RATIONALITY & COMPETITION
CRC TRR 190

Repeated Games with Endogenous Discounting

Asen Kochov (University of Rochester)
Yangwei Song (HU Berlin)

Discussion Paper No. 230

March 2, 2020

Collaborative Research Center Transregio 190 | www.rationality-and-competition.de
Ludwig-Maximilians-Universität München | Humboldt-Universität zu Berlin
Spokesperson: Prof. Dr. Klaus M. Schmidt, University of Munich, 80539 Munich, Germany
+49 (89) 2180 3405 | info@rationality-and-competition.de

Repeated Games with Endogenous Discounting*

Asen Kochov Yangwei Song†

February 22, 2020

In a symmetric repeated game with standard preferences, there are no gains from intertemporal trade. In fact, under a suitable normalization of utility, the payoff set in the repeated game is identical to that in the stage game. We show that this conclusion may no longer be true if preferences are recursive and stationary, but not time separable. If so, the players' rates of time preference are no longer fixed, but may vary endogenously, depending on what transpires in the course of the game. This creates opportunities for intertemporal trade, giving rise to new and interesting dynamics. For example, the efficient and symmetric outcome of a repeated prisoner's dilemma may be to take turns defecting, even though the efficient and symmetric outcome of the stage game is to cooperate. A folk theorem shows that such dynamics can be sustained in equilibrium if the players are sufficiently patient.

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

*We are grateful to Paulo Barelli, John Duggan, Hari Govindan, Narayana Kocherlakota, Takuo Sugaya, Satoru Takahashi, John Yiran Zhu, and seminar participants at Bonn, BYU, Harvard, Hitotsubashi, LSE, UBC, and Osaka for their helpful comments. 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. Song gratefully acknowledges financial support by the German Science Foundation (through Collaborative Research Center TRR 190).

†Department of Economics, University of Rochester, Rochester, NY 14627. Institute for Economic Theory 1, Humboldt University Berlin, Spandauer Str. 1, 10178 Berlin, Germany. E-mail: asen.kochov@rochester.edu, yangwei.song@hu-berlin.de

1 Introduction

In symmetric repeated games with standard preferences, there are no gains from intertemporal trade. In fact, under a suitable normalization of utility, the payoff set in the repeated game is identical to that in the stage game. This conclusion rests on two premises. One is the symmetry of the game, which, among other things, requires that the players' rates of time preference be identical. The second premise is the assumption of time separable utility, which implies that rates of time preference are fixed and unaffected by what transpires in the course of the game. If utility is stationary and recursive, but not time separable, the latter is no longer true. Rates of time preference may vary endogenously, creating opportunities for intertemporal trade. This paper investigates the conditions under which endogenous discounting gives rise to intertemporal trade, the forms such trade could take, and whether it can be sustained in strategic settings with no commitment.

More precisely, we assume that players have intertemporal preferences of the form studied in Uzawa [28] and Epstein [9], with the discounted sum of payoffs satisfying the recursion:

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

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 a. Note that if the functions $\beta_i : A \to (0, 1)$ are constant, one obtains the standard model with exogenous, but potentially unequal, rates of time preference.

The preferences in (1) provide a simple way to capture two assumptions that have been central to the study of endogenous discounting.[1] Say that player i exhibits **decreasing marginal impatience (DMI)** if $1 - \beta_i(a)$ decreases the more desirable he finds the constant path $(a, a, ...)$. Fisher et al. [10, p.72] was an early proponent of this assumption, noting that the needs of the present may bear more heavily on a person whose consumption is low. Friedman [11, p.30], on the other hand, noted that DMI leads to "disequilibrium behavior" and advocated the polar case of **increasing marginal impatience (IMI)**. Later, Lucas and Stokey [17] used IMI in a general equilibrium growth model and showed that it insures the existence and stability of a steady state with a nontrivial distribution of capital. Since then, IMI has become a staple in the growth literature, where, following

[1]Chew and Epstein [8] show that the representation in (1) is implied whenever preferences are recursive, stationary, and exhibit indifference to the timing of resolution of uncertainty. It is in this sense that the model is simple: it allows us to focus on the effects of endogenous discounting, while abstracting from other implications, such as attitudes toward the timing of resolution of uncertainty, that come with relaxations of separability.

Lucas and Stokey [17, p.168], it can be roughly interpreted as "diminishing returns to the accumulation of wealth."

We do not take sides in the ongoing debate which assumption is a more appropriate restriction on behavior. Instead, we study each case in turn, characterizing fully the conditions under which gains from intertemporal trade exist and giving sufficient conditions under which such trade can be sustained in equilibrium.

1.1 IMI and Intertemporal Trade

Intuitively, an intertemporal trade happens when an impatient player "borrows" from a more patient counterpart and then "repays the debt" in a later period. Note, in particular, that an intertemporal trade is attained by means of a non-constant play path $(a^0, a^1, ...)$. If such a path is not Pareto dominated by any constant path $(a, a, ...)$, or a one-time randomization among constant paths, we say that there is a **gain from intertemporal trade**. As remarked previously, symmetric games with standard preferences admit no such gains.

Assuming IMI, our first contribution is to show that gains from intertemporal trade exist in a broad class of games. Namely, it is necessary and sufficient that the game feature some **conflict of interest**, by which we mean that no single play path simultaneously maximizes the utility of every player. Interestingly, the proof exploits a well-known behavioral implication of IMI, known as **correlation aversion**, according to which the players dislike positive autocorrelation in the intertemporal distribution of risk. Such behavior has been the subject of a number of recent decision-theoretic as well as experimental studies.[2] To our knowledge, we are the first to study its implications in the context of a strategic interaction.

In games with no commitment, "the lenders" need credible punishments to insure that "the borrowers pay back." Therefore, our second contribution is to obtain a novel folk theorem for games with endogenous discounting. The result, which covers both IMI and DMI, generalizes a folk theorem of Abreu et al. [1] and insures that gains from intertemporal trade can arise in a **subgame perfect equilbrium (SPE)**, whenever the players are sufficiently patient.

Focusing on the prisoner's dilemma, our third contribution is to characterize the play paths that are efficient (first-best). The goal is to showcase some of the novel dynamics brought about by intertemporal trade. Consider Figure 1(a). The inner rectangle is the convex hull of payoffs attainable by constant play paths $(a, a, ...)$. The feasible set is strictly larger, however, as the Pareto frontier shifts out to incorporate gains from intertemporal

[2] Among others, see Bommier et al. [5], Miao and Zhong [18], Andersen et al. [3].

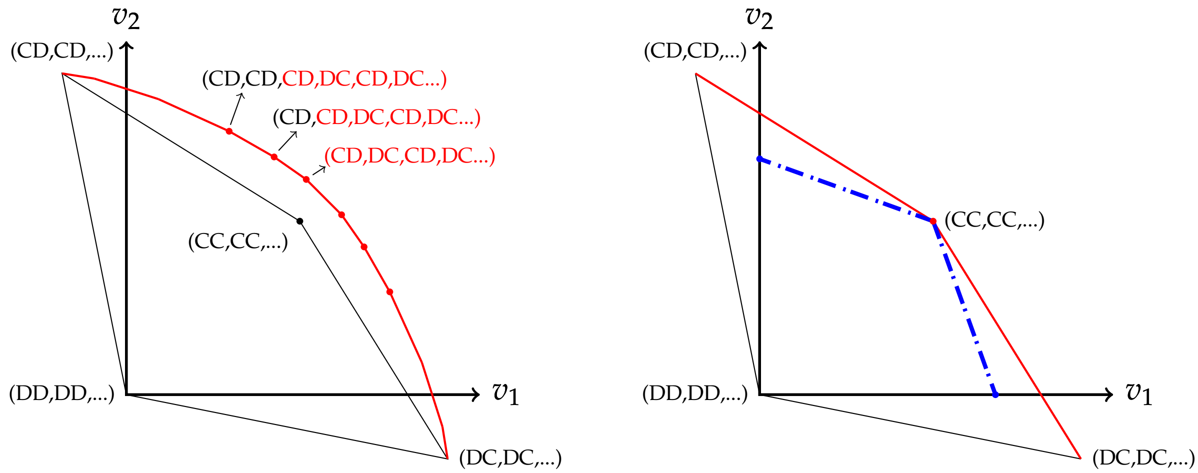

(a) Pareto frontier under IMI

(b) Pareto frontier under DMI

Figure 1: Pareto frontier in a repeated prisoner's dilemma. As usual, C stands for "cooperate" and D for "defect". Extreme points are associated with the play paths that generate them. The dash-dotted lines in Figure 1(b) are independent of the level of patience and give a least-upper bound on the set of SPE payoffs attainable at any level of patience.

trade. Note, in particular, that $(CC, CC, ...)$ need not be efficient in the repeated game, even though it is efficient in the space of constant paths. Instead, the sum of the players' utilities is maximized when they take turns defecting, a form of behavior which we call **intertemporal cooperation**.

Note as well that eventually every efficient path, except the extremes in which one player defects forever, becomes one of intertemporal cooperation. Thus, over time, the continuation utilities (v_1, v_2) induced by such a path converge away from the corners of the Pareto frontier toward an equal division of surplus. From a normative standpoint, this finding presents a curious case in which efficiency and fairness considerations align (at least asymptotically). There are also implications for "the provision of incentives." Recall the usual statement of a folk theorem, according to which if the players are sufficiently patient, a play path can arise in a SPE if and only if it is **sequentially individually rational (SIR)**, which means that the players' security levels are cleared at each point in time. Since efficient paths converge to a symmetric outcome, we see that the SIR constraints become slack over time. This conclusion differs from other studies which we preview momentarily and which typically find that intertemporal trade pushes some players to their security levels.

1.2 DMI and Immiseration

Under DMI, gains from intertemporal trade exist under slightly more stringent conditions. As before, the game must feature conflict of interest, but curiously it should not be the prisoner's dilemma. The nuances of this finding are discussed carefully in Section 6. Here, we focus on another result which presents a more striking contrast between the case of IMI and DMI. Namely, under DMI, the first-best level of intertemporal trade cannot be sustained in a SPE under any level of patience.[3] This finding builds on a literature of so-called "immiseration results," but as we are about to explain the endogeneity of discounting comes with a twist. The earliest antecedent is a conjecture by Ramsey [24] who argued that in a competitive economy with standard preferences, but unequal discount factors, the most patient agents would eventually possess all the capital. This conjecture, confirmed by Rader [23] and Becker [4], rests on the efficiency of competitive markets which mandates that impatient agents borrow to consume early on, while their more patient counterparts lend, deferring their own consumption. More recently, Lehrer and Pauzner [16] considered how Ramsey's conjecture may play out in a strategic setting. Looking at a repeated prisoner's dilemma, they identify a stark tension between the requirements of efficiency and SIR. In parallel with Ramsey [24], efficiency once again requires that the utility of the more patient player be eventually maximized. But in the prisoner's dilemma, this is possible only if the impatient player is pushed below his security level, violating SIR. Thus, first-best outcomes cannot be sustained in a SPE, no matter how patient the players.

Suppose now that marginal impatience is decreasing rather than constant. If at some point in the game player i's continuation utility is higher than j's, i must be more patient. If the play path is efficient, player i must then be rewarded relatively more in the future, which, under DMI, insures that i would sustain the higher level of patience as the game progresses. In this way, DMI and efficiency combine to produce a self-enforcing dynamic that once again pushes one of the players to their most preferred point on the Pareto frontier. What makes our game different, however, is that because the players are a priori symmetric, this dynamic is not unavoidable. If the players coordinate on a play path along which they attain symmetric outcomes, their rates of time preference remain identical and Ramsey's "immiseration dynamic" is never triggered. In the prisoner's dilemma, this means that $(CC, CC, ...)$ is the only first best outcome that can be sustained in a SPE.[4] In sum, the endogeneity of discounting means that Ramsey's immiseration

[3]Some gains from intertemporal trade, though not the first-best level, can still be sustained. See Theorem 6.2 in Section 6.

[4]A suitable generalization applies to all two-player symmetric games. See Theorem 6.1. Unlike Figure 1(a),

dynamic, rather than destroying all first-best equilibria, becomes a form of equilibrium selection.

At this point, one may wonder whether, as the players become increasingly patient, any first-best outcome can be approximated by equilibrium outcomes. Our final result gives an example in which this is not the case: in the prisoner's dilemma, SPE outcomes, other than $(CC, CC, \ldots)$, remain bounded away from the first best frontier even as discount factors converge to 1. See Figure 1(b). Though the formal analysis is limited to the prisoner's dilemma, some of the ideas behind it are more general and reveal additional and interesting ways in which DMI can make a difference in a strategic context. Details are provided in Section 6.

2 The Model

There is a finite set of players: $I := \{1, 2, \ldots, n\}$. In the stage game, player i can choose a pure action a_i in a finite, nonsingleton set A_i. Let $A := \times_{i \in I} A_i$. Player i's mixed actions are denoted by $\alpha_i \in \Delta(A_i)$. In the repeated game, time is discrete and indexed by $t \in \{0, 1, 2, \ldots\} =: \mathcal{T}$. To focus on the effects of endogenous discounting, we keep things as simple as possible and assume perfect monitoring as well as the availability of public randomization. We also assume that deviations from mixed actions are detectable.[5] A complete history thus consists of all past mixed actions, the realized pure action profiles, and the past realizations of the public signal. Given a pure play path $\mathbf{a} = (a^0, a^1, \ldots) \in A^\infty$ and a time period $t > 0$, we let ${}_t\mathbf{a}$ denote the continuation path $(a^t, a^{t+1}, \ldots)$. To describe player i's preferences, we first define a utility function $v_i : A^\infty \to \mathbb{R}$ on the space of pure paths by letting

$$v_i(\mathbf{a}) = g_i(a^0) + \beta_i(a^0) g_i(a^1) + \beta_i(a^0)\beta_i(a^1) g_i(a^2) + \ldots = g_i(a^0) + \beta_i(a^0) v_i({}_1\mathbf{a}). \quad (2)$$

Above, $g_i : A \to \mathbb{R}$ is i's stage payoff and $\beta_i : A \to (0, 1)$ is his discount factor. To define preferences over random paths, let Σ_i be the set of behavioral strategies for player $i \in I$ in the repeated game. A strategy profile $\sigma \in \Sigma = \times_{i \in I} \Sigma_i$ induces a probability distribution on A^∞ in the usual way. Abusing notation, we denote the induced measure by σ as well and let $v_i(\sigma) := \mathbb{E}_\sigma v_i(\mathbf{a})$. We call the preferences thus constructed **Uzawa-Epstein** and

note as well that, under DMI, the first-best outcome in the space of constant play paths remains first-best in the space of all paths.

[5]Except for preferences, our setting is identical to that of Fudenberg and Maskin [12], Abreu et al. [1], and Chen and Takahashi [6]. Note however that the observability of mixed actions is first assumed and then relaxed in the folk theorems of these papers. We conjecture that the same is possible here, but since the proof of our folk theorem is already exceedingly long and since the folk theorem is not our main focus, we leave that step for future work.

identify them with the pair (g_i, β_i).

A repeated game with endogenous discounting is thus a tuple $(A, (g_i, \beta_i)_{i\in I})$. The game is **symmetric** if $A_i = A_j$ for all $i, j \in I$ and the functions $g : a \mapsto (g_1(a), ..., g_n(a))$ and $\beta : a \mapsto (\beta_1(a), ..., \beta_n(a))$ are both symmetric. When defining minmax levels, we assume, as is standard, that the players cannot use correlated strategies. In particular, let $\Sigma_i^* \subset \Sigma_i$ and $\Sigma^* \subset \Sigma$ be the strategy spaces in which the players randomize independently of one another, that is, the public randomization device is not used. Player i's **security level** or **minmax payoff** is the number $\underline{v}_i := \min_{\sigma_{-i}\in \times_{k\neq i}\Sigma_k^*} \max_{\sigma_i \in \Sigma_i^*} v_i(\sigma_i, \sigma_{-i})$. A **minmax strategy against player** i is a strategy profile $\sigma_{-i} \in \times_{k\neq i}\Sigma_k^*$ that attains the minimum in the definition of $\underline{v}_i$.[6]

We use v to denote the function $\sigma \mapsto (v_1(\sigma), ..., v_n(\sigma))$ or a single point in its image. We use v_i^{max} to denote i's maximum feasible payoff in the repeated game. It is known that for each i, there is some $a \in A$ such that $v_i(a, a, ...) = v_i^{max}$. See, e.g., Kochov [13, Lemma 3.4]. Finally, we assume that no player is indifferent among all strategies. This is true if and only if for every i, there are action profiles $a', a'' \in A$ such that $v_i(a'', a'', ...) > v_i(a', a', ...)$ or, equivalently, if the function $a \mapsto g_i(a)/(1 - \beta_i(a))$ is not constant.

3 Security Levels and Patience

In subsequent sections we ask whether gains from intertemporal trade exists and can be sustained in a SPE as the players become sufficiently patient. To answer the latter question, we first need to specify how discount factors approach 1. As we now explain, this requires care when discounting is endogenous. To set the stage, let α^{con} denote the **constant strategy** such that the mixed action $\alpha \in \Delta(A)$ is played after every history and consider the lemma:

Lemma 3.1. *For each i, the minmax strategy against player i and i's best response can be chosen to be constant.*

Given $\alpha \in \Delta(A)$, 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 $a \in A$ by α, and note that

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

As is well-known, note that if discounting is exogenous, that is, if the function β_i is constant, player i's ranking of constant strategies is independent of β_i. In fact, $v_i(\alpha^{con}) >$

[6] Lemma 3.1 shows that if preferences are standard, our definition of minmax values reduces to the minmax values in the stage game. Section 4 explains why we need the above definition.

$v_i(\hat{\alpha}^{con})$ if and only if $g_i(\alpha) > g_i(\hat{\alpha})$. It follows that minmax strategies as well as the players' security levels can be chosen independently of the level of patience. In particular, one can raise the level of patience without affecting what constitutes a punishment in the game. Thus, for high enough levels of patience, the threat of future punishments would eventually outweigh the short-term gains from a deviation, an argument that is at the heart of all folk theorems.

If discounting is endogenous, however, we see that punishments and security levels may vary with the rate of time preference. If, in particular, security levels increase with patience, so that punishments become less severe, it is a priori less clear that raising the level of patience would do anything to deter short-term deviations. Luckily, Lemma 3.1 suggests a solution: minmax strategies would not change if we could restrict the way discount factors approach 1 so that the ranking of constant strategies does not change. To see how this can be done, let $(A, (g_i, \beta_i)_{i\in I})$ be a repeated game and for every $\lambda \in [0, 1)$, let

$$\beta_{i\lambda} := \lambda + (1-\lambda)\beta_i. \tag{3}$$

By construction, $\lambda > \lambda'$ if and only if $\beta_{i\lambda} \gg \beta_{i\lambda'}$,[7] which means that we can interpret λ as a measure of the players' patience. Given λ, we can also scale the stage payoffs by $(1-\lambda)$. As in the case of standard preferences, in which one multiplies by "$(1-\beta)$," this normalization insures that discounted sums do not blow up as the players become increasingly patient. Putting everything together, each λ gives rise to a repeated game $\Gamma(\lambda) = (A, (1-\lambda)g_i, \beta_{i\lambda})$. Letting $v_{i\lambda} : \Sigma \to \mathbb{R}$ be i's utility function in $\Gamma(\lambda)$, we also see that

$$v_{i\lambda}(\alpha^{con}) = \frac{g_i(\alpha)}{1-\beta_i(\alpha)}.$$

As desired, the ranking of constant strategies, in fact the utilities of such strategies, is independent of λ. By Lemma 3.1, so are the minmax strategies against each player and their respective security levels. Eyeing a folk theorem, subsequent sections would thus be concerned with the equilibria of $\Gamma(\lambda)$ as $\lambda \nearrow 1$. We would also write Γ to mean the family of games $\{\Gamma(\lambda) : \lambda \in [0, 1)\}$ as well as the single game $(A, (g_i, \beta_i)_{i\in I})$. To reduce notation and highlight that the ranking of constant strategies is independent of λ, we also write $v_i(\alpha)$ in place of $v_{i\lambda}(\alpha^{con})$.[8]

Another advantage of letting discount factors approach 1 in the above manner is cap-

[7] By $\beta_{i\lambda} \gg \beta_{i\lambda'}$, we mean that $\beta_{i\lambda}(a) > \beta_{i\lambda'}(a)$ for every $a \in A$.

[8] Note that in our notation, $\beta_{i0} = \beta_i$. Similarly, both $\Gamma(0)$ and Γ could designate the single game $(A, (g_i, \beta_i)_{i\in I})$. This overlap in notation makes it possible to switch seamlessly between results concerning a single game and those concerning a family of games.

tured by the equality:

$$\frac{1-\beta_{i\lambda}(a)}{1-\beta_{j\lambda}(a')} = \frac{1-\beta_i(a)}{1-\beta_j(a')} \quad \forall \lambda \in [0,1), i,j \in I, a, a' \in A. \tag{4}$$

Since the fraction on the right-hand side is independent of λ, we see that the relative impatiences across players and action profiles are independent of λ. Keeping fixed the relative (im)patiences *across players* is a requirement familiar from the literature on repeated games with fixed but heterogeneous discount rates. See Chen and Takahashi [6], Lehrer and Pauzner [16], and Sugaya [26], among others. On the other hand, keeping fixed the relative impatiences *across action profiles* is a requirement specific to the study of endogenous discounting. In the context of this paper, it insures that the assumptions of increasing and decreasing marginal impatience, which we study in Sections 5 and 6 respectively, are preserved as we vary λ. This turns out to be especially important since the choice of assumption matters critically for first-best outcomes as well as for equilibrium behavior.[9]

4 A Folk Theorem

Subgame perfections requires that the threat of future punishments be credible. Following Fudenberg and Maskin [12], this can be done by finding strategies that punish a player who deviates, while simultaneously rewarding the players who are supposed to carry out the punishment. Abreu et al. [1] show that such asymmetric treatment is possible under a general condition they call **non-equivalent utilities (NEU)**. It requires that no two players have identical preferences in the stage game, i.e, that for every i, j, there are $\alpha, \alpha' \in \Delta(A)$ such that $g_i(\alpha) > g_i(\alpha')$, but $g_j(\alpha) \leq g_j(\alpha')$.[10] A subtle issue arises in the case of endogenous discounting in that the stage payoffs g_i do not have a well-defined ordinal meaning in terms of the repeated game. To see why, recall that in consumer choice theory one typically speaks of the utility of a *bundle* and that, unless utility is additively separable, it is hard to speak of the utility of a *single good*. Thinking of a play path as a bundle of stage outcomes, we see that an analogous problem arises whenever intertemporal utility is not

[9]It is not difficult to see that (3) and (4) are in fact equivalent restrictions on a family of discount factors. Note as well that our convergence path is sufficient, but not necessary, for preserving the ranking of constant strategies and the assumptions of IMI and DMI. In this respect, Theorem 3 in Sugaya [26] suggests that a more general convergence path with the latter properties would not affect the limit set of equilibrium payoffs as long as (4) holds in the limit as $\beta_i \nearrow 1$. However, to formally extend Sugaya's result to the case of endogenous discounting, one needs to normalize payoffs so they do not blow up as $\beta_i \nearrow 1$. We found such a normalization for the convergence path we selected, but not more generally.

[10]NEU generalizes the stronger assumption of Fudenberg and Maskin [12] that the set $\{g(\alpha) : \alpha \in \Delta(A)\} \subset \mathbb{R}^n$ have full dimension.

time additive, which is precisely how Uzawa-Epstein preferences generalize the standard model. The next lemma, due to Epstein [9], makes this clear.

Lemma 4.1. *Two pairs $(g_i, \beta_i), (g_i', \beta_i')$ induce the same preference relation on $\Delta(A^\infty)$ if and only if $\beta_i' = \beta_i$ and there are constants $\gamma > 0$ and θ such that $g_i' = \gamma g_i + \theta(1-\beta_i)$. In particular, if $\theta \neq 0$, the functions g_i, g_i' need not be monotone transformations of one another.*

We deal with the preceding problem by stating all results and assumptions that concern stage payoffs in terms of the players' preferences over constant play paths. This difference is immaterial in the standard model in which $g_i(\alpha) \geq g_i(\hat{\alpha})$ if and only if $v_i(\alpha) \geq v_i(\hat{\alpha})$. In sum, we can reformulate the NEU condition of Abreu et al. [1] as follows.

Definition 4.1. *A repeated game $(A, (g_i, \beta_i)_i)$ **satisfies NEU** if no two players have identical preferences on the space Σ^{con} of constant strategies.*

The next lemma, due to Chew [7], characterizes NEU in terms of the utility representations (g_i, β_i). The lemma also shows that the condition is generic.[11]

Lemma 4.2. *(g_i, β_i) and (g_j, β_j) induce the same preference relation on the space Σ^{con} of constant strategies if and only if there are constants r, q, s, t such that $qt > rs$ and $g_j = qg_i + r(1-\beta_i)$ and $\beta_j = 1 - sg_i - t(1-\beta_i)$.*

For a path $\boldsymbol{\alpha} = (\alpha^0, \alpha^1, ...) \in (\Delta A)^\infty$ to arise in a SPE, it is necessary that at each point in time the continuation utility of every player exceeds their security level, i.e., that the path be SIR. As is typical, our folk theorem shows that SIR is also sufficient provided that the players are sufficiently patient. To state the result formally, normalize utilities so that the security level of every player is zero[12] and for every $\varepsilon \geq 0$ and λ, let $SIR^\varepsilon(\lambda)$ be the set of all ε**-sequentially individually rational** paths $\boldsymbol{\alpha} \in (\Delta A)^\infty$, i.e., all paths such that $v_{i\lambda}({}_t\boldsymbol{\alpha}) \geq \varepsilon$ for all i, t.

