STABLE PAYOFFS WITH ENDOGENOUS SEPARATION

SEGISMUNDO S. IZQUIERDO
Department of Industrial Organization and BioEcoUva, Universidad de Valladolid

LUIS R. IZQUIERDO
Department of Management Engineering, Universidad de Burgos

ALEJANDRO GUTIÉRREZ-MIELGO
Institute of Mathematics, Universidad de Valladolid

Costless separation allows players to abandon established relationships and restart with anonymous partners. We study finite games with separate populations for each role and strategies depending only on the current partnership's history. A player's opportunity to exploit new partners and restart imposes a payoff threshold. As exogenous separation vanishes, we characterize the limiting payoffs of pure path-protecting profiles, which strictly deter unilateral departures from their prescribed path. They form the closure of feasible payoffs strictly above every role's threshold for some common initial action profile. For two roles, these thresholds reduce to pure minmax payoffs. With more roles, they can conflict. A three-player example has an interior, strictly individually rational payoff vector that remains at a positive distance from every Nash payoff, including polymorphic ones. The constructed profiles are neutrally stable and Lyapunov stable under stationary matching on every fixed finite strategy space containing them.

KEYWORDS: Endogenous separation, multipopulation games, neutral stability, path protection, folk theorem, anonymous rematching.

JEL CLASSIFICATION. C72, C73.

Segismundo S. Izquierdo: segismundo.izquierdo@uva.es
Luis R. Izquierdo: lrizquierdo@ubu.es
Alejandro Gutiérrez-Mielgo: alejandro.gutierrez@uva.es

We acknowledge financial support from the Spanish State Research Agency (PID2024-159461NB-I00/MICIU) and from the Regional Government of Castilla y León + EU-FEDER (CLU-2025-2-05 – BIOECOUVA Unit of Excellence and CLU-2025-1-02 – IMUVA, Universidad de Valladolid).

1. INTRODUCTION

Games with endogenous separation are repeated games in which players may leave their current partners and continue playing in newly formed partnerships (Datta, 1996, Rob and Yang, 2010, Mailath and Samuelson, 2006). Continued association is a choice rather than an externally imposed constraint. A central question is which behavioral conventions can remain stable when players can abandon a relationship and restart with anonymous partners, and which payoffs can be sustained in equilibrium.

The option to restart imposes a common constraint on the roles in a partnership. Suppose every new relationship begins with action profile a. A player in role k can best respond to the other roles' initial actions and immediately leave, repeating this strategy in every new match. Let $b_k(a)$ be the stage payoff from this deviation. A sustainable convention must compensate each role for this repeatable opportunity. The difficulty is that the same initial profile a must satisfy all roles' constraints: an action that lowers one role's restarting payoff may raise another's.

We study finite stage games with any number of roles, one nonatomic population per role, and anonymous rematching after either voluntary or exogenous separation. Strategies depend only on the history of the current partnership. Our main result characterizes the limiting payoffs of pure *path-protecting* profiles, which make every unilateral departure from their prescribed infinite path strictly unprofitable. As exogenous separation becomes rare, these payoffs form exactly the closure of the feasible vectors strictly exceeding $b_k(a)$ for every role k, for some common initial profile a. For two roles, the minimizations defining their pure minmax payoffs can be attained by one initial profile, so every feasible vector strictly above those bounds can be approximated. If such vectors exist and pure and mixed minmax payoffs coincide, the limits of all Nash payoffs are exactly those of pure path-protecting profiles (Corollary 16). With more roles, the restarting constraints can conflict.

This conflict can restrict all Nash payoffs, beyond the class of path-protecting profiles. We construct a three-player game with pure and mixed minmax vector $(1,1,1)$ and an interior feasible payoff vector $(3/2,3/2,3/2)$. Every Nash state, including every polymorphic state in which several strategies coexist within a role, gives some role an average payoff of at least 2. The target therefore remains at least $1/2$ away in the sup norm from every Nash payoff, for every continuation probability, and cannot be approximated by Nash payoffs as exogenous separation vanishes (Proposition 15). The feasible set is full dimensional, so the example meets the individual-rationality and dimensionality hypotheses of the fixed-partner folk theorem (Fudenberg and Maskin, 1986). The obstruction is the need to satisfy restarting incentives jointly, rather than those hypotheses.

