-i}\in\times_{k\neq i}\Sigma_k} \sup_{\sigma_i\in\Sigma_i} v_i(\sigma_i, \sigma_{-i})$. Lemma 3.1 in the next section shows that the inf and sup can be replaced by a min and max respectively. Thus, $\underline{v}_i$ becomes the lowest level of ex ante utility at which player i can be forced to by his opponents. Drawing an analogy with the mechanism design literature - an analogy we employ throughout the paper - one can think of the minmax payoffs $\underline{v}_i$ as imposing a 'participation constraint' on any solution of the game. More formally, one can adapt the arguments in Fudenberg and Tirole [16, p.151] to show that in every Nash equilibrium of the repeated game player i's ex ante utility is at least $\underline{v}_i$. A **minmax strategy against player** i is a strategy profile $\sigma_{-i} \in \times_{k\neq i}\Sigma_k$ that attains the infimum in the definition of $\underline{v}_i$.[5]

[4] In two player games, we can relax the assumption that players observe the *mixed* actions played in the past. With more than two players, we need to strengthen the richness assumption made in Section 5 if Theorem 5.1 is to go through. Details are available upon request. On the other hand, we do not know if one can obtain our folk theorem without the use of public randomization. The techniques developed by Sorin [28] and Fudenberg and Maskin [15] are hard to apply because of the nonadditivity of the preferences we study.

[5] We adhere to the convention that, when minmaxing player i, his opponents cannot condition their choices on a correlating device whose outcome is unobserved by player i. One can then show that it is w.l.o.g. to assume that players $k \neq i$ choose their strategies independently of one another as in the definition of

3 Minmax Strategies

One problem in analyzing games with endogenous discounting is that the minmax strategies against each player and, hence, the implied security levels may change as discount factors converge to one. To understand why this is problematic, recall the usual intuition behind the folk theorem: cooperation can be sustained by the threat of future punishments. If the players are sufficiently patient, such punishments outweigh the short-term gains from a deviation. In fact, there is a monotonic relationship: the more patient the players, the more costly the punishments loom. An implicit assumption in this argument is that what is a 'punishment' can be chosen independently of the players' rates of time preference. It is a priori less clear that the same intuition applies if the punishments vary with the rate of time preference and, in particular, if they become less severe as patience increases. We address this problem in this and the next section. The main observation is that time preferences are summarized by the functions $\beta_i : A \to (0,1)$, which map action profiles into absolute levels of patience, and these functions may approach 1 in many different ways. If the convergence path is chosen appropriately, a choice which we justify in Sections 4 and E , then the problem is effectively circumvented: the minmax strategies against each player and, hence, the implied security levels do not vary with the level of patience.

Our first result shows that the search for minmax strategies can be restricted to strategies that are constant. The result relies on the fact that the preferences we employ are stationary. Intuitively, stationarity implies that the same 'punishment' works no matter at which node of the game tree the players find themselves. The result contributes to our broader goal which is to investigate how familiar results in the study of repeated games depend on the decision theoretic assumptions that govern behavior. To state the result, say that a strategy $\sigma_i, i \in I$, is **constant** if the mixed action $\alpha_i \in \Delta(A_i)$ is played after every history and denote such a strategy by α_i^{con}. A strategy profile $\sigma := (\sigma_1, ..., \sigma_n) \in \Sigma$ is **constant** if each σ_i is constant.

Lemma 3.1. *For every player $i \in I$,*

$$\inf_{\sigma_{-i} \in \times_{k \neq i} \Sigma_k} \sup_{\sigma_i \in \Sigma_i} v_i(\sigma_i, \sigma_{-i}) = \min_{\alpha_{-i} \in \times_{k \neq i} \Delta(A_k)} \max_{\alpha_i \in \Delta(A_i)} v_i(\alpha_i^{con}, \alpha_{-i}^{con}). \quad (3)$$

For future reference, it is helpful to compute the ex ante payoffs from a constant strategy profile. Let $\Delta(A)$ be the space of probability measures on A. For every mixed action pro-

security levels and minmax strategies we have employed. See the related discussion in Fudenberg and Maskin [14, p.536].

file $\alpha \in \Delta(A)$, let α^{con} be the constant strategy profile that prescribes α after every history. For every $i \in I$ and $\alpha \in \Delta(A)$, let $g_i(\alpha) := \sum_{a \in A} g_i(a)\alpha(a)$, and $\beta_i(\alpha) := \sum_{a \in A} \beta_i(a)\alpha(a)$, where $\alpha(a)$ is the probability assigned to action profile a under α. Since each constant strategy induces an IID probability measure on A^∞, the expected ex ante payoff from a constant strategy is equal to the expected payoff after any given history. In particular, we have:

$$v_i(\alpha^{con}) = \mathbb{E}_\alpha[g_i(a) + \beta_i(a)v_i(\alpha^{con})] = g_i(\alpha) + \beta_i(\alpha)v_i(\alpha^{con}) \quad \Leftrightarrow$$
$$v_i(\alpha^{con}) = \frac{g_i(\alpha)}{1 - \beta_i(\alpha)}. \tag{4}$$

Observe that when each β_i is a constant function, as it is in the standard model, the ranking of constant strategies is completely determined by the stage payoffs g_i. This is not so for the preferences we study. As is clear from (4), the level of patience attained by a player may now depend on the play path, which in turn determines the overall utility derived from it.

4 How Discount Factors Converge to 1

The next step is to specify how discount factors converge to 1. Since changes in the discount factors β_i trigger changes in the repeated game, it is convenient to introduce the two types of changes simultaneously. Thus, let $\Gamma(0)$ be a repeated game $(A, (g_i, \beta_{i,0})_{i \in I})$, where we think of the discount factors $\beta_{i,0}$ as capturing a low, initial level of patience. For every $\lambda \in [0, 1)$, let

$$\beta_{i,\lambda} := \lambda + (1 - \lambda)\beta_{i,0} \quad \forall i \in I, \tag{5}$$

and let $\Gamma(\lambda)$ denote the repeated game $(A, (g_i, \beta_{i,\lambda})_{i \in I})$. Given $(g_i, \beta_{i,\lambda})$, player i's utility function over deterministic play paths and its extension to mixed play paths are denoted by $v_{i,\lambda}$. We use Γ to denote the family $\{\Gamma(\lambda) : \lambda \in [0, 1)\}$ of repeated games, although, with a slight abuse of terminology, we sometimes refer to Γ simply as a game. It is clear from (5) that $\beta_{i,\lambda} > \beta_{i,\lambda'}$ for every i and every $\lambda > \lambda'$.[6] Subsequently, we interpret λ as an absolute measure of the players' patience and consider equilibrium behavior as λ converges to one.

[6] Given $i \in I$ and $\lambda \in [0, 1)$, $\beta_{i,\lambda} > \beta_{i,\lambda'}$ means that $\beta_{i,\lambda}(a) > \beta_{i,\lambda'}(a)$ for every $a \in A$. Similarly, when we say that the function $\beta_{i,\lambda}$ converges to 1 as λ converges to 1, we mean that $\beta_{i,\lambda}(a)$ converges to 1 for every $a \in A$.

Whenever a game Γ is specified, we maintain two normalizations as well. First, we scale the utility functions $v_{i,\lambda}$ by multiplying each one of them by $1-\lambda$. As in the benchmark model, in which one multiplies the discounted sum of payoffs by '$(1-\delta)$', this normalization insures that the utilities $v_{i,\lambda}$ do not blow up as the players become increasingly patient. With this normalization, the utility functions $v_{i,\lambda} : A^\infty \to \mathbb{R}$ can be written recursively as:

$$v_{i,\lambda}(\mathbf{a}) = (1-\lambda)g_i(a^0) + \beta_{i,\lambda}(a^0)v_{i,\lambda}({}_1\mathbf{a}). \tag{6}$$

Second, we assume that utilities are normalized so that the security levels $\underline{v}_{i,\lambda}$ are all equal to zero.

The convergence path in (5) has two properties that bear on the subsequent analysis and should be emphasized. The first property, **P1**, is that the ranking of constant strategies, indeed the utility of such strategies, is independent of λ. To see this, plug (5) into (6) to obtain:

$$v_{i,\lambda}(\alpha^{con}) = \frac{g_i(\alpha)}{1-\beta_{i,0}(\alpha)} \quad \forall\lambda \in [0,1), \forall i \in I, \forall \alpha \in \Delta(A), \tag{7}$$

where, as before, $\beta_{i,0}(\alpha) := \sum_{a\in A} \beta_{i,0}(a)\alpha(a)$. It follows immediately from Lemma 3.1 that the minmax strategies against a player and his best response can be chosen independently of λ. As we pointed out earlier, this implication of P1 facilitates the proof of our folk theorem.

The second property, **P2**, is captured by the equality:

$$\frac{1-\beta_{i,\lambda}(a)}{1-\beta_{j,\lambda}(a')} = \frac{1-\beta_{i,0}(a)}{1-\beta_{j,0}(a')} \quad \forall\lambda \in [0,1), i, j \in I, a, a' \in A.$$

The key observation is that the fraction on the right hand side is independent of λ. Thus, the relative impatiences, across players and action profiles, remain fixed as we make the players more patient in absolute terms. Keeping fixed the relative (im)patiences *across players* is a requirement familiar from the literature on repeated games with fixed but unequal discounting. See Compte and Jehiel [8], Lehrer and Pauzner [21], Sugaya [29], and Chen and Takahashi [7], among others.[7] Keeping fixed the relative impatiences *across action profiles* is a requirement specific to this paper and the study of endogenous discounting in particular. It plays a similar role, however, in that it insures that qualitative

[7] Some of these papers differ in whether one keeps the relative impatiences or the relative patiences fixed. Sugaya [29, p.23] shows that this choice becomes immaterial as discount factors converge to 1.

features of the game are preserved as we vary the absolute level of patience. The assumptions of increasing, decreasing marginal impatience which we employ in Sections 7 and 8 fall into this category. We should also point out that properties P1 and P2 effectively characterize the convergence path given by equation (5). Section E in the appendix provides the details.

5 The Folk Theorem

Suppose preferences conform to the standard time additive specification. From Abreu et al. [1] and Chen and Takahashi [7], it is understood that a folk theorem holds if no two players have identical preferences on A^∞. Except in some special cases, this sufficient condition is also necessary. For reasons we explain below, we need a stronger condition. To state it, fix a family Γ of repeated games and note that Γ can be identified with the tuple $(A, (g_i, \beta_{i,0})_{i\in I})$. Also, recall that utilities are normalized so that the security levels are all zero.

Definition 1. The family Γ satisfies **Richness** if for every $i \in I$, there exist action profiles $a^i, \tilde{a}^i \in A$ such that for all $j \neq i$, we have (i) $g_j(a^i) > 0 \geq g_i(a^i)$, and (ii) $g_i(\tilde{a}^i) > 0$ and $g_j(\tilde{a}^i) \geq 0$.

It is helpful to clarify what Richness implies for each of the repeated games $\Gamma(\lambda)$. Consider, for example, part (i) of Richness. The action profile $a^i \in A$ is such that for all $j \neq i$ and λ,

$$v_{j,\lambda}(a^i, a^i, ...) > 0 \geq v_{i,\lambda}(a^i, a^i, ...). \tag{8}$$

Thus, the constant play path $(a^i, a^i, ...)$ 'rewards' every player $j \neq i$ by giving him an ex ante payoff higher than his security level. At the same time, it 'punishes' player i by giving him no more than his security level. The role of such player-specific rewards in proving a folk theorem is well understood: if player i is to deviate, the other players must be given an incentive to carry out a punishment. Where Richness 'bites' is in insuring that we can implement such rewards using constant play paths. This yields two advantages. First, given the convergence path we specified in Section 4, we can vary the level of patience λ, knowing that the incentives in place remain 'incentives'. Second, we do not have to worry about endogenous changes in the players' level of patience along the play path.[8]

[8]The reader may wonder whether we could have replaced Richness with the weaker requirement that the

For a path $\mathbf{a} \in A^\infty$ to arise in equilibrium, it is necessary that at each point in time the continuation payoff of every player exceeds his security level. The folk theorem we state below shows that this requirement, which we refer to as sequential individual rationality, is also sufficient provided that the players are sufficiently patient. A well understood caveat is that we need the continuation payoffs of a path to not only clear the participation constraints of each player but to do so uniformly. This leads us to the following definition. For every $\varepsilon > 0$ and $\lambda \in [0, 1)$, a path $\mathbf{a} \in A^\infty$ is **ε-sequentially individually rational** if $v_{i,\lambda}({}_t\mathbf{a}) \geq \varepsilon$ for all $i \in I, t \in \mathcal{T}$. Let $SIR^\varepsilon(\lambda)$ be the set of all such paths. When $\varepsilon = 0$, we say simply that a path $\mathbf{a} \in SIR^0(\lambda)$ is sequentially individually rational. For future reference, it is also helpful to formalize the weaker notion of individual rationality which requires that the participation constraints of the players are met only at the beginning of the game. Thus, a path $\mathbf{a} \in A^\infty$ is **individually rational given** λ if $v_{i,\lambda}(\mathbf{a}) \geq 0$ for every $i \in I$. Given $\varepsilon > 0$ and λ, $IR^\varepsilon(\lambda)$ denotes the set of all paths $\mathbf{a} \in A^\infty$ such that $v_{i,\lambda}(\mathbf{a}) \geq \varepsilon$ for every $i \in I$.

Theorem 5.1. *Suppose* Γ *satisfies Richness. For every* $\varepsilon > 0$*, there exists* $\underline{\lambda} \in [0, 1)$ *such that for all* $\underline{\lambda} < \lambda < 1$*, every path* $\mathbf{a} \in SIR^\varepsilon(\lambda)$ *can be supported in an equilibrium of the game* $\Gamma(\lambda)$*. In two player games, the same conclusion obtains whether or not Richness holds.*

We discuss two venues for future research. One is to understand how far Richness can be relaxed. The work of Chen and Takahashi [7] is the most relevant in this respect. The paper studies games with fixed but unequal discounting. It shows how to construct equilibrium strategies by exploiting differences in the players' rate of time preference *even when* the players have identical stage payoffs.[9] In such cases, player-specific rewards can no longer be constructed using constant strategies. Instead, players are incentivized to punish a deviator by entering an intertemporal trade whose 'terms' are tilted in their favor. We do not know how to construct such trades when discounting is endogenous, that is, when differences in the players' rates of time preference may change along the play path.

Another extension for future work is to allow for imperfect monitoring. A recent paper by Sugaya [29], which studies games with fixed but unequal discounting, is the most rel-

set of payoffs attainable by constant strategies have full dimension. The problem with this requirement is that, with endogenous discounting, the latter set need not be convex, making full dimensionality a less tractable requirement to work with. By comparison, Richness imposes a more particular structure on that set, which we exploit in the construction of player-specific rewards. See Lemma B 2 in the appendix. It should here be noted that the nonconvexity arises even though we allow for public randomization. This is because, with endogenous discounting, the functions $v_i(\alpha^{con})$ in (4) are nonlinear in the probabilities. It should also be noted that the set of *all* feasible payoffs *is* convex. Working with that set, however, would amount to working with nonconstant strategies, which as we explain above raises its own set of challenges.

[9] See Chen [6] and Guéron et al. [17] as well.

evant in this respect. Building on Fudenberg and Levine [13], the paper characterizes the set of payoffs attainable in a public and perfect equilibrium as discount factors converge to one. A key ingredient in Sugaya's characterization is the recursivity of the standard preference specification. To the extent that the preferences we study are recursive as well, it is reasonable to conjecture that Sugaya's results can be extended to a setting with endogenous discounting.

6 Efficiency

We turn to the main goal of this paper which is to investigate if efficient outcomes can be supported in equilibrium. As we explained in the introduction, the analysis is split in two cases, depending on whether marginal impatience is increasing or decreasing. Below we introduce these assumptions formally, as well as some additional assumptions and notation.

In the rest of the paper, we require that each player's discount factor depends on the action profile only through the player's stage payoff. Keeping the definition of Γ into account, this means that there are functions $f_i : \mathbb{R} \to (0,1)$, $i \in I$, such that $\beta_{i,0}(a) = f_i(g_i(a))$. This assumption is common in the applied, macroeconomic literature on endogenous discounting.[10] It was also maintained by Uzawa [30] with whom the literature on endogenous discounting began. The assumption is not essential, but it simplifies our analysis.

A game Γ is said to be **symmetric** if, in addition to the usual restrictions on the actions sets A_i and the stage payoffs g_i, we have $f_i = f_j$ for all $i, j \in I$. From now on, we restrict attention to two player, symmetric games. The symmetry restriction is particularly important in that it rules out any *a priori* heterogeneity in the way the players discount the future. Such heterogeneity can only emerge *endogenously*, if different players attain different outcomes in the course of the game. In this respect, the games we study are qualitatively different from the games studied by Lehrer and Pauzner [21], in which the heterogeneity is exogenous. We will also show that the switch from exogenous to endoge-

[10] A notable exception is the paper by Becker and Mulligan [4], in which individuals can exert effort to increase their level of patience, e.g., by spending more time thinking about the future. We can model such situations within the more general setting of Section 2, but not once we assume that the discount factors β_i depend on action profiles only through the stage payoffs g_i. It is also worth noting that the analysis in Becker and Mulligan [4] provides a "microfoundation" for the assumption of decreasing marginal impatience, which other papers typically take for given, as we do in Section 8. See Becker and Mulligan [4] for details.

nous heterogeneity matters: it has important implications for the efficient equilibria of a game.

To formalize the assumptions of increasing and decreasing marginal impatience, fix Γ and recall from Section 4 that the ranking of constant play paths does not depend on the level of patience λ. In particular, $v_{i,\lambda}(a, a, ...) = \frac{g_i(a)}{1-\beta_{i,0}(a)}$ for every $\lambda \in [0,1), a \in A$ and $i \in I$. The game Γ is said to satisfy **increasing marginal impatience** if for all $i \in I$ and all $a, a' \in A$,

$$\frac{g_i(a)}{1-\beta_{i,0}(a)} > \frac{g_i(a')}{1-\beta_{i,0}(a')} \quad \text{if and only if} \quad \beta_{i,0}(a) < \beta_{i,0}(a').$$

Decreasing marginal impatience is defined by reversing one of the two inequalities above.

