Intertemporal Hedging and Trade in Repeated Games with Recursive Utility

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

Discussion Paper No. 361

December 28, 2022

Collaborative Research Center Transregio 190 | www.rationality-and-competition.de
Ludwig-Maximilians-Universität München | Humboldt-Universität zu Berlin
Spokesperson: Prof. Georg Weizsäcker, Ph.D., Humboldt University Berlin, 10117 Berlin, Germany
info@rationality-and-competition.de

Intertemporal Hedging and Trade in Repeated Games with Recursive Utility*

Asen Kochov Yangwei Song†

December 22, 2022

Recursive preferences have found widespread application in representative-agent asset-pricing models and general equilibrium. A majority of these applications exploit two decision-theoretic properties not shared by the standard model of intertemporal choice: (i) agents care about the intertemporal distribution of risk and (ii) rates of time preference, rather than being exogenously fixed, may vary with the level of consumption. We investigate what these features imply in the context of a repeated strategic interaction. Specifically, we identify novel opportunities for the players to manage risk and trade *intertemporally*, and characterize when such opportunities lead to an expansion of the feasible set of payoffs. Sharp implications for equilibrium behavior and the folk theorem are also deduced.

KEYWORDS: Recursive utility, repeated games, correlation aversion, endogenous discounting, intertemporal trade, intertemporal hedging.

*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, ITAM, 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 Deutsche Forschungsgemeinschaft through CRC TRR 190 (project number 280092119).

†**Kochov:** Department of Economics, University of Rochester, Rochester, NY 14627. Email: asen.kochov@rochester.edu.

Song: Department of Economics, University of Colorado Boulder, Boulder, CO 80309. Email: yangwei.song@colorado.edu.

1 Introduction

In a repeated game with standard preferences and a common discount factor, there are no gains from intertemporal trade. In fact, under a suitable normalization of utility, the set of payoffs in the repeated game is equal to that of the stage game. At a first glance, the conclusion may appear tautological. If the players are equally patient, where is the heterogeneity that makes trade possible? This paper shows that, far from being tautological, the conclusion rests heavily on the assumption of standard preferences. We do so by adopting a model of recursive utility in the tradition of **?**]. This sets the stage for two conceptually distinct types of interaction to make a difference: (i) *intertemporal trade* based on endogenously generated differences in the players' rates of time preference, and (ii) a form of *intertemporal hedging* based on the fact that risk and time are no longer treated symmetrically. Since there is no commitment, we also study the ability of the players to sustain such interactions in equilibrium and, in particular, the folk theorem.

Relative to the more general class of recursive preferences, standard preferences are predicated on two assumptions – time and state separability. It has been known since **?**] that if time separability is relaxed, then rates of time preference are no longer a fixed and "exogenous" parameter but rather a function of consumption. What has proven less clear is how the covariability between patience and consumption should be signed. The debate spans many authors and fields of inquiry. We give an overview in Section **??**, where we also raise some novel issues concerning games with different *types of outcomes*.

The practical implication of relaxing time separability is that differences in the players' rates of time preference may emerge *endogenously*, in the course of the game, creating opportunities for intertemporal trade. With standard preferences and heterogeneous discounting, the implications of such trade are known to be stark. As was first conjectured by **?**], less patient agents borrow incessantly against any future capital they may be endowed with and are left *immiserated* in the long run. Ramsey's conjecture has been confirmed in a variety of settings, both competitive and strategic, and played a role in numerous applications.[1] It has also bolstered an enormous interest in the "origins" of discounting and any heterogeneity therein. In this regard, and in line with arguments put forth by **?**] and **?**], our contribution is to examine the implications of intertemporal trade in a framework that permits a richer treatment of discounting and does not require any a priori heterogeneity. Two central questions arise: Will differences in discounting emerge in the first place and, if so, will a more patient player be able to sustain the higher level of patience?

[1]**?**] and **?**] confirmed the conjecture in the context of competitive growth economies. Starting with **?**], a growing literature has considered the strategic implications of intertemporal trade. See **?**], **?**], **?**], **?**], **?**], among others.

Before we preview our answers to these questions, we shift attention to our second focus in this paper.

1.1 The Intertemporal Hedging of Risk

Time and state separability jointly imply that the intertemporal distribution of risk is irrelevant. Intimately related is the fact that standard preferences cannot disentangle risk aversion from the degree of intertemporal substitution. The restrictive nature of these implications has been well understood in the literature on asset pricing with recursive utility pioneered by **?**]. In contrast to that literature, which has focused exclusively on representative-agent models, we examine the role of these implications in an interactive setup.

Consider a repeated prisoners' dilemma and let $v(CD)$ be the payoff vector when player 1 cooperates in every period, while player 2 defects. Likewise, let $v(CC)$ be the payoff vector when both players cooperate in every period and consider the average payoff $v' = 0.5v(CD) + 0.5v(CC)$. With standard preferences, v' can be attained in two ways. Flip a coin once and depending on the outcome, play CD forever or CC forever. Call this play the *one-time flip*. Alternatively, the players can flip the coin in each period. Call this *the iid flip*.

With recursive preferences, the iid and one-time flip are typically not indifferent. A preference for the iid flip is an example of what is sometimes called *correlation aversion*,[2] where by correlation one means the positive autocorrelation of the one time-flip. While we use the term as well, it is important to remark that correlation-aversion is not a standalone property of behavior, but a function of one's attitudes toward risk vis-à-viz intertemporal smoothing. The one-time flip, in which the outcome of the initial draw gets propagated forever, offers perfect smoothing across time at the expense of greater risk. The iid flip reverses the stakes: flipping the coin repeatedly offsets the risk in any given period (a bad outcome today need not be repeated tomorrow), but destroys the perfect smoothing across time.

If all players are correlation-averse, the iid flip is a Pareto improvement over the one-time flip. We show that, except in some special cases, this results in an *expansion* of the feasible set as illustrated by the dashed line in Figure **??**. The necessary and sufficient condition is that the game have some **conflict of interest** by which we mean that no single action simultaneously maximizes the utility of every player.

Alternatively, suppose the players are correlation-loving, that is, they care more about

[2]See **?**] for a general definition. Recently, correlation aversion has been the object of increased interest in experimental as well as theoretical work. See **?**], **?**], **?**], **?**].

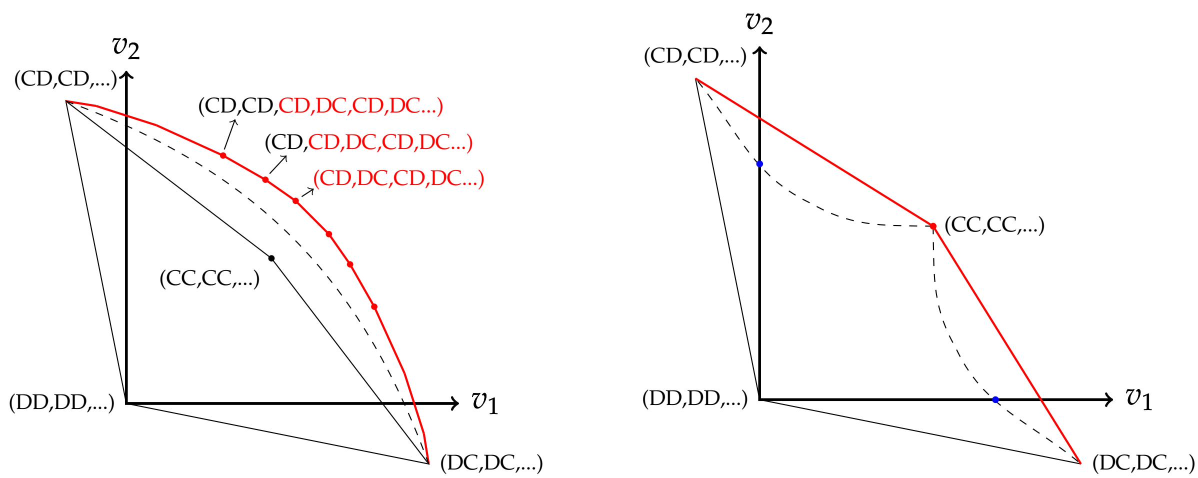

(a) Correlation Aversion/IMI

(b) Correlation Affinity/DMI

Figure 1: As usual, C stands for "cooperate" and D for "defect". Extreme points are associated with the play paths that generate them. The dashed lines in **????** and **????** represent the payoffs from iid flips. (The acronyms IMI/DMI are defined in Section **??**.)

intertemporal smoothing than risk. The dashed line will then curve inward as in Figure **??**. What is more interesting presently are the implications for equilibrium outcomes. Recall that in a subgame perfect equilibrium (SPE), the individual rationality (IR) constraints of each player must hold *after every history*. In a symmetric game with standard preferences, this requirement simplifies in that any payoff can be attained by a stationary strategy (an iid flip) and, consequently, it is enough to check that the IR constraints are met *ex ante*. But presently iid flips are not efficient. In fact, we show that the payoff $v' = 0.5v(CD) + 0.5v(CC)$, which is efficient under correlation affinity, can be attained *only* by a one-time flip. But clearly such behavior cannot be an equilibrium *for any level of patience*: player 1 will deviate in the history in which he has to cooperate forever while the other player defects.

1.2 A Model of Endogenous Discounting

Throughout most of the paper, we formalize the implications of both intertemporal trade and hedging within what is arguably the simplest model of recursive utility. In addition to delivering sharp characterizations, the model has a strong normative foundation, which we believe makes it an important benchmark for the study of strategic interactions. In particular, suppose the discounted sum of payoffs takes the form

$$\begin{aligned} v_i(a^0, a^1, \ldots) &= g_i(a^0) + \beta_i(a^0) g_i(a^1) + \beta_i(a^0)\beta_i(a^1) g_i(a^2) + \ldots \\ &= g_i(a^0) + \beta_i(a^0) v_i(a^1, a^2, \ldots). \end{aligned} \tag{1}$$

Above, $g_i(a)$ is player i's stage payoff from an action profile $a \in A$ and $\beta_i(a) \in (0,1)$ is i's discount factor as a function of a. If mixed strategies are employed, each player i computes the induced distribution over pure paths $(a^0, a^1, ...)$ and takes expectations in the usual way, which one can write as $\mathbb{E}v_i(\tilde{a}^0, \tilde{a}^1, ...)$. We refer to the preferences thus defined as **Uzawa-Epstein (UzE)** and note that if the function $\beta_i : A \to (0,1)$ is constant, one obtains the standard model of preference, which we refer to as one of exogenous discounting.[3]

It is known from **?**] that UzE preferences are the only ones that are recursive, stationary, and *indifferent to the timing of resolution of uncertainty*. Failures of the latter property have been criticized as they imply that agents are willing to pay for information that is of no instrumental value to them. See **?**]. In the context of a repeated game, such failures imply that the players care (for entirely non-strategic reasons) whether a period-t mixed action is implemented using contemporaneous or past signals. By focusing on UzE preferences, we ensure that such behavior plays no part in our results. Additionally, we note that UzE preferences retain the full force of Savage's [**?**] Sure Thing Principle or state separability. In other words, they depart from the standard model only in that they relax *time separability*, arguably the standard model's least appealing feature.

A drawback of UzE preferences is that, while they are sensitive to autocorrelations, attitudes toward risk and intertemporal smoothing are not fully disentangled. Instead, as we explain momentarily, they become entangled with properties of the endogenous discount factor $\beta_i : A \to \mathbb{R}$. It turns out however that our results pertaining to correlation attitudes do not rely on the specifics of the UzE model. In fact, the expansion of the Pareto frontier due to intertemporal hedging (the dashed curve in Figure **??**) does not require *any* restriction on preferences beyond correlation aversion. This is made clear in Section **??**. In Section **??**, we also confirm that the implications of correlation affinity discussed in Section **??** extend to a popular class of preferences introduced by **?**]. These preferences retain standard discounting while permitting what is arguably the most explicit disentanglement of risk aversion from the degree of intertemporal substitution. But first, we discuss two assumptions that have been central to the study of endogenous discounting.

1.3 A Debate on Marginal Impatience

Say that player i exhibits **decreasing marginal impatience (DMI)** if for every $a, a' \in A$

$$v_i(a, a, ...) > v_i(a', a', ...) \quad \Leftrightarrow \quad \beta_i(a) > \beta_i(a').$$

[3]**?**] introduced a special case of (**??**) in the absence of uncertainty. The above formulation is due to **?**] who both generalized the model and extended it to the case of uncertainty, all in a manner that preserves recursivity.

? , 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. **?** , p.30], on the other hand, noted that DMI leads to "disequilibrium behavior" and argued for the polar case of **increasing marginal impatience (IMI)**. Later, **?**] motivated IMI as capturing "a diminishing marginal utility from wealth accumulation," with **?**] arguing that the latter is especially likely at large wealth levels. **?**] also showed that IMI obviates the immiseration dynamics of **?**] and ensures the existence of a steady state with a non-trivial distribution of capital, a result which made IMI a staple in the growth literature.[4]

In the debate about IMI and DMI, attention has so far been focused on consumption savings problems with a single consumption good, whereby only the *level* of consumption matters. Considerations become much more nuanced in settings with different *types* of outcomes (as is the case in many classical games). For instance, building on the discussion in **?**] and **?**], one can argue that reading a book, while less enjoyable than going to the movies, stimulates the imagination and makes the future more salient, or, but to the same effect, that going to the movies produces a "visceral" reaction that biases people toward the present. Of course, for some, reading a book may not only increase patience but be more desirable as well.[5] Accordingly, we believe that one cannot make an a priori judgment in favor of IMI or DMI, which is why investigate each case in turn. First, we note an important connection between marginal impatience and correlation attitudes.

1.4 Marginal Impatience and Correlation Attitudes

It is known from **?**] that UzE preferences satisfy IMI (DMI) if and only if they are correlation-averse (correlation-loving). The intuition is straightforward. Under IMI, low consumption *today* increases the marginal utility from an extra unit of consumption *tomorrow*, boosting the "hedging benefits" afforded by the iid flip, while simultaneously reducing the benefits of intertemporal smoothing.

It follows that IMI leads to an expansion of the feasible set due to the hedging benefits of iid flips (Figure **??**), while DMI leads to a contraction of the equilibrium payoff

[4]See the surveys by **?**] and **?**]. We should also note the recent work of **?**], which examines the effects of happiness on discounting. To the extent that happiness derives from consumption, their results may be viewed as evidence of DMI.