Theorem 4.1. *Assume NEU. For every $\varepsilon > 0$, there exists $\underline{\lambda} \in [0, 1)$ such that for all $\lambda \in (\underline{\lambda}, 1)$, every path $\boldsymbol{\alpha} \in SIR^\varepsilon(\lambda)$ can be supported in a SPE of the game $\Gamma(\lambda)$.*

Remark 4.1. *NEU is not required in two-player games, where deviations can be deterred by the threat of mutual minmaxing. The argument is analogous to that in Fudenberg and Maskin [12].*

---

[11] The connection between the present setup and that of Chew [7] is explained in Appendix B.1.

[12] To do so, let $\hat{g}_i := g_i - \underline{v}_i(1-\beta_i)$. By Lemma 4.1, the pair $(\hat{g}_i, \beta_i)$ induces the same preference on $\Delta(A^\infty)$ as (g_i, β_i). By construction, each game $\hat{\Gamma}(\lambda) := (A, ((1-\lambda)\hat{g}_i, \beta_{i\lambda})_i)$ is strategically equivalent to $\Gamma(\lambda) = (A, ((1-\lambda)g_i, \beta_{i\lambda})_i)$ and all security levels are zero.

Despite the genericity of NEU, the next example presents an important case that is excluded by the condition.

Example 4.1. *Suppose discounting is exogenous and for every i, j, $g_i = g_j$, while $\beta_i \neq \beta_j$. If so, all players have identical preferences on the space Σ^{con} of constant strategies and NEU fails. Yet, no two players have identical preferences on the space Σ of all strategies.*[13] *The latter implication, termed* ***Dynamic NEU*** *by Chen and Takahashi [6], is the most general NEU-type condition one can formulate as it does not restrict in any way the set of strategies that can be ranked differently. Remarkably, when discounting is exogenous, Chen and Takahashi [6] are able to prove a folk theorem under this condition, but we could not extend their theorem to the case of endogenous discounting.*

It is interesting to note that while our NEU condition mimics that of Abreu et al. [1], the proof of our folk theorem shares some important aspects with that of Chen and Takahashi [6]. As usual, the main challenge is finding strategies that punish a deviation, while giving the punishers an incentive to punish. Abreu et al. [1] construct such strategies within the set Σ^{con} of constant strategies by exploiting three ingredients of their setup: (i) the availability of public randomization, which enriches the set Σ^{con}, (ii) the fact that by NEU the players rank this set differently, and (iii) the linearity of the mappings $\alpha \mapsto v_i(\alpha)$, which brings a lot of tractability to the problem. As can be seen from (2), however, that linearity is lost when discounting is endogenous. This makes working with constant strategies much more difficult.[14] To compensate for this, our proof leverages (i) and (ii) together with some useful *intertemporal* properties of the Uzawa-Epstein preferences we study. These properties, first noted by Epstein [9], pertain to behavioral definitions of impatience that can only be fleshed out by the use of non-constant strategies. In this aspect, our proof resembles that of Chen and Takahashi [6], who exploit intertemporal tradeoffs as well. The exact constructions are different however. Ours is designed to deal with the endogeneity of discounting, while maintaining (ii). Chen and Takahashi [6] exploit the time separability of the standard preference specification, while relaxing (ii) completely.

[13] Note that if discounting is endogenous, it is possible for NEU to hold while $g_i = g_j$ for every i, j. In those cases, however, one can invoke Lemma 4.1 and re-normalize utility so that these equalities break down.

[14] Because of the nonlinearity of the functions $\alpha \mapsto v_i(\alpha)$, the set of payoffs attainable by constant strategies need not be convex, even though we assume public randomization. On the other hand, this nonlinearity generates some interesting predictions concerning intertemporal trade. See Sections 5 and 6.

5 Increasing Marginal Impatience

Say that a game Γ satisfies **increasing marginal impatience (IMI)** if for all $i \in I$ and all $a, a' \in A$,

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

In this section, we investigate whether this assumption, advocated by Friedman [11] and made popular by Lucas and Stokey [17], leads to gains from intertemporal trade. Intuitively, an intertemporal trade happens when an impatient player "borrows" from a more patient counterpart and then "repays the debt" in a later period. More generally, such trade requires that distinct actions be taken in at least two periods. With this in mind, let V^{pc} be the convex hull of $\{v(a) : a \in A\}$ and note that in symmetric games with exogenous discounting, V^{pc} is equal to the set of all feasible payoffs, that is, there are no gains from intertemporal trade. In contrast, we are going to show that, under IMI, such gains exist in games with endogenous discounting. Remarkably, this is true even in symmetric games, as differences in discounting can emerge endogenously, in the course of the game. Using our folk theorem, we would also show that intertemporal trade can be sustained in a SPE.

To state our first result, let $V(\lambda)$ be the set of all feasible payoffs in $\Gamma(\lambda)$ and, once again, assume that utilities are normalized so that all security levels are zero, a normalization we maintain throughout the rest of the paper. Given sets $A, B \subset \mathbb{R}^m$, write $B < A$ if $B \subset A \not\subset \{v \in \mathbb{R}^m : \exists v' \in B \text{ s.t. } v' \geq v\}$, that is, if A is a superset of B and A contains at least one element that is not dominated by an element of B. Also, write $v > B$ if there is no $v' \in B$ such that $v' \geq v$, but $v \geq v'$ for some $v' \in B$. We want to show that $V^{pc} < V(\lambda)$. Clearly, a necessary condition is that no action profile $a \in A$ simultaneously maximizes the utility of every player, i.e., there is no $a \in A$ such that $v_i(a) = v_i^{max}$ for every i.[15] Our next result shows that this condition, which we call **conflict of interest (CI)**, is also sufficient.

Theorem 5.1. *Γ satisfies CI if and only if $V^{pc} < V(\lambda)$ for all λ. If CI and NEU hold and V^{pc} contains some payoff $v \gg 0$, then there is a feasible payoff $\hat{v}$ such that $\hat{v} > V^{pc}$ and for all sufficiently high λ, $\hat{v}$ can arise in a SPE.*

The proof of Theorem 5.1 exploits an important behavioral implication of IMI known as **correlation aversion**. Details are provided in the next lemma due to Epstein [9]. To

[15] CI is equivalent to the requirement that no single play path, constant or not, maximizes the utility of every player. This is because a non-constant path can attain v_i^{max}, for some i, if and only if $(a_t, a_t, ...)$ attains v_i^{max} for every t.

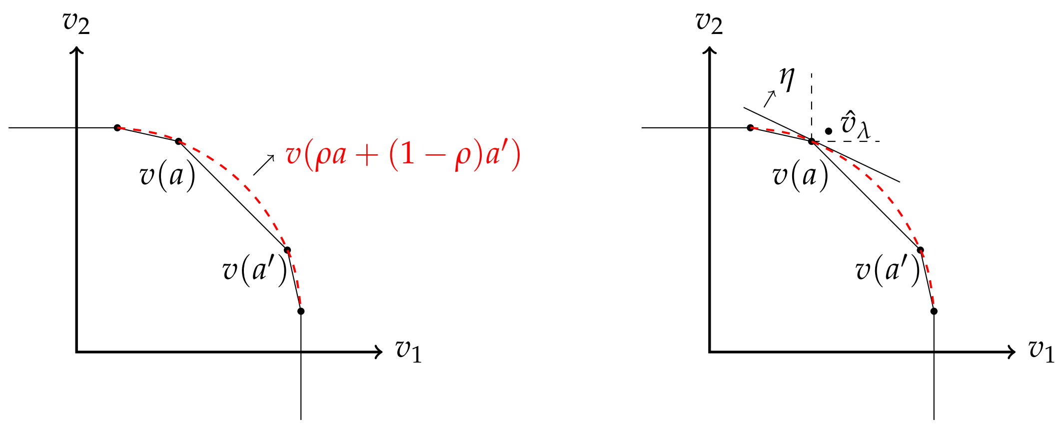

(a) Theorem 5.1 (b) Theorem 5.2

Figure 2: In each figure the solid curve is the Pareto frontier of the set V^{pc}, while the dashed curve shows the Pareto improvements obtained in Theorem 5.1. On the right, we show that mixing the Pareto improvements of Theorem 5.1 need not give a Pareto improvement of the point $v(a)$, which is extreme in V^{pc}. A different construction is thus needed to obtain the Pareto improvement $\hat{v}_\lambda$.

state it, let $\sum_{k=1}^{m} \varrho_k a^k$ be the mixed action profile $\alpha \in \Delta(A)$ such that $a^k \in A$ is played with probability $\varrho_k > 0$.

Lemma 5.1. *IMI holds if and only if for every mixed action $\sum_{k=1}^{m} \varrho_k a^k$ and every i, we have $\sum_{k=1}^{m} \varrho_k v_i(a^k) \leq v_i(\sum_{k=1}^{m} \varrho_k a^k)$. In addition, the inequality is strict if and only if $v_i(a^k) \neq v_i(a^l)$ for some $k, l \leq m$.*

Proof. For $m = 2$, the posited inequality can be rewritten as $(v_i(a^1) - v_i(a^2))(\beta_i(a^1) - \beta_i(a^2)) \leq 0$, which is equivalent to IMI. That IMI implies the desired inequality for $m > 2$ follows from a simple inductive argument. □

To understand the lemma, let $v(\alpha) = (v_1(\alpha), ..., v_n(\alpha))$ and note that a payoff $\varrho v(a) + (1 - \varrho)v(a')$ arises if the players randomize once and, depending on the outcome, play a forever after or a' forever after. Thus, stage play is perfectly positively correlated over time. On the other hand, a payoff $v(\varrho a + (1 - \varrho)a')$ arises if the mixed action $\varrho a + (1 - \varrho)a'$ is played after every history, leading to stage play that is IID over time. It is intuitive that a player may prefer the latter scenario. Positive autocorrelation compounds the uncertainty faced in the initial round of play so that, in particular, a bad outcome in that round is repeated forever after.[16]

Given Lemma 5.1, the proof of Theorem 5.1 is completed by showing that, under CI, the Pareto frontier of V^{pc} has a face F that is perpendicular to some strictly positive direction $\eta \gg 0$ and is not a singleton. This means that moving along the face F benefits some

[16]Experimental evidence supports this intuition. See Miao and Zhong [18] and Andersen et al. [3].

player i, but hurts another player j. We then show that any point v in the relative interior of F is Pareto dominated by a payoff $v(\alpha)$, where $\alpha \in \Delta(A)$ randomizes among actions $a \in A$ whose payoffs $v(a)$ are extreme points of F. Figure 2 provides an illustration.

Example 5.1 (Constant Mixed Strategies and Intertemporal Trade). *At this stage, it is important to clarify the link between the play of a constant mixed strategy α^{con} and intertemporal trade. Consider Figure 2 and let α be a mixed action that randomizes between a and a'. Since α^{con} induces an IID distribution in $\Delta(A^\infty)$, the realized pure play path $(a^0, a^1, ...)$ would be non-constant almost surely. Moreover, the players feel differently about a and a' and, hence, experience different levels of patience along the path. Gains from intertemporal trade are thus factored into the utilities $v_i(\alpha)$. In fact, if $\alpha = \frac{1}{2}a + \frac{1}{2}a'$, it is easy to see that*

$$v_i(\alpha) = \lim_{\lambda \nearrow 1} v_{i\lambda}(a, a', a, a', ...).$$

Thus, as the players become increasingly patient, there is no difference between playing the mixed action α in each period and alternating between a and a' deterministically. To be even more specific, recall Figure 1(a) which depicted a situation in which $(CD, DC, CD, DC, ...)$ is a first-best outcome of the repeated prisoner's dilemma and, hence, specifies an optimal pattern of intertemporal trade. By the law of large numbers, the constant strategy α^{con}, where $\alpha = \frac{1}{2}CD + \frac{1}{2}DC$, induces almost surely a pure play path in which the relative frequency of both CD and DC is $\frac{1}{2}$. However, the realized pattern need not exactly match the repeated alternation between CD and DC that is optimal. Because of this, $v(\alpha)$ would be strictly inside the Pareto frontier, even though it factors in some gains from intertemporal trade. But as the players become infinitely patient, the exact pattern becomes immaterial and only the relative frequencies matter. In particular, $v(\alpha)$ becomes efficient.

Theorem 5.1 does not tells us whether the extreme points on the Pareto frontier of V^{pc} can be dominated. We observed in Section 2 that each player's utility is maximized by some constant path $(a, a, ...)$. Thus, a payoff v such that $v_i = v_i^{max}$ for some i is an extreme point that cannot be dominated. The next section shows that an extreme point v such that $v_i = v_j$ for all $i, j \in I$ need not be dominated either. On the other hand, our next result shows that, generically, all other extreme points are dominated. In addition, if these extreme points are strictly individually rational, then the Pareto improvements can be sustained in a SPE.

Theorem 5.2. *Consider a symmetric game Γ satisfying CI. Let $a \in A$ be such that $v_j^{max} > v_j(a)$ for all $j \in I$ and for some $i \in I$, $v_i(a) > v_j(a)$ for all $j \neq i$. Then, for every λ, there is $\hat{v}_\lambda \in V(\lambda)$ such that $\hat{v}_\lambda \gg v(a)$. If, in addition, NEU holds and $v(a) \gg 0$, then for all sufficiently high λ, the payoff $\hat{v}_\lambda$ can arise in a SPE.*

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

Figure 3: The prisoner's dilemma

Figure 2(b) shows that the payoff $\hat{v}_\lambda$ found in Theorem 5.2 need not be a convex combination of the Pareto improvements found in Theorem 5.1. The construction of $\hat{v}_\lambda$ is thus more intricate, employing strategies that are not constant.[17] However, we do not know whether the use of such strategies is necessary. More importantly, we have not yet confirmed whether there are feasible payoffs outside the convex hull of $\{v(\alpha) : \alpha \in \Delta(A)\}$. Letting V^c denote this convex hull, our next result shows that the answer to the latter question is yes.

Theorem 5.3. *If Γ is symmetric and satisfies CI, then $V^c < V(\lambda)$ for every λ.*

A final question is whether some payoffs $v \notin V^c$ can arise in a SPE. Though we have not been able to answer this question at the present level of generality, there is a sufficient condition that is arguably easy to check. Consider $\alpha \in \Delta(A)$ such that (i) $v(\alpha)$ is on the Pareto frontier of V^c, (ii) for some i, we have $v_i(\alpha) > v_k(\alpha)$ and $\beta_i(\alpha) < \beta_k(\alpha)$ for all $k \neq i$, and (iii) $v_j^{max} > v_j(\alpha) > 0$ for every j. If such α exists, Lemma E23 in the appendix shows that there are Pareto improvements $\hat{v}_\lambda \gg v(\alpha)$ that can arise in a SPE for all λ sufficiently large.

In specific games, it may also be possible to compute directly the play paths that comprise the Pareto frontier of $V(\lambda)$. Section 5.1 does so for the prisoner's dilemma and shows that first-best paths can be sustained in a SPE.

5.1 The Prisoner's Dilemma

Let the action space A and the stage payoffs $g_1, g_2 : A \to \mathbb{R}$ be as in Figure 3 where, as usual, C stands for "cooperate" and D for "defect." To define discounting, let $\beta : \{b, 0, c, d\} \to (0, 1)$ be a function associating each possible stage payoff with a level of patience and for each i, let $\beta_i := \beta \circ g_i$.[18] As is typical in a prisoner's dilemma, we

[17] In addition, our proof of Theorem 5.2 (but not any other results) requires that the public randomization device be able to recommend distinct mixed, as opposed to pure, actions depending on the state of nature. We discuss this point in Appendix E.

[18] Defining discount factors as a function of stage payoffs rather than action profiles is w.l.o.g. in the context of the prisoner's dilemma. Indeed, under IMI, (5) implies that $d > c > 0 > b$. Under DMI, it is possible that $d = c$ but $\beta_1(DC) > \beta_1(CC)$, in which case discount factors cannot be expressed as a function of stage payoffs. It is easy to check however that the proof of Theorem 6.3 goes through in that case as well.

assume that

$$\frac{d}{1-\beta(d)} > \frac{c}{1-\beta(c)} > 0 > \frac{b}{1-\beta(b)}. \tag{5}$$

Consistent with the discussion of Lemma 4.1, note that these inequalities are ordinal restrictions on preferences in the repeated game. 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

$$\frac{c}{1-\beta(c)} > \frac{1}{2}\frac{b}{1-\beta(b)} + \frac{1}{2}\frac{d}{1-\beta(d)}. \tag{6}$$

This means that each player prefers cooperation in every period to receiving his worst or best play path with equal probability. The assumption helps us highlight the different predictions brought about by endogenous discounting. Specifically, if (6) holds within the standard model, then $(CC, CC, ...)$ is the unique play path that maximizes the sum of the players' utilities while, as we show next, this need not be the case when discounting is endogenous.

5.1.1 First-Best Outcomes

Figure 4 depicts two possibilities for the Pareto frontier in the repeated prisoner's dilemma. Consistent with Theorem 5.1, note that in both cases there are gains from intertemporal trade. On the left, the path $\mathbf{a}^C := (CC, CC, ...)$, which we refer to as one of **intratemporal cooperation**, is efficient. The more surprising case is on the right. There, the sum of the players' utilities is maximized not by $\mathbf{a}^C$, but by the play paths in which the players take turns defecting:

$$\mathbf{a}^{\mathcal{A},1} := (DC, CD, DC, CD, ...)$$
$$\mathbf{a}^{\mathcal{A},2} := (CD, DC, CD, DC, ...).$$

These paths represent an intertemporal compromise in which the players alternate between their most preferred outcomes. We refer to these paths as ones of **intertemporal cooperation**.

Another notable implication of IMI is that, eventually, every efficient path, other than the extremes in which one player defects forever, becomes one of cooperation (intra- or inter- temporal). In this sense, IMI and the intertemporal trade it leads to promote an equitable division of surplus. In Section 5.1.2, we would also show that this has implications for the provision of incentives. First, however, we formalize our characterization of the Pareto frontier.

Intratemporal Cooperation. Let $\mathcal{C}_1$ be the set of paths such that DC is played in at most

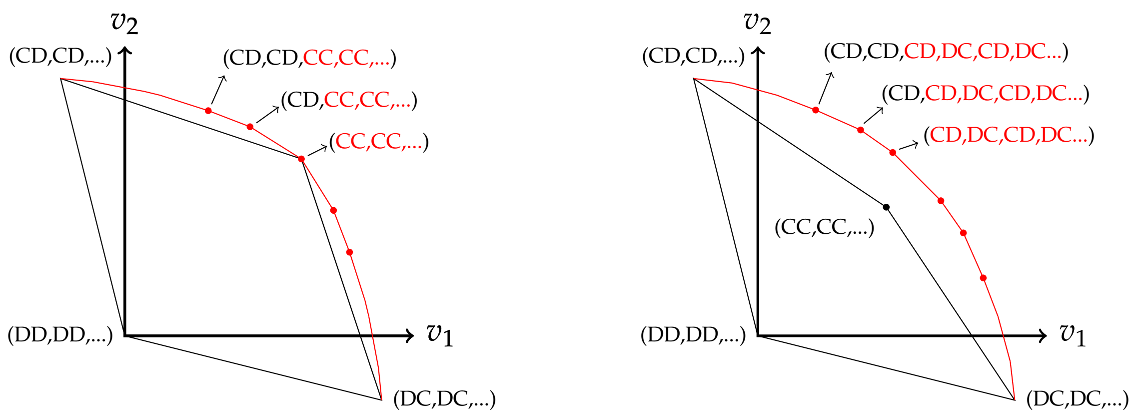

(a) Intratemporal cooperation

(b) Intertemporal cooperation

Figure 4: Two forms of cooperation under IMI

one period while CC is played in all other periods. The subscript "1" is used to designate the fact that the action profile DC, if it occurs, favors player 1. 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, $a^t = DC$ 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. Define the sets $\mathcal{C}_2$ and $\mathcal{E}_2\mathcal{C}_2$ analogously and let $\mathcal{EC} := \mathcal{E}_1\mathcal{C}_1 \cup \mathcal{E}_2\mathcal{C}_2$.[19]

Intertemporal cooperation. Consider the pairs (DC, CD) and (CD, DC) in A^2 and interpret each such pair as a **simple intertemporal trade** in which one of the players gets to defect "today" and, in exchange, lets the other player defect "tomorrow." Let $\mathcal{A}$ be the set of all play paths in which the players make such simple trades in succession. Formally,

$$\mathcal{A} := \{\mathbf{a} \in A^\infty : a^{2t}, a^{2t+1} \in \{DC, CD\} \text{ and } a^{2t} \neq a^{2t+1} \ \ \forall t \in \mathcal{T}\}.$$

One can verify that the sum of utilities, $v_{1\lambda}(\mathbf{a}) + v_{2\lambda}(\mathbf{a})$, is the same for all paths $\mathbf{a} \in \mathcal{A}$, which in the context of Figure 4(b) means that the payoffs from such paths are dispersed along the linear segment of the frontier that is perpendicular to the 45-degree line. Accordingly, we expand the notion of **intertemporal cooperation** to include any path $\mathbf{a} \in \mathcal{A}$, not just the paths $\mathbf{a}^{\mathcal{A},1}$ and $\mathbf{a}^{\mathcal{A},2}$, whose payoffs constitute the extreme points of that segment. It remains to introduce the play paths along which intertemporal cooperation obtains eventually. Thus, let $\mathcal{E}_1\mathcal{A}$ be the set of play paths $\mathbf{a} \in A^\infty$ such that for some $T \geq 0$, depending on the path, $a^t = DC$ for all $t < T$, and ${}_T\mathbf{a} \in \mathcal{A}$. Define $\mathcal{E}_2\mathcal{A}$ analogously and let $\mathcal{EA} := \mathcal{E}_1\mathcal{A} \cup \mathcal{E}_2\mathcal{A}$.

[19] By definition, $\mathcal{C}_1$ contains paths such as $(CC, DC, CC, CC, ...)$ in which player 1 defects in a single period $t > 0$. The payoffs v from such paths are a convex combination of $v(\mathbf{a}^C)$ and $v_\lambda(DC, \mathbf{a}^C)$. Such v are on the Pareto frontier but not an extreme point, which is why they are not shown in Figure 4.

At this point, we should note that the Pareto frontier could take a third form not shown in Figure 4. Namely, for some λ, intra- and inter- temporal cooperation could be simultaneously efficient. The analysis of this case is notationally cumbersome and delivers few additional insights. As, moreover, the case does not arise for any λ sufficiently high, we defer its analysis to Appendix L. To state our current result, say that a path $\boldsymbol{\alpha} \in (\Delta A)^\infty$ is **efficient** if there is no strategy $\sigma \in \Sigma$ that gives each player strictly higher utility and let $P(\lambda)$ be the set of all efficient pure play paths in $\Gamma(\lambda)$. Also, a level of patience λ is **irregular** if intra- and inter- temporal cooperation are both efficient, that is, if $\mathbf{a}^C, \mathbf{a}^{\mathcal{A},1}, \mathbf{a}^{\mathcal{A},2} \in P(\lambda)$. Else, λ is **regular**. Finally, let $\mathbf{a}^{max,i} \in A^\infty$ be a play path which gives i his maximum payoff. In the context of the prisoner's dilemma, this path is unique. For instance, $\mathbf{a}^{max,1} = (DC, DC, ...)$.

Theorem 5.4. *For every regular $\lambda \in [0,1)$, the set $P(\lambda)$ of efficient play paths is equal to either $\mathcal{EC} \cup \{\mathbf{a}^{max,1}, \mathbf{a}^{max,2}\}$ or $\mathcal{EA} \cup \{\mathbf{a}^{max,1}, \mathbf{a}^{max,2}\}$.*

To conclude our characterization of the Pareto frontier, we note that Figure 4(b) could prevail even in the limit as $\lambda \nearrow 1$. That is, intratemporal cooperation could remain inefficient even as absolute differences in discounting disappear. To see this, let $\alpha^* = \frac{1}{2}DC + \frac{1}{2}CD$. A simple calculation shows that the sum of the players' utilities from any alternating path $\mathbf{a} \in \mathcal{A}$ decreases monotonically to $2v_1(\alpha^*)$ as $\lambda \nearrow 1$. If $v_1(\alpha^*) > v_1(\mathbf{a}^C) = v_2(\mathbf{a}^C)$, it follows that the paths $\mathbf{a} \in \mathcal{A}$ (intertemporal cooperation) are efficient for all λ and Pareto dominate $\mathbf{a}^C$ (intratemporal cooperation) for all λ high enough, as well as, in the limit.

5.1.2 Equilibrium Behavior

The folk theorem in Section 4 shows that if the players are sufficiently patients, a play path can arise in a SPE if the security levels of each player are cleared at every stage. Theorem 5.4 suggests that this condition can be weakened in the case of efficient paths. Namely, since all efficient play paths, other than $\mathbf{a}^{max,i}$, converge to cooperation as time progresses, it is enough to check whether each player's security level is cleared at the beginning of the game, i.e., whether the path is **individually rational (IR)**. The only caveat in this line of reasoning arises in the case of intratemporal cooperation when we consider efficient paths such as

$$(CC, CC, DC, CC, CC...). \tag{7}$$

Since a defection occurs in some period $t > 0$, such paths could be IR but not SIR. As is clear, however, a modicum of patience is enough to rule out this possibility. To formalize these observations, let $IR^\varepsilon(\lambda)$ be the set of all **ε-individually rational** paths $\boldsymbol{a} \in A^\infty$, i.e., all pure paths $\boldsymbol{a} \in A^\infty$ such that $v_{i\lambda}(\boldsymbol{a}) \geq \varepsilon$.