Given Γ, one can show that if a play path $\mathbf{a} \in A^\infty$ maximizes the ex ante utility of player i for some level of patience λ, then it does so for all λ.[11] To simplify the exposition, we assume that for every player i, there is a unique play path $\mathbf{a}^{max,i}$ that maximizes his ex ante utility.

Some notation is needed as well. For every Γ and every level of patience $\lambda \in [0,1)$, $P(\lambda)$ denotes the set of all efficient paths $\mathbf{a} \in A^\infty$, while $P_{sym}(\lambda)$ is the subset of paths that maximize the sum of the players' utilities, that is, $P_{sym}(\lambda) := \text{argmax}_{\mathbf{a}\in A^\infty} v_{1,\lambda}(\mathbf{a}) + v_{2,\lambda}(\mathbf{a})$. An equilibrium that induces a play path in $P(\lambda)$ is called a **first best equilibrium**. Finally, $P_{++}(\lambda)$ is the set of all efficient paths other than $\mathbf{a}^{max,1}$ and $\mathbf{a}^{max,2}$.

7 Increasing Marginal Impatience

This section considers the case when marginal impatience is increasing. We focus on the repeated prisoners' dilemma game and obtain two key results. The first one shows that a new form of cooperation, namely one in which the players take turns defecting, may become a first best outcome of the game. The efficiency of this outcome is due to differences in the players' rates of time preference that arise endogenously, in the course of the game. As we explain, these differences imply that the outcome may be viewed as a form of intertemporal trade between the players. The second result shows that cooperation, of the new or usual form, prevails eventually in any first best equilibrium of the repeated

[11] If the optimal path for a given player is unique, one can show that the path is necessarily constant. If there are multiple optimal paths, then some of them are necessarily constant, while all the others are obtained by alternating among the action profiles $a \in A$ that characterize the optimal constant paths. This is true so long as preferences on A^∞ are stationary and continuous. See Kochov [19, Lemma 1] for a proof. The fact that the optimal paths for a given player do not depend on the level of patience follows directly from these observations.

game.

7.1 The Prisoners' Dilemma

There are two players: $I = \{1, 2\}$. The action space A and the stage payoffs $g_1, g_2 : A \to \mathbb{R}$ are as depicted in Figure 1, where, as usual, C stands for 'cooperate' and D for 'defect'. To define the family Γ of repeated games, let $\beta_0 : \{b, 0, c, d\} \to (0, 1)$ be a function associating each possible stage payoff with a level of patience. Also, let $\beta_\lambda := \lambda + (1 - \lambda)\beta_0$ for every λ. Given λ, player i's discount factor $\beta_{i,\lambda} : A \to (0, 1)$ is then defined by letting $\beta_{i,\lambda}(a) := \beta_\lambda(g_i(a))$ for every $a \in A$. This construction of the discount factors $\beta_{i,\lambda}$ insures that they depend on the stage payoffs only and that this dependence is identical across the players.

	C	D
C	c, c	b, d
D	d, b	$0, 0$

Figure 1: The prisoners' dilemma

As is typical in a prisoners' dilemma game, we assume that $\frac{d}{1-\beta_0(d)} > \frac{c}{1-\beta_0(c)} > 0 > \frac{b}{1-\beta_0(b)}$. Note that these inequalities are ordinal restrictions on preferences. For instance, the first one says that each player prefers the constant path in which he defects and the other player cooperates to the play path in which both players cooperate. We also assume that

$$\frac{c}{1 - \beta_0(c)} > \frac{1}{2}\frac{b}{1 - \beta_0(b)} + \frac{1}{2}\frac{d}{1 - \beta_0(d)}. \tag{9}$$

The inequality says that each player prefers cooperation in every period to the mixed play path that yields his least preferred or his most preferred play path with equal probability.[12] This assumption is maintained in order to highlight the different predictions that endogenous discounting may generate relative to the standard model. Specifically, if (9) holds within the standard model, then $((C, C), (C, C), ...)$ is the unique play path that maximizes the sum of the players' utilities. As we are about to show, this need not be the

[12]The reader may wonder why we define the prisoners' dilemma game in terms of the payoffs attained from constant play paths rather than in terms of the stage payoffs $\{b, 0, c, d\}$. We give an answer in Section E of the appendix. Here, we should mention that the two approaches are equivalent in the standard model with fixed discounting.

case when discounting is endogenous, that is, when there are gains from intertemporal trade.

7.2 Two Forms of Cooperation

We now formalize the two forms of cooperation that characterize the efficient outcomes of the repeated prisoners' dilemma game and what it means for cooperation to obtain eventually.

The first form of cooperation requires little discussion: It is given by the path $\mathbf{a}^C := ((C,C),(C,C),...)$, along which (C,C) is played in every period. We refer to this path as one of **intratemporal cooperation**. To introduce the paths along which cooperation prevails eventually, first let $\mathcal{C}_1$ be the set of paths such that (D,C) is played in at most one period while (C,C) is played in all other periods. The subscript '1' is used to designate the fact that the action profile (D,C), if it occurs, favors player 1. Analogously, let $\mathcal{C}_2$ be the set of paths such that (C,D) is played in at most one period while (C,C) is played in all other periods. Next, let $\mathcal{E}_1\mathcal{C}_1$ be the set of paths $\mathbf{a} \in A^\infty$ such that for some $T \geq 0$, depending on the path, $a^t = (D,C)$ for all $t < T$ and ${}_T\mathbf{a} \in \mathcal{C}_1$. Here, the letter $\mathcal{E}$ is mnemonic for the fact that cooperation prevails *eventually*, that is, after some period. The subscript '1' once again designates the fact that all action profiles that are not (C,C) favor player 1. Similarly, define $\mathcal{E}_2\mathcal{C}_2$ to be the set of paths such that (C,D) is played until some period and the continuation path is in $\mathcal{C}_2$. Note that $\mathcal{C}_i \subseteq \mathcal{E}_i\mathcal{C}_i$, $i \in \{1,2\}$. Finally, let $\mathcal{EC} := \mathcal{E}_1\mathcal{C}_1 \cup \mathcal{E}_2\mathcal{C}_2$.

We now define the set of paths which we identify with **intertemporal cooperation**. In particular, let

$$\mathcal{A} := \{\mathbf{a} \in A^\infty : a^{2t} \in \{(D,C),(C,D)\}, a^{2t+1} \in \{(D,C),(C,D)\} \setminus \{a^{2t}\}, \forall t \in \mathcal{T}\}.$$

To understand the definition of $\mathcal{A}$, focus on the pair $((D,C),(C,D))$ of action profiles and its symmetric counterpart $((C,D),(D,C))$. When such a pair designates the behavior of the players in two consecutive time periods, $2t$ and $2t+1$, we interpret the pair as an 'exchange of favors' or an instance of 'intertemporal trade'. A path $\mathbf{a} \in \mathcal{A}$ is then any sequence of such exchanges. For future reference, we wish to single out two special paths in $\mathcal{A}$. In particular, let $\mathbf{a}^{A,1}$ be the alternating path in $\mathcal{A}$ such that $a^{2t} = (D,C)$ for every $t = 0,1,....$ This is the path in which every exchange favors player 1: he is the first to receive what he desires; the other player must wait a period. The path $\mathbf{a}^{A,2}$ is defined analogously.

One may wonder whether we are justified in interpreting a path $\mathbf{a} \in \mathcal{A}$ as a sequence of *separate* exchanges. Shouldn't the whole path be interpreted as a single intertemporal trade instead? To address this question, consider the path $\mathbf{a}^{A,2}$ and the utility each player derives from it :

$$\frac{1}{1-\lambda} v_{1,\lambda}(\mathbf{a}^{A,2}) = b + \boldsymbol{\beta}_{\lambda}(\boldsymbol{b}) d + \boldsymbol{\beta}_{\lambda}(\boldsymbol{b}) \boldsymbol{\beta}_{\lambda}(\boldsymbol{d}) b + ...$$
$$\frac{1}{1-\lambda} v_{2,\lambda}(\mathbf{a}^{A,2}) = d + \boldsymbol{\beta}_{\lambda}(\boldsymbol{d}) b + \boldsymbol{\beta}_{\lambda}(\boldsymbol{d}) \boldsymbol{\beta}_{\lambda}(\boldsymbol{b}) d + ...$$

Notice first that $\beta_\lambda(b) > \beta_\lambda(d)$. This difference in the rate of time preference is what makes the initial exchange of favors appealing to both players and what, as Theorem 7.1 shows, may 'elevate' such alternating paths to first best outcomes of the repeated game. Another important observation, and the one that directly addresses the question at hand, is that both players apply an identical discount factor to their stage payoffs in period $t = 2$: $\beta_\lambda(b)\beta_\lambda(d)$. Thus, after the initial exchange of favors, the ex ante symmetry between the players is restored. They evaluate the subsequent future 'with fresh eyes', ready to enter another exchange. Similarly, after the second exchange takes place in periods $t = 2$ and $t = 3$, the players' discount factors are once again equalised, and so on. One can also verify that these observations apply to *any* of the alternating paths $\mathbf{a} \in \mathcal{A}$, not just $\mathbf{a}^{A,2}$. For this reason we interpret the play paths in $\mathcal{A}$ as sequences of separate, standalone exchanges.

We should also explain why we are justified in interpreting the paths in $\mathcal{A}$ as a form of cooperation, on par with the path $\mathbf{a}^C$. The key observation is that a social planner who maximizes the sum of the players' payoffs, $v_{1,\lambda} + v_{2,\lambda}$, is indifferent between any two paths in $\mathcal{A}$:

$$v_{1,\lambda}(\mathbf{a}) + v_{2,\lambda}(\mathbf{a}) = (1-\lambda)\frac{b + \beta_\lambda(b)d + d + \beta_\lambda(d)b}{1 - \beta_\lambda(b)\beta_\lambda(d)} \quad \forall \lambda \in [0,1), \forall \mathbf{a} \in \mathcal{A}.$$

Thus, whenever the play paths in $\mathcal{A}$ are efficient, each one of them is an optimal choice for a social planner who does not discriminate between the players (assuming away the possibility to randomize). It should also be noted that, as the players become more and more patient, it matters less which player is favored in any of the intertemporal exchanges that comprise a path $\mathbf{a} \in \mathcal{A}$. In particular, $\lim_{\lambda \to 1} v_{1,\lambda}(\mathbf{a}) = \lim_{\lambda \to 1} v_{2,\lambda}(\mathbf{a})$ for every $\mathbf{a} \in \mathcal{A}$.[13]

[13]The ex ante payoff vectors $(v_1, v_2) \in \mathbb{R}^2$ generated by the play paths in $\mathcal{A}$ lie on a linear segment perpendicular to the direction $\eta_{sym} := (1,1)$. As $\lambda \to 1$, this segment collapses to a single point on the 45-degree line.

It remains to introduce the class of play paths along which intertemporal cooperation obtains eventually. Adopting similar notational conventions as before, let $\mathcal{E}_1\mathcal{A}$ be the set of all play paths $\mathbf{a} \in A^\infty$ such that for some $T \geq 0$, depending on the path, $a^t = (D, C)$ for all $t < T$, and ${}_T\mathbf{a} \in \mathcal{A}$. Define $\mathcal{E}_2\mathcal{A}$ analogously. Note that $\mathcal{A} \subseteq \mathcal{E}_i\mathcal{A}$, $i \in \{1, 2\}$ and let $\mathcal{E}\mathcal{A} := \mathcal{E}_1\mathcal{A} \cup \mathcal{E}_2\mathcal{A}$.

7.3 Results

One final observation is needed before we can state our main result of Section 7. Say that a level of patience λ is **irregular** if intra- and inter- temporal cooperation are both efficient, that is, if $\mathbf{a}^C, \mathbf{a}^{A,1}, \mathbf{a}^{A,2} \in P(\lambda)$. Else, λ is **regular**. The proof of Theorem 7.1 makes clear that irregular λ are not only 'rare' but they are bounded away from 1 as well. Hence, such λ become irrelevant as we make the players sufficiently patient. Since high levels of patience are needed anyhow, if one is to sustain first best outcomes in equilibrium, we view the irregular λ as less important for our purposes and restrict attention to regular λ when we state Theorem 7.1. A characterization of the Pareto set $P(\lambda)$ for irregular λ is provided in Section C .4 of the appendix. Here, we should note that the main message of Theorem 7.1 remains valid: along every efficient path in $P_{++}(\lambda)$ the players cooperate eventually.

Theorem 7.1. *Fix a regular $\lambda \in [0, 1)$. For every efficient path $\mathbf{a} \in A^\infty$, other than the extremes $\mathbf{a}^{max,1}$ and $\mathbf{a}^{max,2}$, there is some $T \geq 0$ such that the continuation path ${}_T\mathbf{a}$ is one of cooperation. Whether the cooperation is intra- or inter-temporal depends on the preferences but not on the path $\mathbf{a}$. Specifically, if $\mathbf{a}^C \in P(\lambda)$, then $P(\lambda) = \mathcal{E}\mathcal{C} \cup \{\mathbf{a}^{max,1}, \mathbf{a}^{max,2}\}$. If $\mathbf{a}^{A,1} \in P(\lambda)$, then $P(\lambda) = \mathcal{E}\mathcal{A} \cup \{\mathbf{a}^{max,1}, \mathbf{a}^{max,2}\}$. In the first case, $\mathbf{a}^C$ is the only path that maximizes the sum of the players' payoffs, that is, $P_{sym}(\lambda) = \{\mathbf{a}^C\}$. In the second, any of the alternating paths does so, that is, $P_{sym}(\lambda) = \mathcal{A}$.*

Theorem 7.1 characterizes the set of efficient paths. The question remains if such paths can be sustained in a subgame perfect equilibrium. The usual 'folk wisdom' suggests that the answer is affirmative provided that the players are sufficiently patient. An important lesson from the work of Lehrer and Pauzner [21] is that this intuition becomes circumspect when the players differ in their levels of patience. In fact, their work points to a direct conflict between the requirements of efficiency and sequential individual rationality. What we show next is that such a conflict need not arise when differences in the rate of time preference are endogenous rather than exogenous. In fact, under increasing marginal impatience, the opposite is true: Efficiency and individual rationality imply sequential individual rationality. This is evident from the structure of the Pareto set $P(\lambda)$: since cooperation prevails eventually and since cooperation, whether intratemporal or intertemporal, is individually rational, the participation constraints of the players are more likely to be satisfied as time progresses. We summarize these observations in the next corollary.[14]

Corollary 7.1. *Let $\underline{\lambda} \in [0,1)$ be the smallest λ such that the path $((C,D),(C,C),(C,C),...)$ is individually rational. Then, for every $\underline{\lambda} < \lambda < 1$ and every $\varepsilon > 0$ small enough, if $\mathbf{a} \in P(\lambda) \cap IR^{\varepsilon}(\lambda)$, then $\mathbf{a} \in SIR^{\varepsilon}(\lambda)$.*

The next and final result of this section confirms that first best outcomes can be sustained in equilibrium provided the players are sufficiently patient. Rather than restating the folk theorem from Section 5 however, we opt for a slightly different formulation. In particular, observe that in Theorem 5.1, we first fixed a threshold ε by which a path $\mathbf{a} \in A^{\infty}$ must clear the participation constraints of each player. We then found a patience level λ for which all such paths can be sustained in equilibrium. In applications, it is often helpful to answer the converse question, that is, which paths can be sustained for a given level of patience? We answer this question as follows. First, we fix λ high enough so that some efficient play path can be sustained. We then provide a lower bound on the ex ante payoff of an efficient path $\mathbf{a} \in A^{\infty}$ which, if met, insures that the path can be sustained in equilibrium.

Corollary 7.2. *There exists $\underline{\lambda} \in [0,1)$ such that for every $\underline{\lambda} < \lambda < 1$ and every $\varepsilon > (1-\lambda)d$, every path $\mathbf{a} \in P(\lambda) \cap IR^{\varepsilon}(\lambda)$ can be supported in an equilibrium of the game.*[15]

A final question is whether intertemporal cooperation can be efficient in the limit, as $\lambda \to 1$ and the players become completely patient. In Section C .3 we provide a numerical example showing that this is possible. In fact, in the example every path $\mathbf{a} \in \mathcal{A}$ Pareto dominates the path $\mathbf{a}^C = ((C,C),(C,C),...)$ in the limit, as well as for all λ sufficiently close to one.

[14] In its broad strokes, Theorem 7.1 follows in the footsteps of Lucas and Stokey [22] who show how in a general equilibrium context assuming increasing marginal impatience may overturn the immiseration result of Becker [5], while still allowing for (endogenous) heterogeneity in the rate of time preference. The specifics, however, such as the form of intertemporal cooperation we obtain, have no counterpart in Lucas and Stokey [22].

[15] The result does not deliver a complete characterization of all efficient paths that are sustainable given some λ: the bound $(1-\lambda)d$ is not exact. Observe however that $(1-\lambda)d$ converges to 0 as λ converges to 1. Thus, we reach the familiar conclusion that in the limit every efficient, (sequentially) individually rational path can be sustained.

8 Decreasing Marginal Impatience

Suppose now that marginal impatience is decreasing. The next theorem applies to all symmetric two player games and shows that the efficient play paths take one of two forms.

