Intertemporal Hedging and Trade in Repeated Games with Recursive Utility*

Asen Kochov   Yangwei Song†

August 1, 2023

Two key features distinguish the general class of recursive preferences from the standard model of dynamic choice: (i) agents may care about the intertemporal distribution of risk, and (ii) their rates of time preference, rather than being fixed, may vary with the level of consumption. We investigate what these features imply in the context of a repeated strategic interaction. First, we show that opportunities for intertemporal trade may expand the set of feasible payoffs relative to that in a static interaction. Two distinct sources for such trade are identified: *endogenous heterogeneity* in the players' rates of time preference and a *hedging motive* pertaining to the intertemporal distribution of risk. The set of equilibrium payoffs may on the other hand shrink drastically as many efficient outcomes become unsustainable no matter the level of patience. This "anti-folk" result occurs when the players prefer stage outcomes to be positively correlated rather than independent across time. Intuitively, such preferences make it inefficient to offset short-term losses with future gains, while this is needed to ensure that security levels are met on path. We also establish a folk theorem: if security levels *are* met on path, such play can be sustained in a subgame perfect equilibrium provided that the players are sufficiently patient.

*For their helpful comments, 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. The comments of the editor and four anonymous referees led to significant improvements as well. The paper serves as our response to William Thomson who likes to point out that while fairness sometimes implies efficiency, the converse is never true. Song gratefully acknowledges financial support by the German Science Foundation (through Collaborative Research Center TRR 190).

†**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 the usual normalization of utility, the set of feasible payoffs in the repeated game is equal to that in the stage game. We argue that this conclusion, which may at first sight appear tautological, rests critically on two features of standard preferences: (i) *the players do not care about the intertemporal distribution of risk*, and (ii) *their rates of time preference are exogenously fixed* and unaffected by what transpires in the course of the game. In particular, adopting a more general class of recursive preferences in the tradition of Kreps & Porteus (1978), our first contribution is to identify conditions under which gains from intertemporal trade exist and expand the set of feasible payoffs. Notably, the conditions are generic and do not require a priori heterogeneity across the players.

Yet, existence is only half the story. In the absence of commitment, the outcomes that matter are those that can be sustained in an equilibrium of the game. Our second contribution shows that attitudes toward intertemporal risk play a pivotal role in this respect. The argument has two parts. First, under general recursive utility, many efficient outcomes may necessitate behavior that is history-dependent on path.[1] But then, if the players prefer stage outcomes to be positively correlated rather than independent across time, security levels may be violated in some history, *even if they are met ex ante*.[2] Things come to a head in the prisoners' dilemma where, as a result, cooperation is the *only* efficient outcome sustainable in a subgame perfect equilibrium. Intuitively, other outcomes require the players to suffer the occasional short term loss, which can then be offset by future gains. But under the attitudes in question, such offsets, which are needed to ensure that security levels are met on path, are inefficient. In this way, sensitivity to intertemporal risk takes away a degree of freedom that is needed to reconcile efficiency and individual rationality.

Generalizing a result of Abreu *et al.* (1994), our final contribution is a folk theorem showing that any path of play that *does meet* the security levels of the players can be

[1]The same is not true under standard preferences. Indeed, existing folk theorems such as those of Fudenberg & Maskin (1986), Abreu *et al.* (1994), Lehrer & Pauzner (1999), and Chen & Takahashi (2012) all take advantage of the fact that under standard preferences any payoff can be attained by a strategy that is history-independent on path and, in the special case of a common discount factor, one that is furthermore stationary on path.

[2]Attitudes toward autocorrelations have been the subject of a growing literature, both theoretical and experimental. See, for example, Miao & Zhong (2015), Andersen *et al.* (2018), Kochov (2015), and Bommier *et al.* (2019). To our knowledge, we are the first to examine the role of such attitudes in the context of a strategic interaction.

sustained in a subgame perfect equilibrium, if the players are sufficiently patient.[3] One application, see Theorem 6, is that even when the efficient "level" of intertemporal trade is unsustainable, as is the case under the attitudes described above, *some* level will be. We remark that the use of a more general class of recursive preferences poses some novel challenges not only in terms of the proof but the formulation of the folk theorem. Section 3 expounds on these; at present, we proceed with a detailed example of the role of intertemporal risk.

1.1 The Intertemporal Distribution of Risk

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 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 known as *correlation aversion*,[4] where *correlation* refers to the positive autocorrelation of the one time-flip. To see what is at stake, note that the one-time flip offers perfect smoothing across time at the expense of greater risk: a bad outcome, once drawn, lasts forever. The iid flip reverses the stakes: flipping the coin repeatedly offsets the risk in any given period but destroys the perfect smoothing across time. A person who is more concerned with risk than intertemporal smoothing may thus prefer the iid flip.

If all players are correlation-averse, the iid flip is a Pareto improvement over the one-time flip. Theorem 2 shows that, except for a non-generic case, this results in an *expansion* of the feasible set of payoffs as depicted by the dashed line in Figure 1(a). We refer to this expansion as being the result of **intertemporal hedging**, a special form of intertemporal trade.

If the players are correlation-loving, the dashed line curves inward as in Figure 1(b). What is more interesting in this case are the equilibrium implications. Recall that in a subgame perfect equilibrium (SPE), the individual rationality (IR) constraints of the players must hold *after every history*. In a game with standard preferences and a common discount factor, a major simplification arises in that any payoff can be attained by

[3]More precisely, we generalize *the part* of the folk theorem in Abreu *et al.* (1994) that assumes "observable mixtures." Dispensing with this assumption, as their full result does, will be done in a separate paper.

[4]See Epstein & Tanny (1980) for a general definition of correlation aversion, albeit in a static context.

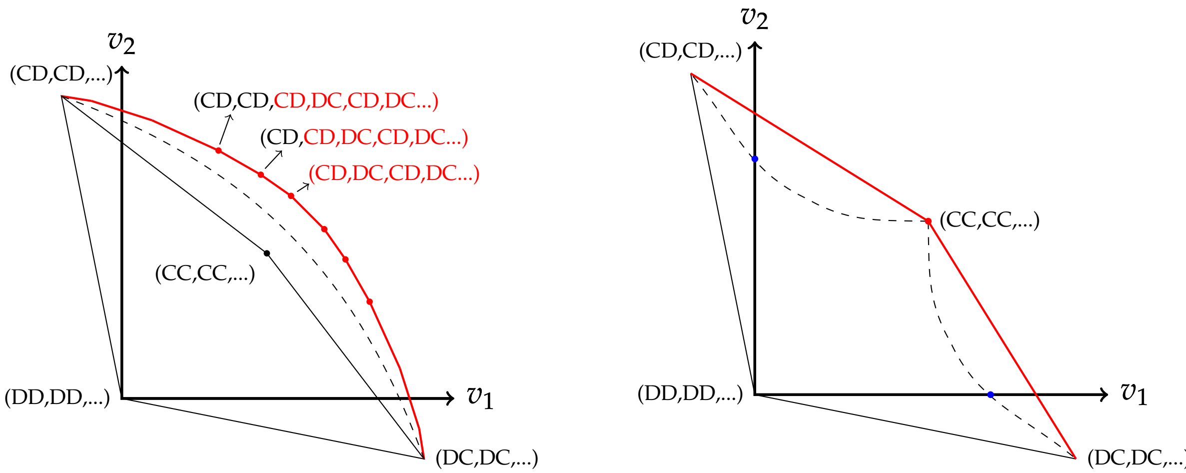

(a) Correlation Aversion/IMI

(b) Correlation Affinity/DMI

Figure 1: Feasible Set of Payoffs in the Prisoners' Dilemma. As usual, C stands for "cooperate" and D for "defect". Extreme points are associated with the play paths that generate them, while the dashed curves represent the payoffs from iid flips. (The acronyms IMI/DMI are defined in Section 1.2.)

a stationary strategy, that is, by an iid flip, which implies that it is enough to check that the IR constraints hold *ex ante*. But presently iid flips are inefficient. In fact, the payoff $v' = 0.5v(CD) + 0.5v(CC)$, which is efficient under correlation affinity, can be attained *only* by the corresponding one-time flip. But clearly the one-time flip cannot be sustained in a SPE: *no matter the punishment and the level of patience*, player 1 will deviate in the history in which he has to cooperate forever while the other player defects. As we show in Sections 7 and 8.2, such logic leaves $(CC, CC, ...)$ as the only efficient outcome sustainable in a SPE.

Next, we explain the concept of endogenous discounting and its own implications for repeated games.

1.2 Endogenous Discounting

Thinking of the rate of time preference as a fixed and immutable characteristic of agents has been common ever since Samuelson (1937) introduced the standard model of discounted utility. Yet Koopmans (1960) showed that the idea does not transcend the model: if time separability is relaxed, the rate of time preference, like any other marginal rate of substitution,[5] may vary with consumption. In a game, this means that differences in dis-

[5]Recall that the rate of time preference is traditionally defined as the marginal rate of substitution between consumption in two successive periods along a constant sequence $(c, c, ...)$. Equivalently, Koopmans (1960, p.307) defines it as the marginal valuation of *continuation utility* along a constant sequence. The latter for-

counting may emerge endogenously, as different players attain different outcomes during the game. Such heterogeneity allows us to prove that gains from intertemporal trade exist generically, *even when the players are correlation loving and ex ante symmetric.*[6] The logic is simple and comparable to the way an endowment point in an Edgeworth box may create differences in *marginal valuations* that deliver gains from trade even when the agents have identical preferences.

Differences in discounting have been of interest not only for their effects on ex ante welfare but their stark long-term implications. Consider a competitive setting with standard preferences. Formalizing a conjecture of Ramsey (1928), Becker (1980) and Rader (1981) showed that less patient agents are left *immiserated* in the long run as they borrow incessantly against any future capital they have. Taking on a strategic setting, Lehrer & Pauzner (1999) deduced a startling implication for the folk theorem: if immiseration pushes borrowers below their security level, they will renege on any promises to repay. The efficient level of trade is then unsustainable, *no matter the level of patience*.[7] To see what endogenous discounting brings up in this respect, it is helpful to introduce our model of recursive utility and a key assumption. Indeed, suppose the utility of a play path takes the form:

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

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 used, each player i computes the induced distribution over pure paths $(a^0, a^1, ...)$ and takes expectations in the usual way. We refer to the preferences thus defined as **Uzawa-Epstein (UzE)**. Their key feature, evidenced by (1), is that the rate at which agents discount future utility depends on the contemporaneous action profile. The standard model, which we refer to as one of exogenous discounting, arises as the special case in which the function $\beta_i : A \to (0,1)$ is constant.[8]

mulation is the one applicable in our context since, even though the action space is discrete, randomization convexifies the range of utility values.

[6]Curiously, the prisoners' dilemma must be excluded in the case of correlation affinity. See Section 7.1.

[7]As noted, the setting of Lehrer & Pauzner (1999) is one in which any payoff can be attained by a strategy that is *history-independent* on path. Their anti-folk implication is thus distinct, both formally and conceptually, from that of correlation affinity discussed in Section 1.1.

[8]See Uzawa (1968) and Epstein (1983) for the origins of the model. Epstein (1983, p.137) verifies the ordinal meaning of $\beta_i(a), a \in A$, as an (endogenous) rate of time preference.

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

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

Fisher (1930, p.72) was an early proponent of this assumption, noting that the needs of the present may bear more heavily on a person whose consumption is low. Friedman (1969, p.30), on the other hand, observed that DMI leads to "disequilibrium behavior" – a claim we address – and advocated for the polar case of **increasing marginal impatience (IMI)**. We investigate each case in turn.[9] At present, note 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*, thus boosting the "hedging benefits" of the iid flip. It follows that under IMI, the set of feasible payoffs expands due to intertemporal hedging; under DMI, the set of equilibrium payoffs may shrink. Yet, these conclusions do not leverage the *heterogeneity* in discounting ($\beta_i(a) \neq \beta_j(a)$) that can emerge under UzE preferences and whose implications, as we now explain, *reinforce* those of correlation sensitivity.

1.3 IMI and Intertemporal Cooperaion

Differences in discounting lead to a trade in which utility is backloaded by some players ("the lenders") and frontloaded by others ("the borrowers"). As we show in Theorem 3, this means that there are gains from intertemporal trade *above and beyond* what can be attained by stationary play (iid flips). Figure 1(a) provides an example in the context of the prisoners' dilemma. The analysis of the game, see Section 6, further reveals that *IMI simplifies the associated incentive constraints:* despite the non-stationarity of efficient outcomes, it suffices to check that no player wants to deviate ex ante. The key is that under IMI no player is able to maintain a higher level of patience; instead, the players' rates of time preference seesaw repeatedly along an efficient play path. This is also why in Figure 1(a) it is efficient to take turns defecting, behavior which we refer to as *intertemporal cooperation*.

