Repeated Games Without Public Randomization: A Constructive Approach*

Ani Dasgupta† Sambuddha Ghosh‡

This version: February 22, 2019

Abstract

We study discounted infinitely repeated games with perfect monitoring and without public randomization. Both symmetric and asymmetric discounting cases are considered; a new geometric construct called 'self-accessibility' is proposed and used to unify the analyses of these two cases. For symmetric discounting, our approach leads to easy computability of a discount factor bound needed to support a specific payoff vector in equilibrium. When discounting is allowed to be asymmetric, we show that any payoff vector that is in the interior of the smallest rectangular region containing the pure-action payoffs is realizable in the repeated game. Next, an easily-verifiable condition, 'strict diagonalizability', is offered as a sufficient and almost necessary condition for a payoff vector to be an equilibrium payoff for *some* discount factor vector. 'Turnpike strategies' that support a target payoff are explicitly constructed. Our results thus encompass and generalize Fudenberg and Maskin (1986, 1991).

Keywords: Repeated Games, Public Randomization, Asymmetric Discounting

*We thank Jimmy Chan, Michael Manove, Eric Maskin, Jawwad Noor, Wojciech Olszewski, Phil Reny, Takuo Sugaya and Satoru Takahashi for helpful comments. We are also grateful to Johannes Hörner, Hsueh-Ling Huynh, Bart Lipman, and Juan Ortner for numerous discussions and to Beixi Zhou for outstanding research assistance. We owe a special debt to Costas Cavounidis; more than one result in this paper owes its current form to his critical inputs. All errors are ours.

†International Maritime Business Department, Massachusetts Maritime Academy and Economics Department, Boston University.

‡Economics Group, Shanghai University of Finance and Economics.

1 Introduction

The pioneering work of Fudenberg and Maskin (1986), hereafter FM 1986, demonstrated that provided the players are patient enough, any FSIR (feasible and strictly individual rational) payoff vector for a game may be supported as the discounted average of payoffs arising from an SPNE (subgame perfect Nash equilibrium) of the corresponding repeated game. Two of the assumptions made in the original FM 1986 paper are of concern to us; the first is the availability of a public randomization device (hereafter PRD) and the second is the assumption that all players have exactly the same discount factor.

Dispensing with PRDs is not a minor addendum because it is problematic to argue that all players have access to a correlating device fine enough to realize all FSIR payoffs. While, following Aumann and Maschler (1995), communication can be used to construct 'jointly controlled lotteries' that serve as correlating devices, in the industrial organization context, this might run foul of antitrust regulations that forbid communication among firms. Similarly, requiring exact equality of discount factors seems too restrictive; if we allow economic agents to have differential preferences over consumption bundles, why should we require them to hold identical time preferences?[1] Lehrer and Pauzner (1999), hereafter LP 1999, have asserted that even when payoffs are monetary and players can borrow in an outside market to smooth their consumption streams, different agents, because of their differential financial standings, may be subject to different interest rates. Finally, it is odd to argue against the importance of asymmetric discounting in complete information games when the literature on incomplete information games is rife with such models.

A plethora of interesting and counterintuitive things happen as soon as one disallows PRDs, even when players are patient. For example, in the symmetric discounting case Yamamoto (2010) shows that for large enough common discount factors, the set of SPNE payoffs can be non-convex and non-monotonic with respect to the common discount factor. Salonen and Vartiainen (2008) show that for large enough but unequal discount factors the feasible payoff set of the repeated game can be totally disconnected,[2] and the Pareto frontier function can be everywhere discontinuous. For a class of two-player *finite* games, Olszewski (1998) shows that the undiscounted folk theorem does not hold.

To establish whether a given point in the payoff space is an equilibrium payoff vector when PRDs are unavailable, the first question that must be answered is: Is this payoff feasible in the repeated game using uncorrelated (but possibly mixed) actions? If this is answered in the affirmative, the second question is: Is this an equilibrium (SPNE) payoff? To fix terminology, we shall henceforth call these questions those of *realizability* and *supportability* respectively.

In the context of symmetric discounting, Fudenberg and Maskin (1991), hereafter FM

[1] See for example Harrington (1989), Obara and Zincenko (2017) and Haag and Lagunoff (2007) who provide interesting economic models with asymmetric discounting, the first two in the context of price-setting oligopolies, and the last in the context of collective effort-expending games.

[2] A set is totally disconnected if any connected subset of it must be either empty or a singleton.

1991, was the first paper to prove a folk theorem without public randomization. They start by addressing realizability and provide a lower bound for discount factors such that any point in the feasible set is representable as the discounted average of an infinite sequence of its vertices. This result goes by the name 'Sorin's lemma' as Sylvain Sorin (1986) had first analyzed a similar representation. Going from realizability to supportability requires work; as the standard text of Mailath and Samuelson (2006) notes, "The difficulty is that some of the continuation values generated by these sequences may fail to be even weakly individual rational." In trying to address this issue FM 1991 relies on complex as well as non-constructive arguments to build on top of Sorin's result which in the end, delivers strategies to support a specific payoff vector (like FM 1986), but is unable to yield a computable discount factor bound (unlike FM 1986).[3] Our key contribution is to use the the recursion-based notion of self-accessibility to address both realizability and individual rationality *simultaneously*. A set of payoffs is self-accessible for a discount factor vector if, for any point in the set, there is a pure action that can be played such that the induced continuation payoff also lies in the set. In the case of symmetric discounting we show that closed balls of small enough radii strictly inside the FSIR set are self-accessible for discount factors above a bound that can be explicitly computed using a simple non-linear program. This insight enables us to strengthen and offer a simpler yet completely constructive argument of the FM 1991 folk theorem.[4]

Why might knowing a discount factor bound be important? From a practical standpoint, the whole argument of repeated interaction being a prime motivator behind co-operation among otherwise selfish individuals is much more plausible when the required discount factor bound is, say, .7 rather than .9999. From a theoretical standpoint, even when we know that the set of equilibrium payoffs approaches the FSIR set as the common discount factor increases, it is of some interest to know if the approach is fast enough.[5]

The first paper to systematically study the asymmetric discounting case was LP 1999. They noted that in this setup, unlike in the symmetric discounting case, as the players become increasingly patient, a) realizable payoffs may lie outside the stage game feasible set and b) the limiting set of supportable payoffs could be very different from the FSIR set. However, they only analyze 2-player games, require fixed ratios of log discount factors, and their arguments crucially rely on existence of PRDs. They conjecture that it might be possible to remove this last restriction using techniques similar to those in FM 1991; however, we show by a simple counterexample that the building block of the FM approach, Sorin's Lemma does not hold in this situation. In the same setting as that of LP, but with possibly more than two players, Chen and Takahashi (2012) show how a sequence of action profiles can be supported in equilibrium if all its continuation payoffs are uniformly

[3]As an example of a non-constructive argument, consider covering a compact set S with an infinite collection of open balls and then choosing some characteristic of the finite subcover of S that is guaranteed because of compactness. But then, how does one figure out *which* finite subcover will do the job?

[4]Although the strategies we use are very similar to the ones used in FM 1991, one of our innovations is to design a punishment phase that does not become arbitrarily long as players become arbitrarily patient.

[5]A recent paper, Hörner and Takahashi (2016) analyzes this issue comprehensively for the first time.

bounded away from minmax values. However, without PRDs, only an 'approximate' folk theorem is obtained; the question of realizability is not tackled either.[6]

Repeated games with imperfect public monitoring are often studied using 'self-generation', a technique originally advanced in Abreu, Pearce, and Stacchetti (1990), and subsequently extended by Fudenberg and Levine (1994) and Fudenberg, Levine, and Maskin (1994).[7] Making use of this notion, Sugaya (2015) extends LP 1999 to prove a comprehensive folk theorem that applies to any finite number of players, perfect and imperfect public monitoring and possibly asymmetric discounting while dispensing with PRDs.[8] He too, works with discount factor vectors exhibiting fixed 'relative patience', specifically assuming that pairwise ratios of *discount rates*[9] are either constants or are converging to constants. He shows that for this setting, as players become more 'absolutely patient', the limiting sets of sequentially individually rational payoffs[10] and equilibrium payoffs are identical; however his methods does not help us compute either of these sets.

This places the folk theorems for repeated games with asymmetric discounting on a very different footing from those with symmetric discounting. One way these theorems inform the applied economist is by ruling out *impossible* payoffs when rational agents interact repeatedly *without him having any knowledge of the exact nature of discounting used by these agents*. Hence, if we cannot describe the potential equilibrium set[11] in terms of the stage game parameters, the existing results have not fulfilled that role.

To elaborate on this point, let F be the set of feasible payoffs in the stage game, and let F^+ be the subset of F where each player receives at least his individual rational payoff. Assume that this set is full-dimensional. Next, in the infinitely repeated game without PRD's, for a discount factor vector $\boldsymbol{\delta}$, let $\mathcal{F}(\boldsymbol{\delta})$ denote the set of realizable payoffs and $\mathcal{V}(\boldsymbol{\delta})$ the subgame-perfect equilibrium payoffs. Letting D be the set of discount factor vectors where all players discount at the same rate, Sorin's lemma gives us $F = \bigcup_{\boldsymbol{\delta}\in D} \mathcal{F}(\boldsymbol{\delta})$, while the Folk Theorem of FM 1991 gives us $F^+ \approx \bigcup_{\boldsymbol{\delta}\in D} \mathcal{V}(\boldsymbol{\delta})$ (here $A \approx B$ means that the sets A and B have zero Hausdorff distance between them). The following diagram then captures some of the inter-connections (with all subset relations being generally improper):

$$\begin{array}{ccc} \bigcup_{\boldsymbol{\delta}\in D} \mathcal{V}(\boldsymbol{\delta}) & \approx & F^+ \\ \cap & & \cap \\ \bigcup_{\boldsymbol{\delta}\in D} \mathcal{F}(\boldsymbol{\delta}) & = & F \end{array}$$

[6]Chen and Fujishige (2013) shows that the set of realizable payoffs in a *finitely* repeated game with unequal discounting is monotonically increasing in the length of the horizon.

[7]The difference between self-generation and self-accessibility is that while the former addresses realizability and supportability simultaneously, the latter addresses only realizability (but with individual-rationality). This decoupling is actually helpful as the rest of the paper demonstrates.

[8]See also Hörner and Olszewski (2005) who prove a folk theorem for almost-perfect private monitoring without using a PRD, though their results do not encompass unequal discounting.

[9]If δ_i is the i-th player's discount factor, his discount rate is $1/\delta_i - 1$.

[10]These are payoffs obtained via paths where each period each player's continuation payoff is at least his minmax payoff.

[11]By this we mean the set of payoffs that are equilibrium payoffs for *some* discount factor vector.

Allowing for asymmetric discounting, thanks to LP, we know that for 2-player games, neither the first nor the second horizontal relation holds (with PRDs though). Also, letting $\mathcal{F}^+(\boldsymbol{\delta})$ denote the set of sequentially individual rational payoffs for a given $\boldsymbol{\delta}$, thanks to Sugaya we know that for n-player games, $\bigcup_{\boldsymbol{\delta}\in D'} \mathcal{V}(\boldsymbol{\delta}) \approx \bigcup_{\boldsymbol{\delta}\in D'} \mathcal{F}^+(\boldsymbol{\delta})$ for certain subsets $D' \subset (0,1)^n$ that fix discount rate ratios. However, we do not know what those two sides of the (approximate) equality sign are. What is going on with asymmetric discounting is that there are sets V and R such that the following hold:

$$\begin{array}{ccccc}
\bigcup_{\boldsymbol{\delta}\in(0,1)^n} \mathcal{V}(\boldsymbol{\delta}) & \approx & V & \supset & F^+ \\
\wr\wr & & \parallel & & \parallel \\
\bigcup_{\boldsymbol{\delta}\in(0,1)^n} \mathcal{F}^+(\boldsymbol{\delta}) & \approx & V & \supset & F^+ \\
\cap & & \cap & & \cap \\
\bigcup_{\boldsymbol{\delta}\in(0,1)^n} \mathcal{F}(\boldsymbol{\delta}) & \approx & R & \supset & F
\end{array}$$

This paper, for the first time allows explicit construction of the sets V and R. As for R (i.e. the set of potentially realizable payoffs), we show that in the repeated game, players can simultaneously obtain essentially anything they can individually get in the stage game. Although this was foreshadowed by LP's work, the paths that we design to realize target points are more versatile than the ones considered by LP in that they do not force continuation payoffs to eventually settle on one of the vertices of F; in fact, they can eventually be made to stay close to any point in the interior of F. Next, we offer necessary and sufficient conditions for payoffs to be supportable. These conditions, respectively referred to as the *weak and strict diagonal conditions*, describe V, a set that we show can be constructed by solving certain linear programs.

As in the case of symmetric discounting, for general discounting, not only does one wish to identify the potentially supportable set, but one also wishes to know how (i.e. using what strategies) a specific payoff vector can be supported. Unfortunately, the existing literature has been silent on this question.[12] In contrast, we explicitly construct the equilibrium path and strategy profile to support a target payoff vector. The strategies we propose for supporting points in V (but outside F^+) are best described as 'turnpike strategies' for the following reason: The equilibrium path consists of two phases: a finite-period turnpike phase when the continuation payoffs are outside the FSIR set and an infinite-period post-turnpike phase when they are inside. If a deviation takes place on the turnpike phase, after punishing the deviant, play returns to the point on the turnpike phase where the deviation took place and continues on the turnpike phase while rewards to the punishers for compliance are offered in the non-turnpike phase only.

The paper is structured as follows. In the next section, we formally introduce the model and notation, define self-accessibility and briefly explain its significance. In section

[12]Sugaya, relying on fixed point arguments in the payoff space does not provide equilibrium paths or strategies while LP and Chen-Takahashi *assume* that we are already given an equilibrium path that is sequentially individual rational and only then construct an equilibrium strategy to support such a path.

3, a numerical example is discussed to illustrate the constructive nature of our arguments and present the flavor of some of our findings. Section 4 analyzes self-accessibility when discounting is symmetric and presents a constructive extension of FM 1991. Section 5 discusses self-accessibility in the asymmetric discounting case and explores the supportability of payoffs within the interior of the FSIR set. Section 6 addresses the issue of realizability in the asymmetric discounting case outside the feasible set. Section 7 defines weak and strict diagonalizability and shows how these conditions relate to supportability of payoff vectors outside the FSIR set. Section 8 concludes. All proofs are collected in an appendix.[13]

2 Preliminaries

2.1 Notation and The Model

We consider a standard infinitely repeated game of perfect monitoring with possibly unequal discounting. At each $t \in \{0, 1, 2, \ldots\}$ the (finite) stage-game $G = \langle I; (A_i)_i; (g_i)_i \rangle$ is played, where $I = \{1, \ldots, n\}$ is the set of players, A_i is player i's finite set of actions, $A := \times_i A_i$ is the set of all pure action profiles, and $g_i : A \to \mathbb{R}$ is player i's (vNM) payoff function. A mixed action of i is $\alpha_i \in \triangle A_i$, where for any set S, $\triangle S$ denotes the set of all probability distributions on the set S. Let $\boldsymbol{a}^{(t)} \in A$ be the (realized) action profile played at time t.[14] When player i discounts future payoffs using the discount factor δ_i, his average discounted utility defined over the infinite sequences of pure actions is

$$u_i(\left\{\boldsymbol{a}^{(t)}\right\}_{t=0}^{\infty}) := (1-\delta_i)\sum_{t=0}^{\infty} {\delta_i}^t g_i\left(\boldsymbol{a}^{(t)}\right).$$

Under perfect monitoring the public history at the end of period t is $h^t = (\boldsymbol{a}^{(0)}, \ldots, \boldsymbol{a}^{(t)}) \in A^{t+1}$. A pure strategy of i is a sequence of maps $s_i(t+1) : H^t \to A_i$ (for $t = -1, 0, 1, \ldots$) where H^t denotes the set of histories at the end of period t (with the convention that H^{-1} is the empty set). Mixed stategies are analogous, except that they map to the corresponding mixed actions $\triangle A_i$. This formulation implies that strategies cannot be conditioned on anything other than the history of actions actually played; in particular, there is no publicly observable random variable on whose realized value actions may be conditioned, and mixed actions are not observable; only their realizations are.

[13] It is worth remarking that although our main objective is to investigate the case of asymmetric discounting without public randomization, the results in Section 4 are not secondary to that goal; both the statement and the proof of the main result in Section 5 depend on the results we prove in Section 4. Similarly, even though we will eventually characterize the entire (possible) equilibrium set in Section 7, to prove those results we first need to establish the 'partial folk theorem' that is obtained in Section 5.

[14] In what follows vectors are boldfaced while scalars and sets are not. Sequence indices are denoted by superscripts and sometimes they are enclosed in parentheses to distinguish them from exponents or from another sequence denoted by the same letter; for example, $\boldsymbol{c}^l$ denotes the l-th vertex of a polytope C, while $\{\boldsymbol{c}^{(t)}\}$ denotes an infinite sequence of vertices each element of which is a $\boldsymbol{c}^l$ for some l. Coordinates of vectors are denoted by subscripts.

This describes the repeated game $G^{\infty}(\boldsymbol{\delta})$, with the discount factor vector $\boldsymbol{\delta} = (\delta_1, \ldots, \delta_n)$. In the special case where each player uses the same discount factor δ, we denote the game as $G^{\infty}(\delta\boldsymbol{\iota})$[15]. As in Section 1, we use $\mathcal{F}(\boldsymbol{\delta})$, $\mathcal{F}^{+}(\boldsymbol{\delta})$ and $\mathcal{V}(\boldsymbol{\delta})$ to denote the set of realizable, sequentially individual rational and supportable payoff vectors in the repeated game.

Let $C = g(A)$. player i's minmax value is $w_i := \min_{\boldsymbol{\alpha}_{-i} \in \times_{j \neq i}(\Delta A_j)} \max_{a_i \in A_i} \; g_i(a_i, \boldsymbol{\alpha}_{-i})$, Let $\boldsymbol{m}^i \in \times_{j=1}^{n}(\Delta A_j)$ be the profile that minmaxes player i, with the latter playing a best response. Whenever it is convenient to do so, we will assume without loss of generality that $w_i = 0$ for all i. The feasible set is $F := co(C)$ (convex hull of C), the feasible and weakly individually rational (FWIR) set is $F^{+} := \{\boldsymbol{x} \in F \,|\, x_i \geqslant w_i \; \forall i\}$ and the feasible strictly individually rational (FSIR) set is $F^{*} := \{\boldsymbol{x} \in F \,|\, x_i > w_i \; \forall i\}$. The lower boundary of F is $\underline{\partial} F := \{\boldsymbol{x} \in F : \nexists \boldsymbol{y} \in F \text{ such that } \boldsymbol{y} << \boldsymbol{x}\}$. We let $M = \max_i \{|g_i(\boldsymbol{a})| : \boldsymbol{a} \in A\}$; with $\boldsymbol{w} = \boldsymbol{0}$, when F^{*} is full-dimensional, this is strictly positive.

For later use, we recall the definitions of a few geometric terms. The affine hull of a set $X \subset \mathbb{R}^n$ is

$$aff(X) := \left\{ \sum_{l=1}^{k} \lambda^l \boldsymbol{x}^l \;\middle|\; \boldsymbol{x}^l \in X, \sum_{l=1}^{k} \lambda^l = 1, \; k \in \mathbb{N} \right\}.$$

For $\boldsymbol{x} \in X$, the affine (closed) ball with center $\boldsymbol{x}$ and radius r is $B_X(\boldsymbol{x}, r) := \{\boldsymbol{y} \in aff(X) : d(\boldsymbol{y}, \boldsymbol{x}) \leqslant r\}$, while $B(\boldsymbol{x}, r)$ denotes the usual (closed) ball in $\mathbb{R}^n$. The relative interior of X is

$$relint(X) := \{\boldsymbol{x} : \exists \; r > 0 \text{ such that } B_X(\boldsymbol{x}, r) \subset X\}.$$

When $X = co(C)$, where $C = \{\boldsymbol{c}^1, \ldots, \boldsymbol{c}^L\}$, and every point in C is an extreme point of X, each point in $relint(X)$ can be expressed as a convex combination of those points with strictly positive weights, i.e. $relint(X) = \left\{ \sum_{l=1}^{L} \lambda^l \boldsymbol{c}^l \;\middle|\; \lambda^l > 0, \; \sum_{l=1}^{L} \lambda^l = 1 \right\}$. The usual interior of a set S is denoted by $int(S)$.

For any set $S \subset \mathbb{R}^n$, and $M \subset \{1, \ldots, n\}$, $Proj_M(S)$ denotes the projection of S along the coordinates in M. For a finite set C of pure action payoffs in an n-player game, we define the corresponding 'feasible set' for players $1, \ldots, l$ ($l \leqslant n$) as $F(1, \ldots, l) := Proj_{\{1,\ldots,l\}} F$. Note that this is also $co(Proj_{\{1,\ldots,l\}}(C))$.

Finally, we introduce a new terminology: the *rectangular hull* of a bounded set S in $\mathbb{R}^n$, denoted as $re(S)$ is the smallest closed rectangle that contains S. Formally,

$$re(S) := \bigcap_{R \in \mathcal{R}} R \quad \text{where } \mathcal{R} = \left\{ \prod_{i=1}^{n} [a_i, b_i] : a_i \leqslant b_i, \prod_{i=1}^{n} [a_i, b_i] \supset S \right\}.$$

2.2 Self-accessibility

We now define *self-accessibility*, for possibly unequal discounting and explain its usefulness.

Definition. Let $C \subset \mathbb{R}^n$ be a finite set. A set $S \subset co(C)$ is said to be **self-accessible** relative to C for a vector $\boldsymbol{\delta} = (\delta_1, \ldots \delta_n) \in [0,1)^n$ if for any $\boldsymbol{x} \in S$ there exists $\boldsymbol{y} \in S$ and $\boldsymbol{c} \in C$ such that $x_j = (1 - \delta_j)c_j + \delta_j y_j$ for $j = 1, \ldots, n$.

[15] $\boldsymbol{\iota}$ is a vector of 1's, while $\boldsymbol{e}^i$, to be used later, is the i'th unit vector.

The definition is particularly intuitive for equal discounting. For some $\boldsymbol{x} \in S \subset co(C)$, suppose that we can find $\boldsymbol{c} \in C$, $\boldsymbol{y} \in S$ and $\delta \in [0,1)$ such that $\boldsymbol{x} = (1-\delta)\,\boldsymbol{c} + \delta\boldsymbol{y}$; this is a 'dynamic programming decomposition' of the target payoff $\boldsymbol{x}$, with the restrictions that the current payoff $\boldsymbol{c}$ is generated by a pure action profile and the continuation payoff $\boldsymbol{y}$ lies in the set S itself. If there is a *uniform* $\delta \in [0,1)$ for which any point $\boldsymbol{x} \in S$ can be written as the $(1-\delta, \delta)$ convex combination of a pure-action payoff and a continuation payoff within the set itself, then S is self-accessible for $\boldsymbol{\delta} = \delta\boldsymbol{\iota}$.

To see the usefulness of the notion with an arbitrary discounting structure, let $S \subset co(C)$ be self-accessible for $\boldsymbol{\delta}$ and let $\boldsymbol{x} \in S$. It follows that there exists $\boldsymbol{c}^{(0)} \in C$ such that $x_j = (1-\delta_j)c_j^{(0)} + \delta_j y_j^1$ for each j where $\boldsymbol{y}^1$ is also in S. Because of the latter, we can write $y_j^1 = (1-\delta_j)c_j^{(1)} + \delta_j y_j^2$ for each j for some $\boldsymbol{c}^1 \in C$ and $\boldsymbol{y}^2 \in S$. By induction there is a sequence of vertices $\{\boldsymbol{c}^{(t)}\}_{t\geqslant 0}$ such that

$$x_j = (1-\delta_j)\sum_{t=0}^{\tau} \delta_j^t c_j^{(t)} + \delta_j^{\tau+1} y_j^{\tau+1} \quad \forall j,\ \forall \tau.$$

Since $\delta_j < 1$ and S is bounded, we have $\parallel \delta_j^{\tau+1} y_j^{\tau+1} \parallel \to 0$ as $\tau \to \infty$. Hence any point $\boldsymbol{x}$ in a self-accessible set S has a representation $x_j = (1-\delta_j)\sum_{t\geqslant 0}\delta_j^t c_j^{(t)}$. Thus, any point lying in a set that is self-accessible for a given $\boldsymbol{\delta}$ vector can be represented as the coordinate-wise discounted average of a sequence of pure action payoffs for that $\boldsymbol{\delta}$ vector.[16]

3 A Numerical Example

In this section we present a numerical example to a) underscore the computational advantages of supporting a specific payoff vector using self-accessibility with symmetric discounting and b) preview our results on asymmetric discounting. Although the general results presented later are valid for any n-player stage game with a full-dimensional FSIR set, in order to abstract away from ancillary issues, we choose an asymmetric version of the Prisoner's dilemma game.

3.1 Supporting the Nash Bargaining Point with Symmetric Discounting

Each of two players simultaneously choose one of two actions (A) and (N) with the payoff matrix displayed below.

	A	N
A	$(4,2)$	$(9,0)$
N	$(0,7)$	$(5,5)$

Henceforth, we refer to the payoff vectors $(4,2)$, $(9,0)$, $(5,5)$ and $(0,7)$ as $\boldsymbol{c}^1, \boldsymbol{c}^2, \boldsymbol{c}^3$

[16] To make this sequence well-defined, when a point in S can be decomposed in more than two ways, we can use some pre-assigned arbitrary ordering among the vertices to decide which current action to choose.

and $\boldsymbol{c}^4$ respectively. The unique dominant strategy equilibrium (A,A) is inefficient. The (efficient) Nash Bargaining payoff vector (where the set of possible agreement payoffs is the feasible set and $\boldsymbol{c}^1$ is the disagreement point) is $\boldsymbol{n} = (5.700, 4.125)$, which is a convex combination of $\boldsymbol{c}^3$ and $\boldsymbol{c}^2$ with weights .825 and .175 respectively (see Figure 1). We will like to obtain this payoff vector in an SPNE.

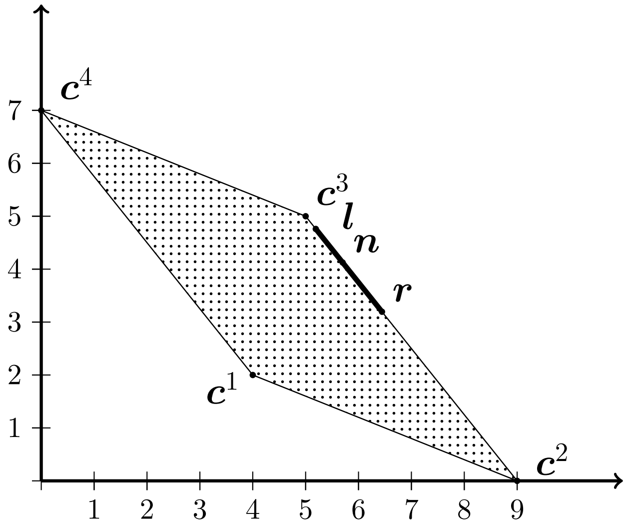

Figure 1: Realizable Payoffs in an Asymmetric Prisoner's Dilemma Game

Let us first examine the realizability of $\boldsymbol{n}$. Choose a one-dimensional closed ball $S \subset co(\{\boldsymbol{c}^2, \boldsymbol{c}^3\})$ containing $\boldsymbol{n}$, with extreme points $\boldsymbol{l}$ and $\boldsymbol{r}$ ($\boldsymbol{l}$ being closer to $\boldsymbol{c}^3$). Let $\boldsymbol{l} = \lambda(5,5) + (1-\lambda)(9,0) = (9-4\lambda, 5\lambda)$ and $\boldsymbol{r} = \mu(5,5) + (1-\mu)(9,0) = (9-4\mu, 5\mu)$ with the requirements

$$1 \geqslant \lambda \geqslant .825 \quad \text{and} \quad 0 \leqslant \mu \leqslant .825 \tag{3.1}$$

to ensure that S contains $\boldsymbol{n}$. We shall find a cutoff $\underline{\delta}$ above which S is self-accessible (relative to $\{\boldsymbol{c}^2, \boldsymbol{c}^3\}$).