Theorem 8.1. *Fix a two player, symmetric game* Γ *and a level of patience* $\lambda \in [0, 1)$*. For every efficient path* $\mathbf{a} \in P(\lambda)$*, one of two things must be true: either there is some* $T \geq 0$ *such that the continuation path* ${}_T\mathbf{a}$ *is the most preferred path for one of the players, that is,* ${}_T\mathbf{a} \in \{\mathbf{a}^{\max,1}, \mathbf{a}^{\max,2}\}$*, or the players attain identical stage payoffs along the entire path, that is,* $g_1(a^t) = g_2(a^t)$ *for every* $t \in \mathcal{T}$*.*

Intuition for Theorem 8.1 was given in the introduction. If one of the players is promised a high outcome at some stage of the game, he will exhibit a greater level of patience when evaluating the future. Efficiency insures that this triggers a 'Ramsey-like' dynamic whose end result is that the player will come to dominate the game, eventually obtaining his most preferred outcome. The only way this can be avoided is if the players coordinate on a play path along which they attain identical payoffs and, hence, identical levels of patience.[16]

The conclusions of Theorem 8.1 become especially powerful when we consider the prisoners' dilemma game from Section 7. Then, Theorem 8.1 tells us that every efficient path is either $\mathbf{a}^C$ or it culminates in one of the players' most preferred play paths, that is, there is some T such that ${}_T\mathbf{a} \in \{\mathbf{a}^{max,1}, \mathbf{a}^{max,2}\}$. The key observation is that in the prisoners' dilemma game the path $\mathbf{a}^{max,i}$ is not individually rational for player $j \neq i$. This leaves us with $\mathbf{a}^C = ((C, C), (C, C), ...)$ as the only efficient path that can arise in an equilibrium of the game.

Corollary 8.1. *The path* $((C, C), (C, C), ...)$ *is the only path that can arise in a first best equilibrium of the prisoners' dilemma game. Such an equilibrium exists provided that the players are sufficiently patient.*

[16]It should also be noted that the logic behind Theorem 8.1 is different from the logic behind the results in Section 7. In the latter case, there was no conflict between efficiency and sequential individual rationality.

Appendix

A Proof of Lemma 3.1

Endow A with the discrete topology and A^∞ with the product topology. Since the set A is finite, A and A^∞ are separable Borel spaces. Endow $\Delta(A_i)$ and $\Delta(A)$ with the usual Euclidean topology. Let $\Delta(A^\infty)$ be the set of Borel probability measures on A^∞. Endow $\Delta(A^\infty)$ with the weak* topology. Endow each Σ_i and Σ with the respective product topologies.

From Kolmogorov's Consistency Theorem, see Parthasarathy [26, Theorem 3.1,3.2], we know that each strategy profile $\sigma \in \Sigma$ induces a probability measure in $\Delta(A^\infty)$. Moreover, the mapping from Σ into $\Delta(A^\infty)$ is continuous. Conclude that player i's utility function $v_i : \Sigma \to \mathbb{R}$ is continuous. Fix some $i \in I$. Since the space Σ_i is compact, the maximization problem

$$\max_{\sigma_i \in \Sigma_i} v_i(\sigma_i, \sigma_{-i})$$

has a solution for every $\sigma_{-i} \in \times_{k \neq i} \Sigma_k$. By the maximum theorem, see Aliprantis and Border [2, Theorem 17.31], $\max_{\sigma_i \in \Sigma_i} v_i(\sigma_i, \sigma_{-i})$ is a continuous function of σ_{-i}. Thus, the minmax payoff

$$\underline{v}_i = \min_{\sigma_{-i} \in \times_{k \neq i} \Sigma_k} \max_{\sigma_i \in \Sigma_i} v_i(\sigma_i, \sigma_{-i})$$

is well defined.

Lemma A 1. *For every $i \in I$ and $\alpha_{-i} \in \times_{k \neq i} \Delta(A_k)$,*

$$\max_{\sigma_i \in \Sigma_i} v_i(\sigma_i, \alpha_{-i}^{con}) = \max_{\alpha_i \in \Delta(A_i)} v_i(\alpha_i^{con}, \alpha_{-i}^{con}).$$

Proof. Fix $i \in I$ and $\alpha_{-i} \in \times_{k \neq i} \Delta(A_k)$. Let $\hat{\alpha}_i \in \operatorname{argmax}_{\alpha_i \in \Delta(A_i)} v_i(\alpha_i^{con}, \alpha_{-i}^{con})$. Note that $\hat{\alpha}_i$ exists since $\Delta(A_i)$ is compact and the function $\alpha_i \mapsto v(\alpha_i^{con}, \alpha_{-i}^{con})$ is continuous. Suppose by way of contradiction that there exists a strategy $\tilde{\sigma}_i \in \Sigma_i$ such that $v_i(\tilde{\sigma}_i, \alpha_{-i}^{con}) > v_i(\hat{\alpha}_i^{con}, \alpha_{-i}^{con})$. By the one shot deviation principle, we may assume that the strategy $\tilde{\sigma}_i$ prescribes the mixed action $\hat{\alpha}_i$ after all but a single history h^t. Moreover, since players $j \neq i$ are using a constant strategy we may assume that h^t is the initial (empty) history. Let $\tilde{\alpha}_i \neq \hat{\alpha}_i$ be the initial action prescribed by $\tilde{\sigma}_i$. Then,

$$v_i(\tilde{\sigma}_i, \alpha_{-i}^{con}) = g_i(\tilde{\alpha}_i, \alpha_{-i}) + \beta_i(\tilde{\alpha}_i, \alpha_{-i}) v_i(\hat{\alpha}_i^{con}, \alpha_{-i}^{con}) > v_i(\hat{\alpha}_i^{con}, \alpha_{-i}^{con}),$$

which, after some rearranging, yields

$$v_i(\tilde{\alpha}_i^{con}, \alpha_{-i}^{con}) = \frac{g_i(\tilde{\alpha}_i, \alpha_{-i})}{1 - \beta_i(\tilde{\alpha}_i, \alpha_{-i})} > v_i(\hat{\alpha}_i^{con}, \alpha_{-i}^{con}).$$

The last inequality contradicts the fact that $\hat{\alpha}_i \in \operatorname{argmax}_{\alpha_i \in \Delta(A_i)} v_i(\alpha_i^{con}, \alpha_{-i}^{con})$. □

Proof of Lemma 3.1. Fix some $i \in I$. Say that a strategy σ_i is **finite** if there is some t such that after every history h^t player i plays a constant strategy. A strategy profile $\sigma = (\sigma_1, ..., \sigma_n) \in \Sigma$ is **finite** if every σ_j is finite. Let $\underline{v}_i^c$ be the minmax payoff of player i when all players are restricted to using constant strategies. By Lemma A 1, the best reply to a constant strategy is a constant strategy. Deduce that $\underline{v}_i \leq \underline{v}_i^c$. To prove the opposite inequality, let $\sigma_{-i} \in \times_{k \neq i} \Sigma_k$ be a minmax strategy against player i and let $\sigma_i \in \Sigma_i$ be a best reply. For every t, let σ_{-i}^t be a finite strategy that coincides with σ_{-i} up to and including period t. By construction, $\sigma_{-i}^t \to_t \sigma_{-i}$. Since the best reply to a constant strategy is constant, a standard backward induction argument shows that each σ_{-i}^t has a best reply $\sigma_i^t \in \Sigma_i$ that is a finite strategy. Using the maximum theorem, deduce that

$$\lim_t v_i(\sigma_i^t, \sigma_{-i}^t) = \lim_t \max_{\sigma_i' \in \Sigma_i} v_i(\sigma_i', \sigma_{-i}^t) = \max_{\sigma_i' \in \Sigma_i} v_i(\sigma_i', \sigma_{-i}) = \underline{v}_i.$$

To show that $\underline{v}_i \geq \underline{v}_i^c$, it is therefore enough to show that $v_i(\sigma_i^t, \sigma_{-i}^t) \geq \underline{v}_i^c$ for every t. Fix some $t \geq 1$ and consider a history h^{t-1}. Let $(\alpha_i, \alpha_{-i}) \in \Delta(A)$ be the action profile prescribed by $(\sigma_i^t, \sigma_{-i}^t)$ at h^{t-1}. For every $a \in A$, let $h^t(a)$ be the history in period t that succeeds h^{t-1} when $a \in A$ is the action profile realized in period $t-1$. Let $v_i(\sigma_i^t, \sigma_{-i}^t | h^{t-1})$ and $v_i(\sigma_i^t, \sigma_{-i}^t | h^t(a))$ be the continuation utilities of player i at each of the histories $h^{t-1}, h^t(a), a \in A$, when $(\sigma_i^t, \sigma_{-i}^t)$ is played. By construction, the continuation strategy prescribed by σ_{-i}^t after each history $h^t(a)$ is constant. Thus, $v_i(\sigma_i^t, \sigma_{-i}^t | h^t(a)) \geq \underline{v}_i^c$ for every $a \in A$.

Then,

$$\begin{aligned}
v_i(\sigma_i^t, \sigma_{-i}^t | h^{t-1}) &= \max_{\alpha_i' \in \Delta(A_i)} \mathbb{E}_{(\alpha_i', \alpha_{-i})}[g_i(a) + \beta_i(a) v_i(\sigma_i^t, \sigma_{-i}^t | h^t(a))] \\
&\geq \max_{\alpha_i' \in \Delta(A_i)} \mathbb{E}_{(\alpha_i', \alpha_{-i})}[g_i(a) + \beta_i(a) \underline{v}_i^c] \\
&\geq \min_{\alpha_{-i}' \in \times_{k \neq i} \Delta(A_k)} \max_{\alpha_i' \in \Delta(A_i)} \mathbb{E}_{(\alpha_i', \alpha_{-i}')}[g_i(a) + \beta_i(a) \underline{v}_i^c] \\
&= \underline{v}_i^c.
\end{aligned}$$

Iterating the argument we see that $v_i(\sigma_i^t, \sigma_{-i}^t) \geq \underline{v}_i^c$, as desired. □

B Proof of Theorem 5.1

Fix Γ. Given $\alpha \in \Delta(A), i \in I$, and $\lambda \in [0,1)$, from now on we write $v_i(\alpha)$ in place of $v_{i,\lambda}(\alpha^{con})$. Similarly, we write $v_i(a)$ in place of $v_{i,\lambda}(a, a, ...)$. We can do so since from (7) we know that the utility of constant strategies (play paths) does not depend on λ. The following basic observation is also used throughout the appendix. For every $\mathbf{a} \in A^\infty$, $\lambda \in [0,1), i \in I, t \in \mathcal{T}$,

$$v_{i,\lambda}({}_t\mathbf{a}) = (1-\lambda)\big(1 - \beta_{i,0}(a^t)\big) v_i(a^t) + \Big(1 - (1-\lambda)\big(1 - \beta_{i,0}(a^t)\big)\Big) v_{i,\lambda}({}_{t+1}\mathbf{a}). \tag{10}$$

In particular, $v_{i,\lambda}({}_t\mathbf{a})$ is a convex combination of $v_i(a^t)$ and $v_{i,\lambda}({}_{t+1}\mathbf{a})$. Observe, however, that the 'weights' depend on i. Thus, for a payoff profile $v_\lambda(\mathbf{a}) := (v_{1,\lambda}(\mathbf{a}), ..., v_{n,\lambda}(\mathbf{a})) \in \mathbb{R}^n$, it is generally not the case that $v_\lambda({}_t\mathbf{a})$ is a convex combination of $v(a^t)$ and $v_\lambda({}_{t+1}\mathbf{a})$.

In this section, we make use of the one shot deviation principle. The latter is applicable since the games we study are continuous at infinity. See Theorem 4.2 in Fudenberg and Tirole [16].

For every $i \in I$, let $M^i := (M_1^i, ..., M_n^i) \in \Sigma$ be a strategy profile whereby player i is minmaxed by his opponents and he plays a best response. By Lemma 3.1, we can choose M^i to be a profile of constant strategies and, hence, identify M^i with an element of $\Delta(A)$. Since utilities are normalized so that each player's minmax payoff is zero, we deduce that $g_i(M^i) = 0$ for every $i \in I$.

Lemma B 2. *Assume Richness. There exists $\overline{\rho} > 0$ such that for every $\rho \in (0, \overline{\rho})$ and every $i \in I$, we can find $\kappa^i \in \Delta(A)$ such that $v_i(\kappa^i) = \rho$ and $v_j(\kappa^i) > \rho$ for all $j \neq i$.*

Proof. Fix $i \in I$. Recall that $v_i(a) = \frac{g_i(a)}{1-\beta_{i,0}(a)}$ for every $a \in A$. Richness implies that there exist action profiles $a^i, \tilde{a}^i \in A$ such that $g_i(\tilde{a}^i) > 0 \geq g_i(a^i)$ and $g_j(a^i) > 0, g_j(\tilde{a}^i) \geq 0$ for all $j \neq i$. Fix $\rho \in (0, v_i(\tilde{a}^i))$ and let

$$\kappa^i(a^i) := \frac{\big(1-\beta_{i,0}(\tilde{a}^i)\big)\big(v_i(\tilde{a}^i)-\rho\big)}{\big(1-\beta_{i,0}(\tilde{a}^i)\big)\big(v_i(\tilde{a}^i)-\rho\big)+\big(1-\beta_{i,0}(a^i)\big)\big(\rho-v_i(a^i)\big)}.$$

The numerator and the first term in the denominator coincide. Moreover, the numerator is positive since $\rho < v_i(\tilde{a}^i)$. The second term in the denominator is positive since $\rho > 0$ and $-g_i(a^i) \geq 0$. Conclude that $\kappa^i(a^i) \in (0,1)$. Let $\kappa^i \in \Delta(A)$ be the mixed action profile such that a^i and $\tilde{a}^i$ are played with probabilities $\kappa^i(a^i)$ and $1-\kappa^i(a^i)$ respectively. By construction, $v_i(\kappa^i) = \rho$.

Next, fix some $j \neq i$. We are going to show that there is $\overline{\rho}_j \in (0, v_i(\tilde{a}^i))$ such that for every $\rho \in (0, \overline{\rho}_j)$, $v_j(\kappa^i) > \rho$. Suppose $g_j(\tilde{a}^i) - g_j(a^i) > 0$ and $\beta_{j,0}(\tilde{a}^i) - \beta_{j,0}(a^i) > 0$. Other cases follow from analogous calculations. For every $\rho \in (0, v_i(\tilde{a}^i))$, the inequality $v_j(\kappa^i) > \rho$ holds if and only if

$$\kappa^i(a^i) < \frac{g_j(\tilde{a}^i) - (1-\beta_{j,0}(\tilde{a}^i))\rho}{g_j(\tilde{a}^i) - g_j(a^i) + \big(\beta_{j,0}(\tilde{a}^i) - \beta_{j,0}(a^i)\big)\rho}.$$

The right hand side is greater than 1 if and only if $\rho < v_j(a^i)$. To obtain the desired $\overline{\rho}_j$, we can thus let $\overline{\rho}_j := \min\{v_i(\tilde{a}^i), v_j(a^i)\}$.

Since the number of players is finite, we can let $\overline{\rho}^i := \min_{j \neq i} \overline{\rho}_j$ and $\overline{\rho} = \min_{i \in I} \overline{\rho}^i$, completing the proof. □

Proof of Theorem 5.1. Take $\varepsilon > 0$ and let $\overline{\rho} > 0$ be as defined in Lemma B 2. Take some ρ such that $0 < \rho < \min\{\overline{\rho}, \varepsilon\}$. Fix $i \in I$ and let $\kappa^i \in \Delta(A)$ be the mixed action profile such that $v_i(\kappa^i) = \rho$ and $v_j(\kappa^i) > \rho$ for all $j \neq i$. Let $\overline{g}_i := \max_a g_i(a)$ and $\overline{\beta}_{i,\lambda} := \lambda + (1-\lambda)\max_a \beta_{i,0}(a)$. Choose an integer μ_i such that $\mu_i > \frac{\overline{g}_i}{v_i(\kappa^i)\big(1-\beta_{i,0}(M^i)\big)}$. Recall that $\beta_{i,\lambda}(M^i) := \lambda + (1-\lambda)\beta_{i,0}(M^i)$ and note that

$$\lim_{\lambda \to 1} \frac{1-[\beta_{i,\lambda}(M^i)]^{\mu_i}}{1-\beta_{i,\lambda}(M^i)} = \lim_{\beta_{i,\lambda}(M^i) \to 1} \frac{1-[\beta_{i,\lambda}(M^i)]^{\mu_i}}{1-\beta_{i,\lambda}(M^i)} = \mu_i.$$

Thus, we can find $\underline{\lambda}'_i \in [0,1)$ such that

$$\frac{\overline{g}_i}{v_i(\kappa^i)\left(1-\beta_{i,0}(M^i)\right)} < \frac{1-[\beta_{i,\lambda}(M^i)]^{\mu_i}}{1-\beta_{i,\lambda}(M^i)} \quad \forall \lambda \in (\underline{\lambda}'_i, 1). \tag{11}$$

Fix $j \neq i$ and an integer μ between 1 and μ_j. Consider the inequality

$$(1-\lambda)\overline{g}_i + \left(v_i(\kappa^i) - [\beta_{i,\lambda}(M^j)]^{\mu} v_i(\kappa^j)\right) - \frac{g_i(M^j)}{1-\beta_{i,0}(M^j)}(1-[\beta_{i,\lambda}(M^j)]^{\mu}) < 0. \tag{12}$$

Note that $\overline{g}_i$ and $\frac{g_i(M^j)}{1-\beta_{i,0}(M^j)}$ are constants that do not depend on λ. Also, the first term and the last term above converge to 0 as $\lambda \to 1$. The second term converges to a negative number since $v_i(\kappa^j) > v_i(\kappa^i)$. Thus, there exists $\underline{\lambda}''_i \in [0,1)$ such that the inequality in (12) is satisfied for all $\lambda \in (\underline{\lambda}''_i, 1)$. Moreover, since there are finitely players $j \neq i$ and finitely many integers μ between between 1 and μ_j, the threshold $\underline{\lambda}''_i$ can be chosen independently of j and μ.

Let $\underline{\lambda}_i := \max\{\underline{\lambda}'_i, \underline{\lambda}''_i\}$ and $\underline{\lambda} := \max_{i \in I} \underline{\lambda}_i$. Take any $\lambda \in (\underline{\lambda}, 1)$ and $\mathbf{a} \in SIR^{\varepsilon}(\lambda)$. By definition, we have $v_{i,\lambda}({}_t\mathbf{a}) \geq \varepsilon$, for all $i \in I$ and $t \in \mathcal{T}$. Consider the following strategy $\sigma_i \in \Sigma_i$ for player i:
(A) play a_i^t in period t as long as a^{t-1} was played last period. If player j deviates from (A), then
(B) play M_i^j for μ_j periods and then
(C) play κ_i^j thereafter.
If player k deviates in phase (B) or (C), then begin phase (B) again with $j = k$.