[9]The debate as to the merits of each assumption is ongoing. For further discussion, see Lucas & Stokey (1984), Epstein (1987), Becker & Mulligan (1997), and Backus *et al.* (2004). More recently, Ifcher & Zarghamee (2011) show that a state of happiness may decrease the rate of time preference. While in their experiment happiness is manipulated *prior to* choice, the results may be loosely viewed as supportive of DMI. Others, such as Loewenstein (1996), have argued that the consumption of some goods (drugs, etc.) may bias people toward the present.

1.4 DMI and Immiseration

Things take a different turn under DMI: if differences in discounting emerge along an efficient path, they propagate over time. This creates an immiseration dynamic akin to that of Ramsey (1928) and Lehrer & Pauzner (1999). In particular, if immiseration pushes players below their security levels, the path is unsustainable in a SPE. The twist, formalized by Theorems 4 and 5, is that under endogenous discounting differences in time preference need not emerge, which is the case if and only if the players attain identical utility at each stage. Mirroring a conclusion reached in the discussion of correlation affinity, such logic implies that $(CC, CC, ...)$ is the only efficient outcome sustainable in a SPE of the prisoners' dilemma. More generally, Section 8.2 shows that the "immiserating effects" of DMI are stronger than those of correlation affinity in that they can "select" the symmetric outcome in a larger class of games.

1.5 Other Correlation Sensitive Preferences

Our choice of UzE preferences is motivated by their strong normative appeal. Relative to the standard model of choice, it is known from Epstein (1983) that they retain state separability while relaxing the less compelling assumption of time separability. Similarly, Chew & Epstein (1991) show that UzE preferences are the only ones that are recursive and *indifferent to the timing of resolution of uncertainty*. As failures of the latter property have been criticized[10] – they imply that agents are willing to pay for information that is of no instrumental value to them – we note that such behavior plays no part in our results. At the same time, Section 8 shows that our results regarding correlation attitudes are not limited to the UzE model and, in particular, extend to the preferences of Epstein & Zin (1989). The analysis of these preferences, which retain standard discounting, allows us to decouple and compare the implications of correlation sensitivity and endogenous discounting.

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 from 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 the strategic environment as simple as possible and assume per-

[10] See, for example, Epstein *et al.* (2014).

fect monitoring, the availability of public randomization, and finally, 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 (Borel measurable) function[11] 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 and let $\omega^t = (\omega_0^t, (\omega_i^t)_i)$. Given $t > 0$, a history $h^t = (\omega^0, \mathfrak{a}^0, \dots, \omega^{t-1}, \mathfrak{a}^{t-1})$ consists of the "mixtures" chosen in the past and the realized public and private signals. A strategy for player i is a sequence $\sigma_i = (\sigma_i^t)_t$ where σ_i^t maps each history h^t into a function $\mathfrak{a}_i^t$. 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 may often suppress the signals and speak of a mixed action $\alpha \in \Delta(A)$ being played after a history h^t.[12]

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

A strategy $\sigma \in \Sigma$ is **stationary** if for each i, there is a function $f_i : [0,1]^2 \to A$ such that $\sigma_i^t(h^t)[\omega_0^t, \omega_i^t] = f_i(\omega_0^t, \omega_i^t)$ for all $\omega_0^t, \omega_i^t, h^t, t$. If the functions f_i can depend on the time period t but not on history, we say that σ is **history-independent.** The path of play induced by a history-independent strategy is a sequence $\boldsymbol{\alpha} = (\alpha^0, \alpha^1, ...) \in (\Delta A)^\infty$ of mixed actions. A constant sequence $(\alpha, \alpha, ...)$, which we refer to as an **iid flip** and denote by α^{iid}, arises if the strategy is furthermore stationary. As we did in Section 1.1, given $\alpha \in \Delta(A)$, we distinguish the iid flip α^{iid} from the corresponding **one-time flip** α^{one}. The latter is an object in $\Delta(A^\infty)$ and arises as the path of play of a *history-dependent* strategy whereby the players randomize once, according to α, and repeat the realized pure action $a \in A$ forever after.[13]

[11]The topological and measure-theoretic conventions we employ are standard and suppressed until the Appendix.

[12]The only role of the private signal is to ensure that when minmaxing an opponent, each player can randomize privately. In particular, it is without loss of generality to assume that the private signals are not used on path.

[13]A general definition of *a path of play*, one that is consistent with any strategy σ and includes both iid and one-time flips as special cases, is given in the Supplemental Material. Presently, note that a sequence $(\alpha^0, \alpha^1, ...) \in (\Delta A)^\infty$ is distinguished from the induced product measure on A^∞. While the latter is enough to compute ex ante utility, the product measure does not encode the timing of randomization and this could matter for the SIR constraints of the players.

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, as shown in Kochov (2017, Lemma 3.4) for example, $v_i^{max} = v_i(a)$ for some $a \in A$. Finally, we assume that no player is indifferent among all strategies, which means that for every i, there are $a', a'' \in A$ such that $v_i(a'') > v_i(a')$.

3 "Sufficient Patience" when Patience is Endogenous

In a folk theorem, it is standard to vary the level of patience *while keeping the stage game fixed*. This delivers a family of repeated games that represent 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. Intuition can be gained from consumer 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 result, due to Epstein (1983), confirms this.

Proposition 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 the stage payoffs g_i lack 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 in terms of the repeated game. We then characterize the class of repeated games with endogenous discounting that share this ordinal input. 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 tells us 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 result provides an analogue for the case of endogenous discounting.

Proposition 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 $\lambda_i < 1$.*

In words, if two UzE preferences agree on $\Delta^{iid} \cup \Delta^{one}$, then it is without loss of generality to assume that they share the same stage payoffs, while their discount factors *must be* related in the specified linear fashion. The first implication is natural. To gain intuition for the second, note that preserving preferences on $\Delta^{iid} \cup \Delta^{one}$ can be viewed as preserving the players' correlation attitudes. But from Section 1.2, we know that the latter are intimately linked to properties of the discount factor. Leveraging this link, Proposition 2 shows that there is in fact a unique way to vary the β_i's without affecting correlation attitudes.

The discussion motivates the following approach to the folk theorem. Starting with a repeated game $(A, (g_i, \beta_i)_i)$, define for each $\lambda \in [0, 1)$ discount factors $\beta_{i\lambda}, i \in I$, and a repeated game Γ_λ by letting

$$\beta_{i\lambda} := \lambda + (1 - \lambda)\beta_i \quad \text{and} \quad \Gamma_\lambda := (A, ((1 - \lambda)g_i, \beta_{i\lambda})_{i \in I}). \tag{2}$$

Then, consider the equilibria of Γ_λ as $\lambda \nearrow 1$. There are a few things to clarify. First, going forward, whenever we speak of differences of patience across players, this will be meant in the context of *a fixed game* Γ_λ and in the *marginal* sense that $\beta_{i\lambda}(a) \neq \beta_{j\lambda}(a)$ for some realized action $a \in A$. Importantly, the parametrization in (2) has no bite in such contexts since, for a fixed λ, it does not restrict the function $a \mapsto \beta_{i\lambda}(a)$ in any way. Likewise, the parameter λ is meaningless when looking at a game Γ_λ in isolation. On the other hand, looking *across games*, we first see that $\lambda > \lambda'$ implies that $\beta_{i\lambda}(a) > \beta_{i\lambda'}(a)$ for all $a \in A, i \in I$. It is in this relational sense, and this sense only, that we will refer to λ as a "level of patience." Second, for any $a, a' \in A$, the ratio of marginal impatiences $(1 - \beta_{i\lambda}(a))/(1 - \beta_{i\lambda}(a'))$ is independent of λ. Complementing Proposition 2, this gives us another way of understanding how the parametrization in (2) restricts the *family* of games Γ_λ.[14]

Next, note that we have scaled the stage payoffs g_i by $(1 - \lambda)$. This is just a normalization. Akin to the "$(1 - \beta_i)$"-normalization one would do if β_i were exogenous, it ensures

[14]This is not unlike the preservation of the ratios $(1 - \beta_i)/(1 - \beta_j)$ of marginal impatiences *across players*, which is often imposed when formulating a folk theorem for games with fixed but heterogeneous rates of time preference. See Sugaya (2015, p.708) for a discussion of this and an interesting result.

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

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

Thus, given the normalization, we not only preserve the players' rankings on $\Delta^{iid} \cup \Delta^{one}$, but the associated utilities as well. In view of this, we will henceforth suppress the λ and write $v_i(\alpha^{iid})$ and $v_i(\alpha^{one})$. In the same vein, since all assumptions made in this paper will concern the players' rankings on $\Delta^{iid} \cup \Delta^{one}$, there should be no confusion to use Γ to mean either a single game $(A, (g_i, \beta_i)_i)$ or a family of games Γ_λ and to say that Γ satisfies a given assumption.

To make a final comment, observe that since UzE preferences are stationary and histories public, the players minmaxing an opponent have no use conditioning on history.[15] That is,

Lemma 1. *For each i, the minmax strategy against player i and i's best response can be chosen to be stationary.*

It follows immediately that minmax strategies can be chosen independently of λ and, given our normalization of utility, that the same is true of the players' security levels. Rescaling the original game $(A, (g_i, \beta_i)_i)$ as necessary, we can furthermore assume that the security levels of all players are zero,[16] a normalization we maintain throughout the rest of the paper.

4 A Folk Theorem

Subgame perfection requires that the threat of punishment be credible. Following Fudenberg & Maskin (1986), this is typically done by finding strategies that punish a deviation while *simultaneously* rewarding the players who carry out the punishment. This dual objective requires some heterogeneity in preference. In the case of standard preferences, a sufficient condition is that of **non-equivalent utilities (NEU)** of Abreu *et al.* (1994), which requires that no two players have identical preferences in the stage game, i.e, that for every i, j, there be $\alpha, \hat{\alpha} \in \Delta(A)$ such that $g_i(\alpha) > g_i(\hat{\alpha})$ and $g_j(\alpha) \leq g_j(\hat{\alpha})$. Under

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

[16]Letting $\hat{g}_i := g_i - \underline{v}_i(1-\beta_i)$, it follows from Proposition 1 that $(A, (\hat{g}_i, \beta_i)_i)$ is strategically equivalent to $(A, (g_i, \beta_i)_i)$ and all security levels are zero.

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: (i) no two players have identical preferences on Δ^{iid}, or (ii) no two players have identical preferences on Δ^{one}. The conditions are logically independent and each one is an extension of that of Abreu *et al.* (1994). We adopt the former because, as we highlight in Appendix B.2, it allows us to leverage some well-known decision-theoretic properties of UzE preferences.

Definition 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* $v_j(\alpha^{iid}) \leq v_j(\hat{\alpha}^{iid})$.

We comment on this condition again at the end of the section. At present, we turn attention to our folk theorem. To keep notation simple, we present a version of the result in which *on-path* behavior is restricted to be history-independent. Such behavior can be identified with a sequence $(\alpha^0, \alpha^1, ...) \in (\Delta(A))^\infty$ of mixed actions, which we refer to as a **play path** or simply a **path**.[17] Notably, both iid flips and pure paths $(a^0, a^1, ...)$ are covered, and this will suffice for our analysis of intertemporal trade. The general formulation of our folk theorem, which allows history-dependent behavior on path, such as a one-time flip, poses no conceptual difficulties and is presented in the Supplemental Material. Currently, let SIR^ε_λ be the set of all **ε-sequentially individually rational (ε-SIR) paths** $\boldsymbol{\alpha} \in (\Delta(A))^\infty$ in Γ_λ, i.e., all paths such that $v_{i\lambda}({}_t\boldsymbol{\alpha}) \geq \varepsilon$ for all i, t, where ${}_t\boldsymbol{\alpha} := (\alpha^t, \alpha^{t+1}, ...)$.

Theorem 1. *Assume NEU. For every* $\varepsilon > 0$, *there exists* $\underline{\lambda}$ *such that for all* $\lambda > \underline{\lambda}$, *every path* $\boldsymbol{\alpha} \in SIR^\varepsilon_\lambda$ *can be supported in a SPE of the game* Γ_λ.[18]

The NEU condition in our folk theorem is generic as can be seen from Lemma A1 in the appendix, which characterizes the condition in terms of the utility representations (g_i, β_i). Presently, we note that the condition is not without loss of generality. 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 $\Delta^{iid} \cup \Delta^{one}$, and NEU fails. Yet, no two players have identical preferences on the space Σ of all strategies. In the context of exogenous discounting, Chen & Takahashi (2012) are able to prove a folk theorem under such a condition, which they term **Dynamic NEU**.[19] Doing the same for the case of endogenous discounting is an open problem.

[17] We reserve the slightly different phrasing, *a path of play*, for the general types of on-path behavior discussed in the Supplemental Material.

[18] As in Fudenberg & Maskin (1986), NEU is not required in two-player games.