Any $\boldsymbol{x} \in S$ can be written as $\theta(5,5) + (1-\theta)(9,0) = (9-4\theta, 5\theta)$, with $\mu \leqslant \theta \leqslant \lambda$. Let $\delta(\boldsymbol{x}, \boldsymbol{c}^3)$ be the lowest value of δ in $[0,1]$ such that $\boldsymbol{x} = (1-\delta)\boldsymbol{c}^3 + \delta\boldsymbol{y}$ for some point $\boldsymbol{y} \in S$. Since the farthest continuation payoff within S is $\boldsymbol{r}$, $\delta(\boldsymbol{x}, \boldsymbol{c}^3)$ solves $\boldsymbol{x} = (1-\delta)\boldsymbol{c}^3 + \delta\boldsymbol{r} = (1-\delta)(5,5) + \delta(9-4\mu, 5\mu)$, and therefore $\delta(\boldsymbol{x}, \boldsymbol{c}^3) = (1-\theta)/(1-\mu)$. Similarly, define $\delta(\boldsymbol{x}, \boldsymbol{c}^2)$ as the the lowest value of δ in $[0,1]$ such that $\boldsymbol{x} = (1-\delta)\boldsymbol{c}^2 + \delta\boldsymbol{y}$ for some point $\boldsymbol{y} \in S$; using an analogous argument this is seen to be θ/λ. If the discount factor δ is at least as much as $\delta^*(\boldsymbol{x}) := \min\{\delta(\boldsymbol{x}, \boldsymbol{c}^2), \delta(\boldsymbol{x}, \boldsymbol{c}^3)\}$, then $\boldsymbol{x}$ can be attained by playing one of the vertices $\boldsymbol{c}^2$ or $\boldsymbol{c}^3$ with the continuation payoff lying in S.[17] Finally, note that the maximum of $\delta^*(\boldsymbol{x})$ as $\boldsymbol{x}$ varies over S is achieved for a point $\bar{\boldsymbol{x}}$ where $\delta(\bar{\boldsymbol{x}}, \boldsymbol{c}^2) = \delta(\bar{\boldsymbol{x}}, \boldsymbol{c}^3)$, i.e. for $\underline{\theta}$ that solves $(1-\theta)/(1-\mu) = \theta/\lambda$; hence, $\underline{\theta} = \frac{\lambda}{1-\mu+\lambda}$. Plugging it in the expression

[17] More accurately, we mean "...by playing the actions corresponding to the vertices $\boldsymbol{c}^2$ or $\boldsymbol{c}^3$...". Here and elsewhere we indulge in this slight abuse of notation for brevity's sake.

of $\delta(\boldsymbol{x}, \boldsymbol{c}^2)$ gives $\underline{\delta}$, a lower bound for the discount factor δ that ensures the self-accessibility of S:

$$\underline{\delta} := \max_{\boldsymbol{x}\in S} \delta^*(\boldsymbol{x}) = \max_{\boldsymbol{x}\in S} \min\{\delta(\boldsymbol{x}, \boldsymbol{c}^2), \delta(\boldsymbol{x}, \boldsymbol{c}^3)\} = \frac{1}{1-\mu+\lambda}. \tag{3.2}$$

When (3.1) and (3.2) hold, any $\boldsymbol{x} = (9-4\theta, 5\theta)$ in S can be realized by playing a sequence of actions from the set $\{\boldsymbol{c}^2, \boldsymbol{c}^3\}$. For $\theta < \underline{\theta}$, $\boldsymbol{c}^2$ starts the sequence, otherwise $\boldsymbol{c}^3$ does.

Next, to support $\boldsymbol{n}$, we wish to deter deviation from the prescribed path via a grim trigger strategy à la Friedman (1971): $\boldsymbol{c}^1$ is played forever as soon as any deviation is detected. To this end we note that if the current (on-equilibrium) continuation payoff is given by a $\theta < \underline{\theta}$, since $\boldsymbol{c}^2$ is to be played next, we need only worry about player 2's incentive compatibility. This is achieved by the following constraint that ensures that for player 2, receiving the worst payoff in the corresponding region of S (which is 5μ) is at least as good as receiving the best deviation payoff once (which is 2) and then being minmaxed forever:

$$5\mu \geqslant 2(1-\delta) + 2\delta. \tag{3.3}$$

On the other hand, if the current (on-equilibrium) continuation payoff is given by a $\theta \geqslant \underline{\theta}$, which means the current prescribed action is $\boldsymbol{c}^3$, we need to deter deviations by both players. Similar reasoning as above shows that player 1's incentive compatibility requires:

$$9 - 4\lambda \geqslant 9(1-\delta) + 4\delta, \tag{3.4}$$

while player 2's incentive compatibility requires:

$$\frac{5\lambda}{1-\mu+\lambda} \geqslant 7(1-\delta) + 2\delta. \tag{3.5}$$

The minimum δ satisfying (3.1), (3.2), (3.3), (3.4) and (3.5) is .731, and the corresponding ball S is given by $\lambda = .914, \mu = .547$. This suggests that above a reasonable discount factor the Nash bargaining payoff $\boldsymbol{n}$ can be implemented in an SPNE.[18]

3.2 Expanding Possibilities With Asymmetric Discounting

If all players use the same discount factor, any discounted average payoff vector must stay inside the feasible set. What if they do not? Aided by the figure below, we preview some of the results to follow by explaining how they apply in the context of the current example.

[18]There is nothing special about the point $\boldsymbol{n}$; the analysis we just did could be carried out for any (individually rational) point on the Pareto frontier. For instance, Mailath Obara and Sekiguchi (2002) concerns itself with finding a bound for the discount factor that gives one player (say player 1) the highest payoff subject to the other player receiving her individually rational payoff. Since nothing in our analysis disallows the target point to be on the boundary of S, in this case, we will require $\boldsymbol{r}$ to be the point $(7.4, 2)$ and then replacing (3.1) by the following two constraints: $\lambda \geqslant \mu$ and $\mu = .4$ we could solve the optimization problem as before to find the relevant bound. It works out to be .781.

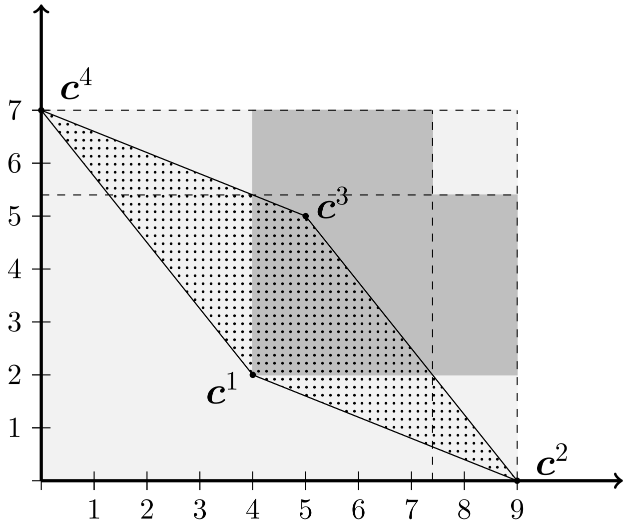

Figure 2: Realizable and Supportable Payoffs with Asymmetric Discounting

Theorem 3 will later show that any payoff vector in $int(re(F))$ (which in this case is the open rectangle $(0,9) \times (0,7)$) can be realized for large enough discount factors if we are allowed to choose the *relative* patience of the two players by fixing the ratio of their discount rates. Which of these are equilibrium payoffs and for what kind of discount factors? Theorem 5 demonstrates that there are some points in the open rectangle $(0,9) \times (0,7)$ that can not be supported in equilibrium no matter what the discount factors are. These are points in the small lightly shaded rectangle in the north-east of Figure 2; a payoff vector such as $(8.9, 6.9)$ where both *both* players receive close to their maximum payoffs is ruled out.

On the positive front, Theorem 2 will show that points in the interior of F^* (the dotted *and* darkly shaded region) can be supported in equilibrium for large enough discount factors and arbitrary discount rate ratios. Moreover, using ideas developed in Theorem 6 we can show that if both players are sufficiently patient and player 2 's discount rate is *sufficiently lower* relative to player 1, then it is possible to support points in the open rectangle $(4,9) \times (2,5.4)$. The coordinates of this latter rectangle are arrived at by letting player 1's payoff to range between his minmax payoff (4) and the maximum he can receive in the game (9), whereas player 2's payoff is allowed to range between his minmax (2) and the maximum he can receive *subject to giving player 1 his minmax amount*. Analogously, if both players are sufficiently patient and player 1 is sufficiently more patient relative to player 2, points in the open rectangle $(4,7.4) \times (2,9)$ can be supported (7.4 is the maximum player 1 can receive subject to giving player 2 his minmax amount). Points that are common to both rectangles may be supported by (large enough) discount factors exhibiting a wide variety of relative patience. To summarize then, the dotted region in Figure 2 is the feasible set, points in $(0,9) \times (0,7)$ are realizable payoffs for some discount factors, while the interior of the darkly shaded region are supportable payoffs for some discount factors.

4 The Case of Symmetric Discounting

This section discusses how self-accessibility simplifies the arguments of FM 1991's main result, makes discount factor bounds computable and delivers some new results as well.

4.1 Self Accessibility Under Symmetric Discounting

The main building block of the FM 1991 is Sorin's Lemma which addresses the question of realizability. The lemma states the following: Suppose $\boldsymbol{x} \in \mathbb{R}^n$ is in the convex hull of $C = \{\boldsymbol{c}^1, \boldsymbol{c}^2, \ldots, \boldsymbol{c}^{\mathrm{L}}\}$. Then, for all $\delta \geqslant 1 - 1/L$ there exists a sequence $\{\boldsymbol{c}^{(t)}\}_{t=0}^{\infty}$ in C such that $\boldsymbol{x} = (1-\delta)\sum_{t=0}^{\infty} \delta^t \boldsymbol{c}^{(t)}$.[19] Armed with this result, FM then tackle the problem of ensuring that continuation payoffs stay close enough to the original payoff (so as to maintain individual rationality) via their Lemma 2, which makes critical use of Sorin's lemma but has a complex argument and offers little computational guidance.

Using our terminology, Sorin's Lemma shows that the entire feasible set is self-accessible relative to its extreme points for large enough discount factors. However, we can show that for any set of points in $\mathbb{R}^n$, affine balls contained in the relative interior of the convex hull of those points are also self-accessible (for large enough discount factors). When a target payoff is in F^*, by placing it anywhere inside a 'small' such ball, we can thus achieve both realizability and (sequential) individual rationality in a single step. This approach bypasses the need for Sorin's Lemma altogether and as a bonus, we can also easily compute a relevant discount factor bound.

Proposition 1. *Suppose $C' \subset C = g(A)$ is a set of points in $\mathbb{R}^n$ where $X = Co(C')$ need not be full-dimensional. Let $S = B_X(\boldsymbol{o}, r) \subset relint(X)$ be some affine ball with center $\boldsymbol{o}$ and radius $r > 0$. Then $\exists\ \underline{\delta} \in (0,1)$ such that S is self-accessible relative to C' for any vector $\delta\boldsymbol{\iota}$ with $\delta \geqslant \underline{\delta}$. This $\underline{\delta}$ is computable by solving a nonlinear maximization problem with linear/quadratic objective and constraint functions.*

As this proposition is central to the computability of discount factor bounds, we provide some intuition behind the construction of $\underline{\delta}$. Fix an affine closed ball $S = B_X(\boldsymbol{o}, r)$ in the relative interior of $co(C')$. Take any point $\boldsymbol{x}$ in S. Let $\delta(\boldsymbol{x}, \boldsymbol{c})$ be the smallest value of $\delta \in [0,1]$ satisfying the dynamic programming decomposition $\boldsymbol{x} = (1-\delta)\,\boldsymbol{c} + \delta\boldsymbol{y}$ for some $\boldsymbol{y} \in S$. The geometrical interpretation of this function is as follows: consider the line connecting $\boldsymbol{c}$ and $\boldsymbol{x}$; it cuts the surface of the ball at two points, one on the same side of $\boldsymbol{x}$ as $\boldsymbol{c}$ and one on the opposite side (they could be same if the line is tangent to the surface). Call this latter point $\boldsymbol{y}$ ($\boldsymbol{y}$ could be $\boldsymbol{x}$ itself). Thus $\boldsymbol{x}$ is a convex combination of $\boldsymbol{c}$ and $\boldsymbol{y}$ with weight $\delta(\boldsymbol{x}, \boldsymbol{c})$ on $\boldsymbol{y}$. It is not hard to find a formula expressing this weight as a continuous function of $\boldsymbol{x}$.

Now we assert that $\delta(\boldsymbol{x}, \boldsymbol{c}) < 1$ for some vertex $\boldsymbol{c}$. If $\boldsymbol{x}$ is in $relint(S)$, any vertex

[19] Actually the bound stated in this lemma is not tight. Using Caratheodory's Theorem it can be shown that the exact tight bound is $1-1/m$ where $m = \min\{L, n+1\}$, rather than $1-1/L$. Also, in his original 1986 paper, Sorin, using a result from Fenchel (1929) obtains the bound $1 - \frac{1}{n}$ when mixed actions are allowed.

$\boldsymbol{c} \in C'$ works but if $\boldsymbol{x}$ is on the (relative) boundary of S, not all vertices do. However, since S lies in the (relative) interior of the convex hull of C', we can use any point in C' that is separated from the ball by the supporting hyperplane to S at $\boldsymbol{x}$. In the accompanying figure, $C' = \{\boldsymbol{c}^1, \boldsymbol{c}^2, \boldsymbol{c}^3\}$ and X is the triangular region with those three points acting as vertices (it could be a face of the feasible set). The shaded region is S; $\boldsymbol{x}$ is a point in S. If we try to extend a line from $\boldsymbol{c}^i$ to $\boldsymbol{x}$ to the farthest point in S, we reach $\boldsymbol{y}^i$. Here, $\delta(\boldsymbol{x}, \boldsymbol{c}^1) = 1$, $\delta(\boldsymbol{x}, \boldsymbol{c}^2) = \frac{\|\boldsymbol{x}-\boldsymbol{c}^2\|}{\|\boldsymbol{y}^2-\boldsymbol{c}^2\|}$ and $\delta(\boldsymbol{x}, \boldsymbol{c}^3) = \frac{\|\boldsymbol{x}-\boldsymbol{c}^3\|}{\|\boldsymbol{y}^3-\boldsymbol{c}^3\|}$.

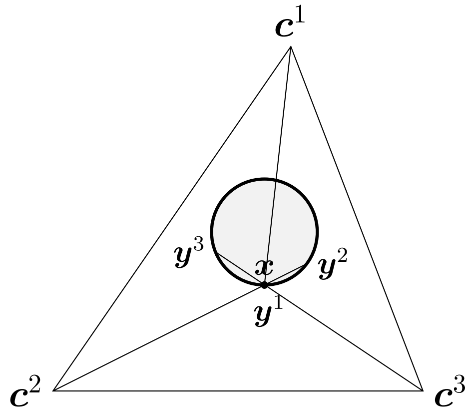

Figure 3: How $\delta^*(\boldsymbol{x})$ is constructed

Let $\delta^*(\boldsymbol{x}) := \min_{\boldsymbol{c} \in C} \delta(\boldsymbol{x}, \boldsymbol{c})$. For our figure, it turns out that $\delta^*(\boldsymbol{x}) = \delta(\boldsymbol{x}, \boldsymbol{c}^3)$. Now consider maximizing δ^* over S; Weierstrass's theorem guarantees that the maximum is attained. Since $\delta^* < 1$ throughout, this maximum $\underline{\delta}$ is less than unity and the convexity of S ensures that S is self-accessible above $\underline{\delta}$.

Once the affine ball $B_X(\boldsymbol{o}, r)$ is self-accessible relative to a set of vertices C' for a discount factor vector $\boldsymbol{\delta} = \delta \boldsymbol{\iota}$, as was explained in section 2.2, any vector $\boldsymbol{x}$ in that ball can be expressed as a discounted average of a sequence of vertices from C'; we let $\{\boldsymbol{a}^{(t)}(\boldsymbol{x}, B_X(\boldsymbol{o}, r), \delta)\}_{t=0}^{\infty}$ denote this sequence.

One of the advantages of the self-accessibility approach over the FM 1991 approach is that we are easily able to provide a *computable* uniform bound on discount factors that guarantees realizability of each point in any geometrically well-described compact set while keeping continuation payoffs within a certain small fixed distance of the original point. Non-constructive arguments involving open covers of compact sets admitting to finite subcovers are not required.[20] Extending ideas that are used to prove Proposition 1, our next result shows that we can find a bound on discount factors that makes a *collection* of balls self-accessible where each has a certain fixed radius (say ω) and a center that lies within a fixed ball with a different radius (say $\bar{r}$); in Figure 4, these balls are colored dark grey, while the ball within which their centers lie is colored light grey.

[20] These uniform bounds become relevant if unobeservable mixed strategies are needed to minmax a player. In such situations, post-punishment plays are not known beforehand; they are calculated based on the realizations of mixed actions during the punishment period.

Proposition 2. *Let* $\boldsymbol{x} = \sum_{l=1}^{K} \lambda^l \boldsymbol{c}^l$ *with* $\sum_{l=1}^{K} \lambda^l = 1$ *and* $\lambda^l > 0$ *for each* l. *Let* $C' = \{\boldsymbol{c}_1, \ldots, \boldsymbol{c}_K\}$ *and* $X = co(C')$. *Let* $\bar{r} > 0$ *and* $\omega > 0$ *be such that* $B_X(\boldsymbol{x}, \bar{r} + \omega) \subset F^* \cap relint(X)$. *Then, we can find* $\underline{\delta} < 1$ *such that for any* $\delta \in (\underline{\delta}, 1)$, *any* $B(\boldsymbol{x}', \omega)$ *where* $\boldsymbol{x}' \in B_X(\boldsymbol{x}, \bar{r})$ *is self-accessible relative to* C' *for* $\boldsymbol{\delta} = \delta \boldsymbol{\iota}$. *Furthermore,* $\underline{\delta}$ *is computable by solving a nonlinear maximization problem.*

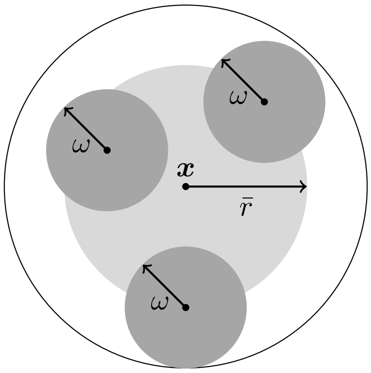

Figure 4: Finding a bound for which a collection of balls is self-accessible

4.2 A Constructive Folk Theorem via Self-Accessibility

As soon as we observe that affine balls with small enough radii are self-accessible, there hardly remains any difference between the two problems where we are trying to support a payoff vector with and without PRDs. If, assuming the existence of PRDs, one can find a discount factor bound (say $\underline{\delta}_1$) for which incentive compatibility conditions are *strict*, then all we have to do in the PRD-less case is to work with small enough balls around points which could potentially be played in the first case, find bounds which will make these balls self-accessible and then we can support the same target payoff vector in the second case using a grand bound which is the maximum of all these bounds and $\underline{\delta}_1$.

This insight drives the main result of this section which is a fully constructive version of the main result in FM 1991. It differs from the original version in three important aspects. First, because of our reliance on self-accessibility, all paths, whether on or off equilibrium become recursively computable. Second, Propositions 1 and 2 and the fact that our method of proof simplifies incentive compatibility conditions, reducing them to only two *linear* inequalities, allow us to compute a discount factor bound needed to support a given payoff vector in $F^* \backslash \underline{\partial} F$. Third, our proof reveals that the number of punishment periods during which a deviating player is minmaxed need not become arbitrarily large as δ goes to one. While in FM 1991, the punishment length is of the order of $(-ln\ \delta)^{-1}$, in our proof, we fix it once and for all.

Theorem 1. *Let* F^* *be full-dimensional and let* $\boldsymbol{v} \in F^* \backslash \underline{\partial} F$. *Then, there exists* $\underline{\delta} \in (0, 1)$ *such that for all* $\delta \in (\underline{\delta}, 1)$ *there is an SPNE that does not use a PRD and has discounted average payoff* $\boldsymbol{v}$. *This bound is computable using the NLPs provided by Propositions 1 and*

2 and two linear inequalities. All paths, on and off equilibrium, are recursively computable as well. Punishment (minmaxing) period lengths are not δ-dependent.

A detailed proof is provided in the appendix for the sake of completeness and also because, the proof applies with little change to a similar theorem in the asymmetric discounting setting. It relies on the standard architecture of equilibrium strategies introduced in FM 1986 characterized in terms of three phases, which is well-understood in the literature. It may be instructive here though to consider what self-accessibility brings to the table that allows us to keep the number of punishment periods δ-independent. To that end, consider the situation where player i has deviated, has been minmaxed, and play now has shifted into the so-called Phase III(i) where players are supposed to receive the (continuation) payoff vector $\boldsymbol{u}$.[21] If one had access to PRD's, one would prescribe a path where in every period, an action generating $\boldsymbol{u}$ would be played. Without PRDs however, players play a sequence of actions that generates $\boldsymbol{u}$ as a discounted average, while the continuation payoffs stay, say, ε-close to $\boldsymbol{u}$. But now suppose, after this path is started, at some point, i's continuation payoff becomes $u_i - \varepsilon$. Then, if the number of punishment periods is a constant independent of the discount factor, a sufficiently patient i might want to re-start his own punishment by deviating! This is why FM 1991 needs to let the punishment phase become unboundedly large as the discount factor approaches one. In our proof, when Phase III(i) is started, the target payoff is not $\boldsymbol{u}$, it is $\boldsymbol{u} - \varepsilon \boldsymbol{e}^i$, the lowest point (from i's perspective) in the self-accessible ball $B(\boldsymbol{u}, \varepsilon)$ (since there is no requirement that a target payoff must be in its center). Thus the perverse situation described above never arises and δ-independent punishment periods can indeed be devised to wipe out any gain from deviation.

5 Asymmetric Discounting: Supporting Points in $int(F^*)$

In this section, we show that any point in the interior of F^* is both realizable and supportable provided each player's discount factor is large enough. In subsection 5.1, we begin by showing that the natural counterpart of Sorin's lemma in the asymmetric case does not hold which negates Lehrer and Pauzner's conjecture and makes the FM 1991 approach to the problem inapplicable. We also show that if the discount rate ratios are not fixed, not even small balls are self-accessible (even when the discount factors are large). Hence, to insure self-accessibility for a set S, either we need to keep discount rates fixed or allow S to assume a geometrical shape which is not a closed ball. In subsection 5.2, we provide two positive results to demonstrate these possibilities. Subsection 5.3 offers an asymmetric counterpart of Theorem 1.

[21] It is during this phase that players $j \neq i$ are rewarded for participating in i's punishment phase.

5.1 Two Negative Results

One might hope that the following 'global' extension of Sorin's lemma to unequal discounting holds: If $C \subset \mathbb{R}^n$ is a finite set there exists $\underline{\delta} \in [0,1)$ such that if $\delta_j \geqslant \underline{\delta}$ for $1 \leqslant j \leqslant n$, and $\boldsymbol{x} \in co(C)$, there exists a sequence of points $\{\boldsymbol{c}^{(t)}\}_{t=0}^{\infty}$ in C for which $x_j = (1-\delta_j)\sum_{t=0}^{\infty} {\delta_j}^t\, c_j^{(t)}$. It is easy to see that as stated, the conjecture cannot be true: for example, when $n = L = 2$ (whereupon $co(C)$ is a one-dimensional set with just two vertices) we need the two discount factors to be equal in order to realize points in $co(C)$. Is this then an artifact of $co(C)$ not being full-dimensional or $\boldsymbol{x}$ lying on the boundary rather than in the interior of $co(C)$? Are points that defy the desired representation non-generic? Unfortunately, the problem runs deeper.

Counterexample 1. Consider the simple, two-player game with the payoff matrix displayed below.

	L	R
U	$(1,0)$	$(0,0)$
D	$(0,0)$	$(0,1)$

Let any $\underline{\delta} \in (0,1)$ be given. We will show the existence of an open set in $co(C)$ and a δ_1, δ_2 pair where the impatient player's discount factor is exactly $\underline{\delta}$, such that no point in that open set is representable using the vertices in C and the given discount factors. To that end, suppose, one can find real numbers $\varepsilon_1, \varepsilon_2, \delta_1, \delta_2$, and an integer T with the following properties:

$$0 < \varepsilon_2 < \varepsilon_1 < 1, \tag{5.1}$$

$$\underline{\delta} = \delta_2 < \delta_1 < 1, \tag{5.2}$$

$${\delta_2}^T < \varepsilon_2, \tag{5.3}$$

$$(1-\delta_1)\,{\delta_1}^{T-1} > \varepsilon_1. \tag{5.4}$$

We assert that the point $(1-\varepsilon_1, \varepsilon_2)$, which is in $int(co(C))$ by (5.1), is not realizable for discount factors (δ_1, δ_2). To prove this, we first prove inductively that if

$$x_1 = 1 - \varepsilon_1 = (1-\delta_1)\sum_{t=0}^{\infty} {\delta_1}^t\, x_1^{(t)},$$

then $\boldsymbol{x}^{(0)} = \boldsymbol{x}^{(1)} = \ldots = \boldsymbol{x}^{(T-1)} = (1,0)$, i.e. $(1,0)$ must be played for the first T periods. If $\boldsymbol{x}^{(0)} \neq (1,0)$, then even if $(1,0)$ were to be played in each subsequent period, x_1 could be at most δ_1. This would mean $1-\varepsilon_1 \leqslant \delta_1$, or $\varepsilon_1 \geqslant (1-\delta_1)$. However, (5.4) rules this out. If $\boldsymbol{x}^{(0)} = (1,0)$ but $\boldsymbol{x}^{(1)} \neq (1,0)$, then $x_1 \leqslant (1-\delta_1) + (\delta_1)^2$, which implies $(1-\delta_1) + (\delta_1)^2 \geqslant 1 - \varepsilon_1$; from this it follows that $\varepsilon_1 \geqslant (1-\delta_1)\,\delta_1$, which violates (5.4). Proceeding this way, $(1,0)$ must be played at least the first T times. But then $x_2 \leqslant (\delta_2)^T$, which violates (5.3) if $x_2 = \varepsilon_2$.