Corollary 5.1. *Let $\underline{\lambda}$ be the smallest $\lambda \in [0,1)$ such that the path $(CD, CC, CC, ...)$ is individually rational. Then, for every $\lambda \in (\underline{\lambda}, 1)$ and every $\varepsilon > 0$ small enough, if $\mathbf{a} \in P(\lambda) \cap IR^\varepsilon(\lambda)$, then $\mathbf{a} \in SIR^\varepsilon(\lambda)$.*

A higher but readily computable threshold is needed to insure that, in addition to being SIR, the path in (7) is attainable in a SPE. Namely, let λ' be such that

$$(1-\lambda')d = (1-\lambda')b + (\lambda' + (1-\lambda')\beta(b))\frac{c}{1-\beta(c)} \tag{8}$$

and let $\underline{\lambda}' = \max\{0, \lambda'\}$. Then,

Corollary 5.2. *For every $\lambda \in (\underline{\lambda}', 1)$ and every $\varepsilon > (1-\lambda)d$, every path $\mathbf{a} \in P(\lambda) \cap IR^\varepsilon(\lambda)$ can be supported in a SPE of the game.*

Consistent with the preceding discussion, we note that if intertemporal cooperation is efficient for all λ, then the thresholds $\underline{\lambda}$ and $\underline{\lambda}'$ in the above corollaries can be set to 0.

6 Decreasing Marginal Impatience

This section studies the implications of decreasing marginal impatience in symmetric two-player games. Our first result shows that efficient paths takes one of two forms: either one player's continuation utility is eventually maximized or both players attain identical continuation utilities along the entire path. In many games, maximizing the utility of one player implies that the other player is pushed below his security level or "immiserated." The result may therefore be viewed as a formalization of Friedman's [11] argument that DMI leads to "disequilibrium behavior." The important twist is that because discounting is endogenous, immiseration need not occur if the players coordinate on a symmetric path.

Theorem 6.1. *For every λ, every efficient path $\mathbf{a} \in P(\lambda)$ is such that either (i) $v_{i\lambda}({}_T\mathbf{a}) = v_i^{max}$ for some $i \in I$ and $T \geq 0$, or (ii) $v_{1\lambda}({}_t\mathbf{a}) = v_{2\lambda}({}_t\mathbf{a})$ for every $t \geq 0$.*

Remark 6.1. *If $v_{i\lambda}({}_T\mathbf{a}) = v_i^{max}$ for some $T \geq 0$, then $v_{i\lambda}({}_t\mathbf{a}) = v_i^{max}$ for all $t \geq T$. Also, under DMI, $v_{1\lambda}({}_t\mathbf{a}) = v_{2\lambda}({}_t\mathbf{a})$ for every $t \geq 0$ if and only if $g_1(a^t) = g_2(a^t)$ for all t. That is, a player's continuation utility is maximized only if all subsequent continuation utilities are maximized as well. In addition, the players continuation utilities are equal throughout the game if and only if the stage payoffs are equal as well.*

Intuition for Theorem 6.1 was given in the introduction. If at some point in time player i's continuation utility is higher than j's, i would exhibit a greater level of patience. Efficiency then requires that j's utility be frontloaded while i 's utility be backloaded. The latter implies that i's continuation utility remains higher as the game progresses and, given DMI, that i sustains the higher level of patience. This self-enforcing dynamic continues until player i's utility cannot be backloaded any further, which is when i's utility is maximized. The only alternative is when the players coordinate on a path along which they attain identical continuation utilities and, hence, identical levels of patience throughout the game.

The conclusions of Theorem 6.1 become especially stark when we consider the prisoner's dilemma. Then, maximizing the utility of player i means that player j is pushed below his security level, leaving $(CC, CC, ...)$ as the only efficient path that can arise in a SPE.

Corollary 6.1. *The path $(CC, CC, ...)$ is the only efficient path that can arise in a SPE of the prisoner's dilemma.*

Theorem 6.1 does not tell us whether gains from intertemporal trade exist under DMI. As in Section 5, a necessary condition is that the game feature some conflict of interest. Remarkably, now we must also require that the game not be the prisoner's dilemma. Formally, say that Γ **has conflict of interest but is not prisoner's dilemma (CINPD)** if the Pareto frontier of V^{pc} has an extreme point v such that $v_i^{max} > v_i > v_j$ for some $i \in \{1, 2\}$ and $j \neq i$.[20] Intuitively, $v_i > v_j$ generates differences in discounting by making player i more patient, while $v_i^{max} > v_i$ insures that there is room for i's utility to grow via intertemporal trade. By comparison, in the prisoner's dilemma, the only points $v(a) \in V^{pc}$ at which differences in discounting obtain are $v(CD)$ and $v(DC)$. But at these points, the utility of the more patient player is fully maximized, leaving no room for intertemporal trade to kick in.

Theorem 6.2. *If Γ satisfies CINPD, then $V^{pc} < V(\lambda)$ for each λ. If CINPD holds and the extreme point v in its statement is strictly individually rational, then for all λ sufficiently large, there is a feasible payoff $\hat{v}_\lambda \in V(\lambda)$ such that $\hat{v}_\lambda > V^{pc}$ and $\hat{v}_\lambda$ can arise in a SPE of $\Gamma(\lambda)$.*

Proof. By CINPD and the symmetry of the game, the Pareto frontier of V^{pc} has an extreme point v such that $v_2^{max} > v_2 > v_1$. Let v^* be the extreme point adjacent to v and such that $v_2^* > v_2$, and let $a, a^* \in A$ be such that $v(a) = v$ and $v(a^*) = v^*$. See Figure 5(a) for an illustration. Also, let $((a)^T, a^*, a, a, ...)$ be the path in which a^* is played in period T and a

[20] Note that, under symmetry, CINPD implies CI.

in all other periods. We claim that for any λ, there is T large enough such that the path generates gains from intertemporal trade. Furthermore, there is λ high enough such that for all T, the path is SIR if and only if $v \gg 0$. The intuition is as follows. Since $v_2 > v_1$, player 2 attains a higher level of patience at the start of the path when $a \in A$ is played. Efficiency then requires that 2's utility be backloaded, which is achieved by playing a^* in period T. In fact, to obtain a first-best outcome, the logic behind Theorem 6.1 tells us that 2's utility should continue to rise until it is fully maximized. Since, however, this may not be SIR for player 1, the path $((a)^T, a^*, a, a, ...)$ sacrifices some efficiency by requiring that play of $a \in A$ be resumed after a single round of $a^* \in A$.

Formally, for any λ, i, and T, let $\varrho_i = [\beta_{i\lambda}(a)]^T(1 - \beta_{i\lambda}(a^*))$. By construction, the Pareto frontier of V^{pc} has a linear segment connecting the extreme points v and v^*, which we chose to be adjacent. On that linear segment is a point $(1 - \varrho_1)v + \varrho_1 v^*$. On the other hand,

$$v_\lambda((a)^T, a^*, a, a, ...) = ((1 - \varrho_1)v_1 + \varrho_1 v_1^*, (1 - \varrho_2)v_2 + \varrho_2 v_2^*).$$

As illustrated in Figure 5(a), to show that $v_\lambda((a)^T, a^*, a, a, ...) > V^{pc}$, it is therefore enough to show that

$$(1 - \varrho_2)v_2 + \varrho_2 v_2^* > (1 - \varrho_1)v_2 + \varrho_1 v_2^*.$$

Since $v_2^* > v_2$, the above is equivalent to $\varrho_2 > \varrho_1$. But $\varrho_2 > \varrho_1$ if and only if

$$\Big[\frac{\beta_{2\lambda}(a)}{\beta_{1\lambda}(a)}\Big]^T > \frac{1 - \beta_{1\lambda}(a^*)}{1 - \beta_{2\lambda}(a^*)}.$$

Given the symmetry of the game and DMI, it is easy to see that $v_2(a) > v_1(a)$ implies $\beta_{2\lambda}(a) > \beta_{1\lambda}(a)$. Hence, $\varrho_2 > \varrho_1$ for all T large enough. To prove the final assertion of the theorem, assume that $v(a) \gg 0$ and pick ε such that $v(a) \gg \varepsilon > 0$. For sufficiently high λ, we have $(a^*, a, a, ...) \in IR^\varepsilon(\lambda)$, which implies that $((a)^T, a^*, a, a, ...) \in SIR^\varepsilon(\lambda)$ for all T. Thus, for λ and T sufficiently high, the path is both a SPE and delivers gains from intertemporal trade. □

We conclude the paper by showing that in the prisoner's dilemma, first-best outcomes, other than $(CC, CC, ...)$, cannot be attained in a SPE even asymptotically as $\lambda \nearrow 1$. An important question for future work is whether a similar result applies in games *with* gains from intertemporal trade. For now, we note that a key step in the present analysis, see the upcoming discussion concerning the Pareto frontier of $\{v(\alpha) : \alpha \in \Delta(A)\}$, holds more generally and may be helpful in such a pursuit. To elaborate, let $V^0(\lambda)$ be the space of SIR payoffs given λ and let $V^*(\lambda)$ be the space of SPE payoffs.[21] Theorem 6.3 shows that the

[21] A SIR payoff is one attainable by a SIR strategy. A SPE payoff is similarly defined.

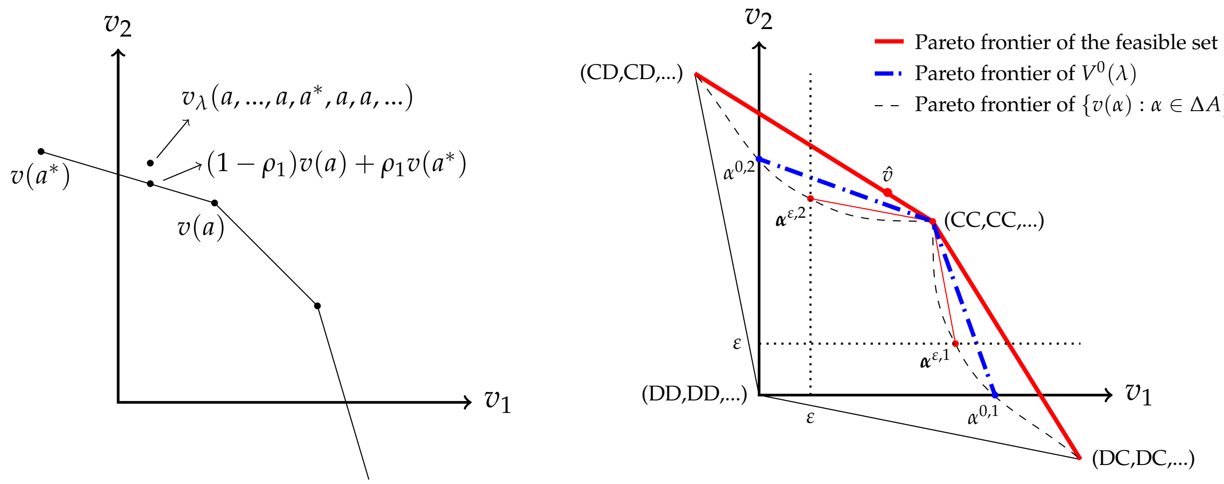

(a) Gains from trade under DMI

(b) The Prisoner's Dilemma

Figure 5: In Figure 5(b), the Pareto frontier of $V^0(\lambda)$, given by the dash-dotted line, is independent of λ and, hence, an upper bound on the set of SPE payoffs attainable at any level of patience λ.

Pareto frontier of $V^0(\lambda)$ is independent of λ and, except for the point $v(CC)$, strictly inside the set of feasible payoffs. This frontier, which we depict as the dash-dotted (blue) line in Figure 5(b), is therefore an upper bound for the set $\cup_\lambda V^*(\lambda)$ of SPE payoffs attainable at any level of patience.

Interestingly, this bound arises for reasons that are conceptually distinct from the immiseration dynamics driving Theorem 6.1. Indeed, recall that a key implication of endogenous discounting is the nonlinearity of the mappings $\alpha \mapsto v_i(\alpha)$. As we showed in Figure 2(a), under IMI, this nonlinearity implies that the Pareto frontier of $\{v(\alpha) : \alpha \in \Delta(A)\}$ bends *outward*, delivering gains from intertemporal trade. Under DMI, the frontier bends *inward*, as shown by the curved dashed line in Figure 5(b). It follows that a point such as $\hat{v}$ in Figure 5(b) cannot be attained by any constant strategy α^{con}. Instead, it can only be attained by a one-time randomization between the paths $(CC, CC, ...)$ and $(CD, CD, ...)$. But since the latter path is not SIR, $\hat{v}$ cannot arise in a SPE no matter the level of patience.

Behaviorally, the inward bend traces back to the discussion of autocorrelations in Section 5. Namely, under DMI, the players prefer perfect positive autocorrelation, as in a one-time randomization $\varrho(a, a, ...) + (1-\varrho)(a', a', ...)$ between the constant paths $(a, a, ...)$ and $(a', a', ...)$, to the IID distribution induced by playing $\varrho a + (1-\varrho)a' \in \Delta(A)$ each period. The advantage of positive autocorrelation is that, with probability 1, the realized outcome stream is smooth over time. The disadvantage is greater risk as the initial draw gets propagated, rather than offset, by future draws. For example, under perfect positive autocorrelation, a bad draw in $t = 0$ implies a bad draw forever after. Recalling Lemma

5.1, we can sum things up as follows: under IMI, the players are more concerned with eliminating risk; under DMI, they are more concerned with obtaining a smooth outcome stream.[22]

Finally, we characterize the upper bound in Figure 5(b) and show that it is in fact a least upper bound for the set $\cup_\lambda V^*(\lambda)$ of all SPE payoffs. Given $\varepsilon \geq 0$, let $\alpha^{\varepsilon,2} \in \Delta A$ be the mixed action $\varrho CC + (1-\varrho)CD$ such that $v_1(\alpha^{\varepsilon,2}) = \varepsilon$. Define $\alpha^{\varepsilon,1}$ analogously and let F_i^ε be the line connecting $v(\alpha^{\varepsilon,i})$ with $v(CC)$. These lines are independent of λ since the utilities $v(\alpha^{\varepsilon,i})$ and $v(CC)$ are independent of λ. Our next and final result shows that $F_1^0 \cup F_2^0$ is the Pareto frontier of $V^0(\lambda)$ for each λ and, hence, an upper bound on $\cup_\lambda V^*(\lambda)$. To see that $F_1^0 \cup F_2^0$ is a least upper bound, note that $F_1^\varepsilon \cup F_2^\varepsilon \subset V^*(\lambda)$ for every $\varepsilon > 0$ and λ such that $\varepsilon > (1-\lambda)d$. Thus, the lines F_i^ε, which we depict as the thin solid lines below F_i^0, give us lower bounds for the Pareto frontier of $\cup_\lambda V^*(\lambda)$. But, as $\lambda \nearrow 1$, we can let $\varepsilon \searrow 0$, so that $F_1^\varepsilon \cup F_2^\varepsilon \to F_1^0 \cup F_2^0$.

Theorem 6.3. *For each λ, the Pareto frontier of the set $V^0(\lambda)$ of SIR payoffs is independent of λ and given by $F_1^0 \cup F_2^0$. Except for $v(CC)$, the space of SPE payoffs is thus bounded away from the first best frontier even as $\lambda \nearrow 1$.*

7 Conclusions and Related Literature

The paper showed that, under IMI, gains from intertemporal trade exist in a broad class of games. In addition, such trade does not lead to immiseration and can be sustained in a SPE. Under DMI, gains from intertemporal trade exist as well. However, the first-best pattern of such trade leads to the immiseration of some player or, in Friedman's [11] words, to "disequilibrium behavior." The twist we highlighted is that immiseration can be prevented if the players coordinate on a path along which they attain symmetric outcomes and consequently remain equally patient.

Our conclusions under IMI owe a conceptual debt to the work Lucas and Stokey [17], who used IMI to circumvent the immiseration results of Rader [23] and Becker [4] and obtain a steady state with a non-trivial distribution of capital. Viewing Lehrer and Pauzner [16] as the strategic analogue of Rader [23] and Becker [4],[23] one can say that our main message under IMI relates to Lehrer and Pauzner [16], the way Lucas and Stokey [17] relate to Becker [4] and Rader [23]. Note however that the repeated games we study are quite different from the growth economies of Lucas and Stokey [17]. As such, there is little formal overlap between the papers. To give one example, recall that one of our main

[22] See Figure 1 in Bommier et al. [5] for an illustration of this trade-off.

[23] Recall that, as in Rader [23] and Becker [4], discounting in Lehrer and Pauzner [16] is exogenous.

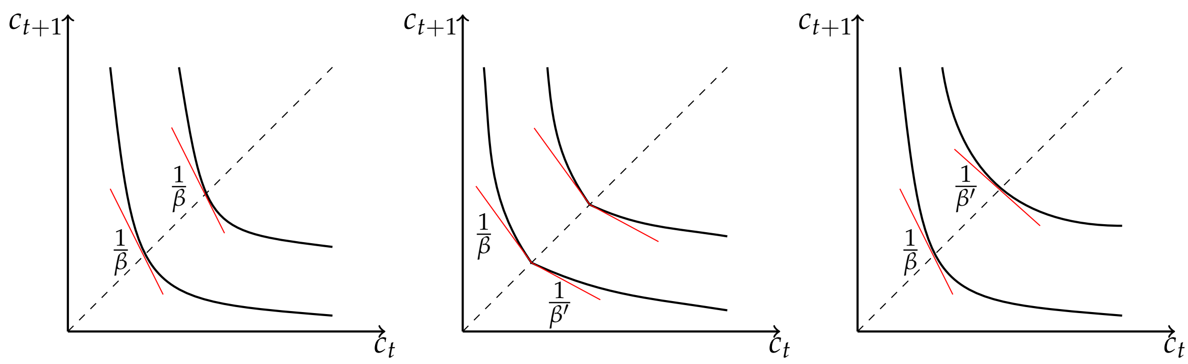

(a) Standard Model (b) Sekiguchi and Wakai [25] (c) DMI

Figure 6: Time preferences in a two-period model with divisible consumption

results, Theorem 5.1, hinges on the player's attitudes toward autocorrelations, that is, on the presence of uncertainty. By comparison, the model of Lucas and Stokey [17] has no uncertainty.

We should also mention two recent papers that study games with non-standard intertemporal preferences. Like us, Sekiguchi and Wakai [25] consider preferences that are recursive and stationary but not time separable. It is known from Koopmans [14] that rates of time preference are endogenous whenever this is the case. As we illustrate in Figure 6 however, the endogeneity can take different forms. In the two-period setup of the figure, discount factors equal the inverse of the marginal rate of substitution at points (c, c) along the 45-degree line. Figure 6(c) depicts the case of DMI in which discount factors increase with the level of consumption c. By comparison, in Sekiguchi and Wakai [25], changes in c do not affect the rate of time preference. Rather, the kink in the indifference curves means that gains in consumption, $(c, c + \varepsilon)$, are discounted differently from losses, $(c, c - \varepsilon)$. Given the different preference specifications and the different setups of the papers — Sekiguchi and Wakai [25] study a Cournot game — there is little overlap, formal or conceptual, between the papers. On the other hand, Obara and Park [21] focus on the present-biased preferences of Laibson [15]. The non-stationarity of these preferences and the associated failures of time consistency lead to very different problems than the ones we consider.

Finally, the paper by Neilson and Winter [20] has a title almost identical to ours, but the setups are completely different. In Neilson and Winter [20], one of the players first chooses the length of time between successive repetitions of the stage game. After that, the players engage in a standard repeated game with no asymmetries related to discounting.

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 topologies. Let $\Delta(A^\infty)$ be the set of Borel probability measures on A^∞, endowed with the weak* topology. Endow each Σ_i and Σ with the respective product topologies. From Kolmogorov's Consistency Theorem, see Parthasarathy [22, Theorem 3.1,3.2], we know that each strategy profile $\sigma \in \Sigma$ induces a probability measure in $\Delta(A^\infty)$ and that 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 A1. *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. Fixing $i \in I$ and $\alpha_{-i} \in \times_{k \neq i} \Delta(A_k)$, take some $\hat{\alpha}_i \in \operatorname{argmax}_{\alpha_i \in \Delta(A_i)} v_i(\alpha_i^{con}, \alpha_{-i}^{con})$, which exists since $\Delta(A_i)$ is compact and the function $\alpha_i \mapsto v(\alpha_i^{con}, \alpha_{-i}^{con})$ 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 can assume that $\tilde{\sigma}_i = \hat{\alpha}_i{}^{con}$, except at a single history h^t. Moreover, since players $j \neq i$ are using a constant strategy, we can 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 $\hat{\alpha}_i \in \operatorname{argmax}_{\alpha_i \in \Delta(A_i)} v_i(\alpha_i^{con}, \alpha_{-i}^{con})$. □

Next, say that a strategy $\sigma_i \in \Sigma_i^*$ 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_i is finite. Let $\underline{v}_i^c$ be i's minmax payoff when all players are restricted to using constant strategies. By Lemma A1, the best reply to a constant strategy is a constant strategy. Hence, $\underline{v}_i \leq \underline{v}_i^c$. 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 the finite strategy that coincides with σ_{-i} up to and including period t. By construction, $\sigma_{-i}^t \to_t \sigma_{-i}$. Using Lemma A1, 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. From the maximum theorem, deduce that

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

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}. 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 i's continuation utilities at h^{t-1} and $h^t(a), a \in A$, 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 4.1

B.1 Payoff Asymmetry

The goal of this section is to prove Lemma B3 below, which is a key step in the proof of our folk theorem and which may be viewed as a generalization of the "payoff-asymmetry" lemma of Abreu et al. [1, Lemma 2]. The proof exploits insights from Chew's [7] work on non-expected utility preferences. To see the connection, note that each pair (g_i, β_i) induces a preference relation $\succeq_i$ on $\Delta(A)$ represented by the utility function $g_i(\alpha)/(1 - \beta_i(\alpha))$. When $\beta_i : A \to (0, 1)$ is constant, $\succeq_i$ is a standard expected utility preference on the simplex $\Delta(A)$. If β_i is not constant, then $\succeq_i$ belongs to the more general class of **weighted-utility preferences** studied in Chew [7]. We begin with a preliminary observation regarding such preferences.

Lemma B2. *Let $\succeq$ be a weighted-utility preference on $\Delta(A)$ and E_1 and E_2 two distinct indifference curves of $\succeq$ intersecting the interior of $\Delta(A)$. Then, $\succeq$ is fully determined by E_1 and E_2 and* *the ranking between them.*

Proof. If $\succeq$ is an expected utility preference, all indifference curves are parallel translations of one another. In addition, the ranking between E_1 and E_2 determines the direction of increasing preference. When $\succeq$ is not expected utility, the proof follows from Figure 1 in Chew [7]. Namely, embed the simplex ΔA into $\mathbb{R}^{|A|-1}$ in the usual way. The indifferences curves E_1 and E_2 are hyperplanes whose intersection is an $(|A|-3)$-dimensional linear subspace L. Rotating the hyperplane E_1 around L generates all indifference curves of $\succeq$. Once again, the ranking between E_1 and E_2 determines the direction of increasing preference. □

Lemma B3. *If the repeated game $(A, (g_i, \beta_i)_i)$ satisfies NEU, then there exist lotteries $\alpha^1, ..., \alpha^n \in \Delta(A)$ such that $v_i(\alpha^j) > v_i(\alpha^i)$ for every $i \neq j$.*

Proof. Say that a vector $(\alpha^1, ..., \alpha^n)$ with the sought after property is a **separation for** $(\succeq_1, ..., \succeq_n)$. Also, let $E_i(\alpha) := \{\alpha' \in \Delta(A) : \alpha' \sim_i \alpha\}$ be player i's indifference set through $\alpha \in \Delta(A)$ and $U_i(\alpha), L_i(\alpha)$ the resp. upper and lower contour sets. If $n = 2$, we claim that one can pick a generic $\alpha \in \Delta(A)$ and α^1, α^2 arbitrarily close to α such that $\alpha^2 \succ_1 \alpha \succ_1 \alpha^1$ and $\alpha^1 \succ_2 \alpha \succ_2 \alpha^2$. Suppose that $\succeq_1$ and $\succeq_2$ have the same indifference sets. Since $\succeq_1 \neq \succeq_2$ by NEU, it follows that $\succeq_1$ is the negation of $\succeq_2$. Hence, for any triple $\alpha^1, \alpha, \alpha^2 \in \Delta(A)$ such that $\alpha^2 \succ_1 \alpha \succ_1 \alpha^1$, we have $\alpha^1 \succ_2 \alpha \succ_2 \alpha^2$. Since the indifference curves of $\succeq_1$ are hyperplanes, the choice of $\alpha^1, \alpha, \alpha^2$ is generic. If $\succeq_1$ and $\succeq_2$ do not have the same indifference sets, then, by Lemma B2, they have in common at most one indifference set E^* intersecting the interior of $\Delta(A)$. Pick any $\alpha \notin E^*$ in the interior of $\Delta(A)$. The hyperplanes $E_1(\alpha)$ and $E_2(\alpha)$ are distinct and partition $\Delta(A)$ into four cones with peak α : $U_1(\alpha) \cap U_2(\alpha), U_1(\alpha) \cap L_2(\alpha), L_1(\alpha) \cap U_2(\alpha), L_1(\alpha) \cap L_2(\alpha)$. Picking any α^2 in the interior of $U_1(\alpha) \cap L_2(\alpha)$ and α^1 in the interior of $L_1(\alpha) \cap U_2(\alpha)$, completes the proof of our claim.