[5]We are far from arguing that UzE preferences fully encapsulate the ideas of **?**] and **?**]. To get a sense of some of the modeling trade-offs, note that the former develop a model in which "imagination" is a stock variable accumulating over time. By comparison, the stationarity of UzE preferences means that any effect on patience is transient. The discussion in **?**] may on the other hand suggest a non-stationary model with (endogenous) present bias. In this context, UzE preferences are best viewed as capturing agents who are aware of the visceral effects of any outcome, not unlike Becker and Murphy's [**?**] model of rational addiction.

set (Figure **??**). Of course, correlation attitudes do not fully leverage the endogeneity of the discount factor and the associated opportunities for intertemporal trade. We turn attention to those next.

1.5 IMI and Intertemporal Cooperation

In Theorem **??**, we show that if conflict of interest holds, then gains from intertemporal trade expand the feasible set above and beyond what is attainable by iid flips. In specific games, we can go a step further and characterize efficient outcomes explicitly. Figure **??** depicts a possible scenario for the prisoners' dilemma. Note first that the unique efficient and symmetric outcome is for the players to take turns defecting, even though $(CC, CC, ...)$ is efficient in the space of one-time flips. The efficiency of such alternation, which we refer to as *intertemporal cooperation*, stems from differences in the players' rates of time preference that emerge along the path. Specifically, under IMI, if the action profile in some period t favors player i, then player i will attach relatively less weight to the future. Efficiency then requires that player $j \neq i$ be rewarded relatively more in the future. By IMI, j will then be relatively *less* patient in period $t + 1$, and so on. In sum, the players' discount factors seesaw repeatedly along the play path, making alternation efficient.

The second takeaway from Figure **??** is that, with the exception of the extremes in which one player defects forever while the other cooperates, eventually every efficient play path becomes one of intertemporal cooperation. An important practical implication is that the IR constraints of the players become slack over time. In fact, as we show formally in Section **??**, under a modicum of patience, it is enough to check that the IR constraints hold ex ante.

1.6 DMI and Immiseration

Recalling Figure **??**, we note that the implications of correlation affinity become less stark in games which, unlike the prisoners' dilemma, have multiple individually rational action profiles. Then, the implied constant paths $(a, a, ...)$ and the one-time flips among them will be sustainable in a SPE (for sufficiently high levels of patience). It turns out however that by leveraging the implications of DMI in terms of intertemporal trade, one can once again conclude that, *no matter the level of patience, the only efficient outcome sustainable in a SPE is a symmetric one, if such exists*. The logic is simple. Inverting the arguments just made in the case of IMI, we see that if a player becomes more patient at any point along an efficient play path, they will sustain the higher level of patience as the game progresses. In effect, efficiency and DMI create an immiseration dynamic not unlike that of

?]. Moreover, as **?**] first observed in the context of a repeated game with standard preferences and heterogeneous discounting, immiseration may push the players below their security levels so that, no matter the level of patience, the outcome cannot be sustained in a SPE. Our key insight relative to **?**] and **?**] is that under endogenous discounting, immiseration is not unavoidable. Coordinating on a symmetric and efficient outcome, if such exists, ensures that the players remain equally patient throughout the game and, hence, that immiseration is never triggered.[6]

2 The Strategic Environment

There is a finite set of players: $I := \{1, 2, ..., n\}$. In the stage game, player i can choose a pure action a_i in a finite, nonsingleton set A_i. Let $A := \times_{i \in I} A_i$. In the repeated game, time is discrete and indexed by $t \in \{0, 1, 2, ...\}$. To focus on the effects of endogenous discounting, we keep things as simple as possible and assume perfect monitoring, the availability of public randomization, and that "mixtures are observable." Formally, suppose that at the start of each period t, nature draws a public signal $\omega_0^t \in [0, 1]$ and, for each player i, a private signal $\omega_i^t \in [0, 1]$. All signals are drawn from the uniform distribution on $[0, 1]$, independent of one another and across time. Let $\mathfrak{a}_i^t : (\omega_0^t, \omega_i^t) \mapsto a_i \in A_i$ be i's action as a function of the observed public and private signal, and let $\mathfrak{a}^t = (\mathfrak{a}_i^t)_i$. Let h^0 be the initial, empty history. Given $t > 0$, a history $h^t = (\omega_0^0, \mathfrak{a}^0, \ldots, \omega_0^{t-1}, \mathfrak{a}^{t-1})$ consists of the "mixtures" chosen in the past and the realized public signals. A strategy for player i is a sequence $\sigma_i = (\sigma_i^t)_t$ where σ_i^t maps h^t into $\mathfrak{a}_i^t$. We let Σ_i be the set of all such strategies and $\Sigma = \times_i \Sigma_i$ be the set of all strategy profiles $\sigma = (\sigma_i)_i$. As is standard, we will often suppress the signals and instead speak of a mixed action $\alpha \in \Delta(A)$ being played after a given history h^t.[7]

Each strategy profile $\sigma \in \Sigma$ induces a probability distribution on A^∞ which, abusing notation slightly, we denote by σ as well. Each player i evaluates this distribution according to an UzE preference defined by a pair (g_i, β_i) as in Section **??**. A repeated game with endogenous discounting is thus a tuple $(A, (g_i, \beta_i)_{i \in I})$, with $v_i(\sigma)$ denoting i's utility from a distribution (strategy) σ. A strategy $\sigma \in \Sigma$ is a **subgame perfect equilibrium (SPE)** of

[6] This analysis, laid out in Theorem **??**, formalizes Friedman's [**?**] intuition that DMI leads to "disequilibrium behavior." Note as well that Friedman's brief remarks do not mention the possibility of coordinating on a symmetric outcome.

[7] In the literature, it is common to start with the above setup for the sake of simplicity and then relax the assumption of "observable mixtures," which is undoubtedly heavy-handed. We note that the assumption, which concerns only our folk theorem, is not needed when minmax strategies are pure. If minmax strategies are not pure, we can replace the assumption with a mild strengthening of the "NEU" assumption in Definition **??**. We pursue this extension in a separate paper.

the game if σ induces a Nash equilibrium in the continuation game associated with each history h^t.

We identify each mixed action $\alpha \in \Delta(A)$ with two distinct distributions on A^∞: (i) an iid distribution over time, denoted as α^{iid}, and (ii) a distribution α^{one} in which the players randomize once, according to α, and repeat the realized pure action $a \in A$ forever after. As in Section **??**, we call these an **iid flip** and a **one-time flip** respectively. We note that an iid flip can be thought of as the outcome of a **stationary strategy** $\sigma \in \Sigma$, that is, one such that for every i, there is a function $f_i : [0,1]^2 \to A$ such that $\sigma_i^t(h^{t-1})[\omega_0^t, \omega_i^t] = f_i(\omega_0^t, \omega_i^t)$ for all t, h^{t-1}, and all signal realizations ω_0^t, ω_i^t. If play depends on the time period t but not on history, then σ is **history-independent.** Any such strategy gives rise to a **play path** $(\alpha^0, \alpha^1, ...) \in (\Delta(A))^\infty$, or equivalently a product measure on A^∞. We use $\boldsymbol{\alpha}$ to denote a generic play path $(\alpha^0, \alpha^1, ...)$ and $\boldsymbol{a}$ to denote a pure play path $(a^0, a^1, ...)$. Given $\boldsymbol{\alpha}$ and $t > 0$, we let ${}_t\boldsymbol{\alpha} = (\alpha^t, \alpha^{t+1}, ...)$.

A game $(A, (g_i, \beta_i)_{i \in I})$ 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. Given $\alpha \in \Delta(A)$, we 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. We use v to denote the function $\sigma \mapsto (v_1(\sigma), ..., v_n(\sigma))$ or a point in its image. We let $v_i^{max} := \max_\sigma v_i(\sigma)$ be i's maximum feasible payoff in the repeated game and $\underline{v}_i := \min_{\sigma_{-i} \in \Sigma_{-i}} \max_{\sigma_i \in \Sigma_i} v_i(\sigma_i, \sigma_{-i})$ be i's **minmax** or **security level**. We also write $v_i(a)$ for $v_i(a, a, ...)$ and note that $v_i^{max} = v_i(a)$ for some $a \in A$.[8] Finally, we assume that no player is indifferent among all strategies or, equivalently, that for every i, there are $a', a'' \in A$ such that $v_i(a'') > v_i(a')$.

3 Varying Patience when Patience is Endogenous

In a folk theorem, it is standard to vary the level of patience *while keeping the stage game fixed*. In this construction, the stage game acts as an anchor ensuring that we have a family of repeated games representing the same strategic situation while differing only in the players' level of patience. 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. Some intuition for this can be gained from consumer choice theory. There, one typically speaks of the utility of a *bundle* and, unless utility is additively separable across goods, it is meaningless 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 issue arises in the case of UzE preferences, which are not time separable. The next lemma, due to **?**], formalizes

[8]See **?** , Lemma 3.4].

the observation.

Lemma 3.1. *Two pairs $(g_i, \beta_i), (g'_i, \beta'_i)$ induce the same UzE preference relation on Σ if and only if $\beta'_i = \beta_i$ and there are constants $\theta > 0$ and γ such that $g'_i = \theta g_i + \gamma(1 - \beta_i)$.*

Thus, if $\gamma \neq 0$ and discounting is endogenous, the functions g_i, g'_i need not be cardinal or even monotone transformations of one another. But if stage payoffs lack a clear ordinal meaning, how does one vary the level of patience while ensuring that the associated repeated games remain meaningfully related? Our answer involves two steps. Assuming exogenous discounting, we first clarify the ordinal meaning of the stage payoffs $(g_i)_i$ in terms of the repeated game. We then characterize the class of repeated games with endogenous discounting that have this ordinal input in common. The first step is clear. Let Δ^{iid} be the set of all iid flips and Δ^{one} be the set of all one-time flips. If discounting is exogenous, the von-Neumann-Morgenstern expected-utility theorem shows that (g_i, β_i) and (g'_i, β'_i) induce the same preference relation on $\Delta^{iid} \cup \Delta^{one}$ if and only if $g'_i = pg_i + q$ for some constants $p > 0$ and q. The next lemma provides an analogous result for the case of endogenous discounting.

Lemma 3.2. *Let (g_i, β_i) be such that $v_i(a) > v_i(a') > v_i(a'')$ for some $a, a', a'' \in A$. The pair (g'_i, β'_i) induces the same preference relation on $\Delta^{iid} \cup \Delta^{one}$ as (g_i, β_i) if and only if (g_i, β'_i) and (g'_i, β'_i) induce the same preference relation on Σ and $\beta'_i = \lambda_i + (1 - \lambda_i)\beta_i$ for some $\lambda_i < 1$.*

Lemma **??** is mouthful but the upshot is simple: if two UzE preferences agree on $\Delta^{iid} \cup \Delta^{one}$, then (i) it is without loss of generality to assume that they share the same stage payoffs and (ii) their discount factors *must be* related in the specified linear fashion.[9] These implications suggest the following approach to the folk theorem. Starting with a repeated game $(A, (g_i, \beta_i)_{i \in I})$, define for each $\lambda \in [0, 1)$ a discount factor $\beta_{i\lambda}$ and a repeated game Γ_λ by letting

$$\beta_{i\lambda} := \lambda + (1 - \lambda)\beta_i \quad \text{and} \quad \Gamma_\lambda := (A, ((1 - \lambda)g_i, \beta_{i\lambda})_{i \in I}).$$

Then, thinking of λ as an overall measure of patience, consider the equilibria of Γ_λ as $\lambda \nearrow 1$. Also, write Γ to mean either a single game $(A, (g_i, \beta_i)_{i \in I})$ or the family of games $\Gamma_\lambda, \lambda \in [0, 1)$, induced by it.

A few additional remarks about this approach are in order. First, note that in the definition of Γ_λ we have scaled the stage payoffs g_i by $(1 - \lambda)$. This is just a normalization.

[9] Of course, the stage payoffs g'_i can also be scaled differently from g_i. Combining Lemmas **??** and **??** gives all the options. We also note that $\beta'_i = \lambda_i + (1 - \lambda_i)\beta_i$ if and only if $(1 - \beta_i(a))(1 - \beta_i(a'))^{-1} = (1 - \beta'_i(a))(1 - \beta'_i(a'))^{-1}$ for all $a, a' \in A$, that is, if and only if the ratios of all marginal impatiences are the same.

Akin to the "$(1-\beta_i)$"-normalization one would do if β_i were exogenous, it ensures that payoffs do not blow up as $\lambda \nearrow 1$. In fact, letting $v_{i\lambda}$ be i's utility function in Γ_λ, observe that for all $\alpha \in \Delta(A)$ and λ:

$$v_{i\lambda}(\alpha^{iid}) = \frac{g_i(\alpha)}{1-\beta_i(\alpha)} \quad \text{and} \quad v_{i\lambda}(\alpha^{one}) = \sum_a \alpha(a)\frac{g_i(a)}{1-\beta_i(a)}. \tag{2}$$

Thus, given the normalization, we not only preserve each player's ranking on $\Delta^{iid} \cup \Delta^{one}$, but the utilities of iid and one-time flips as well.[10] In view of this, we henceforth suppress λ and write $v_i(\alpha^{iid})$ and $v_i(\alpha^{one})$. Second, our approach to the folk theorem implies that the assumptions of IMI/DMI and the corresponding correlation attitudes are all preserved as we vary λ. This is another sense in which the approach preserves qualitative features of the game. In fact, all assumptions imposed in this paper are preserved as we vary λ. Thus, when we say that a game Γ satisfies an assumption, one can understand this to mean that the original game $\Gamma_0 = (A, (g_i, \beta_i)_i)$ satisfies the assumption or equivalently that all games $\Gamma_\lambda, \lambda \geq 0$, do. Finally, observe that since UzE preferences are stationary and histories are public, the players minmaxing an opponent have no use conditioning on history.[11] That is,

Lemma 3.3. *For each i, the minmax strategy against player i and i's best response can be chosen to be stationary.*

It follows immediately that the minmax strategies against a player and the player's best response can be chosen independently of λ and, given the normalization of payoffs, that the security levels of all players are independent of λ as well. Taking advantage of Lemma **??**, we can also re-scale the original game Γ_0 so that all security levels are zero,[12] a normalization we maintain throughout the rest of the paper.

4 A Folk Theorem

Subgame perfection requires that the threat of future punishments be credible. Following **?**], this is typically done by finding strategies that *punish* a deviation, while simultaneously *rewarding* the players who carry out the punishment. In the case of standard preferences, a general condition under which such asymmetric treatment is possible is the

[10] Note that if discount factors are exogenous and identical among the players, one can set $\beta_i = 0$ so that $\beta_{i\lambda} = \lambda$ and $v_{i\lambda}(a, a, ...) = g_i(a)$ for each i, a, λ, that is, utility reduces to the standard discounted average.

[11] Similarly, since the public signal is observed by everyone, minmax strategies can be chosen to be independent of the public signal.