It remains to show that one can indeed satisfy the properties (5.1) – (5.4) by judicious choice of $\varepsilon_1, \varepsilon_2, \delta_1, \delta_2$, and T. Note that the (strict) inequality $(1-\delta_1)\,\delta_1^{T-1} > {\delta_2}^T$ holds when

$$\left(\frac{\delta_1}{\delta_2}\right)^T > \frac{\delta_1}{1-\delta_1} \tag{5.5}$$

holds and hence, if we set $\delta_2 = \underline{\delta}$, and choose δ_1 to be *any* real number in $(\underline{\delta}, 1)$, there will exist an integer T such that (5.5) will hold. We can then choose an open set of $\varepsilon_1, \varepsilon_2$ pairs such that $0 < \delta_2^T < \varepsilon_2 < \varepsilon_1 < (1-\delta_1){\delta_1}^{T-1} < 1$ and we would have fulfilled all our requirements. This shows that, no matter how high the discount factors are forced to be, if we are allowed to choose them unequal, we can find an open set of points in a full-dimensional feasible set of payoffs that cannot be realized without PRDs. The counterexample also demonstrates that the set of realizable payoffs (and even its closure) can fail to be convex in every situation when discounting is not perfectly symmetric. Thus, Sorin's Lemma is an extremely non-generic, knife-edge phenomenon.[22] ■

It then seems natural to invoke self-accessibility. In analogy with the equal-discounting case, one might conjecture the following 'local' version of Sorin's Lemma: If $C \subset \mathbb{R}^n$ is a finite set with an n-dimensional convex hull, for any $\boldsymbol{x}$ in the interior of $co(C)$ there exists a quantity $r > 0$ and a cutoff $\underline{\delta}$ such that the ball $B(\boldsymbol{x}, r)$ is self-accessible for discount factor vectors of the form $\boldsymbol{\delta} = (\delta_1, \ldots, \delta_n)$ with $\delta_i \geqslant \underline{\delta}$. The following shows that this conjecture is false as well.

Counterexample 2. Consider another two-player game with the following payoff matrix:

	L	R
U	(1,-1)	(-1,-1)
D	(-1,0)	(1,1)

Take any ball $B((0,0), r)$ with $r < 1$. Suppose that the conjecture holds for $0 < \underline{\delta} < 1$. In that case, if $\boldsymbol{\delta} = (\delta_1, \delta_2) \geqslant (\underline{\delta}, \underline{\delta})$ and $\boldsymbol{x}$ is in the ball, there exists $\boldsymbol{c} \in C$ such that if we define $\boldsymbol{y}(\boldsymbol{\delta})$ via the equation

$$x_i = (1-\delta_i)\,c_i + \delta_i\,y_i(\boldsymbol{\delta}) \text{ for } i = 1, 2, \tag{5.6}$$

then, $\boldsymbol{y}(\boldsymbol{\delta}) \in B((0,0), r)$. Write $\underline{\delta} = 1/(1+\theta)$ and choose $\delta_2 = \underline{\delta}$, and $\delta_1 = 1/(1+k_1\theta)$ for $k_1 \in (0,1)$; this ensures that $\delta_1 \geqslant \underline{\delta}$. Now specifically let us consider the point $\boldsymbol{x} = (r, 0)$ and ask which $\boldsymbol{c} \in C$ will make the $\boldsymbol{y}(\boldsymbol{\delta})$ given via (5.6) lie in the ball. It is easy to see that the vertices $(-1,1)$ and $(-1,-1)$ are ruled out. By symmetry, $(1,-1)$ works if and only if $(1,1)$ works. For $\boldsymbol{c} = (1,1)$, equation (5.6) gives $\boldsymbol{y}(\boldsymbol{\delta}) = (r + k_1\theta(r-1), -\theta)$. Hence, $\boldsymbol{y}(\boldsymbol{\delta}) \to (r, -\theta)$ as $k_1 \to 0$; for any given θ, this is strictly outside the ball. Hence for every θ, there exists a k_1 (i.e. for every $\delta_2 = \underline{\delta}$, there exists a $\delta_1 > \underline{\delta}$) for which $\boldsymbol{y}(\boldsymbol{\delta})$ is outside

[22] This seems to echo the findings of Salonen and Vartiainen (2008); however, the stage game they use is not full-dimensional, while ours is.

the ball no matter what $\boldsymbol{c}$ is used. ■

5.2 Two Positive Results

We now present two positive results on self-accessibility with asymmetric discounting. The first focuses on the self-accessibility of a *fixed* ball and uses the same parametrization of the discount factors as in Sugaya (2015) where δ_i is written as $1/(1+k_i\theta)$ with $\boldsymbol{k}$ fixed.[23] This result will be used in the next section to examine which points are realizable. The second demonstrates the existence of *flexible*, $\boldsymbol{\delta}$-dependent self-accessible sets but places no restrictions on discount rate ratios (and hence, relative patience). It will be used in this section to support points inside $int(F^*)$.

Proposition 3. *Let $C \subset \mathbb{R}^n$ be finite, and let $X = co(C)$ be full-dimensional and contain in its interior the ball $B(\boldsymbol{o}, r)$ with $r > 0$. For any $\boldsymbol{k} \in \mathbb{R}^n_{++}$, there exists $\bar{\theta}(\boldsymbol{o}, r, \boldsymbol{k}) > 0$ such that for any $\theta \in (\,0,\, \bar{\theta}(\boldsymbol{o}, r, \boldsymbol{k})\,]$, the ball $B(\boldsymbol{o}, r)$ is self-accessible relative to C for any $\boldsymbol{\delta}$ satisfying $\delta_i = 1/(1+k_i\theta)$ for each i. Furthermore, $\bar{\theta}(\boldsymbol{o}, r, \boldsymbol{k})$ is continuous in all its arguments.*

The technique for proving Proposition 3 is similar to that of Proposition 1 though neither follows directly from the other. In particular, note that the current proposition cannot handle affine balls - it must work with full-dimensional balls since, for arbitrary $\boldsymbol{k}$, the continuation payoff vector need not be in the original affine ball, no matter how small θ is. It is possible to compute $\bar{\theta}$ as the solution of an NLP (the details have been omitted and are available from the authors on request). Our next proposition starts out with a ball that is self-accessible for the equal discount factor $\underline{\delta}$, and then for any discount factor vector $\boldsymbol{\delta}$ where each $\delta_i \geqslant \underline{\delta}$, proposes a new, ellipsoidal self-accessible set that is contained within the ball.[24]

Proposition 4. *Let $C \subset \mathbb{R}^n$ be finite, and let $X = co(C)$ be full-dimensional and contain in its interior the ball $B(\boldsymbol{o}, r)$ with $r > 0$. Assume that $B(\boldsymbol{o}, r)$ is self-accessible realtive to C for the discount factor vector $\underline{\delta}\boldsymbol{\iota}$ where $\underline{\delta} \in (0,1)$. For any $\boldsymbol{\delta}$ such that $\delta_i \in [\underline{\delta}, 1)$ for $i = 1, \ldots, n$, there exists an ellipsoid $E(\boldsymbol{o}, r, \boldsymbol{\delta}, \underline{\delta}) \subset B(\boldsymbol{o}, r)$ given by center $\boldsymbol{o}$ and semi-axes lengths $\frac{1-\delta_i}{1-\underline{\delta}}\, r$ such that $E(\boldsymbol{o}, r, \boldsymbol{\delta}, \underline{\delta})$ is self-accessible relative to C for $\boldsymbol{\delta}$.*

The proposition implies that for any $\boldsymbol{v} \in E(\boldsymbol{o}, r, \boldsymbol{\delta}, \underline{\delta})$, there is a sequence of pure actions that realizes $\boldsymbol{v}$ when the discount factor vector $\boldsymbol{\delta}$ is used; we let $\{\boldsymbol{a}^{(t)}(\boldsymbol{v}, E(\boldsymbol{o}, r, \boldsymbol{\delta}, \underline{\delta}))\}_{t=0}^{\infty}$ denote that sequence.

[23] Normalizing k_1 to 1, the $\boldsymbol{k}$ vector captures the ratios of discount rates, i.e. *relative* patience among players. With that fixed, one can let θ tend to 0 so as to simultaneously make all players become *absolutely* very patient.

[24] We gratefully acknowledge the fact that this result was first suggested to us by Costas Cavounidis.

5.3 A (Partial) Folk Theorem

We now present a constructive folk theorem for points in $int(F^*)$ under general discounting and without PRDs. The proof of this result uses Proposition 4 and requires only minor re-writing of the proof of Theorem 1. The self-accessibility led approach thus allows us to provide a unified treatment of folk theorems with or without PRDs and with or without symmetric discounting (for points in $int(F^*)$).

Theorem 2. *Let F^* be full-dimensional. For any $\boldsymbol{v} \in int(F^*)$, let $\underline{\delta}$ be the discount factor bound computed in Theorem 1 such that $\boldsymbol{v}$ can be supported as SPNE when players use the discount factor vector $\delta\boldsymbol{\iota}$ with $\delta \geqslant \underline{\delta}$. Then, $\boldsymbol{v}$ is also an SPNE payoff when the discount factor vector $\boldsymbol{\delta} \geqslant \underline{\delta}\boldsymbol{\iota}$ is used. As in Theorem 1, all paths on and off equilibrium are recursively computable and punishment (minmaxing) periods are independent of the discount factors.*

We remark that the unbridled freedom to have players with any configuration of relative patience works only for supporting points inside F^*; it may not work for points outside it, as the analysis in Section 7 will show. The Fact below extends Theorem 2; its justification follows from the proofs of Theorems 1 and 2 as well as the rich continuity properties that discount factor bounds are endowed with in Theorem 1 thanks to self-accessibiliy.

Fact 1. *Let F^* be full-dimensional and let $\boldsymbol{u}$ and $r > 0$ be such that $B(\boldsymbol{u}, r) \in int(F^*)$. Then there exists a uniform $\underline{\delta} > 0$ such that if $\boldsymbol{\delta}$ is such that $\delta_i \geqslant \underline{\delta}$ for each i, every $\boldsymbol{v} \in B(\boldsymbol{u}, r)$ is an SPNE payoff for $\boldsymbol{\delta}$. Hence, for a fixed $\boldsymbol{k}$ vector, there exists a bound $\bar{\theta}$, such that if $\theta \leqslant \bar{\theta}$, and $\delta_i = 1/(1 + k_i\theta)$, every $\boldsymbol{v} \in B(\boldsymbol{u}, r)$ is an SPNE payoff.*

6 Going Beyond the Feasible Set: Realizability

In non-cooperative game theory, contracts are assumed to be unenforceable, which is why we are only interested in equilibrium outcomes. However, if contracts are actually enforceable and the set of supportable payoffs is smaller than the set of realizable payoffs, identifying the latter becomes important.

Using self-accessibility and results from the previous section, we can show that for n-player games with full-dimensional feasible sets, as long as the players are absolutely patient, for a wide variety of relative patience parameters, any point in $int(re(F))$ is realizable. Thus players can simultaneously obtain virtually anything they can individually get in the stage game. All paths involve pure actions only, making PRDs irrelevant.

As an example, consider a two-player game for which F is the shaded region with four vertices $\boldsymbol{c}^1, \ldots, \boldsymbol{c}^6$ as in the left panel of the figure below. Our next theorem implies that points like $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}$ and $\boldsymbol{\rho}$ are all realizable in the repeated game using continuation payoff paths that can be made to (eventually) settle in the vicinity of any point in F we might choose.

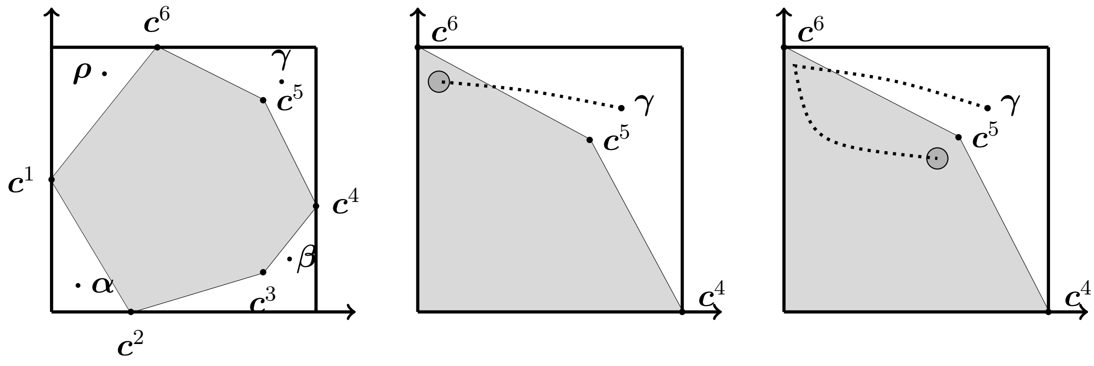

Figure 5: Realizing Payoffs Outside F

Theorem 3. *Let F be full-dimensional. Then*

$$\bigcup_{\boldsymbol{\delta} \in (0,1)^n} \mathcal{F}(\boldsymbol{\delta}) \supset int(re(F)).$$

More precisely, given any target payoff vector $\boldsymbol{v} \in int(re(F))$, any payoff vector $\boldsymbol{u} \in int(F)$ and any $\nu > 0$ such that $B(\boldsymbol{u}, \nu) \subset int(F)$, we can compute positive numbers $k_1, \ldots, k_n$ (not necessarily uniquely) and a $\bar{\theta} > 0$ such that if $\delta_i = \frac{1}{1+k_i\theta}$ with $\theta < \bar{\theta}$ for each i, then, $\boldsymbol{v} \in \mathcal{F}(\boldsymbol{\delta})$ and that along the path that realizes $\boldsymbol{v}$, after a finite sequence of actions, all continuation payoff vectors lie in $B(\boldsymbol{u}, \nu)$.

To obtain some intuition behind this result, suppose we are examining the realizability of a specific point $\boldsymbol{v} \in int(re(F))$. If for some $\boldsymbol{k}$, and low enough θ, playing a certain sequence of vertices brings the continuation payoff (needed to realize $\boldsymbol{v}$) enter a fixed ball $B(\boldsymbol{u}, \nu)$ inside $int(F)$, we are done thanks to our ability to make the ball self-accessible for *any* $\boldsymbol{k}$ if a low enough θ is chosen (see Proposition 3). This implies stringing together a 'pre-entry' finite path into $B(\boldsymbol{u}, \nu)$ and a 'post-entry' infinite path that stays in $B(\boldsymbol{u}, \nu)$, we can generate a path that will realize $\boldsymbol{v}$.

In the middle panel of Figure 5, which enlarges the top-right section of the left panel, for the target payoff $\boldsymbol{\gamma}$, we show an example pre-entry path using a dotted line stretching from $\boldsymbol{\gamma}$ to the small ball just below vertex $\boldsymbol{c}^6$. This path results from the repeated play of $\boldsymbol{c}^4$ until entry inside the ball takes place. The key insight here is to choose the number of periods of playing $\boldsymbol{c}^4$ to be (roughly) inversely proportional to θ; whereupon adjusting the constant of proportionality and the k_i's, the distance traveled by the continuation payoff in both directions can be exactly 'regulated' in the limit (as θ becomes small). For the dashed path in the figure, since player 2's continuation payoff changes much less than that of player 1, k_2 will have to be chosen much smaller than k_1.

Many such balls and many $\boldsymbol{k}$ vectors are compatible with this plan of entry into F, but one can achieve more. In the right panel of Figure 5, we want the continuation payoff to eventually enter the specific ball just below the vertex $\boldsymbol{c}^5$ (centered around some specific $\boldsymbol{u}$). We show that there are many $\boldsymbol{k}$ vectors for which this can be done (for θ small) by playing $\boldsymbol{c}^4$ several times followed by playing $\boldsymbol{c}^6$ several times. The dashed path in the panel describes the trajectory of the continuation payoffs, the kink signifying the point where play switches from $\boldsymbol{c}^4$ to $\boldsymbol{c}^6$. An induction argument generalizes this to any number of players and any ball inside F and we show that there is an n-phase pre-entry path involving a sequence of vertices $c^1, \ldots, c^n$ which can be chosen using a simple rule: $sgn(c_i^i - v_i) = sgn(v_i - u_i)$ for $i = 1, \ldots, n$.

For $n = 2$, Theorem 3 can be sharpened via the next result, which fully characterizes realizable payoffs. Such a full characterization is not possible for $n \geqslant 3$ without additional requirements on the faces of the feasible set.

Theorem 4. *For two-player games, if F is full-dimensional,*

$$\bigcup_{\boldsymbol{\delta} \in (0,1)^2} \mathcal{F}(\boldsymbol{\delta}) = int(re(F)) \bigcup F.$$

Our approach to the question realizability is significantly different from that of LP 1999 and the paths we offer are also very different from the paths that were proposed in LP's paper. For the interested reader, a detailed discussion is provided in Appendix 2.

7 Going Beyond the FSIR Set: Supportability

The path described in the proof of Theorem 3, unfortunately does not guarantee that every continuation payoff on that path satisfies individual rationality, even when the target payoff $\boldsymbol{v}$ is strictly individually rational and the eventual continuation payoff is some point in F^*. We need additional conditions on $\boldsymbol{v}$.

A permutation π is a 1-1 correspondence between I and itself; by π_i we mean $\pi(i)$. We represent π by simply stating the vector $(\pi_1, \ldots, \pi_n)$. The permutation $(1, \ldots, n)$ is called the 'natural permutation' or 'natural order'. An interpretation of these permutations is now suggested: π simply maps the ranks of discount factors into player 'names'; thus, in a 5-player game, if $\pi_2 = 4$, that means player 4's discount factor is the second-lowest. The inverse function maps names to the ranks; thus if $\pi^{-1}(j) \leqslant \pi^{-1}(i)$, we understand that player i is at least as patient as player j. This interpretation will be useful to keep in mind for the two definitions to follow; the proofs of Theorem 5 and 6 will validate it later.

7.1 The Diagonal Conditions

In the definitions below C is any set of points in $\mathbb{R}^n$, $F = co(C)$ and $\boldsymbol{w}$ is some point in $re(C)$. Of course, these objects have their standard interpretations in the context of a game: pure action payoffs, feasible set and minmax point.

Definition A payoff vector $\boldsymbol{v} \in re(C)$ is said to satisfy the weak diagonal condition (WD) if there exists a permutation π, such that $\forall\, i$, $\exists$ a vector $\boldsymbol{u}^i \in F$ with the property

$$u^i_{\pi_i} = v_{\pi_i}$$

and

$$u^i_j \geqslant w_j \quad \text{if } \pi^{-1}(j) \leqslant \pi^{-1}(i).$$

Definition A payoff vector $\boldsymbol{v} \in int(re(C))$ is said to satisfy the strict diagonal condition (SD) if there exists a permutation π, such that $\forall\, i$, $\exists$ a vector $\boldsymbol{u}^i \in int(F)$ with the property

$$u^i_{\pi_i} = v_{\pi_i}$$

and

$$u^i_j > w_j \quad \text{if } \pi^{-1}(j) \leqslant \pi^{-1}(i).$$

When $\boldsymbol{v}$ satisfies the first (second) definition we say that it satisfies WD (SD) in the order π. Also, the set of all points satisfying WD (SD) for a given π will be denoted as $W(\pi)(S(\pi))$. Note that full dimensionality of F^* is needed for $S(\pi)$ to be non-empty, but not so for $W(\pi)$. As for the relation between these sets, it is easy to see that for any π, $cl(S(\pi)) \subset W(\pi)$ and assuming full-dimensionality of F^*, $int(W(\pi)) = S(\pi)$.

In a 2-player game, $\boldsymbol{v}$ satisfies WD in some order if it is weakly individual rational and if, from that point, we can draw a line parallel to one of the axes and make that line intersect F^+. This implies that the the darkly shaded region (including its boundary) in Figure 2 in Section 3.2 is the set of points that satisfy WD for some permutation. Specifically, for example, the point $\boldsymbol{v} = (8, 5)$ satisfies WD in the natural order because $\boldsymbol{u}^1$ can be chosen to be $(8, 1)$ and $\boldsymbol{u}^2$ can be chosen to be $(5, 5)$. For games with 3 or more players, the following Fact, stated without proof, is helpful in discerning which points will satisfy WD or SD (for some order).

Fact 2. *$\boldsymbol{v}$ satisfies WD for permutation π iff $v_i \in [\,\underline{v}_i(\pi), \bar{v}_i(\pi)]$ $\forall i = 1, \ldots, n$, where for any given permutation π, $\underline{\boldsymbol{v}}(\pi)$ and $\bar{\boldsymbol{v}}(\pi)$ are defined as follows:*

$$\underline{v}_i(\pi) := \min\{v_i : \boldsymbol{v} \in F \textit{ such that } v_j \geqslant w_j \;\; \forall j \textit{ such that } \pi^{-1}(j) \leqslant \pi^{-1}(i)\}.$$

$$\bar{v}_i(\pi) := \max\{v_i : \boldsymbol{v} \in F \textit{ such that } v_j \geqslant w_j \;\; \forall j \textit{ such that } \pi^{-1}(j) \leqslant \pi^{-1}(i)\}.$$

The Fact makes is clear that for any π, computing $\underline{v}_i(\pi)$ and $\bar{v}_i(\pi)$ and hence describing $W(\pi)$ is simply a matter of solving a set of linear programs. For symmetric 3-player games, the computations might be even more straightforward as the example below illustrates.

Example (The "stand-out" game): Each of three players can play one of two actions: R (right) or W (wrong). All players get 0 if 0, 2 or 3 players play R. Only if exactly one player plays R, that player gets $1 + 2\eta$ while each of the other two players receive

$-\eta$ where η is some positive number. Thus, $C = \{(0,0,0),(1+2\eta,-\eta,-\eta),(-\eta,1+2\eta,-\eta),(-\eta,-\eta,1+2\eta)\}$. In this game the only way to obtain the superior payoff of $1+2\eta$ is to uniquely 'stand out' (by doing the right thing) whereupon the other players are 'shamed' into getting a negative payoff of $-\eta$. Note that for each player, R weakly dominates W. However, if all players play this weakly dominant strategy the total utility is 0, while each of the three action vectors $(R,W,W),(W,R,W)$ and (W,W,R) gives the players a total utility of 1. Note also that each player can be minmaxed by the other two players playing R, and hence $\boldsymbol{w} = (0,0,0)$.

Since $\boldsymbol{w}$ is in the feasible set, clearly $\underline{v}_i(\pi) = 0\ \forall i$ and $\forall \pi$. Now notice that the Pareto frontier in the feasible set is $co(\{(1+2\eta,-\eta,-\eta),(-\eta,1+2\eta,-\eta),(-\eta,-\eta,1+2\eta)\})$ and thus, all points on this frontier belong to the same plane as that of the unit simplex. Hence, if we were to maximize one player's payoff in F while giving the other two a non-negative payoff, we can give him at most 1. If we were to maximize his payoff while giving one of the others at least 0 and putting no other restriction on the third, apart from requiring that we stay inside F, we can give him at most $1+\eta$. Finally, if we tried to maximize this player's payoff without any restriction on the other two's payoffs apart from requiring that we stay in F, we can give him at most $1+2\eta$. Hence, for the natural permutation π, $\bar{\boldsymbol{v}}(\pi) = (1+2\eta, 1+\eta, 1)$ and $W(\pi) = [0, 1+2\eta] \times [0, 1+\eta] \times [0,1]$.

7.2 Necessity and Weak Diagonalizability

Theorem 5. *Consider an n-player game and let $\mathcal{P}$ denote the set of all permutations of $1,\ldots,n$. Then,*

$$\bigcup_{\pi\in\mathcal{P}} W(\pi) \supset \bigcup_{\boldsymbol{\delta}\in(0,1)^n} \mathcal{V}(\boldsymbol{\delta})$$

i.e., for any n-player game if $\boldsymbol{v}$ is a SPNE payoff vector for some discount factor vector $\boldsymbol{\delta}$, then $\boldsymbol{v} \in W(\pi)$ for some permutation π.

Note that no full-dimensionality assumption is required in the above theorem. Its proof proceeds by asking the counterfactual question: What would the payoff profile look like if the players played the proposed equilibrium path but all of them used one of the player's discount factor to evaluate their normalized payoffs? We show that an increase in a player's discount factor can never let his (normalized) payoff to fall below his individual rational level, from which it follows that the $\boldsymbol{u}$'s needed for the WD property are exactly these freshly evaluated payoff vectors.

We now wish to comment on two insights provided by this theorem, one that shows how equilibrium characterization for multiplayer games bears similarity to that of two-player games and one that shows how it does not.

The proof of Theorem 5 along with Fact 2 rules out arbitrary associations of equilibrium payoffs with relative ordering of patience among the players. For example, going back to the stand-out game, for $\eta = .5$, although the payoff vector $(1.9, 1.4, .9)$ is an equilibrium payoff vector, it is conformable with one and only one ranking of player patience: the one given by the natural order. This suggests that for symmetric stage games the maximum payoff the least patient player can conceivably receive is the largest possible payoff any player can receive in the repeated game. It is noteworthy that in the presence of PRDs and for the two-player case, LP 1999 arrive at the same conclusion.

Theorem 5 and Fact 2 also inform us that in the stand-out game, any payoff vector that does not lie in the union of the six cubes of the form $[0, a] \times [0, b] \times [0, c]$ where a, b, c are some arrangements of the 3 numbers $1, 1 + \eta$ and $1 + 2\eta$, can never be an equilibrium payoff. Thus, again if $\eta = .5$, we can conclude that the payoff vector $(.5, 1.7, 1.9)$ cannot be an equilibrium payoff vector because two of its coordinates lie above 1.5 - it does not matter what the other coordinate is (although, this payoff vector is realizable by virtue of Theorem 3). This is illustrated in Figure 6 below which depicts the set of all weakly diagonalizable payoffs for this game. Clearly, for three-player games, it is *not* true that in analogy with the two-player game in Section 3.2, the only individually rational payoffs that are excluded from being equilibrium payoffs form an 'upper north-east' cuboid.

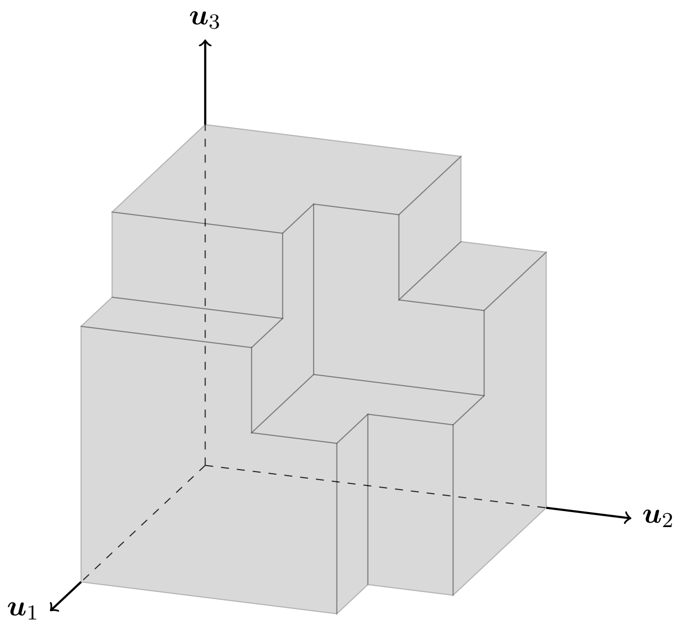

Figure 6: Weakly Diagonalizable Payoffs for the Stand-out Game

7.3 Sufficiency and Strict Diagonalizability

Weak diagonalizability, however, cannot be a sufficient condition for a payoff vector to be an equilibrium payoff vector, even when we have full dimensionality of the FSIR set. One can see this for the game in Counterexample 1, where the payoff vector $(1, 1)$ belongs to $W(\pi)$ for each π, but it is not even realizable for any discount factor vector. Strict diagonalizability, though, 'works'.

Theorem 6. *Consider an n-player game with a full-dimensional FSIR set. Let $\mathcal{P}$ denote the set of all permutations of $1, \ldots, n$. Then,*

$$\bigcup_{\pi \in \mathcal{P}} S(\pi) \subset \bigcup_{\boldsymbol{\delta} \in (0,1)^n} \mathcal{V}(\boldsymbol{\delta})$$

More specifically, if $\boldsymbol{v} \in S(\pi)$ for some π, it is possible to determine $\boldsymbol{k} >> \mathbf{0}$ (not necessarily uniquely), and $\bar{\theta} > 0$ such that if $\theta \in (0, \bar{\theta})$ and $\delta_i = \frac{1}{1+k_i\theta}$ for each i, it is possible to specify a SPNE strategy profile that supports $\boldsymbol{v}$ for those discount factors.