Proceeding by induction, suppose that $(\alpha^1, ..., \alpha^m)$ is a separation for $(\succeq_1, ..., \succeq_m)$ and let $\succeq_{m+1}$ be a distinct weighted-utility preference. Reindexing if necessary, we can suppose that $\alpha^i \succeq_{m+1} \alpha^1$ for all $i < m + 1$. Since $\alpha^2 \succ_1 \alpha^1$ and $\alpha^2 \succeq_{m+1} \alpha^1$, we know that $\succeq_1$ is not the negation of $\succeq_{m+1}$. By perturbing α^1 appropriately, we can assume that $\alpha^i \succ_{m+1} \alpha^1$ for all $i < m + 1$. Since, by Lemma B2, $\succeq_1$ and $\succeq_{m+1}$ have at most one indifference curve in common, we can also assume that $E_1(\alpha^1) \neq E_{m+1}(\alpha^1)$. Using the argument for $n = 2$, we can find lotteries α', α'' such that $\alpha'' \succ_1 \alpha^1 \succ_1 \alpha'$ and $\alpha' \succ_{m+1} \succ \alpha^1 \succ_{m+1} \alpha''$. Moreover, choosing α', α'' sufficiently close to α^1 insures that $(\alpha', \alpha^2, ..., \alpha^m, \alpha'')$ is a separation for $(\succeq_1, \succeq_2, ..., \succeq_{m+1})$. □

B.2 Decision-theoretic preliminaries

Fix i. The next two lemmas, whose simple proofs we omit, capture behavioral implications of impatience that are characteristic of Uzawa-Epstein preferences and that will be useful in the proof of our folk theorem. In particular, let $\alpha^0, \alpha^1, ..., \alpha^K \in \Delta A$ be mixed actions such that $v_i(\alpha^k) \leq v_i(\alpha^{k+1})$ for every $k = 0, ..., K-1$. Lemma B4 shows that player i prefers that more beneficial actions be played first.

Lemma B4. *For every $\boldsymbol{\alpha} \in (\Delta A)^\infty$ and every permutation $\pi : \{0, 1, ..., K\} \to \{0, 1, ..., K\}$, we have $v_i(\alpha^0, \alpha^1, ..., \alpha^K, \boldsymbol{\alpha}) \leq v_i(\alpha^{\pi(0)}, \alpha^{\pi(1)}, ..., \alpha^{\pi(K)}, \boldsymbol{\alpha})$.*

The next lemma says that if the continuation path $\boldsymbol{\alpha}$ is better than each of the actions α^k, it is beneficial to remove some of these actions so as to advance the play of $\boldsymbol{\alpha}$.

Lemma B5. *For every $\boldsymbol{\alpha} \in (\Delta A)^\infty$ such that $v_i(\alpha^K) < v_i(\boldsymbol{\alpha})$ and every subset $\{\hat{\alpha}^0, ..., \hat{\alpha}^{\hat{K}}\} \subset \{\alpha^0, \alpha^1, ..., \alpha^K\}$, we have $v_i(\alpha^0, \alpha^1, ..., \alpha^K, \boldsymbol{\alpha}) \leq v_i(\hat{\alpha}^0, ..., \hat{\alpha}^{\hat{K}}, \boldsymbol{\alpha})$.*

We need one more lemma.

Lemma B6. *If $v_i(\alpha) > v_i(\alpha')$, then $v_i(\alpha) > v_i(\varrho\alpha + (1-\varrho)\alpha') > v_i(\alpha')$ for all $\varrho \in (0, 1)$.*

Proof. This follows from the fact that for all $\rho \in (0, 1), k, l \in \mathbb{R}$, and $s, t \in \mathbb{R}_{++}$, if $ks^{-1} > lt^{-1}$, then $ks^{-1} > (\rho k + (1-\rho)l)(\rho s + (1-\rho)t)^{-1} > lt^{-1}$. □

Finally, we remark that for every path $(\alpha^0, \alpha^1, ...) \in (\Delta A)^\infty$,

$$v_i(\alpha^0, \alpha^1, ...) = \big(1 - \beta_i(\alpha^0)\big)v_i(\alpha^0) + \beta_i(\alpha^0)v_i(\alpha^1, \alpha^2, ...). \tag{9}$$

Thus, $v_i(\alpha^0, \alpha^1, ...)$ is a convex combination of $v_i(\alpha^0)$ and $v_i(\alpha^1, \alpha^2, ...)$.[24]

B.3 Constructing dynamic player-specific punishments

The definition below is adapted from Chen and Takahashi [6].

Definition B1. *For every $\lambda \in [0, 1)$, a play path $\boldsymbol{\alpha} \in (\Delta A)^\infty$ allows **dynamic player-specific punishments (DPSP) with wedge** $\gamma > 0$ if there exists paths $\mathbf{r}^1, ..., \mathbf{r}^n \in (\Delta A)^\infty$ such that for every $i, j \neq i$, and every t, we have (i) $v_{i\lambda}(\mathbf{r}^i) < v_{i\lambda}({}_t\boldsymbol{\alpha}) - \gamma$, (ii) $\gamma < v_{i\lambda}(\mathbf{r}^i) \leq v_{i\lambda}({}_t\mathbf{r}^i)$, and (iii) $v_{i\lambda}(\mathbf{r}^i) < v_{i\lambda}({}_t\mathbf{r}^j) - \gamma$.*

In Section B.4, we use the paths $(\mathbf{r}^i)_i$ in Definition B1 to construct off-path strategies that punish a deviation from $\boldsymbol{\alpha}$, while simultaneously giving the players incentives to

[24]On the other hand, since the "weights" depend on i, the payoff $v(\alpha^0, \alpha^1, ...)$ need not be a convex combination of $v(\alpha^0)$ and $v(\alpha^1, \alpha^2, ...)$.

carry out the punishment. Roughly, condition (i) deters player i from deviating from the target path $\boldsymbol{\alpha}$; condition (ii) insures that the punishment phase is SIR and that no player wants to restart the punishment; finally, condition (iii) provides incentives for player i to carry out a punishment against player j.

The goal of this section is to prove the following lemma.

Lemma B7. *Assume NEU. For every $\varepsilon > 0$, there are $\gamma > 0$ and $\underline{\lambda} \in [0, 1)$ such that for every $\lambda > \underline{\lambda}$, every $\boldsymbol{\alpha} \in SIR^{\varepsilon}(\lambda)$ allows DPSP $\{\mathbf{r}^i_{\lambda}\}_i$ with wedge γ.*

We begin by defining paths $\{\mathbf{r}^i_{\lambda}\}_{i \in I}$ indexed by two free parameters $T_1, T_2 \in \mathbb{N}_{++}$, to be determined later. Fix $\varepsilon > 0$ and $\lambda \in [0, 1)$ such that $SIR^{\varepsilon}(\lambda) \neq \emptyset$, and some $i \in I$. Since the set $SIR^{\varepsilon}(\lambda)$ is compact, we can find a path $\mathbf{w}^i_{\lambda} \in \operatorname{argmin}_{\hat{\boldsymbol{\alpha}} \in SIR^{\varepsilon}(\lambda)} v_{i\lambda}(\hat{\boldsymbol{\alpha}})$. By Lemma B3, there exists $\kappa^1, ..., \kappa^n \in \Delta A$ such that $v_i(\kappa^i) < v_i(\kappa^j)$ for all $j \neq i$. Enumerate the κ's according to i's preferences:

$$v_i(\kappa^{i_0}) \leq v_i(\kappa^{i_1}) \leq ... \leq v_i(\kappa^{i_{n-1}}).$$

By construction, $\kappa^{i_0} = \kappa^i$. For any $\alpha \in \Delta A$ and $T \in \mathbb{N}_{++}$, let $(\alpha)^T \in (\Delta(A))^T$ be the finite sequence such that α is played T times. For every $T_2 \in \mathbb{N}_{++}$, let

$$\boldsymbol{\alpha}^i_{\lambda} := ((\kappa^{i_0})^{T_2}, (\kappa^{i_1})^{T_2}, ..., (\kappa^{i_{n-1}})^{T_2}, \mathbf{w}^i_{\lambda}).$$

Collecting all κ's into a single block $K^i_{\lambda} \in (\Delta A)^{NT_2}$, we can also write $\boldsymbol{\alpha}^i_{\lambda}$ as $(K^i_{\lambda}, \mathbf{w}^i_{\lambda})$. Next, recall from Section 2 that there are action profiles $l^i \in A$ and $h^i \in A$ such that $(l^i, l^i, ...)$ and $(h^i, h^i, ...)$ minimize and resp. maximize $v_{i\lambda}$ among all strategies (and for all λ). Define

$$\mathcal{L}^i_{\lambda} := \{l^j \in A | v_i(l^j) < v_{i\lambda}(\boldsymbol{\alpha}^i_{\lambda}), j \in I\}$$

and let $N^i := |\mathcal{L}^i_{\lambda}|$. Since $v_i(l^i) \leq v_i(\kappa^{i_m})$ for all $m = 0, ..., n-1$, and $v_i(l^i) \leq 0 < \varepsilon \leq v_{i\lambda}(\mathbf{w}^i_{\lambda})$, we have $v_i(l^i) < v_{i\lambda}(\boldsymbol{\alpha}^i_{\lambda})$. Thus, $l^i \in \mathcal{L}^i_{\lambda}$ and $N^i \geq 1$. Enumerate all action profiles in $\mathcal{L}^i_{\lambda}$ according to i's preferences:

$$v_i(l^{i_0}) \leq v_i(l^{i_1}) \leq ... \leq v_i(l^{i_{N^i-1}}). \tag{10}$$

By the definition of l^i, we know that $l^{i_0} = l^i$. For every $T_1 \in \mathbb{N}_{++}$, define the play path

$$\mathbf{r}^i_{\lambda} := ((l^{i_0})^{T_1}, (l^{i_1})^{T_1}, ..., (l^{i_{N^i-1}})^{T_1}, \boldsymbol{\alpha}^i_{\lambda}) \in (\Delta A)^{\infty},$$

Collecting all l's into a block L^i_{λ}, we may also write $\mathbf{r}^i_{\lambda}$ as $(L^i_{\lambda}, \boldsymbol{\alpha}^i_{\lambda})$.

So far, the construction of the paths $\{\mathbf{r}^i_{\lambda}\}_{i \in I}$ did not reference the target path $\boldsymbol{\alpha} \in SIR^{\varepsilon}(\lambda)$. However, since $v_{i\lambda}(\mathbf{w}^i_{\lambda}) \leq v_{i\lambda}({}_t\boldsymbol{\alpha})$ for every t and $\boldsymbol{\alpha} \in SIR^{\varepsilon}(\lambda)$, condition (i) in Definition B1, which is where the target path $\boldsymbol{\alpha}$ appears, would be automatically satisfied if we could show that $v_{i\lambda}(\mathbf{r}^i_{\lambda}) < v_{i\lambda}(\mathbf{w}^i_{\lambda}) - \gamma$ for every i. The rest of the proof calibrates the paths $\{\mathbf{r}^i_{\lambda}\}_{i \in I}$ by choosing T_1 and T_2 appropriately so that the latter condition as well as conditions (ii) and (iii) in Definition B1 are met. To begin, recall the following property of the exponential.

Lemma B8. *For every $\beta \in [0,1)$ and $\theta \in \mathbb{R}$, $\lim_{\lambda\to 1}(\lambda + (1-\lambda)\beta)^{\frac{\theta}{1-\lambda}} = e^{-(1-\beta)\theta}$.*

Let $\overline{\beta}_i := \max_a \beta_i(a)$ and $\underline{\beta}_i := \min_a \beta_i(a)$, and for every $\lambda > 0$, let $\overline{\beta}_{i\lambda} := \lambda + (1-\lambda)\overline{\beta}_i$ and $\underline{\beta}_{i\lambda} := \lambda + (1-\lambda)\underline{\beta}_i$.

Lemma B9. *Take $T_1 = \lceil \frac{\theta(1-\eta)}{1-\lambda} \rceil$ and $T_2 = \lceil \frac{\theta\eta}{1-\lambda} \rceil$, where $\theta > 0, 0 < \eta < 1$. There exist $\theta^* > 0$, $\gamma' > 0$, and $\underline{\lambda}' \in [0,1)$ such that if $\theta = \theta^*$, then for every $i \in I$, $\lambda \in (\underline{\lambda}', 1)$, and $\eta \in (0,1)$,*

$$(1 - [\underline{\beta}_{i\lambda}]^{n(T_1+T_2)})v_i(l^i) + [\underline{\beta}_{i\lambda}]^{n(T_1+T_2)}\varepsilon > \gamma'.$$

Proof. By Lemma B8,

$$\lim_{\lambda\to 1}(1 - [\underline{\beta}_{i\lambda}]^{n(T_1+T_2)})v_i(l^i) + [\underline{\beta}_{i\lambda}]^{n(T_1+T_2)}\varepsilon = (1 - \frac{1}{e^{(1-\underline{\beta}_i)n\theta}})v_i(l^i) + \frac{1}{e^{(1-\underline{\beta}_i)n\theta}}\varepsilon.$$

Let $f_i(\theta)$ denote the above limit and notice that $f_i(0) = \varepsilon > 0$ for every $i \in I$. Since $v_i(l^i) \leq 0 < \varepsilon$, f_i is decreasing and continuous in θ. Thus, there exists $\theta_i > 0$, small enough, such that $f_i(\theta) > 0$ for all $\theta \in (0, \theta_i]$. Take $\theta^* := \min_i \theta_i$ and choose $\gamma' > 0$ such that $f_i(\theta^*) > \gamma'$ for all $i \in I$. Finally, pick $\underline{\lambda}'_i > 0$ such that

$$(1 - [\underline{\beta}_{i\lambda}]^{n(T_1+T_2)})v_i(l^i) + [\underline{\beta}_{i\lambda}]^{n(T_1+T_2)}\varepsilon > \gamma' \quad \forall \lambda \in (\underline{\lambda}'_i, 1),$$

and let $\underline{\lambda}' := \max_i \underline{\lambda}'_i$ to complete the proof. □

Lemma B10. *Let θ^* be defined as in Lemma B9. Take $T_1 = \lceil \frac{\theta^*(1-\eta)}{1-\lambda} \rceil$ and $T_2 = \lceil \frac{\theta^*\eta}{1-\lambda} \rceil$ where $0 < \eta < 1$. There exist $0 < \eta^* < 1$, $\gamma'' > 0$, and $\underline{\lambda}'' \in [0,1)$ such that if $\eta = \eta^*$, then for every $i \in I$ and $\lambda \in (\underline{\lambda}'', 1)$*

$$(1 - [\overline{\beta}_{i\lambda}]^{T_1}[\underline{\beta}_{i\lambda}]^{nT_2})\varepsilon - (1 - [\overline{\beta}_{i\lambda}]^{T_1})v_i(l^i) - [\overline{\beta}_{i\lambda}]^{T_1}(1 - [\underline{\beta}_{i\lambda}]^{nT_2})v_i(h^i) > \gamma''.$$

Proof. For every $i \in I$, define

$$f_i(\eta) := \frac{(1 - e^{-(1-\overline{\beta}_i)\theta^*(1-\eta)})v_i(l^i) + e^{-(1-\overline{\beta}_i)\theta^*(1-\eta)}(1 - e^{-(1-\underline{\beta}_i)n\theta^*\eta})v_i(h^i)}{1 - e^{-(1-\overline{\beta}_i)\theta^*(1-\eta)-(1-\underline{\beta}_i)n\theta^*\eta}}.$$

The function f_i is continuous, strictly increasing, and such that $f_i(0) = v_i(l^i) \leq 0 < \varepsilon$. Thus, there exists $\eta_i > 0$, small enough, such that $f_i(\eta) < \varepsilon$ for all $\eta \in (0, \eta_i]$. Taking $\eta^* := \min_i \eta_i$, we have $f_i(\eta^*) < \varepsilon$ for every $i \in I$. Thus, there exists $\gamma'' > 0$ such that

$$(1 - \frac{1}{e^{(1-\overline{\beta}_i)\theta^*(1-\eta^*)+(1-\underline{\beta}_i)n\theta^*\eta^*}})\varepsilon - (1 - \frac{1}{e^{(1-\overline{\beta}_i)\theta^*(1-\eta^*)}})v_i(l^i)$$
$$- \frac{1}{e^{(1-\overline{\beta}_i)\theta^*(1-\eta^*)}}(1 - \frac{1}{e^{(1-\underline{\beta}_i)n\theta^*\eta^*}})v_i(h^i) > \gamma'' \quad \forall i \in I.$$

Lemma B8 implies that

$$\lim_{\lambda\to 1}(1-[\overline{\beta}_{i\lambda}]^{T_1}[\underline{\beta}_{i\lambda}]^{nT_2})\varepsilon-(1-[\overline{\beta}_{i\lambda}]^{T_1})v_i(l^i)-[\overline{\beta}_{i\lambda}]^{T_1}(1-[\underline{\beta}_{i\lambda}]^{nT_2})v_i(h^i)$$
$$=(1-\frac{1}{e^{(1-\overline{\beta}_i)\theta^*(1-\eta^*)+(1-\underline{\beta}_i)n\theta^*\eta^*}})\varepsilon-(1-\frac{1}{e^{(1-\overline{\beta}_i)\theta^*(1-\eta^*)}})v_i(l^i)$$
$$-\frac{1}{e^{(1-\overline{\beta}_i)\theta^*(1-\eta^*)}}(1-\frac{1}{e^{(1-\underline{\beta}_i)n\theta^*\eta^*}})v_i(h^i).$$

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

$$(1-[\overline{\beta}_{i\lambda}]^{T_1}[\underline{\beta}_{i\lambda}]^{nT_2})\varepsilon-(1-[\overline{\beta}_{i\lambda}]^{T_1})v_i(l^i)-[\overline{\beta}_{i\lambda}]^{T_1}(1-[\underline{\beta}_{i\lambda}]^{nT_2})v_i(h^i)>\gamma''.$$

Taking $\underline{\lambda}'':=\max_i\underline{\lambda}''_i$ completes the proof. □

Let $T_1=\lceil\frac{\theta^*(1-\eta^*)}{1-\lambda}\rceil$ and $T_2=\lceil\frac{\theta^*\eta^*}{1-\lambda}\rceil$, where θ^* is defined as in Lemma B9 and η^* is defined as in Lemma B10.

Lemma B11. *There exist $\gamma'>0$ and $\underline{\lambda}'\in[0,1)$ such that $v_{i\lambda}(\mathbf{r}^i_\lambda)>\gamma'$ for all $i\in I,\lambda\in(\underline{\lambda}',1)$.*

Proof. By Lemma B9, there exist $\gamma'>0$ and $\underline{\lambda}'\in[0,1)$ such that

$$(1-[\underline{\beta}_{i\lambda}]^{n(T_1+T_2)})v_i(l^i)+[\underline{\beta}_{i\lambda}]^{n(T_1+T_2)}\varepsilon>\gamma'\quad\forall i\in I,\forall\lambda\in(\underline{\lambda}',1).\tag{11}$$

Take $\lambda\in(\underline{\lambda}',1)$ and $i\in I$. Since $v_i(l^i)\le v_i(l^{i_m})$ for all $m=0,...,N^i-1$ and $v_i(l^i)\le v_i(\kappa^{i_m})$ for all $m=0,...,n-1$, we have

$$v_{i\lambda}(\mathbf{r}^i_\lambda)\ge(1-[\underline{\beta}_{i\lambda}]^{n(T_1+T_2)})v_i(l^i)+[\underline{\beta}_{i\lambda}]^{n(T_1+T_2)}v_{i\lambda}(\mathbf{w}^i_\lambda).$$

Since $v_{i\lambda}(\mathbf{w}^i_\lambda)\ge\varepsilon$, we obtain

$$v_{i\lambda}(\mathbf{r}^i_\lambda)\ge(1-[\underline{\beta}_{i\lambda}]^{n(T_1+T_2)})v_i(l^i)+[\underline{\beta}_{i\lambda}]^{n(T_1+T_2)}\varepsilon>\gamma'.$$

The last inequality follows from (11) and $\lambda\in(\underline{\lambda}',1)$. □

Lemma B12. *There exist $\gamma''>0$ and $\underline{\lambda}''\in[0,1)$ such that $v_{i\lambda}(\mathbf{r}^i_\lambda)<v_{i\lambda}(\mathbf{w}^i_\lambda)-\gamma''$ for all $i\in I$ and $\lambda\in(\underline{\lambda}'',1)$.*

Proof. Fix $i\in I$. Since $v_i(h^i)\ge v_i(\kappa^{i_m})$ for all $m=0,...,n-1$, we obtain

$$v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda)\le(1-[\underline{\beta}_{i\lambda}]^{nT_2})v_i(h^i)+[\underline{\beta}_{i\lambda}]^{nT_2}v_{i\lambda}(\mathbf{w}^i_\lambda).$$

By Lemma B5, $v_{i\lambda}(\mathbf{r}^i_\lambda)$ reaches its maximum when $\mathcal{L}^\lambda_i=\{l^i\}$. Since $v_i(l^i)<v_{i\lambda}(\mathbf{w}^i_\lambda)\le v_i(h^i)$, we have $v_{i\lambda}(\mathbf{r}^i_\lambda)\le x$

$$x=(1-[\overline{\beta}_{i\lambda}]^{T_1})v_i(l^i)+[\overline{\beta}_{i\lambda}]^{T_1}(1-[\underline{\beta}_{i\lambda}]^{nT_2})v_i(h^i)+[\overline{\beta}_{i\lambda}]^{T_1}[\underline{\beta}_{i\lambda}]^{nT_2}v_{i\lambda}(\mathbf{w}^i_\lambda).$$

Since $v_{i\lambda}(\mathbf{w}^i_\lambda)\ge\varepsilon$, Lemma B10 implies that there are $\gamma''>0$ and $\underline{\lambda}''\in[0,1)$ such that for all $i\in I$ and $\lambda\in(\underline{\lambda}'',1)$,

$$(1-[\overline{\beta}_{i\lambda}]^{T_1}[\underline{\beta}_{i\lambda}]^{nT_2})v_{i\lambda}(\mathbf{w}^i_\lambda)-(1-[\overline{\beta}_{i\lambda}]^{T_1})v_i(l^i)-[\overline{\beta}_{i\lambda}]^{T_1}(1-[\underline{\beta}_{i\lambda}]^{nT_2})v_i(h^i)>\gamma''.$$

This is equivalent to $x < v_{i\lambda}(\mathbf{w}^i_\lambda) - \gamma''$. Thus, $v_{i\lambda}(\mathbf{r}^i_\lambda) \leq x < v_{i\lambda}(\mathbf{w}^i_\lambda) - \gamma''$. □

Lemma B13. *For all $i \in I$ and all $\lambda \in (\underline{\lambda}'', 1)$, $v_{i\lambda}(\mathbf{r}^i_\lambda) \leq v_{i\lambda}({}_t\mathbf{r}^i_\lambda)$ for all t.*

Proof. Take $\lambda \in (\underline{\lambda}'', 1)$ and $i \in I$. Since $v_i(l^{i_m}) < v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda)$ for all $m = 0, ..., N^i - 1$, it follows from (9) and (10) that

$$v_{i\lambda}(\mathbf{r}^i_\lambda) \leq v_{i\lambda}({}_1\mathbf{r}^i_\lambda) \leq ... \leq v_{i\lambda}({}_{N^iT_1-1}\mathbf{r}^i_\lambda) \leq v_{i\lambda}({}_{N^iT_1}\mathbf{r}^i_\lambda) = v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda). \tag{12}$$

Thus, $v_{i\lambda}(\mathbf{r}^i_\lambda) \leq v_{i\lambda}({}_t\mathbf{r}^i_\lambda)$ for all $t \leq N^iT_1$. To prove the same for $t > N^iT_1$, suppose first that

$$v_i(\kappa^{i_m}) < v_{i\lambda}({}_{(m+1)T_2}\boldsymbol{\alpha}^i_\lambda) \quad \forall m = 0, ..., n-1. \tag{13}$$

The construction of $\boldsymbol{\alpha}^i_\lambda$ implies that for every $m = 0, ..., n-1$,

$$v_{i\lambda}({}_{mT_2}\boldsymbol{\alpha}^i_\lambda) = v_i(\kappa^{i_m})(1 - [\beta_{i\lambda}(\kappa^{i_m})]^{T_2}) + [\beta_{i\lambda}(\kappa^{i_m})]^{T_2} v_{i\lambda}({}_{(m+1)T_2}\boldsymbol{\alpha}^i_\lambda). \tag{14}$$

It follows from (9) and (13) that $v_{i\lambda}({}_{mT_2}\boldsymbol{\alpha}^i_\lambda) < v_{i\lambda}({}_{(m+1)T_2}\boldsymbol{\alpha}^i_\lambda)$ for all $m = 0, ..., n-1$. Hence, $v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda) < v_{i\lambda}({}_t\boldsymbol{\alpha}^i_\lambda)$ for all $t > 0$. Together with (12), this implies $v_{i\lambda}(\mathbf{r}^i_\lambda) \leq v_{i\lambda}({}_t\mathbf{r}^i_\lambda)$ for all $t > N^iT_1$.