If player i deviates in phase (A) and then conforms, he receives at most $\overline{g}_i$ the period he deviates, 0 for μ_i periods, and a continuation payoff $v_i(\kappa^i)$. Thus, his continuation payoff after deviating is no greater than $(1-\lambda)\overline{g}_i + \overline{\beta}_{i,\lambda}[\beta_{i,\lambda}(M^i)]^{\mu_i} v_i(\kappa^i)$ and the gain from deviating is less than

$$(1-\lambda)\overline{g}_i + \overline{\beta}_{i,\lambda}[\beta_{i,\lambda}(M^i)]^{\mu_i} v_i(\kappa^i) - v_{i,\lambda}({}_t\mathbf{a}).$$

The latter is less than

$$(1-\lambda)\overline{g}_i + \left([\beta_{i,\lambda}(M^i)]^{\mu_i} - 1\right) v_i(\kappa^i) \tag{13}$$

because $\overline{\beta}_{i,\lambda} < 1$ and $v_i(\kappa^i) < \varepsilon \leq v_{i,\lambda}({}_t\mathbf{a})$. Direct verification shows that the term in (13) is less than 0 if and only if (11) is satisfied. Since $\underline{\lambda}'_i < \lambda < 1$, the gain from deviating is less than 0.

If player i deviates in phase (B) in which he is being punished, he obtains at most 0 in the period in which he deviates, and then only lengthens his punishment, postponing the positive continuation payoff $v_i(\kappa^i)$. If player i deviates in phase (B) in which player j is being punished and then conforms, he receives at most $(1-\lambda)\overline{g}_i + \overline{\beta}_{i,\lambda}[\beta_{i,\lambda}(M^i)]^{\mu_i} v_i(\kappa^i)$. If he does not deviate, he receives at least $\frac{g_i(M^j)\left(1-[\beta_{i,\lambda}(M^j)]^{\mu}\right)}{1-\beta_{i,0}(M^j)} + [\beta_{i,\lambda}(M^j)]^{\mu} v_i(\kappa^j)$, where μ is some integer between 1 and μ_j. Conclude that the gain from deviating is at most

$$(1-\lambda)\overline{g}_i + v_i(\kappa^i) - [\beta_{i,\lambda}(M^j)]^{\mu} v_i(\kappa^j) - \frac{g_i(M^j)\left(1-[\beta_{i,\lambda}(M^j)]^{\mu}\right)}{1-\beta_{i,0}(M^j)},$$

which is the expression in (12). Since $\underline{\lambda}_i'' < \lambda < 1$, the potential gain from deviating is less than 0. The argument why the players don't deviate in phase (C) is analogous to that for phase (A). □

C Results from Section 7

Let $V(\lambda) \subseteq \mathbb{R}^2$ denote the set of feasible ex ante payoffs given a level of patience λ and a two player game Γ. Because we allow for public randomization, the set $V(\lambda)$ is convex. It follows that every efficient path maximizes a weighted sum of the players' utilities. Since we use this observation repeatedly, it is helpful to establish the requisite formalism. First, define

$$s_\lambda(\mathbf{a}, \eta) := \eta_1 v_{1,\lambda}(\mathbf{a}) + \eta_2 v_{2,\lambda}(\mathbf{a}),$$

where $\eta := (\eta_1, \eta_2) \in \mathbb{R}^2_+$ is a vector of 'Pareto weights', one for each player. Say that $\eta, \eta' \in \mathbb{R}^2_+$ **determine the same direction** if there exists $\xi \in \mathbb{R}_{++}$ such that $\eta = \xi\eta'$. Given any λ and η, let $P(\lambda, \eta)$ be the set of paths $\mathbf{a} \in A^\infty$ that maximize the function $s_\lambda(\mathbf{a}, \eta)$. Note that $P(\lambda, \eta) \subset P(\lambda)$, that is, the paths in $P(\lambda, \eta)$ are efficient. Furthermore, if η and η' determine the same direction, then $P(\lambda, \eta) = P(\lambda, \eta')$. Similarly, for every $\eta = (\eta_1, \eta_2) \in \mathbb{R}^2_+$ and λ, let $F(\lambda, \eta)$ denote the set of all payoff profiles $(v_1, v_2) \in V(\lambda)$ such that $\eta_1 v_1 + \eta_2 v_2 \geq \eta_1 v_1' + \eta_2 v_2'$ for all $(v_1', v_2') \in V(\lambda)$. The set $F(\lambda, \eta)$ is the linear segment of the Pareto frontier of $V(\lambda)$ that is perpendicular to the direction η. Note that given a level of patience $\lambda \in [0, 1)$, if a path $\mathbf{a} \in A^\infty$ is strictly Pareto dominated by a path $\mathbf{a}' \in A^\infty$, then $s_\lambda(\mathbf{a}', \eta) > s_\lambda(\mathbf{a}, \eta)$ for all $\eta \in \mathbb{R}^2_+$. For every $\eta \in \mathbb{R}^2_+$, $\mathbf{a} \in A^\infty$, and $t \geq 1$, let

$$\eta^t_\lambda(\mathbf{a}) := \left(\eta_1 \prod_{\tau=0}^{t-1} \beta_{1,\lambda}(a^\tau), \eta_2 \prod_{\tau=0}^{t-1} \beta_{2,\lambda}(a^\tau)\right) \in \mathbb{R}^2_+.$$

Given $\mathbf{a} \in P(\lambda, \eta)$, **player i's relative weight in period** t is defined as $\frac{\eta^t_{i,\lambda}(\mathbf{a})}{\eta^t_{j,\lambda}(\mathbf{a})}$, $i \in \{1, 2\}$ and $j \neq i$. When the play path $\mathbf{a}$ is clear from the context, we may also write η^t_λ in place of $\eta^t_\lambda(\mathbf{a})$.

The following lemmas, whose familiar proofs we omit, are used repeatedly from now on.

Lemma C 3. *Take any $\lambda \in [0, 1)$, $\eta \in \mathbb{R}^2_+$ and $\mathbf{a} \in P(\lambda, \eta)$. Then, ${}_t\mathbf{a} \in P(\lambda, \eta^t_\lambda(\mathbf{a}))$ for all $t > 0$. Also, if $\hat{\mathbf{a}} \in P(\lambda, \eta^t_\lambda(\mathbf{a}))$ for some $t > 0$, then $(a^0, ..., a^{t-1}, \hat{\mathbf{a}}) \in P(\lambda, \eta)$.*

Lemma C 4. *Take any $\lambda \in [0, 1)$, $\eta, \eta' \in \mathbb{R}^2_+$, $\mathbf{a} \in P(\lambda, \eta)$, and $\mathbf{a}' \in P(\lambda, \eta')$. For every $i \in \{1, 2\}$ and $j \neq i$, if $\frac{\eta'_i}{\eta'_j} > \frac{\eta_i}{\eta_j}$, then $v_{i,\lambda}(\mathbf{a}') \geq v_{i,\lambda}(\mathbf{a})$ and $v_{j,\lambda}(\mathbf{a}') \leq v_{j,\lambda}(\mathbf{a})$.*

In the next two lemmas, we assume either decreasing or increasing marginal impatience, henceforth DMI and IMI. Let A^E be the set of action profiles $a \in A$ such that $v_1(a) = v_2(a)$. When the set A^E is nonempty, we assume that $\arg\max_{a \in A^E} v_1(a)$ consists of a single element $a^* \in A^E$. This assumption simplifies the exposition without affecting the results.

Lemma C 5. *For every $\lambda \in [0, 1)$, $\eta \in \mathbb{R}^2_+$ and $\mathbf{a} \in P(\lambda, \eta)$, if $a^0 \in A^E$, then $(a^0, a^0, ...) \in P(\lambda, \eta)$. Moreover, $v_\lambda({}_1\mathbf{a}), v_\lambda(\mathbf{a}), v_\lambda(a^0) \in F(\lambda, \eta)$.*

Proof. Note that under both IMI and DMI, $a^0 \in A^E$ if and only if $g_1(a^0) = g_2(a^0)$ and $\beta_{1,\lambda}(a^0) = \beta_{2,\lambda}(a^0)$. Since $\beta_{1,\lambda}(a^0) = \beta_{2,\lambda}(a^0)$, the direction $\eta^1_\lambda = \left(\eta_1\beta_{1,\lambda}(a^0), \eta_2\beta_{2,\lambda}(a^0)\right)$ is the same as the direction η. By Lemma C 3, ${}_1\mathbf{a} \in P(\lambda, \eta)$. Then, the fact that $\mathbf{a} = (a^0, {}_1\mathbf{a}) \in P(\lambda, \eta)$ implies that

$$s_\lambda(\mathbf{a}, \eta) = s_\lambda({}_1\mathbf{a}, \eta) \tag{14}$$

Since $v_{i,\lambda}(\mathbf{a}) = (1 - \lambda)g_i(a^0) + \beta_{i,\lambda}(a^0)v_{i,\lambda}({}_1\mathbf{a})$ for $i \in \{1, 2\}$, and $\beta_{1,\lambda}(a^0) = \beta_{2,\lambda}(a^0)$, we can rewrite (14) as

$$\eta_1 \frac{g_1(a^0)}{1 - \beta_{1,0}(a^0)} + \eta_2 \frac{g_2(a^0)}{1 - \beta_{2,0}(a^0)} = s_\lambda({}_1\mathbf{a}, \eta). \tag{15}$$

Observe that the left-hand side of (15) is just $\eta_1 v_1(a_0) + \eta_2 v_2(a_0)$. Moreover, since ${}_1\mathbf{a} \in P(\lambda, \eta)$, we may conclude that $(a^0, a^0, ...) \in P(\lambda, \eta)$. Since the paths ${}_1\mathbf{a}$, $\mathbf{a}$, and $(a^0, a^0, ...)$

are all efficient given η, we may conclude that $v_\lambda({}_1\mathbf{a}), v_\lambda(\mathbf{a}), v(a^0) \in F(\lambda,\eta)$, as desired. □

Lemma C 6. *Suppose A^E is nonempty. For every $\lambda \in [0,1), \eta \in \mathbb{R}^2_+, \mathbf{a} \in P(\lambda,\eta)$, if $a^t \in A^E$ for some t, then $a^t = a^*$.*

Proof. If A^E is singleton, the result holds trivially. Else take a path $\mathbf{a} \in P(\lambda,\eta)$ and suppose that $a^t \in A^E \setminus \{a^*\}$ for some t. By Lemma C 5, $(a^t, a^t, ...) \in P(\lambda, \eta^t_\lambda)$. However, since $v_i(a^t) < v_i(a^*)$, $i \in \{1,2\}$, we have $\eta^t_{1,\lambda} v_1(a^t) + \eta^t_{2,\lambda} v_2(a^t) < \eta^t_{1,\lambda} v_1(a^*) + \eta^t_{2,\lambda} v_2(a^*)$. Thus, the path $(a^t, a^t, ...)$ is strictly Pareto dominated by the path $(a^*, a^* ...)$, contradicting the efficiency of $(a^t, a^t, ...)$. □

C .1 Proof of Theorem 7.1

Assume IMI and focus on the prisoners' dilemma game. Recall that $\mathbf{a}^{max,i}$ denotes player i's most preferred play path. In particular, $a^{max,1} = ((D,C),(D,C),...)$ and $a^{max,2} = ((C,D),(C,D),...)$.

Lemma C 7. *There are no $\lambda \in [0,1)$ and $\eta \in \mathbb{R}^2_{++}$ such that $\mathbf{a}^{max,1} \in P(\lambda,\eta)$ or $\mathbf{a}^{max,2} \in P(\lambda,\eta)$.*

Proof. By way of contradiction, suppose $\lambda \in [0,1)$ and $\eta \in \mathbb{R}^2_{++}$ are such that $\mathbf{a}^{max,1} \in P(\lambda,\eta)$. Since $\beta_\lambda(d) < \beta_\lambda(b)$, there is some $T \in \mathcal{T}$ large enough such that player 1's relative weight $\frac{\eta_1}{\eta_2}\frac{[\beta_\lambda(d)]^T}{[\beta_\lambda(b)]^T}$ is almost zero, while player 2's relative weight $\frac{\eta_2}{\eta_1}\frac{[\beta_\lambda(b)]^T}{[\beta_\lambda(d)]^T}$ is almost infinity. But then $\mathbf{a}^{max,1} \notin P(\lambda, \eta^T_\lambda)$, contradicting Lemma C 3. Similar arguments apply for the path $\mathbf{a}^{max,2}$. □

Lemma C 8. *For every $\lambda \in [0,1)$, $\eta \in \mathbb{R}^2_+$, and $\mathbf{a} \in P(\lambda,\eta)$, if $a^0 = (C,D)$ and $a^1 = (D,C)$, then $\mathbf{a}^{A,2} \in P(\lambda,\eta)$. Similarly, if $a^0 = (D,C)$ and $a^1 = (C,D)$, then $\mathbf{a}^{A,1} \in P(\lambda,\eta)$.*

Proof. It is enough to prove the first assertion. Fix $\lambda \in [0,1)$ and $\eta \in \mathbb{R}^2_+$. Take a path $\mathbf{a} \in P(\lambda,\eta)$ such that $a^0 = (C,D)$ and $a^1 = (D,C)$. Note that

$$\eta^2_\lambda = (\eta_1 \beta_\lambda(b)\beta_\lambda(d), \eta_2 \beta_\lambda(d)\beta_\lambda(b)).$$

Thus, the direction η^2_λ is the same as η. By Lemma C 3, ${}_2\mathbf{a} \in P(\lambda,\eta)$ and, in particular,

$$s_\lambda(\mathbf{a},\eta) = s_\lambda({}_2\mathbf{a},\eta). \tag{16}$$

But

$$v_{1,\lambda}(\mathbf{a}) = (1-\lambda)\big(b + \beta_\lambda(b)d\big) + \beta_\lambda(b)\beta_\lambda(d)v_{1,\lambda}({}_2\mathbf{a})$$
$$v_{2,\lambda}(\mathbf{a}) = (1-\lambda)\big(d + \beta_\lambda(d)b\big) + \beta_\lambda(b)\beta_\lambda(d)v_{2,\lambda}({}_2\mathbf{a}).$$

Plugging these expressions into (16) and rearranging terms gives

$$\eta_1(1-\lambda)\frac{b+\beta_\lambda(b)d}{1-\beta_\lambda(b)\beta_\lambda(d)} + \eta_2(1-\lambda)\frac{d+\beta_\lambda(d)b}{1-\beta_\lambda(b)\beta_\lambda(d)} = s_\lambda({}_2\mathbf{a},\eta). \tag{17}$$

But note that $v_1(\mathbf{a}^{A,2}) = (1-\lambda)\frac{b+\beta_\lambda(b)d}{1-\beta_\lambda(b)\beta_\lambda(d)}$ and $v_2(\mathbf{a}^{A,2}) = (1-\lambda)\frac{d+\beta_\lambda(d)b}{1-\beta_\lambda(b)\beta_\lambda(d)}$. Conclude that $s_\lambda(\mathbf{a}^{A,2},\eta) = s_\lambda({}_2\mathbf{a},\eta)$ and, hence, that $\mathbf{a}^{A,2} \in P(\lambda,\eta)$. □

Lemma C 9. *For every $\lambda \in [0,1)$, $\eta \in \mathbb{R}^2_+$, and $\mathbf{a} \in P(\lambda,\eta)$, if $\frac{\eta_1}{\eta_2} < 1$, then $v_{1,\lambda}(\mathbf{a}) \le v_{2,\lambda}(\mathbf{a})$ and $a^0 \ne (D,C)$; if $\frac{\eta_1}{\eta_2} > 1$, then $v_{1,\lambda}(\mathbf{a}) \ge v_{2,\lambda}(\mathbf{a})$ and $a^0 \ne (C,D)$.*

Proof. Fix $\lambda \in [0,1)$ and $\eta \in \mathbb{R}^2_+$. It is enough to consider the case when $\frac{\eta_1}{\eta_2} < 1$. The inequality $v_{1,\lambda}(\mathbf{a}) \le v_{2,\lambda}(\mathbf{a})$ follows directly from the symmetry of the game. To prove the second assertion, suppose by way of contradiction that $a^0 = (D,C)$. Let $T \ge 1$ be the first period t such that $a^t \ne (D,C)$. Such T exists because $v_{1,\lambda}(\mathbf{a}) \le v_{2,\lambda}(\mathbf{a})$. Suppose $a^T = (C,C)$. Consider the path $\hat{\mathbf{a}}$ such that $\hat{a}^t = (D,C)$ for all $0 \le t < T$ and $\hat{a}^t = (C,C)$ for all $t \ge T$. From Lemma C 5, $\hat{\mathbf{a}} \in P(\lambda,\eta)$. But, by construction, $v_{1,\lambda}(\hat{\mathbf{a}}) > v_{2,\lambda}(\hat{\mathbf{a}})$, contradicting the first assertion in the lemma. Thus, a^T can only be (C,D). But then, by Lemma C 8, $\mathbf{a}^{A,1} \in P(\lambda,\eta^{T-1}_\lambda(\mathbf{a}))$. Also,

$$\frac{\eta^{T-1}_{1,\lambda}(\mathbf{a})}{\eta^{T-1}_{2,\lambda}(\mathbf{a})} = \frac{[\beta_\lambda(d)]^{T-1}}{[\beta_\lambda(b)]^{T-1}}\frac{\eta_1}{\eta_2} \le \frac{\eta_1}{\eta_2} < 1,$$

where the first inequality follows from the fact that $\beta_\lambda(d) < \beta_\lambda(b)$. But then $v_{1,\lambda}(\mathbf{a}^{A,1}) > v_{2,\lambda}(\mathbf{a}^{A,1})$, contradicting the first assertion in the lemma. Thus, $a^0 \ne (D,C)$. □

Some additional notation is now needed. Let $\mathbf{a}^C(0) := \mathbf{a}^C$. For every $T \ge 1$, let $\mathbf{a}^C(T)$ be the path such that $a^t = (C,D)$ for all $0 \le t < T$ and ${}_T\mathbf{a} = \mathbf{a}^C$. Recall that $\eta_{sym} := (1,1)$ and $P_{sym}(\lambda) := P(\lambda,\eta_{sym})$. Let $s_\lambda(\mathbf{a})$ be the sum of payoffs, that is, $s_\lambda(\mathbf{a}) := s_\lambda(\mathbf{a},\eta_{sym})$ for all $\mathbf{a} \in A^\infty$ and $\lambda \in [0,1)$. The next lemma shows that there is no λ such that $\mathbf{a}^C(1) \in P_{sym}(\lambda)$.