Since the proof of Theorem 6 uses several new ideas, we now provide some intuition behind our arguments.

We will borrow a technique from the proof of Theorem 3: Devise relative patience parameters ($\boldsymbol{k}$ vectors) and an action sequence that will generate a path for the continuation payoffs which will start with the target payoff and for θ small, will end in a ball inside the interior of the FSIR set (rather than the feasible set as in the case of Theorem 3). As before, we will call this the pre-entry path. Once our path enters that ball, Fact 1 takes over, since that result guarantees that any point in such a ball is an equilibrium payoff provided each player's discount factor exceeds a certain bound. There are two new challenging tasks here: first to create a pre-entry path that stays strictly individual rational throughout and second, to design an equilibrium strategy that takes care of incentive compatibilty along the pre-entry path.

Where strict diagonalizability helps is with the first task. To see this for the two-player case, suppose our target payoff $\boldsymbol{v}$ is in $S(\pi)$ where π is the natural order. Then, strict diagonalizability of $\boldsymbol{v}$ guarantees us the existence of a vector $\boldsymbol{u}^2$ in $int(F^*)$ where $u_2^2 = v_2$ and u_1^2 (strictly) exceeds player 1's minmax payoff. Also, there must exist a vertex $\boldsymbol{c}$ such that each time this vertex is played the continuation payoff of player 1 shifts towards u_1^2 (from v_1). If we play this vertex T times where T is $\lceil \frac{b}{\theta} \rceil$), letting $k_1 = 1$ and by choosing a suitable $b > 0$ and θ small, limiting arguments used in the proof of Theorem 3, show us that it is possible to take the first coordinate of the continuation payoff arbitrarily close to u_1^2. On top of this if k_2 is chosen small, the second coordinate of the continuation payoff barely changes from v_2 during these T periods and the continuation payoff vector after T periods gets close to $\boldsymbol{u}^2$ and hence, enters the desired ball.

For three or more players, unfortunately, extending this idea runs into difficulties. Suppose $n = 3$ and $\boldsymbol{v} \in S(\pi)$ for the natural order π. Now, there is a point in $int(F^*)$, namely $\boldsymbol{u}^3$, the third element of which is v_3. A small neighborhood of this point suggests itself as the location of entry into $int(F^*)$. This is what we might want to do to accomplish our goal: first we 'fix' player 1's payoff by moving it to u_1^3 (via the play of some suitable vertex). If we keep players 2 and 3 patient *relative to player 1*, their payoffs do not fall below their minmax values during this first phase. Next, we try and fix player 2's payoff by taking it close to u_2^3 playing some other suitable vertex and thereby try and enter $int(F^*)$. The problem with this strategy is that in the second phase, we can ensure player

3's continuation payoff to stay above his minmax levels by keeping him relatively patient vis-a'-vis player 2, but we can't do the same for player 1 (because by design, player 1 is now *impatient* relative to player 2). Consequently, when we are trying to 'fix' player 2's payoff, player 1 can find his payoff getting 'unfixed'. We could try and fix both player 1 and 2's payoffs simultaneously following the method used in the proof of Theorem 3, but then, we would not have any guarantee that the path will maintain individual rationality throughout.

We, therefore, need a new tool: how to move one strictly individually rational payoff vector to another's neighborhood by playing a sequence of pure actions without violating individual rationality along the path. This can indeed be done for sufficiently patient players, with *any* relative patience configuration as long as the two payoff vectors are both inside $int(F^*)$; moreover, the length of the sequence can be bounded in terms of θ.

In m-dimensional Euclidean space, let $[\boldsymbol{y}, \boldsymbol{z}]$ represent the line segment joining the points $\boldsymbol{y}$ and $\boldsymbol{z}$ and let $\mathcal{C}(\boldsymbol{y}, \boldsymbol{z}, r)$ denote the set of points that are at most distance r from $[\boldsymbol{y}, \boldsymbol{z}]$, i.e.

$$\mathcal{C}(\boldsymbol{y}, \boldsymbol{z}, r) = \bigcup_{\boldsymbol{o} \in [\boldsymbol{y}, \boldsymbol{z}]} B(\boldsymbol{o}, r)$$

From now on, we will refer to sets of these types as 'capsules'.

Proposition 5. (*The Capsule Lemma*) *Let D be a finite set in m-dimensional Euclidean space. For $\boldsymbol{y}, \boldsymbol{z}$ in $int(co(D))$, let $r > 0$ be such that the capsule $\mathcal{C}(\boldsymbol{y}, \boldsymbol{z}, r)$ is in the interior of $co(D) \cap \mathbb{R}^m_{++}$, which is assumed to be m-dimensional. Then,*

a) For any $\boldsymbol{u} \in B(\boldsymbol{y}, r)$ and any vector $\boldsymbol{k} \in \mathbb{R}^m_{++}$, there exists $\theta^\dagger(\boldsymbol{k}) > 0$ such that for any $\theta \in (0, \theta^\dagger(\boldsymbol{k}))$ we can find a finite sequence of points $\{\boldsymbol{c}^t\}_{t=0}^{T-1}$ in D and a finite sequence of points $\{\boldsymbol{x}^t\}_{t=0}^{T}$ in $\mathcal{C}(\boldsymbol{y}, \boldsymbol{z}, r)$ such that
i) $\boldsymbol{x}^0 = \boldsymbol{u}$,
ii) $\boldsymbol{x}_i^{t+1} - \boldsymbol{c}_i^t = (1 + k_i\theta)(\boldsymbol{x}_i^t - \boldsymbol{c}_i^t) \quad$ for $i = 1, \ldots, m$, $t = 0, \ldots, T-1$ and
iii) $\boldsymbol{x}^T \in B(\boldsymbol{z}, r)$.

b) Furthermore, there exists a $\theta^\ddagger(\boldsymbol{k}) \leqslant \theta^\dagger(\boldsymbol{k})$, a strictly positive number m_1 and a strictly negative number m_2 (both dependent on $\boldsymbol{k}$) such that when $\theta \in (0, \theta^\ddagger)$, the T given in part a) is less than or equal to $\left\lceil \frac{||y-z||}{r - \sqrt{m_1\theta^2 + m_2\theta + r^2}} \right\rceil$.

The Capsule lemma's proof offers the remarkable insight that the notion of self-accessibility is not just useful for keeping continuation payoffs tethered to a point; it can also be used to take them for a 'walk' inside $int(F^*)$. Along with strict diagonalizability it can be used to design a $\boldsymbol{k}$, a bound on θ and a (continuation payoff) path that starts at the target payoff $\boldsymbol{v}$ and ends inside $int(F^*)$, close to $\boldsymbol{u}^n$. We sketch how for the $n = 3$ case.

While $\boldsymbol{u}^3$ is the (approximate) point of entry, our construction will also make use of $\boldsymbol{u}^2$ given by SD. First, let us replace the first coordinate of $\boldsymbol{v}$ by the first coordinate of $\boldsymbol{u}^2$ and create the vector $\boldsymbol{z} = (u_1^2, v_2, v_3) \in R^3_{++}$ (thus $\boldsymbol{z}$ and $\boldsymbol{u}^2$ have the same first two coordinates). Next, we design the pre-entry path of the continuation payoff by splitting it into two phases. In the first phase the path stays very close to the line segment joining $\boldsymbol{v}$ and $\boldsymbol{z}$ while in the second phase it stays very close to the line segment joining $\boldsymbol{z}$ and $\boldsymbol{u}^3$. Since all three vectors $\boldsymbol{v}$, $\boldsymbol{z}$ and $\boldsymbol{u}^3$ belong to $\mathbb{R}^3_{++}$, this will guarantee individual rationality along the entire path.

For the first phase we choose $k_1 = 1$ and a vertex such that playing that vertex $T_1(\theta)$ times makes player 1's continuation payoff moves from v_1 to close to u_1^2 for θ small. As has been argued before, this can be done by making $T_1(\theta)$ inversely proportional to θ and choosing a suitable constant of proportionality. Next, we choose k_2 small enough, so that during the first phase, player 2's continuation payoff does not change much from v_2. Hence, by the end of phase 1, the continuation payoff (sub)-vector of players 1 and 2 will be (for θ small) very close to (z_1, z_2).

Next, notice that the projection of any vector in F on the first two coordinates belongs to the convex hull of the projection of the extreme points of F on the same two coordinates; further, if the original vector is strictly positive, so is the projected subvector. Hence, we can use the Capsule lemma (in two dimension) to create a sequence of actions that takes the continuation payoff (subvector) from a neighborhood of (z_1, z_2) to one of (u_1^3, u_2^3). If $T_2(\theta)$ is the number of steps required in the process, now we can determine k_3 small enough such that during none of the $T_1(\theta) + T_2(\theta)$ combined steps player 3's payoff changes much from $u_3^3 = v_3$ for θ small (this is possible because $T_2(\theta)$ is also of the order of $\frac{1}{\theta}$). Thus, at the end of both phases the continuation payoff gets very close to $\boldsymbol{u}^3$ without violating individual rationality in any period (for the chosen $\boldsymbol{k}$ and θ small).

Now that the equilibrium path has been described, the equilibrium strategy can be qualitatively described; of course, only the equilibrium prescription along the pre-entry path needs to be spelled out. Play starts and continues on the pre-entry path until one of the players unilaterally deviates whereupon he is minmaxed by the other players for a certain number of periods. *After minmaxing, play 'returns' to the same point on the pre-entry path where the deviation took place* (this is why we referred to these strategies in the introduction as 'turnpike strategies'). At the completion of the pre-entry path, players play an action sequence corresponding to a point which shifts the target post-entry equilibrium continuation payoff by two 'adjustment' terms; one to make any punishing player indifferent among his minmaxing pure strategies, and a second to 'reward' the punishers for participating in the punishment. During pre-entry, any deviation from a punishment phase or a new deviation after play has returned to the 'pre-entry' path is treated as if a fresh deviation just took place.

Why does this "stick now (for the deviant), carrot later (for the punishers)" strategy ensure incentive compatibility for patient players? With rising patience, bad outcomes for a finite number of periods still can wipe out one period gains (if nothing else changes). That is how the 'stick now' threat keeps players from deviating. The prospective punishers may also suffer during the punishment periods but they are compensated by a reward coming later. As they become more patient, that future reward overrides the sacrifices they make for a finite number of periods but counteracting this, the time taken to to get one's reward also grows infinitely large. Fortunately, the second part of the Capsule Lemma bounds the time spent on the pre-entry path in a way that the PDV of the reward tends to a limit. This allows us to incentivize patient players suitably to participate in the punishment of the deviant.

Before leaving this section, it is worthwhile mentioning that together, the proofs of Theorems 5 and 6 also help us establish an assertion we made in the introduction: for stage games where F^+ is full-dimensional, the set of sequentially individually rational payoffs has zero Hausdorff distance from the set of equilibrium payoffs, or in notation used there: $\bigcup_{\boldsymbol{\delta}\in(0,1)^n} \mathcal{F}^+(\boldsymbol{\delta}) \approx \bigcup_{\boldsymbol{\delta}\in(0,1)^n} \mathcal{V}(\boldsymbol{\delta})$. This is true because in the proof of Theorem 5, the only property of an equilibrium payoff that is used is sequential individual rationality. Hence, from that theorem we have $\bigcup_{\pi\in\mathcal{P}} W(\pi) \supset \bigcup_{\boldsymbol{\delta}\in(0,1)^n} \mathcal{F}^+(\boldsymbol{\delta})$ while from Theorem 6 we have $\bigcup_{\boldsymbol{\delta}\in(0,1)^n} \mathcal{V}(\boldsymbol{\delta}) \supset \bigcup_{\pi\in\mathcal{P}} S(\pi)$. Since, $\bigcup_{\boldsymbol{\delta}\in(0,1)^n} \mathcal{F}^+(\boldsymbol{\delta}) \supset \bigcup_{\boldsymbol{\delta}\in(0,1)^n} \mathcal{V}(\boldsymbol{\delta})$, and since for every π, $int(W(\pi)) = S(\pi)$ which implies $\bigcup_{\pi\in\mathcal{P}} W(\pi) \approx \bigcup_{\pi\in\mathcal{P}} S(\pi)$, the relation follows. This assertion is comparable to Theorem 1 of Sugaya (2015) which essentially makes the same claim, but having restricted relative patience parameters.

8 Conclusion

This paper provides a unified treatment of discounted repeated games with perfect monitoring and without PRDs. The scope of our inquiry follows a logical chain, successively allowing for wider target payoff sets and less restrictive discounting structures. The glue that holds all the results together is the simple geometrical notion: self-accessibility. The analysis culminates in Theorem 6, where we show that any point $\boldsymbol{v}$ satisfying the Strict Diagonal Condition is an SPNE payoff for some possibly asymmetric discounting profile. This easy to check condition translates into the following: there exists an ordering π of the players such that for player $\pi(i)$ there is a payoff vector in the interior of the feasible set at which $\pi(i)$ gets the payoff $v_{\pi(i)}$ and everyone before $\pi(i)$ in the ordering gets more than their respective minmax values. Our result can be viewed as a new folk theorem for repeated games with unrestricted discounting patterns that is built on fully constructive foundations.

9 Appendix 1: All Proofs

Proof of Proposition 1

Existence of $\underline{\delta}$: We assume wlog that the center of the ball is the origin. Also wlog, we assume that $C' = \{\boldsymbol{c}^1, \ldots \boldsymbol{c}^K\}$ is in fact, the set of extreme points of X.

Fix $\boldsymbol{c} \in C'$, $\boldsymbol{x} \in S$ and let $\delta(\boldsymbol{x}, \boldsymbol{c})$ be defined as the solution of the following problem[25]:

$$\text{Min } \delta \in [0,1] \text{ subject to } \boldsymbol{x} = (1-\delta)\boldsymbol{c} + \delta\, \boldsymbol{y}, \text{ for some } \boldsymbol{y} \in S.$$

We now characterize $\delta(\boldsymbol{x}, \boldsymbol{c})$. Note that if $\boldsymbol{y}$ satisfies the equation $\boldsymbol{x} = (1 - \delta(\boldsymbol{x}, \boldsymbol{c}))\boldsymbol{c} + \delta(\boldsymbol{x}, \boldsymbol{c})\, \boldsymbol{y}$, then $\boldsymbol{y}$ must be at the boundary of the ball. Hence, it must be that $\boldsymbol{y}.\boldsymbol{y} = r^2$ which implies $\boldsymbol{x} - (1 - \delta(\boldsymbol{x}, \boldsymbol{c}))\boldsymbol{c}$'s dot product with itself is $\delta^2 r^2$. Upon rearranging this shows that $\delta(\boldsymbol{x}, \boldsymbol{c})$ must be a root of the following quadratic equation in δ :

$$\delta^2(\boldsymbol{c}.\boldsymbol{c} - r^2) + 2\delta\boldsymbol{c}.(\boldsymbol{x} - c) + (\boldsymbol{x} - \boldsymbol{c}).(\boldsymbol{x} - \boldsymbol{c}) = 0. \tag{9.1}$$

For any $\boldsymbol{c}$, given that the left hand side of (9.1) is a *convex* quadratic (since $\boldsymbol{c}.\boldsymbol{c} - r^2 > 0$), with a strictly positive value at 0 (since $(\boldsymbol{x} - \boldsymbol{c}).(\boldsymbol{x} - \boldsymbol{c}) > 0$) and a non-positive value at 1 (since $\boldsymbol{x}.\boldsymbol{x} - r^2 \leqslant 0$), there are two roots: one is greater than 0 and less than or equal to 1 while the other is greater than or equal to 1. We are seeking the smaller root, which is continuous in $\boldsymbol{x}$, making $\delta(\boldsymbol{x}, \boldsymbol{c})$ continuous in $\boldsymbol{x}$.

Furthermore, we assert that for one of the $\boldsymbol{c}$'s, the smaller root must be *strictly* less than 1. Clearly this will be true if the quadratic at 1 is strictly negative, i.e. $\boldsymbol{x}.\boldsymbol{x} - r^2 < 0$. So assume that $\boldsymbol{x}.\boldsymbol{x} = r^2$. It now suffices to show that the slope of the quadratic at 1 is strictly positive for some $\boldsymbol{c}$, which will be true if for at least one l, $\boldsymbol{c}^l.\boldsymbol{x} > r^2$. If not, then

$$\boldsymbol{c}^l.\boldsymbol{x} \leqslant r^2 \quad \text{for } \ l = 1, \ldots, K \tag{9.2}$$

If that is the case, we claim that each of these inequalities must actually be an equality. To see this note that since $\boldsymbol{x}$ is in the relative interior of X, $\boldsymbol{x} = \sum_{l=1}^K \lambda^l\, \boldsymbol{c}^l$ where λ^l's are strictly positive weights summing to 1. Multiplying each inequality in (9.2) by λ^l and summing over l, on the left hand side we will have $\left(\sum_{l=1}^K \lambda^l \boldsymbol{c}^l\right).\boldsymbol{x} = \boldsymbol{x}.\boldsymbol{x}$ while on the right hand side we will have $\left(\sum_{l=1}^K \lambda^l\right) r^2 = r^2$ and since these two are equal, the claim follows. But now, if for each l, $\boldsymbol{c}^l.\boldsymbol{x} = r^2$, since the center of the ball (the origin) can also be expressed as a convex cobination of the vertices, i.e. $\boldsymbol{o} = \sum \theta^l \boldsymbol{c}^l$ for a set of weights θ^l summing to one, this will imply $\boldsymbol{o}.\boldsymbol{x} = 0 = r^2$, a contradiction. Hence, for every $\boldsymbol{x} \in B_X(\mathbf{0}, r)$, there exists a vertex $\boldsymbol{c}$ such that the quadratic in (9.1) has a strictly positive slope at 1, and hence for that vertex, $\delta(\boldsymbol{x}, \boldsymbol{c}) \in (0, 1)$.

Let $\delta^*(\boldsymbol{x}) := \min\{\delta(\boldsymbol{x}, \boldsymbol{c}) : \boldsymbol{c} \in C'\}$, with the minimum attained at $c^*(\boldsymbol{x})$.[26]Clearly, δ^* is continuous, being the minimum of continuous functions and lies in $(0, 1)$. Finally, define $\underline{\delta} := \max\{\delta^*(\boldsymbol{x}) : \boldsymbol{x} \in S\}$. Since S is compact, this maximum is attained at some $\boldsymbol{x}^*$. Since for any $\boldsymbol{x}$, $\delta^*(\boldsymbol{x}) \in (0, 1)$ we must have $\underline{\delta} \in (0, 1)$. It is now easily verifiable that for this $\underline{\delta}$ and any common discount factor above this value S is self-accessible relative to C'.

[25] For the first part of the proof, to keep the notation simple, when a function depends on the location of the ball, we will drop the center and the radius as arguments (thus, for example $\delta(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, r)$ will be simply referred to as $\delta(\boldsymbol{x}, \boldsymbol{c})$).

[26] To achieve well-definition, in case of ties, use any arbitrary preference ordering among the vertices.

Computability of $\underline{\delta}$: Having described the problem of determining $\underline{\delta}(\boldsymbol{o}, r)$ as

$$\underset{\boldsymbol{x} \in B_X(\boldsymbol{o}, r)}{\text{Max}} \quad \underset{\boldsymbol{c} \in C'}{\text{Min}} \quad \delta(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, r)$$

we now show that this two-stage nested optimization problem can be written as one large optimization problem. This is seen from noting the fact that the 'inside' problem can be stated as a maximization, rather than a minimization problem as shown:

$$\begin{aligned}
&\text{Max} \quad \delta \\
&\textit{subject to} \\
&\delta \leqslant \delta(\boldsymbol{x}, \boldsymbol{c}^l, \boldsymbol{o}, r) \quad \forall \boldsymbol{c}^l \in C'.
\end{aligned}$$

For affine balls $B_X(\boldsymbol{o}, r)$ in the relative interior of X, an explicit formula for $\delta(\boldsymbol{x}, \boldsymbol{c}^l, \boldsymbol{o}, r)$ exists in terms of the smaller root of equation (9.1); a numerically simpler way to characterize that root is to just require that the slope of the quadratic is less than or equal to zero, in addition to stating that the quadratic vanishes. For any given $\boldsymbol{o}$ (not necessarily the origin), this leads to the NLP below the solution of which gives us $\underline{\delta}(\boldsymbol{o}, r)$:

$$\begin{aligned}
&\text{Max} \quad \delta \\
&\textit{subject to} \\
&\delta \leqslant \delta^l \quad \forall\, l = 1, \ldots, K \qquad (1) \\
&(\delta^l)^2\{(\boldsymbol{c}^l - \boldsymbol{o}).(\boldsymbol{c}^l - \boldsymbol{o}) - r^2\} + 2\delta^l(\boldsymbol{c}^l - \boldsymbol{o}).(\boldsymbol{x} - \boldsymbol{c}^l) + \|\boldsymbol{x} - \boldsymbol{c}^l\|^2 = 0 \quad \forall\, l = 1, \ldots, K \qquad (2) \\
&(\delta^l)\{(\boldsymbol{c}^l - \boldsymbol{o}).(\boldsymbol{c}^l - \boldsymbol{o}) - r^2\} + (\boldsymbol{c}^l - \boldsymbol{o}).(\boldsymbol{x} - \boldsymbol{c}^l) \leqslant 0 \quad \forall\, l = 1, \ldots, K \qquad (3) \\
&(\boldsymbol{x} - \boldsymbol{o}).(\boldsymbol{x} - \boldsymbol{o}) \leqslant r^2 \qquad (4) \\
&\boldsymbol{x} = \textstyle\sum_{l=1}^K \lambda^l \boldsymbol{c}^l \qquad (5) \\
&\textstyle\sum_{l=1}^K \lambda^l = 1 \qquad (6) \\
&\lambda^l \geqslant 0 \quad \forall\, l = 1, \ldots, K \qquad (7).
\end{aligned}$$

In the NLP, constraints (2) and (3) characterize the $\delta(\boldsymbol{x}, \boldsymbol{c}^i, \boldsymbol{o}, r)$'s (for $i = 1, \ldots, K$) while constraint (1) finds the minimum of these. Of course, the maximization is also over $\boldsymbol{x}$ and constraints (4) - (7) ensure that each feasible $\boldsymbol{x}$ belongs to the affine ball $B_X(\boldsymbol{o}, r)$.[27] ■

Proof of Proposition 2

The function $\underline{\delta}(\boldsymbol{o}, r)$ defining a discount factor bound that makes $B_X(\boldsymbol{o}, r)$ self-accessible is continuous in both its arguments - as follows from the continuity of $\delta(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, r)$ in $(\boldsymbol{o}, r)$, the proof of Proposition 1 and a straightforward application of the Maximum Theorem. The set $B_X(\boldsymbol{x}, \bar{r})$ being compact (wherein the centers of the smaller balls must lie), the existence of the required

[27]Note that if the ball was full-dimensional, we could have dropped constraints (5) - (7).

uniform bound follows. Now we can state the required NLP.

$$\begin{aligned}
&\text{Max } \ \delta \\
&\textit{subject to} \\
&\delta \leqslant \delta^l \quad \forall\, i = 1, \ldots, K \qquad (1) \\
&(\delta^l)^2\{(\boldsymbol{c}^l - \boldsymbol{x}')(\boldsymbol{c}^l - \boldsymbol{x}') - \omega^2\} + 2\delta^l(\boldsymbol{c}^l - \boldsymbol{x}')(\boldsymbol{x}'' - \boldsymbol{c}^l) + \|\boldsymbol{x}'' - \boldsymbol{c}^l\|^2 = 0 \quad \forall\, l = 1, \ldots K \qquad (2) \\
&(\delta^l)\{(\boldsymbol{c}^l - \boldsymbol{x}')(\boldsymbol{c}^l - \boldsymbol{x}') - \omega^2\} + (\boldsymbol{c}^l - \boldsymbol{x}')(\boldsymbol{x}'' - \boldsymbol{c}^l) \leqslant 0 \quad \forall\, l = 1, \ldots, K \qquad (3) \\
&(\boldsymbol{x} - \boldsymbol{x}').(\boldsymbol{x} - \boldsymbol{x}') \leqslant \bar{r}^2 \qquad (4) \\
&(\boldsymbol{x}' - \boldsymbol{x}'').(\boldsymbol{x}' - \boldsymbol{x}'') \leqslant \omega^2 \qquad (5) \\
&\boldsymbol{x}' = \textstyle\sum_{l=1}^K \lambda^l \boldsymbol{c}^l \qquad (6) \\
&\boldsymbol{x}'' = \textstyle\sum_{l=1}^K \theta^l \boldsymbol{c}^l \qquad (7) \\
&\textstyle\sum_{l=1}^K \lambda^l = 1 \qquad (8) \\
&\textstyle\sum_{l=1}^K \theta^l = 1 \qquad (9) \\
&\lambda^l, \theta^l \geqslant 0 \quad \forall\, l = 1, \ldots, K \qquad (10)
\end{aligned}$$

The construction of the given NLP follows directly from the proof of Proposition 1 and the observation that the problem of finding a uniform bound as $\boldsymbol{x}'$ ranges in the set $B_X(\boldsymbol{x}, \bar{r})$ adds another layer of nesting to the optimization problem in that proposition, making the current problem representable as

$$\underset{\boldsymbol{x}' \in B_X(\boldsymbol{x},\bar{r})}{\text{Max}} \ \underset{\boldsymbol{x}'' \in B_X(\boldsymbol{x}',\omega)}{\text{Max}} \ \underset{\boldsymbol{c} \in C'}{\text{Min}} \ \delta(\boldsymbol{x}'', \boldsymbol{c}, \boldsymbol{x}', \omega).$$

■

Proof of Theorem 1

Because F^* is full-dimensional and $\boldsymbol{v} \in F^* \backslash \underline{\partial} F$, there exists $\boldsymbol{v}'$ such that it is in $int(F^*)$ and $\boldsymbol{v}' << \boldsymbol{v}$. First, we define two constants Δ and N that we need to define the equilibrium strategy.[28]

We choose $\Delta > 0$ such that

$$B(\boldsymbol{v}', 4\sqrt{n-1}\Delta) \subset int(F^*). \tag{9.3}$$

Further, if $\boldsymbol{v}$ is not a payoff vector associated with a pure strategy profile, and can be expressed as $\sum_{l=1}^K \lambda^l \boldsymbol{c}^l$ where each $\lambda^l > 0$, Δ should be small enough so that

$$B_X(\boldsymbol{v}, \Delta) \ \subset \ relint(X) \tag{9.4}$$

where $X = co\{\boldsymbol{c}^1, \ldots, \boldsymbol{c}^K\}$. Note that from (9.3), it follows that $v_i' > \Delta$. Now define $N \in \mathbb{N}$ by

$$N = \left\lceil \max_i \frac{M}{v_i' - \Delta} \right\rceil, \tag{9.5}$$

[28] 2Δ serves as the 'reward' for the punishers; while Δ serves as the radius of all self-accessible balls we will be dealing with in this proof. N is the number of minmaxing periods.

implying that $N+1 > M/(v_i' - \Delta)$ for all i. Let $\underline{\delta}_1$ be such that for $\delta \geqslant \underline{\delta}_1$,

$$\delta^N \geqslant \frac{N}{N+1}. \tag{9.6}$$

The last two inequalities guarantee the following inequality which will be critical later:

$$1 + \delta + \ldots + \delta^N \geqslant \frac{M}{v_i' - \Delta} \quad \forall i \,. \tag{9.7}$$

Next, given N, there is a $\underline{\delta}_2$, such that for $\delta \geqslant \underline{\delta}_2$,

$$\delta^N \geqslant \frac{M}{M+\Delta} \quad \forall i. \tag{9.8}$$

Note that this implies, for $\delta \geqslant \underline{\delta}_2$,

$$\delta^N \geqslant \frac{v_i'}{v_i' + \Delta} \geqslant \frac{v_i' - \Delta}{v_i' + \Delta} \quad \forall i. \tag{9.9}$$

If $\boldsymbol{v}$ is a payoff vector for a pure action profile, let $\underline{\delta}_3 = 0$. Otherwise, $B_X(\boldsymbol{v}, \Delta)$ referred to in (9.4) is self-accessible for all δ larger than some bound; let $\underline{\delta}_3$ be that bound computable via Proposition 1. Lastly define

$$\boldsymbol{v}'(i) := (v_1' + 2\Delta, \ldots, v_{i-1}' + 2\Delta,\ v_i',\ v_{i+1}' + 2\Delta, \ldots, v_n' + 2\Delta)$$

Consider the set of all (full-dimensional) balls the centers of which are at most $\sqrt{n-1}\ \Delta$ away from $\boldsymbol{v}'(i)$, each with radius Δ. It may be checked because of (9.3), each of these small balls are fully contained in the interior of F^*. Using Proposition 2, a uniform bound can be computed such that each of these balls is self-accessible when the common discount factor is as large as the bound. Call this bound δ_{4i}. Now define

$$\underline{\delta} = \max(\underline{\delta}_1,\ \underline{\delta}_2,\ \underline{\delta}_3,\ \max_i \underline{\delta}_{4i}). \tag{9.10}$$

For any discount factor δ exceeding this bound, $\boldsymbol{v}$ may be supported by a three-phase strategy which we now describe.