[12] Let $\hat{g}_i := g_i - \underline{v}_i(1-\beta_i)$. By Lemma **??**, $(A, (\hat{g}_i, \beta_i)_i)$ is strategically equivalent to $(A, (g_i, \beta_i)_i)$ and all security levels are zero.

assumption of **non-equivalent utilities (NEU)** of **?**]. As formulated in that paper, NEU says that no two players have identical preferences in the stage game, i.e, for every i, j, there are $\alpha, \alpha' \in \Delta(A)$ such that $g_i(\alpha) > g_i(\alpha')$ and $g_j(\alpha) \le g_j(\alpha')$. Under endogenous discounting, some modification is necessary since, as previously explained, stage payoffs do not have a well-defined ordinal meaning. The analysis so far suggests two options: either we assume that no two players have identical preferences on Δ^{iid} or that no two players have identical preferences on Δ^{one}. The conditions are logically independent and each one of them is an extension of the condition in **?**]. We chose the former because, as seen in Appendix **??**, the approach allows us to explore some well-known decision-theoretic properties of UzE preferences.

Definition 4.1. *A repeated game Γ **satisfies Non-Equivalent Utilities (NEU)** if for every $i, j \in I$, $i \neq j$, there are $\alpha, \hat{\alpha} \in \Delta(A)$ such that $v_i(\alpha^{iid}) > v_i(\hat{\alpha}^{iid})$ and $v_j(\alpha^{iid}) \le v_j(\hat{\alpha}^{iid})$.*

The next lemma, due to **?**], characterizes NEU in terms of the utility representations (g_i, β_i) and shows that the condition is generic.

Lemma 4.1. *(g_i, β_i) and (g_j, β_j) induce the same preference relation on the set Δ^{iid} of iid flips 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 notational simplicity, we state a folk theorem for the case when on-path behavior is history-independent, which means that it can be identified with a path $\boldsymbol{\alpha} = (\alpha^0, \alpha^1, ...) \in (\Delta(A))^\infty$. Such play is broad enough to encompass both iid flips and non-constant pure paths $(a^0, a^1, ...)$, which are used to implement intertemporal hedging and intertemporal trade respectively. History-dependent play, such as one-time flips, can be easily handled as well; all one has to do is to require that the IR constraints hold history by history. To state the theorem, for every $\varepsilon \ge 0$ 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 $v_{i\lambda}({}_t\boldsymbol{\alpha}) \ge \varepsilon$ 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 Γ_λ.*

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 **?**].*

We conclude this section with an important case excluded by our NEU condition (as well as by the alternative condition in terms of the players' preferences on Δ^{one}). Suppose discounting is exogenous and for every i, j, $g_i = g_j$ and $\beta_i \neq \beta_j$. Then, all players have identical preferences on Δ^{iid} and NEU fails. Yet, no two players have identical preferences on the space Σ of all strategies. When the latter is true, **?**] say that the game satisfies **Dynamic NEU**, which in the present context can be stated as: for every λ and i, j, there exist $\sigma, \sigma' \in \Sigma$ such that $v_{i\lambda}(\sigma) > v_{i\lambda}(\sigma')$ and $v_{j\lambda}(\sigma') \geq v_{j\lambda}(\sigma)$. Note that σ and σ' are not restricted in any way, nor is it required that they be independent of λ. When discounting is exogenous, **?**] are able to prove a folk theorem under this general condition. Doing the same for the case of endogenous discounting is an open problem.

5 Increasing Marginal Impatience

Say that a game Γ satisfies IMI if for each player i and every $a, a' \in A$,

$$v_i(a) > v_i(a') \quad \Leftrightarrow \quad \beta_i(a) < \beta_i(a').$$

As was illustrated in Figure **??**, IMI expands the Pareto frontier by delivering gains from both intertemporal hedging and intertemporal trade. The goal of this section is to formalize these observations. Beginning with hedging, let V^{one} be the payoff set from one-time flips:

$$V^{one} := \{v(\alpha^{one}) : \alpha \in \Delta(A)\}.$$

If a game has standard preferences and a common discount factor, then V^{one} is equal to the entire feasible set and, under the discounted average representation, to the payoff set in the stage game. Thus, with standard preferences, the only way to "improve upon" V^{one} is to assume ex ante heterogeneity in discounting, which then delivers gains from intertemporal trade. By comparison, our results do not require any a priori heterogeneity. We also note that UzE preferences cannot be distinguished from standard preferences if attention is restricted to one-time flips.[13] This is another sense in which V^{one} provides an appropriate benchmark against which to measure the implications of recursive preferences.

To state our first result, say that Γ satisfies **Conflict of Interest (CI)** if there is no action $a \in A$ such that $v_i(a) = v_i^{max}$ for every i.[14] Also, given a set $A \subset \mathbb{R}^n$ and a vector $x \in \mathbb{R}^n$, write $x >^* A$ if there is no $y \in A$ such that $y \geq x$ and $x \gg y$ for some $y \in A$. Finally, define the **Pareto frontier** of a set A to be the set of $x \in A$ for which there is no $y \in A$ such that $y \gg x$.

Theorem 5.1. *Assume IMI. Then, CI holds if and only if there is $\alpha \in \Delta(A)$ such that $v(\alpha^{iid}) >^* V^{one}$. In addition, if CI and NEU hold and V^{one} contains some payoff $v \gg 0$, then α can be chosen so that α^{iid}, or equivalently the path $(\alpha, \alpha, ...)$, can be sustained in a SPE for all λ large enough.*

[13] In fact, the same is true of all recursive preferences as defined in **?**].

[14] Recall that for every i, there is $a \in A$ such that $v_i(a) = v_i^{max}$. Thus, v_i^{max} is independent of λ.

The next lemma is key to the proof of Theorem **??**. Say that player i is **correlation-averse** if $v_i(\alpha^{iid}) \geq v_i(\alpha^{one})$ for each $\alpha \in \Delta(A)$, with a strict preference whenever $v_i(a) \neq v_i(a')$ for some $a, a' \in A$ in the support of α.

Lemma 5.1. *If player i's preferences (g_i, β_i) satisfy IMI, then the player is correlation-averse. The reverse is true if we assume that $v_i(a) \neq v_i(a')$ for all $a, a' \in A$.*

Proof. Assume IMI. If α has two actions in its support, then $v_i(\alpha^{iid}) \geq v_i(\alpha^{one})$ if and only if $(v_i(a) - v_i(a'))(\beta_i(a) - \beta_i(a')) \leq 0$. The latter inequality is automatically true when $v_i(a) = v_i(a')$. If $v_i(a) \neq v_i(a')$, the inequality is strict and follows from IMI. For more than two actions, the argument follows by induction. The reverse direction follows by analogous arguments.[15] □

Under IMI, any iid flip is a Pareto improvement over the corresponding one-time flip. The caveat we need to address is that the existence of a Pareto improvement need not deliver an *expansion* of the set V^{one}. In a two-player game, for example, the Pareto improvement could simply be a "movement" along a vertical segment (one orthogonal to the vector $(1, 0)$) of the frontier of V^{one}. In fact, if the Pareto frontier of V^{one} consists of exactly one vertical and one horizontal segment, or is a singleton, there are no gains from intertemporal hedging or trade. This is precisely what CI rules out: As we show formally in Appendix **??**, CI holds if and only if the Pareto frontier of V^{one} has at least one segment which is orthogonal to a strictly positive vector and isn't a singleton. In two-player games, the segment can be visualized as one that is strictly downward sloping but not vertical.

The proof sketch reveals that apart from CI, which we view as a structural assumption on the game rather than a preference assumption, Theorem **??** relies only on the assumption of correlation aversion. One may ask: what other well-known preferences exhibit correlation aversion? In Section **??**, we answer this question for a class of preferences studied in **?**], giving us another case in which the conclusions of Theorem **??** apply.

Going back to UzE preferences, it i clear that Theorem **??** does not fully leverage the endogeneity of discount factors as it does not take into account opportunities for intertemporal trade. To see this formally, let V^{iid} be the payoff set from iid flips and let $conv(V^{iid})$ be its convex hull.

Theorem 5.2. *Consider a symmetric game satisfying IMI and CI. For every λ, there exists a feasible payoff $\hat{v}$ such that $\hat{v} >^* conv(V^{iid})$.*

[15]**?**] was the first to observe the relationship between IMI and correlation aversion, without giving a proof. Our statement and proof is tailored to the present case of a discrete outcome space.

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

Figure 2: The prisoners' dilemma

The proof of Theorem **??** entails two main steps. First, we show that one can always find $\alpha \in \Delta(A)$ such that $v(\alpha^{iid})$ is on the Pareto frontier of V^{iid} and $v_i(\alpha^{iid}) > v_k(\alpha^{iid})$ and $\beta_i(\alpha) > \beta_k(\alpha)$ for some i, k. Then, exploiting the heterogeneity in discounting induced by α, Lemma **??** constructs a path $\boldsymbol{\alpha}_\lambda \in (\Delta(A))^\infty$ such that $v_\lambda(\boldsymbol{\alpha}_\lambda) >^* conv(V^{iid})$. Unfortunately, at this levels of generality we do not know if $\boldsymbol{\alpha}_\lambda$ can be sustained in a SPE (even if λ is large enough). This would be the case if, in addition to the properties already mentioned, α is such that $v(\alpha^{iid}) \gg 0$. Letting ε be such that $v(\alpha^{iid}) \gg \varepsilon > 0$, it is then immediate from the construction that $\boldsymbol{\alpha}_\lambda \in SIR^\varepsilon(\lambda)$ for all λ, which is enough for our folk theorem to kick in.

In specific games, we can characterize efficient outcomes explicitly and say much more about intertemporal trade and its sustainability in equilibrium. We do so next for the prisoners' dilemma.

5.1 The Prisoners' Dilemma

Let the action space A and the stage payoffs $g_1, g_2 : A \to \mathbb{R}$ be as in Figure **??** where, as usual, C stands for "cooperate" and D for "defect." For notational simplicity, we define discount factors as a function of stage payoffs rather than action profiles.[16] In particular, suppose $\beta_i = \beta \circ g_i$ for some function $\beta : \{b, 0, c, d\} \to (0, 1)$, where β is independent of i because we are interested in a symmetric game. As is typical in a prisoners' dilemma, we assume that

$$\frac{d}{1-\beta(d)} > \frac{c}{1-\beta(c)} > 0 > \frac{b}{1-\beta(b)}. \tag{3}$$

Consistent with the discussion in Section **??**, we 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 they defect and the other player cooperates to the play path in which both players cooperate. We also assume that

$$\frac{c}{1-\beta(c)} > \frac{1}{2}\frac{b}{1-\beta(b)} + \frac{1}{2}\frac{d}{1-\beta(d)}. \tag{4}$$

[16]Take some i and suppose, as is typical in a prisoners' dilemma, that $v_i(a) \neq v_i(a')$ for all $a, a' \in A$. Then, under IMI, or DMI, $\beta_i(a) \neq \beta_i(a)'$ for all $a, a' \in A$. By Lemma **??**, we can then normalize i's utility so that $g_i(a) \neq g_i(a')$ for all $a, a' \in A$, which means that β_i can be viewed either as a function on A or on $g_i(A)$.

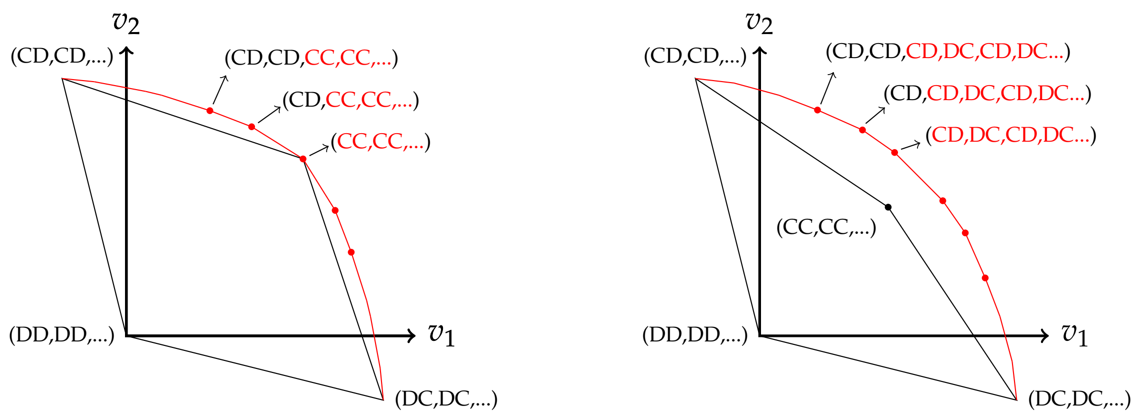

(a) Intratemporal cooperation

(b) Intertemporal cooperation

Figure 3: Two forms of cooperation under IMI

Thus, each player prefers cooperation in every period to receiving their worst or best play path with equal probability. The assumption helps us highlight the different predictions brought about by endogenous discounting and, specifically, the possibility that $(CC, CC, ...)$ may be Pareto dominated *even when* it is efficient in the space of one-time flips.

5.1.1 First-Best Outcomes

Figure **??** depicts two possibilities for the Pareto frontier in the repeated prisoners' dilemma. Consistent with Theorem **??**, 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. On the right, 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:

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

We refer to these paths as ones of **intertemporal cooperation**. Next, we characterize all efficient play paths for the two scenarios depicted in Figure **??**.

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

$\mathcal{EC} := \mathcal{E}_1\mathcal{C}_1 \cup \mathcal{E}_2\mathcal{C}_2$.[17]

Intertemporal cooperation. Consider the pairs (DC, CD) and (CD, DC) in A^2 and interpret each such pair as a "simple trade" in which the players swap turns defecting. Let $\mathcal{A}$ be the set of all play paths in which the players make such simple trades in succession:

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

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 **????** means that the payoffs from such paths are dispersed along the linear segment of the frontier 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 $\mathcal{EA} := \mathcal{E}_1\mathcal{A} \cup \mathcal{E}_2\mathcal{A}$.

We note that the Pareto frontier can take a third form not shown in Figure **??**. Thus, for some λ, intra- and inter- temporal cooperation can be simultaneously efficient. The analysis of this case is notationally cumbersome and delivers few additional insights. Since, in addition, the case does not arise for any λ sufficiently high, we defer its description to the Online Appendix. Presently, 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 Γ_λ. 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 attaining i's maximum payoff. In the context of the prisoners' dilemma, the path is unique. For instance, $\mathbf{a}^{max,1} = (DC, DC, ...)$.