The positive construction begins with a sufficiently long repetition of an initial profile a, followed by an optional finite intermediate phase and an indefinitely repeated finite pattern. The initial phase creates an opportunity cost of restarting: leaving carries no fee, but a new relationship begins before the higher-payoff phase has been reached. Its contribution to complying players' average payoffs vanishes as exogenous separations become rare, while voluntary restarting continues to reintroduce it. The initial profile need not be Nash, and the construction uses neither full dimensionality nor the nonequivalent-utilities condition of Abreu, Dutta, and Smith (1994). The same constructed profile works for every continuation probability above a suitable threshold.

The distinction between exact attainment and approximation matters. For example, a convention that starts every partnership with mutual cooperation in the Prisoner's Dilemma invites defection followed by immediate exit. An initial phase that deters this deviation can nevertheless allow cooperation payoffs to be approximated arbitrarily well as exogenous separation vanishes. Our two-role result permits this recovery of the usual asymptotic region provided that some feasible vector strictly exceeds both pure minmax payoffs and the pure and mixed minmax bounds coincide. The three-player example establishes the stronger failure of approximation even in the limit.

We also connect the constructed conventions to evolutionary stability. Endogenous separation makes the distribution of available partners differ from the distribution of strategies in the population: strategies associated with shorter partnerships are overrepresented among new matches. Under stationary matching, the tensor of expected partnership durations determines a unique consistent pool. The pool–population map is bi-Lipschitz on each fixed finite strategy space, giving Lipschitz payoff functions and well-defined multipopulation replicator dynamics. Path protection implies neutral stability, which in turn implies Lyapunov stability on every finite strategy space containing the resident profile. Allowing strategy frequencies to vary independently across roles can change stability even for a symmetric stage game.

The voluntarily repeated Prisoner's Dilemma has been studied over unrestricted strategy spaces (Carmichael and MacLeod, 1997, Fujiwara-Greve and Okuno-Fujiwara, 2009) and for particular strategy families (Vanberg and Congleton, 1992, Schuessler, 1989, Hayashi and Yamagishi, 1998, Izquierdo, Izquierdo, and Vega-Redondo, 2010, 2014). Initial phases of mutual defection or other costs followed by cooperation are central to this literature (Datta, 1996, Ghosh and Ray, 1996, Carmichael and MacLeod, 1997, Kranton, 1996). For symmetric two-player games, Fujiwara-Greve and Okuno-Fujiwara (2009) analyze stationary matching with endogenous relationship durations, including trust-building phases and neutral stability. We extend the payoff analysis to general finite games with separate role populations and identify the resulting compatibility constraints. The initial phase is related to starting small (Watson, 2002), but here it deters restarting without learning about partners' types or changing the stakes. Experimental and simulation studies also examine how the option to leave affects cooperation (Barclay and Raihani, 2016, Rand, Arbesman, and Christakis, 2011, Graser, Fujiwara-Greve, García, and van Veelen, 2025, Izquierdo, Izquierdo, and Boyd, 2026).

Our closest predecessor is Izquierdo and Izquierdo (2026), who develop a single-population framework for symmetric two-player games with endogenous separation, introduce path protection, and prove a folk theorem for that setting. They also discuss neutral stability and connect it to evolutionary dynamics. Separate role populations turn the scalar payoff problem into a vector problem in which the roles' restarting incentives must be jointly compatible. The roles can represent buyers and sellers or complementary members of a team; the stage game itself may be symmetric or asymmetric.

A different route to cooperation uses histories across matches. In finite-population repeated matching games, community enforcement can sustain cooperation when players observe only their own encounters (Kandori, 1992, Ellison, 1994). Deb, Sugaya, and Wolitzky (2020) establish a general folk theorem and consider extensions with endogenous separation and multiple communities. Their strategies use experience across encounters and calendar time. Here strategies start afresh at rematching, populations are nonatomic, and payoffs are evaluated at stationary matching distributions. The restarting restrictions concern this framework. They complement fixed-partner folk theorems, including their counterparts under imperfect public monitoring (Fudenberg, Levine, and Maskin, 1994), where continuation play can respond to the history of the ongoing relationship.