[19] In the present context, Dynamic NEU can be stated as: for every λ and i, j, there exist $\sigma, \hat{\sigma} \in \Sigma$ such that $v_{i\lambda}(\sigma) > v_{i\lambda}(\hat{\sigma})$ and $v_{j\lambda}(\sigma) \leq v_{j\lambda}(\hat{\sigma})$

5 Gains from Intertemporal Trade: A Necessary Condition

Define the **(strong) Pareto frontier** of a set $X \subset \mathbb{R}^n$ as the subset of $x \in X$ for which there is no $x' \in X$ such that $x' \gg x$ ($x' > x$).[20] When x is on the Pareto frontier of X, we also say that x is **efficient in** X. Given $y \in \mathbb{R}^n$, we write $y >^* X$ if $y \gg x$ for some $x \in X$ and $x' \geq y$ for no $x' \in X$. One can visualize such y as lying "above" the Pareto frontier of X. Given a game $(A, (g_i, \beta_i)_i)$, we say that a feasible payoff v is **efficient** if v is efficient in the space of all feasible payoffs and that v is a **gain from intertemporal trade**, or simply **a gain from trade**, if

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

In the rest of the paper, we study when gains from intertemporal trade exist and whether they can be sustained in a SPE. First, we need to explain why V^{one} is the appropriate benchmark against which to define gains from trade. In Section 3, we already explained why a more obvious benchmark – the set of stage payoffs – does not work. We also observed that V^{one} reduces to the set of stage payoffs when discounting is exogenous, making it a possible replacement. Presently, we add that recursive preferences cannot be distinguished from standard ones when attention is restricted to one-time flips.[21] This makes V^{one}, and not the payoff set from iid flips, the right benchmark against which to measure the effects of recursive preferences, including any ensuing opportunities for intertemporal trade.

As defined, gains from trade do not exist in every game. To see this, recall from Section 2 that for every i, there is $a \in A$ such that $v_i(a) = v_i^{max}$. If a can be chosen independently of i, then $v(a) \geq v_\lambda(\sigma)$ for every $\sigma \in \Sigma$ and λ. Thus, a necessary condition for existence is that there be no $a \in A$ such that $v_i(a) = v_i^{max}$ for every i. Next, we show that under IMI this condition, which we call **Conflict of Interest (CI)**, is not only necessary but sufficient as well.

6 Increasing Marginal Impatience

Recall that player i's preferences satisfy IMI if for every $a, a' \in A$,

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

[20] Given $x, x' \in \mathbb{R}^n$, we write $x' > x$ if $x' \geq x$ and $x' \neq x$; if $x'_i > x_i$ for all i, we write $x' \gg x$.

[21] This is true not only for the UzE preferences we study, but in fact for all recursive preferences as defined in Kreps & Porteus (1978).

If all players' preferences satisfy IMI, we say simply that IMI holds. IMI captures a situation in which good outcomes create a "live in the moment" attitude that biases agents toward the present. In Section 1.2, we gave intuition why, as a result, the players prefer the "hedging benefits" of iid flips to the "smoothing benefits" of one-time flips. To be more formal, 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. The next result confirms that IMI implies correlation aversion and that the converse is true generically.

Proposition 3. *If player i's preferences (g_i, β_i) satisfy IMI, then the player is correlation-averse. The converse 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, a and a', 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 if $v_i(a) = v_i(a')$. If $v_i(a) \neq v_i(a')$, the inequality is strict and follows from IMI. If there are more than two actions in the support of α, the proof is by induction. The converse is proved analogously.[22] □

We can now state our first result establishing the existence of gains from trade.

Theorem 2. *Assume IMI and CI. Then, there is $\alpha \in \Delta(A)$ such that $v(\alpha^{iid}) >^* V^{one}$. If, in addition, NEU holds and V^{one} contains some payoff $v \gg 0$, then α can be chosen so that α^{iid} can be sustained in a SPE for all λ large enough.*

Under correlation aversion, each iid flip α^{iid} Pareto dominates its counterpart α^{one}.[23] In general, this is not sufficient to deliver gains from trade since each α^{iid} can itself be dominated by another one-time flip $\hat{\alpha}^{one}$. This is where CI kicks in. Specifically, Lemma C18 in the appendix shows that CI holds if and only if the Pareto frontier of V^{one} has at least one "downward-sloping" segment, by which we mean a face which is not a singleton and which is orthogonal to some strictly positive direction in $\mathbb{R}^n$. Such segments capture a zero-sum situation in which any movement along the segment benefits some player at the expense of another. As a result, no $v \in V^{one}$ in the interior of such a segment can be strictly dominated by the payoff of another one-time flip. At the same time, any such v is *strictly* dominated by the corresponding iid flip. Put together, these observations show that gains from trade exist. That these gains can be sustained in a SPE follows from our folk theorem.

[22] Epstein (1983) was the first to observe the relationship between IMI and correlation aversion, though without providing details.

[23] Given, $x, y \in \mathbb{R}^n$, x Pareto dominates y if $x \geq y$; x strictly Pareto dominates y if, in addition, $x \neq y$.

We should emphasize that the *existence* of gains from trade established by Theorem 2 does not depend on the specifics of the UzE model. All that matters, apart from CI, is that the players be correlation averse and their preferences on Δ^{one} be consistent with expected utility. Existence can thus be deduced for other preferences. Notably, as we show in Section 8, this is the case for the preferences of Epstein & Zin (1989), which satisfy the expected-utility requirement automatically and are correlation averse under the commonly made assumption that agents care more about risk than intertemporal smoothing.[24]

Because the preferences of Epstein & Zin (1989) retain standard discounting, the analysis in Section 8 will also make clear a flipside to the generality of Theorem 2: the established existence (due to intertemporal hedging) does not leverage the *heterogeneity* in marginal patience, $\beta_{i\lambda}(a) \neq \beta_{j\lambda}(a)$, that can emerge under UzE preferences. If so, less patient players want to "borrow" or frontload utility, while their more patient counterparts want to "lend" or backload utility. Intuitively, this means that there must be gains from trade *above and beyond* what can be attained by stationary play (iid flips). Our next result confirms this, while also showing that the additional gains from trade persist even in the limit as $\lambda \nearrow 1$. To state the result, let V^{iid} be the payoff set from iid flips and let $conv(V^{iid})$ be its convex hull.

Theorem 3. *Consider a symmetric game satisfying IMI and CI. For each λ, there is $\boldsymbol{\alpha}_\lambda \in (\Delta(A))^\infty$ such that $v_\lambda(\boldsymbol{\alpha}_\lambda) >^* conv(V^{iid})$. In addition, the paths $\boldsymbol{\alpha}_\lambda$ can be chosen so that $\lim_{\lambda \nearrow 1} v_\lambda(\boldsymbol{\alpha}_\lambda) >^* conv(V^{iid})$.*

We note that at this level of generality we do not know if the payoffs $v >^* conv(V^{iid})$ can be sustained in a SPE. While a relatively simple sufficient condition can be deduced from the proof of the theorem,[25] we prefer to shift attention to the analysis of a specific game. This allows us to characterize efficient outcomes explicitly and draw sharper insights about the effects of IMI on incentives.

6.1 The Prisoners' Dilemma

Consider a symmetric prisoners' dilemma game. Thus, $A_1 = A_2 = \{C, D\}$, where as usual C stands for "cooperate" and D for "defect," and the preferences (g_1, β_1) of player

[24] The requirement that preferences be consistent with expected utility on Δ^{one} is used to show that V^{one} is a convex set. See Lemma C18.

[25] The condition is that there be some $\alpha \in \Delta(A)$ such that (i) $v(\alpha^{iid})$ is on the Pareto frontier of V^{iid}, (ii) $v(\alpha^{iid}) \gg 0$, and (iii) $v_i(\alpha^{iid}) > v_j(\alpha^{iid})$ and $\beta_{i\lambda}(\alpha) < \beta_{j\lambda}(\alpha)$ for some i, j. Observe that IMI does not imply the equivalence of the latter two inequalities as they concern mixed rather than pure actions.

1, say, are such that:

$$\begin{aligned} &v_1(DC) > v_1(CC) > v_1(DD) > v_1(CD), \quad \text{and} \\ &v_1(CC) > 0.5v_1(CD) + 0.5v_1(DC). \end{aligned} \tag{3}$$

In particular, the set V^{one} looks as in Figure 1. We remark that these inequalities are stated in terms of V^{one}, *rather than stage payoffs*, for the reasons outlined in Section 3. We also note that the inequality in (3), which posits that $v(CC)$ be efficient in V^{one} and, in fact, an extreme point thereof, is not typically viewed as a defining feature of the prisoners' dilemma. We impose it because it helps us highlight a key implication of endogenous discounting, namely, that $v(CC)$ may be Pareto dominated *even though* it is efficient in V^{one}. Before we delve into these specifics, however, it is helpful to recall that under UzE preferences $v_i(\sigma)$ is defined as an expectation over the utility of pure paths. It follows that the feasible set of payoffs in any game $(A, (g_i, \beta_i)_i)$ is equal to the convex hull of $v(A^\infty)$. To characterize the Pareto frontier of the feasible set, it is thus sufficient to characterize the set of efficient pure paths $(a^0, a^1, ...)$, where a pure path is **efficient** if it yields an efficient payoff.

6.1.1 Efficient Outcomes

Figure 2 depicts two possibilities for the Pareto frontier of the prisoners' dilemma.[26] In the more interesting case on the right, $(CC, CC, \dots)$ is dominated. Instead, the efficient and (almost) symmetric outcome is to take turns defecting. To give some intuition for these scenarios, consider an efficient play path that begins with CD. Under IMI, this renders player 2 less patient than his opponent: $\beta_2(CD) < \beta_1(CD)$. Then, unless the path in question is $(CD, CD, ...)$, a corner of the frontier, efficiency requires that the difference in marginal patience be baked into the path and, in particular, that player 1's utility be backloaded. But this means that play must eventually switch from CD to an action profile that is more favorable to player 1. If, as in Figure 2(a), it is efficient to switch to CC, the players' rates of time preference are equalized and there are no further switches. On the other hand, if the initial difference in marginal patience is sufficiently large, and $v_2(CC)$ sufficiently low, switching to CC may not be enough of a "repay" to player 1. Then, play switches to DC as in Figure 2(b). But this means that the roles of the players reverse again: as DC renders player 1 less patient, play must switch back to CD, and so on ad

[26]For some parameters, it is possible that the paths $(CC, CC, ...)$ and $(CD, DC, ...), (DC, CD, ...)$ are *simultaneously* efficient. Such cases, while not depicted in Figure 2, are fully described in the Supplemental Material.

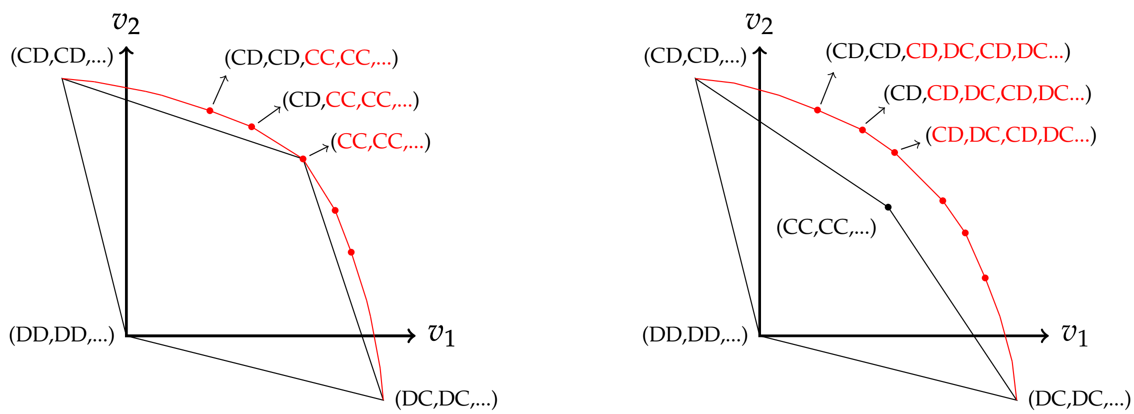

(a) Intratemporal cooperation

(b) Intertemporal cooperation

Figure 2: Two forms of cooperation under IMI

infinitum.[27]

In sum, the "live in the moment" effect of IMI leads to a "reversal of fortunes:" the permanent switch from CD (or DC) to CC in Figure 2(a) or the continual alternation between CD and DC in Figure 2(b). Next, we explain how these reversals affect equilibrium incentives.

6.1.2 The Incentives to Deviate

Lehrer & Pauzner (1999) showed forcefully that as differences in discounting lead to non-stationary play, the incentives to deviate will vary over time. E.g., "borrowers" want to deviate when it's time to "repay," not before. This is why, as our folk theorem stipulated, the IR constraints must be checked *at each point in time*, not merely ex ante. But Figure 2 suggests that IMI may deliver a non-trivial simplification in this regard: since efficient play converges to the symmetric outcome over time, all IR constraints except the initial one become slack.

A minor caveat arises in the context of Figure 2(a).[28] Consider a pure path such as

$$(CC, CC, CD, CC, CC, ...) \tag{4}$$

in which a single defection occurs at some period $t \neq 0$. Such paths *are* efficient in the

[27] To understand why the initial switch, unlike subsequent ones, need *not* occur immediately, think of the length of the initial phase during which CD is played as the fraction of "trade surplus" captured by player 2. Then, given a positive rate of time preference, it is intuitive that the division of surplus be settled through a single "transfer" at the start of play, that is, through a single, initial phase during which CD is repeated.