Theorem 5.3. *Assume IMI. For every regular $\lambda \in [0, 1)$, the set $P(\lambda)$ of efficient play paths in the prisoners' dilemma is either $\mathcal{EC} \cup \{\mathbf{a}^{max,1}, \mathbf{a}^{max,2}\}$ or $\mathcal{EA} \cup \{\mathbf{a}^{max,1}, \mathbf{a}^{max,2}\}$.*

A natural question is whether gains from intertemporal trade persist in the limit as $\lambda \nearrow 1$. It turns out that the answer is yes: there are pure paths $\mathbf{a} \in A^\infty$ such that $\lim_{\lambda \nearrow 1} v_\lambda(\mathbf{a}) >^* conv(V^{iid})$.[18] The Online Appendix provides an example along with the proof of Theorem **??**.

[17] By definition, $\mathcal{C}_1$ contains paths such as $(CC, DC, CC, CC, ...)$ in which there is a single defection in some period $t > 0$. Such paths are not shown in Figure **??**, since their payoffs are not an extreme point.

[18] These conclusions should be contrasted with the conclusions reached at the end of Section **??**, where we examine the class of Epstein-Zin preferences.

5.1.2 Equilibrium Behavior

With intertemporal trade, the incentives to deviate vary over time, i.e., "borrowers" want to deviate when it's time to "repay," not before. This is why the folk theorem asks that the IR constraints be met *at each point in time.* As verifying the latter could be quite challenging in general, it is important to note that Theorem **??** delivers an important simplification. Namely, since efficient play paths converge to a symmetric outcome over time, the IR constraints become slack. The only caveat is when cooperation is *intratemporal* and one considers paths such as $(CC, CC, DC, CC, CC, ...)$ in which a single defection takes place at some $t > 0$. Clearly, for λ low enough, such paths are IR but not SIR. To deal with this problem, let λ' be such that

$$(1 - \lambda')d = v_{2\lambda'}(DC, CC, CC, ...) \tag{5}$$

and let $\underline{\lambda} = \max\{0, \lambda'\}$. With grim-trigger strategies, $(1 - \lambda)d$ is the maximum continuation utility of a player who deviates along any path, while $v_{2\lambda}(DC, CC, CC, ...)$ is the minimum continuation utility of *any* player at *any* point in time along a path with a single defection. Thus, $\underline{\lambda}$ is a threshold above which all paths with at most one defection can be sustained in a SPE. The discussion is summarized in the next corollary, where $IR^{\varepsilon}(\lambda)$ is 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. *Assume IMI. For every $\lambda \in (\underline{\lambda}, 1)$ and $\varepsilon > (1 - \lambda)d$, every path $\mathbf{a} \in P(\lambda) \cap IR^{\varepsilon}(\lambda)$ can be supported in a SPE of the prisoners' dilemma. If intertemporal cooperation is efficient for all λ, the restriction that $\lambda > \underline{\lambda}$ can be dropped.*

Note that, by Corollary **??**, one can first fix λ and then find ε such that a path $\mathbf{a} \in P(\lambda) \cap IR^{\varepsilon}(\lambda)$ can be sustained in a SPE. The result complements our folk theorem which, like other folk theorems, first fixes ε and then finds a threshold λ above which the respective SIR paths can be sustained.

6 Decreasing Marginal Impatience

Under DMI, an analogue of Lemma **??** shows that iid flips are Pareto inferior to one-time flips. Any expansion of the feasible set must therefore come from intertemporal trade, not hedging. We begin by formalizing the long-term of implications of such trade.

Theorem 6.1. *Consider a two-player, symmetric game satisfying DMI. 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 $i \in \{1, 2\}$ and $t > 0$, or (ii) $v_1(a^t) = v_2(a^t)$ for all $t \geq 0$.*

Intuition for the theorem was given in Section **??**. If at some point in time player i attains higher utility than j, then i will exhibit a greater level of patience. Given an efficient play path, this means that j's lifetime utility should be frontloaded, while i's utility should be backloaded. It follows that i's utility will remain higher as the game progresses and, given DMI, that i will sustain the higher level of patience. This self-enforcing dynamic continues until i's utility cannot be backloaded any further, which is when i's utility is maximized. If the game has conflict of interest, this also means that j's utility is at its minimum along the Pareto frontier, a conclusion not unlike that of Ramsey's [**?**] immiseration dynamic. The key difference is that under endogenous discounting the conclusion is not an inevitable consequence of efficiency. If there exists a constant play path $(a, a, ...)$ that is symmetric ($v_1(a) = v_2(a)$) and efficient, coordinating on this path ensures that differences in discounting do not emerge and that immiseration is effectively forestalled.[19]

Theorem **??** has obvious implications for equilibrium behavior. Suppose, as is true in many games, that a player's minimum on the Pareto frontier is below their security level. It follows that, *no matter the level of patience*, an efficient play path $(a_0, a_1, ...)$ can arise in a SPE *only if* $v_1(a_t) = v_2(a_t)$ for all t. Now, in general there may be no efficient play path that meets this condition. This is so *even when* there is an action $a \in A$ such that $v_1(a) = v_2(a)$ and $v(a)$ is on the Pareto frontier of V^{one}: the path $(a, a, ...)$ may be Pareto dominated by a one-time randomization among *non-constant* play paths featuring gains from intertemporal trade. As we establish in Section **??** however, (CC,CC,...) *is* efficient in the prisoners' dilemma. Thus,

Corollary 6.1. *Under DMI,* $(CC, CC, ...)$ *is the only efficient path that can arise in a SPE of the prisoners' dilemma.*[20]

6.1 Does the Feasible Set Expand? A Sufficient Condition

Theorem **??** does not show whether intertemporal trade results in an expansion of the feasible set. As was the case under IMI, CI is a necessary condition. Interestingly, CI is not sufficient under DMI. Consider the prisoners' dilemma. The actions that generate differences in discounting are CD and DC. If CD is played, the logic behind Theorem **??** tells us that 2's utility should be backloaded. But since CD is already as good as it gets

[19]If there are multiple paths $(a, a, ...)$ that attain the same symmetric payoff, then one can also alternate among the respective pure actions without triggering any differences in discounting. Case (ii) in Theorem **??** factors in this possibility.

[20]In contrast to the discussion in Section **??**, note that Corollary **??** is deduced as an implication of intertemporal trade, not correlation affinity. The implications of the latter are studied in Section **??**.

for player 2, this can only be done by repeating CD forever after, *leading to a constant play path*. Because intertemporal hedging is not a factor either, we see that the feasible set is equal V^{one}, which apropos proves the efficiency of $(CC, CC, ...)$ asserted in the lead-up to Corollary **??**.

To ensure that the feasible set expands, we need an action profile $a \in A$ that generates differences in discounting without automatically maximizing the utility of any player. Formally, say that a two-player game Γ satisfies **Richness** if there is $a^r \in A$ such that $v(a^r)$ is on the Pareto frontier of V^{one} and $v_i^{max} > v_i(a^r) > v_j(a^r)$ for some $i \in \{1, 2\}$ and $j \neq i$.[21]

Theorem 6.2. *Consider a two-player, symmetric game satisfying DMI and Richness. For each λ, there exists a feasible payoff v_λ such that $v_\lambda >^* V^{one}$. If the action profile a^r in the statement of Richness is such that $v(a^r) \gg 0$, then for all λ sufficiently large, the payoff v_λ can be chosen so it arises in a SPE.*

We give a sketch of the proof. By symmetry, it is without loss of generality to assume that a^r is such that $v_2^{max} > v_2(a^r) > v_1(a^r)$. As in Figure **??**, let $v(a^*)$ be the extreme point of V^{one} immediately to the left of $v(a^r)$ and consider the path $(a^r_{-T}, a^*_T) \in A^\infty$ such that a^* is played in period T and a^r in all other periods. Fixing λ, we first claim that for T large enough, the path generates gains from intertemporal trade, that is, $v_\lambda(a^r_{-T}, a^*_T) >^* V^{one}$. The intuition is simple. Since $v_2(a^r) > v_1(a^r)$, player 2 attains a higher level of patience at the start of the path when $a^r \in A$ is played. Efficiency then requires that 2's utility be backloaded, which is achieved by playing a^* in period T. This gives us *some* gains from intertemporal trade. To obtain a first-best outcome, the logic behind Theorem **??** tells us that 2's utility should continue to rise until it is fully maximized. However, as this may violate the IR constraints of player 1 and we want to show that gains from trade can be sustained in a SPE, the path (a^r_{-T}, a^*_T) sacrifices some efficiency by switching back to $a^r \in A$ after $a^* \in A$ is played. The switch may appear extreme but serves an important purpose. Suppose as in the second part of the theorem that $v(a^r) \gg 0$ and recall that, by construction, $v_1(a^*) < v_1(a^r)$. It follows that for a given λ and T, (a^r_{-T}, a^*_T) is SIR if and only if $(a^*, a^r, a^r, ...)$ is IR. That is, *the level of patience at which one can sustain (a^r_{-T}, a^*_T) is independent of T*. Therefore, by *first* choosing a sufficiently high λ and *then* a sufficiently large T, we can ensure that the path (a^r_{-T}, a^*_T) is *both* an equilibrium outcome and delivers gains from intertemporal trade.

To conclude this section, consider Figure **??** again, where $v(a^r)$ is on the frontier of V^{one} but not an extreme point. Clearly, gains from intertemporal trade will continue to exist if $v(a^r)$ is moved slightly below the frontier. Thus, despite being a rather weak requirement,

[21] In a symmetric game, Richness implies CI. Indeed, suppose there is $a \in A$ such that $v(a) = (v_1^{max}, v_2^{max})$. By symmetry, $v_1^{max} = v_2^{max} > v_2(a^r)$. But then $v(a) \gg v(a^r)$ and $v(a^r)$ cannot be on the frontier of V^{one}.

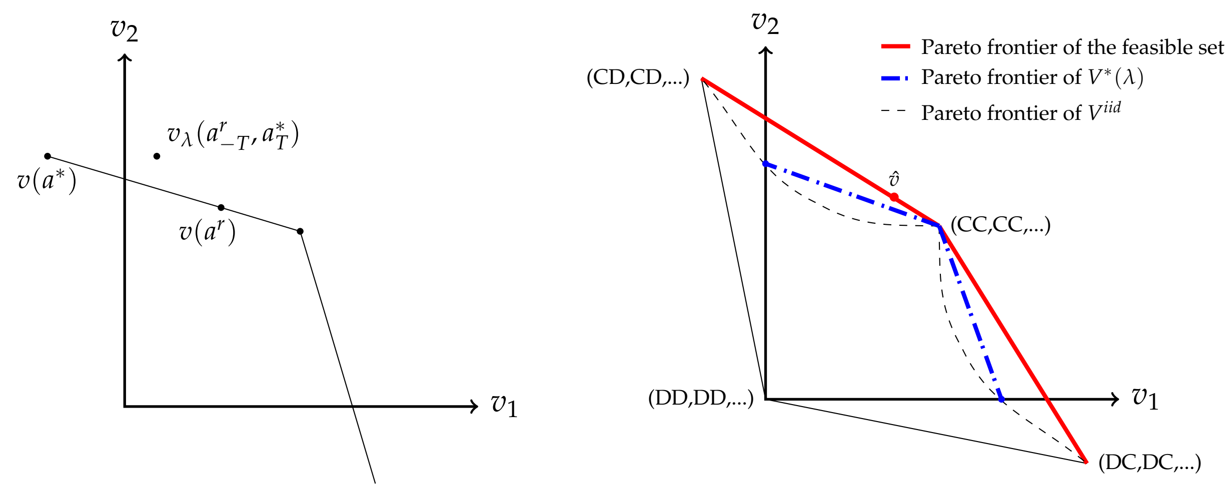

(a) Richness and gains from trade

(b) A bound on equilibrium payoffs

Figure 4

Richness is *not necessary* for the existence of such gains. Note however that exactly how far $v(a^r)$ can be moved is a quantitative question: one must weigh the benefits of the induced heterogeneity in discounting, $\beta_1(a^r) \neq \beta_2(a^r)$, against the inefficiency of playing a^r. Accordingly, we believe that Richness and the present analysis paint a reasonably complete picture as to when and whether gains from intertemporal trade exist under DMI.

6.2 Prisoners' Dilemma: The Limit Set of Equilibrium Payoffs

Returning to the prisoners' dilemma, consider Figure **??** and note that, as argued in Section **??**, the feasible set is equal to V^{one}. In Section **??**, we also argued that because iid flips are inefficient under correlation affinity, a payoff such as $\hat{v}$ can be attained only by a one-time flip between CC and CD and *such play cannot be an equilibrium for any level of patience*.[22] Presently, we go a step further and show that iid flips, whose stationarity helps ensure that the IR constraints are met after every history, fully delineate what can be attained in a SPE. In other words, and oversimplifying slightly, if the IR constraints cannot be met in a stationary way, they cannot be met in any other way. To state the result formally, let $v_{i\lambda}(\sigma \mid h^t)$ be i's utility in the subgame given a history h^t and a strategy $\sigma \in \Sigma$, and let

$$V^*(\lambda) = \{v_\lambda(\sigma) : \sigma \in \Sigma \text{ s.t. } v_\lambda(\sigma \mid h^t) \geq 0 \ \forall h^t, t\}.$$

[22]This is another way to deduce the conclusion of Corollary **??**.

$V^*(\lambda)$ is the payoff set from strategies that meet the IR constraints *after every history*, a necessary condition for subgame perfection.

Theorem 6.3. *In the prisoners' dilemma under DMI,* $V^*(\lambda) = conv(V^{iid} \cap \mathbb{R}^2_+)$ *for each* λ*.*

Figure **??** shows the Pareto frontier of $conv(V^{iid} \cap \mathbb{R}^2_+)$. We note that any payoff in $conv(V^{iid} \cap \mathbb{R}^2_+)$ can be attained by an individually rational iid flip or by a one-time randomization among such flips.[23] It follows from our folk theorem that any payoff $v \in conv(V^{iid} \cap \mathbb{R}^2_+)$ such that $v \gg 0$ can be sustained in a SPE for λ large enough. In particular, $conv(V^{iid} \cap \mathbb{R}^2_+)$ is the limit of the set of equilibrium payoffs as $\lambda \nearrow 1$.

7 Epstein-Zin Preferences

In this final section, we shift attention to a popular class of preferences introduced by **?**]. The goal is to deliver another formalization of the implications of correlation attitudes discussed in Section **??**. We begin with some conceptual issues that arise when the preferences in question, henceforth **EZ preferences**, are used in a game-theoretic context.