Section 2 defines the model and stationary payoffs. Section 3 develops restarting constraints, path protection, and an incentive test. Section 4 gives the payoff characterization and shows that strict individual rationality does not ensure approximation by Nash payoffs. Section 5 connects neutral stability to replicator dynamics and studies population structure and finite breakup. Section 6 concludes. Proofs of the technical results appear in Appendix A.

2. GAMES WITH ENDOGENOUS SEPARATION

2.1. *Stage game, matching, and strategies*

Let $n \geq 2$. The stage game has finite nonempty action sets $A^{(k)}$, $k = 1, \dots, n$, and payoff functions

$$U^{(k)} : A \longrightarrow \mathbb{R}, \qquad A := \prod_{k=1}^{n} A^{(k)}.$$

Write $U(a) := (U^{(1)}(a), \dots, U^{(n)}(a))$. There is one unit-mass population for each role. Each partnership contains one player from every population. Time is discrete, and every player plays the stage game once per period.

After each stage, a partnership survives exogenous separation with probability $\delta \in (0,1)$, independently of its composition and history. It also ends if any player chooses to leave. All members of a dissolved partnership enter the matching pool and are randomly rematched at the beginning of the next period. Matching draws independently from the role-specific pool distributions. There are no separation costs, no waiting periods, and no information flow between partnerships: new partners are anonymous (Ghosh and Ray, 1996). All players observe the complete action history within their current partnership. A higher δ means that partnerships are less likely to dissolve exogenously; it is not a separately specified time-preference parameter. Because partnerships dissolve as a whole, the profile of strategies in a partnership is fixed when it forms and does not change while it lasts.

A history of length t is an element $h \in A^t$. Put

$$\mathcal{H} := \bigcup_{t=0}^{\infty} A^t, \qquad A^0 := \{\emptyset\}.$$

A pure strategy for role k is a map $i_k : \mathcal{H} \to A^{(k)} \cup \{break\}$ with $i_k(\emptyset) \in A^{(k)}$. Thus new partners must play at least once. After a nonempty history, the players simultaneously prescribe an action or a breakup. If anyone prescribes $break$, no further stage is played in that partnership; otherwise their prescribed actions form the next stage profile. Let $\Omega^{(k)}$ denote the set of all such strategies. A pure strategy is a complete contingent plan, not a fixed stage-game action: it can prescribe different actions after different histories. Strategies depend only on the history of the current partnership, not on calendar time or previous partners.

For a strategy profile $I = (i_1, \ldots, i_n) \in \prod_k \Omega^{(k)}$, let $T_I \in \{1, 2, \ldots\} \cup \{\infty\}$ be the number of stages played before endogenous separation, in the absence of exogenous separation. Its path is $a_I^{[1,T_I]} = (a_I^{[1]}, \ldots, a_I^{[T_I]}) \in A^{T_I}$; when $T_I = \infty$ this denotes an infinite sequence. Paths record actions and their length, so a finite path and an infinite path are different even if the former is a prefix of the latter. For finite sequences, concatenation and repetition have their usual meanings; Φ^∞ denotes indefinite repetition of a nonempty finite sequence Φ.

2.2. *Population states and the stationary pool*

A population state is $x = (x^{(1)}, \ldots, x^{(n)})$, where each $x^{(k)}$ is a probability distribution with finite support in $\Omega^{(k)}$. Write $S_x^{(k)} := \operatorname{supp} x^{(k)}$ and let $\mathbb{D}$ be the set of all such states. A state is *monomorphic* if every role uses a single strategy, and *polymorphic* otherwise. The monomorphic state associated with a pure profile $s = (s_1, \ldots, s_n)$ is denoted by $\mathbf{e}_s$: everyone in role k uses s_k. For $n = 2$ we also write i-j for (i, j) and $\mathbf{e}_{ij}$ for its monomorphic state.

The matching-pool state $p = (p^{(1)}, \ldots, p^{(n)})$ consists of the corresponding role-specific distributions among players about to be matched.