[28] An analogous caveat arises in the context of the "mixed cases" (see ft.26) not shown in Figure 2 but is handled in the same way. See the full proof of Proposition 4 given in the Supplemental Material.

context of the figure but were not depicted as their payoffs are not an extreme point of the frontier. Clearly, the trouble with such paths is that the IR constraints might fail at the time of defection, *even if* they hold ex ante. However, since there is only a single defection along the entire path, it is also clear that a modicum of patience should rule out this possibility. Taking advantage of the specifics of the prisoners' dilemma, we can in fact calculate the incentives to deviate explicitly and find the exact level of patience λ that would deter deviations. Thus, under grim-trigger strategies, the maximum lifetime utility of any player who deviates is $(1-\lambda)g_1(DC)$. (Recall that we normalize utility so that $v_i(DD) = 0$ for each i.) On the other hand, the minimum lifetime utility of any player along one of the "troublesome" paths is $v_{1\lambda}(CD, CC, CC, ...)$. Stacking the two together, let λ' be such that

$$v_{1\lambda'}(CD, CC, CC, ...) = (1-\lambda')g_1(DC),$$

and let $\underline{\lambda} = \max\{0, \lambda'\}$. We deduce that:

Proposition 4. *Assume IMI. For every $\lambda > \underline{\lambda}$, every efficient pure path $\mathbf{a} \in A^\infty$ such that $v_\lambda(\mathbf{a}) \gg (1-\lambda)g_1(DC)$ can be sustained in a SPE of the prisoners' dilemma.*

Whether IMI can simplify the equilibrium analysis of intertemporal trade in general games is an interesting open problem.

7 Decreasing Marginal Impatience

Our analysis of DMI is focused on two-player symmetric games. Reversing the order of Section 6, we begin by studying the long-term implications of intertemporal trade.

Theorem 4. *Consider a two-player, symmetric game satisfying DMI. For every λ, every efficient path $\mathbf{a} = (a^0, a^1, ...) \in A^\infty$ is such that either (i) there are $i \in \{1,2\}$ and T such that $v_{i\lambda}({}_t\mathbf{a}) = v_i^{max}$ for all $t \geq T$, or (ii) there is $a \in A$ such that $v_1(a) = v_2(a)$ and $v(a^t) = v(a)$ for all $t \geq 0$.*

To understand the statement of the theorem, note first that in a symmetric two-player game maximizing the utility of one player implies that the other is at, or below, their minimum along the strong Pareto frontier. In this sense, case (i) of the theorem reflects the "immiseration" of one of the players. The only alternative, as described by case (ii) of the theorem, is when the path yields identical utility across players and time periods. To see why efficiency delivers such extreme implications, suppose $v_i(a^t) > v_j(a^t)$ at some t. Under DMI, this means that player i attains a higher level of patience. Efficiency then requires that i's utility be backloaded, which, invoking DMI again, implies that i will sustain the higher level of patience as the game progresses. As anticipated by Friedman (1969, p.30), and unlike "the reversal of fortunes" we deduced in Section 6.1.1, this creates an immiseration dynamic akin to that of Ramsey (1928) and Lehrer & Pauzner (1999). The twist, highlighted by our theorem, is that under endogenous discounting differences in patience need not emerge in the first place. This is so if and only if the players attain identical utility at any given point in time. Given that the path is efficient, this further implies that utility must be the identical across time periods, leading to case (ii) of the theorem.

Theorem 4 has obvious implications for equilibrium behavior: as in Lehrer & Pauzner (1999), if immiseration pushes players below their security levels, they will deviate. An efficient path $\mathbf{a} \in A^\infty$ can thus be sustained *only if* it is symmetric as in case (ii) of the theorem. Two questions must be addressed in this context. First, what about mixed strategies? Can randomization help the players avoid immiseration and sustain an efficient payoff v such that $v_1 \neq v_2$? Our next result shows that the answer is no. Intuitively, efficiency does not allow the players to place positive probability on a pure path that is inefficient. To state the result, let $v_{i\lambda}(\sigma \,|\, h^t)$ denote i's payoff from $\sigma \in \Sigma$ in the subgame given history h^t, and let

$$V_\lambda^{SIR} := \{v_\lambda(\sigma) \,:\, \sigma \in \Sigma \text{ such that } v_\lambda(\sigma \,|\, h^t) \geq 0 \text{ for all } t \text{ and on path histories } h^t\}.$$

Theorem 5. *Consider a two-player symmetric game satisfying DMI and such that $v_1(a) < 0$ whenever $v_2(a) = v_2^{max}$. For every λ, a payoff $v \in V_\lambda^{SIR}$ is efficient only if $v_1 = v_2$ and $v = v(a)$ for some $a \in A$.*

Second, we note that in many games there could be no action $a \in A$ that yields a symmetric and efficient payoff. In such cases, immiseration will be an unavoidable consequence of efficiency. More subtly, immiseration might be unavoidable *even if* some $a \in A$ yields a symmetric payoff which is efficient in V^{one}. This is because, given the possibility of intertemporal trade, efficiency in V^{one} does not guarantee efficiency in the full set of payoffs. However, we will see momentarily that $v(CC)$ *is* efficient in the prisoners' dilemma under DMI. Taking this for granted for now, we have the following corollary of Theorem 5:

Corollary 1. *In the prisoners' dilemma under DMI, if an efficient payoff v can be sustained in a SPE, then $v = v(CC)$.*

7.1 Sustaining *Some* Gains from Trade

Theorems 4 and 5 raise the question whether *some* gains from trade, but not the efficient level, can be sustained. One must also address the question of existence. Because the players are correlation-loving under DMI, intertemporal hedging is not a factor. On the other hand, thinking of the heterogeneity in marginal patience that can arise under UzE preferences, it becomes clear that CI is no longer sufficient for existence. Consider the prisoners' dilemma. The actions that generate differences in discounting are CD and DC. If CD is played, the logic behind Theorem 4 tells us that player 2's utility should be backloaded. But since CD is already as good as it gets for player 2, this can only be done by repeating CD forever after, *leading to a constant play path*. In a sense, intertemporal trade takes a trivial form and does not generate any gains. The feasible set of payoffs is thus V^{one}, which incidentally proves the efficiency of $v(CC)$ asserted in the lead-up to Corollary 1.

For gains from trade to exist, one thus needs an action $a \in A$ that generates differences in discounting without automatically maximizing the utility of any player. The action cannot be "too" inefficient either as this will offset any gains from the induced heterogeneity. Since the exact "permissible" level of inefficiency will depend on the game, we focus on a *sufficient condition* that simply requires $a \in A$ to be efficient in the space of one-time flips.

Definition 2. *A two-player game Γ satisfies **Richness** if there is some $a^r \in A$ such that $v(a^r)$ is efficient in V^{one} and $v_i^{max} > v_i(a^r) > v_j(a^r)$ for some $i \in \{1,2\}$ and $j \neq i$.*[29]

Our next result establishes the existence of gains from trade and, under an obvious strengthening of Richness and a sufficiently high λ, that those can be sustained in a SPE.

Theorem 6. *Consider a two-player, symmetric game satisfying DMI and Richness. For each λ, there exists a feasible payoff v_λ such that $v_\lambda >^* V^{one}$. Additionally, if the action profile a^r in the statement of Richness is such that $v(a^r) \gg 0$, there exists $\underline{\lambda}$ such that for all $\lambda > \underline{\lambda}$, there exist payoffs $v_\lambda >^* V^{one}$ that can be sustained in a SPE.*

We give a sketch of the proof. By the symmetry of the game, 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 3(a), 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 trade, i.e., $v_\lambda(a^r_{-T}, a^*_T) >^*$

[29] In symmetric two-player games, Richness implies CI.

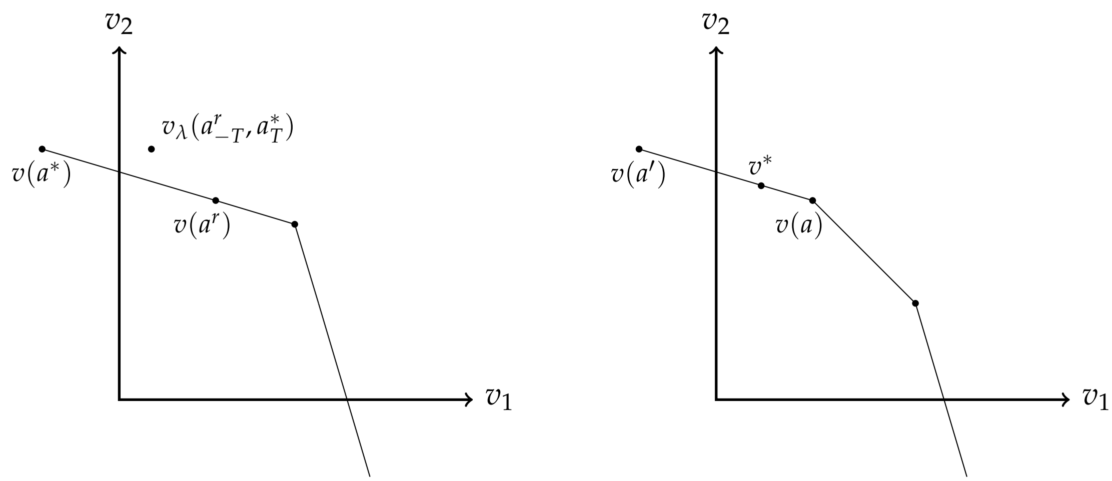

(a) Richness and gains from trade

(b) Correlation Affinity under EZ preferences

Figure 3

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 trade. To obtain a first-best outcome, the logic behind Theorem 4 tells us that 2's utility should continue to rise until it is fully maximized. But, as we have seen, this may violate the IR constraints of player 1, which is why the path (a^r_{-T}, a^*_T) sacrifices some efficiency by switching back to $a^r \in A$ after $a^* \in A$ is played. While the switch is extreme, it helps us prove the second part of the theorem. In particular, suppose $v(a^r) \gg 0$ and recall that, by construction, $v_1(a^*) < v_1(a^r)$. Then, observe that for *any* λ and T, (a^r_{-T}, a^*_T) is SIR if and only if $(a^*, a^r, a^r, ...)$ is IR, that is, whether or not (a^r_{-T}, a^*_T) is SIR does not depend on T. Picking λ so that $(a^*, a^r, a^r, ...)$ is IR and then a sufficiently high T, we can thus ensure that the path (a^r_{-T}, a^*_T) is *both* a SPE outcome and delivers gains from trade.

8 Epstein-Zin Preferences

In this final section, we shift attention to another prominent class of recursive preferences, those of Epstein & Zin (1989). The preferences, henceforth **EZ**, are sensitive to autocorrelations but retain standard discounting. As such, the analysis allows us to decouple the effects of correlation attitudes from those of endogenous discounting. Of course, EZ preferences come with their own set of assumptions and may generate insights that are of independent interest. We begin with two tweaks to our setup needed to accommodate EZ preferences.

First, recall that 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, these two aspects of preference are fully independent of one another. Accordingly, we need two "cardinal scales" and a way to convert between them. A simple way to accomplish this is to assume that stage outcomes take the form of a physical, infinitely divisible good in terms of which one can compute certainty equivalents. In particular, let $g_i : A \to \mathbb{R}_+$ describe i's consumption levels as a function of action profiles $a \in A$. Then, assuming for the sake of simplicity that the game is symmetric, let $C := conv(g_i(A))$ 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.

Second, because EZ preferences are sensitive to the timing of resolution of uncertainty, we cannot define them by looking at the distribution over pure paths induced by a strategy σ; instead, we must encode the timing of resolution of uncertainty and look at the induced **infinite probability tree** or what Kreps & Porteus (1978) call **a temporal lottery**. To define the concept, ignore for the moment the possibility of randomization in the initial period. An 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 the infinite probability trees that may prevail in period 1. In other words, the space D of such trees can be defined as the unique (up to a homeomorphism) set satisfying $D = A \times \Delta(D)$. Since randomization in the initial period is of course possible, the set of all infinite probability trees is $\Delta(D)$ rather than D.

With the preceding in mind, define $v_i : D \to \mathbb{R}$ as the solution to the equation:

$$v_i(a, \mu) = rs^{-1}\Big((1-\beta)s(g_i(a)) + \beta sr^{-1}(\mathbb{E}_\mu v_i)\Big) \quad \forall (a, \mu) \in A \times \Delta(D) = D. \tag{5}$$

Above, $\beta \in (0, 1)$ is the common and fixed discount factor, rs^{-1} is the composition of r and s^{-1}, and $\mathbb{E}_\mu v_i$ is the expectation of v_i with respect to μ.[30] To understand the equation, note that we have set it up so that v_i is denominated in r-utils. Then, starting backwards, we apply sr^{-1} to the continuation utility $\mathbb{E}_\mu v_i$ so as to convert it into s-utils. We then aggregate across time by computing the discounted average with current utility. Applying rs^{-1} to that average converts utility back into r-utils. Finally, having defined $v_i : D \to \mathbb{R}$

[30] Here and throughout, it is implicitly assumed that r, s are such that equation (5) *has* a solution. A well-known example is when r, s are homothetic. Other cases can be found in Marinacci & Montrucchio (2010).

in terms of r-utils, the utility of some $\mu \in \Delta(D)$ is simply $\mathrm{E}_\mu v_i$.