Phase I: If $\boldsymbol{v}$ is a pure action profile, play that pure action profile forever. Otherwise, play the action sequence $\{\boldsymbol{a}^{(t)}(\boldsymbol{v}, B_X(\boldsymbol{v}, \Delta), \delta)\}_{t=0}^{\infty}$. If there is a unilateral deviation by player i in Phase I, move to Phase II(i).

Phase II(i): For each of N periods play $\boldsymbol{m}^i$, the (possibly mixed) action profile that minmaxes i. If player j unilaterally deviates from this phase (i.e. he is observed to play an action that is not in the support of m_j^i), start Phase II(j). Otherwise, at the completion of this phase, go to Phase III(i).

Phase III(i): Let $\tilde{\boldsymbol{a}}^{(t)}$, $t = 1, \ldots, N$ be the realized actions during Phase II(i). In Phase III(i), play the sequence of actions given by $\{\boldsymbol{a}^{(t)}(\boldsymbol{v}'(i) - \boldsymbol{z}^i - \Delta \boldsymbol{e}^i, B(\boldsymbol{v}'(i) - \boldsymbol{z}^i, \Delta), \delta)\}_{t=0}^{\infty}$, where $\boldsymbol{z}^i$ is an adjustment vector defined by the following two equations

$$z_j^i = \begin{cases} \frac{(1-\delta^N)}{\delta^N} r_j^i & \text{if } j \neq i \\ 0 & \text{otherwise.} \end{cases} \tag{9.11}$$

$$r_j^i = \frac{(1-\delta)}{(1-\delta^N)} \sum_{t=1}^{N} \delta^{t-1}\, g_i(\tilde{\boldsymbol{a}}^{(t)}), \tag{9.12}$$

If there is any unilateral deviation from Phase III(i) by player j, start Phase II(j).

Inequality (9.8) implies $\frac{1-\delta^N}{\delta^N}\ M \leqslant \Delta$ and since, $|r_j^i| \leqslant M$, $|z_j^i| \leqslant \Delta$ for $j \neq i$. This implies that $\boldsymbol{v}^{'}(i) - \boldsymbol{z}^i$ is at most $\sqrt{n-1}\ \Delta$ away from $\boldsymbol{v'}(i)$. Hence, given the construction of $\underline{\delta}_{4i}$ earlier, $B(\boldsymbol{v}^{'}(i) - \boldsymbol{z}^i, \Delta)$ is indeed self-accessible for discount factors above that bound. Let us now examine conditions for player i's strategy to be unimprovable.

- For unimprovability from Phase I, it suffices to have

$$(1-\delta)M + \delta^{N+1}(v_i' - \Delta) \leqslant v_i - \Delta, \tag{9.13}$$

- For unimprovability from Phase II(i) with τ periods left in the phase, it suffices to have

$$0 + \delta^{N+1}(v_i' - \Delta) \leqslant 0 + \delta^{\tau}(v_i' - \Delta) \quad \text{for } \tau = 1, \ldots, N. \tag{9.14}$$

- For unimprovability from Phase III(i), it suffices to have

$$(1-\delta)M + \delta^{N+1}(v_i' - \Delta) \leqslant v_i' - \Delta. \tag{9.15}$$

- To analyze unimprovability from Phase II(j) we note that because of the adjustment term $\boldsymbol{z}$ in Phase IIIj's target point, player i is indifferent in Phase IIj between playing any action that is in the support of $\boldsymbol{m}^j$. The question is whether he wishes to play an action which is not in the support of $\boldsymbol{m}^j$. Letting $\{\tilde{\boldsymbol{a}}^{(t)}\}_{t=1}^N$ denote the sequence of actions that would be realized in Phase II(j) on equilibrium path, with τ periods left in that phase, if the following inequality holds (no matter what the last τ entries of that sequenc $\{\tilde{\boldsymbol{a}}^{(t)}\}$ are), it would deter deviation:

$$(1-\delta)M + \delta^{N+1}(v_i' - \Delta) \leqslant (1-\delta)[g_i(\tilde{\boldsymbol{a}}^{(N-\tau+1)}) + \ldots + \delta^{\tau-1} g_i(\tilde{\boldsymbol{a}}^{(N)})] + \delta^{\tau}(v_i' + 2\Delta - z_i^j). \tag{9.16}$$

Using (9.11) and (9.12), the right hand side of (9.16) becomes:

$$\delta^{\tau}(v_i' + 2\Delta) - \frac{1-\delta}{\delta^{N-\tau}} \left(g_i(\tilde{\boldsymbol{a}}^{(1)}) + \ldots + \delta^{N-\tau-1} g_i(\tilde{\boldsymbol{a}}^{(N-\tau)}) \right), \tag{9.17}$$

which is bounded from below by $\delta_i^{\tau} \left(v_i' + 2\Delta - \frac{1-\delta^N}{\delta^N} M \right)$. Hence by (9.8), inequality (9.16) is satisfied for any sequence of $\boldsymbol{a}$'s if

$$(1-\delta)M + \delta_i^{N+1}(v_i' - \Delta) \leqslant \delta^N \left(v_i' + \Delta \right). \tag{9.18}$$

- Lastly, for unimprovability from Phase III(j) it suffices to have:

$$(1-\delta)M + \delta^{N+1}(v_i' - \Delta) \leqslant v_i' + 2\Delta - z_i^j - \Delta \tag{9.19}$$

for all possible values of z_i^j. Because $|z_i^j| \leqslant \Delta$, this is satisfied if the following holds:

$$(1-\delta)M + \delta^{N+1} v_i' \leqslant v_i'. \tag{9.20}$$

Examination of these conditions shows that while (9.14) is trivially true, (9.15) directly implies (9.13) and (9.20). Because of (9.9), it also implies (9.18). Thus to ensure all incentive compatibility conditions one only needs to satisfy equation (9.15), which is (9.7). Thus the given strategy profile is indeed an SPNE. ■

Proof of Proposition 3

Without loss of generality, we can and henceforth do discard points in C that are not extreme points of X, and relabel it as $C' = \{c^1, \cdots, c^{L'}\}$. For $\boldsymbol{x} \in B(\boldsymbol{o}, r)$, and $\boldsymbol{c} \in C'$ define a vector $\boldsymbol{y}(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, r, \boldsymbol{k}, \theta)$ where

$$y_i(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, r, \boldsymbol{k}, \theta) := \theta k_i(x_i - c_i) + x_i \quad \text{for } i = 1, \ldots, n. \tag{9.21}$$

In terms of discount factors, the above is just $\frac{1}{\delta_i} x_i - \frac{1-\delta_i}{\delta_i} c_i$, i.e. it is the i'th coordinate of the'continuation point' given the target $\boldsymbol{x}$, the current action $\boldsymbol{c}$ and the discount factor vector $\boldsymbol{\delta}$. Let

$$\begin{aligned} f(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, r, \boldsymbol{k}, \theta) &:= ||\boldsymbol{y}(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, r, \boldsymbol{k}, \theta) - \boldsymbol{o}||^2 - r^2 \\ &= \sum_{i=1}^{n} (\theta\, k_i(x_i - c_i) + (x_i - o_i))^2 - r^2. \\ &= \theta^2 \sum_{i=1}^{n} k_i^2 (x_i - c_i)^2 + 2\theta \sum_{i=1}^{n} k_i(x_i - c_i)(x_i - o_i) + (\sum_{i=1}^{n} (x_i - o_i)^2 - r^2) \end{aligned} \tag{9.22}$$

Because $B(\boldsymbol{o}, r)$ is full-dimensional, $f \leqslant 0 \implies \boldsymbol{y}(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, r, \boldsymbol{k}, \theta) \in B(\boldsymbol{o}, r)$. To prove the proposition, we will show that there exists $\bar{\theta}(\boldsymbol{o}, r, \boldsymbol{k}) > 0$ such that if $0 < \theta \leqslant \bar{\theta}(\boldsymbol{o}, r, \boldsymbol{k})$, then for every $\boldsymbol{x} \in B(\boldsymbol{o}, r)$, there exists a $\boldsymbol{c}$ such that $f \leqslant 0$.

The expression in (9.22) is a strictly convex quadratic in θ with $f(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, r, \boldsymbol{k}, 0) \leqslant 0$ (and so has at least one non-negative real root). Let $\theta(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, r, \boldsymbol{k})$ denote its larger root, which is continuous in $(\boldsymbol{x}, \boldsymbol{o}, r, \boldsymbol{k})$.

For $\boldsymbol{x} = \boldsymbol{o}$, for every $\boldsymbol{c}$, $f(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, r, \boldsymbol{k}, 0) = -r^2 < 0$ and hence, the larger root is strictly positive and hence, for every $\boldsymbol{c}$, $f \leqslant 0$ for $\theta \in (0, \theta(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, r, \boldsymbol{k})]$.

For $\boldsymbol{x} \neq \boldsymbol{o}$, if we can show that there exists a $\boldsymbol{c}$ such that $\frac{\partial f}{\partial \theta} < 0$ at $\theta = 0$ then we can assert that for that $\boldsymbol{x}$ there exists a $\boldsymbol{c}$, such that $\theta(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, r, \boldsymbol{k}) > 0$ and for $\theta \in (0, \theta(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, r, \boldsymbol{k})]$, f is non-positive. It suffices to show that that for every $\boldsymbol{x}$, there exists a $\boldsymbol{c}$ such that $\sum_{i=1}^{n} k_i(x_i - c_i)(x_i - o_i) < 0$, or

$$\sum_{i=1}^{n} k_i(x_i - o_i)x_i < \sum_{i=1}^{n} k_i(x_i - o_i)c_i \tag{9.23}$$

Consider the hyperplane $H = \{\boldsymbol{y} : \boldsymbol{p}\boldsymbol{y} = \alpha\}$, where $p_i = k_i(x_i - o_i)$ and $\alpha = \sum_{i=1}^{n} k_i(x_i - o_i)x_i$ (since $\boldsymbol{x} \neq \boldsymbol{o}$, $\boldsymbol{p}$ is a non-zero vector). If inequality (9.23) is false for every vertex $\boldsymbol{c}$, that would mean that every vertex lies on one side of the hyperplane, while clearly $\boldsymbol{x}$ is situated *on* that hyperplane. This can not be true since $\boldsymbol{x}$ is in a ball which lies in the *interior* of $co(C')$.

The foregoing analysis implies $\theta^*(\boldsymbol{x}, \boldsymbol{o}, r, \boldsymbol{k}) := \max_{\boldsymbol{c} \in C} \theta(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, r, \boldsymbol{k})$ is a strictly positive number, and is continuous in $\boldsymbol{k}, \boldsymbol{o}, r$ and $\boldsymbol{x}$ (being the maximum of continuous functions), and furthermore, for every $\boldsymbol{x} \in B(\boldsymbol{o}, r)$, if $\theta \in (0, \theta^*(\boldsymbol{x}, \boldsymbol{o}, r, \boldsymbol{k})]$, $f \leqslant 0$. Finally, define the required bound as

$$\bar{\theta}(\boldsymbol{o}, r, \boldsymbol{k}) := \min_{\boldsymbol{x} \in B(\boldsymbol{o}, r)} \theta^*(\boldsymbol{x}, \boldsymbol{o}, r, \boldsymbol{k}). \tag{9.24}$$

Since $B(\boldsymbol{o}, r)$ is compact, this minimum is achieved at some $\boldsymbol{x}$, is strictly positive-valued, and because of the Maximum Theorem is continuous in $\boldsymbol{o}$, r and $\boldsymbol{k}$. Clearly it satisfies the desired requirements of the bound. ■

Proof of Proposition 4

Define $r_i := \frac{1-\delta_i}{1-\underline{\delta}} r$ for each i. With these being the lengths of the semi-axes of the desired ellipsoid, the latter can be written as: $E(\boldsymbol{o}, r, \boldsymbol{\delta}, \underline{\delta}) = \{\boldsymbol{x} : \sum_{i=1}^{n} \frac{(x_i - o_i)^2}{r_i^2} \leqslant 1\}$. It is easy to see that $E(\boldsymbol{o}, r, \boldsymbol{\delta}, \underline{\delta})$ is contained in $B(\boldsymbol{o}, r)$ and can be written as $f(B(\boldsymbol{o}, r))$ where f is a 1-1 correspondence from $B(\boldsymbol{o}, r)$ to $E(\boldsymbol{o}, r, \boldsymbol{\delta}, \underline{\delta})$ given by

$$f_i(\boldsymbol{x}) = o_i + (x_i - o_i)\,\frac{1-\delta_i}{1-\underline{\delta}} \qquad \forall i. \tag{9.25}$$

Now let $\boldsymbol{x} \in E(\boldsymbol{o}, r, \boldsymbol{\delta}, \underline{\delta})$; to prove the proposition, we need to show that there exists $\boldsymbol{c} \in C$ such that $\boldsymbol{y} \in E(\boldsymbol{o}, r, \boldsymbol{\delta}, \underline{\delta})$ where

$$y_i = \frac{1}{\delta_i} x_i - \frac{1-\delta_i}{\delta_i} c_i \qquad \forall i. \tag{9.26}$$

To see this, let $\boldsymbol{x} = f(\boldsymbol{x}')$ where $\boldsymbol{x}' \in B(\boldsymbol{o}, r)$. By equation (9.25),

$$(1-\underline{\delta})x_i = (1-\delta_i)x_i' + (\delta_i - \underline{\delta})o_i \quad \forall i. \tag{9.27}$$

and because of the self-accessibility of $B(\boldsymbol{o}, r)$ for $\underline{\delta}\boldsymbol{\iota}$, there exists $\boldsymbol{c}$ such that

$$\sum_{i=1}^{n} \frac{1}{r^2}\left(\frac{1}{\underline{\delta}}\left[x_i' - (1-\underline{\delta})c_i\right] - o_i\right)^2 \leqslant 1 \tag{9.28}$$

Using this particular $\boldsymbol{c}$ in the defintion of $\boldsymbol{y}$ in (9.26),

$$\begin{aligned}
\frac{y_i - o_i}{r_i} &= \frac{1}{r}\frac{1-\underline{\delta}}{1-\delta_i}\left\{\frac{1}{\delta_i}[x_i - (1-\delta_i)c_i] - o_i\right\} && (9.29)\\
&= \frac{1}{r}\left\{\frac{1}{\delta_i}\left[\frac{1-\underline{\delta}}{1-\delta_i}x_i - (1-\underline{\delta})c_i\right] - \frac{1-\underline{\delta}}{1-\delta_i}o_i\right\} && (9.30)\\
&= \frac{1}{r}\left\{\frac{1}{\delta_i}\left[x_i' + \frac{\delta_i - \underline{\delta}}{1-\delta_i}o_i - (1-\underline{\delta})c_i\right] - \frac{1-\underline{\delta}}{1-\delta_i}o_i\right\} && (9.31)\\
&= \frac{1}{r}\left\{\frac{1}{\delta_i}\left[x_i' - (1-\underline{\delta})c_i\right] - \frac{\underline{\delta}}{\delta_i}o_i\right\} && (9.32)\\
&= \frac{1}{r}\left\{\frac{\underline{\delta}}{\delta_i}\left(\frac{1}{\underline{\delta}}\left[x_i' - (1-\underline{\delta})c_i\right] - o_i\right)\right\} && (9.33)
\end{aligned}$$

where we have used (9.27) to go from (9.30) to (9.31). Now, equation (9.28) and the fact that $\delta_i \geqslant \underline{\delta}$ for each i, allow us to conclude that $\sum_{i=1}^{n}\left(\frac{y_i - o_i}{r_i}\right)^2 \leqslant 1$ and we are done. ■

Proof of Theorem 2

The reader is requested to refer once again to the proof of Theorem 1 as we point out the parallels and dissmilarities between that proof and the current one. We choose $\boldsymbol{v}'$, Δ, N, $\boldsymbol{v}_i'$, $\underline{\delta}_1, \underline{\delta}_2, \underline{\delta}_3, \underline{\delta}_{4i}$ and hence, $\underline{\delta}$ exactly as before. This guarantees that equation (9.7) is valid with δ being replaced by δ_i (since δ_i is at least as large as $\underline{\delta}$ and hence $\underline{\delta}_1$), Thus, we have

$$1 + \delta_i + \ldots + {\delta_i}^N \geqslant \frac{M}{v_i' - \Delta} \quad \forall i \, . \tag{9.34}$$

Similarly, we may argue, since $\delta_i \geqslant \underline{\delta} \geqslant \underline{\delta}_2$,

$$\delta_i^N \geqslant \frac{M}{M + \Delta} \quad \forall i. \tag{9.35}$$

Define $\Delta_i := \Delta \frac{1-\delta_i}{1-\underline{\delta}}$. From the previous inequality it also follows that

$$\delta_i^N \geqslant \frac{v_i'}{v_i' + \Delta} \geqslant \frac{v_i' - \Delta_i}{v_i' + \Delta} \quad \forall i.. \tag{9.36}$$

Next, we describe the strategies which follow the standard three-phase pattern used previously. Since $\boldsymbol{\delta} \geqslant \underline{\delta}\boldsymbol{\iota}$, and $B(\boldsymbol{v}, \Delta)$ is self-accessible for $\underline{\delta}\boldsymbol{\iota}$, via Proposition 4, we know that the ellipsoid $E(\boldsymbol{v}, \Delta, \boldsymbol{\delta}, \underline{\delta}) \subset B(\boldsymbol{v}, \Delta)$ is self-accessible for $\boldsymbol{\delta}$. In Phase I, it is prescribed that the players play the sequence $\{\boldsymbol{a}^{(t)}(\boldsymbol{v}, E(\boldsymbol{v}, \Delta, \boldsymbol{\delta}, \underline{\delta}))\}_{t=0}^{\infty}$. Note that at anytime during this phase, the worst continuation payoff for player i is $v_i - \Delta_i$. Phase II(i)'s play does not change at all. To describe Phase III(i), define the quantities z_j^i and r_j^i's as before except to use δ_j rather that δ in their expressions given by equations (9.11) and (9.12). With these new definitions in place, now in Phase III(i), let the players play the action sequence $\{\boldsymbol{a}^{(t)}(\boldsymbol{v}'(i) - \boldsymbol{z}^i - \Delta_i \boldsymbol{e}_i, E(\boldsymbol{v}'(i) - \boldsymbol{z}^i, \Delta, \boldsymbol{\delta}, \underline{\delta}))\}_{t=0}^{\infty}$. The transitions between the phases follow the same pattern as before.

It now remains to verify the incentive-compatibility conditions which are exactly the same as before except that each occurrence of δ is now subscripted with an i and *some* occurrences of Δ are subscripted with an i.

- For unimprovability from Phase I, it suffices to have

$$(1 - \delta_i)M + {\delta_i}^{N+1}(v_i' - \Delta_i) \leqslant v_i - \Delta_i, \tag{9.37}$$

- For unimprovability from Phase II(i) with τ periods left in the phase, it suffices to have

$$0 + {\delta_i}^{N+1}(v_i' - \Delta_i) \leqslant 0 + {\delta_i}^{\tau}(v_i' - \Delta_i) \quad \text{for } \tau = 1, \ldots, N. \tag{9.38}$$

- For unimprovability from Phase III(i), it suffices to have

$$(1 - \delta_i)M + {\delta_i}^{N+1}(v_i' - \Delta_i) \leqslant v_i' - \Delta_i. \tag{9.39}$$

- Unimprovability from Phase II(j), after identical analysis undertaken before, is assured by:

$$(1 - \delta_i)M + \delta_i^{N+1}(v_i' - \Delta_i) \leqslant \delta_i^N \left(v_i' + \Delta\right). \tag{9.40}$$

- Lastly, for unimprovability from Phase III(j) it suffices to have:

$$(1 - \delta_i)M + \delta_i^{N+1} v_i' \leqslant v_i'. \tag{9.41}$$

As in the previous proof, the nontrivial inequalities (9.37), (9.41) are directly guaranteed by (9.39) while (9.40) is guaranteed by (9.39) because of (9.36). (9.39) is itself guaranteed by (9.34) and hence.... ■

Proof of Fact 1

First, we prove an analogous 'uniform' version of Theorem 1: Let F^* be full-dimensional. Suppose $\boldsymbol{u}$ and $r > 0$ be such that $B(\boldsymbol{u}, r) \subset int(F^*)$. Then, there exists a uniform discount factor bound $\underline{\delta}$ such that when $\delta \in [\underline{\delta}, 1)$, every point $\boldsymbol{v} \in B(\boldsymbol{u}, r)$ is an SPNE payoff for $\boldsymbol{\delta} = \delta\boldsymbol{\iota}$.

To see this, consider the set $V(\varepsilon) = \{\boldsymbol{v'} : \boldsymbol{v'} = \boldsymbol{u'} - \varepsilon\boldsymbol{\iota}, \boldsymbol{u'} \in \underline{\partial}B(\boldsymbol{u}, r)\}$ where $\underline{\partial}B(\boldsymbol{u}, r)$ is the lower boundary of $B(\boldsymbol{u}, r)$. Clearly, there exists a ε small enough, say $\bar{\varepsilon}$ such that $V(\bar{\varepsilon})$ is inside $int(F^*)$. For every point $\boldsymbol{v}$ in $B(\boldsymbol{u}, r)$, we can then find a $\boldsymbol{v'} \in V$ such that $\boldsymbol{v'} << \boldsymbol{v}$ (here and in the rest of this proof, we use the same notation we used in the proof of Theorem 1). In addition, there is also a uniform $\Delta > 0$ such that for every such pair of pair $\boldsymbol{v}$ and $\boldsymbol{v'}$ the conditions (9.3) and (9.4) hold. Now N, a uniform number of punishment periods can be chosen as $\left\lceil \max_{\boldsymbol{v'} \in V(\bar{\varepsilon})} \max_i \frac{M}{v'_i - \Delta} \right\rceil$. Having defined N, $\underline{\delta}_1$ and $\underline{\delta}_2$ can be defined as before. A uniform bound can can be chosen for $\underline{\delta}_3$ as $\boldsymbol{v}$ varies over the compact set $B(\boldsymbol{u}, r)$ since the proof of Proposition 1 plus an application of Maximum Theorem shows that the discount factor bound found in that proposition is continuous in the center of the relevant ball. Similarly, for each i, using the proof of Proposition 2 and the compactness of $V(\bar{\varepsilon})$ over which $\boldsymbol{v'}$ varies, a uniform bound for $\underline{\delta}_{4i}$ can be found. Choosing $\underline{\delta}$ to be maximum of $\underline{\delta}_1, \underline{\delta}_2$, and the last two uniform bounds, works as a common discount factor bound for supporting all points in $B(\boldsymbol{u}, r)$.

Now Fact 1 follows from this the same way Theorem 2 follows from Theorem 1. ■

Proof of Theorem 3

Step 1. Let $\boldsymbol{k} >> \boldsymbol{0}$ and player i's discount factor be given by $\delta_i = \frac{1}{1+k_i\theta}$ where for now, θ, a positive number, is unspecified. For a given payoff vector $\boldsymbol{v}$, consider the problem of designing a path such that $\boldsymbol{v}$ is realized through an $m+1$ phase path described as follows. For given non-negative numbers $b_1, \ldots, b_m$, in phase 1, lasting for $T_1 = \lceil (b_1/\theta) \rceil$ periods, a certain vertex $\boldsymbol{c}^{(1)}$ will be played, then in phase 2, lasting for the next $T_2 = \lceil (b_2/\theta) \rceil$ periods some vertex $\boldsymbol{c}^{(2)}$ will be played, etc., and for phase m, vertex $\boldsymbol{c}^{(m)}$ will be played for $T_m = \lceil (b_m/\theta) \rceil$ periods. In the $m+1$'th phase, an action sequnce that generates a continuation payoff will be played so that the whole path indeed realizes $\boldsymbol{v}$. Call this continuation payoff $\tilde{\boldsymbol{v}}(\theta)$. Because of Proposition 3, our strategy will succeed if at least for small values of θ, at the end of the first m phases, the continuation payoff enters $int(F)$. This suggests that we need to know what $\tilde{\boldsymbol{v}} := \lim_{\theta \to 0} \tilde{\boldsymbol{v}}(\theta)$ is.

If $\boldsymbol{x}^t$ denotes the continuation payoff during the t'th period of any path, and in the t'th period vertex $\boldsymbol{c}$ is played, then for each i the following holds: $x_i^t = (1 - \delta_i)c_i + \delta_i x_i^{t+1}$. Since $\delta_i = \frac{1}{1+k_i\theta}$, we can rewrite this as

$$x_i^{t+1} - c_i = (1 + k_i\theta)(x_i^t - c_i). \tag{9.42}$$

Hence, reasoning recursively, if $\boldsymbol{c}$ is played for T periods in periods $t, t+1, \ldots t+T-1$, then for each i,

$$x_i^{t+T} - c_i = (1 + k_i\theta)^T (x_i^t - c_i). \tag{9.43}$$

If $T = \lceil \frac{b}{\theta} \rceil$, where b is some non-negative number, then as θ tends to 0, the i'th coordinate of the limiting continuation payoff vector will satisfy

$$\lim_{\theta \to 0} x_i^{t+T} - c_i = e^{k_i b}(x_i^t - c_i). \tag{9.44}$$

Hence, in the context of the $m+1$ phase path discussed above if $m=1$,

$$\tilde{v}_i - c_i^{(1)} = e^{k_i b_1}(v_i - c_i^{(1)}) \tag{9.45}$$

and therefore using the same idea twice, if $m=2$,

$$\begin{aligned}\tilde{v}_i - c_i^{(2)} &= e^{k_i b_2}(c_i^{(1)} + e^{k_i b_1}(v_i - c_i^{(1)}) - c_i^{(2)}) \\ &= e^{k_i b_2}(c_i^{(1)} - c_i^{(2)}) + e^{k_i(b_2+b_1)}(v_i - c_i^{(1)})\end{aligned} \tag{9.46}$$

Proceeding inductively, we conclude that for arbitrary integer m, for each i we will have

$$\tilde{v}_i - c_i^{(m)} = e^{k_i b_m}(c_i^{(m-1)} - c_i^{(m)}) + e^{k_i(b_m+b_{m-1})}(c_i^{(m-2)} - c_i^{(m-1)}) + \cdots + e^{k_i(b_m+\cdots+b_1)}(v_i - c_i^{(1)}) \tag{9.47}$$

Step 2. Let $\hat{\boldsymbol{v}}, \boldsymbol{v}$ be any two points in $int(re(F))$ such that for all i, $\hat{v}_i \neq v_i$. We claim that for any $\varepsilon > 0$, there exist positive numbers $k_1, \ldots, k_n$, $b_1, \ldots, b_n$, and vertices $\boldsymbol{c}^{(1)}, \ldots, \boldsymbol{c}^{(n)}$, such that the system

$$\begin{aligned}\hat{v}_1 - c_1^{(n)} &= e^{k_1 b_n}(c_1^{(n-1)} - c_1^{(n)}) + e^{k_1(b_n+b_{n-1})}(c_1^{(n-2)} - c_1^{(n-1)}) + \cdots + e^{k_1(b_n+\cdots+b_1)}(v_1 - c_1^{(1)}) \\ \hat{v}_2 - c_2^{(n)} &= e^{k_2 b_n}(c_2^{(n-1)} - c_2^{(n)}) + e^{k_1(b_n+b_{n-1})}(c_2^{(n-2)} - c_2^{(n-1)}) + \cdots + e^{k_1(b_n+\cdots+b_1)}(v_2 - c_2^{(1)}) \\ &\vdots \\ \hat{v}_n - c_n^{(n)} &= e^{k_n b_n}(c_n^{(n-1)} - c_n^{(n)}) + e^{k_1(b_n+b_{n-1})}(c_n^{(n-2)} - c_n^{(n-1)}) + \cdots + e^{k_n(b_n+\cdots+b_1)}(v_n - c_n^{(1)})\end{aligned} \tag{9.48}$$

has an ε-solution, in the sense that the two sides of each equation differ by at most ε.