ASSUMPTION 1—Stationary matching: *Payoffs are evaluated at a stationary partnership distribution conditional on the current population state. Partnership turnover is treated as fast relative to changes in population composition.*

The matching calculation holds strategy frequencies fixed. The replicator dynamics introduced below describe their slower evolution, as players revise the contingent plans they use across partnerships. Stationary matching is the maintained approximation linking these two levels. We establish existence and uniqueness of the stationary pool; convergence of the underlying matching process is not analyzed.

Fix finite nonempty sets $S_k \subset \Omega^{(k)}$, and set

$$\Delta_{S_k} := \left\{ z \in [0,1]^{S_k} : \sum_{i_k \in S_k} z_{i_k} = 1 \right\}, \qquad X := \prod_{k=1}^{n} \Delta_{S_k}, \qquad \mathcal{I} := \prod_{k=1}^{n} S_k.$$

Distributions in X are identified with elements of $\mathbb{D}$ by extending them by zero. For a finite vector $z = (z_i)_{i \in S}$, write $\|z\|_1 := \sum_{i \in S} |z_i|$. For $I = (i_1, \ldots, i_n) \in \mathcal{I}$, define the duration tensor

$$B_I := \sum_{t=1}^{T_I} \delta^{t-1} = \frac{1 - \delta^{T_I}}{1 - \delta}, \qquad \delta^{\infty} := 0.$$

In particular, $1 \leq B_I \leq (1-\delta)^{-1}$. For arbitrary vectors $d = (d^{(1)}, \ldots, d^{(n)})$, define the polynomials

$$q_I(d) := B_I \prod_{k=1}^{n} d^{(k)}_{i_k}, \quad Z(d) := \sum_{I \in \mathcal{I}} q_I(d), \quad M^{(k)}_{i_k}(d) := \sum_{i_{-k}} B_I \prod_{\ell \neq k} d^{(\ell)}_{i_\ell}, \tag{1}$$

where $i_{-k} := (i_\ell)_{\ell \neq k}$ and the last sum runs over $\prod_{\ell \neq k} S_\ell$.

At stationarity, the mass of partnerships of type I at stage $t \leq T_I$ is proportional to $\delta^{t-1} \prod_k p^{(k)}_{i_k}$. Summing over their ages and normalizing shows that the fraction of partnerships of type I is $q_I(p)/Z(p)$. Taking the role-specific marginals gives the pool–population map

$$x^{(k)}_{i_k} = f^{(k)}_{i_k}(p) := \frac{p^{(k)}_{i_k} M^{(k)}_{i_k}(p)}{Z(p)} = \frac{\displaystyle\sum_{i_{-k}} q_I(p)}{Z(p)}. \tag{2}$$

Both $Z(p)$ and $M^{(k)}_{i_k}(p)$ are positive on X, being convex combinations of entries of B. Thus $f : X \to X$ is well defined and preserves each population's support. Pool shares describe the composition of new matches, whereas population shares describe all players, including those in continuing partnerships. Equation (2) reweights the former by expected partnership duration. Strategies whose partnerships are shorter than average are therefore overrepresented among new matches, whether separation is initiated by them or by their partners.

If every $I \in \mathcal{I}$ has $T_I = \infty$, then $B_I = (1-\delta)^{-1}$ for every type of partnership. Hence $Z(p) = M_{i_k}^{(k)}(p) = (1-\delta)^{-1}$ and (2) gives $f(p) = p$ on X. This simplification need not survive an enlargement of the strategy sets by mutants that induce endogenous separation.

PROPOSITION 1—The stationary pool: *On every fixed finite strategy space X, the pool–population map f in (2) is a bi-Lipschitz homeomorphism from X onto itself. In particular, every population state in X has a unique consistent stationary pool.*

The proof in Appendix A.1 uses positive tensor scaling (Friedland, 2011) for existence and uniqueness, followed by a derivative and compactness argument for regularity. The appendix also gives the scaling equations for computing the inverse. The inverses on different finite strategy sets are consistent. When the sets are enlarged, support preservation forces the inverse of a state on the original face to remain on that face; the formulas for f coincide there, so uniqueness gives the same inverse. Thus every $x \in \mathbb{D}$ has a unique consistent $p \in \mathbb{D}$. Stationary matching therefore gives determinate payoff predictions: each population state selects one distribution of available partners. On each fixed finite strategy space, small changes in either distribution imply controlled changes in the other. The Lipschitz constants may depend on the strategy sets and on δ.