Some examples may cast further light on this construction. First, consider the utility of a pure path:

$$v_i(a^0, a^1, \ldots) = rs^{-1}\big((1-\beta)\textstyle\sum_t \beta^t s(g_i(a^t))\big), \quad (a^0, a^1, \ldots) \in A^\infty. \tag{6}$$

Since rs^{-1} is just an increasing transformation, we see that preferences over pure paths conform to the standard model of discounted utility, with s capturing the desire to smooth consumption over time. On the other hand, the utility of a one-time flip is given by:

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

We see that the curvature of r reflects risk aversion and that $v_i(\alpha^{one})$ coincides with the stage payoff of α. The latter is the main advantage of denominating $v_i : D \to \mathbb{R}$ in r-utils: it ensures that payoffs in the repeated game are in the same units as stage payoffs.

8.1 Correlation Aversion

In Section 1.1, we gave intuition why a player who is more concerned with risk than intertemporal smoothing will be correlation-averse. EZ preferences allow us to formalize this observation. In fact, Proposition 5 below shows that the converse implication is also true generically. First, some notation. We 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. Such a game is taken as given throughout the rest of the section. For one direction of the result, we need to allow A to be an infinite set and indeed such that $g_i(A) = C$ for some i, the choice of i being immaterial by the symmetry of the game. If so, we say that the game is **connected**.[31] Also, r is said to be **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 $c \in C$. As is well understood, such a relationship between r and s expresses the fact that risk matters more than intertemporal smoothing.

Proposition 5. *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.*[32]

[31] As we continue to assume that no player is indifferent among all strategies, C cannot be a singleton.

[32] The first claim of the proposition may be viewed as a corollary of Strzalecki (2013, Theorem 6). The second claim is new and made possible by the fact that we deal with a more special class of preferences than Strzalecki (2013). See also Stanca (2022), who contemporaneously gives sufficient conditions for a *more special* class of EZ preferences to exhibit a *stronger* notion of correlation aversion.

Proposition 5 gives us another instance in which we can invoke Theorem 2 and obtain an expansion of V^{one} due to intertemporal hedging. Taking advantage of the specifics of the EZ model, we now go a step further and characterize the Pareto frontier of the feasible payoff set. First, 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. *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 of the game is equal to the Pareto frontier of $V^{pure} := \{v(a^0, a^1, ...) : (a^0, a^1, ...) \in A^\infty\}$. In addition, a payoff on the strong Pareto frontier of V can be attained only by a trivially randomized μ.*

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 the restriction on β, recall from (6) that the utility of a pure path is the usual discounted average (modulo the rs^{-1} change of units). It follows from Fudenberg & Maskin (1991, Lemma 1) that the set of payoffs from pure paths is convex whenever $\beta > 1 - |A|^{-1}$. This richness allows us to prove that *every* $\mu \in \Delta(D)$ is Pareto dominated by *some* pure path.[33]

To draw some additional insights, recall from Fudenberg & Maskin (1991) that V^{pure} is independent of β for all $\beta > 1 - |A|^{-1}$. In fact,

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

On the other hand, V^{iid} varies with β, even as β nears 1. This is because EZ preferences, unlike their UzE counterpart, are sensitive to the late resolution of uncertainty implied by iid flips and this sensitivity itself depends on β. In characterizing the latter dependence, a major difficulty is that there is no closed-form expression for the utility of iid flips under EZ preferences. However, assuming preferences to be homothetic, i.e., that $r(c) = c^\gamma$ and $s(c) = c^\rho$ for some $\gamma, \rho \in (0, 1)$, Al-Najjar & Shmaya (2019) recently obtained just such an expression for $\lim_{\beta \nearrow 1} v_i(\alpha^{iid})$. It follows immediately from their formula that the

[33]If $\beta \leq 1 - |A|^{-1}$, non-trivial randomization can be efficient. In such cases, it matters whether the players prefer early or late resolution in the sense of Kreps & Porteus (1978). In the former case, all randomization takes place in period 0; in the latter, the exact timing depends on the specific μ.

Pareto frontier of V^{iid} converges to that of the feasible set as $\beta \nearrow 1$. This occurs because in the limit of perfect patience, iid flips are no longer penalized for the risk they carry – in a sense, intertemporal hedging works perfectly in the limit – or for the late resolution of uncertainty. Presently, the convergence is noteworthy as it stands in sharp contrast to the conclusions of Theorem 3, where the endogeneity of discounting meant that iid flips remain inefficient even as $\lambda \nearrow 1$.

8.2 Correlation Affinity

In direct juxtaposition to Theorem 7, our next result characterizes efficient outcomes when the players care more about intertemporal smoothing than risk.

Theorem 8. *If s is a strictly concave transformation of r, then the Pareto frontier of the set V of feasible payoffs of the game is equal to the Pareto frontier of V^{one}.*

As with Theorem 7, if we exclude some non-generic cases, an efficient payoff can be attained *only* by a one-time flip. We omit the details, which are not difficult to fill in.[34] Instead, it is useful to compare Theorem 8 to its "cousins" under DMI, Theorems 4 and 5. Notably, the comparison reveals that in the absence of (endogenous) differences in time preference and the associated immiseration dynamics, the *anti-folk* implications of correlation affinity do not go as far as they did under UzE preferences. Consider first the payoff v^* in Figure 3(b). While v^* is individually rational, the one-time flip that attains it must put positive probability on $(a', a', ...)$ and, since $v_1(a') < 0$, neither the one-time flip nor v^* can be sustained in a SPE. On the other hand, no payoff in the middle segment of the frontier entails a similar failure of SIR. Thus, unlike the case of DMI, a *non-trivial range* of efficient payoffs meet the necessary conditions for subgame perfection. Of course, no such range exists in the prisoners' dilemma, leading to the following analogue of Corollary 4.

Corollary 2. *Consider a prisoners' dilemma such that s is a strictly concave transformation of r. If an efficient payoff v can be sustained in a SPE, then $v = v(CC)$.*[35]

Appendix

Given a repeated game $(A, (g_i, \beta_i)_i)$, the finite set A is endowed with the discrete topology, making it a compact metrizable space. Product spaces are endowed with the product

[34] 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' delivers the payoff $r(g(a))$ as well.

[35] The prisoners' dilemma is as defined in Section 6.1.

topology, and all functions are understood to be Borel measurable. Given a topological space X, $\Delta(X)$ denotes the space of Borel probability measures endowed with the weak* convergence topology. We remark that these technical assumptions, while listed here for ease of reference, play no role until the proof of Theorem 5 and, more significantly, until the extension of our folk theorem to general paths of play given in the Supplemental Material.

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

A Proof of Proposition 2

To begin, we state a lemma, due to Chew (1983), which characterizes when two UzE preferences agree on the set of iid flips.

Lemma A1. *(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 $\beta_j = 1 - sg_i - t(1-\beta_i)$.*

By Lemma A1, if (g_i, β_i) and (g'_i, β'_i) agree on Δ^{iid}, then there are q, r, s, t, $qt > sr$, such that $g'_i = qg_i + r(1-\beta_i)$ and $1 - \beta'_i = sg_i + t(1-\beta_i)$. If, furthermore, (g_i, β_i) and (g'_i, β'_i) agree on Δ^{one}, there are $\theta > 0$ 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 $\beta'_i = \lambda_i + (1-\lambda_i)\beta_i$. To prove the other assertion, note that $t > 0$ and $qt > sr = 0$ imply $q > 0$. By Proposition 1, for every $\hat{\theta} > 0$ and $\hat{\gamma}$, (g'_i, β'_i) 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 1

B.1 Payoff Asymmetry

Each pair (g_i, β_i) induces a preference relation $\succeq_i$ on $\Delta(A)$ represented by the utility function $\alpha \mapsto v_i(\alpha)$. If $\beta_i : A \to (0,1)$ is constant, $\succeq_i$ is a standard expected utility preference. If not, then $\succeq_i$ belongs to the more general class of **weighted-utility preferences** studied in Chew (1983). We begin with some preliminary observations regarding such preferences.

Lemma B2. *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 B3. *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. If not, the proof can be deduced from Figure 1 in Chew (1983). Thus, 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 Abreu *et al.* (1994).

Lemma B4. *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 B3, 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 B3, $\succeq_1$ 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 reminding the reader of two useful properties of UzE preferences.[36] The proofs are obvious and omitted. Fix some $i \in I$ and let $\alpha^0, \alpha^1, ..., \alpha^K \in \Delta(A)$ be such that $v_i(\alpha^k) \leq v_i(\alpha^{k+1})$ for every $k = 0, ..., K-1$. Lemma B5 shows that player i prefers more beneficial actions to be played first.

Lemma B5. *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 B6. *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 $(\alpha^0, \alpha^1, ...) \in (\Delta(A))^\infty$,

$$v_i(\alpha^0, \alpha^1, ...) = (1 - \beta_i(\alpha^0))v_i(\alpha^0) + \beta_i(\alpha^0)v_i(\alpha^1, \alpha^2, ...). \tag{7}$$

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

B.3 Constructing Dynamic Player-Specific Punishments

The definition below is adapted from Chen & Takahashi (2012). Recall that given $\boldsymbol{\alpha} \in (\Delta(A))^\infty$ and t, ${}_t\boldsymbol{\alpha} = (\alpha^t, \alpha^{t+1}, ...)$.

Definition B3. *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) $v_{i\lambda}(\mathbf{r}^i) < v_{i\lambda}({}_t\mathbf{r}^j) - \gamma$.*

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

[36] For another application of these properties, see Bommier *et al.* (2019).

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

We begin by defining paths $\{\mathbf{r}^i_\lambda\}_{i \in I}$ indexed by two 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 B4, 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 $v_i(l^i) \leq v_{i\lambda}(\sigma) \leq v_i(h^i)$ for all $\sigma \in \Sigma$. Let $\mathcal{L}^i_\lambda$ be the set of all $l^j \in A, j \in I$, such that $v_i(l^j) < v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda)$, and let $N^i := |\mathcal{L}^i_\lambda|$. Note that $l^i \in \mathcal{L}^i_\lambda$. Enumerate the elements of $\mathcal{L}^i_\lambda$ according to i's preferences:

$$v_i(l^{i_0}) \leq v_i(l^{i_1}) \leq ... \leq v_i(l^{i_{N^i-1}}). \tag{8}$$

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 $(L^i_\lambda, \boldsymbol{\alpha}^i_\lambda)$. We note that the paths $\mathbf{r}^i_\lambda, i \in I$, do not reference 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 B3, 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. This and conditions (ii) and (iii) of Definition B3 will be established by choosing T_1 and T_2 appropriately. First, recall the following property of the exponential.

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

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

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

Proof. By Lemma B8,

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

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

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

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

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

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

Proof. For every η, define

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

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

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

By Lemma B8,

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

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

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

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

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

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

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

$$(1-[\underline{\beta}_{i\lambda}]^{n(T_1+T_2)})v_i(l^i)+[\underline{\beta}_{i\lambda}]^{n(T_1+T_2)}\varepsilon>\gamma'\quad\forall i\in I,\forall\lambda\in(\underline{\lambda}',1).\tag{9}$$

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

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

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

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

The last inequality follows from (9) and $\lambda\in(\underline{\lambda}',1)$. □

Lemma B12. *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 $v_i(h^i)\ge v_i(\kappa^{i_m})$ for all $m=0,...,n-1$, we obtain

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

By Lemma B6, $v_{i\lambda}(\mathbf{r}^i_\lambda)$ 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

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

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

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

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

Lemma B13. *For all i 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 (7) and (8) that

$$v_{i\lambda}(\mathbf{r}^i_\lambda) \leq v_{i\lambda}({}_1\mathbf{r}^i_\lambda) \leq ... \leq v_{i\lambda}({}_{N^iT_1-1}\mathbf{r}^i_\lambda) \leq v_{i\lambda}({}_{N^iT_1}\mathbf{r}^i_\lambda) = v_{i\lambda}(\boldsymbol{\alpha}^i_\lambda). \tag{10}$$

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

$$v_i(\kappa^{i_m}) < v_{i\lambda}({}_{(m+1)T_2}\boldsymbol{\alpha}^i_\lambda) \quad \forall m = 0, ..., n-1. \tag{11}$$

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

$$v_{i\lambda}({}_{mT_2}\boldsymbol{\alpha}^i_\lambda) = v_i(\kappa^{i_m})(1 - [\beta_{i\lambda}(\kappa^{i_m})]^{T_2}) + [\beta_{i\lambda}(\kappa^{i_m})]^{T_2}v_{i\lambda}({}_{(m+1)T_2}\boldsymbol{\alpha}^i_\lambda). \tag{12}$$