Alternatively, suppose that there is an index k such that $v_i(\kappa^{i_k}) \geq v_{i\lambda}({}_{(k+1)T_2}\boldsymbol{\alpha}^i_\lambda)$ and $v_i(\kappa^{i_m}) < v_{i\lambda}({}_{(m+1)T_2}\boldsymbol{\alpha}^i_\lambda)$ for all $m < k$. It follows from (9) and (12) that

$$v_{i\lambda}(\mathbf{r}^i_\lambda) \leq v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda) < v_{i\lambda}({}_t\boldsymbol{\alpha}^i_\lambda) \quad \forall t = 1, ..., kT_2.$$

Since $v_i(\kappa^{i_k}) \geq v_{i\lambda}({}_{(k+1)T_2}\boldsymbol{\alpha}^i_\lambda)$, (14) and (9) yield

$$v_{i\lambda}({}_{kT_2}\boldsymbol{\alpha}^i_\lambda) \geq v_{i\lambda}({}_t\boldsymbol{\alpha}^i_\lambda) \quad t = kT_2 + 1, ..., (k+1)T_2.$$

By construction,

$$v_{i\lambda}({}_{(k+1)T_2}\boldsymbol{\alpha}^i_\lambda) = v_i(\kappa^{i_{k+1}})(1 - [\beta_{i\lambda}(\kappa^{i_{k+1}})]^{T_2}) + [\beta_{i\lambda}(\kappa^{i_{k+1}})]^{T_2} v_{i\lambda}({}_{(k+2)T_2}\boldsymbol{\alpha}^i_\lambda).$$

Since $v_i(\kappa^{i_{k+1}}) \geq v_i(\kappa^{i_k}) \geq v_{i\lambda}({}_{(k+1)T_2}\boldsymbol{\alpha}^i_\lambda)$, we have $v_i(\kappa^{i_{k+1}}) \geq v_{i\lambda}({}_{(k+2)T_2}\boldsymbol{\alpha}^i_\lambda)$. The latter implies that

$$v_{i\lambda}({}_{(k+1)T_2}\boldsymbol{\alpha}^i_\lambda) \geq v_{i\lambda}({}_t\boldsymbol{\alpha}^i_\lambda) \quad \forall t = (k+1)T_2 + 1, ..., (k+2)T_2.$$

Repeating the arguments above, we can show that for every $t = kT_2 + 1, ..., nT_2 - 1$,

$$v_{i\lambda}({}_{kT_2}\boldsymbol{\alpha}^i_\lambda) \geq v_{i\lambda}({}_t\boldsymbol{\alpha}^i_\lambda) \geq v_{i\lambda}({}_{nT_2}\boldsymbol{\alpha}^i_\lambda) = v_{i\lambda}(\mathbf{w}^i_\lambda). \tag{15}$$

For all $t > nT_2$, we have ${}_t\boldsymbol{\alpha}^i_\lambda = {}_\tau\mathbf{w}^i_\lambda \in SIR^\varepsilon(\lambda)$, where $\tau = t - nT_2$. Hence, $v_{i\lambda}(\mathbf{w}^i_\lambda) \leq v_{i\lambda}({}_t\boldsymbol{\alpha}^i_\lambda)$. Combined with (15), this yields

$$v_{i\lambda}(\mathbf{w}^i_\lambda) = v_{i\lambda}({}_{nT_2}\boldsymbol{\alpha}^i_\lambda) \leq v_{i\lambda}({}_t\boldsymbol{\alpha}^i_\lambda) \quad \forall t \geq kT_2 + 1.$$

Since $\lambda \in (\underline{\lambda}'', 1)$, Lemma B12 shows that $v_{i\lambda}(\mathbf{r}^i_\lambda) < v_{i\lambda}(\mathbf{w}^i_\lambda) \leq v_{i\lambda}({}_t\boldsymbol{\alpha}^i_\lambda)$ for all $t \geq kT_2 + 1$, completing the proof. □

Lemma B14. *There exist $\gamma''' > 0$ and $\underline{\lambda}''' \in [0,1)$ such that for every $i, j \in I$, $i \neq j$, and $\lambda \in (\underline{\lambda}''', 1)$, we have $[\underline{\beta}_{i\lambda}]^{nT_1}(v_i(\kappa^j) - v_i(\kappa^i))(1 - [\overline{\beta}_{i\lambda}]^{T_2})^2 > \gamma'''$.*

Proof. By Lemma B8,

$$\lim_{\lambda \to 1} [\underline{\beta}_{i\lambda}]^{nT_1}(v_i(\kappa^j) - v_i(\kappa^i))(1 - [\overline{\beta}_{i\lambda}]^{T_2})^2 = \frac{1}{e^{(1-\underline{\beta}_i)n\theta(1-\eta)}}(v_i(\kappa^j) - v_i(\kappa^i))(1 - \frac{1}{e^{(1-\overline{\beta}_i)\theta\eta}})^2,$$

which is strictly greater than 0 since $v_i(\kappa^j) - v_i(\kappa^i) > 0$ for all $j \neq i$. □

Given a list $B = (x^0, ..., x^{T-1})$ in a product space X^T and $k < T-1$, we write ${}_kB$ for the list $(x^k, x^{k+1}, ..., x^{T-1}) \in X^{T-k}$. Given lists $B = (x^0, ..., x^{T-1}) \in X^T$ and $B' = (y^0, ..., y^{K-1}) \in X^K$, we write $B \subset B'$ if $\{x^0, ..., x^{T-1}\} \subset \{y^0, ..., y^{K-1}\}$. Given a list $B = (\alpha^0, ..., \alpha^{T-1})$ of action profiles, we let $\pi_i^{\uparrow}(B) := (x^{\pi(0)}, ..., x^{\pi(T-1)})$ be the permutation of B such that $v_i(\alpha^{\pi(t)}) \leq v_i(\alpha^{\pi(t+1)})$ for all $t = 0, ..., T-2$.

Lemma B15. *For all $i, j \in I, i \neq j, \lambda \in (\underline{\lambda}''', 1)$, and $t \leq N^j T_1$, $v_{i\lambda}({}_t\mathbf{r}^j_\lambda) - v_{i\lambda}({}_tL^j_\lambda, \boldsymbol{\alpha}^i_\lambda) > \gamma'''$.*

Proof. For all $t \leq N^j T_1$, we have ${}_t\mathbf{r}^j_\lambda = ({}_tL^j_\lambda, \boldsymbol{\alpha}^j_\lambda)$ and, hence,

$$\begin{aligned} &v_{i\lambda}({}_t\mathbf{r}^j_\lambda) - v_{i\lambda}({}_tL^j_\lambda, \boldsymbol{\alpha}^i_\lambda) \geq v_{i\lambda}(L^j_\lambda, \boldsymbol{\alpha}^j_\lambda) - v_{i\lambda}(L^j_\lambda, \boldsymbol{\alpha}^i_\lambda) = \\ &= \prod_{m=0}^{N^j-1} [\beta_{i\lambda}(l^{j_m})]^{T_1}(v_{i\lambda}(\boldsymbol{\alpha}^j_\lambda) - v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda)) \geq [\underline{\beta}_{i\lambda}]^{nT_1}(v_{i\lambda}(\boldsymbol{\alpha}^j_\lambda) - v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda)). \end{aligned} \tag{16}$$

Thus, we seek a lower bound for $v_{i\lambda}(\boldsymbol{\alpha}^j_\lambda) - v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda)$. By the construction of $\boldsymbol{\alpha}^i_\lambda$, there is an index $k \neq 0$ such that $\kappa^{i_k} = \kappa^j$. Let

$$K^{i\setminus j} := ((\kappa^{i_0})^{T_2}, ..., (\kappa^{i_{k-1}})^{T_2}, (\kappa^{i_{k+1}})^{T_2}, ..., (\kappa^{i_{n-1}})^{T_2}) \text{ and } K^{j\setminus j} := ((\kappa^{j_1})^{T_2}, (\kappa^{j_2})^{T_2}, ..., (\kappa^{j_{n-1}})^{T_2}).$$

Thus, $K^{i\setminus j}$ and $K^{j\setminus j}$ are obtained from K^i and K^j respectively by removing the κ^j's. The list $K^{i\setminus j}$, like K^i, orders its elements in a way that is unfavorable to player i. Thus, by Lemma B4, $v_{i\lambda}(K^{j\setminus j}, \mathbf{w}^i_\lambda) \geq v_{i\lambda}(K^{i\setminus j}, \mathbf{w}^i_\lambda)$ and, by stationarity,

$$v_{i\lambda}(K^j, \mathbf{w}^i_\lambda) = v_{i\lambda}((\kappa^j)^{T_2}, K^{j\setminus j}, \mathbf{w}^i_\lambda) \geq v_{i\lambda}((\kappa^j)^{T_2}, K^{i\setminus j}, \mathbf{w}^i_\lambda).$$

Since $v_{i\lambda}(\mathbf{w}^j_\lambda) \geq v_{i\lambda}(\mathbf{w}^i_\lambda)$,

$$v_{i\lambda}(\boldsymbol{\alpha}^j_\lambda) = v_{i\lambda}(K^j, \mathbf{w}^j_\lambda) \geq v_{i\lambda}(K^j, \mathbf{w}^i_\lambda) \geq v_{i\lambda}((\kappa^j)^{T_2}, K^{i\setminus j}, \mathbf{w}^i_\lambda).$$

Next, let $\tilde{K}$ be the list obtained from K^i by moving the block $(\kappa^j)^{T_2}$ immediately after the initial block $(\kappa^i)^{T_2}$. By Lemma B4, we have $v_{i\lambda}(\tilde{K}, \mathbf{w}^i_\lambda) \geq v_{i\lambda}(K^i, \mathbf{w}^i_\lambda) = v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda)$. We conclude that

$$\begin{aligned} &[\underline{\beta}_{i\lambda}]^{nT_1}(v_{i\lambda}(\boldsymbol{\alpha}^j_\lambda) - v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda)) \geq [\underline{\beta}_{i\lambda}]^{nT_1}(v_{i\lambda}((\kappa^j)^{T_2}, K^{i\setminus j}, \mathbf{w}^i_\lambda) - v_{i\lambda}(\tilde{K}, \mathbf{w}^i_\lambda)) \\ =&[\underline{\beta}_{i\lambda}]^{nT_1}(v_i(\kappa^j) - v_i(\kappa^i))(1 - [\beta_{i\lambda}(\kappa^j)]^{T_2})(1 - [\beta_{i\lambda}(\kappa^i)]^{T_2}) \end{aligned}$$

$$\geq [\underline{\beta}_{i\lambda}]^{nT_1}(v_i(\kappa^j) - v_i(\kappa^i))(1 - [\overline{\beta}_{i\lambda}]^{T_2})^2 > \gamma''',$$

where the equality follows by a direct calculation and the last inequality by Lemma B14. Together with (16), the last chain of inequalities completes the proof. □

Lemma B16. *For all $i, j \in I, i \neq j$, and $\lambda \in (\underline{\lambda}''', 1)$, $v_{i\lambda}({}_t\mathbf{r}^j_\lambda) - v_{i\lambda}(\mathbf{r}^i_\lambda) > \gamma'''$ for all $t \leq N^j T_1$.*

Proof. Write ${}_t\mathbf{r}^j_\lambda$ as $({}_tL^j_\lambda, \boldsymbol{\alpha}^j_\lambda)$. By Lemma B4, $v_{i\lambda}({}_tL^j_\lambda, \boldsymbol{\alpha}^i_\lambda) \geq v_{i\lambda}(\pi^\uparrow_i({}_tL^j_\lambda), \boldsymbol{\alpha}^i_\lambda)$. Hence, by Lemma B15, $v_{i\lambda}\left({}_t\mathbf{r}^j_\lambda\right) - v_{i\lambda}(\pi^\uparrow_i({}_tL^j_\lambda), \boldsymbol{\alpha}^i_\lambda) > \gamma'''$. It is therefore enough to show that $v_{i\lambda}(\pi^\uparrow_i({}_tL^j_\lambda), \boldsymbol{\alpha}^i_\lambda) \geq v_{i\lambda}(\mathbf{r}^i_\lambda)$. Recall that $\mathbf{r}^i_\lambda = (L^i_\lambda, \boldsymbol{\alpha}^i_\lambda)$. Since ${}_tL^j_\lambda \subset L^j_\lambda$, we can write $\pi^\uparrow_i({}_tL^j_\lambda)$ as (L', L'') where $L' \subset L^i_\lambda$ and $L'' \subset L^j_\lambda \setminus L^i_\lambda$. We claim that

$$v_{i\lambda}(L', L'', \boldsymbol{\alpha}^i_\lambda) \geq v_{i\lambda}(L', \boldsymbol{\alpha}^i_\lambda) \geq v_{i\lambda}(L^i_\lambda, \boldsymbol{\alpha}^i_\lambda) =: v_{i\lambda}(\mathbf{r}^i_\lambda) \tag{17}$$

By stationarity, or if $L' = \emptyset$, the first inequality is equivalent to $v_{i\lambda}(L'', \boldsymbol{\alpha}^i_\lambda) \geq v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda)$, which follows since $v_i(l'') \geq v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda)$ for all $l'' \in L''$. The second inequality in (17) follows from Lemma B5. □

Lemma B17. *For all $i \neq j, \lambda \in (\underline{\lambda}'', 1)$, $v_{i\lambda}({}_t\mathbf{r}^j_\lambda) - v_{i\lambda}(\mathbf{r}^i_\lambda) > \gamma''$ for all $t > N^j T_1$.*

Proof. The desired inequality is equivalent to $v_{i\lambda}({}_t\boldsymbol{\alpha}^j_\lambda) \geq v_{i\lambda}(\mathbf{r}^i_\lambda)$ for all $t > 0$. If $t \geq nT_2$, then ${}_t\boldsymbol{\alpha}^j_\lambda = {}_\tau\mathbf{w}^j_\lambda \in SIR^\varepsilon(\lambda)$ where $\tau = t - nT_2$. Hence, $v_{i\lambda}({}_t\boldsymbol{\alpha}^j_\lambda) \geq v_{i\lambda}(\mathbf{w}^i_\lambda)$. By Lemma B12, $v_{i\lambda}(\mathbf{w}^i_\lambda) - \gamma'' > v_{i\lambda}(\mathbf{r}^i_\lambda)$ and we are done. Suppose now that $t < nT_2$ and write ${}_t\boldsymbol{\alpha}^j_\lambda$ as $({}_tK^j, \mathbf{w}^j_\lambda)$. Lemmas B5 and B4 imply that

$$v_{i\lambda}(\mathbf{r}^i_\lambda) \leq v_{i\lambda}((l^i)^{T_1}, K^i, \mathbf{w}^i_\lambda) \leq v_{i\lambda}((l^i)^{T_1}, {}_tK^j, K^i \setminus {}_tK^j, \mathbf{w}^i_\lambda).$$

These inequalities, together with the construction of $\mathbf{w}^i_\lambda$, yield

$$v_{i\lambda}({}_tK^j, \mathbf{w}^j_\lambda) - v_{i\lambda}(\mathbf{r}^i_\lambda) \geq v_{i\lambda}({}_tK^j, \mathbf{w}^i_\lambda) - v_{i\lambda}((l^i)^{T_1}, {}_tK^j, K^i \setminus {}_tK^j, \mathbf{w}^i_\lambda) =: x.$$

Lengthy but straightforward calculations show that

$$x \geq (1 - [\overline{\beta}_{i\lambda}]^{T_1}[\underline{\beta}_{i\lambda}]^{nT_2})\varepsilon - (1 - [\overline{\beta}_{i\lambda}]^{T_1})v_i(l^i) - [\overline{\beta}_{i\lambda}]^{T_1}(1 - [\underline{\beta}_{i\lambda}]^{nT_2})v_i(h^i).$$

By Lemma B10, $x > \gamma''$ whenever $\lambda \in (\underline{\lambda}'', 1)$. □

Take $\gamma := \min\{\gamma', \gamma'', \gamma'''\}$ and $\underline{\lambda} := \max\{\underline{\lambda}', \underline{\lambda}'', \underline{\lambda}'''\}$, where $\gamma', \gamma'', \gamma'''$ and $\underline{\lambda}', \underline{\lambda}'', \underline{\lambda}'''$ are defined as in Lemmas B11, B12, and B16. Then, Lemmas B12, B13, B16, and B17 show that for all $\lambda \in (\underline{\lambda}, 1)$ and $\boldsymbol{\alpha} \in SIR^\varepsilon(\lambda)$, the paths $\{\mathbf{r}^i_\lambda\}_{i \in I}$ meet all the conditions in Definition B1 .

B.4 Equilibrium Strategies

Let $m^i := (m^i_1, ..., m^i_n) \in \Sigma$ be a strategy profile in which player i best-responds to a minmax strategy by the opponents. 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 $\Delta(A)$. Utilities are normalized so that $g_i(m^i) = 0$ for every $i \in I$. Take $\varepsilon > 0$. By Lemma B7, there exist $\gamma > 0$ and $\underline{\lambda}' \geq 0$ such that for every $\lambda > \underline{\lambda}'$, every $\boldsymbol{\alpha} \in SIR^{\varepsilon}(\lambda)$ allows DPSP with wedge γ. Let $\overline{g}_i := \max_a g_i(a)$ and choose an integer μ_i such that $\mu_i > \frac{\overline{g}_i}{\gamma\left(1-\beta_i(m^i)\right)}$. Recall that $\beta_{i\lambda}(m^i) := \lambda + (1-\lambda)\beta_i(m^i)$ 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

$$\frac{\overline{g}_i}{\gamma\left(1-\beta_i(m^i)\right)} < \frac{1-[\beta_{i\lambda}(m^i)]^{\mu_i}}{1-\beta_{i\lambda}(m^i)} \quad \forall \lambda > \underline{\lambda}''_i. \tag{18}$$

Fix $j \neq i$ and an integer μ between 1 and μ_j. Let $\overline{m} := \max_{i,a} \frac{g_i(a)}{1-\beta_i(a)}$ and consider the inequality

$$(1-\lambda)\overline{g}_i + \left(\overline{m} - [\beta_{i\lambda}(m^j)]^{\mu}(\overline{m}+\gamma)\right) - \frac{g_i(m^j)}{1-\beta_i(m^j)}\left(1-[\beta_{i\lambda}(m^j)]^{\mu}\right) < 0. \tag{19}$$

Note that $\overline{g}_i$ and $\frac{g_i(m^j)}{1-\beta_i(m^j)}$ are constants that do not depend on λ. Also, the first and last term converge to 0 as $\lambda \to 1$. The second term converges to a negative number. Thus, there exists $\underline{\lambda}'''_i$ such that the inequality in (19) is satisfied for all $\lambda > \underline{\lambda}'''_i$. 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 $j \neq i$ and μ.

Let $\underline{\lambda}_i := \max\{\underline{\lambda}''_i, \underline{\lambda}'''_i\}$, $\underline{\lambda}'' := \max_i \underline{\lambda}_i$, and $\underline{\lambda} := \max\{\underline{\lambda}', \underline{\lambda}''\}$. Take any $\lambda > \underline{\lambda}$ and $\boldsymbol{\alpha} \in SIR^{\varepsilon}(\lambda)$. Let $\{\mathbf{r}^i_{\lambda}\}_{i \in I}$ be the DPSP with wedge γ. By definition, we have $v_{i\lambda}({}_t\boldsymbol{\alpha}) \geq \varepsilon$, for all $i \in I$ and $t \in \mathcal{T}$. Consider the following strategy $\sigma_i \in \Sigma_i$ for player i: (A) play α_i as long as α was played last period. If player j deviates from (A), then (B) play m^j_i for μ_j periods, and then (C) play $\mathbf{r}^j_{\lambda}$ thereafter. If player k deviates in phase (B) or (C), begin phase (B) again with $j = k$. It remains to show that, given the choice of $\underline{\lambda}$, no player has an incentive to deviate. We leave this part of the proof, which follows from simple calculations, as an exercise.

C Convexity of the Feasible Set

Fix an n-player game $(A, (g_i, \beta_i)_i)$. The next two lemmas establish that the set $V \subset \mathbb{R}^n$ of all feasible payoffs is convex, which will be useful in subsequent proofs.

Lemma C18. *For every measure $\sigma \in \Delta(A^{\infty})$ with finite support, there exists a behavioral strat-*

egy that induces σ.[25]

Proof. Let $\{\mathbf{a}_1, ..., \mathbf{a}_m\} \subset A^\infty$ be the support of $\sigma \in \Delta(A^\infty)$. W.l.o.g., we can assume that $\mathbf{a}_k \neq \mathbf{a}_l$ for every $k, l \leq m$. Let Π'_t be the partition of $S := \{1, ..., m\}$ such that $k, \tilde{k}$ belong to an element of Π'_t if and only if $a^t_k = a^t_{\tilde{k}}$, and let Π_t be the join of the partitions $\Pi'_0, ..., \Pi'_t$. View σ as a probability measure on S and for each $t \in \mathcal{T}, A \in \Pi_t, B \in \Pi_{t+1}$ such that $B \subset A$, we can define the conditional probability $\sigma(B|A, t)$. But each $B \in \Pi_{t+1}$ corresponds to a pure action $a \in A$. We can thus define a mixed action $\alpha_{A,t} \in \Delta(A)$ by letting $\alpha(a|A, t) = \sigma(B|A, t)$. Moreover, each pair $(t, A), A \in \Pi_t$ corresponds to a history in the repeated game. The on-path behavioral strategy "play $\alpha_{t,A}$ in history (t, A)" induces the distribution $\sigma \in \Delta(A^\infty)$. □

Lemma C19. *The set of all feasible payoffs is convex.*

Proof. Let V^p be the convex hull of $v(A^\infty)$. Note that the set $v(A^\infty) \subset \mathbb{R}^n$ is compact since A^∞ is compact in the product topology and each function v_i is continuous. As the convex hull of a finite-dimensional compact set, V^p is compact as well. By Lemma C18, $V^p \subset V$. Thus, the proof would be complete if we can show the converse inclusion. Take any Borel measure $\sigma \in \Delta(A^\infty)$ and let $\sigma^m \in \Delta(A^\infty), m \in \mathbb{N}$, be measures with finite support such that $\sigma^m \to_m \sigma$ in the weak* topology. By construction, $v(\sigma^m) \in V^p$ for each m. Since V^p is compact, $v(\sigma) \in V^p$ and $V \subset V^p$. □

We need one more lemma, whose straightforward proof we omit.

Lemma C20. *Assume the game is symmetric and take $a \in A$ and $i, j \in I$.*

1. *Under IMI, $v_i(a) \geq v_j(a)$ if and only if $\beta_i(a) \leq \beta_j(a)$.*

2. *Under DMI, $v_i(a) \geq v_j(a)$ if and only if $\beta_i(a) \geq \beta_j(a)$.*

D Proof of Theorem 5.1

Necessity of CI is obvious. Turn to sufficiency. For every $\eta \in \mathbb{R}^n_+$, define the **η-face** of V^{pc} to be the set

$$F(\eta) = \{v \in V^{pc} : \eta \cdot v \geq \eta \cdot v' \ \forall v' \in V^{pc}\}.$$

Lemma D21. *For some $\eta \in \mathbb{R}^n_{++}$, the set $F(\eta)$ is not a singleton.*

[25] The lemma is closely related to Kuhn's theorem. We provide a proof as we are unaware of this particular formulation.

Proof. By way of contradiction, suppose $F(\eta)$ is a singleton for each $\eta \in \mathbb{R}^n_{++}$. Since V^{pc} is a polytope, $E := \{F(\eta) : \eta \in \mathbb{R}^n_{++}\}$ is a finite set of extreme points. By CI, the set E is not a singleton. For every $v \in E$, let $N(v) = \{\eta \in \mathbb{R}^n_{++} : F(\eta) = \{v\}\}$. By construction, each set $N(v)$ is closed in $\mathbb{R}^n_{++}$ and $N(v) \cap N(v') = \emptyset$ for all distinct $v, v' \in E$. But then $\{N(v) : v \in E\}$ is a finite partition of $\mathbb{R}^n_{++}$ into disjoint relatively closed subsets, which is impossible since $\mathbb{R}^n_{++}$ is connected. □

Pick $\eta \in \mathbb{R}^n_{++}$ such that $F(\eta)$ is not a singleton. Let I^* be the set of players i such that $v_i \neq v'_i$ for some $v, v' \in F(\eta)$. For any v in the relative interior of $F(\eta)$, we are going to find a feasible payoff v' such that $v'_i > v_i$ for all $i \in I^*$ and $v'_i = v_i$ for all $i \notin I^*$. Since $\eta \gg 0$, such a v' cannot belong to V^{pc}, proving the first part of Theorem 5.1. Let $v^1, ..., v^m$ be the extreme points of $F(\eta)$ and $a^1, ..., a^m \in A$ be the action profiles that generate them, so that $v_i(a^k) = v^k_i$ for every $i \in I$ and $k \leq m$. If v is in the relative interior of $F(\eta)$, there are weights $(\varrho^1, ..., \varrho^m) \in (0,1)^m$ such that $\sum_{k=1}^m \varrho^k = 1$ and $v = \sum_{k=1}^m \varrho^k v^k$. Let $\alpha := \sum_{k=1}^m \varrho^k a^k \in \Delta(A)$ be the mixed action profile such that a^k is played with probability ϱ^k. By Lemma 5.1,

$$v_i := \sum_{k=1}^{m} \varrho^k v_i(a^k) \leq v_i(\sum_{k=1}^{m} \varrho^k a^k) =: v_i(\alpha) \quad \forall i \in I. \tag{20}$$

In addition, the inequality is strict if and only if $i \in I^*$, that is, if and only if $v_i(a^k) \neq v_i(a^l)$ for some $k, l \leq m$.

Next, suppose there is $v \in V^{pc}$ such that $v \gg 0$. Let $\varepsilon > 0$ be such that $v \gg \varepsilon$ and let V^{pc}_ε be the set of all $v' \in V^{pc}$ such that $v' \geq \varepsilon$. We claim that there is no $\overline{v} \in V^{pc}_\varepsilon$ such that $\overline{v} \geq v'$ for all $v' \in V^{pc}_\varepsilon$. If there were, then, by CI, there is $i \in I$ and $v^i \in V^{pc}$ such that $v^i_i > \overline{v}_i$. But then for all $\varrho \in (0,1)$ sufficiently high, $\varrho \overline{v}_i + (1-\varrho) v^i_i > \overline{v}_i$ and $\varrho \overline{v} + (1-\varrho) v^i \in V^{pc}_\varepsilon$, contradicting the definition of $\overline{v}$. Next, we can apply Lemma D21 to show that V^{pc}_ε has a face $F(\eta)$, with $\eta \gg 0$, that is not a singleton, and the arguments following the lemma to show that for any v' in the relative interior of $F(\eta)$, there is $\alpha \in \Delta(A)$ such that $v(\alpha) \geq v'$ and $v(\alpha) \neq v'$. By construction, $v(\alpha) \geq \varepsilon$. Hence, by Theorem 4.1, there is $\underline{\lambda}$ such that for all $\lambda \geq \underline{\lambda}$, $v(\alpha)$ can be sustained in a SPE.