2.3. *Payoff functions*

For any pure profile I, define its expected cumulative partnership payoff and its per-period payoff vector $(F_I^{(k)})_{k=1}^n$ by

$$W_I^{(k)} := \sum_{t=1}^{T_I} \delta^{t-1} U^{(k)}(a_I^{[t]}), \qquad F_I^{(k)} := \frac{W_I^{(k)}}{B_I}. \tag{3}$$

All sums converge because the stage game is finite and $\delta < 1$. For any strategy $i_k \in \Omega^{(k)}$, including a strategy absent from the population, its payoff against pool state p is

$$\widehat{F}_{i_k}^{(k)}(p) := \frac{\displaystyle\sum_{i_{-k}} W_I^{(k)} \prod_{\ell \neq k} p_{i_\ell}^{(\ell)}}{\displaystyle\sum_{i_{-k}} B_I \prod_{\ell \neq k} p_{i_\ell}^{(\ell)}}. \tag{4}$$

Here the sums run over the supports of the other populations' pool distributions, so the payoff does not depend $p^{(k)}$. For a resident strategy, the formula follows by dividing its total stationary payoff by its stationary mass. For an absent strategy, it gives the limiting payoff of a rare entrant. The population-state payoff is

$$F_{i_k}^{(k)}(x) := \widehat{F}_{i_k}^{(k)}(f^{-1}(x)). \tag{5}$$

Unlike (4), it can depend on all population distributions, including $x^{(k)}$, through the pool map.

The payoff (4) is a convex combination of the profile payoffs $F_I^{(k)}$, with positive weights

$$\frac{B_I \prod_{\ell \neq k} p_{i_\ell}^{(\ell)}}{\sum_{j_{-k}} B_{i_k, j_{-k}} \prod_{\ell \neq k} p_{j_\ell}^{(\ell)}}$$

on the opponents' support. On every fixed X, each payoff function in (5) is Lipschitz: the numerator and denominator in (4) are polynomials, its denominator is bounded below by a positive constant, and f^{-1} is Lipschitz by Proposition 1. This also holds for any fixed absent strategy, by adjoining it to the finite strategy set.

For a monomorphic state $\mathbf{e}_s$, the pool is $\mathbf{e}_s$ and $F_{s_k}^{(k)}(\mathbf{e}_s) = F_s^{(k)}$. A unilateral mutant i_k receives $F_{(i_k, s_{-k})}^{(k)}$. If $T_I < \infty$, let the repeated path of I be $h_I = (a_I^{[1,T_I]})^\infty$; if $T_I = \infty$, let $h_I = a_I^{[1,\infty]}$. The geometric-series identity gives

$$F_I^{(k)} = (1-\delta) \sum_{t=1}^{\infty} \delta^{t-1} U^{(k)}(h_I^{[t]}). \tag{6}$$

Thus stationary average payoffs can be calculated as normalized discounted values of repeated paths. This is an algebraic identity: the model's payoff criterion is the stationary per-period average, not a separately assumed preference for discounted lifetime consumption. When the dependence on δ matters, we write $F_\delta(s) := (F_s^{(1)}, \dots, F_s^{(n)})$.

Finally, for $x, y \in \mathbb{D}$, define

$$E^{(k)}(y,x) := \sum_{i_k \in S_y^{(k)}} y_{i_k}^{(k)} F_{i_k}^{(k)}(x), \qquad E(y,x) := \sum_{k=1}^{n} E^{(k)}(y,x). \tag{7}$$

This aggregate is linear in the component distributions of its first argument, but generally nonlinear in its second. It evaluates role-specific strategy distributions y against the same resident environment x; it is not the payoff from matching all members of y together. We use the given cardinal payoffs and assign equal weight to the roles in E. The aggregate neutral-stability condition below is relative to this specified normalization and need not be invariant under independent positive rescalings of different roles' payoffs. Nash and path-protection comparisons, in contrast, are role-specific.