7.1 Setup: Physical Outcomes and Temporal Lotteries

In the standard model of repeated games, stage payoffs are cardinal payoffs typically interpreted as encoding the players' risk attitudes. In fact, they serve a double purpose in that they also encode the players' attitudes toward intertemporal smoothing. With EZ preferences, however, attitudes toward risk and intertemporal smoothing are fully independent of one another. Accordingly, we need two "cardinal scales" and a way to convert between them. To achieve that, it is convenient to assume that stage outcomes take the form of a physical, infinitely divisible good, in terms of which one can compute certainty equivalents. More precisely, let i's stage outcomes be consumption levels described by a function $g_i : A \to \mathbb{R}_+$. Then, assuming for the sake of simplicity that the game is symmetric, let $C := conv(g_i(A))$ to be the convex hull of consumption levels that can arise in the context of a game and let $r, s : C \to \mathbb{R}$ be strictly increasing functions whose curvatures represent the players' risk attitudes and respectively the desire to smooth consumption over time.

[23]To see what we mean by this, take $\alpha, \alpha' \in \Delta(A)$ and $p \in (0,1)$. Then, partition $[0,1]$, the range of the period 0 public signal, into $[0,p]$ and $(p,1]$, and partition those intervals further so as to implement α and α' respectively. Finally, consider the play such that if $\omega_0^0 \in [0,p]$, α is played forever after; else α' is played forever after.

We need one more tweak to the standard setup before we can introduce EZ preferences formally. Namely, since EZ preferences are sensitive to the timing of resolution of uncertainty, we cannot identify a strategy $\sigma \in \Sigma$ with the induced distribution over play paths. Instead, we must encode the timing of resolution of uncertainty and look at what is known as the induced *temporal lottery* or *probability tree*. To avoid unnecessary technicalities, we give only a heuristic definition of these objects; the formal details are well-known and can be found in **?**]. As a first step, let us for the moment ignore the possibility of randomization in the initial period. Each infinite probability tree can then be visualized as a pair (a, μ), where $a \in A$ is the action played in period 0 and μ is a distribution over infinite probability trees, one of which will prevail in period 1. This recursive structure means that we can define the space D of infinite probability trees as the unique (up to homeomorphism) set satisfying $D = A \times \Delta(D)$. Since, of course, randomization in the initial period is possible, the actual set of infinite probability trees is $\Delta(D)$ rather than D.

7.2 EZ Utility

We are ready to define the class of EZ preferences. Given a function $v_i : D \to \mathbb{R}$ and a distribution $\mu \in \Delta(D)$, let $\mathbb{E}_\mu v_i$ be the expectation of v_i. Then, $v_i : D \to \mathbb{R}$ is implicitly defined as the solution to the equation:

$$v_i(a, \mu) = rs^{-1}\big((1-\beta)s(g_i(a)) + \beta sr^{-1}(\mathbb{E}_\mu v_i)\big) \quad \forall (a, \mu) \in A \times \Delta(D) = D, \tag{6}$$

where $\beta \in (0, 1)$ is the players' common and fixed discount factor and rs^{-1} denotes the composition of r and s^{-1}, etc.[24] To understand equation (**??**), note first that we have written it so that the utility function v_i is denominated in r-utils. Then, starting backward, we apply sr^{-1} to the continuation utility $\mathbb{E}_\mu v_i$ so as to convert the latter into s-utils. This allows us to aggregate across time by computing the discounted average with current utility. Applying rs^{-1} to that average converts utility back into r-utils. Finally, note that having defined utility in r-utils, the utility of $\mu \in \Delta(D)$ is simply $\mathrm{E}_\mu v_i$.

Some examples may cast further light on (**??**). The utility of a path $(a^0, a^1, ...) \in A^\infty$ of pure actions is

$$v_i(a^0, a^1, ...) = rs^{-1}\big((1-\beta)\textstyle\sum_t \beta^t s(g_i(a^t))\big). \tag{7}$$

Since rs^{-1} is just an increasing transformation, we see that preferences over pure play paths conform to the standard model of discounted utility, with s capturing the desire to smooth consumption over time. On the other hand, given $\alpha \in \Delta(A)$, the utility of the

[24]We assume throughout that r, s are such that equation (**??**) has a unique solution. A well-known case is when both r, s are homothetic. Other cases can be found in **?**].

one-time flip α^{one} is:

$$v_i(\alpha^{one}) = \sum_a \alpha(a) r(g_i(a)). \tag{8}$$

We see that the curvature of r reflects risk aversion, as intended. We also note that (**??**) is the cardinal payoff in the stage game. Thus, denominating $v_i : D \to \mathbb{R}$ in r-utils ensures that payoffs in the repeated game are measured in the same units as stage payoffs.

7.3 Correlation Aversion

Write (β, r, s) for an EZ preference on $\Delta(D)$ and $(A, (g_i)_i, \beta, r, s)$ for a symmetric repeated game with EZ preferences. The game is **connected** if $g_i(A) = C$ for some i (the choice of i is immaterial by symmetry). Also, say that r **is a strictly concave transformation of** s if there is a strictly concave function $f : s(C) \to r(C)$ such that $r(c) = f(s(c))$ for all $c \in C$. Our first result in this section characterizes when EZ preferences are correlation-averse.

Lemma 7.1. *If r is a strictly concave transformation of s, then the EZ preference (β, r, s) is correlation-averse. The converse is true as well if we assume that the game is connected and both r and s are twice continuously differentiable.*[25]

The intuition behind Lemma **??** is simple. The greater curvature of r implies that a player is more concerned with risk than intertemporal smoothing which, as explained in Section **??**, makes iid flips preferable to one-time flips. The converse is true under natural conditions.

Lemma **??** gives us another case in which we can invoke Theorem **??** and obtain an expansion of V^{one} due to intertemporal hedging. Our next result goes a step further and characterizes the Pareto frontier of the feasible payoff set. First, define the **strong Pareto frontier** of a set $A \subset \mathbb{R}^n$ to be the set of $x \in A$ for which there is no $y \in A$ such that $y \geq x$ and $y \neq x$. Also, say that $\mu \in \Delta(D)$ is **trivially randomized** if $v_i(a', \mu') = v_i(a'', \mu'')$ for all i and $(a', \mu'), (a'', \mu'') \in D$ in the support of μ; the same is true for all elements in the support of each μ' such that (a', μ') is in the support of μ for some a', and so on. In other words, μ is trivially randomized if *after each history*, the players are indifferent about what happens next.

Theorem 7.1. *Suppose r is a strictly concave transformation of s and $\beta > 1 - |A|^{-1}$. Then, the Pareto frontier of the set V of feasible payoffs in a repeated game $(A, (g_i)_i, \beta, r, s)$ is equal to the Pareto frontier of $V^{pure} := \{v(a^0, a^1, ...) : (a^0, a^1, ...) \in A^\infty\}$. In addition, any payoff on the strong Pareto frontier of V can be attained only by a pure play path or a trivially randomized μ.*

[25] Contemporaneously with us, **?**] has obtained a related but logically independent result which gives a sufficient condition for a *more special* class of EZ preferences to exhibit a *more general* notion of correlation aversion.

The intuition is once again simple. Since the players are more concerned with risk than intertemporal smoothing, they avoid randomization in any efficient outcome. To understand why the restriction on β is needed, recall from (**??**) that the utility of a pure path is just the standard discounted average (subject to the rs^{-1} change of units).[26] Moreover, it is known from **?**] that if $\beta > 1 - |A|^{-1}$, the set of discounted-average payoffs from pure paths is convex. This allows us to prove that any $\mu \in \Delta(D)$ is Pareto dominated by a pure play path. If $\beta \leq 1 - |A|^{-1}$, non-trivial randomization may be efficient. But then, since EZ preferences are sensitive to the timing of resolution of uncertainty, one must know whether early or late resolution is preferred.[27] If early resolution is preferred, all randomization will take place in period 0; if late resolution is preferred, the exact timing will depend on the specific paths one is randomizing among. Unfortunately, for the general class of EZ preferences we have defined, we do not know if r being a concave transformation of s pins down the players' attitudes toward the timing of resolution of uncertainty. Thus, a full analysis of the case $\beta \leq 1 - |A|^{-1}$ awaits further development.

On the other hand, consider the popular case of *homothetic EZ preferences*: $r(c) = c^\gamma$ and $s(c) = c^\rho$ for some $\gamma, \rho \in (0,1)$. Then, as shown in **?**], r is a strictly concave transformation of s ($\gamma < \rho$) if and only if there is a preference for early resolution of uncertainty. Thus, if non-trivial randomization is to be efficient, it must take place in period 0.

We conclude this section with one additional insight based on the homothetic case. Invoking **?**] again, we first note that V^{pure} is independent of the discount factor β. In fact,

$$V^{pure} = \Big\{ \Big(rs^{-1}(v_1), ..., rs^{-1}(v_n) \Big) : (v_1, ..., v_n) \in V^s \Big\}, \quad \text{where}$$
$$V^s := conv\Big\{ \Big(s(g_1(a)), ..., s(g_n(a)) \Big) : a \in A \Big\}.$$

On the other hand, V^{iid} depends on the level of patience β. This arises because EZ preferences, unlike UzE preferences, are sensitive to the late resolution of uncertainty implied by iid flips and because this sensitivity depends on β. Precise characterizations of the latter dependence have proved difficult to obtain as there is no closed-form expression for the utility of iid flips. However, the analysis in **?**], **?**], **?** , p.1448] clearly show that such a dependence exists. In addition, **?**] have successfully computed $\lim_{\beta \nearrow 1} v_i(\alpha^{iid})$ for the case of homothetic EZ utility. It follows immediately from their formula that the Pareto frontier of V^{iid} converges to the Pareto frontier of the feasible set. The intuition, also revealed by the formula, is that as $\beta \nearrow 1$, iid flips are no longer penalized for the risk they carry (intertemporal hedging works perfectly in the limit) or for the late resolution of uncertainty. Whether these conclusions extend to all correlation-averse EZ preferences

[26] Because of this, it is also straightforward to characterize which pure paths are efficient. We omit the details.
[27] See **?**] for a formal definition of preference for early (late) resolution.

is an interesting open problem.

7.4 Correlation Affinity

In direct juxtaposition to Theorem **??**, our next result characterizes efficient outcomes when the players care more about intertemporal smoothing than risk.

Theorem 7.2. *If s is a strictly concave transformation of r, then the Pareto frontier of the feasible set of a repeated game $(A, (g_i)_i, \beta, r, s)$ is equal to the Pareto frontier of V^{one}.*

As was the case with Theorem **??**, if we exclude some non-generic cases, an efficient payoff can be attained *only* by a one-time flip. To see what needs to be excluded, suppose $g(a) = g(a')$ and $r(g(a))$ is on the Pareto frontier of V^{one}. Then, a path that alternates between a and a' will deliver the payoff $r(g(a))$ and be efficient as well. Likewise, such paths can be in the support of one-time flips. We omit the formal details for the sake brevity and since they are not a factor in our last result, which concerns the prisoners' dilemma.[28] The result gives us another formalization of the implications of correlation affinity first expressed in Section **??**.

Corollary 7.1. *In the prisoners' dilemma with EZ preferences, if s is a strictly concave transformation of r, then $(CC, CC, ...)$ is the only efficient outcome that can be sustained in a SPE.*

Appendix

In the appendix, we write $v_i(\alpha)$ instead of $v_i(\alpha^{iid})$.

A Proof of Lemma ??

Since (g_i, β_i) and (g_i', β_i') agree on Δ^{iid}, it follows Lemma **??** that there are q, r, s, t, $qt > sr$, such that $g_i' = qg_i + r(1-\beta_i)$ and $1-\beta_i' = sg_i + t(1-\beta_i)$. Since (g_i, β_i) and (g_i', β_i') agree on Δ^{one}, there are constants $\theta > 0$ and γ such that $v_i(a) = \theta v_i'(a) + \gamma$ for every $a \in A$. Plugging the former restrictions into the latter gives $sv_i(a)^2 + (t - \theta q - \gamma s)v_i(a) - (\theta r + \gamma t) = 0$ for all $a \in A$. Since a quadratic equation has at most two solutions and $v_i(a) > v_i(a') > v_i(a'')$ for some $a, a', a'' \in A$, it must be that $s = 0$. (Also, $t = \theta q$ and $\theta r = -\gamma t$.) Thus, $1 - \beta_i' = t(1-\beta_i)$ and $t > 0$. Letting $\lambda_i = 1 - t$ gives $\beta_i' = \lambda_i + (1-\lambda_i)\beta_i$. Next, note that $t > 0$ and $qt > sr = 0$ imply $q > 0$. By Lemma **??**, for every $\hat{\theta} > 0$ and $\hat{\gamma}$, (g_i', β_i')

[28] Presently, by a prisoners' dilemma, we mean that $r \circ g$ is given by the payoff matrix in Figure **??** and that $d > c > 0 > b$ and $c > 0.5b + 0.5d$. Interpreted as restrictions on the rankings of constant pure paths, the latter inequalities are the exact analogues of (**??**) and (**??**).

and $(\hat{\theta}g_i' + \hat{\gamma}(1-\beta_i'), \beta_i')$ induce the same preference relation on Σ. Letting $\hat{\theta} = q^{-1}$ and $\hat{\gamma} = -\hat{\theta}rt^{-1}$ implies that $\hat{\theta}g_i' + \hat{\gamma}(1-\beta_i') = g_i$. The opposite direction is trivial.

B Proof of Theorem ??

B.1 Payoff Asymmetry

Each pair (g_i, β_i) induces a preference relation $\succeq_i$ on the simplex $\Delta(A)$ represented by the utility function $v_i(\alpha)$. If $\beta_i : A \to (0,1)$ is constant, $\succeq_i$ is a standard expected utility preference. If β_i is not constant, then $\succeq_i$ belongs to the more general class of **weighted-utility preferences** studied in **?**]. We begin with some preliminary observations regarding such preferences.

Lemma B1. *If $v_i(\alpha) > v_i(\alpha')$, then $v_i(\alpha) > v_i(\varrho\alpha + (1-\varrho)\alpha') > v_i(\alpha')$ for all $\varrho \in (0,1)$. If $v_i(\alpha) = v_i(\alpha')$, then $v_i(\alpha) = v_i(\varrho\alpha + (1-\varrho)\alpha')$ for all $\varrho \in (0,1)$ (i.e., the indifference sets of $\succeq_i$ are hyperplanes).*

Proof. The first part 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 $ks^{-1} > (\rho k + (1-\rho)l)(\rho s + (1-\rho)t)^{-1} > lt^{-1}$. The second part is proved analogously. □

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

Proof. The result is clear if $\succeq$ is an expected utility preference. When $\succeq$ is not expected utility, the proof follows from Figure 1 in **?**]. Embedding the simplex $\Delta(A)$ into $\mathbb{R}^{|A|-1}$, we see that 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 $\succeq$, with the ranking between E_1 and E_2 determining the direction of increasing preference. □

Next is a generalization of the "payoff-asymmetry lemma" of **?**].