Lemma C 10. $s_\lambda\big(\mathbf{a}^C(1)\big) < \max\{s_\lambda(\mathbf{a}^C), s_\lambda(\mathbf{a}^{A,2})\}$ *for all $\lambda \in [0,1)$.*

Proof. Fix $\lambda \in [0,1)$. First, compute the sums of payoffs for the paths $\mathbf{a}^C$, $\mathbf{a}^{A,2}$, and $\mathbf{a}^C(1)$:

$$
\begin{aligned}
s_\lambda(\mathbf{a}^C) &= \frac{2c}{1-\beta_0(c)}, \\
s_\lambda(\mathbf{a}^{A,2}) &= (1-\lambda)\frac{b+\beta_\lambda(b)d+d+\beta_\lambda(d)b}{1-\beta_\lambda(b)\beta_\lambda(d)}, \\
s_\lambda\big(\mathbf{a}^C(1)\big) &= (1-\lambda)b+\beta_\lambda(b)\frac{c}{1-\beta_0(c)}+(1-\lambda)d+\beta_\lambda(d)\frac{c}{1-\beta_0(c)}.
\end{aligned}
$$

Note that $s_\lambda(\mathbf{a}^C) > s_\lambda\big(\mathbf{a}^C(1)\big)$ if and only if

$$
\frac{c}{1-\beta_0(c)} > \frac{b+d}{1-\beta_0(b)+1-\beta_0(d)}. \tag{18}
$$

If the inequality in (18) holds, the proof is complete. Suppose now that (18) holds with equality. Then,

$$
s_\lambda\big(\mathbf{a}^C(1)\big) = \frac{2(b+d)}{1-\beta_0(b)+1-\beta_0(d)}. \tag{19}
$$

Also, since $s_{\lambda'}(\mathbf{a}^{A,2})$ is increasing in λ', we have

$$
s_\lambda(\mathbf{a}^{A,2}) > \lim_{\lambda'\to 1} s_{\lambda'}(\mathbf{a}^{A,2}) = \frac{2(b+d)}{1-\beta_0(b)+1-\beta_0(d)}. \tag{20}
$$

Combining (19) and (20) yields: $s_\lambda(\mathbf{a}^{A,2}) > s_\lambda\big(\mathbf{a}^C(1)\big)$. Finally, consider the case when the strict inequality in (18) is reversed. Then,

$$
\begin{aligned}
s_\lambda\big(\mathbf{a}^C(1)\big) &< (1-\lambda)(b+d)+\big(\beta_\lambda(b)+\beta_\lambda(d)\big)\frac{b+d}{1-\beta_0(b)+1-\beta_0(d)} \\
&= \frac{2(b+d)}{1-\beta_0(b)+1-\beta_0(d)} < s_\lambda(\mathbf{a}^{A,2}).
\end{aligned}
$$

The equality follows from direct simplification. The last inequality follows from the inequality in (20). $\square$

Consider the inequalities

$$
1 \le f(\lambda) := \frac{\frac{c}{1-\beta_0(c)} - (1-\lambda)\frac{b+\beta_\lambda(b)d}{1-\beta_\lambda(b)\beta_\lambda(d)}}{(1-\lambda)\frac{d+\beta_\lambda(d)b}{1-\beta_\lambda(b)\beta_\lambda(d)} - \frac{c}{1-\beta_0(c)}} \le \sqrt{\frac{\beta_\lambda(b)}{\beta_\lambda(d)}}. \tag{21}
$$

Lemma C 11. *The inequalities in (21) hold if and only if there is $\eta \in \mathbb{R}^2_+$ such that $\mathbf{a}^{A,2}, \mathbf{a}^C \in$*

$P(\lambda, \eta)$, that is, if and only if λ is irregular.

Proof. Take an irregular $\lambda \in [0, 1)$ and assume that $\mathbf{a}^{A,2}, \mathbf{a}^C \in P(\lambda, \eta)$ for some $\eta \in \mathbb{R}^2_+$. If $f(\lambda) < 0$, then either $\mathbf{a}^C$ or $\mathbf{a}^{A,2}$ is strictly Pareto dominated. If $f(\lambda) \in [0, 1)$, then $v_\lambda(\mathbf{a}^C)$ is strictly inside the Pareto frontier $F(\lambda, \eta_{sym})$. Hence, $\mathbf{a}^C$ is strictly Pareto dominated by some path in $\mathcal{A}$. Conclude that $f(\lambda) \geq 1$. Turn to the second inequality. Since $\mathbf{a}^{A,2}, \mathbf{a}^C \in P(\lambda, \eta)$, Lemma C 3 implies that $\big((C, D), (D, C), \mathbf{a}^C\big) \in P(\lambda, \eta)$. By Lemma C 3, the paths $\big((D, C), \mathbf{a}^C\big)$ and $\big((D, C), \mathbf{a}^{A,2}\big) = \mathbf{a}^{A,1}$ are efficient given the direction $\big(\eta_1 \beta_\lambda(b), \eta_2 \beta_\lambda(d)\big)$. By the symmetry of the game, the path $\big((C, D), \mathbf{a}^C\big)$ and the path $\big((C, D), (D, C), \mathbf{a}^{A,2}\big) = \mathbf{a}^{A,2}$ are efficient given the direction

$$\eta' := \big(\eta_2 \beta_\lambda(d), \eta_1 \beta_\lambda(b)\big). \tag{22}$$

Thus, the path $\mathbf{a}^{A,2}$ is efficient under both η and η'. By the convexity of the feasible set, we have $\frac{\eta'_1}{\eta'_2} \leq \frac{\eta_1}{\eta_2}$. Observe now that we can compute η explicitly. In particular, if $\mathbf{a}^{A,2}, \mathbf{a}^C \in P(\lambda, \eta)$, then

$$\eta = (v_{2,\lambda}(\mathbf{a}^{A,2}) - v_{2,\lambda}(\mathbf{a}^C), v_{1,\lambda}(\mathbf{a}^C) - v_{1,\lambda}(\mathbf{a}^{A,2})).$$

Taking this and (22) into account, the inequality $\frac{\eta'_1}{\eta'_2} \leq \frac{\eta_1}{\eta_2}$ is seen to be equivalent to the second inequality in (21). Identical arguments show that $\mathbf{a}^{A,2}, \mathbf{a}^C \in P(\lambda, \eta)$ whenever (21) holds. □

Given a path $\mathbf{a} \in A^\infty$ and some $T \in \mathcal{T}$, say that $(T, T+1)$ is an **alternation for a** if $a^T, a^{T+1} \in \{(C, D), (D, C)\}$ and $a^T \neq a^{T+1}$. In the rest of this section, we fix a regular λ and characterize the Pareto set $P(\lambda)$. A characterization for irregular λ is provided in Section C .4. One more preliminary observation: If $\eta_i = 0$ and $\eta_j > 0$ for some $i \in \{1, 2\}$ and $j \neq i$, then $P(\lambda, \eta) = \{\mathbf{a}^{max,j}\}$. Taking Lemma C 7 into account, we see that $P_{++}(\lambda) = \cup_{\eta \in \mathbb{R}^2_{++}} P(\lambda, \eta)$. The proof of Theorem 7.1 will thus be complete once we characterize the set $P_{++}(\lambda)$.

Lemma C 12. *For every $\eta \in \mathbb{R}^2_{++}$ and $\mathbf{a} \in P(\lambda, \eta)$, if $(T, T+1)$ is an alternation for $\mathbf{a}$, then $a^t \neq (C, C)$ for every $t \in \mathcal{T}$.*

Proof. It is w.l.o.g. to assume that $a^T = (C, D)$ and $a^{T+1} = (D, C)$. By Lemma C 8, $\mathbf{a}^{A,2} \in P(\lambda, \eta^T_\lambda(\mathbf{a}))$.

First we are going to show that $a^t \neq (C, C)$ for every $t < T$. If $T = 0$, there is nothing to prove. So assume that $T \geq 1$. By way of contradiction, let T' be the greatest integer

$k < T$ such that $a^k = (C, C)$. By Lemma C 5, we know that $\mathbf{a}^C \in P(\lambda, \eta_\lambda^{T'}(\mathbf{a}))$. The latter is possible only if

$$v_{1,\lambda}(\mathbf{a}^C) = v_{2,\lambda}(\mathbf{a}^C) > v_{1,\lambda}(\mathbf{a}^{A,2}). \tag{23}$$

Otherwise, we would have

$$v_{1,\lambda}(\mathbf{a}^C) = v_{2,\lambda}(\mathbf{a}^C) \leq v_{1,\lambda}(\mathbf{a}^{A,2}) < v_{2,\lambda}(\mathbf{a}^{A,2}),$$

and, hence, $\mathbf{a}^C$ would be strictly Pareto dominated by $\mathbf{a}^{A,2}$, a contradiction. Observe now that, by construction, $T' \leq T - 1$. There are two cases to consider. Suppose first that $T' = T - 1$. Recall that $a^{T'} = (C, C)$. Hence, the direction $\eta_\lambda^T(\mathbf{a})$ satisfies

$$(\eta_{1,\lambda}^T(\mathbf{a}), \eta_{2,\lambda}^T(\mathbf{a})) = (\eta_{1,\lambda}^{T'}(\mathbf{a})\beta_\lambda(c), \eta_{2,\lambda}^{T'}(\mathbf{a})\beta_\lambda(c)).$$

Thus, $\eta_\lambda^{T'}(\mathbf{a})$ and $\eta_\lambda^T(\mathbf{a})$ determine the same direction. Hence, $P(\lambda, \eta_\lambda^T(\mathbf{a})) = P(\lambda, \eta_\lambda^{T'}(\mathbf{a}))$. But then $\mathbf{a}^C, \mathbf{a}^{A,2} \in P(\lambda, \eta_\lambda^T(\mathbf{a}))$, contradicting the assumption that λ is regular. Suppose now that $T' < T - 1$. It is w.l.o.g. to assume that $a^t = (C, D)$ for all $T' < t < T$. Else, there would be an alternation $(k, k+1)$ where $T' < k < T$ and we can use the latter alternation in place of $(T, T+1)$. The direction $\eta_\lambda^{T'+2}(\mathbf{a})$ satisfies

$$(\eta_{1,\lambda}^{T'+2}(\mathbf{a}), \eta_{2,\lambda}^{T'+2}(\mathbf{a})) = ((\eta_{1,\lambda}^{T'}(\mathbf{a})\beta_\lambda(c)\beta_\lambda(b), \eta_{2,\lambda}^{T'}(\mathbf{a})\beta_\lambda(c)\beta_\lambda(d)).$$

Since $\beta_\lambda(b) > \beta_\lambda(d)$, we have $\frac{\eta_{1,\lambda}^{T'+2}(\mathbf{a})}{\eta_{2,\lambda}^{T'+2}(\mathbf{a})} > \frac{\eta_{1,\lambda}^{T'}(\mathbf{a})}{\eta_{2,\lambda}^{T'}(\mathbf{a})}$. But then, $\mathbf{a}^C \in P(\lambda, \eta_\lambda^{T'}(\mathbf{a}))$ implies that

$$v_{1,\lambda}(\mathbf{a}') \geq v_{1,\lambda}(\mathbf{a}^C) \quad \forall \mathbf{a}' \in P(\lambda, \eta_\lambda^{T'+2}(\mathbf{a})). \tag{24}$$

Let $\hat{\mathbf{a}} \in A^\infty$ be a path such that $\hat{a}^t = a^t = (C, D)$ for all $T' + 2 \leq t < T$ and ${}_T\hat{\mathbf{a}} = \mathbf{a}^{A,2}$. By Lemma C 3, the fact that ${}_{T'+2}\mathbf{a} \in P(\lambda, \eta_\lambda^{T'+2}(\mathbf{a}))$ and $\mathbf{a}^{A,2} \in P(\lambda, \eta_\lambda^T(\mathbf{a}))$ implies that ${}_{T'+2}\hat{\mathbf{a}} \in P(\lambda, \eta_\lambda^{T'+2}(\mathbf{a}))$. Note now that, by construction, the following inequalities are true:

$$v_{1,\lambda}({}_{T'+2}\hat{\mathbf{a}}) \leq v_{1,\lambda}(\mathbf{a}^{A,2}) < v_{1,\lambda}(\mathbf{a}^C).$$

The first inequality follows from the fact that ${}_{T'+2}\hat{\mathbf{a}}$ begins with a repetitive play of the action profile (C, D) which hurts player 1 and which is succeeded by the more desirable path $\mathbf{a}^{A,2}$. The second inequality follows from (23). Together, the inequalities contradict (24).

Next, we are going to show that $a^t \neq (C,C)$ for every $t > T+1$. Note that

$$(\eta_{1,\lambda}^{T+2}(\mathbf{a}), \eta_{2,\lambda}^{T+2}(\mathbf{a})) = (\eta_{1,\lambda}^{T}(\mathbf{a})\beta_\lambda(b)\beta_\lambda(d), \eta_{2,\lambda}^{T}(\mathbf{a})\beta_\lambda(d)\beta_\lambda(b)).$$

Thus, $\eta_\lambda^T(\mathbf{a})$ and $\eta_\lambda^{T+2}(\mathbf{a})$ determine the same direction, from where we may conclude that $P(\lambda, \eta_\lambda^T(\mathbf{a})) = P(\lambda, \eta_\lambda^{T+2}(\mathbf{a}))$. By way of contradiction, suppose first that $a^{T+2} = (C,C)$. By Lemma C 5, $\mathbf{a}^C \in P(\lambda, \eta_\lambda^{T+2}(\mathbf{a}))$. But then, $\mathbf{a}^C, \mathbf{a}^{A,2} \in P(\lambda, \eta_\lambda^T(\mathbf{a}))$, contradicting Lemma C 11. Suppose now that $a^k = (C,C)$ for some $k > T+2$. Let T' be the smallest such k. It is w.l.o.g. to assume that $a^t = (D,C)$ for all $T+1 < t < T'$. Else, there would be an alternation $(k, k+1)$ where $T < k < T'$ and we can use the latter alternation in place of $(T, T+1)$. Since $(T, T+1)$ is an alternation, $\eta_\lambda^T(\mathbf{a})$ and $\eta_\lambda^{T+2}(\mathbf{a})$ determine the same direction, from where it follows that $P(\lambda, \eta_\lambda^T(\mathbf{a})) = P(\lambda, \eta_\lambda^{T+2}(\mathbf{a}))$. If $T' = T+2$, then $a^{T+2} = (C,C)$ and Lemma C 5 would imply that $\mathbf{a}^C \in P(\lambda, \eta_\lambda^{T+2}(\mathbf{a}))$. But then $\mathbf{a}^C, \mathbf{a}^{A,2} \in P(\lambda, \eta_\lambda^T(\mathbf{a}))$, contradicting Lemma C 11. Hence, $T' > T+2$. Now, since $a^T = (C,D)$, Lemma C 9 shows that $\eta_{1,\lambda}^T(\mathbf{a}) \leq \eta_{2,\lambda}^T(\mathbf{a})$. And, since $a^{T+2} = (D,C)$, Lemma C 9 shows that $\eta_{1,\lambda}^{T+2}(\mathbf{a}) \geq \eta_{2,\lambda}^{T+2}(\mathbf{a})$. Conclude that both $\eta_\lambda^T(\mathbf{a})$ and $\eta_\lambda^{T+2}(\mathbf{a})$ determine the same direction as η_{sym}. To complete the proof, suppose first that $T' = T+3$. Hence, $a^{T+3} = (C,C)$ and, by Lemma C 5, we know that $\mathbf{a}^C \in P(\lambda, \eta_\lambda^{T+3}(\mathbf{a}))$. Then, by Lemma C 3, $((D,C), \mathbf{a}^C) \in P(\lambda, \eta_\lambda^{T+2}(\mathbf{a}))$. But recall that $\eta_\lambda^{T+2}(\mathbf{a})$ and η_{sym} determine the same direction. Thus, $((D,C), \mathbf{a}^C) \in P_{sym}(\lambda)$, contradicting Lemma C 10. Alternatively, suppose that $T' > T+3$. Then, $a^{T+2} = a^{T+3} = (D,C)$ and, hence,

$$\eta_\lambda^{T+3}(\mathbf{a}) = (\eta_{1,\lambda}^{T+2}(\mathbf{a})\beta_\lambda(d), \eta_{2,\lambda}^{T+2}(\mathbf{a})\beta_\lambda(b)).$$

Since $\eta_{1,\lambda}^{T+2}(\mathbf{a}) = \eta_{2,\lambda}^{T+2}(\mathbf{a})$, we may conclude that $\eta_{1,\lambda}^{T+3}(\mathbf{a}) < \eta_{2,\lambda}^{T+3}(\mathbf{a})$. But then, Lemma C 9 shows that a^{T+3} cannot be (D,C), a contradiction. □

Lemma C 13. *For every $\eta \in \mathbb{R}^2_{++}$ such that $\eta_1 < \eta_2$ and every path $\mathbf{a} \in P(\lambda, \eta)$, if $a^0 = (C,C)$, then $\mathbf{a} \in \mathcal{C}_2$.*