The feasible stage-payoff set and the maximum joint stage payoff are

$$\mathcal{F} := \mathrm{co}\{U(a) : a \in A\}, \qquad J := \max_{a\in A} \sum_{k=1}^{n} U^{(k)}(a).$$

Every pure profile's payoff vector belongs to $\mathcal{F}$, by (6). The same holds for the vector of average payoffs $\left(E^{(k)}(x,x)\right)_{k=1}^{n}$ of an arbitrary state.

LEMMA 2—Average payoffs of a state: *Let $x \in \mathbb{D}$, let $p = f^{-1}(x)$ be its pool, and take $S_k := S_x^{(k)}$ in (1). Then*

$$E^{(k)}(x,x) = \sum_{I\in\mathcal{I}} w_I\, F_I^{(k)} \qquad (k = 1,\ldots,n), \qquad w_I := \frac{q_I(p)}{Z(p)}, \tag{8}$$

where the weights w_I are nonnegative, sum to one, and do not depend on k. Consequently $\left(E^{(k)}(x,x)\right)_{k=1}^{n} \in \mathcal{F}$ and $E(x,x) \le J$.

PROOF: By (2) and (4), the factor $M_{i_k}^{(k)}(p)$ cancels:

$$x_{i_k}^{(k)} \widehat{F}_{i_k}^{(k)}(p) = \frac{p_{i_k}^{(k)} M_{i_k}^{(k)}(p)}{Z(p)} \cdot \frac{\displaystyle\sum_{i_{-k}} W_I^{(k)} \prod_{\ell\neq k} p_{i_\ell}^{(\ell)}}{M_{i_k}^{(k)}(p)} = \frac{1}{Z(p)} \sum_{i_{-k}} W_I^{(k)} \prod_{\ell=1}^{n} p_{i_\ell}^{(\ell)}.$$

Summing over $i_k \in S_k$ and using $W_I^{(k)} = B_I F_I^{(k)}$ gives (8). The weights w_I are the stationary frequencies of partnership types, so they are nonnegative and sum to one. They are the same for every role, and each F_I lies in $\mathcal{F}$ by (6), since one common sequence h_I serves all roles. Hence the vector of average payoffs is a convex combination of points of $\mathcal{F}$, and its coordinate sum is at most J. *Q.E.D.*

The role-independent weights in (8) reflect the fact that all roles participate in the same partnerships. Their average payoffs therefore form one feasible vector, rather than independent convex combinations chosen separately for each role.

3. RESTARTING INCENTIVES AND PATH PROTECTION

3.1. *Nash states and first-stage constraints*

A state y is a best response to x if $E(y,x) \geq E(z,x)$ for every $z \in \mathbb{D}$. By (7), this is equivalent to requiring that every strategy in $S_y^{(k)}$ maximize $F_{i_k}^{(k)}(x)$ over $\Omega^{(k)}$, for every role k.

DEFINITION 1—Nash state: A state $x \in \mathbb{D}$ is Nash if $E(x,x) \geq E(y,x)$ for every $y \in \mathbb{D}$. A pure strategy profile s is Nash if e_s is Nash.

The Nash condition is role-specific despite its aggregate notation: no player can earn more by adopting another contingent plan against the same stationary resident environment. It compares complete plans for current-partnership histories under the stationary average-payoff criterion. We do not impose a separate requirement of sequential optimality at every off-path history; the incentive test below checks every possible first departure against the constructed residents.

For a stage-game action profile $a \in A$, let $a_{-k} := (a^{(\ell)})_{\ell \neq k}$ and write

$$b_k(a) := \max_{\widehat{a}^{(k)} \in A^{(k)}} U^{(k)}(\widehat{a}^{(k)}, a_{-k}), \qquad m_k := \min_{a_{-k} \in \prod_{\ell \neq k} A^{(\ell)}} b_k(a). \tag{9}$$

Here m_k is the *pure minmax payoff*; $b_k(a)$ does not depend on $a^{(k)}$. Write $b(a) := (b_1(a), \ldots, b_n(a))$. A stage-game profile a is Nash precisely when $U^{(k)}(a) = b_k(a)$ for every k. Leaving provides no exogenous outside payoff: it returns the player to the matching pool. Against a monomorphic resident profile beginning with a, repeatedly playing a best response at the first stage and leaving earns $b_k(a)$, independently of δ. This is one available restarting strategy, not necessarily the best deviation overall. We distinguish these pure minmax bounds from mixed minmax payoffs; the two need not coincide.