Lemma B3. *Under NEU, there exist $\alpha^1, ..., \alpha^n \in \Delta(A)$ such that $v_i(\alpha^j) > v_i(\alpha^i)$ for every $i \neq j$.*

Proof. Call the sought after $(\alpha^i)_i$ a **separation for** $(\succeq_i)_i$. Let $E_i(\alpha) := \{\alpha' \in \Delta(A) : \alpha' \sim_i \alpha\}$ be player i's indifference curve through $\alpha \in \Delta(A)$ and let $U_i(\alpha), L_i(\alpha)$ be the 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 $\alpha^1 \succ_2 \alpha \succ_2 \alpha^2$. If $\succeq_1$ and $\succeq_2$ share the same indifference curves, then , by NEU, $\succeq_1$ must be the negation of $\succeq_2$ and the claim follows. If $\succeq_1$ and $\succeq_2$ do not share the same indifference curves, then, by Lemma **??**, they have in common at most one indifference curve E^* intersecting the interior of $\Delta(A)$. Pick any $\alpha \notin E^*$ in the interior of $\Delta(A)$. The hyperplanes $E_1(\alpha)$ and $E_2(\alpha)$ partition $\Delta(A)$ into four cones with peak $\alpha : U_1(\alpha) \cap U_2(\alpha), U_1(\alpha) \cap L_2(\alpha), L_1(\alpha) \cap U_2(\alpha), L_1(\alpha) \cap L_2(\alpha)$. Picking any α^2 in the interior of $U_1(\alpha) \cap L_2(\alpha)$ and α^1 in the interior of $L_1(\alpha) \cap U_2(\alpha)$ proves the claim.

Proceeding inductively, suppose $(\alpha^1, ..., \alpha^m)$ is a separation for $(\succeq_1, ..., \succeq_m)$ and let $\succeq_{m+1}$ be a distinct weighted-utility preference. Reindexing if necessary, we can assume that $\alpha^i \succeq_{m+1} \alpha^1$ for all $i < m+1$. Since $\alpha^2 \succ_1 \alpha^1$ and $\alpha^2 \succeq_{m+1} \alpha^1$, $\succeq_1$ cannot be 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 **??**, $\succeq_1$ and $\succeq_{m+1}$ have at most one indifference curve in common, we can also assume that $E_1(\alpha^1) \neq E_{m+1}(\alpha^1)$. By the argument for $n = 2$, we can find α', α'' such that $\alpha'' \succ_1 \alpha^1 \succ_1 \alpha'$ and $\alpha' \succ_{m+1} \succ \alpha^1 \succ_{m+1} \alpha''$. Choosing α', α'' sufficiently close to α^1 ensures that $(\alpha', \alpha^2, ..., \alpha^m, \alpha'')$ is a separation for $(\succeq_1, \succeq_2, ..., \succeq_{m+1})$. □

B.2 Decision-theoretic preliminaries

We continue by stating two known and useful intertemporal properties of UzE preferences.[29] The proofs are obvious and omitted. Fix some $i \in I$ and 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 **??** shows that player i prefers more beneficial actions to 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 $v_i(\alpha^0, \alpha^1, ..., \alpha^K, \boldsymbol{\alpha}) \leq v_i(\hat{\alpha}^0, ..., \hat{\alpha}^{\hat{K}}, \boldsymbol{\alpha})$.*

Finally, we note 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 $v_i(\alpha^1, \alpha^2, ...)$.[30]

[29] For another application of these properties, see **?**].

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

B.3 Constructing dynamic player-specific punishments

The definition below is adapted from **?**].

Definition B1. *Given $\lambda \in [0,1)$, 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) $v_{i\lambda}(\mathbf{r}^i) < v_{i\lambda}({}_t\mathbf{r}^j) - \gamma$.*

The paths $(\mathbf{r}^i)_i$, which we are about to construct, will be used to punish deviations from the target path $\boldsymbol{\alpha}$. Roughly, condition (i) deters player i from deviating from $\boldsymbol{\alpha}$; condition (ii) ensures that the punishment phase is SIR and that no player wants to restart the punishment; and condition (iii) provides incentives for player i to carry out a punishment against player j.

Lemma B6. *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 parameters $T_1, T_2 \in \mathbb{N}_{++}$ (to be determined later). Fix $\varepsilon > 0$ and λ such that $SIR^\varepsilon(\lambda) \neq \emptyset$. Fix $i \in I$. 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 **??**, there exist $\kappa^1, ..., \kappa^n \in \Delta(A)$ 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 $(K^i_\lambda, \mathbf{w}^i_\lambda)$. Let $l^i, h^i \in A$ be such that $v_i(l^i) \leq v_{i\lambda}(\sigma) \leq v_{i\lambda}(h^i)$ for all $\sigma \in \Sigma$. Let $\mathcal{L}^i_\lambda$ be the set of all $l^j \in A, j \in I$, s.t. $v_i(l^j) < v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda)$, and let $N^i := |\mathcal{L}^i_\lambda|$. The set is nonempty since $l^i \in \mathcal{L}^i_\lambda$. Enumerate all action profiles in $\mathcal{L}^i_\lambda$ according to i's preferences:

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

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

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

Collecting all l's into a block L^i_λ, we may also write $\mathbf{r}^i_\lambda$ as $(L^i_\lambda, \boldsymbol{\alpha}^i_\lambda)$. Notice that we constructed the paths $\{\mathbf{r}^i_\lambda\}_i$ without referencing the target path $\boldsymbol{\alpha} \in SIR^\varepsilon(\lambda)$. 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 **??**, which is where the target path $\boldsymbol{\alpha}$ appears, will be automatically satisfied if we can 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 $\{\mathbf{r}^i_\lambda\}_i$ by choosing T_1 and T_2 appropriately so that all conditions of Definition **??** are met. To begin, recall the following property of the exponential.

Lemma B7. *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 $\underline{\beta}_i := \min_a \beta_i(a)$. For every λ, let $\overline{\beta}_{i\lambda} := \lambda + (1-\lambda)\overline{\beta}_i$ and $\underline{\beta}_{i\lambda} := \lambda + (1-\lambda)\underline{\beta}_i$.

Lemma B8. *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 **??**,

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

Let $f_i(\theta)$ denote the above limit and notice that $f_i(0) = \varepsilon > 0$ for every $i \in I$. Since $v_i(l^i) \le 0 < \varepsilon$, 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 B9. *Let θ^* be defined as in Lemma **??**. 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 $f_i(0) = v_i(l^i) \le 0 < \varepsilon$. 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 **??** 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 **??** and η^* is defined as in Lemma **??**.

Lemma B10. *There exist $\gamma' > 0$ and $\underline{\lambda}'$ such that $v_{i\lambda}(\mathbf{r}^i_\lambda) > \gamma'$ for all $\lambda > \underline{\lambda}'$ and i.*

Proof. By Lemma **??**, 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 $v_i(l^i) \leq v_i(l^{i_m})$ for all $m = 0, ..., N^i - 1$ and $v_i(l^i) \leq v_i(\kappa^{i_m})$ for all $m = 0, ..., n - 1$, we have

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

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

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

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

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

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

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

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

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

Since $v_{i\lambda}(\mathbf{w}^i_\lambda) \geq \varepsilon$, Lemma **??** 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 B12. *For all i and all $\lambda > \underline{\lambda}''$, $v_{i\lambda}(\mathbf{r}^i_\lambda) \leq v_{i\lambda}({}_t\mathbf{r}^i_\lambda)$ 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 (**??**) and (**??**) 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 (**??**) and (**??**) 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 (**??**), 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 (**??**) and (**??**) 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 $v_i(\kappa^{i_k}) \geq v_{i\lambda}({}_{(k+1)T_2}\boldsymbol{\alpha}^i_\lambda)$, (**??**) and (**??**) 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 (**??**), 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 **??** 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 B13. *There exist $\gamma''' > 0$ and $\underline{\lambda}'''$ such that for every $i, j \in I$, $i \neq j$, and $\lambda > \underline{\lambda}'''$, 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 **??**,

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

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 B14. *For all $i, j \in I, i \neq j, \lambda > \underline{\lambda}'''$, and $t \leq N^j T_1$, $v_{i\lambda}\big({}_t\mathbf{r}^j_\lambda\big) - v_{i\lambda}\big({}_tL^j_\lambda, \boldsymbol{\alpha}^i_\lambda\big) > \gamma'''$.*

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,

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

Thus, we seek a lower bound for $v_{i\lambda}\big(\boldsymbol{\alpha}^j_\lambda\big) - v_{i\lambda}\big(\boldsymbol{\alpha}^i_\lambda\big)$. 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 **??**, $v_{i\lambda}\big(K^{j\setminus j}, \mathbf{w}^i_\lambda\big) \geq v_{i\lambda}(K^{i\setminus j}, \mathbf{w}^i_\lambda)$ and, by stationarity,

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

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 **??**, 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

$$\begin{aligned}
&[\underline{\beta}_{i\lambda}]^{nT_1}(v_{i\lambda}(\boldsymbol{\alpha}^j_\lambda) - v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda)) \geq [\underline{\beta}_{i\lambda}]^{nT_1}(v_{i\lambda}((\kappa^j)^{T_2}, K^{i\setminus j}, \mathbf{w}^i_\lambda) - v_{i\lambda}(\tilde{K}, \mathbf{w}^i_\lambda)) \\
=&[\underline{\beta}_{i\lambda}]^{nT_1}(v_i(\kappa^j) - v_i(\kappa^i))(1 - [\beta_{i\lambda}(\kappa^j)]^{T_2})(1 - [\beta_{i\lambda}(\kappa^i)]^{T_2}) \\
\geq&[\underline{\beta}_{i\lambda}]^{nT_1}(v_i(\kappa^j) - v_i(\kappa^i))(1 - [\overline{\beta}_{i\lambda}]^{T_2})^2 > \gamma''',
\end{aligned}$$

where the equality follows by a direct calculation and the last inequality by Lemma **??**. Together with (**??**), the last chain of inequalities completes the proof. ☐

Lemma B15. *For all $i, j \in I, i \neq j$, and $\lambda > \underline{\lambda}'''$, $v_{i\lambda}({}_t\mathbf{r}^j_\lambda) - v_{i\lambda}(\mathbf{r}^i_\lambda) > \gamma'''$ 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 **??**, $v_{i\lambda}({}_tL^j_\lambda, \boldsymbol{\alpha}^i_\lambda) \geq v_{i\lambda}(\pi^\uparrow_i({}_tL^j_\lambda), \boldsymbol{\alpha}^i_\lambda)$. Hence, by Lemma **??**, $v_{i\lambda}({}_t\mathbf{r}^j_\lambda) - v_{i\lambda}(\pi^\uparrow_i({}_tL^j_\lambda), \boldsymbol{\alpha}^i_\lambda) > \gamma'''$. 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 the stationarity of UzE preferences, 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 (**??**) follows from Lemma **??**. ☐

Lemma B16. *For all $i, j \in I, i \neq j, \lambda > \underline{\lambda}''$, $v_{i\lambda}({}_t\mathbf{r}^j_\lambda) - v_{i\lambda}(\mathbf{r}^i_\lambda) > \gamma''$ 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 **??**, $v_{i\lambda}(\mathbf{w}^i_\lambda) - \gamma'' > v_{i\lambda}(\mathbf{r}^i_\lambda)$ 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 **??** and **??** 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 **??**, $x > \gamma''$ 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 **??**, **??**, and **??**. Then, Lemmas **??**, **??**, **??**, and **??** show that for all $\lambda > \underline{\lambda}$ and $\boldsymbol{\alpha} \in SIR^\varepsilon(\lambda)$, the paths $\{\mathbf{r}^i_\lambda\}_i$ meet the conditions in Definition **??**.

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 min-max strategy by the opponents. By Lemma **??**, 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 **??**, 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 $\mu_i > \frac{\overline{g}_i}{\gamma\left(1-\beta_i(m^i)\right)}$. Since

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

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

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

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

$$(1 - \lambda)\overline{g}_i + \big(\overline{m} - [\beta_{i\lambda}(m^j)]^{\mu}(\overline{m} + \gamma)\big) - v_i(m^j)(1 - [\beta_{i\lambda}(m^j)]^{\mu}) < 0. \tag{19}$$

Since $\overline{g}_i$ and $v_i(m^j)$ are constants that do not depend on λ, 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 (**??**) is satisfied for all $\lambda > \underline{\lambda}'''_i$. Since there are finitely many players and finitely many integers 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. 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. The calculations are straightforward and omitted.

C More Preliminaries

Fix an n-player game $(A, (g_i, \beta_i)_i)$. Let $V(\lambda)$ be the set of all feasible payoffs in Γ_λ.

Lemma C17. *The set $V(\lambda)$ is convex.*

Proof. By the definition of UzE preferences, $V(\lambda) \subset conv(V^{pure})$. The converse inclusion will follow if one can show that any distribution on A^∞ with finite support can be induced by some behavioral strategy σ. This is trivial given that the public signal is continuously distributed.[31] $\square$

The proof of the next lemma is straightforward and omitted.

Lemma C18. *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 ??

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

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

Lemma D19. *For some $\eta \in \mathbb{R}^n_{++}$, the set $F(\eta)$ is not a singleton.*

Proof. By way of contradiction, suppose $F(\eta)$ is a singleton for each $\eta \in \mathbb{R}^n_{++}$. Since V^{one} 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. $\square$

Pick $\eta \in \mathbb{R}^n_{++}$ as in Lemma **??**. For any v in the relative interior of $F(\eta)$, there is $\alpha \in \Delta(A)$ such that $v(\alpha^{one}) = v$. By Lemma **??**, $v(\alpha) \geq v(\alpha^{one})$. In addition, by the choice of η, there is $i \in I$ and actions a, a' in the support of α such that $v_i(a) \neq v_i(a')$. By Lemma **??** again, $v_i(\alpha) > v_i(\alpha^{one})$. It follows that $v(\alpha) >^* V^{one}$.

Next, suppose there is $\varepsilon > 0$ and $v \in V^{one}$ such that $v \gg \varepsilon$. Let V^{one}_ε be the set of all $v' \in V^{one}$ such that $v' \geq \varepsilon$. We claim that there is no $\overline{v} \in V^{one}_\varepsilon$ such that $\overline{v} \geq v'$ for all $v' \in V^{one}_\varepsilon$. If such a $\overline{v}$ existed, then, by CI, there would be $i \in I$ and $v^i \in V^{one}$ 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 $\varrho\overline{v} + (1-\varrho)v^i \in V^{one}_\varepsilon$, contradicting the definition of $\overline{v}$. Using the same arguments as above, we can then show that V^{one}_ε has a face $F(\eta)$, $\eta \gg 0$, that is not a singleton and, in addition, 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, $v(\alpha) \geq \varepsilon$ and $v(\alpha) \notin V^{one}$. By our folk theorem, $v(\alpha)$ is an SPE payoff for all λ sufficiently high.