Proof. Let $\mathbf{a}$ be as in the statement of the lemma. If $\mathbf{a} = \mathbf{a}^C$, we are done. Suppose that $\mathbf{a} \neq \mathbf{a}^C$. We are going to show that ${}_1\mathbf{a} \in \mathcal{C}_2$ and, hence, $\mathbf{a} \in \mathcal{C}_2$. Let T be the first period t such that $a^t \neq (C,C)$. Since $a^0 = (C,C)$, we know that $T > 0$. By the choice of T, we know that the direction $\eta_\lambda^t(\mathbf{a})$ is the same as η for every $0 < t \leq T$. Since $\frac{\eta_1}{\eta_2} < 1$, Lemma C 9 shows that $a^T \neq (D,C)$. Thus, $a^T = (C,D)$. Next, we are going to show that $a^{T+1} = (C,C)$. By Lemma C 7, the constant path $((C,D),(C,D),...) \notin P(\lambda, \eta_\lambda^T(\mathbf{a}))$. Hence, there exists $t > T$ such that $a^t \neq (C,D)$. Let T' be the smallest such t. By construction, $a^{T'-1} = (C,D)$. Since $a^0 = (C,C)$, Lemma C 12 implies that $a^{T'} \neq (D,C)$. Else, $(T'-1, T')$ would be an alternation for a path that contains (C, C). Conclude that $a^{T'} = (C, C)$. Next, observe that

$$(\eta_{1,\lambda}^{T+1}(\mathbf{a}), \eta_{2,\lambda}^{T+1}(\mathbf{a})) = (\eta_1[\beta_\lambda(c)]^T \beta_\lambda(b), \eta_2[\beta_\lambda(c)]^T \beta_\lambda(d)).$$

Since $\beta_\lambda(b) > \beta_\lambda(d)$, we have $\frac{\eta_{1,\lambda}^{T+1}(\mathbf{a})}{\eta_{2,\lambda}^{T+1}(\mathbf{a})} > \frac{\eta_1}{\eta_2}$. Since $a^0 = (C, C)$, Lemma C 5 shows that $\mathbf{a}^C \in P(\lambda, \eta)$. Combining the last two observations, we may conclude that

$$v_{1,\lambda}(\mathbf{a}') \geq v_{1,\lambda}(\mathbf{a}^C) \quad \forall \mathbf{a}' \in P(\lambda, \eta_\lambda^{T+1}(\mathbf{a})). \tag{25}$$

Recall that $a^{T'} = (C, C)$. By Lemma C 5, $\mathbf{a}^C \in P(\lambda, \eta_\lambda^{T'}(\mathbf{a}))$. Define the path $\hat{\mathbf{a}} \in A^\infty$ such that $\hat{a}^t = a^t = (C, D)$ for $T + 1 \leq t < T'$ and ${}_{T'}\hat{\mathbf{a}} = \mathbf{a}^C$. Lemma C 3 implies that ${}_{T+1}\hat{\mathbf{a}} \in P(\lambda, \eta_\lambda^{T+1}(\mathbf{a}))$. If $T' > T + 1$, then

$$v_{1,\lambda}({}_{T+1}\hat{\mathbf{a}}) < v_{1,\lambda}(\mathbf{a}^C),$$

contradicting (25). Hence, $T' = T + 1$, that is, $a^{T+1} = (C, C)$. To summarize, we have shown that for every $\mathbf{a} \in P(\lambda, \eta)$ such that $a^t = (C, C)$ for all $t < T$ and $a^T = (C, D)$, we have $a^{T+1} = (C, C)$.

Next, we are going to show that, in fact, ${}_{T+1}\mathbf{a} = \mathbf{a}^C$. If not, we can find $k > T + 1$ such that $a^k \neq (C, C)$. Let T'' be the smallest such k. By the choice of T'', we know that $\eta_\lambda^{T''}(\mathbf{a})$ and $\eta_\lambda^{T+1}(\mathbf{a})$ determine the same direction so that $P(\lambda, \eta_\lambda^{T+1}(\mathbf{a})) = P(\lambda, \eta_\lambda^{T''}(\mathbf{a}))$. By Lemma C 3, ${}_{T''}\mathbf{a} \in P(\lambda, \eta_\lambda^{T''}(\mathbf{a}))$. Thus, ${}_{T''}\mathbf{a} \in P(\lambda, \eta_\lambda^{T+1}(\mathbf{a}))$. But then, by Lemma C 3, $\tilde{\mathbf{a}} := (a^0, a^1, ..., a^T, {}_{T''}\mathbf{a}) \in P(\lambda, \eta)$. By construction, $\tilde{\mathbf{a}}$ is such that $\tilde{a}^t = (C, C)$ for all $t < T$, $\tilde{a}^T = (C, D)$, and $\tilde{a}^{T+1} \neq (C, C)$, contradicting the arguments from the first part of the proof. □

Lemma C 14. *For every $\eta \in \mathbb{R}^2_{++}$ and $\mathbf{a} \in P(\lambda, \eta)$, if $a^0 = (C, D)$, and $a^1 = (C, C)$, then ${}_2\mathbf{a} \notin \mathcal{C}_1$.*

Proof. The proof is by contradiction. Suppose ${}_2\mathbf{a} \in \mathcal{C}_1$. Then there exists $T > 1$ such that $a^T = (D, C)$ and ${}_{T+1}\mathbf{a} = \mathbf{a}^C$. By the choice of T, we know that $\eta_\lambda^T(\mathbf{a})$ and $\eta_\lambda^1(\mathbf{a})$ determine the same direction, so that $P(\lambda, \eta_\lambda^T(\mathbf{a})) = P(\lambda, \eta_\lambda^1(\mathbf{a}))$. But, by Lemma C 3, ${}_T\mathbf{a} \in P(\lambda, \eta_\lambda^T(\mathbf{a}))$. Thus, ${}_T\mathbf{a} \in P(\lambda, \eta_\lambda^1(\mathbf{a}))$. But then, by Lemma C 3, $\hat{\mathbf{a}} := \big((C, D), {}_T\mathbf{a}\big) = \big((C, D), (D, C), \mathbf{a}^C\big) \in P(\lambda, \eta)$. Thus, $\hat{\mathbf{a}}$ contains an alternation followed by a play of (C, C) in contradiction of Lemma C 12. □

Lemma C 15. *For every $\mathbf{a} \in P_{sym}(\lambda)$, if $a^0 = (C, D)$, then $\mathbf{a} \in \mathcal{A}$.*

Proof. Let $\eta := \eta_{sym}$. Since $a^0 = (C, D)$ and $\beta_\lambda(b) > \beta_\lambda(d)$,

$$\frac{\eta^1_{1,\lambda}(\mathbf{a})}{\eta^1_{2,\lambda}(\mathbf{a})} = \frac{\eta_1 \beta_\lambda(b)}{\eta_2 \beta_\lambda(d)} = \frac{\beta_\lambda(b)}{\beta_\lambda(d)} > 1.$$

Since ${}_1\mathbf{a} \in P(\lambda, \eta^1_\lambda(\mathbf{a}))$, we can apply Lemma C 9 to deduce that $a^1 \in \{(C, C), (D, C)\}$. If $a^1 = (C, C)$, it follows from Lemma C 5 that $\mathbf{a}^C \in P(\lambda, \eta^1_\lambda(\mathbf{a}))$. But then, by Lemma C 3, $((C, D), \mathbf{a}^C) \in P_{sym}(\lambda)$, contradicting Lemma C 10. Therefore, a^1 can only be (D, C). Deduce that $\eta^2_\lambda(\mathbf{a})$ and η determine the same direction and, by Lemma C 8, that $\mathbf{a}^{A,2} \in P(\lambda, \eta)$. Now, since (0,1) is an alternation for the path $\mathbf{a}$, we know from Lemma C 12 that $a^t \neq (C, C)$ for all $t > 1$. Thus, $a^2 \in \{(C, D), (D, C)\}$. By Lemma C 3, ${}_2\mathbf{a} \in P(\lambda, \eta^2_\lambda(\mathbf{a}))$. But, since $\eta^2_\lambda(\mathbf{a})$ and η determine the same direction, we have ${}_2\mathbf{a} \in P(\lambda, \eta)$. We also know that $a^2 \in \{(C, D), (D, C)\}$. Thus, the same arguments that were used to show that $a^1 = (D, C)$ can now be used to show that $a^3 \in \{(C, D), (D, C)\} \setminus \{a^2\}$. Proceeding like this, we may conclude that $\mathbf{a} \in \mathcal{A}$. □

Lemma C 16. *For every $\eta \in \mathbb{R}^2_{++}$ and $\mathbf{a} \in P(\lambda, \eta)$, if $a^0 = (C, D)$ and $a^1 = (D, C)$, then $\mathbf{a} \in \mathcal{E}_2\mathcal{A}$.*

Proof. Since $a^0 = (C, D)$, it follows from Lemma C 9 that $\eta_1 \leq \eta_2$. Suppose $\eta_1 = \eta_2$. We know from Lemma C 15 that $\mathbf{a} \in \mathcal{A}$ and, hence, that $\mathbf{a} \in \mathcal{E}_2\mathcal{A}$. Next, suppose $\eta_1 < \eta_2$. Since $a^0 = (C, D)$ and $a^1 = (D, C)$, η and η^2_λ determine the same direction. Hence, $\eta^2_{1,\lambda} < \eta^2_{2,\lambda}$. Since the path $\mathbf{a}$ has an alternation $(0, 1)$, we know from Lemma C 12 that $a^t \neq (C, C)$ for all $t > T$. Hence, Lemma C 9 implies that $a^2 = (C, D)$. Moreover, since $a^1 = (D, C)$, Lemma C 9 shows that $\eta^1_{1,\lambda} \geq \eta^1_{2,\lambda}$. If $\eta^1_{1,\lambda} = \eta^1_{2,\lambda}$, we know from Lemma C 15 that $P(\lambda, \eta^1_\lambda) \subseteq \mathcal{A}$. Therefore, $\mathbf{a} \in \mathcal{E}_2\mathcal{A}$, as desired. Now suppose $\eta^1_{1,\lambda} > \eta^1_{2,\lambda}$. Recall that $a^1 = (D, C)$ and $a^2 = (C, D)$. Thus, η^1_λ and η^3_λ determine the same direction. As a result, we have $\eta^3_{1,\lambda} > \eta^3_{2,\lambda}$. Recall that $a^t \neq (C, C)$ for all $t > T$. Hence, Lemma C 9 implies that $a^3 = (D, C)$. Proceeding like this, we have $\eta^{2t}_{1,\lambda} < \eta^{2t}_{2,\lambda}$ and $\eta^{2t+1}_{1,\lambda} > \eta^{2t+1}_{2,\lambda}$ for all $t \in \mathcal{T}$. Lemma C 9 implies that $a^{2t} = (C, D)$ and $a^{2t+1} = (D, C)$ for all $t \in \mathcal{T}$. That is, $\mathbf{a} = \mathbf{a}^{A,2} \in \mathcal{E}_2\mathcal{A}$. □

Lemma C 17. *$P_{++}(\lambda) \subseteq \mathcal{EC} \cup \mathcal{EA}$.*

Proof. Fix any $\mathbf{a} \in P_{sym}(\lambda)$. It is convenient to write η for η_{sym}. First, we are going to show that $\mathbf{a}$ is either one of intertemporal or intratemporal cooperation, that is, $\mathbf{a} \in \mathcal{A} \cup \{\mathbf{a}^C\}$. By Lemma C 6, (D, D) cannot be played along any efficient path. Hence, $a^0 \in \{(C, C), (D, C), (C, D)\}$. Suppose $a^0 = (C, D)$. Lemma C 15 shows that $\mathbf{a} \in \mathcal{A}$. A symmetric argument shows that if $a^0 = (D, C)$, then $\mathbf{a} \in \mathcal{A}$. Alternatively, suppose that $a^0 = (C, C)$. By Lemma C 5, the path $\mathbf{a}^C$ is efficient. Assume that $\mathbf{a} \neq \mathbf{a}^C$. Let T be the first period t such that $a^t \neq (C, C)$. By construction, for any $0 < t \leq T$, the direction η^t_λ is the same as η. In particular, ${}_T\mathbf{a} \in P(\lambda, \eta)$. W.l.o.g, assume $a^T = (C, D)$. The proof in Lemma C 15 shows that $a^{T+1} = (D, C)$. Thus, $(T, T+1)$ is an alternation for $\mathbf{a}$. Since $a^0 = (C, C)$, Lemma C 12 is contradicted. Conclude that $P(\lambda, \eta) = \{\mathbf{a}^C\}$. The next step is to characterize the efficient paths for any direction $\eta \in \mathbb{R}^2_{++}$ such that $\eta_1 \neq \eta_2$. It is w.l.o.g. to assume $0 < \eta_1 < \eta_2$. Take any $\mathbf{a} \in P(\lambda, \eta)$. Since $\eta_1 < \eta_2$, Lemma C 9 shows that $v_{1,\lambda}(\mathbf{a}) \leq v_{2,\lambda}(\mathbf{a})$ and $a^0 \neq (D, C)$. By Lemma C 6, (D, D) cannot be played along any efficient path. Hence, $a^0 \in \{(C, C), (C, D)\}$. If $a^0 = (C, C)$, Lemma C 13 shows that ${}_1\mathbf{a} \in \mathcal{C}_2$ and, hence, $\mathbf{a} \in \mathcal{C}_2$.

Alternatively, suppose $a^0 = (C, D)$. By Lemma C 7, the constant path $\big((C, D), (C, D), ...\big)$ is not efficient. Let T be the first period t such that $a^t \neq (C, D)$. Suppose $a^T = (C, C)$. If $\eta^T_{1,\lambda} < \eta^T_{2,\lambda}$, then Lemma C 13 shows that ${}_T\mathbf{a} \in \mathcal{C}_2$. If $\eta^T_{1,\lambda} = \eta^T_{2,\lambda}$, we have already shown that ${}_T\mathbf{a} = \mathbf{a}^C$. If $\eta^T_{1,\lambda} > \eta^T_{2,\lambda}$, Lemma C 13 implies that ${}_T\mathbf{a} \in \mathcal{C}_1$. Moreover, Lemma C 14 implies that $a^t \neq (D, C)$ for all $t > T$. Therefore, ${}_T\mathbf{a} = \mathbf{a}^C$. Finally, suppose $a^T = (D, C)$. Lemma C 16 shows that ${}_{T-1}\mathbf{a} \in \mathcal{E}_2\mathcal{A}$ and, hence, $\mathbf{a} \in \mathcal{E}_2\mathcal{A}$. □

Lemma C 18. *If $s_\lambda(\mathbf{a}^C) > s_\lambda(\mathbf{a}^{A,2})$, then $P_{++}(\lambda) \supseteq \mathcal{E}\mathcal{C}$. Else, $P_{++}(\lambda) \supseteq \mathcal{E}\mathcal{A}$.*

Proof. We prove the first assertion. Also, we only show that $\mathcal{E}_2\mathcal{C}_2 \subseteq P_{++}(\lambda)$. All other parts of the proof follow from analogous arguments. Recall the definition of the paths $\mathbf{a}^C(T)$, $T \in \mathcal{T}$, which we gave prior to Lemma C 10. Note that $\mathbf{a}^C(T) \in \mathcal{E}_2\mathcal{C}_2$ for every T. The first step of the proof is to show that every such path is efficient. Let $\eta(0) := (1, 1)$ and for every $T \geq 1$, let

$$\eta(T) := \Big(v_{2,\lambda}\big(\mathbf{a}^C(T)\big) - v_{2,\lambda}\big(\mathbf{a}^C(T-1)\big), v_{1,\lambda}\big(\mathbf{a}^C(T-1)\big) - v_{1,\lambda}\big(\mathbf{a}^C(T)\big)\Big).$$

We are going to show that $\mathbf{a}^C(T) \in P(\lambda, \eta(T))$ for every $T \in \mathcal{T}$. The proof is by induction. From Lemma C 17, we know that $\mathbf{a}^C(0) \in P(\lambda, \eta(0))$. Suppose that $\mathbf{a}^C(T) \in P(\lambda, \eta(T))$ for some $T > 0$. We have to show that $\mathbf{a}^C(T+1) \in P(\lambda, \eta(T+1))$. From Lemma C 17, we know that $P(\lambda, \eta(T+1)) \subseteq \mathcal{E}\mathcal{C}$. It is therefore enough to show that

$$s_\lambda\big(\mathbf{a}^C(T+1), \eta(T+1)\big) \geq s_\lambda\big(\mathbf{a}, \eta(T+1)\big) \quad \forall \mathbf{a} \in \mathcal{E}\mathcal{C}. \tag{26}$$

Observe that, by construction, $\frac{\eta_1(T+1)}{\eta_2(T+1)} < 1$. Lemma C 9 implies that

$$v_{1,\lambda}(\mathbf{a}) \leq v_{2,\lambda}(\mathbf{a}) \quad \forall \mathbf{a} \in P\big(\lambda, \eta(T+1)\big).$$

Hence, it is enough to show that the inequality in (26) is satisfied for all paths $\mathbf{a} \in \mathcal{E}_2\mathcal{C}_2$. First, we verify that the inequality is satisfied for all paths in the set $\{\mathbf{a}^C(T') : T' \in \mathcal{T}\} \subseteq \mathcal{E}_2\mathcal{C}_2$. If $T' > T+1$, then the inequality in (26) is equivalent to

$$\beta_\lambda(d) + ... + [\beta_\lambda(d)]^{T'-T-1} \leq \beta_\lambda(b) + ... + [\beta_\lambda(b)]^{T'-T-1},$$

which holds since $\beta_\lambda(d) < \beta_\lambda(b)$. If $T' = T$, then (26) holds since, by the definition of $\eta(T+1)$, we have

$$s_\lambda\big(\mathbf{a}^C(T+1), \eta(T+1)\big) = s_\lambda\big(\mathbf{a}^C(T), \eta(T+1)\big). \tag{27}$$