It follows from (7) and (11) 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 (10), 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 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 (7) and (10) 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 $v_i(\kappa^{i_k}) \geq v_{i\lambda}({}_{(k+1)T_2}\boldsymbol{\alpha}^i_\lambda)$, (12) and (7) 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 $t = kT_2+1, ..., nT_2-1$,

$$v_{i\lambda}({}_{kT_2}\boldsymbol{\alpha}^i_\lambda) \geq v_{i\lambda}({}_t\boldsymbol{\alpha}^i_\lambda) \geq v_{i\lambda}({}_{nT_2}\boldsymbol{\alpha}^i_\lambda) = v_{i\lambda}(\mathbf{w}^i_\lambda). \tag{13}$$

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 (13), this yields

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

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

Lemma B14. *There exist $\gamma''' > 0$ and $\underline{\lambda}'''$ 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 B8,

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

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

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

Lemma B15. *For all $i, j \in I, i \neq j, \lambda > \underline{\lambda}'''$, and $t \leq N^jT_1$, $v_{i\lambda}({}_t\mathbf{r}^j_\lambda) - v_{i\lambda}({}_tL^j_\lambda, \boldsymbol{\alpha}^i_\lambda) > \gamma'''$.*

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

$$\begin{aligned} & v_{i\lambda}\big({}_t\mathbf{r}^j_\lambda\big) - v_{i\lambda}\big({}_tL^j_\lambda, \boldsymbol{\alpha}^i_\lambda\big) \geq v_{i\lambda}\big(L^j_\lambda, \boldsymbol{\alpha}^j_\lambda\big) - v_{i\lambda}\big(L^j_\lambda, \boldsymbol{\alpha}^i_\lambda\big) = \\ = & \prod_{m=0}^{N^j-1} [\beta_{i\lambda}(l^{j_m})]^{T_1}\big(v_{i\lambda}\big(\boldsymbol{\alpha}^j_\lambda\big) - v_{i\lambda}\big(\boldsymbol{\alpha}^i_\lambda\big)\big) \geq [\underline{\beta}_{i\lambda}]^{nT_1}\big(v_{i\lambda}\big(\boldsymbol{\alpha}^j_\lambda\big) - v_{i\lambda}\big(\boldsymbol{\alpha}^i_\lambda\big)\big). \end{aligned} \tag{14}$$

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 B5, $v_{i\lambda}\big(K^{j\setminus j}, \mathbf{w}^i_\lambda\big) \geq v_{i\lambda}(K^{i\setminus j}, \mathbf{w}^i_\lambda)$ 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 B5, we have $v_{i\lambda}(\tilde{K}, \mathbf{w}^i_\lambda) \geq v_{i\lambda}(K^i, \mathbf{w}^i_\lambda) = v_{i\lambda}\big(\boldsymbol{\alpha}^i_\lambda\big)$. We conclude that

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

where the equality follows by a direct calculation and the last inequality by Lemma B14. Together with (14), the last chain of inequalities completes the proof. □

Lemma B16. *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 B5, $v_{i\lambda}({}_tL^j_\lambda, \boldsymbol{\alpha}^i_\lambda) \geq v_{i\lambda}(\pi^\uparrow_i({}_tL^j_\lambda), \boldsymbol{\alpha}^i_\lambda)$. Hence, by Lemma B15, $v_{i\lambda}\big({}_t\mathbf{r}^j_\lambda\big) - v_{i\lambda}(\pi^\uparrow_i({}_tL^j_\lambda), \boldsymbol{\alpha}^i_\lambda) > \gamma'''$. 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

$\pi_i^{\uparrow}({}_tL_\lambda^j)$ as (L', L'') where $L' \subset L_\lambda^i$ and $L'' \subset L_\lambda^j \setminus L_\lambda^i$. We claim that

$$v_{i\lambda}(L', L'', \boldsymbol{\alpha}_\lambda^i) \geq v_{i\lambda}(L', \boldsymbol{\alpha}_\lambda^i) \geq v_{i\lambda}(L_\lambda^i, \boldsymbol{\alpha}_\lambda^i) =: v_{i\lambda}(\mathbf{r}_\lambda^i) \tag{15}$$

By the stationarity of UzE preferences, or if $L' = \emptyset$, the first inequality is equivalent to $v_{i\lambda}(L'', \boldsymbol{\alpha}_\lambda^i) \geq v_{i\lambda}(\boldsymbol{\alpha}_\lambda^i)$, which follows since $v_i(l'') \geq v_{i\lambda}(\boldsymbol{\alpha}_\lambda^i)$ for all $l'' \in L''$. The second inequality in (15) follows from Lemma B6. $\square$

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

Proof. The desired inequality is equivalent to $v_{i\lambda}({}_t\boldsymbol{\alpha}_\lambda^j) \geq v_{i\lambda}(\mathbf{r}_\lambda^i)$ for all $t > 0$. If $t \geq nT_2$, then ${}_t\boldsymbol{\alpha}_\lambda^j = {}_\tau\mathbf{w}_\lambda^j \in SIR_\lambda^\varepsilon$ where $\tau = t - nT_2$. Hence, $v_{i\lambda}({}_t\boldsymbol{\alpha}_\lambda^j) \geq v_{i\lambda}(\mathbf{w}_\lambda^i)$. By Lemma B12, $v_{i\lambda}(\mathbf{w}_\lambda^i) - \gamma'' > v_{i\lambda}(\mathbf{r}_\lambda^i)$ and we are done. Suppose now that $t < nT_2$ and write ${}_t\boldsymbol{\alpha}_\lambda^j$ as $({}_tK^j, \mathbf{w}_\lambda^j)$. Lemmas B6 and B5 imply that

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

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

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

Lengthy but straightforward calculations show that

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

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

Take $\gamma := \min\{\gamma', \gamma'', \gamma'''\}$ and $\underline{\lambda} := \max\{\underline{\lambda}', \underline{\lambda}'', \underline{\lambda}'''\}$, where $\gamma', \gamma'', \gamma'''$ and $\underline{\lambda}', \underline{\lambda}'', \underline{\lambda}'''$ are defined as in Lemmas B11, B12, and B16. Then, Lemmas B12, B13, B16, and B17 show that for all $\lambda > \underline{\lambda}$ and $\boldsymbol{\alpha} \in SIR_\lambda^\varepsilon$, the paths $\{\mathbf{r}_\lambda^i\}_i$ meet the conditions in Definition B3.

B.4 Equilibrium Strategies

Let $m^i := (m_1^i, ..., m_n^i) \in \Sigma$ be a strategy profile in which player i best-responds to a min-max strategy by the opponents. By Lemma 1, we can choose m^i to be a profile of stationary strategies and, hence, identify m^i with an element of $\Delta(A)$. Utilities are normalized so that $g_i(m^i) = 0$ for every $i \in I$. Take $\varepsilon > 0$. By Lemma B7, there exist $\gamma > 0$ and $\underline{\lambda}' \geq 0$ such that for every $\lambda > \underline{\lambda}'$, every $\boldsymbol{\alpha} \in SIR_\lambda^\varepsilon$ allows DPSP with wedge γ. Let $\overline{g}_i := \max_a g_i(a)$

and choose an integer χ_i such that $\chi_i > \frac{\overline{g}_i}{\gamma\left(1-\beta_i(m^i)\right)}$. Since

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

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

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

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

$$(1-\lambda)\overline{g}_i + \left(\overline{m} - [\beta_{i\lambda}(m^j)]^{\chi}(\overline{m}+\gamma)\right) - v_i(m^j)(1-[\beta_{i\lambda}(m^j)]^{\chi}) < 0. \tag{16}$$

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 (16) is satisfied for all $\lambda > \underline{\lambda}_i'''$. Since there are finitely many players and finitely many integers between between 1 and χ_j, the threshold $\underline{\lambda}_i'''$ can be chosen independently of $j \neq i$ and χ.

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 $\boldsymbol{\alpha} \in SIR_\lambda^\varepsilon$. Let $\{\mathbf{r}_\lambda^i\}_{i\in I}$ 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) follow $\boldsymbol{\alpha}$ as long as no player deviates. If player j deviates from (A), then (B) play m_i^j for χ_j periods, and then (C) play $\mathbf{r}_\lambda^j$ thereafter. If player k deviates in phase (B) or (C), begin phase (B) again with $j = k$. Given the choice of $\underline{\lambda}$, it is easy to show no player has an incentive to deviate; we omit the details.

C Proof of Theorem 2

For every $\eta \in \mathbb{R}_+^n$, define the η-**face** of a convex set $X \subset \mathbb{R}^n$ to be the set $F(\eta) = \{v \in X : \eta \cdot v \geq \eta \cdot v' \ \forall v' \in X\}$. In what follows we focus on the faces of V^{one}.

Lemma C18. *For some $\eta \in \mathbb{R}_{++}^n$, the set $F(\eta)$ is not a singleton.*

Proof. Suppose not. Then, since V^{one} is a polytope, $E := \{F(\eta) : \eta \in \mathbb{R}_{++}^n\}$ is a finite set of extreme points. By CI, E is not a singleton. For every $v \in E$, let $H_{++}(v) = \{\eta \in \mathbb{R}_{++}^n : F(\eta) = \{v\}\}$. By construction, each set $H_{++}(v)$ is closed in $\mathbb{R}_{++}^n$ and $H_{++}(v) \cap$

$H_{++}(v') = \emptyset$ for all distinct $v, v' \in E$. But then $\{H_{++}(v) : v \in E\}$ is a finite partition of $\mathbb{R}^n_{++}$ into relatively closed subsets, which is impossible since $\mathbb{R}^n_{++}$ is connected. □

Pick $\eta \in \mathbb{R}^n_{++}$ and $\alpha \in \Delta(A)$ such that $v(\alpha^{one})$ is in the relative interior of $F(\eta)$. Then, $v_i(a) \neq v_i(a')$ for some i and a, a' in the support of α. By Lemma 3, $v(\alpha) \geq v(\alpha^{one})$ and $v_i(\alpha) > v_i(\alpha^{one})$. Since $\eta \gg 0$, $v(\alpha) >^* V^{one}$. 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 not, then, by CI, there is $i \in I$ 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 $\overline{v}$. As before, 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, $v(\alpha)$ can be sustained for all λ sufficiently high.

D Proof of Theorem 3

Lemma D19. *Let $\alpha \in \Delta(A)$ and $i, k \in I$ be such that $v_i(\alpha) > v_k(\alpha)$ and $\beta_i(\alpha) < \beta_k(\alpha)$. Then, for every $\eta \in \mathbb{R}^n_+$ such that $\eta_k > 0$ and every λ, there is $\boldsymbol{\alpha}_\lambda \in (\Delta(A))^\infty$ such that $\eta \cdot v_\lambda(\boldsymbol{\alpha}_\lambda) > \eta \cdot v(\alpha)$ and $\eta \cdot \lim_{\lambda \nearrow 1} v_\lambda(\boldsymbol{\alpha}_\lambda) > \eta \cdot v(\alpha)$.*

Proof. Fix η such that $\eta_k > 0$. By symmetry, there is $\alpha_k \in \Delta(A)$ such that $v_i(\alpha_k) = v_k(\alpha), v_i(\alpha) = v_k(\alpha_k)$, $\beta_i(\alpha_k) = \beta_k(\alpha), \beta_i(\alpha) = \beta_k(\alpha_k)$, and for all $j \neq i, k$, $v_j(\alpha) = v_j(\alpha_k)$ and $\beta_j(\alpha) = \beta_j(\alpha_k)$. Take $\theta > \max\{0, \frac{\ln(\eta_i/\eta_k)}{\beta_k(\alpha)-\beta_i(\alpha)}\}$ and $T_\lambda = \lceil \frac{\theta}{1-\lambda} \rceil$. Let $\boldsymbol{\alpha}_\lambda = (\alpha^0, \alpha^1, ...)$ be such that $\alpha^t = \alpha$ for all $t \leq T_\lambda$ and $\alpha^t = \alpha_k$ for all $t > T_\lambda$. By construction, $v_{j\lambda}(\boldsymbol{\alpha}_\lambda) = v_j(\alpha)$ for all $j \neq i, k$. Then,

$$\eta \cdot v_\lambda(\boldsymbol{\alpha}_\lambda) - \eta \cdot v(\alpha) = (\eta_k[\beta_{k\lambda}(\alpha)]^{T_\lambda} - \eta_i[\beta_{i\lambda}(\alpha)]^{T_\lambda})(v_i(\alpha) - v_k(\alpha)). \tag{17}$$

Since $v_i(\alpha) > v_k(\alpha)$, we first want to show that $\eta_k[\beta_{k\lambda}(\alpha)]^{T_\lambda} - \eta_i[\beta_{i\lambda}(\alpha)]^{T_\lambda} > 0$. If $\eta_i \leq \eta_k$, the inequality is trivially true since $\beta_k(\alpha) > \beta_i(\alpha)$. Suppose $\eta_i > \eta_k$. Since $\theta > \frac{\ln(\eta_i/\eta_k)}{\beta_k(\alpha)-\beta_i(\alpha)}$, a sufficient condition for the desired inequality is that

$$f(\lambda) := \ln \frac{\lambda + (1-\lambda)\beta_k(\alpha)}{\lambda + (1-\lambda)\beta_i(\alpha)} - (1-\lambda)(\beta_k(\alpha) - \beta_i(\alpha)) > 0..$$

But the latter follows since $f(1) = 0$ and since f is strictly decreasing in λ whenever

$\beta_k(\alpha) > \beta_i(\alpha)$. Finally, using (17) and Lemma B8, note that

$$\eta \cdot \lim_{\lambda \to 1} v_\lambda(\boldsymbol{\alpha}_\lambda) - \eta \cdot v(\alpha) = (\eta_k e^{-(1-\beta_k(\alpha))\theta} - \eta_i e^{-(1-\beta_i(\alpha))\theta})(v_i(\alpha) - v_k(\alpha)) > 0,$$

where inequality follows from the choice of θ and $v_i(\alpha) > v_k(\alpha)$. □

The proof of the next lemma is straightforward and omitted.