We give an induction-type argument to justify our claim. First, we will show that if we consider just two players, in fact an exact solution is possible. Wlog, let these be players 1 and 2; we will specifically show that there exist k_1, k_2, b_1, b_2 (all positive), and vertices $\boldsymbol{c}^{(1)}, \boldsymbol{c}^{(2)} \in C$, such that the following two equations hold:

$$\hat{v}_1 - c_1^{(2)} = e^{k_1 b_2}(c_1^{(1)} - c_1^{(2)}) + e^{k_1(b_2+b_1)}(v_1 - c_1^{(1)}) \tag{9.49}$$

$$\hat{v}_2 - c_2^{(2)} = e^{k_2 b_2}(c_2^{(1)} - c_2^{(2)}) + e^{k_2(b_2+b_1)}(v_2 - c_2^{(1)}) \tag{9.50}$$

Now choose vertices $\boldsymbol{c}^{(1)}, \boldsymbol{c}^{(2)}$ satisfying the conditions below.

$$sgn(\hat{v}_1 - v_1) = sgn(v_1 - c_1^{(1)}) \tag{9.51}$$

$$sgn(\hat{v}_2 - v_2) = sgn(v_2 - c_2^{(2)}) \tag{9.52}$$

Note that since $\hat{v}_i \neq v_i$ for all i, and both $\hat{\boldsymbol{v}}$ and $\boldsymbol{v}$ are in $int(re(C))$, it is possible to find two such vertices (and they could be the same vertex), where none of the signum functions above return 0. Let us set $k_2 = 1$ and make the following substitutions: $e^{b_2} = 1 + p$ and $e^{b_1+b_2} = 1 + p + q$ in equations (9.49) and (9.50). After some cancellations, the second equation can be rewritten as

$$\hat{v}_2 - v_2 = p(v_2 - c_2^{(2)}) + q(v_2 - c_2^{(1)}) \tag{9.53}$$

while the first equation becomes

$$\hat{v}_1 - c_1^{(2)} = (1+p)^{k_1}(c_1^{(1)} - c_1^{(2)}) + (1+p+q)^{k_1}(v_1 - c_1^{(1)}). \tag{9.54}$$

We are interested in positive k_1, p, q that will solve this pair of equations. In (9.53), we can choose q to be very small and strictly positive such that $sgn((\hat{v}_2 - v_2) - q(v_2 - c_2^{(1)})) = sgn(\hat{v}_2 - v_2)$, and then because of (9.52), we can find a positive p that solves (9.53). From p and q, b_1 and b_2 may be extracted. Turning to equation (9.54), we note that if $k_1 \to 0$, $sgn(RHS - LHS)$ is the same as $sgn(v_1 - \hat{v}_1)$. On the other hand, as $k_1 \to \infty$, because of the positivity of p and q, $sgn(RHS - LHS)$ is the same as $sgn(v_1 - c_1^{(1)})$. Hence, by the Intermediate Value Theorem, there exists $k_1 > 0$ which solves the equation if $sgn(v_1 - \hat{v}_1) = -sgn(v_1 - c_1^{(1)})$ or $sgn(\hat{v}_1 - v_1) = sgn(v_1 - c_1^{(1)})$. But this is just (9.51).[29]

Now, we show that for any $m < n$, and for any $\varepsilon > 0$ if an ε-solution exists for $m-1$ equation version of (9.48), where the vertices are chosen according to the rule $sgn(\hat{v}_i - v_i) = sgn(v_i - c_i^{(i)})$ for each player i then a solution exists for the m equation version as well where the additional vertex is chosen using the same rule (used for the additional player). To see this, consider (9.48) with n replaced by m; suppose we set $b_1 = 0$ in all those m equations. If we consider the second through the m-th equation of the system, they become exactly the system for an $m-1$ player scenario (where the players are indexed 2 through m). This is because in the equation pertaining to player i ($i = 2, \ldots, m$) the sum of the last two terms $e^{k_i(b_m+\cdots+b_2)}(c_i^{(1)} - c_i^{(2)}) + e^{k_i(b_m+\cdots+b_1)}(v_i - c_i^{(1)})$ collapses to the single term $e^{k_m(b_m+\cdots+b_2)}(v_i - c_i^{(2)})$ on setting b_1 to zero (note that $\boldsymbol{c}^{(1)}$ disappears from the system as a result as well). We will call this particular system the 'revised system'. Choose $k_2, \ldots, k_m$, $\boldsymbol{c}^{(2)}, \ldots \boldsymbol{c}^{(m)}$ and $b_2, \ldots, b_m$, so that the left hand side and the right hand side of each equation in this revised system differ by at most $\frac{\varepsilon}{2}$. Next, choose b_1 small enough so that the right hand sides of the original equations for players 2 through m and right hand sides of the corresponding revised equations differ by at most $\frac{\varepsilon}{2}$, *no matter what* $\boldsymbol{c}^{(1)}$ *is*, which can indeed be ensured since the right hand sides of the original equations are continuous in b_1. Thus, we have chosen now $k_2, \ldots, k_m$ and $b_1, \ldots, b_m$ and vertices $\boldsymbol{c}^{(2)}, \ldots, \boldsymbol{c}_2^{(m)}$ such that equations 2 through m satisfy the desired property. It remains to tackle the first equation and determine k_1. Indeed k_1 can now be chosen to satisfy the equation exactly. Once again, as can be easily checked, this is an application of the Intermediate Value Theorem, as long as $\boldsymbol{c}^{(1)}$ satisfies the condition $sgn(\hat{v}_1 - v_1) = sgn(v_1 - c_1^{(1)})$. This proves the claim.

Step 3. Now, choose any $\boldsymbol{u}$ in $int(F)$ such that there exists a ν-ball around $\boldsymbol{u}$ which lies fully inside $int(F)$ (guaranteed as a consequence of the full-dimensionality assumption). Let $\hat{\boldsymbol{v}}$ be such that the distance between $\hat{\boldsymbol{v}}$ and $\boldsymbol{u}$ is at most $\nu/3$, and for each i, $\hat{v}_i \neq v_i$. For this $\hat{\boldsymbol{v}}$ and the given $\boldsymbol{v}$, choose the k_i's and the b_i's and the $\boldsymbol{c}^i$'s using Step 2 such that for these parameters, $\tilde{\boldsymbol{v}}$, the limit point after the first n phases of the path described in Step 1 is at most $\nu/3$ away from $\hat{\boldsymbol{v}}$. Let $\bar{\theta}_1$ be such that for $\theta < \bar{\theta}_1$, the actual required continuation payoff $\tilde{\boldsymbol{v}}(\theta)$ is at most $\nu/3$ away from its limit point $\tilde{\boldsymbol{v}}$. This will ensure that when $\theta < \bar{\theta}_1$, for the k_i's chosen, and for $\delta_i = \frac{1}{1+k_i\theta}$, at the end of the first n phases, the required continuation payoff $\tilde{\boldsymbol{v}}(\theta)$ will be within a distance of ν from $\boldsymbol{u}$. For the chosen k_i's let $\bar{\theta}_2$ be the cutoff on θ that is required to make $B(\boldsymbol{u}, \nu)$ self-accessible as demonstrated in Proposition 3. Then, when $\theta < \bar{\theta} = \min(\bar{\theta}_1, \bar{\theta}_2)$, we can *simultaneously* ensure

[29] Writing equation (9.54) as $a = bx^{k_1} + cy^{k_1}$ where $y > x > 1$, it is easily seen that for $k_1 \geqslant ln(\frac{|a|+|b|}{|c|})/ln(\frac{y}{x})$, sgn(RHS - LHS) =sgn(c). Now the method of bisection can be used to identify the solution to any arbitrary desired degree of precision.

that $\tilde{\boldsymbol{v}}(\theta)$ is inside $B(\boldsymbol{u}, \nu)$ and that there is a sequence of actions that generate $\tilde{\boldsymbol{v}}(\theta)$. Combining this last phase with the n phases described in Step 1, it follows that for discount factors given by $\delta_i = \frac{1}{1+k_i\theta}$ with $\theta < \bar{\theta}$, the $n+1$ phase path described in that step will indeed realize $\boldsymbol{v}$. ■

Proof of Theorem 4

We first demonstrate the inclusion: $\bigcup\{\mathcal{F}(\boldsymbol{\delta}) \mid \boldsymbol{\delta} \in (0,1)^2\} \subset int(re(F))\bigcup F$. Let $\boldsymbol{x} \in F(\boldsymbol{\delta})$ for some $\boldsymbol{\delta}$. Obviously, $\boldsymbol{x} \notin (re(F))^c$, for otherwise there exists a player i such that x_i is either strictly greater or strictly less than what player i can achieve in the stage game – an impossibility. Next we show that if $\boldsymbol{x}$ is on the boundary of $re(F)$, then $\boldsymbol{x} \in F$. In this case, there exists i such that x_i is an extremal (either maximum or minimum) payoff for player i (in the stage game). Let $\{\boldsymbol{a}^{(t)} \mid t \in \mathbb{Z}_+\}$ be the sequence of actions played to realize $\boldsymbol{x}$, and define $C_i^e := \{(g(\boldsymbol{a}^{(t)}) \mid t \in \mathbb{Z}_+\}$ to denote the set of all payoff profiles earned in any period. Since x_i is an extremal payoff of i, we must have $g_i(\boldsymbol{a}^{(t)}) = x_i$ for all t; therefore not only is it true that $x_i = (1-\delta_i)\sum \delta_i^t g_i(\boldsymbol{a}^{(t)})$, but it is also true that $x_i = (1-\delta_j)\sum \delta_j^t g_i(\boldsymbol{a}^{(t)})$ for any δ_j. For player $j = 3-i$, of course we have $x_j = (1-\delta_j)\sum \delta_j^t g_j(\boldsymbol{a}^{(t)})$. Therefore, we can write the *vector* equality using j's discount factor: $\boldsymbol{x} = (1-\delta_j)\sum \delta_j^t g(\boldsymbol{a}^{(t)})$, where each $g(\boldsymbol{a}^{(t)}) \in C_i^e$. This implies $\boldsymbol{x} \in co(C_i^e) \subset F$.[30]

To demonstrate the other inclusion, $\bigcup\{F(\boldsymbol{\delta}) \mid \boldsymbol{\delta} \in (0,1)^2\} \supset int(re(F))\bigcup F$, we appeal to Theorem 3 (making use of the full-dimensionality assumption) and further note that for any $\boldsymbol{x} \in F$, Proposition 1 guarantees that $\boldsymbol{x} \in \mathcal{F}(\delta\boldsymbol{\iota})$ for sufficiently large δ. ■

Proof of Theorem 5

Let us wlog assume that $w_j = 0$ for all j. We begin by asserting that on any path for any player, if he is using discount factor δ, then provided all his continuation payoffs are nonnegative, increasing the discount factor to $\tilde{\delta} > \delta$ would keep all his continuation payoffs still nonnegative. To see this let s^t be the continuation payoff from t onwards when δ is used, i.e.

$$s^t = (1-\delta)[v^t + \delta v^{t+1} + \delta^2 v^{t+2} + \cdots] \tag{9.55}$$

where v^t is the player's actual payoff in period t. Equation (9.55) implies $\frac{s^t}{1-\delta} - v^t = \delta\frac{s^{t+1}}{1-\delta}$ and hence,

$$v^t = \frac{s^t}{1-\delta} - \delta\frac{s^{t+1}}{1-\delta}. \tag{9.56}$$

Similarly let $\tilde{s}^t$ be the continuation payoff from t onwards with the discount factor $\tilde{\delta}$. Using equations (9.55) and (9.56), we can write that as

$$\begin{aligned}
\tilde{s}^t &= (1-\tilde{\delta})\left[\left(\frac{s^t}{1-\delta} - \frac{\delta s^{t+1}}{1-\delta}\right) + \tilde{\delta}\left(\frac{s^{t+1}}{1-\delta} - \frac{\delta s^{t+2}}{1-\delta}\right) + \tilde{\delta}^2\left(\frac{s^{t+2}}{1-\delta} - \frac{\delta s^{t+3}}{1-\delta}\right) + \cdots\right] \\
&= \left(\frac{1-\tilde{\delta}}{1-\delta}\right)\left[s^t + (\tilde{\delta}-\delta)s^{t+1} + \tilde{\delta}(\tilde{\delta}-\delta)s^{t+2} + \cdots\right] \\
&\geqslant 0
\end{aligned} \tag{9.57}$$

Now assume for the moment that $\boldsymbol{v}$ was achieved as an equilibrium payoff vector with a discount factor vector where $\delta_1 \leqslant \delta_2 \leqslant \ldots \leqslant \delta_n$. Consider what payoff vector would realize if we stayed with

[30] When there are more than 2 players a little thought should convince the reader that this logic will not extend unless all players other than i have the same discount factor.

the same played path but increased each of player $1, 2, \ldots, n-1$'s discount factors to δ_n. Since, $\boldsymbol{v}$ is an equilibrium payoff it is weakly individual rational, and all continuation payoffs for all players for all periods must be nonnegative as well. Hence, after this adjustment of discount factors, the payoff vector we obtain has the first $n-1$ components non-negative, the last component is the same as v_n and moreover, since with equal discounting any play must result in a payoff vector that is in F, this particular payoff vector satisfies all the requirements of $\boldsymbol{u}^n$ in the WD condition (under the natural order).

Next consider the effect of changing all the discount factors to δ_{n-1}; this involves increasing the first $n-2$ discount factors, keeping the $(n-1)$'th discount factor same and *decreasing* the last discount factor. We cannot predict what happens to player n's payoff as a result, but we can surely claim that the first $n-2$ players' payoffs continue to remain non-negative, player $n-1$'s payoff remains at v_{n-1} and the whole payoff vector, taken together is in F. But then this new payoff vector satisfies the requirements of $\boldsymbol{u}^{n-1}$ in the WD condition. Proceeding similarly, each of the conditions imposed in WD can be seen to be satisfied. Finally, If the ordering of the discount factors is not 'natural', but it is the case that $\delta_{\pi_1} \leqslant \delta_{\pi_2} \leqslant \ldots \leqslant \delta_{\pi_n}$, for some permutation π, the same argument can be easily adapted to show that $\boldsymbol{v} \in W(\pi)$. ■

Proof of Proposition 5

Part a): In this proof we borrow notation used and results derived in the proof of Proposition 3. For a given ball $B(\boldsymbol{o}, r)$, let $\boldsymbol{y}(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, \boldsymbol{k}, \theta)$ refer to the continuation point when we decompose the current payoff vector $\boldsymbol{x}$ using the current action $\boldsymbol{c}$ while $\boldsymbol{k}, \theta$ parametrize the discount factor vector. We let $d(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, \boldsymbol{k}, \theta) := ||\boldsymbol{y} - \boldsymbol{o}||$ and $f(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, \boldsymbol{k}, \theta) = d^2 - r^2$.[31] The arguments used in Proposition 3 shows that for the fixed ball $B(\boldsymbol{o}, r)$ there exists a strictly positive-valued function $\bar{\theta}(\boldsymbol{o}, \boldsymbol{k})$, continuous in its arguments such that if $\theta \in (0, \bar{\theta}(\boldsymbol{o}, \boldsymbol{k}))$, for each $\boldsymbol{x} \in B(\boldsymbol{o}, r)$ there exists a vertex $\boldsymbol{c}^*(\boldsymbol{x}, \boldsymbol{o}, \boldsymbol{k})$ with the property $\sum_{i=1}^{n} k_i(x_i - c^*{}_i)(x_i - o_i) < 0$ and for that vertex, $f < 0$ and hence, $d < r$.[32] Define

$$\theta^\dagger(\boldsymbol{k}) = \min_{\boldsymbol{o} \in [\boldsymbol{y}, \boldsymbol{z}]} \bar{\theta}(\boldsymbol{o}, \boldsymbol{k}) \tag{9.58}$$

which is well-defined and strictly positive because $\bar{\theta}$ is strictly positive and continuous in its arguments and because $[\boldsymbol{y}, \boldsymbol{z}]$ is compact. Hence, if $\theta \in (0, \theta^\dagger(\boldsymbol{k}))$, for any $\boldsymbol{o} \in [\boldsymbol{y}, \boldsymbol{z}]$ and $\boldsymbol{x} \in B(\boldsymbol{o}, r)$, there exists $\boldsymbol{c}$, such that $f < 0$ or $d < r$. Now define $\bar{d}(\boldsymbol{k}, \theta)$ by

$$\bar{d}(\boldsymbol{k}, \theta) := \max_{\boldsymbol{o} \in [\boldsymbol{y}, \boldsymbol{z}]} \max_{\boldsymbol{x} \in B(\boldsymbol{o}, r)} \min_{\boldsymbol{c} \in C} d(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, \boldsymbol{k}, \theta) \tag{9.59}$$

which is well-defined and $< r$ because of continuity of $\boldsymbol{d}$ in $\boldsymbol{x}$ and $\boldsymbol{o}$, the Maximum Theorem and compactness of the relevant feasible sets. In $\bar{d}(\boldsymbol{k}, \theta)$, we now have created a uniform bound on the distance of the continuation payoff from the center of any ball that makes up the capsule and any target point in the ball by choosing the current action to be $\hat{\boldsymbol{c}}(\boldsymbol{x}, \boldsymbol{o}, \boldsymbol{k}, \theta) := \operatorname{argmin}_{\boldsymbol{c} \in C} d(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, \boldsymbol{k}, \theta)$ (ties being broken using any arbitrary preference ordering among vertices).[33]

Now, for $\theta < \theta^\dagger(\boldsymbol{k})$, we recursively construct the sequences $\{\boldsymbol{c}^{(t)}\}$ and $\{\boldsymbol{x}^t\}$ via a third sequence $\{\boldsymbol{o}^t\}$. Define $\boldsymbol{o}^0 = \boldsymbol{y}$ and $\boldsymbol{x}^0 = \boldsymbol{u}$. Of course, if $\boldsymbol{u}$ is already in $B(\boldsymbol{z}, r)$, we can stop

[31] $\boldsymbol{y}, d$ and f also depend on r, but we ignore this for brevity's sake as r stays fixed once we fix our capsule. This shortcut is used for other functions as well.

[32] Recall that $\boldsymbol{c}^*$ was chosen with a view to maximize the range of θ over which f stays non-positive.

[33] We note in passing that $\hat{\boldsymbol{c}}$ depends on θ and need not be the same as $\boldsymbol{c}^*$.

immediately with $T = 0$, so we assume that this is not the case. Let $\boldsymbol{c}^{(0)} = \hat{\boldsymbol{c}}(\boldsymbol{x}^0, \boldsymbol{o}^0, \boldsymbol{k}, \theta)$. Let $\boldsymbol{x}^1 = \boldsymbol{y}(\boldsymbol{x}^0, \boldsymbol{c}^{(0)}, \boldsymbol{o}^0, \boldsymbol{k}, \theta)$. Again, we stop with $T = 1$ if $\boldsymbol{x}^1 \in B(\boldsymbol{z}, r)$. Otherwise, we define $\boldsymbol{o}^1 = \operatorname{argmin}_{\substack{\boldsymbol{x}\in[\boldsymbol{y},\boldsymbol{z}] \\ ||\boldsymbol{x}-\boldsymbol{x}^1||=r}} ||\boldsymbol{x} - \boldsymbol{z}||$, and $\boldsymbol{c}^{(1)} = \hat{\boldsymbol{c}}(\boldsymbol{x}^1, \boldsymbol{o}^1, \boldsymbol{k}, \theta)$.

In general, given $\boldsymbol{x}^t, \boldsymbol{o}^t, \boldsymbol{c}^{(t)}$ we define

$$\boldsymbol{x}^{t+1} = \boldsymbol{y}(\boldsymbol{x}^t, \boldsymbol{c}^{(t)}, \boldsymbol{o}^t, \boldsymbol{k}, \theta) \tag{9.60}$$

$$\boldsymbol{o}^{t+1} = \operatorname{argmin}_{\substack{\boldsymbol{x}\in[\boldsymbol{y},\boldsymbol{z}] \\ ||\boldsymbol{x}-\boldsymbol{x}^{t+1}||=r}} ||\boldsymbol{x} - \boldsymbol{z}|| \tag{9.61}$$

$$\boldsymbol{c}^{(t+1)} = \hat{\boldsymbol{c}}(\boldsymbol{x}^{t+1}, \boldsymbol{o}^{t+1}, \boldsymbol{k}, \theta) \tag{9.62}$$

and we stop the recurrence with $T = t$ as soon as $\boldsymbol{x}^t \in B(\boldsymbol{z}, r)$. Indeed, we are assured of stopping in a finite number of steps because, $||\boldsymbol{o}^t - \boldsymbol{x}^{t+1}|| \leqslant \bar{d}(\boldsymbol{k}, \theta)$, $||\boldsymbol{x}^{t+1} - \boldsymbol{o}^{t+1}|| = r$ which on applying Triangle Inequality ensures that $||\boldsymbol{o}^t - \boldsymbol{o}^{t+1}||$ is at least $r - \bar{d}(\boldsymbol{k}, \theta)$. This implies $T \leqslant \left\lceil \frac{||y-z||}{r-\bar{d}(\boldsymbol{k},\theta)} \right\rceil$ and completes the proof of Part a).

Part b): To prove this part, we need to bound $\bar{d}(\boldsymbol{k}, \theta)$ by some suitable function of θ. We start by observing that the square of the d function

$$d^2(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, \boldsymbol{k}, \theta) = \theta^2 \sum_{i=1}^{n} k_i^2 (x_i - c_i)^2 + 2\theta \sum_{i=1}^{n} k_i (x_i - c_i)(x_i - o_i) + \sum_{i=1}^{n} (x_i - o_i)^2 \tag{9.63}$$

is convex in $\boldsymbol{x}$.[34] Now for any $\boldsymbol{x} \neq \boldsymbol{o}$ in $B(\boldsymbol{o}, r)$, let $\bar{\boldsymbol{x}}$ be the point on the surface of the ball that is intersected by the ray emanating from $\boldsymbol{o}$ and going towards $\boldsymbol{x}$, i.e. $\bar{\boldsymbol{x}} = \boldsymbol{o} + \frac{r}{||\boldsymbol{x}-\boldsymbol{o}||}(\boldsymbol{x} - \boldsymbol{o})$. Then, since, $\boldsymbol{x}$ is a convex combination of $\bar{\boldsymbol{x}}$ and $\boldsymbol{o}$, we have:

$$d^2(\boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, \boldsymbol{k}, \theta) \leqslant \max\{d^2(\bar{\boldsymbol{x}}, \boldsymbol{c}, \boldsymbol{o}, \boldsymbol{k}, \theta), d^2(\boldsymbol{o}, \boldsymbol{c}, \boldsymbol{o}, \boldsymbol{k}, \theta)\} \tag{9.64}$$

This in turn implies that

$$\begin{aligned} d^2(\boldsymbol{x}, \hat{\boldsymbol{c}}(\boldsymbol{x}, \boldsymbol{o}, \boldsymbol{k}, \theta), \boldsymbol{o}, \boldsymbol{k}, \theta) &\leqslant d^2(\boldsymbol{x}, \hat{\boldsymbol{c}}(\bar{\boldsymbol{x}}, \boldsymbol{o}, \boldsymbol{k}, \theta), \boldsymbol{o}, \boldsymbol{k}, \theta) \\ &\leqslant \max\{d^2(\bar{\boldsymbol{x}}, \hat{\boldsymbol{c}}(\bar{\boldsymbol{x}}, \boldsymbol{o}, \boldsymbol{k}, \theta), \boldsymbol{o}, \boldsymbol{k}, \theta), d^2(\boldsymbol{o}, \hat{\boldsymbol{c}}(\bar{\boldsymbol{x}}, \boldsymbol{o}, \boldsymbol{k}, \theta), \boldsymbol{o}, \boldsymbol{k}, \theta)\} \end{aligned} \tag{9.65}$$

We will now bound each component in the max function in equation (9.65). For the first component, we note that

$$d^2(\bar{\boldsymbol{x}}, \hat{\boldsymbol{c}}(\boldsymbol{k}, \bar{\boldsymbol{x}}, \boldsymbol{o}, \theta), \boldsymbol{o}, \boldsymbol{k}, \theta) \leqslant d^2(\bar{\boldsymbol{x}}, \check{\boldsymbol{c}}(\boldsymbol{k}, \bar{\boldsymbol{x}}, \boldsymbol{o}), \boldsymbol{o}, \boldsymbol{k}, \theta) \tag{9.66}$$

where $\check{\boldsymbol{c}}(\boldsymbol{k}, \bar{\boldsymbol{x}}, \boldsymbol{o})$ minimizes $\sum_{i=1}^{n} 2k_i(x_i - c_i)(x_i - o_i)$ where the minimand, as has been shown in the the proof of Proposition 3, is strictly negative. Now by referring to equation (9.63) we see that $d^2(\boldsymbol{k}, \bar{\boldsymbol{x}}, \hat{\boldsymbol{c}}(\boldsymbol{k}, \bar{\boldsymbol{x}}, \boldsymbol{o}, \theta), \boldsymbol{o}, \theta)$ is less than or equal to $m_1(\boldsymbol{k})\theta^2 + m_2(\boldsymbol{k})\theta + r^2$ where

$$m_1(\boldsymbol{k}) := \max_{\boldsymbol{x}\in\mathcal{C}(\boldsymbol{y},\boldsymbol{z},r), \boldsymbol{c}\in C} \sum_{i=1}^{n} k_i^2 (x_i - c_i)^2 > 0 \tag{9.67}$$

and

$$m_2(\boldsymbol{k}) := \max_{\substack{\boldsymbol{o}\in[\boldsymbol{y},\boldsymbol{z}] \\ \boldsymbol{x}\in B(\boldsymbol{o},r)}} \min_{\boldsymbol{c}\in C} \sum_{i=1}^{n} 2k_i (x_i - c_i)(x_i - o_i) < 0. \tag{9.68}$$

On the other hand, for the second of the two components in the right hand side of equation

[34] This can be easily verified by checking the positive definiteness of the Hessian.