[31] More generally, the result is true as long as the public signal is rich enough to implement any $\alpha \in \Delta(A)$.

E Proof of Theorem ??

Lemma E20. *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 $j \neq i, k$.*

Proof. Fix λ and η such that $\eta_k > 0$. By symmetry, there is $\alpha_k \in \Delta(A)$ 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

$$v_k(\alpha_k) - v_k(\alpha) > \frac{\eta_i}{\eta_k} [\frac{\beta_{i\lambda}(\alpha)}{\beta_{k\lambda}(\alpha)}]^T (v_i(\alpha) - v_i(\alpha_k)). \tag{20}$$

Let $\boldsymbol{\alpha} = (\alpha^0, \alpha^1, ...)$ be such that $\alpha^t = \alpha$ for all $t \leq T$ and $\alpha^t = \alpha_k$ for all $t > T$. From (**??**), deduce that $\eta \cdot v_\lambda(\boldsymbol{\alpha}) > \eta \cdot v(\alpha)$. The other (in)equalities follow by construction. □

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 $conv(V^{iid})_+$ be the Pareto frontier of $conv(V^{iid})$ and $e^i \in \mathbb{R}^n_+$ be the vector whose i^{th}-coordinate is 1 and all other coordinates are 0. Note that $v(a^{max,i}) \in F(e^i)$.

Lemma E21. *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 **??**. By Lemma **??**, if $v(\alpha) \in F(e^i)$, then every $a \in \operatorname{supp} \alpha$ is such that $v(a) \in F(e^i)$. By Lemma **??**, $\beta_i(a) = \min\{\beta_i(a') : a' \in A\}$ for all $a \in \operatorname{supp} \alpha$. By the symmetry of the game, $v_i(\alpha) \geq v_k(\alpha)$ and $\beta_i(\alpha) \leq \beta_k(\alpha)$ for all k. 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 **??**, $\beta_i(a) < \beta_k(a)$. 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 $conv(V^{iid})$ such that $v \in conv(V^{iid})_+$. 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 $conv(V^{iid})_+ \setminus F(e^i)$. Suppose $Z \cap \operatorname{cl} Y \neq \emptyset$. One can think of this as a situation in which the face $F(e^i)$ connects smoothly with the rest of the frontier of $conv(V^{iid})$. 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, assume that $\alpha^m \to_m \alpha^*$, where $v(\alpha^*) = \hat{v} \in Z$. By Lemma **??**, $\alpha^* \in S_i$ 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 $conv(V^{iid})_+$, with η such that $\eta \in \mathbb{R}^n_+ \setminus \{0, e^i\}$. But, by Lemma **??**, $v(\alpha^m)$ is not on the corresponding η-face of $V(\lambda)$ and the theorem is proved. Alternatively, suppose $F(e^i)$ connects "nonsmoothly" to the rest of the frontier $conv(V^{iid})_+$, by which we mean that there is $v^* \in X$ belonging to both $F(e^i)$ and a face $F(\eta)$ of $conv(V^{iid})_+$ such that $\eta \in \mathbb{R}^n_+ \setminus \{0, e^i\}$. Letting $\alpha \in \Delta(A)$ be such that $v(\alpha) = v^*$, similar use of Lemma **??** completes the proof. Finally, the next lemma, an adaptation of Lemma 2.2 in **?**], confirms that the two scenarios we considered are exhaustive.

Lemma E22. *If $Z \setminus \mathrm{cl}\, Y \neq \emptyset$, then there is $v^* \in X \cap F(e^i)$ such that $N(v^*)$ is not a singleton.*

F Two-Player Games: Preliminary Lemmas

This section introduces some notation and results about two-player games which will be useful later on. Fix λ and $\eta \in \mathbb{R}^2_+$. Given $\mathbf{a} \in A^\infty$, let $s_\lambda(\mathbf{a}, \eta) := \eta \cdot v_\lambda(\mathbf{a})$ and let $P(\lambda, \eta)$ be the set of pure play paths $\mathbf{a} \in A^\infty$ that maximize $s_\lambda(\cdot, \eta)$. Also, say that η' **determine the same direction as** η 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 $\eta^t_\lambda(\mathbf{a})$. Finally, when indices $i, j \in I$ appear in the same context, it will be understood that $i \neq j$. The next two results are standard so we omit the proofs.

Lemma F23. *If $\mathbf{a} = (a^0, a^1, ...) \in P(\lambda, \eta)$, 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 $(a^0, ..., a^{t-1}, \hat{\mathbf{a}}) \in P(\lambda, \eta)$.*

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 there is a unique $a^* \in A^E$ such that $v_1(a^*) = \max_{a \in A^E} v_1(a)$. The next two lemmas assume either IMI or DMI.

Lemma F24. *For every $\mathbf{a} \in P(\lambda, \eta)$, if $a^0 \in A^E$, then ${}_1\mathbf{a} \in P(\lambda, \eta)$ and $(a^0, a^0, ...) \in P(\lambda, \eta)$.*

Proof. Under both IMI and DMI, $a^0 \in A^E$ if and only if $g_1(a^0) = g_2(a^0)$ and $\beta_{1\lambda}(a^0) = \beta_{2\lambda}(a^0)$. Thus, η and $\eta^1_\lambda = \big(\eta_1\beta_{1\lambda}(a^0), \eta_2\beta_{2\lambda}(a^0)\big)$ determine the same direction and, by Lemma **??**, ${}_1\mathbf{a} \in P(\lambda, \eta)$. Since $\mathbf{a} = (a^0, {}_1\mathbf{a}) \in P(\lambda, \eta)$, we get $s_\lambda(\mathbf{a}, \eta) = s_\lambda({}_1\mathbf{a}, \eta)$. 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 get

$$s_\lambda(\mathbf{a}, \eta) = s_\lambda({}_1\mathbf{a}, \eta) = \eta_1 v_1(a^0) + \eta_2 v_2(a^0). \tag{21}$$

Since ${}_1\mathbf{a} \in P(\lambda, \eta)$, it follows that $(a^0, a^0, ...) \in P(\lambda, \eta)$. □

Lemma F25. *For every $\mathbf{a} \in P(\lambda, \eta)$, if $a^t \in A^E$ for some t, then $a^t = a^*$.*

Proof. Obvious given Lemma **??**. □

G Proof of Theorem ??

To state the next four lemmas, fix $\lambda \in [0,1)$, $\eta \in \mathbb{R}^2_{++}$, and $\mathbf{a} \in P(\lambda, \eta)$.

Lemma G26. *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 ${}_1\mathbf{a} \in P(\lambda, \eta^1_\lambda)$,

$$v_{2\lambda}({}_1\mathbf{a}) \leq v_{2\lambda}(\mathbf{a}) \quad \text{and} \quad v_{1\lambda}({}_1\mathbf{a}) \geq v_{1\lambda}(\mathbf{a}). \tag{22}$$

From (**??**), 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 (**??**) 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 **??**, $\beta_{2\lambda}(a^0) < \beta_{1\lambda}(a^0)$ implies $v_2(a^0) < v_1(a^0)$. Hence, $v_{2\lambda}(\mathbf{a}) < v_{1\lambda}(\mathbf{a})$. □

Lemma G27. *If $v_{1\lambda}(\mathbf{a}) = v_{2\lambda}(\mathbf{a})$, then $\mathbf{a} = (a^*, a^*, ...)$.*

Proof. By Lemma **??**, $\beta_{1\lambda}(a^0) = \beta_{2\lambda}(a^0)$ and, hence, $a^0 \in A^E$ by Lemma **??**. 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 **??**, $\mathbf{a} = (a^*, a^*, ...)$. □

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

Lemma G28. *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 G29. *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, since $\beta_{1\lambda}(a^0) < \beta_{2\lambda}(a^0)$, Lemma **??** implies that $v_{1\lambda}(\mathbf{a}) < v_{2\lambda}(\mathbf{a})$. Conclude that

$$v_{1\lambda}(\hat{\mathbf{a}}) < v_{2\lambda}(\hat{\mathbf{a}}) \quad \forall \hat{\mathbf{a}} \in P(\lambda, \eta^T_\lambda(\mathbf{a})). \tag{23}$$

By Lemma **??**, ${}_T\mathbf{a} \in P(\lambda, \eta^T_\lambda(\mathbf{a}))$ and, hence, $v_{1\lambda}({}_T\mathbf{a}) < v_{2\lambda}({}_T\mathbf{a})$. By Lemma **??**, $\beta_{1\lambda}(a^T) \leq \beta_{2\lambda}(a^T)$. By the choice of T, it must be that $\beta_{1\lambda}(a^T) = \beta_{2\lambda}(a^T)$. By Lemma **??**, $v_1(a^T) = v_2(a^T)$ so that $a^T \in A^E$. It follows from Lemmas **??** and **??** 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 (**??**). □

We can now complete the proof of Theorem **??**. 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 λ. From now on, assume $\mathbf{a}^{max,1} \neq \mathbf{a}^{max,2}$. Take λ, $\eta \in \mathbb{R}^2_+$, and $\mathbf{a} \in P(\lambda, \eta)$. If $\eta_i = 0$ and $\eta_j > 0$ for some $i, j \in I$, 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 **??** shows that $\mathbf{a} = (a^*, a^*, ...)$, as desired. Assume $v_{1\lambda}(\mathbf{a}) < v_{2\lambda}(\mathbf{a})$. By Lemma **??**, $\beta_{1\lambda}(a^0) \leq \beta_{2\lambda}(a^0)$. 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 **??** 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 **??**, 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 **??** implies that $a^t = a^*$ for every $0 \leq t < T$. Since $a^0 = a^*$, Lemma **??** 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 **??** repeatedly, conclude that $v_{1\lambda}({}_t\mathbf{a}) < v_{2\lambda}({}_t\mathbf{a})$ for every $t \leq T$. By Lemma **??**, $\beta_{1\lambda}(a^T) \leq \beta_{2\lambda}(a^T)$ and, by the choice of T, $\beta_{1\lambda}(a^T) < \beta_{2\lambda}(a^T)$. By Lemma **??**, $\beta_{1\lambda}(a^t) < \beta_{2\lambda}(a^t)$ 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.

H Proof of Theorem ??

Let a^* be as defined in the text and let F be the face of V^{one} containing $v(a^r)$ and $v(a^*)$. We claim that F is downward sloping, that is, orthogonal to a vector $\eta \gg 0$. That $\eta_1 > 0$ follows from $v_2(a^r) < v_2^{max}$. On the other hand, if $\eta_2 = 0$, then $v_1^{max} = v_1(a^r) < v_2(a^r) < v_2^{max}$, contradicting symmetry.

Next, for any λ, i, and T, let $\varrho_i = [\beta_{i\lambda}(a^r)]^T(1 - \beta_{i\lambda}(a^*))$. Note that $(1 - \varrho_1)v(a^r) + \varrho_1 v(a^*) \in F$, while

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

Since F is downward sloping, $v_\lambda(a^r_{-T}, a^*_T) >^* V^{one}$ if and only if $\varrho_2 > \varrho_1$. But

$$\varrho_2 > \varrho_1 \quad \Leftrightarrow \quad \left[\frac{\beta_{2\lambda}(a^r)}{\beta_{1\lambda}(a^r)}\right]^T > \frac{1-\beta_{1\lambda}(a^*)}{1-\beta_{2\lambda}(a^*)} \equiv \frac{1-\beta_1(a^*)}{1-\beta_2(a^*)}.$$

By Lemma **??**, $v_2(a^r) > v_1(a^r)$ implies $\beta_{2\lambda}(a^r) > \beta_{1\lambda}(a^r)$. Hence, $\varrho_2 > \varrho_1$ for all T large enough. The second assertion of the theorem was proved in the main text.

I Proof of Theorem ??

By construction, $conv(V^{iid} \cap \mathbb{R}^2_+) \subset V^*(\lambda)$ for each λ. We prove the opposite inclusion by showing that the Pareto frontier of $V^*(\lambda)$ coincides with the Pareto frontier of $conv(V^{iid} \cap \mathbb{R}^2_+)$. We start with some notation. Let $\alpha^{0,2} \in \Delta(A)$ be the mixed action $\varrho CC + (1-\varrho)CD$ such that $v_1(\alpha^{0,2}) = 0$. Define $\alpha^{0,1}$ analogously and let F^0_i be the line connecting $v(\alpha^{0,i})$ with $v(CC)$. These lines are independent of λ since the utilities $v(\alpha^{0,i})$ and $v(CC)$ are independent of λ. By construction, $F^0_1 \cup F^0_2$ is the Pareto frontier of $conv(V^{iid} \cap \mathbb{R}^2_+)$. Fixing λ, let $\hat{P}(\lambda, \eta)$ be the set of strategies $\sigma \in V^*(\lambda)$ that maximize $\eta_1 v_1 + \eta_2 v_2$ and let $\hat{P}(\lambda) := \cup_{\eta \in \mathbb{R}^2_+} \hat{P}(\lambda, \eta)$. If a strategy $\hat{\sigma}$ does not depend on h^1, we write it as (α, σ), where $\alpha \in \Delta(A)$ is the mixed action played in period $t = 0$ and strategy σ is played starting from period $t = 1$. Note that

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

which is an analogue of (**??**). We also have the following analogue of Lemma **??**:

Lemma I30. *If $(\alpha, \sigma) \in \hat{P}(\lambda, \eta)$, then $\sigma \in \hat{P}(\lambda, \eta')$ where $\eta'_i = \beta_{i\lambda}(\alpha)\eta_i$.*

For the prisoners' dilemma, we also have the following extension of Lemma **??**.

Lemma I31. *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 $g_1(\alpha) \geq g_2(\alpha)$. The desired conclusion follows. □

Lemma I32. *If $(\alpha, \sigma) \in \hat{P}(\lambda)$, then $\alpha^{iid} \in V^*(\lambda)$. 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 **??**, $v_i(\alpha) = v_j(\alpha)$. Thus, $v(CC) \geq v(\alpha)$ and, unless $\alpha = CC$, some convex combination of $v(CC)$ and $v_\lambda(\sigma)$ will strictly Pareto dominate $v_\lambda(\alpha, \sigma)$, a contradiction. Next, suppose that $\beta_1(\alpha) < \beta_2(\alpha)$. If $(\alpha, \sigma) \in \hat{P}(\lambda, \eta)$, we know from Lemma **??** 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 **??**, $v_2(\alpha) > v_1(\alpha) \geq 0$. □

Lemma I33. $\max\{v_2(\alpha) : \alpha^{iid} \in V^*(\lambda)\} = v_2(\alpha^{0,2}) > v_2(CC)$.