E Proof of Theorem 5.2

The proof of Theorem 5.2 uses strategies in which the public randomization device (PRB) may recommend distinct *mixed* actions depending on the state of nature. This allows us to implement strategies in which, for example, the players follow a path $(\alpha^0, \alpha^1, ...) \in (\Delta A)^\infty$ with some probability ϱ or $(\hat{\alpha}^0, \hat{\alpha}^1, ...) \in (\Delta A)^\infty$ with probability $(1-\varrho)$. We call such play **random mixed paths (RMP)**. They can be implemented by letting the PRB be a uniform draw from $[0,1]$. In the present example, one can first partition $[0,1]$ into $[0,\varrho]$ and $(\varrho,1]$, and then partition the intervals $[0,\varrho]$ and $(\varrho,1]$ so as to implement the mixed actions α^0 and $\hat{\alpha}^0$ respectively.

Note that use of RMPs is not necessary when preferences are standard since the single mixed path $(\varrho\alpha^0 + (1-\varrho)\hat{\alpha}^0, \varrho\alpha^1 + (1-\varrho)\hat{\alpha}^1, ...) \in (\Delta A)^\infty$ would deliver an equivalent payoff. That is,

$$\varrho v_\lambda(\alpha^0, \alpha^1, ...) + (1-\varrho) v_\lambda(\hat{\alpha}^0, \hat{\alpha}^1, ...) = v_\lambda(\varrho\alpha^0 + (1-\varrho)\hat{\alpha}^0, \varrho\alpha^1 + (1-\varrho)\hat{\alpha}^1, ...). \tag{21}$$

In addition, if the two paths on the left are SIR, so is the RMP on the right.[26] What happens when discounting is endogenous? First, the identity in (21) breaks down. To gain some intuition why, note that the RMP, whose utility is on left-hand-side of (21), induces a distribution on A^∞ that is autocorrelated and, as the discussion following Theorem 5.1 showed, autocorrelations matter when discounting is endogenous. Also note that, by Lemmas C18 and C19, any feasible payoff, including that of an RMP, can be replicated without the use of RMPs. The problem is that when this replication is carried out, it is difficult to guarantee that play would be SIR. By comparison, if each of the single paths $(\alpha^0, \alpha^1, ...)$, $(\hat{\alpha}^0, \hat{\alpha}^1, ...)$ are SIR, then so is the RMP that randomizes between them.[27] This makes it easy to insure that the payoff attained by an RMP can be attained in a SPE.

We should also remark that RMPs are only used in the proof of Theorem 5.2. In particular, the off-path punishments we construct in the proof of our folk theorem take the form of single paths $(\alpha^0, \alpha^1, ...)$, not RMPs. Moreover, if RMPs appear "on path," we can use the same punishments to sustain the behavior in a SPE, provided, of course, that the appropriate SIR constraints are met and the players are sufficiently patient. That is, our folk theorem extends to the case of RMPs.[28] Finally, we note that allowing for RMPs works against us in Theorem 6.3 since the bounds derived in that theorem have to cover a larger class of strategies.

Lemma E22. *If $\alpha \in \Delta(A)$ is such that $v_i(\alpha) > v_k(\alpha)$ and $\beta_i(\alpha) < \beta_k(\alpha)$ for some $i, k \in I$, then for every $\eta \in \mathbb{R}^n_+$ such that $\eta_k > 0$ and every λ, there is $\boldsymbol{\alpha} \in (\Delta A)^\infty$ such that $\eta \cdot v_\lambda(\boldsymbol{\alpha}) > \eta \cdot v(\alpha)$. In addition, $v_{i\lambda}(\boldsymbol{\alpha}) < v_i(\alpha)$, $v_{k\lambda}(\boldsymbol{\alpha}) > v_k(\alpha)$, and $v_{j\lambda}(\boldsymbol{\alpha}) = v_j(\alpha)$ for all $j \neq i, k$.*

[26] Consistent with these observations, note that, while using a uniform distribution on $[0,1]$ as PRB is common, more minimalist formulations of the standard repeated-game model, such as Myerson [19, p.332], do not allow the use of RMPs.

[27] The SIR constraints for the RMP are: (i)$\varrho v_{i\lambda}(\alpha^0, \alpha^1, ...) + (1-\varrho) v_{i\lambda}(\hat{\alpha}^0, \hat{\alpha}^1, ...) \geq 0$ and, for every $t > 0$, $v_{i\lambda}(\alpha_t, \alpha_{t+1}, ...) \geq 0$ and $v_{i\lambda}(\hat{\alpha}_t, \hat{\alpha}_{t+1}, ...) \geq 0$.

[28] One can ask whether the use of RMPs could simplify the proof of our folk theorem. We suspect that the answer is no, as the single paths in the support of an RMP must have properties not unlike those in Definition B1. In addition, as we noted above, some specifications of the standard model do not allow for RMPs. Accordingly, we wanted to make sure that our folk theorem covers those cases and generalizes existing results.

Proof. Fix λ and η such that $\eta_k > 0$. By symmetry, there exists $\alpha_k \in \Delta(A)$ such that $v_i(\alpha_k) = v_k(\alpha), v_i(\alpha) = v_k(\alpha_k)$, $\beta_{i\lambda}(\alpha_k) = \beta_{k\lambda}(\alpha), \beta_{i\lambda}(\alpha) = \beta_{k\lambda}(\alpha_k)$, and for all $j \neq i, k$, $v_j(\alpha) = v_j(\alpha_k)$ and $\beta_{j\lambda}(\alpha) = \beta_{j\lambda}(\alpha_k)$. Since $\beta_{i\lambda}(\alpha) < \beta_{k\lambda}(\alpha)$, there is T large enough such that

$$v_k(\alpha_k) - v_k(\alpha) > \frac{\eta_i}{\eta_k}[\frac{\beta_{i\lambda}(\alpha)}{\beta_{k\lambda}(\alpha)}]^T(v_i(a) - v_i(\alpha_k)). \tag{22}$$

Let $\boldsymbol{\alpha}$ be the path $(\alpha^0, \alpha^1, ...)$ such that $\alpha^t = \alpha$ for all $t \leq T$ and $\alpha^t = \alpha_k$ for all $t > T$. It follows from (22) that $\eta \cdot v_\lambda(\boldsymbol{\alpha}) > \eta \cdot v(\alpha)$. The rest of the (in)equalities follow by construction. □

Lemma E23. *Suppose $\alpha \in \Delta(A)$ is such that for some $\varepsilon > 0$, $v_j^{max} > v_j(\alpha) \geq \varepsilon$ for every $j \in I$ and for some $i \in I$, $v_i(\alpha) > v_k(\alpha)$ and $\beta_i(\alpha) < \beta_k(\alpha)$ for every $k \neq i$. Then, for every λ, there is a payoff $\hat{v}_\lambda \gg v(\alpha)$. In addition, there is $\underline{\lambda}$ such that for all $\lambda \geq \underline{\lambda}$, the payoff $\hat{v}_\lambda$ can arise in a SPE.*

Proof. Let $a^{max,i} \in A$ be some action such that $v_i(a^{max,i}) = v_i^{max}$. By Lemma B6, $v_i(\rho\alpha + (1-\rho)a^{max,i}) > v_i(\alpha) \geq \varepsilon$ for all $\rho \in (0,1)$. Take $0 < \varepsilon' \leq \varepsilon$ and $\rho \in (0,1)$ close enough to 1 such that

$$v_k(\rho\alpha + (1-\rho)a^{max,i}) > \varepsilon' \quad \forall k \in I. \tag{23}$$

Let $v := (v_1(\alpha), ..., v_n(\alpha)), \alpha^\rho := \rho\alpha + (1-\rho)a^{max,i}$, and $v^\rho := (v_1(\alpha^\rho), ..., v_n(\alpha^\rho))$. By construction, $v_i^\rho - v_i > 0$. Fix λ. By Lemma E22, for every $k \neq i$ and $\eta \in \mathbb{R}^n_{++}$, there is a payoff $v^{k,\eta}$ (depending on λ) such that

$$\eta \cdot (v^{k,\eta} - v) > 0, \tag{24}$$

$$v_i^{k,\eta} - v_i < 0 < v_k^{k,\eta} - v_k \quad \text{and} \quad [v_j^{k,\eta} = v_j \quad \forall j \neq i, k] \tag{25}$$

Given (25), we can rewrite (24) as

$$\frac{|v_i^{k,\eta} - v_i|}{v_k^{k,\eta} - v_k} < \frac{\eta_k}{\eta_i}. \tag{26}$$

Choose $\eta^m \in \mathbb{R}^n_{++}$ such that $\frac{\eta_k^m}{\eta_i^m} \to_m 0$. By (26), we must have

$$\frac{|v_i^{k,\eta^m} - v_i|}{v_k^{k,\eta^m} - v_k} \to_m 0. \tag{27}$$

For each m, let A^m be the convex hull of the set $\{v^\rho\} \cup \{v^{k,\eta^m} : k \neq i\}$ and let $D := \{v' \in \mathbb{R}^n : v'_j > v_j \forall j\}$. We claim that $A^m \cap D \neq \emptyset$ for some m. If not, it follows from the separating hyperplane theorem that for every m, there is a probability vector $\rho^m \in [0,1]^n$

such that

$$\rho^m \cdot (v^\rho - v) \leq 0 \quad \text{and} \quad [\rho^m \cdot (v^{k,\eta^m} - v) \leq 0 \quad \forall k \neq i]. \tag{28}$$

It follows from the latter inequality and the inequalities in (25) that $\rho_i^m > 0$. Fix $k \neq i$. Using the equalities in (25), we can rewrite the second inequality in (28) as

$$\frac{\rho_k^m}{\rho_i^m} \leq \frac{|v_i^{k,\eta^m} - v_i|}{v_k^{k,\eta^m} - v_k}.$$

From (27), it follows that $\rho_k^m \to_m 0$ for every $k \neq i$. But then, $\rho^m \cdot (v^\rho - v) \to_m (v_i^\rho - v_i) > 0$, contradicting the first inequality in (28), which holds for every m. Thus, $A^m \cap D \neq \emptyset$ for some m, which means that there exists a probability vector $\gamma \in [0,1]^n$ such that

$$\hat{v}_\lambda := \gamma_i v^\rho + \sum\nolimits_{k \neq i} \gamma_k v^{k,\eta^m} \gg v.$$

Let $\boldsymbol{\alpha}^k \in (\Delta A)^\infty$ be the path that generates v^{k,η^m}. The payoff $\hat{v}_\lambda$ is obtained by the RMP in which $(\alpha^\rho, \alpha^\rho, ...)$ is played with probability γ_i and each $\boldsymbol{\alpha}^k \in A^\infty$ with probability γ_k, $k \neq i$. For each λ, all paths in the support of the RMP belong to $SIR^{\varepsilon'}(\lambda)$. As noted at the start of the section, Theorem 4.1 guarantees that the RMP can be sustained in a SPE for all λ sufficiently high. □

To complete the proof of Theorem 5.2, take $a \in A$ such that $v_j^{max} > v_j(a)$ for every j and, for some i, $v_i(a) > v_k(a)$ for all $k \neq i$. By Lemma C20, $\beta_i(a) < \beta_k(a)$ for all $k \neq i$. It follows from Lemma E22 that there are no λ and $\eta \in \mathbb{R}_+^n \setminus \{0\}$ such that $v(a)$ belongs to a face $F(\eta)$ of $V(\lambda)$. Since $V(\lambda)$ is convex, it follows that there is $\hat{v}_\lambda \in V(\lambda)$ such that $\hat{v}_\lambda \gg v(a)$. If, in addition, $v(a) \geq \varepsilon$ for some $\varepsilon > 0$, then Lemma E23 shows that we can choose the Pareto improvements $\hat{v}_\lambda$ such that they can arise in a SPE for all λ sufficiently high.

F Proof of Theorem 5.3

Let S_i be the set of $\alpha \in \Delta(A)$ such that $v_i(\alpha) > v_k(\alpha)$ and $\beta_i(\alpha) < \beta_k(\alpha)$ for some $k \neq i$. Let V_+^c be the Pareto frontier of V^c, that is, the set of points $v \in V^c$ such that there is no $v' \in V^c$ such that $v' \gg v$. Also, let e^i be the vector in $\mathbb{R}_+^n$ whose i^{th}-coordinate is 1 and all other coordinates are 0, and note that $v(a^{max,i}) \in F(e^i)$.

Lemma F24. *If $\alpha \in \Delta(A)$ is such that $v(\alpha) \in F(e^i)$, then $\alpha \in S_i$.*

Proof. If $a \in A$ is such that $v(a) \in F(e^i)$, then $a \in S_i$ by Lemma C20. By Lemma B6, if $\alpha \in \Delta(A) \in F(e^i)$, then every $a \in \operatorname{supp} \alpha$ is such that $v(a) \in F(e^i)$. By Lemma C20, $\beta_i(a) = \min\{\beta_i(a') : a' \in A\}$ for all $a \in \operatorname{supp} \alpha$. By the symmetry of the game, $v_i(\alpha) \geq$

$v_k(\alpha)$ and $\beta_i(\alpha) \leq \beta_k(\alpha)$. Finally, by CI, for every $a \in \operatorname{supp} \alpha$, there is $k \in I$ such that $v_i(a) > v_k(a)$ and, by Lemma C20, $\beta_i(a) < \beta_k(a)$. It follows that $v_i(\alpha) > v_k(\alpha)$ and $\beta_i(\alpha) < \beta_k(\alpha)$ for some k. □

Let X be the set of extreme points v of V^c such that $v \in V^c_+$. Let $Y := X \setminus F(e^i)$ and let Z be the set of $v \in X \cap F(e^i)$ such that every open neighborhood O of v intersects $V^c_+ \setminus F(e^i)$. Suppose $Z \cap \operatorname{cl} Y \neq \emptyset$, which one can think of a situation in which the face $F(e^i)$ connects smoothly with the rest of the frontier V^c_+. Let $\hat{v} \in Z$ and $y^m \in Y$ be such that $y^m \to_m \hat{v}$. Let α^m be such that $v(\alpha^m) = y^m$ for each m. Passing onto a subsequence if necessary, we may assume that $\alpha^m \to_m \alpha^*$, where $v(\alpha^*) = \hat{v} \in Z$. By Lemma F24, $\alpha^* \in S_i$ and, hence, $\alpha^m \in S_i$ for some m large enough. By construction, $v(\alpha^m)$ belongs to a face $F(\eta)$ of the frontier V^c_+ for some $\eta \in \mathbb{R}^n_+ \setminus \{0, e^i\}$. But, by Lemma E22, $v(\alpha^m)$ is not on the corresponding η-face of $V(\lambda)$. Thus, $V^c \subsetneq V(\lambda)$. Alternatively, suppose $F(e^i)$ connects nonsmoothly to the rest of the frontier V^c_+. Then, there is $v^* \in X$ belonging to both $F(e^i)$ and a face $F(\eta)$ of V^c_+ such that $\eta \in \mathbb{R}^n_+ \setminus \{0, e^i\}$. Again, by Lemma E22, $V^c \subsetneq V(\lambda)$. The next lemma, an adaptation of Lemma 2.2 in Toth [27], confirms that the two scenarios we considered are exhaustive.

Lemma F25. *If $Z \setminus \operatorname{cl} Y \neq \emptyset$, then there is $v^* \in X \cap F(e^i)$ such that $N(v^*)$ is not a singleton.*

G Two Player Games: Preliminary Lemmas

This section introduces some preliminary notation and results concerning two player games which are used in the remainder of the paper. First, when the indices i, j appear in the same context, it would be understood that $i \neq j$, i.e., that they refer to the two distinct players of the game. Fix $\lambda \in [0, 1)$ and $\eta \in \mathbb{R}^2_+$. Given $\mathbf{a} \in A^\infty$, define $s_\lambda(\mathbf{a}, \eta) := \eta \cdot v_\lambda(\mathbf{a})$ and let $P(\lambda, \eta)$ be the set of pure play paths $\mathbf{a} \in A^\infty$ that maximize the function $s_\lambda(\cdot, \eta)$. Clearly, $P(\lambda, \eta) \subset P(\lambda)$, that is, the paths in $P(\lambda, \eta)$ are efficient. Similarly, $F(\lambda, \eta)$ denotes the set of payoffs $(v_1, v_2) \in V(\lambda)$ such that $\eta_1 v_1 + \eta_2 v_2 \geq \eta_1 v_1' + \eta_2 v_2'$ for all $(v_1', v_2') \in V(\lambda)$. Also, say that η and η' **determine the same direction** if there is $\xi > 0$ such that $\eta' = \xi \eta$. If true, this implies that $P(\lambda, \eta) = P(\lambda, \eta')$. Finally, given $\mathbf{a} \in A^\infty$ and $t \geq 1$, let

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

When the path $\mathbf{a}$ is clear from the context, we may also write η^t_λ in place of $\eta^t_\lambda(\mathbf{a})$.

The next two results are standard and we omit the proofs.

Lemma G26. *If* $\mathbf{a} \in P(\lambda, \eta)$*, 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 G27. *If* $\mathbf{a} \in P(\lambda, \eta), \mathbf{a}' \in P(\lambda, \eta')$ *and* $\frac{\eta'_i}{\eta'_j} > \frac{\eta_i}{\eta_j}$*, then* $v_{i\lambda}(\mathbf{a}') \geq v_{i\lambda}(\mathbf{a})$ *and* $v_{j\lambda}(\mathbf{a}') \leq v_{j\lambda}(\mathbf{a})$*.*

Let $A^E := \{a \in A : v_1(a) = v_2(a)\}$. For the sake of simplicity, we assume that if $A^E \neq \emptyset$, then $\arg\max_{a \in A^E} v_1(a)$ consists of a single element $a^* \in A^E$. The next two lemmas assume either DMI or IMI.

Lemma G28. *For every* $\mathbf{a} \in P(\lambda, \eta)$*, 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. 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 η. By Lemma G26, ${}_1\mathbf{a} \in P(\lambda, \eta)$. Since $\mathbf{a} = (a^0, {}_1\mathbf{a}) \in P(\lambda, \eta)$ as well, we have $s_\lambda(\mathbf{a}, \eta) = s_\lambda({}_1\mathbf{a}, \eta)$. Since $v_{i\lambda}(\mathbf{a}) = (1-\lambda)g_i(a^0) + \beta_{i\lambda}(a^0)v_{i\lambda}({}_1\mathbf{a})$ and $\beta_{1\lambda}(a^0) = \beta_{2\lambda}(a^0)$, we conclude that

$$s_\lambda(\mathbf{a}, \eta) = s_\lambda({}_1\mathbf{a}, \eta) = \eta_1 \frac{g_1(a^0)}{1 - \beta_1(a^0)} + \eta_2 \frac{g_2(a^0)}{1 - \beta_2(a^0)} = \eta_1 v_1(a^0) + \eta_2 v_2(a^0). \tag{29}$$

Moreover, since ${}_1\mathbf{a} \in P(\lambda, \eta)$, it follows 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, it follows that $v_\lambda({}_1\mathbf{a})$, $v_\lambda(\mathbf{a}), v(a^0) \in F(\lambda, \eta)$, as desired. □

Lemma G29. *Suppose* $A^E \neq \emptyset$*. For every* $\mathbf{a} \in P(\lambda, \eta)$*, 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 $a^t \in A^E \setminus \{a^*\}$ for some t. By Lemma G28, $(a^t, a^t, ...) \in P(\lambda, \eta^t_\lambda)$. However, since $v_i(a^t) < v_i(a^*)$, $i \in I$, 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 $(a^t, a^t, ...)$. □

H Proof of Theorem 5.4

Write v_c for $c(1-\beta(c))^{-1}$ and define v_d and v_b similarly. Since $v_d > v_c > 0 > v_b$, IMI implies $\beta(d) < \beta(c) < \beta(b)$. Given λ, write $\beta_\lambda(d)$ for $\lambda + (1-\lambda)\beta(d)$ etc., and note that $\beta_\lambda(d) < \beta_\lambda(c) < \beta_\lambda(b)$ for all λ. Finally, recall that in the prisoner's dilemma, $\mathbf{a}^{max,1} = (DC, DC, ...)$ and likewise for player 2.

Lemma H30. *For every* $i \in I, \lambda \in [0, 1)$ *and* $\eta \in \mathbb{R}^2_{++}$*, we have* $\mathbf{a}^{max,i} \notin P(\lambda, \eta)$*.*

Proof. Since $\beta_\lambda(d) < \beta_\lambda(b)$, there is $T > 0$ large enough such that i's weight $\frac{\eta_i}{\eta_j}[\frac{\beta_\lambda(d)}{\beta_\lambda(b)}]^T$ is almost zero. But then $\mathbf{a}^{max,i} \notin P(\lambda, \eta_\lambda^T)$ and, by Lemma G26, $\mathbf{a}^{max,i} \notin P(\lambda, \eta)$. $\square$

Lemma H31. *For every $\lambda \in [0,1)$, $\eta \in \mathbb{R}^2_+$, and $\mathbf{a} \in P(\lambda,\eta)$, if $a^0 = CD$ and $a^1 = DC$, then $\mathbf{a}^{\mathcal{A},2} \in P(\lambda,\eta)$. Similarly, if $a^0 = DC$ and $a^1 = CD$, then $\mathbf{a}^{\mathcal{A},1} \in P(\lambda,\eta)$.*

Proof. If $\mathbf{a} \in P(\lambda,\eta)$ is such that $a^0 = CD$ and $a^1 = DC$, then

$$\eta_\lambda^2 = \big(\eta_1\beta_\lambda(b)\beta_\lambda(d), \eta_2\beta_\lambda(d)\beta_\lambda(b)\big).$$

Thus, the direction η_λ^2 is the same as η. By Lemma G26, ${}_2\mathbf{a} \in P(\lambda,\eta)$ and, hence, $s_\lambda(\mathbf{a},\eta) = s_\lambda({}_2\mathbf{a},\eta)$. Deduce that

$$s_\lambda({}_2\mathbf{a},\eta) = \eta_1(1-\lambda)\frac{b+\beta_\lambda(b)d}{1-\beta_\lambda(b)\beta_\lambda(d)} + \eta_2(1-\lambda)\frac{d+\beta_\lambda(d)b}{1-\beta_\lambda(b)\beta_\lambda(d)} = s_\lambda(\mathbf{a}^{\mathcal{A},2},\eta). \quad (30)$$

Hence, $\mathbf{a}^{\mathcal{A},2} \in P(\lambda,\eta)$. $\square$

Lemma H32. *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 $a^0 \ne DC$; if $\frac{\eta_1}{\eta_2} > 1$, then $v_{1\lambda}(\mathbf{a}) \ge v_{2\lambda}(\mathbf{a})$ and $a^0 \ne CD$.*

Proof. 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 $a^0 = DC$. Let $T \ge 1$ be the first period t such that $a^t \ne DC$. Such T exists because $v_{1\lambda}(\mathbf{a}) \le v_{2\lambda}(\mathbf{a})$. Suppose $a^T = CC$. Consider the path $\hat{\mathbf{a}}$ such that $\hat{a}^t = DC$ for all $0 \le t < T$ and $\hat{a}^t = CC$ for all $t \ge T$. From Lemma G28, $\hat{\mathbf{a}} \in P(\lambda,\eta)$. 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 CD. But then, by Lemma H31, $\mathbf{a}^{\mathcal{A},1} \in P(\lambda, \eta_\lambda^{T-1}(\mathbf{a}))$. Also,

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

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

Next, let $\mathbf{a}^C(0) := \mathbf{a}^C$ and for every $T \ge 1$, let $\mathbf{a}^C(T)$ be the path such that $a^t = CD$ for all $0 \le t < T$ and ${}_T\mathbf{a} = \mathbf{a}^C$. Recall that $\eta_{sym} := (1,1)$ and $P_{sym}(\lambda) := P(\lambda, \eta_{sym})$. For simplicity, write $s_\lambda(\mathbf{a})$ instead of $s_\lambda(\mathbf{a},\eta_{sym})$. The next lemma shows that there is no λ such that $\mathbf{a}^C(1) \in P_{sym}(\lambda)$.