Finally, take $T' < T$. By the induction hypothesis, $\mathbf{a}(T) \in P(\lambda, \eta(T))$ and, hence,

$$s_\lambda\big(\mathbf{a}^C(T), \eta(T)\big) \geq s_\lambda\big(\mathbf{a}^C(T'), \eta(T)\big).$$

The above inequality is equivalent to

$$\frac{\eta_2(T)}{\eta_1(T)} \geq \frac{v_{1,\lambda}(\mathbf{a}^C(T')) - v_{1,\lambda}(\mathbf{a}^C(T))}{v_{2,\lambda}(\mathbf{a}^C(T)) - v_{2,\lambda}(\mathbf{a}^C(T'))}.$$

Also, by construction, $\frac{\eta_2(T+1)}{\eta_1(T+1)} > \frac{\eta_2(T)}{\eta_1(T)}$. Hence, we have

$$\frac{\eta_2(T+1)}{\eta_1(T+1)} > \frac{v_{1,\lambda}(\mathbf{a}^C(T')) - v_{1,\lambda}(\mathbf{a}^C(T))}{v_{2,\lambda}(\mathbf{a}^C(T)) - v_{2,\lambda}(\mathbf{a}^C(T'))}. \tag{28}$$

Combining (27) and (28) yields

$$s_\lambda\big(\mathbf{a}^C(T+1), \eta(T+1)\big) \geq s_\lambda\big(\mathbf{a}^C(T'), \eta(T+1)\big),$$

as desired. Now, we are going to show that the inequality in (26) is satisfied for every path $\mathbf{a} \in \mathcal{E}_2\mathcal{C}_2 \setminus \{\mathbf{a}^C(T') : T' \in \mathcal{T}\}$. For such a path $\mathbf{a}$, there are periods $T^* < T^{**}$ such that (C, D) is played in all periods $t < T^*$, (C, D) is played in period T^{**} as well, and (C, C) is played in all other periods. Letting $\phi := 1 - [\beta_\lambda(c)]^{T^{**}-T^*}$, observe that

$$v_{i,\lambda}(\mathbf{a}) = \phi v_{i,\lambda}(\mathbf{a}^C(T^*)) + (1-\phi) v_{i,\lambda}(\mathbf{a}^C(T^*+1)) \quad \forall i \in \{1, 2\}.$$

Thus, the vector $v_\lambda(\mathbf{a}) \in \mathbb{R}^2$ is a convex combination of the vectors $v_\lambda(\mathbf{a}^C(T^*))$ and $v_\lambda(\mathbf{a}^C(T^*+1))$. Conclude that (26) holds for all paths $\mathbf{a} \in \mathcal{E}_2\mathcal{C}_2$ and, hence, that every path $\mathbf{a}^C(T')$, $T' \in \mathcal{T}$, is efficient. This completes the first step of the proof. The second and final step is to show that every path $\mathbf{a} \in \mathcal{E}_2\mathcal{C}_2 \setminus \{\mathbf{a}^C(T') : T' \in \mathcal{T}\}$ is efficient. But, as we just showed, the payoff vector from any such path $\mathbf{a}$ is a convex combination of the payoff vectors from the paths $\mathbf{a}^C(T)$ and $\mathbf{a}^C(T+1)$ for some T. Moreover, we have already established that $\mathbf{a}^C(T), \mathbf{a}^C(T+1) \in P(\lambda, \eta(T+1))$. It follows that $\mathbf{a} \in P(\lambda, \eta(T+1))$ as well. □

C .2 Proof of Corollary 7.2

Proof of Corollary 7.2. Let λ' be such that

$$(1-\lambda')d = (1-\lambda')b + \beta_{\lambda'}(b)\frac{c}{1-\beta_0(c)}. \tag{29}$$

Let $\underline{\lambda} = \max\{0, \lambda'\}$. Observe that $(1-\lambda)b + \beta_\lambda(b)\frac{c}{1-\beta_0(c)}$ increases as λ increases. Hence, for any $\underline{\lambda} < \lambda < 1$, we have $(1-\lambda)d < (1-\lambda)b + \beta_\lambda(b)\frac{c}{1-\beta_0(c)}$.

Take λ, ε, and a path $\mathbf{a}$ as in the hypothesis of the corollary. Take any ε' such that

$$(1-\lambda)d < \varepsilon' \leq \min\{\varepsilon, (1-\lambda)b + \beta_\lambda(b)\frac{c}{1-\beta_0(c)}\}.$$

Since $\mathbf{a} \in IR^\varepsilon(\lambda)$ and $\varepsilon' \leq \varepsilon$, we have $\mathbf{a} \in IR^{\varepsilon'}(\lambda)$. Corollary 7.1 implies that $\mathbf{a} \in SIR^{\varepsilon'}(\lambda)$. To support $\mathbf{a}$ in a subgame perfect equilibrium, we consider the following strategy $\sigma_i \in \Sigma_i$ for player i: (A) play a_i^t in period t as long as a^{t-1} was played last period. After any deviation from (A), then (B) play (D, D) forever after. If there are any deviations while in phase (B), then begin phase (B) again. If player i deviates in phase (A) and then conforms, he receives at most d the period he deviates, and zero afterwards. Thus, his total payoff is no greater than $(1-\lambda)d$ and the gain from deviating is less than $(1-\lambda)d - \varepsilon'$, which is less than zero by the choice of ε'. Thus, no player has an incentive to deviate in phase (A). Since playing (D, D) after any history is a Nash equilibrium of the repeated game, no player wants to deviate in phase (B). □

C .3 A Numerical Example under IMI

This section gives a numerical example in which intertemporal cooperation is efficient for all $\lambda \in [0, 1)$, while intratemporal cooperation is never efficient. Moreover, every alternating path $\mathbf{a} \in \mathcal{A}$ strictly Pareto dominates the path $\mathbf{a}^C = ((C,C),(C,C),...)$ in the limit as $\lambda \to 1$.

Let $b = -1, c = 1.1, d = 3$ and $\beta_0(b) = 0.8, \beta_0(c) = 0.72, \beta_0(d) = 0.7$. It is readily seen that IMI is satisfied. Also,

$$\frac{c}{1-\beta_0(c)} \approx 3.93,$$
$$\frac{1}{2}\frac{b}{1-\beta_0(b)} + \frac{1}{2}\frac{d}{1-\beta_0(d)} = 2.5.$$

Hence, the inequality in (9) is satisfied. The latter inequality and the fact that $2c > b + d$ both stack the deck against intertemporal cooperation being efficient. For instance, the latter inequality says that (C,C) is the unique symmtric first best outcome in the stage game. Yet, things are different in the repeated game. In particular, fix any alternating path $\mathbf{a} \in \mathcal{A}$ and note that:

$$\lim_{\lambda' \to 1} s_{\lambda'}(\mathbf{a}) = 2\frac{b+d}{1-\beta_0(b)+1-\beta_0(d)} = 8 > s_\lambda(\mathbf{a}^C) = 7.86.$$

Since $s_{\lambda'}(\mathbf{a})$ is decreasing in λ', we can conclude that intertemporal cooperation is efficient for all $\lambda \in [0,1)$. Moreover, since

$$\lim_{\lambda' \to 1} v_{i,\lambda'}(\mathbf{a}) = 4 > 3.93 \approx v_i(\mathbf{a}^C) \quad \forall i \in \{1,2\},$$

intertemporal cooperation Pareto dominates the path $\mathbf{a}^C$ for all λ sufficiently close to 1. The latter path remains dominated even in the limit when both players are completely patient.

C .4 When Both Intra- and Inter-temporal Cooperation are Efficient

In this section, we characterize the Pareto set for irregular λ, that is, for λ such that $\mathbf{a}^{A,2}, \mathbf{a}^C \in P(\lambda)$. Some additional notation is needed first. Let $\mathcal{CA}_1$ be the set of all paths $\mathbf{a} \in A^\infty$ such that $a^t \in \{(C,C),(D,C),(C,D)\}$ for every t and, in addition, $a^t = (D,C)$ for some t if and only if $a^{t+1} = (C,D)$. Note that if we let $X := (C,C)$ and $Y = ((D,C),(C,D))$, then we can identify $\mathcal{CA}_1$ with the Cartesian product $\{X,Y\}^\infty$. Define $\mathcal{CA}_2$ analogously, with the difference that (C,D) is played in period t if and only if (D,C) is played in period $t+1$. Note that $\mathbf{a}^C, \mathbf{a}^{A,i} \in \mathcal{CA}_i, i \in \{1,2\}$. Let $\mathcal{CA} := \mathcal{CA}_1 \cup \mathcal{CA}_2$. Define $\mathcal{ECA}_1$ to be the set of paths such that (D,C) is played until some period $T \geq 0$ and

${}_T\mathbf{a} \in \mathcal{CA}$. Define $\mathcal{ECA}_2$ analogously and let $\mathcal{ECA} := \mathcal{ECA}_1 \cup \mathcal{ECA}_2$. Next, let $X := (C,C)$, $Y := ((D,C),(C,D))$, and $Z := ((C,D),(D,C))$. Let $\mathcal{CA}^M := \{X,Y,Z\}^\infty$ and identify $\mathcal{CA}^M$ with a subset of A^∞ in the obvious manner. Observe that $\mathcal{CA}_1 \cup \mathcal{CA}_2 \subseteq \mathcal{CA}^M$. Also, for any path $\mathbf{a} \in \mathcal{CA}^M$, note that $\mathbf{a} \in P_{sym}(\lambda)$ if and only if $s_\lambda(\mathbf{a}^{A,2}) = s_\lambda(\mathbf{a}^C)$. Define $\mathcal{ECA}_1^M$ to be the set of paths such that (D,C) is played until some period $T \geq 0$ and ${}_T\mathbf{a} \in \mathcal{CA}^M$. Define $\mathcal{ECA}_2^M$ analogously. Finally, let $\mathcal{ECA}^M := \mathcal{ECA}_1^M \cup \mathcal{ECA}_1^M$. The proof of the next theorem follows from arguments analogous to those used in the proof of Theorem 7.1.

Theorem C 2. *Suppose (21) holds. If the first inequality is strict, then $P(\lambda) = \mathcal{ECA}$; else, $P(\lambda) = \mathcal{ECA}^M$.*

D Results in Section 8

Assume DMI. Recall that we write $v_i(a)$ for $v_{i,\lambda}(a,a,...)$, for all $\lambda \in [0,1)$ and $a \in A$, and that the utility of a constant path does not depend on λ. Also, A^E denotes the set of actions $a \in A$ such that $v_1(a) = v_2(a)$. We maintain the assumption that the set $\arg\max_{a \in A^E} v_1(a)$ is a singleton whenever the set A^E is nonempty and we let a^* denote the action that attains the maximum. The players' most preferred paths, $\mathbf{a}^{max,1}$ and $\mathbf{a}^{max,2}$, are assumed to be unique as well.

Lemma D 19. *For every $\lambda \in [0,1)$, $\eta \in \mathbb{R}^2_+$, and $\mathbf{a} \in P(\lambda,\eta)$, if $v_{1,\lambda}(\mathbf{a}) \leq v_{2,\lambda}(\mathbf{a})$, then $\beta_{1,\lambda}(a^0) \leq \beta_{2,\lambda}(a^0)$.*

Proof. Suppose by way of contradiction that $\beta_{1,\lambda}(a^0) > \beta_{2,\lambda}(a^0)$. Then, $\frac{\eta^1_{1,\lambda}}{\eta^1_{2,\lambda}} > \frac{\eta_1}{\eta_2}$. Since $\mathbf{a} \in P(\lambda,\eta)$ and ${}_1\mathbf{a} \in P(\lambda,\eta^1_\lambda)$, it must then be the case that

$$v_{1,\lambda}({}_1\mathbf{a}) \geq v_{1,\lambda}(\mathbf{a}), \quad \text{and} \quad v_{2,\lambda}({}_1\mathbf{a}) \leq v_{2,\lambda}(\mathbf{a}). \tag{30}$$

From (10), we know that $v_{i,\lambda}(\mathbf{a})$ is a convex combination of $v_i(a^0)$ and $v_{i,\lambda}({}_1\mathbf{a})$ for every $i \in \{1,2\}$. Thus, the inequalities in (30) are possible only if

$$v_1(a^0) \leq v_{1,\lambda}(\mathbf{a}) \quad \text{and} \quad v_{2,\lambda}(\mathbf{a}) \leq v_2(a^0).$$

Since, by assumption, $v_{1,\lambda}(\mathbf{a}) \leq v_{2,\lambda}(\mathbf{a})$, we have $v_1(a^0) \leq v_2(a^0)$. By DMI, $\beta_{1,\lambda}(a^0) \leq \beta_{2,\lambda}(a^0)$, a contradiction. □

Lemma D 20. *For every $\lambda \in [0,1)$, $\eta \in \mathbb{R}^2_+$, and $\mathbf{a} \in P(\lambda,\eta)$, if $v_{1,\lambda}(\mathbf{a}) = v_{2,\lambda}(\mathbf{a})$, then $\mathbf{a} = (a^*,a^*,...)$.*

Proof. Since $v_{1,\lambda}(\mathbf{a}) = v_{2,\lambda}(\mathbf{a})$, Lemma D 19 implies that $\beta_{1,\lambda}(a^0) = \beta_{2,\lambda}(a^0)$. Hence, $a^0 \in A^E$. We are now going to show that $a^1 \in A^E$ as well. Since $v_1(a^0) = v_2(a^0)$ and $v_{1,\lambda}(\mathbf{a}) = v_{2,\lambda}(\mathbf{a})$, equation (10) implies that $v_{1,\lambda}({}_1\mathbf{a}) = v_{2,\lambda}({}_1\mathbf{a})$. Also, by Lemma C 3, ${}_1\mathbf{a} \in P(\lambda, \eta^1_\lambda)$. But then Lemma D 19 implies that $a^1 \in A^E$. Repeating the argument, we may conclude that $a^t \in A^E$ for every $t \in \mathcal{T}$. By Lemma C 6, $\mathbf{a} = (a^*, a^*, ...)$. □

Lemma D 21. *For every $\lambda \in [0,1)$, $\eta \in \mathbb{R}^2_{++}$, and $\mathbf{a} \in P(\lambda, \eta)$, if $v_{1,\lambda}(\mathbf{a}) < v_{2,\lambda}(\mathbf{a})$ and $a^0 \in A^E$, then $v_{1,\lambda}({}_1\mathbf{a}) < v_{1,\lambda}(\mathbf{a})$ and $v_{2,\lambda}({}_1\mathbf{a}) > v_{2,\lambda}(\mathbf{a})$.*

Proof. The proof is by contradiction. Suppose that the first inequality fails while the second holds so that: $v_{1,\lambda}({}_1\mathbf{a}) \geq v_{1,\lambda}(\mathbf{a})$ and $v_{2,\lambda}({}_1\mathbf{a}) > v_{2,\lambda}(\mathbf{a})$. But then the path ${}_1\mathbf{a}$ strictly Pareto dominates $\mathbf{a}$, contradicting the efficiency of $\mathbf{a}$. Similarly, if the first inequality held while the second failed, then the path ${}_1\mathbf{a}$ would be strictly Pareto dominated by $\mathbf{a}$, contradicting the fact that ${}_1\mathbf{a} \in P(\lambda, \eta^1_\lambda)$, which we proved in Lemma C 3. Finally, consider the case when both of the desired inequalities fail so that:

$$v_{1,\lambda}({}_1\mathbf{a}) \geq v_{1,\lambda}(\mathbf{a}) \quad \text{and} \quad v_{2,\lambda}({}_1\mathbf{a}) \leq v_{2,\lambda}(\mathbf{a}). \tag{31}$$

Since $a^0 \in A^E$, equation (10) and the inequalities in (31) together imply that

$$v_{1,\lambda}(\mathbf{a}) \geq v_1(a^0) = v_2(a^0) \geq v_{2,\lambda}(\mathbf{a}),$$

contradicting the assumption that $v_{1,\lambda}(\mathbf{a}) < v_{2,\lambda}(\mathbf{a})$. □

Lemma D 22. *For every $\lambda \in [0,1)$, $\eta \in \mathbb{R}^2_{+}$, and $\mathbf{a} \in P(\lambda, \eta)$, if $\beta_{1,\lambda}(a^0) < \beta_{2,\lambda}(a^0)$, then $\beta_{1,\lambda}(a^t) < \beta_{2,\lambda}(a^t)$ for all $t > 0$.*

Proof. Suppose by way of contradiction that there is some t such that $\beta_{1,\lambda}(a^t) \geq \beta_{2,\lambda}(a^t)$. Let T be the smallest such t. Since $\beta_{1,\lambda}(a^t) < \beta_{2,\lambda}(a^t)$ for all $t < T$,

$$\frac{\eta^T_{1,\lambda}(\mathbf{a})}{\eta^T_{2,\lambda}(\mathbf{a})} = \frac{\eta_1 \prod_{0 \leq t < T} \beta_{1,\lambda}(a^t)}{\eta_2 \prod_{0 \leq t < T} \beta_{2,\lambda}(a^t)} < \frac{\eta_1}{\eta_2}.$$

Thus, any path $\hat{\mathbf{a}} \in P(\lambda, \eta^T_\lambda(\mathbf{a}))$ should satisfy

$$v_{1,\lambda}(\hat{\mathbf{a}}) \leq v_{1,\lambda}(\mathbf{a}) \quad \text{and} \quad v_{2,\lambda}(\mathbf{a}) \leq v_{2,\lambda}(\hat{\mathbf{a}}).$$

Also, because $\beta_{1,\lambda}(a^0) < \beta_{2,\lambda}(a^0)$, Lemma D 19 implies that $v_{1,\lambda}(\mathbf{a}) < v_{2,\lambda}(\mathbf{a})$. Conclude that

$$v_{1,\lambda}(\hat{\mathbf{a}}) < v_{2,\lambda}(\hat{\mathbf{a}}) \quad \forall \hat{\mathbf{a}} \in P(\lambda, \eta^T_\lambda(\mathbf{a})). \tag{32}$$