Proof. Let $\varrho^* \in (0,1)$ be such that $v_1(\varrho^* CD + (1-\varrho^*)DC) = 0$. One checks 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 $v_2(\alpha^{0,2}) > v_2(CC)$. 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 (**??**). □

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 signal determines whether the players follow σ or σ', with the probability of the former being ϱ. Note that

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

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

Lemma I34. *If $\sigma \in V^*(\lambda)$ is such that $v_{2\lambda}(\sigma) = v^*_{2\lambda}$, then $v_{1\lambda}(\sigma) = 0$.*

Proof. If $v_{1\lambda}(\sigma) > 0$, then there is $\varrho \in (0,1)$ such that $\varrho(CD, \sigma) + (1-\varrho)\sigma \in V^*(\lambda)$ 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 $v(CC)$. Let R be the ray originating at $v(CC)$ and passing through $v(\alpha^{0,2})$. Let $F_2(\lambda) := \{v_\lambda(\sigma) : \sigma \in \hat{P}(\lambda) \text{ and } v_{2\lambda}(\sigma) \geq v_{1\lambda}(\sigma)\}$. The next two lemmas collect several facts about the geometry of the feasible set (under DMI). The simple, but tedious, proofs are omitted.

Lemma I35. $v_\lambda(CD, \alpha^{0,2}) \in R$. *Also, if $\sigma \in \Sigma$ is such that $v_{1\lambda}(\sigma) = 0$ and $v_\lambda(\sigma) = \varrho v_\lambda(CD, \sigma) + (1-\varrho)v(CC)$ 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 $v(CC)$ and passing through $v_\lambda(\sigma)$.*

Lemma I36. *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 I37. *If $v^*_{2\lambda} = v_{2\lambda}(\alpha, \sigma)$ for some strategy $(\alpha, \sigma) \in V^*(\lambda)$, 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 **??** and **??** show that $v^*_{2\lambda} = v_2(\alpha^{0,2})$. Assuming the latter, suppose that for some $\sigma \in V^*(\lambda)$, $v_\lambda(\sigma) \in F_2(\lambda) \setminus F^0_2$. Since $v_\lambda(CD, \alpha^{0,2}) \in R$, there is $\varrho \in (0,1)$ such that $\varrho(CD, \alpha^{0,2}) + (1-\varrho)\sigma \in V^*(\lambda)$ 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 $\hat{\sigma}$ of the form $\varrho(\alpha', \sigma') + (1-\varrho)(\alpha, \sigma)$ where $v_{1\lambda}(\alpha', \sigma') < 0 < v_{1\lambda}(\alpha, \sigma)$ and $\sigma', (\alpha, \sigma), \sigma \in V^*(\lambda)$. Also, there must be some $\eta \in \mathbb{R}^2_+$ such that $\hat{\sigma}, (\alpha, \sigma) \in \hat{P}(\lambda, \eta)$. Let $L(\eta)$ be the linear segment connecting $v_\lambda(\hat{\sigma})$ and $v_\lambda(\alpha, \sigma)$. $L(\eta)$ is orthogonal to η and part of the frontier $F_2(\lambda)$. By Lemma **??**, $\sigma \in \hat{P}(\lambda, \eta')$, where $\eta'_i = \beta_{i\lambda}(\alpha)\eta_i$.

Case 1: Suppose $v_2(\alpha) > v_1(\alpha)$. By Lemma **??**, $\beta_2(\alpha) > \beta_1(\alpha)$ 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 as well that $\alpha(DD) = 0$. Otherwise, replacing DD with CC in α would lead to a strict Pareto improvement and contradict $v_{2\lambda}(\hat{\sigma}) = v^*_{2\lambda}$. With this in mind, observe that

$$v_\lambda(\alpha, \hat{\sigma}) = \alpha(CD)v_\lambda(CD, \hat{\sigma}) + \alpha(CC)v_\lambda(CC, \hat{\sigma}) + \alpha(DC)v_\lambda(DC, \hat{\sigma}). \tag{24}$$

Let L be the line passing through $v(CC)$ and $v_\lambda(\hat{\sigma})$. By construction, $v_\lambda(CC, \hat{\sigma}) \in L$. By Lemma **??**, $v_\lambda(DC, \hat{\sigma})$ is below the line connecting $v(DC)$ and $v_\lambda(\hat{\sigma})$, and hence, below L. Finally, by Lemma **??**, $v_\lambda(CD, \hat{\sigma})$ 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 (**??**) 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 **??**, $F_2(\lambda) = F^0_2$.

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 $v_\lambda(\sigma) = v(CC)$, then $v_\lambda(\alpha, \sigma) = v(CC)$ and the claim follows. If $v_{1\lambda}(\sigma) < v_1(CC)$, 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 $v(CC)$ 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 $v(CC)$.

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 $v_\lambda(\alpha, \sigma) \in F_2(\lambda)$, $v_\lambda(\alpha', \sigma')$ must lie on the line L'' defined by $F_2(\lambda)$. As before,

$$v_\lambda(\alpha', \sigma') = \alpha'(CC)v_\lambda(CC, \sigma') + \alpha'(CD)v_\lambda(CD, \sigma') + \alpha'(DC)v_\lambda(DC, \sigma'). \tag{25}$$

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

Case 2.1: Suppose $v_\lambda(\sigma') = \varrho v_\lambda(\hat{\sigma}) + (1-\varrho)v(CC)$ for some $\varrho \in [0,1]$. Then, we have $v_\lambda(CC, \sigma') \in F_2(\lambda)$. By Lemma **??**, $v_\lambda(DC, \sigma')$ 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 **??**, $v_\lambda(CD,\hat{\sigma})$ 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 $v(CC)$ and $v(\alpha^{0,2})$. Summarizing, we see from (**??**) that $v_\lambda(\alpha',\sigma') \in L''$ if and only if $v_\lambda(\hat{\sigma}) = v(\alpha^{0,2})$.

Case 2.2: Letting $\check{\sigma}$ be the symmetric analogue of $\hat{\sigma}$, so that $v_{i\lambda}(\check{\sigma}) = v_{j\lambda}(\hat{\sigma})$, suppose $v_\lambda(\sigma') = \varrho v_\lambda(\check{\sigma}) + (1-\varrho)v(CC)$ for some $\varrho \in [0,1)$. We are going to obtain a contradiction. Recall that, by definition, $v^*_{2\lambda} = v_{2\lambda}(\hat{\sigma})$ and note that $v_\lambda(CD,\hat{\sigma}) = ((1-\lambda)b, (1-\lambda)d + \beta_\lambda(d)v^*_{2\lambda})$ and $v_\lambda(CD,\check{\sigma}) = ((1-\lambda)b + \beta_\lambda(b)v^*_{2\lambda}, (1-\lambda)d)$, where $\beta_\lambda(d) := \lambda + (1-\lambda)\beta(d)$ and $\beta_\lambda(b)$ is similarly defined. Since $v_{1\lambda}(\alpha',\sigma') < 0$, deduce from (**??**) that $v_{1\lambda}(CD,\sigma') < 0$ and, hence, that $v_{1\lambda}(CD,\check{\sigma}) = (1-\lambda)b + \beta_\lambda(b)v^*_{2\lambda} < 0$.

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

$$\hat{S} := \frac{(1-\lambda)d + \beta_\lambda(d)v^*_{2\lambda} - v_2(CC)}{(1-\lambda)b - v_1(CC)} \quad \text{and} \quad \check{S} := \frac{(1-\lambda)d - v_2(CC)}{(1-\lambda)b + \beta_\lambda(b)v^*_{2\lambda} - v_1(CC)}.$$

Since $v_\lambda(\sigma') = \varrho v_\lambda(\check{\sigma}) + (1-\varrho)v(CC)$, we get $|\check{S}| \geq |\hat{S}|$ and $v_{2\lambda}(CD,\check{\sigma}) = (1-\lambda)d > v_2(CC)$. Since $v^*_{2\lambda} > v_2(CC)$ and $(1-\lambda)b + \beta_\lambda(b)v^*_{2\lambda} < 0$, we get

$$v^*_{2\lambda} \geq \frac{1}{\beta_\lambda(b)\beta_\lambda(d)}[v_2(CC)\beta_\lambda(b) + v_2(CC)\beta_\lambda(d) - \beta_\lambda(b)(1-\lambda)d - \beta_\lambda(d)(1-\lambda)b]. \tag{26}$$

Since $\mathbf{a}^C$ maximizes the sum of the players' utilities, $2v_2(CC) \geq v^*_{2\lambda}$. Combining with (**??**) and simplifying gives

$$x := (v_2(CD) - v_2(CC))\big(v_1(CC) - v_1(CD)\big)^{-1} \geq \beta_\lambda(d)(1-\beta(b))\big(\beta_\lambda(b)(1-\beta(d)\big)^{-1}.$$

By DMI, $\beta_\lambda(d) > \beta_\lambda(b)$ and $1-\beta(b) > 1-\beta(d)$. Thus, $x > 1$, which contradicts (**??**).

J Proof of Lemma ??

Note that r is a strictly concave transformation of s if and only if rs^{-1} is strictly concave. To see that the latter implies correlation aversion, observe that

$$\begin{aligned} v_i(\alpha^{iid}) &= \sum_a \alpha(a) rs^{-1}[(1-\beta)s(g_i(a)) + \beta s r^{-1}(v_i(\alpha^{iid}))] \\ &\geq (1-\beta)\sum_a \alpha(a) r(g_i(a)) + \beta v_i(\alpha^{iid}). \end{aligned} \tag{27}$$

Thus, $v_i(\alpha^{iid}) \geq \sum_a \alpha(a) r(g_i(a)) = v_i(\alpha^{one})$, with a strict inequality if $v_i(a) \neq v_i(a')$ for some $a, a' \in A$ in the support of α.

To prove the converse, suppose the game is connected. We can then imagine that EZ preferences are defined on the space of infinite probability trees of *consumption outcomes* and that they inherit correlation aversion on that space. Accordingly, for any $c, c' \in C$ and $\theta \in (0,1)$, let $\theta^{iid}_{c,c'}$ be the iid flip such that in each period consumption is c with probability θ and c' with probability $1-\theta$. Let $\theta^{one}_{c,c'}$ be the corresponding one-time flip.

Lemma J38. *For all $c, c' \in C$ such that $c < c'$ and $\theta \in (0,1)$, we have $r(c) < v_i(\theta^{iid}_{c,c'}) < r(c')$.*

Proof. Suppose $v_i(\theta^{iid}_{c,c'}) \geq r(c')$. Since sr^{-1} is strictly increasing and $c' > c$, we have $sr^{-1}(v_i(\theta^{iid}_{c,c'})) \geq s(c') > s(c)$. Letting $x := v_i(\theta^{iid}_{c,c'})$, observe that

$$\begin{aligned} x &= \theta rs^{-1}[(1-\beta)s(c) + \beta sr^{-1}(x))] + (1-\theta)rs^{-1}[(1-\beta)s(c') + \beta sr^{-1}(x))] \\ &< \theta x + (1-\theta)x = x, \end{aligned}$$

a contradiction. The case $v_i(\theta^{iid}_{c,c'}) \leq r(c)$ can be similarly ruled out. □

Since r and s are twice continuously differentiable, so is rs^{-1}. Next, observe that rs^{-1} cannot be convex on any interval. Else, one can use Lemma **??** and an argument analogous to that behind (**??**) to obtain a violation of correlation aversion. It follows that $(rs^{-1})'' \leq 0$ and the set of points at which $rs^{''} < 0$ is nonempty and open. Being continuously differentiable, $(rs^{-1})'$ is absolutely continuous. Using the fundamental theorem of calculus, we see that $(rs^{-1})'(x) - (rs^{-1})'(y) = \int_y^x (rs^{-1})'' < 0$ for all $x > y$. Thus, rs^{-1} is strictly concave.

K Proof of Theorem ??

We first claim that for any $\mu \in \Delta(D)$, there is $\alpha \in \Delta(A)$ such that

$$sr^{-1}(v_i(\mu)) \leq \sum_{a \in A} \alpha(a) s(g_i(a)) \quad \forall i. \tag{28}$$

At period 0, μ induces a distribution over $m \geq 1$ probability trees in the form of $(a^l, \mu^l) \in A \times \Delta(D)$, where $l = 1, ..., m$. Let $\kappa^l > 0$ denote the probability of (a^l, μ^l). Since rs^{-1} is strictly concave, we obtain that

$$\begin{aligned} v_i(\mu) &= \sum_{l=1}^m \kappa^l rs^{-1}[(1-\beta)s(g_i(a^l)) + \beta sr^{-1}(v_i(\mu^l))] \\ &\leq rs^{-1}[(1-\beta)\sum_{l=1}^m \kappa^l s(g_i(a^l)) + \beta \sum_{l=1}^m \kappa^l sr^{-1}(v_i(\mu^l))] \quad \forall i. \end{aligned} \tag{29}$$

Since r and s are strictly increasing functions, we get

$$sr^{-1}(v_i(\mu)) \leq (1-\beta)\sum_{l=1}^m \kappa^l s(g_i(a^l)) + \beta \sum_{l=1}^m \kappa^l sr^{-1}(v_i(\mu^l)) \quad \forall i. \tag{30}$$

If each μ^l is a constant pure path $(\hat{a}^l, \hat{a}^l, ...)$, then $sr^{-1}(v_i(\mu^l)) = s(g_i(\hat{a}^l))$, and we are done. For general μ, the claim is proved by iterating the argument and invoking continuity at infinity. Combining our claim with **?** , Lemma 1] shows that any payoff on the Pareto frontier of V can be attained by a pure play path. Finally, if $\mu \in \Delta(D)$ is not trivially randomized, then the inequalities (**??**) and (**??**) must be strict for at least some player i and thus $v(\mu)$ cannot be on the strong Pareto frontier.

L Proof of Theorem ??

Adopt the same notation as in the proof of Theorem **??**. Since rs^{-1} is convex,

$$\begin{aligned} v_i(\mu) &= \sum_{l=1}^{m} \kappa^l rs^{-1}[(1-\beta)s(g_i(a^l)) + \beta sr^{-1}(v_i(\mu^l))] \leq \sum_{l=1}^{m} \kappa^l [(1-\beta)r(g_i(a^l)) + \beta v_i(\mu^l)] \\ &= \sum_{l=1}^{m} \kappa^l [(1-\beta)v_i(a^l, a^l, ...) + \beta v_i(\mu^l)] =: \hat{v}_i \quad \forall i. \end{aligned}$$

If each μ^l is a constant pure path $(\hat{a}^l, \hat{a}^l, ...)$, then $\hat{v}_i$ is the utility of a one-time flip. As before, iterating the argument and using continuity at infinity completes the proof.