Lemma H33. $s_\lambda\big(\mathbf{a}^C(1)\big) < \max\{s_\lambda(\mathbf{a}^C), s_\lambda(\mathbf{a}^{\mathcal{A},2})\}$ *for all $\lambda \in [0,1)$.*

Proof. By construction, $s_\lambda(\mathbf{a}^C) > s_\lambda\big(\mathbf{a}^C(1)\big)$ if and only if

$$v_c > \frac{b+d}{1-\beta(b)+1-\beta(d)}. \quad (31)$$

If (31) holds, the proof is complete. Suppose that (31) holds with equality. Then,

$$s_\lambda\big(\mathbf{a}^C(1)\big) = \frac{2(b+d)}{1-\beta(b)+1-\beta(d)}. \tag{32}$$

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

$$s_\lambda(\mathbf{a}^{\mathcal{A},2}) > \lim_{\lambda'\to 1} s_{\lambda'}(\mathbf{a}^{\mathcal{A},2}) = \frac{2(b+d)}{1-\beta(b)+1-\beta(d)}. \tag{33}$$

Combining (32) and (33) gives $s_\lambda(\mathbf{a}^{\mathcal{A},2}) > s_\lambda\big(\mathbf{a}^C(1)\big)$. Finally, if the strict inequality in (31) is reversed, then

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

The equality follows from direct simplification. The last inequality follows from (33). □

Consider the inequalities

$$1 \leq f(\lambda) := \frac{\frac{c}{1-\beta(c)} - (1-\lambda)\frac{b+\beta_\lambda(b)d}{1-\beta_\lambda(b)\beta_\lambda(d)}}{(1-\lambda)\frac{d+\beta_\lambda(d)b}{1-\beta_\lambda(b)\beta_\lambda(d)} - \frac{c}{1-\beta(c)}} \leq \sqrt{\frac{\beta_\lambda(b)}{\beta_\lambda(d)}}. \tag{34}$$

Lemma H34. *The inequalities in (34) hold if and only if there is $\eta \in \mathbb{R}^2_+$ such that $\mathbf{a}^{\mathcal{A},2}, \mathbf{a}^C \in P(\lambda,\eta)$, that is, if and only if λ is irregular.*

Proof. Suppose $\mathbf{a}^{\mathcal{A},2}, \mathbf{a}^C \in P(\lambda,\eta)$ for some λ and $\eta \in \mathbb{R}^2_+$. If $f(\lambda) < 0$, then either $\mathbf{a}^C$ or $\mathbf{a}^{\mathcal{A},2}$ is strictly Pareto dominated. If $f(\lambda) \in [0,1)$, then $\mathbf{a}^C$ is strictly Pareto dominated by some path in $\mathcal{A}$. Conclude that $f(\lambda) \geq 1$. Turn to the second inequality. Since $\mathbf{a}^{\mathcal{A},2}, \mathbf{a}^C \in P(\lambda,\eta)$, Lemma G26 implies that $\big(CD, DC, \mathbf{a}^C\big) \in P(\lambda,\eta)$. By Lemma G26, the paths $\big(DC, \mathbf{a}^C\big)$ and $\big(DC, \mathbf{a}^{\mathcal{A},2}\big) = \mathbf{a}^{\mathcal{A},1}$ 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 $\big(CD, \mathbf{a}^C\big)$ and the path $\big(CD, DC, \mathbf{a}^{\mathcal{A},2}\big) = \mathbf{a}^{\mathcal{A},2}$ are efficient given the direction

$$\eta' := \big(\eta_2\beta_\lambda(d), \eta_1\beta_\lambda(b)\big). \tag{35}$$

Thus, the path $\mathbf{a}^{\mathcal{A},2}$ is efficient under both η and η'. By the convexity of the feasible set, we have $\frac{\eta_1'}{\eta_2'} \leq \frac{\eta_1}{\eta_2}$. Since $\mathbf{a}^{\mathcal{A},2}, \mathbf{a}^C \in P(\lambda,\eta)$, it must be that

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

Taking this and (35) into account, deduce that $\frac{\eta_1'}{\eta_2'} \leq \frac{\eta_1}{\eta_2}$ is equivalent to the second inequality in (34). Identical arguments show that $\mathbf{a}^{\mathcal{A},2}, \mathbf{a}^C \in P(\lambda,\eta)$ whenever (34) 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 $a^T, a^{T+1} \in \{CD, DC\}$ and $a^T \neq a^{T+1}$. In the rest of this section, we fix a regular λ. The case of irregular λ is deferred until Section L.

Lemma H35. *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 $a^t \neq CC$ for every $t \in \mathcal{T}$.*

Proof. It is w.l.o.g. to assume that $a^T = CD$ and $a^{T+1} = DC$. By Lemma H31, $\mathbf{a}^{\mathcal{A},2} \in P(\lambda, \eta^T_\lambda(\mathbf{a}))$. Assuming $T \geq 1$, we are going to show that $a^t \neq CC$ for every $t < T$. If not, let T' be the greatest integer $k < T$ such that $a^k = CC$. By Lemma G28, we know that $\mathbf{a}^C \in P(\lambda, \eta^{T'}_\lambda(\mathbf{a}))$. The latter is possible only if

$$v_c > v_{1\lambda}(\mathbf{a}^{\mathcal{A},2}). \tag{36}$$

Otherwise, we would have

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

and, hence, $\mathbf{a}^C$ would be strictly Pareto dominated by $\mathbf{a}^{\mathcal{A},2}$. Next, observe that, by construction, $T' \leq T-1$. Suppose first that $T' = T-1$. Since $a^{T'} = CC$,

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

Thus, $\eta^{T'}_\lambda(\mathbf{a})$ and $\eta^T_\lambda(\mathbf{a})$ determine the same direction and so $P(\lambda, \eta^T_\lambda(\mathbf{a})) = P(\lambda, \eta^{T'}_\lambda(\mathbf{a}))$. But then $\mathbf{a}^C, \mathbf{a}^{\mathcal{A},2} \in P(\lambda, \eta^T_\lambda(\mathbf{a}))$, contradicting the regularity of λ. Suppose now that $T' < T-1$. It is w.l.o.g. to assume that $a^t = CD$ 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 $\eta^{T'+2}_\lambda(\mathbf{a})$ satisfies

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

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

$$v_{1\lambda}(\mathbf{a}') \geq v_c \quad \forall \mathbf{a}' \in P(\lambda, \eta^{T'+2}_\lambda(\mathbf{a})). \tag{37}$$

Let $\hat{\mathbf{a}} \in A^\infty$ be a path such that $\hat{a}^t = a^t = CD$ for all $T'+2 \leq t < T$ and ${}_T\hat{\mathbf{a}} = \mathbf{a}^{\mathcal{A},2}$. By Lemma G26, the fact that ${}_{T'+2}\mathbf{a} \in P(\lambda, \eta^{T'+2}_\lambda(\mathbf{a}))$ and $\mathbf{a}^{\mathcal{A},2} \in P(\lambda, \eta^T_\lambda(\mathbf{a}))$ implies that ${}_{T'+2}\hat{\mathbf{a}} \in P(\lambda, \eta^{T'+2}_\lambda(\mathbf{a}))$. We claim that $v_{1\lambda}({}_{T'+2}\hat{\mathbf{a}}) \leq v_{1\lambda}(\mathbf{a}^{\mathcal{A},2}) < v_c$. The first inequality follows since ${}_{T'+2}\hat{\mathbf{a}}$ begins with a repetitive play of the action profile CD, which hurts player 1, followed by the more desirable path $\mathbf{a}^{\mathcal{A},2}$. The second inequality follows from (36). Together, the inequalities contradict (37).

Next, we show that $a^t \neq CC$ for every $t > T+1$. Note that

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

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 $a^{T+2} = CC$. By Lemma G28, $\mathbf{a}^C \in P(\lambda, \eta_\lambda^{T+2}(\mathbf{a}))$. But then, $\mathbf{a}^C, \mathbf{a}^{\mathcal{A},2} \in P(\lambda, \eta_\lambda^T(\mathbf{a}))$, contradicting Lemma H34. Suppose now that $a^k = CC$ for some $k > T+2$. Let T' be the smallest such k. It is w.l.o.g. to assume that $a^t = DC$ 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 $a^{T+2} = CC$ and Lemma G28 would imply that $\mathbf{a}^C \in P(\lambda, \eta_\lambda^{T+2}(\mathbf{a}))$. But then $\mathbf{a}^C, \mathbf{a}^{\mathcal{A},2} \in P(\lambda, \eta_\lambda^T(\mathbf{a}))$, contradicting Lemma H34. Hence, $T' > T+2$. Now, since $a^T = CD$, Lemma H32 shows that $\eta_{1\lambda}^T(\mathbf{a}) \leq \eta_{2\lambda}^T(\mathbf{a})$. And, since $a^{T+2} = DC$, Lemma H32 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, $a^{T+3} = CC$ and, by Lemma G28, we know that $\mathbf{a}^C \in P(\lambda, \eta_\lambda^{T+3}(\mathbf{a}))$. Then, by Lemma G26, $(DC, \mathbf{a}^C) \in P(\lambda, \eta_\lambda^{T+2}(\mathbf{a}))$. But recall that $\eta_\lambda^{T+2}(\mathbf{a})$ and η_{sym} determine the same direction. Thus, $(DC, \mathbf{a}^C) \in P_{sym}(\lambda)$, contradicting Lemma H33. Alternatively, suppose that $T' > T+3$. Then, $a^{T+2} = a^{T+3} = DC$ 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 H32 shows that a^{T+3} cannot be DC, a contradiction. □

Lemma H36. *For every $\eta \in \mathbb{R}^2_{++}$ such that $\eta_1 < \eta_2$ and every path $\mathbf{a} \in P(\lambda, \eta)$, if $a^0 = CC$, then $\mathbf{a} \in \mathcal{C}_2$.*

Proof. 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 $a^t \neq CC$. Since $a^0 = CC$, 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 H32 shows that $a^T \neq DC$. Thus, $a^T = CD$. Next, we are going to show that $a^{T+1} = CC$. By Lemma H30, the constant path $(CD, CD, ...) \notin P(\lambda, \eta_\lambda^T(\mathbf{a}))$. Hence, there exists $t > T$ such that $a^t \neq CD$. Let T' be the smallest such t. By construction, $a^{T'-1} = CD$. Since $a^0 = CC$, Lemma H35 implies that $a^{T'} \neq DC$. Else, $(T'-1, T')$ would be an alternation for a path that contains CC. Conclude that $a^{T'} = CC$. 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 $a^0 = CC$, Lemma G28 shows that

$\mathbf{a}^C \in P(\lambda, \eta)$. Combining the last two observations, conclude that

$$v_{1\lambda}(\mathbf{a}') \geq v_c \quad \forall \mathbf{a}' \in P(\lambda, \eta_\lambda^{T+1}(\mathbf{a})). \tag{38}$$

Recall that $a^{T'} = CC$. By Lemma G28, $\mathbf{a}^C \in P(\lambda, \eta_\lambda^{T'}(\mathbf{a}))$. Define the path $\hat{\mathbf{a}} \in A^\infty$ such that $\hat{a}^t = a^t = CD$ for $T+1 \leq t < T'$ and ${}_{T'}\hat{\mathbf{a}} = \mathbf{a}^C$. Lemma G26 implies that ${}_{T+1}\hat{\mathbf{a}} \in P(\lambda, \eta_\lambda^{T+1}(\mathbf{a}))$. If $T' > T+1$, then $v_{1\lambda}({}_{T+1}\hat{\mathbf{a}}) < v_c$, contradicting (38). Hence, $T' = T+1$, that is, $a^{T+1} = CC$. To summarize, we have shown that for every $\mathbf{a} \in P(\lambda, \eta)$ such that $a^t = CC$ for all $t < T$ and $a^T = CD$, we have $a^{T+1} = CC$.

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 $a^k \neq CC$. 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 G26, ${}_{T''}\mathbf{a} \in P(\lambda, \eta_\lambda^{T''}(\mathbf{a}))$. Thus, ${}_{T''}\mathbf{a} \in P(\lambda, \eta_\lambda^{T+1}(\mathbf{a}))$. But then, by Lemma G26, $\tilde{\mathbf{a}} := (a^0, a^1, ..., a^T, {}_{T''}\mathbf{a}) \in P(\lambda, \eta)$. By construction, $\tilde{\mathbf{a}}$ is such that $\tilde{a}^t = CC$ for all $t < T$, $\tilde{a}^T = CD$, and $\tilde{a}^{T+1} \neq CC$, contradicting the first part of the proof. □

Lemma H37. *For every $\eta \in \mathbb{R}^2_{++}$ and $\mathbf{a} \in P(\lambda, \eta)$, if $a^0 = CD$, and $a^1 = CC$, then ${}_2\mathbf{a} \notin \mathcal{C}_1$.*

Proof. If ${}_2\mathbf{a} \in \mathcal{C}_1$, there is $T > 1$ such that $a^T = DC$ 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 G26, ${}_T\mathbf{a} \in P(\lambda, \eta_\lambda^T(\mathbf{a}))$. Thus, ${}_T\mathbf{a} \in P(\lambda, \eta_\lambda^1(\mathbf{a}))$. But then, by Lemma G26, $\hat{\mathbf{a}} := (CD, {}_T\mathbf{a}) = (CD, DC, \mathbf{a}^C) \in P(\lambda, \eta)$. Thus, $\hat{\mathbf{a}}$ contains an alternation followed by a play of CC, contradicting Lemma H35. □

Lemma H38. *For every $\mathbf{a} \in P_{sym}(\lambda)$, if $a^0 = CD$, then $\mathbf{a} \in \mathcal{A}$.*

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

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

Since ${}_1\mathbf{a} \in P(\lambda, \eta_\lambda^1(\mathbf{a}))$, we can apply Lemma H32 to deduce that $a^1 \in \{CC, DC\}$. If $a^1 = CC$, it follows from Lemma G28 that $\mathbf{a}^C \in P(\lambda, \eta_\lambda^1(\mathbf{a}))$. But then, by Lemma G26, $(CD, \mathbf{a}^C) \in P_{sym}(\lambda)$, contradicting Lemma H33. Thus, $a^1 = DC$. Deduce that $\eta_\lambda^2(\mathbf{a})$ and η determine the same direction and, by Lemma H31, that $\mathbf{a}^{\mathcal{A},2} \in P(\lambda, \eta)$. Since (0,1) is an alternation for the path $\mathbf{a}$, we know from Lemma H35 that $a^t \neq CC$ for all $t > 1$. Thus, $a^2 \in \{CD, DC\}$. By Lemma G26, ${}_2\mathbf{a} \in P(\lambda, \eta_\lambda^2(\mathbf{a}))$. But, since $\eta_\lambda^2(\mathbf{a})$ and η determine the same direction, we have ${}_2\mathbf{a} \in P(\lambda, \eta)$. We also know that $a^2 \in \{CD, DC\}$. Thus, the same arguments that showed that $a^1 = DC$ now show that $a^3 \in \{CD, DC\} \setminus \{a^2\}$. Proceeding like this, conclude that $\mathbf{a} \in \mathcal{A}$. □

Lemma H39. *For every $\eta \in \mathbb{R}^2_{++}$ and $\mathbf{a} \in P(\lambda, \eta)$, if $a^0 = CD$ and $a^1 = DC$, then $\mathbf{a} \in \mathcal{E}_2\mathcal{A}$.*

Proof. Since $a^0 = CD$, it follows from Lemma H32 that $\eta_1 \leq \eta_2$. Suppose $\eta_1 = \eta_2$. We know from Lemma H38 that $\mathbf{a} \in \mathcal{A}$ and, hence, that $\mathbf{a} \in \mathcal{E}_2\mathcal{A}$. Next, suppose $\eta_1 < \eta_2$. Since $a^0 = CD$ and $a^1 = DC$, 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 H35 that $a^t \neq CC$ for all $t > T$. Hence, Lemma H32 implies that $a^2 = CD$. Moreover, since $a^1 = DC$, Lemma H32 shows that $\eta^1_{1\lambda} \geq \eta^1_{2\lambda}$. If $\eta^1_{1\lambda} = \eta^1_{2\lambda}$, we know from Lemma H38 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 $a^1 = DC$ and $a^2 = CD$. 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 $a^t \neq CC$ for all $t > T$. Hence, Lemma H32 implies that $a^3 = DC$. Proceeding like this, we get $\eta^{2t}_{1\lambda} < \eta^{2t}_{2\lambda}$ and $\eta^{2t+1}_{1\lambda} > \eta^{2t+1}_{2\lambda}$ for all $t \in \mathcal{T}$. Lemma H32 implies that $a^{2t} = CD$ and $a^{2t+1} = DC$ for all $t \in \mathcal{T}$. That is, $\mathbf{a} = \mathbf{a}^{\mathcal{A},2} \in \mathcal{E}_2\mathcal{A}$. □

Let $P_{++}(\lambda) := \cup_{\eta \in \mathbb{R}^2_{++}} P(\lambda, \eta)$.

Lemma H40. $P_{++}(\lambda) \subseteq \mathcal{EC} \cup \mathcal{EA}$.

Proof. First, we show that if $\mathbf{a} \in P_{sym}(\lambda)$, then $\mathbf{a} \in \mathcal{A} \cup \{\mathbf{a}^C\}$. By Lemma G29, DD cannot be played along any efficient path. Hence, $a^0 \in \{CC, DC, CD\}$. If $a^0 \in \{CD, DC\}$, Lemma H38 shows that $\mathbf{a} \in \mathcal{A}$. Alternatively, suppose that $a^0 = CC$. By Lemma G28, the path $\mathbf{a}^C$ is efficient. Assume that $\mathbf{a} \neq \mathbf{a}^C$. Let T be the first period t such that $a^t \neq CC$. By construction, for any $t \leq T$, the direction $(\eta^{sym})^t_\lambda$ is the same as η^{sym}. Thus, ${}_T\mathbf{a} \in P(\lambda, \eta)$. W.l.o.g, assume $a^T = CD$. The proof in Lemma H38 shows that $a^{T+1} = DC$. Thus, $(T, T+1)$ is an alternation for $\mathbf{a}$. Since $a^0 = CC$, Lemma H35 is contradicted. Conclude that $P(\lambda, \eta^{sym}) = \{\mathbf{a}^C\}$.

Next, take any $\mathbf{a} \in P(\lambda, \eta)$ where $0 < \eta_1 < \eta_2$. Since $\eta_1 < \eta_2$, Lemma H32 shows that $v_{1\lambda}(\mathbf{a}) \leq v_{2\lambda}(\mathbf{a})$ and $a^0 \neq DC$. By Lemma G29, DD cannot be played along any efficient path. Hence, $a^0 \in \{CC, CD\}$. If $a^0 = CC$, Lemma H36 shows that ${}_1\mathbf{a} \in \mathcal{C}_2$ and, hence, $\mathbf{a} \in \mathcal{C}_2$. Alternatively, suppose $a^0 = CD$. By Lemma H30, the constant path $(CD, CD, \ldots)$ is not efficient. Let T be the first period t such that $a^t \neq CD$. Suppose $a^T = CC$. If $\eta^T_{1\lambda} < \eta^T_{2\lambda}$, then Lemma H36 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 H36 implies that ${}_T\mathbf{a} \in \mathcal{C}_1$. Moreover, Lemma H37 implies that $a^t \neq DC$ for all $t > T$. Therefore, ${}_T\mathbf{a} = \mathbf{a}^C$. Finally, suppose $a^T = DC$. Lemma H39 shows that ${}_{T-1}\mathbf{a} \in \mathcal{E}_2\mathcal{A}$ and, hence, $\mathbf{a} \in \mathcal{E}_2\mathcal{A}$. □

Lemma H41. *If $s_\lambda(\mathbf{a}^C) > s_\lambda(\mathbf{a}^{\mathcal{A},2})$, then $P_{++}(\lambda) \supseteq \mathcal{EC}$. Else, $P_{++}(\lambda) \supseteq \mathcal{EA}$.*

Proof. We prove that $\mathcal{E}_2\mathcal{C}_2 \subseteq P_{++}(\lambda)$. Everything else follows from analogous arguments. Recall the paths $\mathbf{a}^C(T)$, $T \in \mathcal{T}$, defined prior to Lemma H33. Note that $\mathbf{a}^C(T) \in \mathcal{E}_2\mathcal{C}_2$ for every T. Let $\eta(0) := (1,1)$ and, for every $T \geq 1$,

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

First, 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 H40, 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 H40, we know that $P(\lambda, \eta(T+1)) \subseteq \mathcal{EC}$. It is therefore enough to show that

$$s_\lambda\big(\mathbf{a}^C(T+1), \eta(T+1)\big) \geq s_\lambda\big(\mathbf{a}, \eta(T+1)\big) \quad \forall \mathbf{a} \in \mathcal{EC}. \tag{39}$$

Observe that, by construction, $\frac{\eta_1(T+1)}{\eta_2(T+1)} < 1$. Lemma H32 implies that $v_{1\lambda}(\mathbf{a}) \leq v_{2\lambda}(\mathbf{a})$ for all $\mathbf{a} \in P\big(\lambda, \eta(T+1)\big)$. Hence, it is enough to show that the inequality in (39) 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 (39) 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 (39) holds since, by the definition of $\eta(T+1)$, we have

$$s_\lambda\big(\mathbf{a}^C(T+1), \eta(T+1)\big) = s_\lambda\big(\mathbf{a}^C(T), \eta(T+1)\big). \tag{40}$$

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

$$\frac{\eta_2(T+1)}{\eta_1(T+1)} > \frac{v_{1\lambda}(\mathbf{a}^C(T')) - v_{1\lambda}(\mathbf{a}^C(T))}{v_{2\lambda}(\mathbf{a}^C(T)) - v_{2\lambda}(\mathbf{a}^C(T'))}. \tag{41}$$

Combining (40) and (41) 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 (39) 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 CD is played in all periods $t < T^*$, CD is played in period T^{**} as well, and CC is played in all other periods. Letting $\varrho := 1 - [\beta_\lambda(c)]^{T^{**}-T^*}$, observe that

$$v_\lambda(\mathbf{a}) = \varrho v_\lambda(\mathbf{a}^C(T^*)) + (1-\varrho) v_\lambda(\mathbf{a}^C(T^*+1)).$$

Conclude that (39) 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. It remains 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, $v(\mathbf{a})$ is a convex combination of $v(\mathbf{a}^C(T))$ and $v(\mathbf{a}^C(T+1))$ for some T. Since $\mathbf{a}^C(T), \mathbf{a}^C(T+1) \in P(\lambda, \eta(T+1))$, we see that $\mathbf{a} \in P(\lambda, \eta(T+1))$. □

I Proof of Corollary 5.2

Since $(1-\lambda)b + \beta_\lambda(b)v_c$ is increasing in λ, we have $(1-\lambda)d < (1-\lambda)b + \beta_\lambda(b)v_c$ for any $\lambda > \underline{\lambda}$. Take any ε' such that

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

Since $\mathbf{a} \in IR^{\varepsilon}(\lambda)$ and $\varepsilon' \leq \varepsilon$, we have $\mathbf{a} \in IR^{\varepsilon'}(\lambda)$. Corollary 5.1 implies that $\mathbf{a} \in SIR^{\varepsilon'}(\lambda)$. To support $\mathbf{a}$ in a SPE, consider the following grim trigger 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 DD 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 DD after any history is a SPE, no player wants to deviate in phase (B) either.

J Proof of Theorem 6.1

As in Section G, let $A^E := \{a \in A : v_1(a) = v_2(a)\}$. For simplicity, we continue to assume that if $A^E \neq \emptyset$, then $\arg\max_{a \in A^E} v_1(a)$ consists of a single element $a^* \in A^E$. To state the next four lemmas, fix $\lambda \in [0,1)$, $\eta \in \mathbb{R}^2_{++}$, and $\mathbf{a} \in P(\lambda, \eta)$.

Lemma J42. *If $\beta_{1\lambda}(a^0) > \beta_{2\lambda}(a^0)$, then $v_{1\lambda}(\mathbf{a}) > v_{2\lambda}(\mathbf{a})$.*

Proof. Since $\beta_{1\lambda}(a^0) > \beta_{2\lambda}(a^0)$, $\frac{\eta^1_{1\lambda}}{\eta^1_{2\lambda}} > \frac{\eta_1}{\eta_2}$ and, since ${}_1\mathbf{a} \in P(\lambda, \eta^1_\lambda)$,

$$v_{2\lambda}({}_1\mathbf{a}) \leq v_{2\lambda}(\mathbf{a}) \quad \text{and} \quad v_{1\lambda}({}_1\mathbf{a}) \geq v_{1\lambda}(\mathbf{a}). \tag{42}$$

From (9), 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 $i \in I$. Thus, the inequalities in (42) are possible only if $v_{2\lambda}(\mathbf{a}) \leq v_2(a^0)$ and $v_1(a^0) \leq v_{1\lambda}(\mathbf{a})$. By Lemma C20, $\beta_{2\lambda}(a^0) < \beta_{1\lambda}(a^0)$ implies $v_2(a^0) < v_1(a^0)$. Hence, $v_{2\lambda}(\mathbf{a}) < v_{1\lambda}(\mathbf{a})$. □

Lemma J43. *If $v_{1\lambda}(\mathbf{a}) = v_{2\lambda}(\mathbf{a})$, then $\mathbf{a} = (a^*, a^*, ...)$.*

Proof. By Lemma J42, $\beta_{1\lambda}(a^0) = \beta_{2\lambda}(a^0)$ and, hence, $a^0 \in A^E$ by Lemma C20. It follows that $v_{1\lambda}({}_1\mathbf{a}) = v_{2\lambda}({}_1\mathbf{a})$. Since, ${}_1\mathbf{a} \in P(\lambda, \eta^1_\lambda)$, the exact same argument shows that $a^1 \in A^E$ and, inductively, that $a^t \in A^E$ for every t. By Lemma G29, $\mathbf{a} = (a^*, a^*, ...)$. □

The proof of the next lemma follows from similar arguments and is omitted.