LEMMA 3—First-stage constraint: *Let s be a pure Nash profile and let $a_s^{[1]}$ denote its initial action profile. Then*

$$F_s^{(k)} \geq b_k(a_s^{[1]}) \geq m_k \qquad (k = 1, \dots, n).$$

Consequently, $\sum_k b_k(a_s^{[1]}) \leq J$.

PROOF: A player in role k can play a best response to $a_{s,-k}^{[1]}$ in every new partnership and break after that stage. Every such partnership lasts one stage and yields $b_k(a_s^{[1]})$. The Nash condition gives the first inequality; the definition of m_k gives the second. Summing and using $F_s \in \mathcal{F}$ proves the last assertion. *Q.E.D.*

At a polymorphic Nash state x, let σ_k be the distribution of initial actions induced by the pool $p = f^{-1}(x)$:

$$\sigma_k(a_k) := \sum_{i_k : i_k(\emptyset) = a_k} p_{i_k}^{(k)}, \qquad a_k \in A^{(k)}.$$

The same best-response-and-leave deviation yields

$$E^{(k)}(x,x) \geq \widetilde{b}_k(\sigma) := \max_{\widehat{a}_k \in A^{(k)}} \sum_{a_{-k}} U^{(k)}(\widehat{a}_k, a_{-k}) \prod_{\ell \neq k} \sigma_\ell(a_\ell). \tag{10}$$

All roles face one common profile of initial-action distributions σ. As in Lemma 3, these simultaneous constraints need not reduce to separate minmax bounds. Combining them with Lemma 2 gives the following restriction.

COROLLARY 4—First-stage bound at every Nash state: *Let $x \in \mathbb{D}$ be Nash and let σ be the profile of initial-action distributions induced by its pool. Then $v := \big(E^{(k)}(x,x)\big)_{k=1}^n$ satisfies*

$$v \in \mathcal{F} \qquad \textit{and} \qquad v_k \geq \widetilde{b}_k(\sigma) \quad (k = 1, \dots, n).$$

This is a necessary restriction on the average payoffs of every Nash state; it is not an equilibrium characterization.

Consider the three two-player games in Table I: an asymmetric Prisoner's Dilemma (PD), Hawk–Dove, and Split the Pie. Entries give row and column payoffs. In the first two games, C denotes cooperation or Dove and D denotes defection or Hawk. In the third, H offers the column player the higher share, L offers the lower share, and A and R mean accept and reject.

TABLE I

EXAMPLE TWO-PLAYER STAGE GAMES.

	C'	D'
C	$2,3$	$0,4$
D	$3,1$	$1,2$

Asymmetric PD

	C'	D'
C	$4,3$	$2,4$
D	$5,1$	$1,0$

Hawk–Dove

	A	R
H	$2,3$	$0,0$
L	$3,2$	$0,0$

Split the Pie

Their pure minmax vectors are $(1,2)$, $(2,1)$, and $(0,2)$, respectively. In the asymmetric Prisoner's Dilemma, $b(CC') = (3,4)$ has sum $7 > J = 5$. In Hawk–Dove, $b(CC') = (5,4)$ has sum $9 > J = 7$. Thus no pure Nash profile in either repeated game can start with CC'. In particular, perpetual mutual cooperation cannot be the path of a pure Nash profile in this model.

3.2. *Neutral stability*

We extend the classical definition of neutral stability (Maynard Smith, 1982) as follows. The second condition is imposed on every state that is a best response to x, not merely on pure alternatives.

DEFINITION 2—Neutrally stable state: A state $x \in \mathbb{D}$ is neutrally stable if

$$
\begin{aligned}
E(x,x) &\geq E(y,x) && \text{for every } y \in \mathbb{D}, \\
E(x,y) &\geq E(y,y) && \text{for every } y \in \mathbb{D} \text{ with } E(y,x) = E(x,x).
\end{aligned}
$$

A pure profile s is neutrally stable if $\mathbf{e}_s$ is neutrally stable.