(9.65), we notice that both the θ term and the constant term in (9.63) drop out and hence, that component is less than or equal to $m_1(\boldsymbol{k})\theta^2$. There exists a $\theta^\ddagger(\boldsymbol{k}) \leqslant \theta^\dagger(\boldsymbol{k})$ when $m_2(\boldsymbol{k})\theta + r^2 \geqslant 0$, and hence, for $\theta \in (0, \theta^\ddagger(\boldsymbol{k}))$, for all $\boldsymbol{o} \in [\boldsymbol{y}, \boldsymbol{z}]$, all $\boldsymbol{x} \in B(\boldsymbol{o}, r)$ (including the centers of the balls), $\min_{\boldsymbol{c} \in C} d^2(\boldsymbol{k}, \boldsymbol{x}, \boldsymbol{c}, \boldsymbol{o}, \theta) \leqslant m_1(\boldsymbol{k})\theta^2 + m_2(\boldsymbol{k})\theta + r^2$ and hence, $\bar{d}(\boldsymbol{k}, \theta))^2 \leqslant m_1(\boldsymbol{k})\theta^2 + m_2(\boldsymbol{k})\theta + r^2$. This shows $r - \bar{d}(\boldsymbol{k}, \theta) \geqslant r - \sqrt{m_1(\boldsymbol{k})\theta^2 + m_2(\boldsymbol{k})\theta + r^2}$ and the claim follows using the bound derived on T in the proof of part a). ■

Proof of Theorem 6

Wlog we assume that $\boldsymbol{w} = \boldsymbol{0}$ and $\boldsymbol{v}$ is strictly diagonalizable for the natural order, letting $\boldsymbol{u}^1, \ldots, \boldsymbol{u}^n$ be as in the definition of the property. We will show that $\boldsymbol{v}$ is an SPNE payoff.

We will first specify the equilibrium path since the equilibrium strategy is based off it. The equilibrium path will entail a 'pre-entry path' based on a sequence of actions lasting for $T(\theta)$ periods that will transition the (required) continuation payoff from $\boldsymbol{v}$ to a point inside a ball in $int(F^*)$, with $\boldsymbol{v}^t(\theta)$ denoting the continuation payoff at period t. The center of this ball will be $\boldsymbol{u}^n$ which we know, resides in $int(F^*)$. Thereafter, the equilibrium path will coincide with the path created by playing the SPNE strategy to support $\boldsymbol{v}^{T(\theta)}(\theta)$. Of course, to specify this strategy exactly, and hence the 'post-entry path' it leads to, we need to specify the radius of the ball of entry, the $\boldsymbol{k}$ vector and ensure that θ will be below a certain bound - we will do all that in due course.

The pre-entry path will be broken down into $n-1$ stages, each stage witnessing an application of the Capsule Lemma. In the l'th stage ($l = 1, \ldots, n-1$) we will be operating with a capsule that is situated inside the payoff space of the first l players for $T^l(\theta)$ periods and throughout this stage, these players' payoff (sub)vector will stay inside that capsule. We now describe these capsules. In what follows, if $\boldsymbol{x}$ is an n-dimensional vector, the subvector consisting of its first m coordinates will be denoted as $\boldsymbol{x}[m]$.

By strict diagonalizability, for every $l = 1, \ldots, n-1$, $\boldsymbol{u}^l[l]$ and $\boldsymbol{u}^{l+1}[l]$ both belong to $int(F(1, \ldots, l)) \cap \mathbb{R}^l_{++}$. Hence, there exists $\bar{r}_l$ such that if $r_l < \bar{r}_l$, the entire capsule $\mathcal{C}(\boldsymbol{u}^l[l], \boldsymbol{u}^{l+1}[l], r_l)$ also lies in $int(F(1, \ldots, l)) \cap \mathbb{R}^l_{++}$. This is the capsule we will work with in the l'th stage with r_l to be further specified later.

The purpose of operating the l'th stage is to change the continuation payoffs of players 1 through l (from what they were at the end of $l-1$'th stage). But it is not that during stages 1 through l, the continuation payoffs of player i, where $l+1 \leqslant i \leqslant n$, will stay put. However, we can make these players relatively patient so that their payoffs will not change by much and thus we can maintain strict individual rationality for those players (since v_i was strictly positive for all i). In particular, let ε be any number strictly below $\min_i v_i$. We will show that k_l's can be chosen so that during each of the stages $1, \ldots l-1$, player l's continuation payoff changes by at most ε/n (from the previous stage). This will ensure that at the beginning of stage l, his continuation payoff stays strictly positive. From stages l onwards we do not have to worry about the strict individual rationality of his continuation payoffs because for these stages they are in capsules within which every vector is strictly positive. What works here is that we choose a player's k *before* his payoff becomes part of any capsule, and once done, his payoffs can be transitioned through any capsule since capsules can handle any arbitrary $\boldsymbol{k}$ vector.

Given the vectors $\boldsymbol{u}^l, l = 1, \ldots, n$, we define the following *anchor* vectors: $\boldsymbol{z}^1 = \boldsymbol{v}$ and for

$l = 2, \dots, n$,

$$\boldsymbol{z}_i^l = \begin{cases} u_i^l & \text{for } i \leqslant l-1 \\ v_i & \text{for } i > l-1 \end{cases}$$

Notice that the the first $\boldsymbol{z}$ is the target payoff vector and the last $\boldsymbol{z}$ is $\boldsymbol{u}^n$, which is in $int(F^*)$. The point behind the terminology 'anchor' should be clear now: until the continuation payoff enters $int(F^*)$, the entire on equilibrium continuation payoff path will stay close to the following 'piecewise-linear' path:

$$\boldsymbol{z}^1 \longrightarrow \boldsymbol{z}^2 \longrightarrow \cdots \longrightarrow \boldsymbol{z}^n.$$

In the l'th stage, all continuation payoffs $\boldsymbol{v}^t(\theta)$, will be zig-zagging around the line segment joining $\boldsymbol{z}^l$ and $\boldsymbol{z}^{l+1}$.

We need to make sure that the 'starting ball' of the $l+1$'th stage will accommodate the continuation payoffs of players $1, \dots l$ arriving 'transformed' via the 'ending ball' of the previous capsule *as well as* the continuation payoff of player $l+1$. This is ensured by the following relation between the radii of one capsule and the next: $r_{l+1}^2 = r_l^2 + \varepsilon^2$ where ε is an upper bound on by how much player $l+1$'s payoff can change up until period $T^1(\theta) + \cdots + T^l(\theta)$. As our equilibrium strategy will show, we need 'room' around the final ball of entry inside F^* to take care of off-equilibrium behavior. Letting the maximum of the absolute value of the adjustment term needed to be Δ, [35] and the amount of reward for each punishing player to be 2Δ, we need to have a ball of center $\boldsymbol{u}^n$ and radius $r_n + 3\sqrt{n-1}\Delta$ fit inside $int(F^*)$. To summarize then, we impose the following restrictions on the sequence of capsule radii:

$$B(\boldsymbol{u}^l[l], r_l) \subset int(F(1,\dots,l)) \cap \mathbb{R}^l_{++} \qquad \text{for } l = 1,\dots,n-1 \tag{9.69}$$
$$B(\boldsymbol{u}^{l+1}[l], r_l) \subset int(F(1,\dots,l)) \cap \mathbb{R}^l_{++} \qquad \text{for } l = 1,\dots,n-1 \tag{9.70}$$
$$B(\boldsymbol{u}^n, r_n + 3\sqrt{n-1}\Delta) \subset int(F^*) \qquad \text{for some } \Delta > 0 \tag{9.71}$$
$$r_{l+1}^2 = r_l^2 + \varepsilon^2 \qquad \text{for } l = 1,\dots,n-1, \tag{9.72}$$

where $\varepsilon < \min_i v_i$. Note that the first two constraints ensure that the capsule $\mathcal{C}(\boldsymbol{u}^l[l], \boldsymbol{u}^{l+1}[l], r_l)$ lies in $int(F(1,\dots,l)) \cap \mathbb{R}^l_{++}$.

Let m' be the minimum any player receives in any point in any of the capsules, and let $m = \min(m', \min_i v_i - \varepsilon) > 0$. Hence, m is a lower bound on any player's continuation payoff at any point on the pre-entry path. Now define N, which we will use as the number of punishment periods, such that

$$N = \left\lceil \frac{M}{m} \right\rceil. \tag{9.73}$$

Next we turn our attention to permissible patterns of discount factor vectors. We start by specifying the k_i's. This is done recursively using part b) of the Capsule Lemma. Set $k_1 = 1$. Next, for any l, assuming that we already know $k_1, \dots, k_l$ we will determine k_{l+1}. Apply the Capsule Lemma to the j'th capsule $\mathcal{C}(\boldsymbol{u}^j[j], \boldsymbol{u}^{j+1}[j], r_j)$where $j \leqslant l$. Using the notation from that result let $\bar{\theta}'_{1j}$ be $\theta^{\ddagger}(k_1, \dots, k_j)$. With $T^j(\theta)$ being the number of periods needed to execute the procedure described there, let $\{\boldsymbol{c}^{(1)}(j,\theta), \cdots, \boldsymbol{c}^{(T^j(\theta))}(j,\theta)\}$ be the vertices *in the original game* to be played

[35] As in the proofs of Theorems 1 and 2, the purpose of this adjustment is to make punishers indifferent on the various pure actions in the support of any mixed strategy that may be needed to minmax a deviating player.

to carry out the procedure.[36,37] Now, we assert that there exists a $\bar{\theta}_{1jl} \leqslant \bar{\theta}'_{1j}$ and a $\bar{k}_{l+1}$, such that if $k_{l+1} < \bar{k}_{l+1}$ and $\theta < \bar{\theta}_{1jl}$, the maximum absolute difference between player l's continuation payoff at the beginning of the procedure compared to that at the end of the procedure is ε/n. To see this recall that for any fixed k_{l+1}, $v_{l+1}^{t+1}(\theta) - c_{l+1}^{(t)} = (1 + k_{l+1}\theta)(v_{l+1}^{t}(\theta) - c_{l+1}^{(t)})$ where $\boldsymbol{c}^{(t)}$ is the vertex played and $v^t(\theta)$ is player l's continuation payoff for the t'th period during the operation. We can rewrite this as

$$v_{l+1}^{t+1}(\theta) - v_{l+1}^{t}(\theta) = (k_{l+1}\theta)(v_{l+1}^{t}(\theta) - c_{l+1}^{(t)}) \tag{9.74}$$

Hence,

$$\begin{aligned} |v_{l+1}^{t+1}(\theta) - v_{l+1}^{t}(\theta)| &= (k_{l+1}\theta)|(v_{l+1}^{t}(\theta) - c_{l+1}^{(t)})| \\ &\leqslant (k_{l+1}\theta)2M \end{aligned} \tag{9.75}$$

from which it follows that the absolute difference between beginning and end payoffs during stage j for player l is at most $2k_{l+1}M\theta T^j(\theta)$ the θ-dependent part of which is bounded by the expression

$$\theta\left(\frac{\alpha_j}{r_j - \sqrt{m_{1j}\theta^2 + m_{2j}\theta + r_j}} + 1\right),$$

with α_j being $d(\boldsymbol{u}^j[j], \boldsymbol{u}^{j+1}[j])$ and the m's are constants depending on the capsule specification, $k_1, \ldots, k_l$ but not on k_{l+1}. The limit of the above expression as θ tends to 0, using L'Hospital's rule is the constant $\frac{2\alpha_j\sqrt{r_j}}{-m_{2j}}$ and hence choosing $\bar{k}_{l+1} < \frac{-m_{2j}\varepsilon}{4\alpha_j nM\sqrt{r_j}}$ suffices for the assertion. From the above it is clear if $k_{l+1} < \bar{k}_{l+1}$ and $\theta < \bar{\theta}_{1l} := \min_{j\leqslant l}\bar{\theta}_{1jl}$ after l stages, i.e. after $T^1(\theta)+\cdots+T^l(\theta)$ periods, player $l+1$'s continuation payoff could not change by more than ε from its original target value v_l. If $\bar{\theta}_1 := \min_{1\leqslant l\leqslant n-1}\bar{\theta}_{1l}$, then given the $\boldsymbol{k}$ vector we have chosen, for $\theta < \bar{\theta}_1$, the above statement is true for each player.

For the chosen $\boldsymbol{k}$ vector, from Fact 1 we know that there is another positive bound $\bar{\theta}_2$ such that when $\theta < \bar{\theta}_2$, for any $\boldsymbol{x} \in B(\boldsymbol{u}^n, r_n + 3\sqrt{n-1}\Delta)$ there is a SPNE strategy $\sigma^*(\boldsymbol{x})$ that supports and hence, realizes $\boldsymbol{x}$. Hence, for $\theta < \min(\bar{\theta}_1, \bar{\theta}_2)$ we have designed a path that realizes $\boldsymbol{v}$. This path involves playing the following sequence of vertices along its pre-entry segment:

$$\boldsymbol{c}^{(1)}(1,\theta),\cdots,\boldsymbol{c}^{(T^1(\theta))}(1,\theta),\boldsymbol{c}^{(1)}(2,\theta),\cdots,\boldsymbol{c}^{(T^2(\theta))}(2,\theta),\cdots,\boldsymbol{c}^{(1)}(n-1,\theta),\cdots,\boldsymbol{c}^{(T^{n-1}(\theta))}(n-1,\theta)$$

followed by the path yielded by $\sigma^*(\boldsymbol{v}^{T(\theta)})$ with $T(\theta)$ being $T^1(\theta)+\cdots+T^{n-1}(\theta)$. For notational ease we will henceforth refer to the sequence of vertices on the pre-entry path simply as $\tilde{\boldsymbol{c}}^{(1)}, \ldots, \tilde{\boldsymbol{c}}^{(T(\theta))}$.

Now, we can formally describe the equilibrium strategy in the language of automata (Rubinstein 1986) as shown below. There are three types of (common) states, each identifed by a set of state variables:

- A$[\tau, i, \boldsymbol{z}]$ where $1 \leqslant \tau \leqslant T(\theta)$, $0 \leqslant i \leqslant n$, $\boldsymbol{z} \in \mathbb{R}^n$
- B$[\tau, \tau', i, \boldsymbol{r}]$ where $1 \leqslant \tau \leqslant T(\theta)$, $1 \leqslant \tau' \leqslant N$, $1 \leqslant i \leqslant n$, $\boldsymbol{r} \in \mathbb{R}^n$
- C$[\boldsymbol{x}]$ where $\boldsymbol{x} \in \mathbb{R}^n$

[36] It is important to note that these are n-dimensional vertices. Though the capsule is in the projected space of players' payoffs on the first j coordinates, playing a vertex in the projected space requires the participation of *all* players.

[37] Note that the starting payoff for the first capsule is $\boldsymbol{u}^1[1]$. Given a θ, the vertices to be played in each stage are now recursively defined.

The interpretation of an A-type state is that going forward, we have τ periods left of going through the pre-entry path, i was the last deviator (if $i = 0$, no deviation ever took place), and $\boldsymbol{z}$ is the adjustment vector in the ball $B(\boldsymbol{u}^n, r_n + 3\sqrt{n-1}\Delta)$ that we will need to subtract from the equilibrium point of entry (besides giving a 'reward' to player(s) $j \neq i$) once we are done with the pre-entry path. The interpretation of a B-type state is that we are on a punishment path where player i, the last deviator is being minmaxed and τ' periods of minmaxing still needs to be done while τ denotes from what type of A state we have (eventually) arrived here, and the j'th component of $\boldsymbol{r}$ denotes the normalized payoff for $j \neq i$ based on the past realizations of the $N - \tau'$ periods of minmaxing i ($r_i = 0$).[38] The interpretation of a Type C state is that it is an 'absorbing' state where $\sigma^*(\boldsymbol{x})$ is played from that point onwards.

The game starts at the state $\mathrm{A}[T(\theta), 0, \mathbf{0}]$. For any state $\mathrm{A}[\tau, i, \boldsymbol{z}]$, $\tilde{\boldsymbol{c}}^{(T(\theta)-\tau+1)}$ is to be played next. If in the observed action profile, there is a unilateral deviation by player j, play switches to the state $\mathrm{B}[\tau, N, j, \mathbf{0}]$. Otherwise, play switches to $\mathrm{A}[\tau-1, i, \boldsymbol{z}]$ if $\tau > 1$ and to $\mathrm{C}[\boldsymbol{v}^{T(\theta)}]$ if $\tau = 1$ and $i = 0$ and to $\mathrm{C}[\boldsymbol{v}^{T(\theta)} - \boldsymbol{z} + 2\Delta(\boldsymbol{\iota} - \boldsymbol{e}_i)]$ if $\tau = 1$ and $i \neq 0$. For any state $\mathrm{B}[\tau, \tau', i, \boldsymbol{r}]$, $\boldsymbol{m}^i$ is to be played next. If j is the only player whose action is not observed to be in support of $\boldsymbol{m}^i_j$, play switches to $\mathrm{B}[\tau, N, j, \mathbf{0}]$. Otherwise, if $\tau' > 1$, play next moves to $\mathrm{B}[\tau, \tau'-1, i, \tilde{\boldsymbol{r}}]$ where

$$\tilde{r}_j = \begin{cases} \frac{(1+\delta_j+\cdots+\delta_j^{N-\tau'-1})r_j + \delta_j^{N-\tau'} g_j(\boldsymbol{a})}{1+\delta_j+\cdots+\delta_j^{N-\tau'}} & \text{if } j \neq i \\ 0 & \text{if } j = i \end{cases} \tag{9.76}$$

with $\boldsymbol{a}$ being the last action profile observed. If on the other hand, $\tau' = 1$, play switches to $\mathrm{A}[\tau, i, \boldsymbol{z}]$ where $z_j = \frac{1-\delta_j^N}{\delta_j^{N+\tau}} \tilde{r}_j$, where $\tilde{r}_j$ is defined above. The behavior of the automata at a C type stage has already been described.

We need to ensure that $|z_j|$ is suitably bounded otherwise the adjustment term could take us out of the last ball. Note that $|z_j| \leqslant \frac{1}{\delta_j^{T(\theta)}} \frac{1-\delta_j^N}{\delta_j^N} M$. The limit of ${\delta_j}^{T(\theta)}$ (as θ goes to 0) can be written as $\lim_{\theta \to 0} {\delta_j}^{T_1(\theta)} \times \cdots \times \lim_{\theta \to 0} {\delta_j}^{T_{n-1}(\theta)}$. It may be easily checked that that if $\lim_{\theta \to 0} \theta T^l(\theta)$ is some constant, say b_l (as we have shown previously), then $\lim_{\theta \to 0} {\delta_j}^{T_l(\theta)} = e^{-k_j b_l}$, and hence, there exists a bound $\bar{\theta}_3$, such that if $\theta < \bar{\theta}_3$, $|z^i_j| < \Delta$. Given that each 'punisher' is rewarded by the amount 2Δ, and $||\boldsymbol{v}^{T(\theta)} - \boldsymbol{u}^n|| \leqslant r_n$, this shows that $||\boldsymbol{v}^{T(\theta)} - \boldsymbol{z} + 2\Delta(\boldsymbol{\iota} - \boldsymbol{e}_i) - \boldsymbol{u}^n|| \leqslant r_n + 3\sqrt{n-1}\Delta$.

Next, we examine the requirements for incentive compatibility. As σ^* is an SPNE by construction, we only need to check for incentive compatibility at A and B type states.

If the current state is $\mathrm{A}[\tau, 0, \mathbf{0}]$ or $\mathrm{A}[\tau, i, \boldsymbol{z}]$ for some $\boldsymbol{z}$, and player i did not deviate, the game will follow a certain path and he will receive a certain stage game payoff sequence. Let his normalized payoff from this be y_i. If he unilaterally deviates, he receives 0 for the next N periods and thereafter, he will receive exactly the same sequence of payoffs had he not deviated at all (note that he receives neither a reward nor an adjustment post entry). Unimprovability from the prescription at a Type A state in this case is then ensured by

$$(1 - \delta_i)M + {\delta_i}^{N+1} y_i \leqslant y_i \tag{9.77}$$

or, since $y_i \geqslant m$ by

$$M/m \leqslant 1 + \delta_i + \cdots + \delta_i^N \tag{9.78}$$

which equation (9.73) assures us will hold for high enough δ_i and hence θ below a certain bound.

[38] Note: $\boldsymbol{r} \neq \boldsymbol{z}$. The latter will depend on τ.

Our next bound $\bar{\theta}_4$ is precisely this bound.

The argument for i's unimprovability from states of the form $\mathrm{A}[\tau, j, \boldsymbol{z}]$, $(j \neq i)$ is even stronger than the argument given in favor of unimprovability in the just argued case, because by sticking to the equilibrium prescription, player i would have received y_i in normalized payoff plus at least an extra amount of Δ forever after τ periods.

Next consider i's incentive to deviate from a state of the form $\mathrm{B}[\tau, \tau', i, \boldsymbol{r}]$. This can't be profitable because it will simply postpone playing the same path that has a strictly positive (normalized) payoff.

Lastly, consider the prospect of i deviating from a state of the form $\mathrm{B}[\tau, \tau', j, \boldsymbol{r}]$, $j \neq i$. The design of the adjustment term makes him indifferent between playing any of the vertices on the support of $\boldsymbol{m}^i$ as can be easily verified in the usual manner. Hence, the only thing that remains to consider is his deviation to an action not in the support of $\boldsymbol{m}^i$. If equilibrium prescription is followed, i will receive at worst,

$$(1 - \delta_i^{\tau'}). - M + \delta_i^{\tau'}(y_i + \delta_i^{\tau}\Delta) \tag{9.79}$$

where y_i is his continuation (equilibrium path) payoff at $\mathrm{A}[\tau, 0, \mathbf{0}]$. if he deviates he will receive at best

$$(1 - \delta_i)M + \delta_i^{N+1}y_i \tag{9.80}$$

Hence the difference in i's payoff between conforming and deviating is at least

$$\begin{aligned} &(1 - \delta_i^{\tau'}). - M - (1 - \delta_i)M + (\delta_i^{\tau'} - \delta_i^{N+1})y_i + \delta_i^{\tau'+\tau}\Delta \\ &\geqslant (1 - \delta_i^{N}). - M - (1 - \delta_i)M + (\delta_i^{N} - \delta_i^{N+1})m + \delta_i^{N+T(\theta)}\Delta \end{aligned} \tag{9.81}$$

The last term in the above expression converges to the positive number $e^{-k_i b}\Delta$ (where $b = \sum_{l=1}^{n-1} b_l$) while other terms go to 0 as δ_i goes to 1. Hence, there exists a positive bound $\bar{\theta}_5$, such that if $\theta < \bar{\theta}_5$, there is no incentive for i to deviate at any $\mathrm{B}[\tau, \tau', j, \boldsymbol{r}^j]$ type state. We can now conclude that for the chosen $\boldsymbol{k}$ vector if $\theta < \min(\bar{\theta}_1, \bar{\theta}_2, \bar{\theta}_3, \bar{\theta}_4, \bar{\theta}_5)$, the prescribed strategy is an SPNE. ■

10 Appendix 2: LP 1999, Intertemporal Tradeoffs and the Realizability Problem

In this appendix we contrast our findings with those of LP 1999 regarding realizability, paying special attention to realizable payoff vectors that are also Pareto-optimal. As noted before, that paper considers only two-player games and assumes that PRDs exist. LP assume that Player i values \$1 received at the end of one time unit as \$$\bar{\delta}$ now, and fixes these $\bar{\delta}$'s. Next, assuming that the interactions take place Δ time units apart, they first characterize the realizable set and then examine that set as Δ becomes small. This is equivalent to fixing the time interval between interactions to one, and letting the discount factor for i to be $\delta_i = \bar{\delta}_i^{\Delta}$, which in turn amounts to fixing the log-discount ratios in our set-up. Our Theorems 3 and 4 impose no such restrictions on discount factors.

LP's technique of characterizing the realizable set for a given Δ (and also the limiting set) consists of characterizing its extreme points: they assert that these are normalized repeated game payoff pairs (u_1, u_2) which maximize $\lambda_1 u_1 + \lambda_2 u_2$ for a given direction vector $\boldsymbol{\lambda}$ within the realizable set. This immediately allows them to show *how* a certain payoff vector on the boundary of the realizable set is to be realized. As an illustration, re-consider the two-player game in Figure 5. To realize $\boldsymbol{\gamma}$, which gives high payoffs to both players, assuming $\bar{\delta_1} < \bar{\delta_2}$ (i.e. Player 2 is more patient than Player 1), LP's method would suggest playing $\boldsymbol{c}^4$ (Player 1's most preferred vertex) for a certain number of periods (say T_1), followed by a possibly mixed action once (with weights on $\boldsymbol{c}^4$ and $\boldsymbol{c}^5$ only), then $\boldsymbol{c}^5$ for another few periods (say T_2), followed by another possibly mixed action once (with weights on $\boldsymbol{c}^5$ and $\boldsymbol{c}^6$ only) and finally $\boldsymbol{c}^6$ (Player 2's most preferred vertex) forever. These paths signal some kind of intertemporal tradeoff; as if the more patient player is 'loaning' to the less patient player at the start and then getting paid back later.

It is instructive to see the mathematical foundation behind LP's prescriptions. Allowing for mixing in every period, let a path be defined by the sequence $\alpha^t(.)$, $t = 0, 1, \ldots$ where $\alpha^t : A \mapsto \mathbb{R}_+$ with $\sum_{\boldsymbol{a} \in A} \alpha^t(\boldsymbol{a}) = 1$. LP make the point that when we choose a path to maximize $\lambda_1 u_1 + \lambda_2 u_2 = \sum_{i=1,2} \lambda_i \sum_{t=0}^{\infty} (1 - \delta_i)\, {\delta_i}^t \sum_{\boldsymbol{a} \in A} \alpha^t(\boldsymbol{a}) g_i(\boldsymbol{a})$, all we need to do is to pick that α^t in period t which maximizes $\sum_{i=1,2} \lambda_i (1 - \delta_i)\, {\delta_i}^t \sum_{\boldsymbol{a} \in A} \alpha^t(\boldsymbol{a}) g_i(\boldsymbol{a})$. Next, suppose we are given positive λ_i's and $\bar{\delta_2} > \bar{\delta_1}$. Hence, for any Δ, as t increases, the direction vector $(\lambda_1(1 - \delta_1){\delta_1}^t, \lambda_2(1 - \delta_2){\delta_2}^t)$ starts moving anti-clockwise, with the second term of the vector eventually dominating the first, which immediately provides a justification of the proposed path(s). Note that *nothing in this argument requires players to be absolutely patient* (*i.e.* Δ *to be small*); the argument is supposed to hold for any set of unequal discount factors. Also note that if one or more of the λ_i's are negative, a similar argument is applicable. For example, to maximize $\lambda_1 u_1 + \lambda_2 u_2$ with $\lambda_1 > 0, \lambda_2 < 0$, LP's suggested path will shift weight progressively over time from $\boldsymbol{c}^4$ to $\boldsymbol{c}^3$ to $\boldsymbol{c}^2$, eventually settling on the last vertex forever.

For the rest of this appendix, we will generalize LP's notion to any number of players and different randomization set-ups. For n players, by "LP-type paths" we will mean paths that maximize $\sum_{i=1}^{n} \lambda_i \sum_{t=0}^{\infty} (1 - \delta_i)\, {\delta_i}^t \sum_{\boldsymbol{a} \in A} \alpha^t(\boldsymbol{a}) g_i(\boldsymbol{a})$, for some non-zero vector $\boldsymbol{\lambda} \in \mathbb{R}^n$. When PRDs are not available, and only independent mixing is allowed, we will impose the additional restriction that for each t and each $\boldsymbol{a} \in A$, $\alpha^t(\boldsymbol{a}) = \prod_{i=1}^{n} \alpha_i^t(a_i)$ where for each i, $\alpha_i^t : A_i \mapsto \mathbb{R}_+$ with $\sum_{\boldsymbol{a}_i \in A} \alpha_i^t(a_i) = 1$.