Lemma J44. *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 $v_{2\lambda}({}_1\mathbf{a}) > v_{2\lambda}(\mathbf{a})$.*

Lemma J45. *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 t such that $\beta_{1\lambda}(a^t) \geq \beta_{2\lambda}(a^t)$ and 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 J42 implies that $v_{1\lambda}(\mathbf{a}) < v_{2\lambda}(\mathbf{a})$. Conclude that

$$v_{1\lambda}(\hat{\mathbf{a}}) < v_{2\lambda}(\hat{\mathbf{a}}) \quad \forall \hat{\mathbf{a}} \in P(\lambda, \eta^T_\lambda(\mathbf{a})). \tag{43}$$

By Lemma G26, ${}_T\mathbf{a} \in P(\lambda, \eta^T_\lambda(\mathbf{a}))$ and, hence, $v_{1\lambda}({}_T\mathbf{a}) < v_{2\lambda}({}_T\mathbf{a})$. By Lemma J42, $\beta_{1\lambda}(a^T) \leq \beta_{2\lambda}(a^T)$. By the choice of T, it must be that $\beta_{1\lambda}(a^T) = \beta_{2\lambda}(a^T)$. By Lemma C20, $v_1(a^T) = v_2(a^T)$ so that $a^T \in A^E$. It follows from Lemmas G28 and G29 that $\mathbf{a}' := (a^*, a^*, ...) \in P(\lambda, \eta^T_\lambda(\mathbf{a}))$. But then, $v_{1\lambda}(\mathbf{a}') = v_{2\lambda}(\mathbf{a}')$, contradicting (43). □

We can now complete the proof of Theorem 6.1. For simplicity, assume that for each i, the path $\mathbf{a}^{max,i} \in A^\infty$ that attains i's maximum payoff is unique. If $\mathbf{a}^{max,1} = \mathbf{a}^{max,2}$, then $P(\lambda) = \{\mathbf{a}^{max,1}\} = \{\mathbf{a}^{max,2}\}$ for all $\lambda \in [0, 1)$. From now on, assume $\mathbf{a}^{max,1} \neq \mathbf{a}^{max,2}$. 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 $i \in I$ and $j \neq i$, then $\mathbf{a} = \mathbf{a}^{max,j}$. Thus, assume $\eta \in \mathbb{R}^2_{++}$. If $v_{1\lambda}(\mathbf{a}) = v_{2\lambda}(\mathbf{a})$, then Lemma J43 shows that $\mathbf{a} = (a^*, a^*, ...)$, as desired. Assume $v_{1\lambda}(\mathbf{a}) < v_{2\lambda}(\mathbf{a})$. By Lemma J42, $\beta_{1\lambda}(a^0) \leq \beta_{2\lambda}(a^0)$. We claim that there is 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 J45 shows that $\beta_{1\lambda}(a^t) < \beta_{2\lambda}(a^t)$ for all $t > 0$, as desired. Assume $\beta_{1\lambda}(a^0) = \beta_{2\lambda}(a^0)$ and let $T \geq 1$ be the first period t such that $\beta_{1\lambda}(a^t) \neq \beta_{2\lambda}(a^t)$. By Lemma C20, such T exists since $v_{1\lambda}(\mathbf{a}) < v_{2\lambda}(\mathbf{a})$. By construction, $\beta_{1\lambda}(a^t) = \beta_{2\lambda}(a^t)$ for every $0 \leq t < T$. Lemma G29 implies that $a^t = a^*$ for every $0 \leq t < T$. Since $a^0 = a^*$, Lemma J44 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, $v_{1\lambda}(\mathbf{a}) < v_{2\lambda}(\mathbf{a})$, conclude that $v_{1\lambda}({}_1\mathbf{a}) < v_{2\lambda}({}_1\mathbf{a})$. Applying Lemma J44 repeatedly, conclude that $v_{1\lambda}({}_t\mathbf{a}) < v_{2\lambda}({}_t\mathbf{a})$ for every $t \leq T$. By Lemma J42, $\beta_{1\lambda}(a^T) \leq \beta_{2\lambda}(a^T)$ and, by the choice of T, $\beta_{1\lambda}(a^T) < \beta_{2\lambda}(a^T)$. By Lemma J45, $\beta_{1\lambda}(a^t) < \beta_{2\lambda}(a^t)$ for all $t > T$.

Finally, let $B := \{a \in A : \beta_{1\lambda}(a) < \beta_{2\lambda}(a)\}$ and $l := \min_{a\in B} \frac{\beta_{2\lambda}(a)}{\beta_{1\lambda}(a)}$. 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 $l > 1$, $l^{t-T} \to +\infty$ 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.

K Proof of Theorem 6.3

In Appendix E, we constructed a SIR payoff v using a random mixed path (RMP) and highlighted that it is an open question whether v could be generated by a single, SIR path $(\alpha_0, \alpha_1, ...) \in (\Delta A)^\infty$. Accordingly, we want the bounds posited by Theorem 6.3 to cover payoffs from RMPs and, in fact, from any SIR strategy σ. Fixing λ, we thus let $\widehat{SIR}^0(\lambda)$ be the set of all SIR strategies $\sigma \in \Sigma$ given λ. Similarly, we let $\hat{P}(\lambda, \eta)$ be the set of strategies $\sigma \in \widehat{SIR}^0(\lambda)$ that maximize $\eta_1 v_1 + \eta_2 v_2$ and let $\hat{P}(\lambda) := \cup_{\eta \in \mathbb{R}^2_+} \hat{P}(\lambda, \eta)$. Of special interest would be strategies σ' in which play in periods $t \geq 1$ is independent of what transpired in period $t = 0$. Such strategies can be written in the form (α, σ), where $\alpha \in \Delta(A)$ is the action played in period $t = 0$, strategy σ is played starting from period $t = 1$. Given the independence of period-0 history, we get

$$v_{i\lambda}(\alpha, \sigma) = (1 - \beta_{i\lambda}(\alpha))v_i(\alpha) + \beta_{i\lambda}(\alpha)v_{i\lambda}(\sigma),$$

which is an analogue of (9). We also have the following analogue of Lemma G26:

Lemma K46. *If $(\alpha, \sigma) \in \hat{P}(\lambda, \eta)$, then $\sigma \in \hat{P}(\lambda, \eta')$ where $\eta'_i = \beta_{i\lambda}(\alpha)\eta_i$.*

The next lemma shows that in the prisoner's dilemma, the conclusions of Lemma C20 extend to mixed actions.

Lemma K47. *Under DMI, $\beta_i(\alpha) \geq \beta_j(\alpha)$ if and only if $v_i(\alpha) \geq v_j(\alpha)$.*

Proof. Recall that $d > b$ and, under DMI, $\beta(d) > \beta(b)$. Hence, $\beta_1(\alpha) \geq \beta_2(\alpha)$ if and only if $\alpha(DC) \geq \alpha(CD)$ if and only if $g_1(\alpha) \geq g_2(\alpha)$. Combining these equivalences gives the desired conclusion. □

Lemma K48. *If $(\alpha, \sigma) \in \hat{P}(\lambda)$, then $\alpha^{con} \in \widehat{SIR}^0(\lambda)$. In addition, if $\beta_1(\alpha) = \beta_2(\alpha)$, then $\alpha = CC$.*

Proof. If $\beta_i(\alpha) = \beta_j(\alpha)$, then $v_\lambda(\alpha, \sigma)$ is a convex combination of $v(\alpha)$ and $v_\lambda(\sigma)$. Moreover, by Lemma K47, $v_i(\alpha) = v_j(\alpha)$. Thus, $(v_c, v_c) \geq v(\alpha)$ and, unless $\alpha = CC$, some convex combination of (v_c, v_c) and $v_\lambda(\sigma)$ would strictly Pareto dominate $v_\lambda(\alpha, \sigma)$. Suppose that $\beta_1(\alpha) < \beta_2(\alpha)$. If $(\alpha, \sigma) \in \hat{P}(\lambda, \eta)$, we know from Lemma K46 that $\sigma \in \hat{P}(\lambda, \eta')$ where $\eta'_i = \beta_{i\lambda}(\alpha)\eta_i$. Since $\frac{\eta'_1}{\eta'_2} < \frac{\eta_1}{\eta_2}$, we must have $v_{1\lambda}(\sigma) \leq v_{1\lambda}(\alpha, \sigma)$. This implies that $v_1(\alpha) \geq v_{1\lambda}(\alpha, \sigma) \geq 0$. Also, by Lemma K47, $v_2(\alpha) > v_1(\alpha) \geq 0$. □

Lemma K49. $\max\{v_2(\alpha) : \alpha^{con} \in \widehat{SIR}^0(\lambda)\} = v_2(\alpha^{0,2}) > v_c$.

Proof. Let $\varrho^* \in (0,1)$ be such that $v_1(\varrho^* CD + (1-\varrho^*)DC) = 0$. Direct calculations show that

$$v_2(\alpha^{0,2}) = c(d-b)(c(1-\beta(d)) - b(1-\beta(c)))^{-1}$$
$$v_2(\varrho^* CD + (1-\varrho^*)DC) = (d^2-b^2)(d(1-\beta(d) - b(1-\beta(b))^{-1}$$

It is immediate that $v_2(\alpha^{0,2}) > v_c$. To show the other assertion, it is enough to show that $v_2(\alpha^{0,2}) > v_2(\varrho^* CD + (1-\varrho^*)DC)$. Using DMI and the above expressions, the latter inequality can be reduced to (6). □

Given strategies $\sigma, \sigma' \in \Sigma$ and $\varrho \in [0,1]$, let $\varrho\sigma + (1-\varrho)\sigma'$ be the strategy in which the period-0 public randomization device determines whether the players follow σ or σ', with the probability of the former being ϱ. The construction is analogous to that of an RMP in Appendix E. As in that context, note that

$$v_{i\lambda}(\varrho\sigma + (1-\varrho)\sigma') = \varrho v_{i\lambda}(\sigma) + (1-\varrho)v_{i\lambda}(\sigma').$$

Also, if $\sigma, \sigma' \in \widehat{SIR}^0(\lambda)$, then $\varrho\sigma + (1-\varrho)\sigma' \in \widehat{SIR}^0(\lambda)$. Finally, any strategy $\hat{\sigma}$ can be expressed as a distribution over strategies of the form (α, σ). To state the next lemma, let $v^*_{2\lambda} := \max\{v_{2\lambda}(\sigma) : \sigma \in \widehat{SIR}^0(\lambda)\}$.

Lemma K50. *If $\sigma \in \widehat{SIR}^0(\lambda)$ is such that $v_{2\lambda}(\sigma) = v^*_{2\lambda}$, then $v_{1\lambda}(\sigma) = 0$.*

Proof. If $v_{1\lambda}(\sigma) > 0$, then for some $\varrho \in (0,1)$, the strategy $\varrho(CD, \sigma) + (1-\varrho)\sigma$ would belong to $\widehat{SIR}^0(\lambda)$ and $v_{2\lambda}(\varrho(CD,\sigma) + (1-\varrho)\sigma) > v_{2\lambda}(\sigma)$, contradicting $v_{2\lambda}(\sigma) = v^*_{2\lambda}$. □

Recall that F^0_2 is the linear segment connecting $v(\alpha^{0,2})$ and (v_c, v_c). Let R be the ray originating at (v_c, v_c) and passing through $v(\alpha^{0,2})$, and let $F_2(\lambda) := \{v_\lambda(\sigma) : \sigma \in \hat{P}(\lambda)$ and $v_{2\lambda}(\sigma) \geq v_{1\lambda}(\sigma)\}$. The next two lemmas collect several facts about the geometry of the feasible set (under DMI). The simple, but tedious, proofs are omitted.

Lemma K51. $v_\lambda(CD, \alpha^{0,2}) \in R$. *Also, if $\sigma \in \Sigma$ is such that $v_{1\lambda}(\sigma) = 0$ and $v_\lambda(\sigma) = \varrho v_\lambda(CD, \sigma) + (1-\varrho)(v_c, v_c)$ for some $\varrho \in (0,1)$, then $v_{2\lambda}(\sigma) = v_2(\alpha^{0,2})$. Finally, if $v_{1\lambda}(\sigma) = 0$ and $v_{2\lambda}(\sigma) > v_2(\alpha^{0,2})$, then $v_\lambda(CD, \sigma)$ lies strictly below the ray originating from (v_c, v_c) and passing through $v_\lambda(\sigma)$.*

Lemma K52. *If (α, σ) is such that $v_i(\alpha) < v_j(\alpha)$ and $v_{i\lambda}(\sigma) > v_{j\lambda}(\sigma)$, then $v_\lambda(\alpha, \sigma)$ lies strictly below the straight line passing through $v(\alpha)$ and $v_\lambda(\sigma)$.*

Lemma K53. *If* $v^*_{2\lambda} = v_{2\lambda}(\alpha,\sigma)$ *for some strategy* $(\alpha,\sigma) \in \widehat{SIR}^0(\lambda)$*, then* $v^*_{2\lambda} = v_2(\alpha^{0,2})$*. Moreover, if* $v^*_{2\lambda} = v_2(\alpha^{0,2})$*, then* $F_2(\lambda) = F^0_2$*.*

Proof. Lemmas K48 and K49 show that $v^*_{2\lambda} = v_2(\alpha^{0,2})$. Assuming the latter, suppose that for some $\sigma \in \widehat{SIR}^0(\lambda)$, $v_\lambda(\sigma) \in F_2(\lambda) \setminus F^0_2$. Since $v_\lambda(CD,\alpha^{0,2}) \in R$, there is $\varrho \in (0,1)$ such that $\varrho(CD,\alpha^{0,2}) + (1-\varrho)\sigma \in \widehat{SIR}^0(\lambda)$ and $v_{2\lambda}(\varrho(CD,\alpha^{0,2}) + (1-\varrho)\sigma) > v_2(\alpha^{0,2}) = v^*_{2\lambda}$, a contradiction. □

If $v^*_{2\lambda}$ cannot be attained by a strategy of the form (α,σ), then it must be attainable by a strategy $\hat{\sigma} = \varrho(\alpha',\sigma') + (1-\varrho)(\alpha,\sigma)$ where $v_{1\lambda}(\alpha',\sigma') < 0 < v_{1\lambda}(\alpha,\sigma)$ and $\sigma', (\alpha,\sigma), \sigma \in \widehat{SIR}^0(\lambda)$. Also, there must be some direction $\eta \in \mathbb{R}^2_+$ such that $\hat{\sigma}, (\alpha,\sigma) \in \hat{P}(\lambda,\eta)$. Let $L(\eta)$ be the linear segment connecting $v_\lambda(\hat{\sigma})$ and $v_\lambda(\alpha,\sigma)$, and note that $L(\eta)$ is a part of the frontier $F_2(\lambda)$ that is orthogonal to η. By Lemma K46, $\sigma \in \hat{P}(\lambda,\eta')$, where $\eta'_i = \beta_{i\lambda}(\alpha)\eta_i$.

Case 1: Suppose $v_2(\alpha) > v_1(\alpha)$. By Lemma K47, $\beta_2(\alpha) > \beta_1(\alpha)$ and $\frac{\eta'_1}{\eta'_2} < \frac{\eta_1}{\eta_2}$. It follows that $v_\lambda(\sigma) = v_\lambda(\hat{\sigma})$ and w.l.o.g. that we can express $\hat{\sigma}$ as $\varrho(\alpha',\sigma') + (1-\varrho)(\alpha,\hat{\sigma})$. Note that $\alpha(DD) = 0$. Otherwise, replacing DD with CC in α would lead to a strict Pareto improvement, contradicting the fact that $v_{2\lambda}(\hat{\sigma}) = v^*_{2\lambda}$. With this in mind, observe that

$$v_\lambda(\alpha,\hat{\sigma}) = \alpha(CD)v_\lambda(CD,\hat{\sigma}) + \alpha(CC)v_\lambda(CC,\hat{\sigma}) + \alpha(DC)v_\lambda(DC,\hat{\sigma}). \tag{44}$$

Let L be the line passing through (v_c, v_c) and $v_\lambda(\hat{\sigma})$. By construction, $v_\lambda(CC,\hat{\sigma}) \in L$. By Lemma K52, $v_\lambda(DC,\hat{\sigma})$ is below the line connecting $v(DC)$ and $v_\lambda(\hat{\sigma})$, and hence, below L. Finally, by Lemma K51, $v_\lambda(CD,\hat{\sigma})$ is on or below L. Moreover, $v_\lambda(CD,\hat{\sigma}) \in L$ if and only if $v_{2\lambda}(\hat{\sigma}) = v_2(\alpha^{0,2})$. Putting everything together, we see from (44) that $v_\lambda(\alpha,\hat{\sigma}) \in \hat{P}(\lambda)$ is possible only if $v_{2\lambda}(\hat{\sigma}) = v_2(\alpha^{0,2})$ and $\alpha(DC) = 0$. By Lemma K53, $F_2(\lambda) = F^0_2$.

Case 2: Suppose $v_2(\alpha) = v_1(\alpha)$, which implies that $\alpha = CC$. We claim that the frontier $F_2(\lambda)$ is linear. If $v_\lambda(\sigma) = (v_c, v_c)$, then $v_\lambda(\alpha,\sigma) = (v_c, v_c)$ and the claim follows. If $v_{1\lambda}(\sigma) < v_c$, then $v_{1\lambda}(\sigma) < v_{1\lambda}(CC,\sigma) = v_{1\lambda}(\alpha,\sigma)$. Thus, $v_\lambda(\sigma)$ belongs to the linear segment $L(\eta)$ connecting $v_\lambda(\hat{\sigma})$ and $v_\lambda(\alpha,\sigma)$. But since $\alpha = CC$, $v_\lambda(\alpha,\sigma)$ lies on a linear segment L' connecting (v_c, v_c) and $v_\lambda(\sigma)$. Putting everything together, we see that $v_\lambda(\sigma) \neq v_\lambda(CC,\sigma)$ and $v_\lambda(\sigma), v_\lambda(CC,\sigma) \in L(\eta) \cap L'$. This implies that $L' \subset L(\eta)$ and, hence, that the frontier $F_2(\lambda)$ is a single linear segment connecting $v_\lambda(\hat{\sigma})$ with (v_c, v_c).

It remains to show that $v_\lambda(\hat{\sigma}) = v(\alpha^{0,2})$. Since $F_2(\lambda)$ is linear and since $v_\lambda(\hat{\sigma})$ is a convex combination of $v_\lambda(\alpha',\sigma')$ and the point $v_\lambda(\alpha,\sigma) \in F_2(\lambda)$, $v_\lambda(\alpha',\sigma')$ must lie on the line L'' defined by $F_2(\lambda)$. As before,

$$v_\lambda(\alpha',\sigma') = \alpha'(CC)v_\lambda(CC,\sigma') + \alpha'(CD)v_\lambda(CD,\sigma') + \alpha'(DC)v_\lambda(DC,\sigma'). \tag{45}$$

We know that $\sigma' \in \hat{P}(\lambda)$.

Case 2.1: Suppose $v_\lambda(\sigma') = \varrho v_\lambda(\hat{\sigma}) + (1-\varrho)(v_c, v_c)$ for some $\varrho \in [0,1]$. Then, we have $v_\lambda(CC, \sigma') \in F_2(\lambda)$. By Lemma K52, $v_\lambda(DC, \sigma')$ is below the line connecting $v(DC)$ and $v_\lambda(\sigma')$, and, hence, below L''. By construction, it is also the case that

$$v_\lambda(CD, \sigma') = \varrho v_\lambda(CD, \hat{\sigma}) + (1-\varrho) v_\lambda(CD, \mathbf{a}^C).$$

By Lemma K51, $v_\lambda(CD, \hat{\sigma})$ is on or below L''. Moreover, $v_\lambda(CD, \hat{\sigma}) \in L''$ if and only if $v_\lambda(\hat{\sigma}) = v(\alpha^{0,2})$. Direct verification shows that $v_\lambda(CD, \mathbf{a}^C)$ is below the line connecting (v_c, v_c) and $v(\alpha^{0,2})$. Summarizing, we see from (45) that $v_\lambda(\alpha', \sigma') \in L''$ if and only if $v_\lambda(\hat{\sigma}) = v(\alpha^{0,2})$.

Case 2.2: Let $\check{\sigma}$ be the symmetric analogue of $\hat{\sigma}$ so that, in particular, $v_{i\lambda}(\check{\sigma}) = v_{j\lambda}(\hat{\sigma})$. Suppose $v_\lambda(\sigma') = \varrho v_\lambda(\check{\sigma}) + (1-\varrho) v_c$ for some $\varrho \in [0,1)$. We are going to obtain a contradiction, which would effectively complete the proof of Theorem 6.3. Recall that, by definition, $v^*_{2\lambda} = v_{2\lambda}(\hat{\sigma})$ and note that $v_\lambda(CD, \hat{\sigma}) = (b, d + \beta(d) v^*_{2\lambda})$ and $v_\lambda(CD, \check{\sigma}) = (b + \beta(b) v^*_{2\lambda}, d)$. Since $v_{1\lambda}(\alpha', \sigma') < 0$, deduce from (45) that $v_{1\lambda}(CD, \sigma') < 0$ and, hence, that $v_{1\lambda}(CD, \check{\sigma}) = b + \beta(b) v^*_{2\lambda} < 0$.

Next, let $\hat{R}$ be the ray originating at (v_c, v_c) and passing through $v_\lambda(CD, \hat{\sigma})$ and let $\check{R}$ be the ray originating at (v_c, v_c) and passing through $v_\lambda(CD, \check{\sigma})$. The slopes of these rays are respectively:

$$\check{S} := (d - v_c)/(b + \beta(b) v^*_{2\lambda} - v_c) \quad \text{and} \quad \hat{S} := (d + \beta(d) v^*_{2\lambda} - v_c)/(b - v_c)$$

Since $v_\lambda(\sigma') = \varrho v_\lambda(\check{\sigma}) + (1-\varrho)(v_c, v_c)$, deduce that $|\check{S}| \geq |\hat{S}|$ and that $v_{2\lambda}(CD, \check{\sigma}) = d > v_c$. The latter, together with the fact that $v^*_{2\lambda} > v_c$ and $b + \beta(b) v^*_{2\lambda} < 0$, implies that the inequality $|\check{S}| \geq |\hat{S}|$ is equivalent to

$$v^*_{2\lambda} \geq \frac{1}{\beta(b)\beta(d)} [v_c \beta(b) + v_c \beta(d) - \beta(b) d - \beta(d) b]. \tag{46}$$

Since $\mathbf{a}^C$ is a first-best outcome, we know that $2v_c \geq v^*_{2\lambda}$. Combining the latter with (46) and simplifying gives

$$\frac{v_d - v_c}{v_c - v_b} \geq \frac{\beta(d)}{\beta(b)}.$$

It follows from DMI that $(v_d - v_c)/(v_c - v_b) > 1$, but this contradicts (6).

L When Both Intra- and Inter-temporal Cooperation are Efficient

This section describes the Pareto set for irregular λ, i.e., such that $\mathbf{a}^{\mathcal{A},2}, \mathbf{a}^C \in P(\lambda)$. Let $\mathcal{CA}_1$ be the set of paths $\mathbf{a} \in A^\infty$ such that $a^t \in \{CC, DC, CD\}$ for every t and, in addition, $a^t = DC$ for some t if and only if $a^{t+1} = CD$. If we let $X := CC$ and $Y = (DC, CD)$, we can identify $\mathcal{CA}_1$ with the Cartesian product $\{X, Y\}^\infty$. Define $\mathcal{CA}_2$ analogously, with the difference that CD is played in period t if and only if DC is played in period $t+1$. Note that $\mathbf{a}^C, \mathbf{a}^{\mathcal{A},i} \in \mathcal{CA}_i$, $i \in I$. Let $\mathcal{CA} := \mathcal{CA}_1 \cup \mathcal{CA}_2$. Define $\mathcal{ECA}_1$ to be the set of paths such that DC 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 $X := CC$, $Y := (DC, CD)$, and $Z := (CD, DC)$. 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 $\mathbf{a} \in \mathcal{CA}^M$, $\mathbf{a} \in P_{sym}(\lambda)$ if and only if $s_\lambda(\mathbf{a}^{\mathcal{A},2}) = s_\lambda(\mathbf{a}^C)$. Define $\mathcal{ECA}_1^M$ to be the set of paths such that DC 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 parallels that of Theorem 5.4 and is omitted.

Theorem L1. *Suppose (34) holds. If the first inequality is strict,* $P(\lambda) = \mathcal{ECA} \cup \{\mathbf{a}^{max,1}, \mathbf{a}^{max,2}\}$*; else,* $P(\lambda) = \mathcal{ECA}^M \cup \{\mathbf{a}^{max,1}, \mathbf{a}^{max,2}\}$.

References

[1] D. Abreu, P. K. Dutta, and L. Smith. The folk theorem for repeated games: a NEU condition. *Econometrica*, 62(4):939–948, 1994.

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

[3] S. Andersen, G. W. Harrison, M. I. Lau, and E. E. Rutström. Multiattribute utility theory, intertemporal utility, and correlation aversion. *International Economic Review*, 59(2):537–555, 2018.

[4] 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.

[5] A. Bommier, A. Kochov, and F. Le Grand. Ambiguity and endogenous discounting. *Journal of Mathematical Economics*, 83:48–62, 2019.

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

[7] S. H. Chew. A generalization of the quasilinear mean with applications to the measurement of income inequality and decision theory resolving the Allais paradox. *Econometrica*, 51(4): 1065–1092, 1983.

[8] S. H. Chew and L. G. Epstein. Recursive utility under uncertainty. In A. M. Khan and N. C. Yannelis, editors, *Equilibrium Theory in Infinite Dimensional Spaces*, volume 1 of *Studies in Economic Theory*, pages 352–369. Springer Berlin Heidelberg, 1991.

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

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

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

[12] D. Fudenberg and E. Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3):533–554, 1986.

[13] A. Kochov. A separation result for stationary preferences. *Journal of Mathematical Economics*, 70:123–126, 2017.

[14] T. Koopmans. Stationary ordinal utility and impatience. *Econometrica*, 28(2):287–309, 1960.

[15] D. Laibson. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112 (2):443–478, 1997.

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

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

[18] B. Miao and S. Zhong. Risk preferences are not time preferences: separating risk and time preference: comment. *American Economic Review*, 105(7):2272–86, 2015.

[19] R. B. Myerson. *Game theory*. Harvard University Press, 1991.

[20] W. S. Neilson and H. Winter. Infinitely-repeated games with endogenous discounting. *Economics letters*, 52(2):163–169, 1996.

[21] I. Obara and J. Park. Repeated games with general discounting. *Journal of Economic Theory*, 172:348–375, 2017.

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

[23] T. Rader. Utility over time: the homothetic case. *Journal of Economic Theory*, 25:219–236, 1981.

[24] F. Ramsey. A mathematical theory of saving. *The Economic Journal*, 38(152):543–559, 1928.

[25] T. Sekiguchi and K. Wakai. Repeated games with recursive utility: Cournot duopoly under gain/loss asymmetry. 2016.

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

[27] G. Toth. Asymmetry of convex sets with isolated extreme points. *Proceedings of the American Mathematical Society*, 137(1):287–295, 2009.

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