Lemma D20. *Consider a symmetric game and take $a \in A$ and $i, j \in I$. Under IMI, $v_i(a) \geq v_j(a)$ if and only if $\beta_i(a) \leq \beta_j(a)$. Under DMI, $v_i(a) \geq v_j(a)$ if and only if $\beta_i(a) \geq \beta_j(a)$.*

Next, fix some i and 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 $k \neq i$. For each $\eta \in \mathbb{R}^n$, let $F(\eta)$ be the corresponding face of $conv(V^{iid})$ and let $e^i \in \mathbb{R}^n$ be the vector whose i^{th}-coordinate is 1 and all other coordinates are 0.

Lemma D21. *If $\alpha \in \Delta(A)$ is such that $v(\alpha) \in F(e^i)$, then $\alpha \in S_i$.*

Proof. By Lemma B2, if $v(\alpha) \in F(e^i)$, then every $a \in \operatorname{supp} \alpha$ is such that $v(a) \in F(e^i)$. By IMI, $\beta_i(a) = \min\{\beta_i(a') : a' \in A\}$ 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 D20, $\beta_i(a) < \beta_k(a)$. It follows that $v_i(\alpha) > v_k(\alpha)$ and $\beta_i(\alpha) < \beta_k(\alpha)$ for some k. □

Take a sequence $(\eta^m)_m$ such that $\eta^m \in \mathbb{R}^n_{++}$ for all m and $\eta^m \to_m e^i$. For each m, pick α^m such that $v(\alpha^m)$ belongs to the face $F(\eta^m)$ of $conv(V^{iid})$. Passing onto a subsequence if necessary, assume $\alpha^m \to_m \alpha^*$. By the theorem of the maximum, $v(\alpha^*) \in F(e^i)$ and, by Lemma D21, $\alpha^* \in S_i$. Since S_i is open, we can pick m large enough such that $\alpha^m \in S_i$. Since also $\eta^m_k > 0$, we can invoke Lemma D19 to deduce that (i) for any given λ, $v(\alpha^m)$ is not on the corresponding η^m-face of V_λ and (ii) the Pareto frontier of V_λ does not collapse to that of $conv(V^{iid})$ as $\lambda \nearrow 1$.

E 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 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_λ^t in place of $\eta_\lambda^t(\mathbf{a})$. 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 and we omit the proofs.

Lemma E22. *If $\mathbf{a} = (a^0, a^1, ...) \in P_\lambda(\eta)$, then ${}_t\mathbf{a} \in P_\lambda(\eta_\lambda^t(\mathbf{a}))$ for all $t > 0$. Also, if $\hat{\mathbf{a}} \in P_\lambda(\eta_\lambda^t(\mathbf{a}))$ 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 $a^{sym} \in A^E$ such that $v_1(a^{sym}) = \max_{a \in A^E} v_1(a)$. The next two lemmas assume either IMI or DMI.

Lemma E23. *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 $\eta_\lambda^1 = \left(\eta_1\beta_{1\lambda}(a^0), \eta_2\beta_{2\lambda}(a^0)\right)$ determine the same direction and, by Lemma E22, ${}_1\mathbf{a} \in P_\lambda(\eta)$. 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

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

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

Lemma E24. *For every $\mathbf{a} \in P_\lambda(\eta)$, if $a^t \in A^E$ for some t, then $a^t = a^{sym}$.*

Proof. Obvious given Lemma E23. □

F Proof of Theorem 4

For the sake of simplicity, assume that for each i the action $a^{max,i} \in A$ such that $v_i(a^{max,i}) = v_i^{max}$ is unique. Let $\mathbf{a}^{max,i} = (a^{max,i}, a^{max,i}, ...)$. If $a^{max,1} = a^{max,2}$, then the unique efficient path is $\mathbf{a}^{max,1}$ and the proof is complete. From now on, assume $a^{max,1} \neq a^{max,2}$. By the symmetry of the game, $v_i(a^{max,i}) > v_j(a^{max,i})$ and, by Lemma D20, $\beta_i(a^{max,i}) > \beta_j(a^{max,i})$. Fix $\lambda, \eta \in \mathbb{R}_+^2 \setminus \{0\}$, and $\mathbf{a} \in P_\lambda(\eta)$. If $\eta_i = 0$, then $\mathbf{a} = \mathbf{a}^{max,j}$. Thus, assume $\eta \gg 0$.

Lemma F25. *If $\beta_{1\lambda}(a^0) > \beta_{2\lambda}(a^0)$, then $v_{1\lambda}(\mathbf{a}) > v_{2\lambda}(\mathbf{a})$.*

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

$$v_{2\lambda}({}_1\mathbf{a}) \leq v_{2\lambda}(\mathbf{a}) \quad \text{and} \quad v_{1\lambda}({}_1\mathbf{a}) \geq v_{1\lambda}(\mathbf{a}). \tag{18}$$

From (7), 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 (18) 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 D20, $\beta_{2\lambda}(a^0) < \beta_{1\lambda}(a^0)$ implies $v_2(a^0) < v_1(a^0)$. Hence, $v_{2\lambda}(\mathbf{a}) < v_{1\lambda}(\mathbf{a})$. □

Lemma F26. *If $a^0 = a^{max,2}$, then $\mathbf{a} = \mathbf{a}^{max,2}$.*

Proof. Suppose $a^0 = a^{max,2}$. Then we have $\beta_{2\lambda}(a^0) > \beta_{1\lambda}(a^0)$. It follows from the proof of Lemma F25 that $v_{2\lambda}(\mathbf{a}) \geq v_2(a^0) = v_2(a^{max,2})$. Thus, $\mathbf{a} = \mathbf{a}^{max,2}$. □

Lemma F27. *If $v_{1\lambda}(\mathbf{a}) = v_{2\lambda}(\mathbf{a})$, then $\mathbf{a} = (a^{sym}, a^{sym}, ...)$.*

Proof. By Lemma F25, $\beta_{1\lambda}(a^0) = \beta_{2\lambda}(a^0)$ and, hence, $a^0 \in A^E$ by Lemma D20. 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 E24, $\mathbf{a} = (a^{sym}, a^{sym}, ...)$. □

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

Lemma F28. *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 F29. *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 F25 implies that $v_{1\lambda}(\mathbf{a}) < v_{2\lambda}(\mathbf{a})$. Conclude that

$$v_{1\lambda}(\hat{\mathbf{a}}) < v_{2\lambda}(\hat{\mathbf{a}}) \quad \forall \hat{\mathbf{a}} \in P_\lambda(\eta^T_\lambda(\mathbf{a})). \tag{19}$$

By Lemma E22, ${}_T\mathbf{a} \in P_\lambda(\eta^T_\lambda(\mathbf{a}))$ and, hence, $v_{1\lambda}({}_T\mathbf{a}) < v_{2\lambda}({}_T\mathbf{a})$. By Lemma F25, $\beta_{1\lambda}(a^T) \leq \beta_{2\lambda}(a^T)$. By the choice of T, it must be that $\beta_{1\lambda}(a^T) = \beta_{2\lambda}(a^T)$. By Lemma D20, $v_1(a^T) = v_2(a^T)$ so that $a^T \in A^E$. It follows from Lemmas E23 and E24 that $\mathbf{a}' := (a^{sym}, a^{sym}, ...) \in P_\lambda(\eta^T_\lambda(\mathbf{a}))$. But then, $v_{1\lambda}(\mathbf{a}') = v_{2\lambda}(\mathbf{a}')$, contradicting (19). □

If $v_{1\lambda}(\mathbf{a}) = v_{2\lambda}(\mathbf{a})$, then, by Lemma F27, $\mathbf{a} = (a^{sym}, a^{sym}, ...)$ and we are done. Assume $v_{1\lambda}(\mathbf{a}) < v_{2\lambda}(\mathbf{a})$. By Lemma F25, $\beta_{1\lambda}(a^0) \leq \beta_{2\lambda}(a^0)$. We claim that there is T such that

$\beta_{1\lambda}(a^T) < \beta_{2\lambda}(a^T)$. If $\beta_{1\lambda}(a^0) < \beta_{2\lambda}(a^0)$, we are done. Assume $\beta_{1\lambda}(a^0) = \beta_{2\lambda}(a^0)$ and let $T \geq 1$ be the first period t such that $\beta_{1\lambda}(a^t) \neq \beta_{2\lambda}(a^t)$. Since $v_{1\lambda}(\mathbf{a}) < v_{2\lambda}(\mathbf{a})$, such T exists by Lemma D20. Since $\beta_{1\lambda}(a^t) = \beta_{2\lambda}(a^t)$ for every $t < T$, Lemma E24 implies that $a^t = a^{sym}$ for all such t. Since $a^0 = a^{sym}$, Lemma F28 implies that $v_{1\lambda}({}_1\mathbf{a}) < v_{1\lambda}(\mathbf{a})$ and $v_{2\lambda}(\mathbf{a}) < v_{2\lambda}({}_1\mathbf{a})$. 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 F28 repeatedly, conclude that $v_{1\lambda}({}_t\mathbf{a}) < v_{2\lambda}({}_t\mathbf{a})$ for every $t \leq T$. By Lemma F25, $\beta_{1\lambda}(a^T) \leq \beta_{2\lambda}(a^T)$ and, by the choice of T, $\beta_{1\lambda}(a^T) < \beta_{2\lambda}(a^T)$. By Lemma F29, it now follows that $\beta_{1\lambda}(a^t) < \beta_{2\lambda}(a^t)$ for all $t > T$. Let $l := \min \frac{\beta_{2\lambda}(a)}{\beta_{1\lambda}(a)}$, with the minimum taken over all $a \in A$ such that $\beta_{1\lambda}(a) < \beta_{2\lambda}(a)$. By construction, $l > 1$ and

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

Conclude that, as $t \nearrow \infty$, $\frac{\eta^t_{2\lambda}(\mathbf{a})}{\eta^t_{1\lambda}(\mathbf{a})} \nearrow \infty$ and, hence, $v_{2\lambda}({}_t\mathbf{a}) \nearrow v_2^{max}$. If $a^t \neq a^{max,2}$ for all $t > T$, then $v_{2\lambda}({}_t\mathbf{a}) \leq \max_{\tau > T} v_2(a^\tau) < v_2^{max}$ for all $t > T$, a contradiction. Thus, $a^t = a^{max,2}$ for some $t > T$ and, by Lemma F26, ${}_t\mathbf{a} = \mathbf{a}^{max,2}$, completing the proof.

G Proof of Theorem 5

As in the preceding proof, assume for the sake of simplicity that the actions $a^{max,i}$ and a^{sym} are uniquely defined. Given $\alpha \in \Delta A$ and $\sigma \in \Sigma$, write (α, σ) for a strategy such that α is played in $t = 0$ and σ is played after every history h^1. Say that $\sigma \in \Sigma$ is η-efficient if $\eta \cdot v(\sigma) \geq \eta \cdot v$ for all $v \in V_\lambda$. Fix λ and define $H^{sym} = \{\eta \in \mathbb{R}^2_+ : (a^{sym}, a^{sym}, ...) \in P_\lambda(\eta)\}$. If there is no symmetric action a^{sym} such that $(a^{sym}, a^{sym}, ...)$ is efficient, then $H^{sym} = \varnothing$.