We now make two observations about LP-type paths.

First note that the economic intuition behind using these paths for realizing efficient outcomes does not extend to three or more players. Suppose $\boldsymbol{\lambda}$ is a strictly positive vector when obviously, all along an LP-type path, only vertices on the stage game Pareto frontier will be played. For two-player games, since on that frontier, one player's preference ranking over vertices is *exactly* the reverse of the other player's ranking, along such paths, one player (the impatient one) does get progressively worse payoffs while the other gets better payoffs. But this makes it clear that the idea of 'loaning' hinges on having two players, because the same argument cannot be made for three or more players: a vertex that gives the most impatient player his highest possible amount does not necessarily offer the second-most impatient player more than it does to the third-most impatient player. Moreover, while for two players with unequal discount factors, it is never possible that an LP-type path will result in a play of vertex $\boldsymbol{c}$ in some period t, then a play of a different vertex $\boldsymbol{c}'$ in period $t+1$ and then again a play of $\boldsymbol{c}$ in period $t+2$, this can happen for three or more players; the weights on various vertices along LP-type paths are not guaranteed to move 'monotonically'. For example, let F have these four vertices: $(0, 0, 0)$, $(1, 1, .5)$, $(20, 20, 20)$ and $(4.7, 52.3, 2.9)$. Let $\boldsymbol{\delta} = (0.7, 0.8, 0.9)$. For $\boldsymbol{\lambda} = (\frac{10}{3}, 5, 10)$, the reader can verify that in $t = 0$ and 2, $(20, 20, 20)$ is to

[39]This is not an original point; LP themselves were aware of it (see subsection 5.4 of their paper, where after discussing the problem they conclude: "...Therefore, our method fails in the n-player case.").

be played, while for $t = 1$, $(4.7, 52.3, 2.9)$ is to be played.[39]

A second feature of LP-type paths is that they could be bottom-heavy for some players, in that play will eventually settle on one of the vertices of F and that vertex could potentially result in a very low payoff for those players. From the point of view of the economist attempting to design stable cooperative agreements, this could be undesirable.

Next, we offer a detailed contrast between LP's approach and ours with regards to two different questions that may be raised regarding the realizability problem. The first questions is: Suppose a point $\boldsymbol{v} \in int(re(F) \cap F^c$ is given. Is it possible to describe a set of discount factor vectors for which $\boldsymbol{v}$ is realizable and if so, via which path? We will call this the Folk Theorem question. Another, somewhat more pragmatic question is: Suppose both the target point $\boldsymbol{v}$ and the discount factor vector $\boldsymbol{\delta}$ are given. Is it possible to realize $\boldsymbol{v}$ given $\boldsymbol{\delta}$, and if so, via what sort of path? We will call this the fixed $\boldsymbol{\delta}$ question.

Even when we assume the existence of PRDs, the LP approach has difficulty answering the Folk Theorem question, particularly for $n \geqslant 3$, because their approach works by fixing relative patience parameters. Once that is fixed their approach will 1) work out the boundary of the realizable set for a given Δ (or as Δ goes to 0) and then 2) figure out two boundary points such that the target point is a convex combination of these two points and then 3) use the same convex combination of the paths needed to realize the boundary point. But for arbitrary relative patience configurations, the target may not belong to the realizable set even as Δ goes to 0. Further, for more than two players, the relative patience configurations that are needed to realize a certain target as Δ goes to 0 can be quite complex to derive, a question that the LP approach does not shed much light on.

If we do not have PRD's, the LP approach becomes inapplicable because characterizing a set by pinning down its extreme points works for a convex set, but not for non-convex sets; for such sets this method may miss many efficient payoff vectors, not to mention the fact that for a non-convex set characterizing just the boundary of the set does not characterize the full set. Convexity of the realizable set for any discount factor vector is trivial when PRDs are allowed, but from Counterexample 1 in Section 5 we know that even for large discount factors, convexity fails without PRDs. If additionally we fix relative patience of players either the way LP do (by letting $\delta_i = \bar{\delta_i}^{\Delta}$ while varying Δ) or the way Sugaya and we do (by letting $\delta_i = \frac{1}{1+k_i\theta}$ while varying θ), and let players become more and more absolutely patient, this failure still cannot be remedied. Specifically, examining equation 5.5 in Counterexample 1 makes it clear that there is no $\bar{\Delta}$ such that $\mathcal{F}(\bar{\delta_1}^{\Delta}, \bar{\delta_2}^{\Delta})$ is convex if $\Delta < \bar{\Delta}$. Hence, the validity of LP's method of characterizing the limiting set of realizable payoffs by characterizing its boundary does not extend to the PRD-less situation.

Our approach, on the other hand, which does not rely on whether PRDs are present or not, is geared exactly to answer the Folk Theorem question. The proof of Theorem 3 shows that we can choose the vertices to be played on the pre-entry path in many different ways, along with the ball inside F where the continuation payoff will eventually settle, and for each such choice, we can provide many relative patience configurations (i.e. $\boldsymbol{k}$ vectors) and an absolute patience parameter bound (i.e. a $\bar{\theta}$) that are consistent with the realization of a given $\boldsymbol{v}$. A suite of three optimization problems for doing this is available from the authors on request.

Turning now to the fixed $\boldsymbol{\delta}$ question, assuming PRDs, the LP approach works very nicely with $n = 2$, but becomes quite complicated for more than two players. Because the monotonicity of vertex weights along the LP-type paths is now lost, it is not easy to characterize the boundary of the realizable set; for every $\boldsymbol{\lambda}$ and every t an optimization problem is to be solved and the whole boundary can be fully described only when this solution is known for every possible $\boldsymbol{\lambda}$.

Furthermore, when PRDs are unavailable, and we are dealing with fixed, low discount factors, the non-convexity of the realizable set not only makes the first step of LP's mathematical argument invalid, (that every efficient point is an argmax of a linear function of the repeated game payoffs), but the conclusion becomes invalid as well (that every such payoff *must* be realizable via some LP-type path). Alternate vertex switching (AVS) paths where switching among vertices takes place *every* period can be efficient for many situations where all players are almost equally myopic. To illustrate, the following holds:

Proposition A. 1) *For the stage game described in Counterexample 1, when PRD's are disallowed but* (*independent*) *mixing is not, there exists an open ball D of discount factor vectors with center* $(\underline{\delta}, \underline{\delta})$ *for some* $\underline{\delta} \in (0,1)$ *such that when* $\boldsymbol{\delta} \in D$*, an AVS path where players play the vertices* $(1,0)$ *and* $(0,1)$ *in alternate periods, starting with either* $(1,0)$ *or* $(0,1)$*, is Pareto-efficient and* 2) *for any* $\boldsymbol{\delta} \in D$*, where* $\delta_1 \neq \delta_2$*, the payoff vector generated by such a path cannot be realized by any LP-type path.*

The intuition behind the result is as follows. With low discount factors, an AVS path cannot be Pareto-improved by converting a pure-action payoff on its path, such as $(1,0)$, to a different pure-action payoff such as $(0,1)$ because the lumpiness of the current loss may not be remediable even by offering the sacrificing player 'everything' later. If PRDs are available, the patient player might be persuaded to release a small amount of payoff to his impatient counterpart in return of getting paid back later (with interest), leading to a Pareto-improvement. However, without PRDs, small changes to a particular period's utility can be effected by independent mixing only, which is 'wasteful' in the sense that it reduces the total utilities players could have in that period. If the discount factors are very close, the prospective gain from intertemporal tradeoff becomes small and is outweighed by this 'waste', thus making Pareto-improvements impossible.

How does our approach work for the fixed $\boldsymbol{\delta}$ question? Unfortunately, not very well either. It is possible to write a Mixed Integer NLP that can search if there exists a sequence of vertices to be used along the pre-entry path and a compatible ball of entry, but admittedly, it is neither simple to solve nor does it settle the fixed $\boldsymbol{\delta}$ question completely, since there can be many different ways of entry into $int(F)$ - not just via the paths used in the proof of Theorem 3. If the target point happens to be Pareto optimal (which a priori, without PRDs, there is no simple way of determining[40]), the paths suggested by our approach may never be able to realize it exactly, since, once inside $int(F)$, the path may feature the play of an inefficient vertex. The fixed $\boldsymbol{\delta}$ question, therefore, remains a computationally challenging open problem, worthy of future investigations.

Proof of Proposition A

1) The path where players alternate between the two vertices $(0,1)$ and $(1,0)$ starting with $(1,0)$ will be hereafter called AVS(1), and when the starting vertex is $(0,1)$, we will call that path AVS(2).

When players mix in a period, if one of the players gets the (undiscounted) utility x, the other player can get at most $f(x) := 1 + x - 2\sqrt{x}$.[41] Henceforth, without loss of generality, we will focus attention on those paths where in each period the player's utilities are the pair $(x, f(x))$ for some $x \in [0,1]$. For future reference, we note the following four properties of f:

[40] As was noted in the introduction, even for a two-player game with only two vertices in the stage game, Salonen and Vartiainen (2008) show how 'irregular' the Pareto frontier can be.

[41] $f(x)$ is the maximum of $(1-p)(1-q)$ subject to $pq \geqslant x$ and $p, q \in [0,1]$.

a) For $x \in (0,1)$, $f' < 0$ and $f'' > 0$,
b) $x + f(x)$ is a strictly convex function which evaluates to 1 at 0 and 1 only within the interval [0,1] with $f'(.25) = 0$,
c)

$$\lim_{x \to 1} \frac{f(x)}{(1-x)^2} = .25,$$

d)

$$f^{-1}(.) = f(.).$$

Step 1: First we construct D. Choose any $\bar{\delta}$ be such that $c := (1 - \bar{\delta}(1+\bar{\delta}))(1-\bar{\delta}^2) > .25$. Indeed, this condition may be satisfied by choosing $\bar{\delta}$ small. Note that it implies the following condition: $(1 - \bar{\delta}(1+\bar{\delta})) > .25$. Let $\varepsilon_1 = c - .25$ and let r_1 be such that if $\boldsymbol{\delta} \in B(\bar{\delta}\boldsymbol{\iota}, r_1)$, the following hold:

$$1 - \delta_i\,(1+\delta_i) > .25 \qquad \text{for} \quad i = 1,2 \tag{10.1}$$

$$\delta_i > \delta_j^2 \qquad \text{for} \quad i = 1,2,\ j \neq i \tag{10.2}$$

$$\frac{(1-\delta_i(1+\delta_i))(1-\delta_i^2)^2}{1-\delta_j^2} > c - .5\varepsilon_1 \qquad \text{for} \quad i = 1,2,\ j \neq i \tag{10.3}$$

$$\frac{\delta_i}{1-\delta_i^2} + f\left(\frac{\delta_j}{1-\delta_j^2}\right) < 1 \qquad \text{for} \quad i = 1,2,\ j \neq i \tag{10.4}$$

Clearly, a small enough r_1 can be chosen to satisfy the above given the choice of $\bar{\delta}$ and property b) of f. Because of property c) of f, there is an $\eta < .75$ such that

$$\frac{f(x)}{(1-x)^2} < .25 + .5\varepsilon_1 \quad \text{for } x > 1 - \eta. \tag{10.5}$$

Define ε_2 via the following:

$$(1-\eta) + f(1-\eta) = 1 - \varepsilon_2. \tag{10.6}$$

Let r_2 be such that when $\boldsymbol{\delta} \in B(\bar{\delta}\boldsymbol{\iota}, r_2)$

$$\frac{1}{1-\delta_i^2} + \frac{\delta_j}{1-\delta_j^2} - \frac{\delta_m}{1-\delta_m} > 1 - \varepsilon_2 \quad \text{for} \quad i = 1,2,\ j \neq i \tag{10.7}$$

where $\delta_m = \max(\delta_1, \delta_2)$. Since, the left hand side of the above tends to 1 as both δ_1 and δ_2 tend to $\bar{\delta}$, it is possible to satisfy (10.7) by choosing r_2 small. Let $r = \min(r_1, r_2)$. Then $D = B(\bar{\delta}\boldsymbol{\iota}, r)$.

Step 2. We show that for any (δ_1, δ_2) pairs satisfying a less stringent condition than (10.1), namely

$$1 > \delta_i(1+\delta_i) \quad \text{for} \quad i = 1,2 \tag{10.8}$$

both AVS(1) and AVS(2) are Pareto-optimal within the class of paths that employ pure actions only. To see this for AVS(1), note that playing the prescribed strategy offers Player 1 the amount $1 + \delta_1^2 + \delta_1^4 + \ldots$ in undiscounted payoff; if $(1,0)$ is not played at $t = 0$, even if it was played for all $t > 0$, he would earn at most $\delta_1 + \delta_1^2 + \delta_1^3 + \ldots$; which is less than his prescription payoff if

$\frac{1}{1-\delta_1^2} > \frac{\delta_1}{1-\delta_1}$ or $1 > \delta_1(1+\delta_1)$. Hence, if any other payoff vector Pareto-dominates the payoff from the prescribed path, it must also start with the play of $(1,0)$. Given this, an analogous argument from Player 2's perspective shows that $(0,1)$ must be played in $t=1$ and now this argument can be carried forward ad infinitum. A similar logic establishes the efficiency of AVS(2) within the limited class of pure action paths.

Step 3: Next, we consider (independent) mixing. In this step, focusing on AVS(1), we show that for discount factors obeying (10.8), if AVS(1) is not Pareto-optimal, then either AVS(1) or AVS(2) is Pareto-improvable by a path that involves strict mixing only in $t=0$. To see this consider the following optimization problem: Choose a sequence $\{x_t\}_{t=0}^{\infty}$ to maximize

$$\sum_{t=0}^{\infty} \delta_1^t x_t + \sum_{t=0}^{\infty} \delta_1^t f(x_t)$$

subject to

$$\sum_{t=0}^{\infty} \delta_1^t x_t \geqslant \frac{1}{1-\delta_1^2} \tag{10.9}$$

$$\sum_{t=0}^{\infty} \delta_2^t f(x_t) \geqslant \frac{\delta_2}{1-\delta_2^2} \tag{10.10}$$

$$x_t \in [0,1] \quad t = 0,1,2... \tag{10.11}$$

A standard compactness-continuity argument (in the product topology) shows that the problem has a solution. If AVS(1) is not Pareto-optimal, the existence of a solution to the above problem provides us with an alternate path that Pareto improves on it *and is itself Pareto-optimal.* Call it the challenger path.

Because of its Pareto optimality, we now claim that along the challenger path, a mixed action can be played in at most one period. Suppose, by contradiction, in period t, players receive the utility pair $(x_t^*, f(x_t^*))$ and in period $s > t$, they receive another utility pair $(x_s^*, f(x_s^*))$ where both $x_t^*, x_s^* \in (0,1)$. In the $x_t - x_s$ space, Player 1 has linear indifference curves with direction of increasing utility north and east, (his utility function, having fixed his payoffs in all other periods is $\delta_1^t x_t + \delta_1^s x_s$). Player 2 on the other hand has indifference curves that are convex to the origin, with the direction of increasing utility being south and west (his utility function is $\delta_2^t f(x_t) + \delta_2^s f(x_s)$, which is quasi-convex, a claim that can be verified using property a) of the function f stated earlier). Clearly a strictly mutually beneficial readjustment of the two actions can be organized if the two players' indifference curves are not tangent at (x_t^*, x_s^*), but even in the case of tangency, it is easy to verify pictorially that there will be room for Pareto-improvement by moving away from the point of tangency in either direction along Player 1's indifference curve. This proves the claim.

Suppose the challenger path involves mixed action only in period T. We could then, using exactly the same argument used in Step 1 conclude that as long as the condition (10.8) holds, the challenger path coincides with AVS(1) for all periods up to and including $T-1$. This in turn shows that there is a challenger path that involves playing a strictly mixed action on $t=0$ and pure actions thereafter that Pareto-improves either on AVS(1) or on AVS(2) (which it is depends on whether T is even or odd).

Step 4. We now will show that for $\boldsymbol{\delta} \in D$, AVS(1) could not be Pareto- improved by a challenger path that mixes only in $t = 0$ (the argument for AVS(2) being symmetric will be omitted). For this challenger path, let π_i^τ denote what Player i receives in unnormalized payoff on a path the t'th period action of which concides $t + \tau$'th period action of the challenger path. Then, there exists $x \in (0, 1)$ such that the following equations hold:

$$x + \delta_1 \pi_1^1 \geqslant \frac{1}{1 - \delta_1^2} \tag{10.12}$$

$$f(x) + \delta_2 \pi_2^1 \geqslant \frac{\delta_2}{1 - \delta_2^2} \tag{10.13}$$

with at least one of the inequalities being strict. If we add the two above inequalities we obtain:

$$\begin{aligned} x + f(x) &> \frac{1}{1 - \delta_1^2} + \frac{\delta_2}{1 - \delta_2^2} - (\delta_1 \pi_1^1 + \delta_2 \pi_2^1) && (10.14) \\ &\geqslant \frac{1}{1 - \delta_1^2} + \frac{\delta_2}{1 - \delta_2^2} - \delta_m (\pi_1^1 + \pi_2^1) && (10.15) \\ &\geqslant \frac{1}{1 - \delta_1^2} + \frac{\delta_2}{1 - \delta_2^2} - \frac{\delta_m}{1 - \delta_m} && (10.16) \\ &> 1 - \varepsilon_2, && (10.17) \end{aligned}$$

the last inequality following from (10.7). Now $x + f(x)$ can be above a certain value if either x is low ($<.25$, the argmin) or high (above .25). We rule out the first possibility by noting that $\pi_1^1 \leqslant \frac{1}{1-\delta_1}$, which paired with (10.12) gives us a lower bound on x: $\frac{1-\delta_1(1+\delta_1)}{1-\delta_1^2} > (1 - \delta_1(1 + \delta_1) > .25$, because of (10.1). Now, because of the definition of η in (10.6) allows us to conclude that $1 - x < \eta$.

The challenger path clearly must differ from AVS(1) not just in $t = 0$, but in some subsequent period as well. Let T be the *first time after* $t = 0$ when this change occurs. Taking advantage of the fact that the challenger path from T onward uses only pure actions, we can see that T can't be even because then we will have $x + \delta_1^T \pi_1^T \geqslant 1 + \delta_1^T \frac{1}{1-\delta_1^2}$, or equivalently

$$\begin{aligned} 1 - x &\leqslant \delta_1^T \left(\pi_1^T - \frac{1}{1 - \delta_1^2} \right) && (10.18) \\ &\leqslant \delta_1^T \left(\frac{\delta_1}{1 - \delta_1} - \frac{1}{1 - \delta_1^2} \right) && (10.19) \\ &\leqslant \delta_1^T \, \frac{\delta_1(1 + \delta_1) - 1}{1 - \delta_1^2} && (10.20) \\ &< 0, && (10.21) \end{aligned}$$

an impossibility. If on the other hand, T is odd, a similar argument provides a positive upper bound on $1 - x$:

$$\begin{aligned} 1 - x &\leqslant \delta_1^T \left(\pi_1^T - \frac{\delta_1}{1 - \delta_1^2} \right) && (10.22) \\ &\leqslant \delta_1^T \left(\frac{1}{1 - \delta_1} - \frac{\delta_1}{1 - \delta_1^2} \right) && (10.23) \\ &\leqslant \delta_1^T \, \frac{1}{1 - \delta_1^2}. && (10.24) \end{aligned}$$

Also, if T is odd, for the second player, we will have $f(x) + \delta_2^T \pi_2^T \geqslant 0 + \delta_2^T \frac{1}{1-\delta_2^2}$, which can be used to find a lower bound on $f(x)$:

$$f(x) \geqslant \delta_2^T \left(\frac{1}{1-\delta_2^2} - \pi_2^T \right) \tag{10.25}$$

$$\geqslant \delta_2^T \left(\frac{1}{1-\delta_2^2} - \frac{\delta_2}{1-\delta_2} \right) \tag{10.26}$$

$$\geqslant \delta_2^T \frac{1-\delta_2(1+\delta_2)}{1-\delta_2^2}. \tag{10.27}$$

Now combining (10.24) and (10.27), we observe that

$$\frac{f(x)}{(1-x)^2} \geqslant \left(\frac{\delta_2}{\delta_1^2} \right)^T \frac{(1-\delta_2(1+\delta_2))(1-\delta_1^2)^2}{1-\delta_2^2} \tag{10.28}$$

Because of (10.2), no matter what T is $\left(\frac{\delta_2}{\delta_1^2} \right)^T$ will exceed 1, and because of (10.3), $\frac{(1-\delta_2(1+\delta_2))(1-\delta_1^2)^2}{1-\delta_2^2}$ will exceed $c - .5\varepsilon_1$. On the other hand, because $1 - x < \eta$, and (10.5), $\frac{f(x)}{(1-x)^2} < .25 + .5\varepsilon_1$. Given the definition of ε_1, this creates a contradiction.

2) For fixed $\boldsymbol{\lambda}$, we will call a path that is a solution to the LP problem when $\delta_1 > \delta_2$ an LP(1) path; similarly an LP(2) path is one that solves LP's problem when $\delta_2 > \delta_1$. We show that neither an LP(1) path nor an LP(2) path can attain the payoffs realized by the AVS(1) path. The case of AVS(2) is similar and will be omitted.

Clearly, if an LP-type path was to realize the payoffs attained by either an AVS(1) or an AVS(2) path $\boldsymbol{\lambda}$ must be a strictly positive vector. Since in that case, the path must eventually settle on either $(1, 0)$ or $(0, 1)$ and since any path involving only pure actions cannot realize the payoffs obtained by AVS(1) without coinciding with AVS(1), for an LP path to attain AVS(1) payoffs there must be mixing involved; further given $\delta_1 \neq \delta_2$, such mixing can take place in at most one period.

Suppose an LP(2) path uses independent mixing for the first and only time on $t = T$; thereafter $(0, 1)$ is played. T can't be 0 because the maximum Player 1 can then receive on the LP(2) path (which is 1) is strictly less than what he was receiving in the AVS(1) path (which is $\frac{1}{1-\delta_1^2}$). Also, T can't be 2 because then the maximum Player 2 would receive in the PVS(1) path ($\frac{\delta_2^2}{1-\delta_2}$) is strictly less than what he receives in AVS(1) path ($\frac{\delta_2}{1-\delta_2^2}$), because of (10.8). Hence, $T = 1$, and there is some $x \in (0, 1)$ such that the following equations hold:

$$1 + \delta_1 x = \frac{1}{1-\delta_1^2} \tag{10.29}$$

$$0 + \delta_2 f(x) + \frac{\delta_2^2}{1-\delta_2} = \frac{\delta_2}{1-\delta_2^2} \tag{10.30}$$

Equation (10.29) gives $x = \frac{\delta_1}{1-\delta_1^2}$, while equation (10.30) can be solved for $f(x)$ to show $f(x) = 1 - \frac{\delta_2}{1-\delta_2^2}$. Together they imply

$$f\left(\frac{\delta_1}{1-\delta_1^2} \right) = 1 - \frac{\delta_2}{1-\delta_2^2}. \tag{10.31}$$

But this contradicts (10.4).

Next, suppose LP(1) path attains the payoffs realized by the AVS(1) path. Suppose the mixing period is T; for $t > T$, $(1, 0)$ is played. Then T has to be period 0, because we have already established that even if Player 1 gets 1 for all periods from $t = 1$ onward, his target still cannot be met. Hence, there exists $x \in (0,1)$ such that

$$x + \frac{\delta_1}{1-\delta_1} = \frac{1}{1-\delta_1^2} \tag{10.32}$$

$$f(x) = \frac{\delta_2}{1-\delta_2^2} \tag{10.33}$$

Equation (10.32) shows $x = 1 - \frac{\delta_1}{1-\delta_1^2}$, while equation (10.33) upon using property d) of the f function shows that $x = f\left(\frac{\delta_2}{1-\delta_2^2}\right)$. Hence,

$$f\left(\frac{\delta_2}{1-\delta_2^2}\right) = 1 - \frac{\delta_1}{1-\delta_1^2}. \tag{10.34}$$

which again violates (10.4). This completes the proof. ■

References

[1] ABREU, D., D. PEARCE and E. STACCHETTI (1990): "Towards a Theory of Discounted Repeated Games with Imperfect Monitoring," Econometrica, 58 (5), 1041-1063.

[2] AUMANN, R., and M. MASCHLER (1995): "Repeated Games with Incomplete Information," with the collaboration of R. E. Stearns, MIT Press, Cambridge, MA.

[3] CHEN, B. and S. FUJISHIGE (2013): "On The Feasible Payoff Set of 2-player Repeated Games with Unequal Discounting," International Journal of Game Theory, 42, 295-303.

[4] CHEN, B. and S. TAKAHASHI (2012) : "A Folk Theorem for Repeated Games with Unequal Discounting," Games and Economic Behavior, 76 (2), 571-581.

[5] FENCHEL, W (1929): "Uber Kriimmung und Windung Geschlossener Raumkurven," Math. Ann., 101, 238-252.

[6] FUDENBERG, D. and E. MASKIN (1986): "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," Econometrica, 54 (3), 533-554.

[7] FUDENBERG, D. and E. MASKIN (1991): "On the Dispensability of Public Randomization in Discounted Repeated Games," Journal of Economic Theory, 53, 428-438.

[8] HAAG, M. and R. LAGUNOFF (2007) : "On the Structure and Size of Group Cooperation", Journal of Economic Theory 135, 68-89.

[9] HARRINGTON, J. (1989): "Collusion among Asymmetric Firms: The Case of Different Discount Factors," International Journal of Industrial Organization, 7(2), 289-307.

[10] HÖRNER, J. and W. OLSZEWSKI (2005): "The Folk Theorem for Games with Private Almost-Perfect Monitoring", Econometrica, 74(6), 1499–1544.

[11] HÖRNER, J. and S. TAKAHASHI (2016): "How Fast Do Equilibrium Payoff Sets Converge in repeated Games?," Journal of Economic Theory, Elsevier, Vol 165(c), 332-359.

[12] LEHRER, E. and A. PAUZNER (1999): "Repeated Games with Differential Time Preferences," Econometrica, 67, 393-412.

[13] MAILATH, G. J., I. OBARA and T. SEKIGUCHI (2002): "The Maximum Efficient Equilibrium Payoff in the Repeated Prisoner's Dilemma", Games and Economic Behavior, 40 (1), 99-122.

[14] MAILATH, G. and L. SAMUELSON (2006): "Repeated Games and Reputations: Long-Run Relationships," Oxford.

[15] OBARA, I and F. ZINCENKO (2017): "Collusion and Heterogeneity of Firms", Rand Journal of Economics, 48, 230-249.

[16] OLSZEWSKI, W. (1998): "Perfect Folk Theorems. Does Public Randomization Matter?," International Journal of Game Theory, 27, 147 - 156.

[17] RUBINSTEIN, A. (1986): "Finite Automata Play the Repeated Prisoners' Dilemma", Journal of Economic Theory, 39 (1), 83-96.

[18] SALONEN, H. and H. VARTIAINEN (2008): "Valuing Payoff Streams Under Unequal Discount Factors," Economic Letters, 99, 595-598.

[19] SORIN, S. (1986): "On Repeated Games with Complete Information," Math. Oper. Res., 11 (1), 147-160.

[20] SUGAYA, T. (2015): "Characterizing the Limit Set of PPE Payoffs with Unequal Discounting", Theoretical Economics, 10(3), 691-717

[21] YAMAMOTO, Y. (2010): "The Use of Public Randomization in Discounted Repeated Games," International Journal of Game Theory, 39, 441-443.