By Lemma C 3, ${}_T\mathbf{a} \in P(\lambda, \eta^T_\lambda(\mathbf{a}))$ and, hence, $v_{1,\lambda}({}_T\mathbf{a}) < v_{2,\lambda}({}_T\mathbf{a})$. By Lemma D 19, $\beta_{1,\lambda}(a^T) \leq \beta_{2,\lambda}(a^T)$. But, by the choice of T, $\beta_{1,\lambda}(a^T) \geq \beta_{2,\lambda}(a^T)$. Hence, $\beta_{1,\lambda}(a^T) = \beta_{2,\lambda}(a^T)$. By DMI, we may conclude that $v_1(a^T) = v_2(a^T)$ and, hence, that $a^T \in A^E$. It follows from Lemmas C 5 and C 6 that $\mathbf{a}' := (a^*, a^*, ...) \in P(\lambda, \eta^T_\lambda(\mathbf{a}))$. But $v_{1,\lambda}(\mathbf{a}') = v_{2,\lambda}(\mathbf{a}')$, contradicting (32). □

Proof of Theorem 8.1. If $\mathbf{a}^{max,1} = \mathbf{a}^{max,2}$, it is clear that $P(\lambda) = \{\mathbf{a}^{max,1}\} = \{\mathbf{a}^{max,2}\}$ for all $\lambda \in [0,1)$. From now on, assume that $\mathbf{a}^{max,1} \neq \mathbf{a}^{max,2}$. Take $\lambda \in [0,1)$, $\eta \in \mathbb{R}^2_+$, and $\mathbf{a} \in P(\lambda, \eta)$. If $\eta_i = 0$ and $\eta_j > 0$ for some $i \in \{1,2\}$ and $j \neq i$, it is clear that $\mathbf{a} = \mathbf{a}^{max,j}$. For the rest of the proof, assume that $\eta \in \mathbb{R}^2_{++}$. If $v_{1,\lambda}(\mathbf{a}) = v_{2,\lambda}(\mathbf{a})$, then Lemma D 20 shows that $\mathbf{a} = (a^*, a^*, ...)$, as desired. Next, consider the case when $v_{1,\lambda}(\mathbf{a}) < v_{2,\lambda}(\mathbf{a})$. The case when when $v_{1,\lambda}(\mathbf{a}) > v_{2,\lambda}(\mathbf{a})$ follows from analogous arguments and is omitted. By Lemma D 19, $\beta_{1,\lambda}(a^0) \leq \beta_{2,\lambda}(a^0)$. We are going to show that there exists T such that $\beta_{1,\lambda}(a^t) < \beta_{2,\lambda}(a^t)$ for all $t > T$. If $\beta_{1,\lambda}(a^0) < \beta_{2,\lambda}(a^0)$, Lemma D 22 shows that $\beta_{1,\lambda}(a^t) < \beta_{2,\lambda}(a^t)$ for all $t > 0$, as desired. Now, suppose that $\beta_{1,\lambda}(a^0) = \beta_{2,\lambda}(a^0)$. Let $T \geq 1$ be the first period t such that $\beta_{1,\lambda}(a^t) \neq \beta_{2,\lambda}(a^t)$. By DMI, such T exists since $v_{1,\lambda}(\mathbf{a}) < v_{2,\lambda}(\mathbf{a})$. By the choice of T, $\beta_{1,\lambda}(a^t) = \beta_{2,\lambda}(a^t)$ for every $0 \leq t < T$. Lemma C 6 implies that $a^t = a^*$ for every $0 \leq t < T$. Since $a^0 = a^*$, Lemma D 21 implies that

$$v_{1,\lambda}({}_1\mathbf{a}) < v_{1,\lambda}(\mathbf{a}) \quad \text{and} \quad v_{2,\lambda}(\mathbf{a}) < v_{2,\lambda}({}_1\mathbf{a}).$$

Since, by assumption, $v_{1,\lambda}(\mathbf{a}) < v_{2,\lambda}(\mathbf{a})$, we may conclude that $v_{1,\lambda}({}_1\mathbf{a}) < v_{2,\lambda}({}_1\mathbf{a})$. Applying Lemma D 21 repeatedly, we have $v_{1,\lambda}({}_t\mathbf{a}) < v_{2,\lambda}({}_t\mathbf{a})$ for every $0 < t \leq T$ and, in particular, $v_{1,\lambda}({}_T\mathbf{a}) < v_{2,\lambda}({}_T\mathbf{a})$. By Lemma D 19, we may conclude that $\beta_{1,\lambda}(a^T) \leq \beta_{2,\lambda}(a^T)$ and, by the choice of T, that $\beta_{1,\lambda}(a^T) < \beta_{2,\lambda}(a^T)$. It follows from Lemma D 22 that $\beta_{1,\lambda}(a^t) < \beta_{2,\lambda}(a^t)$ for all $t > T$.

Thus, in both cases, we know that there is some T such that $\beta_{1,\lambda}(a^t) < \beta_{2,\lambda}(a^t)$ for every $t \geq T$. Let $B := \{a \in A : \beta_{1,\lambda}(a) < \beta_{2,\lambda}(a)\}$. The set B is finite since the set A is finite. Let $l := \min_{a \in B} \frac{\beta_{2,\lambda}(a)}{\beta_{1,\lambda}(a)}$. By construction, $l > 1$ and, for every $t \geq T$,

$$\frac{\eta^t_{2,\lambda}(\mathbf{a})}{\eta^t_{1,\lambda}(\mathbf{a})} = \frac{\eta^T_{2,\lambda}(\mathbf{a})}{\eta^T_{1,\lambda}(\mathbf{a})} \times \prod_{T \leq \tau < t} \frac{\beta_{2,\lambda}(a^\tau)}{\beta_{1,\lambda}(a^\tau)} \geq \frac{\eta^T_{2,\lambda}(\mathbf{a})}{\eta^T_{1,\lambda}(\mathbf{a})} \times l^{t-T}.$$

Since $l > 1$, we know that $l^{t-T} \to +\infty$ as $t \to +\infty$. Thus, player 2's relative weight $\frac{\eta^t_{2,\lambda}(\mathbf{a})}{\eta^t_{1,\lambda}(\mathbf{a})}$ increases to infinity. Conclude that there is some T' such that ${}_{T'}\mathbf{a} = \mathbf{a}^{max,2}$, completing the proof. □

E Strategically Equivalent Games

When formulating a folk theorem, one typically fixes a stage game and then constructs a collection of repeated games by varying the players' level of patience. The fact that the stage game is kept fixed serves an important purpose: the *stage game* becomes an anchor insuring that the *repeated games* one constructs are appropriately related. In this section, we argue that the stage game is an inappropriate, or at least insufficient, anchor when analyzing repeated games with endogenous discounting. We derive the appropriate anchors for such games and show that they amount to fixing a specific convergence path along which discount factors must approach 1. If we add the requirement that the relative impatiences across players and action profiles remain fixed as we vary the level of patience, we arrive at the convergence path which we adopted in Section 4 and in our own folk theorem.

To proceed, consider a repeated game $(A, (g_i, \beta_i)_{i \in I})$ as defined prior to imposing the parametrization in (5). With a slight abuse of notation compared to the conventions made in Section 4, denote this game by Γ. We associate two one shot games, G and S, with each game Γ. The game G has an action space A and payoffs $(\frac{g_i}{1-\beta_i})_i$. It can be viewed as a constrained version of Γ whereby the players choose an action profile $a \in A$ once, and that choice remains binding thereafter. The game S is what is typically called the **stage game**: it has an action space A and payoffs $(g_i)_i$. It is helpful to observe that when the functions β_i are constant, as they are in the standard model, the games S and G are strategically equivalent. Hence, any distinction between S and G is immaterial. This is not the case when discounting is endogenous. Moreover, as we explain next, the one shot game G, rather than S, is the appropriate anchor when one proceeds to vary the level of patience.

First, we need to consider when two repeated games, $(A, (g_i, \beta_i)_{i \in I})$ and $(A, (g'_i, \beta'_i)_{i \in I})$, are strategically equivalent, that is, when each player i has identical preferences over the possible outcomes of the games. To that end, fix a player i and consider his preference relation over the space $\Delta(A^\infty)$ of mixed play paths. As constructed in Section 2, this preference relation is defined by the pair (g_i, β_i) where $\beta_i : A \to (0,1)$ is the player's discount factor and $g_i : A \to \mathbb{R}$ represents his stage payoffs. It is known from Epstein [9] that another pair (g_i', β_i') represents the same preference relation on $\Delta(A^\infty)$ if and only if $\beta_i = \beta_i'$ and

$$\frac{g_i}{1-\beta_i} = \theta_i \frac{g_i'}{1-\beta_i'} + \gamma_i \tag{33}$$

for some constants $\theta_i \in \mathbb{R}_{++}$ and $\gamma_i \in \mathbb{R}$.[17] Equation (33) says that the utility function $v_i : A^\infty \to \mathbb{R}$ over pure play paths is cardinally unique. What concerns us here is the equation does *not* imply that the functions g_i and g_i' are cardinal transformations of one another. In fact, the functions g_i and g_i' need not even be *ordinal* transformations of one another. Going back to the analysis of repeated games, this means that we can find two repeated games, Γ and Γ', that are strategically equivalent, yet whose stage games, S and S', are not. The games G and G' are, on the other hand, strategically equivalent whenever Γ and Γ' are.

More generally, the last paragraph tells us that the functions $(g_i)_i$ do not have a clear ordinal meaning once we allow for the possibility of endogenous discounting. It is for this reason that we view the stage game S as an inappropriate anchor when formulating a folk theorem.[18] Luckily, the last paragraph suggests a solution to the problem it raises: one can vary the level of patience while insuring that the games G remain strategically equivalent. Two comments about this solution are useful. First, since the games G and S are strategically equivalent when discounting is fixed, this solution does not *contradict* the traditional analysis in such settings; it *extends* the analysis to the more general class of games we study. Second, using G as an anchor does not preclude the possibility of also keeping the stage payoffs $(g_i)_i$ fixed. In fact, we did so in Sections 4 through 8 and we do so below in order to simplify one of our derivations. Keeping the stage payoffs fixed may also be natural in applications in which the stage payoffs g_i may have a meaning that precedes the specification of preferences. E.g., they may measure an agent's wages or a firm's profits.

It remains to show that using G as an anchor amounts to fixing a convergence path along which discount factors approach 1. Take a sequence $(\Gamma_m)_{m\in\mathbb{N}}$ of repeated games such

[17]Technically, the uniqueness of the discount factors β_i proved in Epstein [9] requires that A be a connected topological set rather than a finite set as it is in this paper. The uniqueness of the discount factors, however, is not essential for the discussion in this section. The same conclusions apply if we treat the discount factors as given.

[18]The fact that the stage payoffs $(g_i)_i$ need not have a clear ordinal meaning also explains why in Section 7 we defined the repeated prisoners' dilemma game in terms of the ex ante utility from constant play paths rather than in terms of the stage payoffs $(g_i)_i$. In the language of this section, we chose the preference parameters so that the one shot game G is a repeated prisoners' dilemma game. Though not necessary for our results, we could have also chosen the parameters so that both G and S are a prisoners' dilemma game.

that $\Gamma_m = (A, (g_{i,m}, \beta_{i,m})_{i \in I})$ for every $m \in \mathbb{N}$ and $\beta_{i,m} \to 1$ as $m \to +\infty$. It is w.l.o.g. to assume that $g_{i,1} > 0$ for every i. Fix some $m > 1$ and let G_1 and G_m be the one shot games associated with Γ_1 and Γ_m respectively. For the game G_m to be strategically equivalent to G_1, it must be the case that for every $i \in I$ we can find constants $\theta_{i,m} \in \mathbb{R}_{++}, \gamma_{i,m} \in \mathbb{R}$ such that

$$\frac{g_{i,1}}{1-\beta_{i,1}} = \theta_{i,m}\frac{g_{i,m}}{1-\beta_{i,m}} + \gamma_{i,m}.$$

Rearranging, we obtain

$$\beta_{i,m} = 1 - (1-\beta_{i,1})(\theta_{i,m}\frac{g_{i,m}}{g_{i,1}} + \frac{\gamma_{i,m}}{g_{i,1}}). \tag{34}$$

If we assume that $g_{i,m} = g_{i,1}$ for every m and i, then (34) simplifies to

$$\beta_{i,m} = 1 - (1-\beta_{i,1})(\theta_{i,m} + \frac{\gamma_{i,m}}{g_{i,1}}). \tag{35}$$

Equation (34) gives the most general way to specify a convergence path so that the games G_m are all strategically equivalent. One may worry that this specification features both the discount factors $\beta_{i,m}$ and the stage payoffs $g_{i,m}$. Observe, however, that one can insure that the discount factors $\beta_{i,m}$ converge to 1 without imposing any restrictions on the stage payoffs: it is sufficient to choose the parameters $\theta_{i,m}$ and $\gamma_{i,m}$ so that $\theta_{i,m} \to_m 0$ and $\gamma_{i,m} \to_m 0$. In this sense, (34) specifies a convergence path for $\beta_{i,m}$ that is independent of the stage payoffs.

Suppose now that $g_{i,m} = g_{i,1}$ for every m and i. Suppose further that, for a given $i \in I$, we want to insure that the relative impatiences across action profiles $a \in A$ do not very as we vary the level of patience, that is, as $\beta_{i,m} \to_m 1$. This is possible only when $\gamma_{i,m} = 0$ for every m. If, in addition, we want to preserve the relative impatiences across players, then the $\theta_{i,m}$'s must be independent of $i \in I$. Imposing these restrictions on (35), we deduce that $\beta_{i,m} = 1 - \theta_m(1-\beta_{i,1})$ for every i and m. Letting $\lambda_{i,m} := 1 - \theta_{i,m}$, we see that this is exactly the convergence path adopted in Section 4. If we do not require that the stage payoffs remain fixed, i.e., that $g_{i,m} = g_{i,1}$ for every m and i, then we obtain the same convergence path for generic specifications of the stage payoffs $g_{i,m}, i \in I, m \in \mathbb{N}$. We omit the details.

References

[1] D. Abreu, P. K. Dutta, and L. Smith. The folk theorem for repeated games: a NEU condition. *Econometrica: Journal of the Econometric Society*, pages 939–948, 1994.

[2] C. Aliprantis and K. Border. *Infinite Dimensional Analysis*. Springer-Verlag, 2 edition, 1999.

[3] D. K. Backus, B. R. Routledge, and S. E. Zin. Exotic preferences for macroeconomists. In *NBER Macroeconomics Annual*, volume 19, pages 319–390. MIT Press, 2004.

[4] G. Becker and C. Mulligan. The endogenous determination of time preference. *The Quarterly Journal of Economics*, 112(3):729–758, 1997.

[5] R. A. Becker. On the long-run steady state in a simple dynamic model of equilibrium with heterogeneous households. *The Quarterly Journal of Economics*, 95:375–382, 1980.

[6] B. Chen. On effective minimax payoffs and unequal discounting. *Economics Letters*, 100(1): 105–107, 2008.

[7] B. Chen and S. Takahashi. A folk theorem for repeated games with unequal discounting. *Games and Economic Behavior*, 76(2):571–581, 2012.

[8] O. Compte and P. Jehiel. On the role of outside options in bargaining with obstinate parties. *Econometrica*, 70(4):1477–1517, 2002.

[9] L. Epstein. Stationary cardinal utility and optimal growth under uncertainty. *Journal of Economic Theory*, 31(1):133–152, 1983.

[10] L. Epstein. A simple dynamic general equilibrium model. *Journal of Economic Theory*, 41: 68–95, 1987.

[11] I. Fisher et al. *The theory of interest*. Macmillan New York, 1930.

[12] M. Friedman. *The optimum quantity of money and other essays*. Aldine, Chicago, 1969.

[13] D. Fudenberg and D. Levine. Efficiency and observability with long-run and short-run players. *Journal of Economic Theory*, 62.

[14] D. Fudenberg and E. Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, pages 533–554, 1986.

[15] D. Fudenberg and E. Maskin. On the dispensability of public randomization in discounted repeated games. *Journal of Economic Theory*, 53(2):428–438, 1991.

[16] D. Fudenberg and J. Tirole. Game theory, 1991. *Cambridge, Massachusetts*, 1991.

[17] Y. Guéron, T. Lamadon, and C. D. Thomas. On the folk theorem with one-dimensional payoffs and different discount factors. *Games and Economic Behavior*, 73(1):287–295, 2011.

[18] G. W. Harrison, M. I. Lau, and M. B. Williams. Estimating individual discount rates in denmark: A field experiment. *American economic review*, 92:1606–1617, 2002.

[19] A. Kochov. Stationary cardinal utility. University of Rochester, February 2015.

[20] E. C. Lawrance. Poverty and the rate of time preference: evidence from panel data. *Journal of Political economy*, pages 54–77, 1991.

[21] E. Lehrer and A. Pauzner. Repeated games with differential time preferences. *Econometrica*, 67(2):393–412, 1999.

[22] R. Lucas and N. Stokey. Optimal growth with many consumers. *Journal of Economic Theory*, 32(1):139–171, 1984.

[23] I. Obara and J. Park. Repeated games with general discounting. Working Paper, 2015.

[24] I. Obara and F. Zincenko. Tacit collusion and heterogeneity of firms. Working Paper, 2015.

[25] M. M. Opp and J. Y. Zhu. Impatience versus incentives. *Econometrica*, 83(4):1601–1617, 2015.

[26] K. Parthasarathy. *Probability measures on metric spaces*. AMS Chelsea Publishing, 1967.

[27] F. Ramsey. A mathematical theory of saving. *The economic journal*, pages 543–559, 1928.

[28] S. Sorin. On repeated games with complete information. *Mathematics of Operations Research*, 11(1):147–160, 1986.

[29] T. Sugaya. Characterizing the limit set of PPE payoffs with unequal discounting. *Theoretical Economics*, 10(3):691–717, 2015.

[30] H. Uzawa. Time preference, the consumption function, and optimum asset holdings. In J. Wolfe, editor, *Value, Capital and Growth: Papers in Honour of Sir John Hicks*. Aldine, Chicago, 1968.