Lemma G30. *There exists $T > 0$ such that $\min\{v_{1\lambda}({}_t\mathbf{a}), v_{2\lambda}({}_t\mathbf{a})\} < 0$ for all $t > T, \eta \notin H^{sym}$ and $\mathbf{a} \in P_\lambda(\eta)$.*

Proof. Since $v_1(a^{max,2}) < 0$, there exists a payoff $(\hat{v}_1, \hat{v}_2)$ on the Pareto frontier of V_λ such that $v_1(a^{max,2}) < \hat{v}_1 < 0$, $\hat{v}_2 < v_2(a^{max,2})$, and $(\hat{v}_1, \hat{v}_2)$ is efficient for some strictly positive direction $\hat{\eta}$. Take T such that $l^T \geq \frac{\hat{\eta}_2}{\hat{\eta}_1}$, where l is as defined toward the end of the proof of Theorem 4. For each $\eta \notin H^{sym}$ and $\mathbf{a} \in P_\lambda(\eta)$, Lemmas E23 and E24 imply that $\beta_1(a^0) \neq \beta_2(a^0)$. Suppose $\beta_1(a^0) < \beta_2(a^0)$, the other case being handled analogously. Then, by Lemma F29, $\beta_1(a^t) < \beta_2(a^t)$ for all $t > 0$ and, hence, $v_{1\lambda}(\mathbf{a}) < v_{2\lambda}(\mathbf{a})$. It follows that $\eta_2 \geq \eta_1$ and, given the definition of l, that $\eta^t_\lambda(\mathbf{a}) \geq l^t$. It follows from the choice of T that $\eta^t_\lambda(\mathbf{a}) \geq \frac{\hat{\eta}_2}{\hat{\eta}_1}$ for all $t > T$. By Lemma E22, ${}_t\mathbf{a} \in P_\lambda(\eta^t_\lambda(\mathbf{a}))$ and, hence, $v_{1\lambda}({}_t\mathbf{a}) \leq \hat{v}_1 < 0$ for all $t > T$. □

Now, take an efficient strategy σ such that $v_{1\lambda}(\sigma) < v_{2\lambda}(\sigma)$ and let η be such that σ is η-efficient. Since $v_{1\lambda}(\sigma) < v_{2\lambda}(\sigma)$, we have $\eta_2 \geq \eta_1$. If $\eta_1 = 0$, the result is obvious, so suppose $\eta_1 > 0$ and consider the case when $\eta \notin H^{sym}$. Thinking of σ as a (Borel) distribution on A^∞, define its support $S_\sigma \subset A^\infty$ as in Aliprantis & Border (1999, p.374). Then, every open set $O \subset A^\infty$ that intersects S_σ has strictly positive probability under σ. Since σ is η-efficient, it follows that $S_\sigma \subset P_\lambda(\eta)$. Then, by Lemma G30, there is $T > 0$ such that $v_{1\lambda}(\sigma \,|\, h^t) < 0$ for all h^t (σ-a.s.) and $t > T$, as desired. It remains to consider the case when $\eta \in H^{sym}$. If σ induces a pure path $\mathbf{a} \in A^\infty$, the desired conclusion follows from Theorem 4. If σ involves randomization, there is an on-path history h^t such that with strictly positive probability (in terms of the public signal ω_0^t) some $a \in A$ is played such that $\beta_1(a) \neq \beta_2(a)$. Otherwise, $v_{1\lambda}(\sigma) = v_{2\lambda}(\sigma)$. Fix such h^t and $a \in A$, and focus on the states ω_0^t in which a is played.[37] Any such state gives rise to an on-path history h^{t+1}. Let $\sigma(\omega_0^t) \in \Sigma$ be the strategy representing the restriction of σ to the corresponding subgame. Let $(a, \sigma(\omega_0^t)) \in \Sigma$ be the strategy in which a is played first, followed by $\sigma(\omega_0^t)$ independently of history. The efficiency of σ implies that the strategies $(a, \sigma(\omega_0^t))$ are efficient (σ-a.s). Pick ω_0^t such that $(a, \sigma(\omega_0^t))$ is efficient. Then, $\sigma(\omega_0^t)$ is efficient for some $\hat{\eta}$. If $\hat{\eta} \in H^{sym}$, then, by an obvious extension of Lemma E22, $(a, a^{sym}, a^{sym}, ...) \in P_\lambda(\hat{\eta})$. But this contradicts Lemma F29 since $\beta_1(a) \neq \beta_2(a)$. Thus, $\hat{\eta} \notin H^{sym}$ and the proof reduces to a previous case.

H Proof of Theorem 6

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 orthogonal to some $\eta \gg 0$. Since $v_2(a^r) < v_2^{max}$, we have $\eta_1 > 0$. 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 $\eta \gg 0$, $v_\lambda(a^r_{-T}, a^*_T) >^* V^{one}$ 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 D20, $v_2(a^r) > v_1(a^r)$ 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.

[37]From Section 2, ft.12, recall that, without loss of generality, the private signals are not used on path.

I Proof of Proposition 5

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

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

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, i.e., $g_i(A) = C$. Take some $c < c'$ in C. By the first part of the proof, rs^{-1} cannot be convex on $[c, c'] \subset C$ for otherwise correlation aversion will be contradicted. Conclude that $(rs^{-1})'' \leq 0$ on $[c, c']$ and for some $c'' \in (c, c')$, $(rs^{-1})'' < 0$. Since r and s are twice continuously differentiable, so is rs^{-1}. It follows that the set of points at which $(rs^{-1})'' < 0$ is open. Being continuously differentiable, $(rs^{-1})'$ is also absolutely continuous. Applying the fundamental theorem of calculus, deduce that $(rs^{-1})'(c') - (rs^{-1})'(c) = \int_c^{c'} (rs^{-1})'' < 0$. Thus, $(rs^{-1})'$ is strictly decreasing, which shows that rs^{-1} is strictly concave.

J Proof of Theorem 7

For each $\mu \in \Delta(D)$, which can be visualized as an infinite probability tree, and each t, let $\mu^t \in \Delta(A)$ be the marginal induced over pure actions in that period and let $\mu^t_+ \in \Delta(D)$ be the marginal over the continuation trees. Since rs^{-1} is strictly concave:[38]

$$\begin{aligned} v_i(\mu) =& \mathbb{E}_\mu rs^{-1}\Big((1-\beta)[s \circ g_i] + \beta[s \circ r^{-1} \circ v_i]\Big) \\ \leq& rs^{-1}\Big((1-\beta)\mathbb{E}_{\mu^0}[s \circ g_i] + \beta \mathbb{E}_{\mu^0_+}[s \circ r^{-1} \circ v_i]\Big) \quad \forall i. \end{aligned}$$

Moreover, for each i, the inequality is strict unless $(a', \mu') \mapsto v_i(a', \mu')$ is constant μ-a.s. Applying the same argument to each μ' in the support of μ^0_+, deduce that

$$v_i(\mu) \leq rs^{-1}\Big((1-\beta)\big(\mathbb{E}_{\mu^0}[s \circ g_i] + \beta \mathbb{E}_{\mu^1}[s \circ g_i]\big) + \beta^2 \mathbb{E}_{\mu^1_+}[s \circ r^{-1} \circ v_i]\Big) \quad \forall i.$$

[38] As convenient, the composition of two functions f and g will be denoted either as $f \circ g$ or fg.

Iterating the argument and noting that $\beta^t \to 0$, deduce that

$$v_i(\mu) \leq rs^{-1}\big(\mathbb{E}_\alpha[s \circ g_i]\big) \quad \forall i,$$

where $\alpha = (1-\beta)\sum_t \beta^t \mu^t \in \Delta(A)$. Since $\beta > 1 - |A|^{-1}$, it follows from Lemma 1 in Fudenberg & Maskin (1991) that there is $\mathbf{a} \in A^\infty$ such that $v_i(\mathbf{a}) = rs^{-1}\big(\mathbb{E}_\alpha[s \circ g_i]\big)$ for all i, which completes the proof of the theorem.

K Proof of Theorem 8

Adopt the same notation as in the proof of Theorem 7. Since rs^{-1} is convex,

$$v_i(\mu) = \mathbb{E}_\mu rs^{-1}\Big((1-\beta)[s \circ g_i] + \beta[s \circ r^{-1} \circ v_i]\Big) \leq (1-\beta)\mathbb{E}_{\mu^0}[r \circ g_i] + \beta \mathbb{E}_{\mu^0_+} v_i \quad \forall i.$$

Iterating the argument as in the proof of Theorem 7, deduce that

$$v_i(\mu) \leq (1-\beta)\sum_{t=0}^{\infty} \beta^t \mathbb{E}_{\mu^t}[r \circ g_i] = v_i(\alpha^{one}) \quad \forall i,$$

where $\alpha = (1-\beta)\sum_t \beta^t \mu^t$. This completes the proof.

References

Abreu, Dilip, Dutta, Prajit K, & Smith, Lones. 1994. The folk theorem for repeated games: a NEU condition. *Econometrica*, **62**(4), 939–948.

Al-Najjar, Nabil I, & Shmaya, Eran. 2019. Recursive utility and parameter uncertainty. *Journal of Economic Theory*, **181**, 274–288.

Aliprantis, Charalambos, & Border, Kim. 1999. *Infinite Dimensional Analysis*. 2 edn. Springer-Verlag.

Andersen, Steffen, Harrison, Glenn W, Lau, Morten I, & Rutström, E Elisabet. 2018. Multiattribute Utility Theory, Intertemporal Utility, and Correlation Aversion. *International Economic Review*, **59**(2), 537–555.

Backus, David K, Routledge, Bryan R, & Zin, Stanley E. 2004. Exotic preferences for macroeconomists. *Pages 319–390 of: NBER Macroeconomics Annual*, vol. 19. MIT Press.

Becker, Gary, & Mulligan, Casey. 1997. The Endogenous Determination of Time Preference. *The Quarterly Journal of Economics*, **112**(3), 729–758.

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

Bommier, Antoine, Kochov, Asen, & Le Grand, François. 2019. Ambiguity and endogenous discounting. *Journal of Mathematical Economics*, **83**, 48–62.

Chen, Bo, & Takahashi, Satoru. 2012. A folk theorem for repeated games with unequal discounting. *Games and Economic Behavior*, **76**(2), 571–581.

Chew, Soo H., & Epstein, Larry G. 1991. Recursive Utility Under Uncertainty. *Pages 352–369 of:* Khan, Ali M., & Yannelis, Nicholas C. (eds), *Equilibrium Theory in Infinite Dimensional Spaces*. Studies in Economic Theory, vol. 1. Springer Berlin Heidelberg.

Chew, Soo Hong. 1983. A generalization of the quasilinear mean with applications to the measurement of income inequality and decision theory resolving the Allais paradox. *Econometrica*, **51**(4), 1065–1092.

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

Epstein, Larry. 1987. A simple dynamic general equilibrium model. *Journal of Economic theory*, **41**(1), 68–95.

Epstein, Larry, & Zin, Stanley. 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica*, **57**(4), 937–969.

Epstein, Larry G, Farhi, Emmanuel, & Strzalecki, Tomasz. 2014. How much would you pay to resolve long-run risk? *American Economic Review*, **104**(9), 2680–97.

Epstein, L.G., & Tanny, S.M. 1980. Increasing generalized correlation: a definition and some economic consequences. *Canadian Journal of Economics*, 16–34.

Fisher, Irving. 1930. *The theory of interest*. Macmillan New York.

Friedman, Milton. 1969. *The optimum quantity of money and other essays*. Aldine, Chicago.

Fudenberg, Drew, & Maskin, Eric. 1986. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, **54**(3), 533–554.

Fudenberg, Drew, & Maskin, Eric. 1991. On the dispensability of public randomization in discounted repeated games. *Journal of Economic Theory*, **53**(2), 428–438.

Ifcher, John, & Zarghamee, Homa. 2011. Happiness and time preference: The effect of positive affect in a random-assignment experiment. *American Economic Review*, **101**(7), 3109–29.

Kochov, Asen. 2015. Time and no lotteries: An axiomatization of maxmin expected utility. *Econometrica*, **83**(1), 239–262.

Kochov, Asen. 2017. A separation result for stationary preferences. *Journal of Mathematical Economics*, **70**, 123–126.

Koopmans, T.C. 1960. Stationary ordinal utility and impatience. *Econometrica*, **28**(2), 287–309.

Kreps, David, & Porteus, Evan. 1978. Temporal Resolution of Uncertainty and Dynamic Choice Theory. *Econometrica*, **46**(1), 185–200.

Lehrer, Ehud, & Pauzner, Ady. 1999. Repeated games with differential time preferences. *Econometrica*, **67**(2), 393–412.

Loewenstein, George. 1996. Out of control: Visceral influences on behavior. *Organizational behavior* *and human decision processes*, **65**(3), 272–292.

Lucas, Robert, & Stokey, Nancy. 1984. Optimal growth with many consumers. *Journal of Economic Theory*, **32**(1), 139–171.

Marinacci, Massimo, & Montrucchio, Luigi. 2010. Unique solutions for stochastic recursive utilities. *Journal of Economic Theory*, **145**(5), 1776–1804.

Miao, Bin, & Zhong, Songfa. 2015. Risk preferences are not time preferences: separating risk and time preference: comment. *American Economic Review*, **105**(7), 2272–86.

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

Ramsey, Frank. 1928. A mathematical theory of saving. *The Economic Journal*, **38**(152), 543–559.

Samuelson, Paul A. 1937. A note on measurement of utility. *The Review of Economic Studies*, **4**(2), 155–161.

Stanca, Lorenzo Maria. 2022. *Essays on Recursive Preferences*. Ph.D. thesis, Northwestern University.

Strzalecki, Tomasz. 2013. Temporal resolution of uncertainty and recursive models of ambiguity aversion. *Econometrica*, **81**(3), 1039–1074.

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

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