The first condition rules out profitable alternatives in any role. The second compares the residents with alternatives that tie against them, now using those alternatives as the environment and the specified aggregate of role payoffs.

The use of a weak inequality is important in an unrestricted repeated-game strategy space. Distinct strategies can generate the same behavior against every resident strategy and differ only after histories that never arise in those interactions. Strict superiority against every distinct best-response alternative would exclude these behaviorally equivalent strategies; related difficulties arise in standard repeated games (Selten and Hammerstein, 1984). We focus on the precise neutral-stability condition above rather than identifying it with other invasion criteria for nonlinear population games.

3.3. *Finite endogenous breakup*

We next ask when a stable convention can prescribe ending a partnership after finitely many stages. Nash equilibrium first imposes a restriction on its last stage.

LEMMA 5: *If a pure Nash profile s has $T_s < \infty$, its last action profile $a_s^{[T_s]}$ is a pure Nash profile of the stage game.*

Indeed, a player with a strictly profitable stage-game deviation at the last stage can make that deviation and then leave, preserving the partnership length and increasing that player's payoff. The formal comparison is given in Appendix A. If the stage game has no pure Nash profile, no pure Nash profile of the repeated game can have a finite endogenous breakup period. Conversely, perpetual play of any pure stage-game Nash profile, with separation after an action deviation, is a Nash profile of the repeated game. A deviator's stage payoff never exceeds its equilibrium value.

Call a finite endogenous breakup of a profile s *unanimous* if $s_k(a_s^{[1,T_s]}) = \mathit{break}$ for every role k. Thus every member prescribes breakup after the last on-path history. For two players this is bilateral breakup.

LEMMA 6: *If a pure neutrally stable profile s has a finite unanimous breakup period, then every action profile on its path maximizes the joint stage payoff $\sum_k U^{(k)}(a)$.*

The proof exploits neutral alternatives that continue the relationship when all their co-players also agree to continue; it is given in Appendix A. Together with Lemma 5, this rules out finite unanimous breakup whenever no pure Nash action profile attains the maximum joint stage payoff. This applies to both the asymmetric Prisoner's Dilemma and Hawk–Dove examples in Table I. The unanimity hypothesis is essential to this argument: the lemma makes no claim about all finite breakups. For λ-neutral stability as in Remark 4, the same proof requires every on-path action profile to maximize $\sum_k \lambda_k U^{(k)}(a)$, so this restriction depends on the relative payoff weights. If the same profile with unanimous breakup were λ-neutrally stable for every $\lambda > 0$, each action on its path would maximize every player's stage payoff separately, by letting the weights tend to a coordinate vector.

3.4. *Path-protecting profiles*

A profile can be neutrally stable even though alternative strategies earn the same payoff. Path protection identifies a useful class for which every departure from an agreed infinite path is strictly disadvantageous. The definition extends the symmetric construction of Izquierdo and Izquierdo (2026) to role-specific strategies.

DEFINITION 3—Path-protecting profile: A pure profile s with $T_s = \infty$ and path $\pi := a_s^{[1,\infty]}$ is path-protecting if, for every role k and every $i_k \in \Omega^{(k)}$,

$$a_{(i_k,s_{-k})}^{[1,T_{(i_k,s_{-k})}]} \neq \pi \quad \Longrightarrow \quad F_{(i_k,s_{-k})}^{(k)} < F_s^{(k)}. \tag{11}$$

Departures include both an action different from the prescribed action and a decision to end the partnership. An alternative that follows the entire infinite path earns the same payoff as the resident. Equivalently, for every pure profile I,

$$a_I^{[1,T_I]} \neq \pi \quad \Longrightarrow \quad E(\mathbf{e}_I, \mathbf{e}_s) < E(\mathbf{e}_s, \mathbf{e}_s). \tag{12}$$

To see this, if I has a different path, its first departure involves a component strategy that also departs when paired with residents in all other roles. That component earns strictly less by (11), and no other component earns more. Conversely, taking $I = (i_k, s_{-k})$ in (12) gives (11).

PROPOSITION 7: *Every path-protecting profile is neutrally stable.*

PROOF: Let s be path-protecting. In each population, every alternative strategy either follows the entire resident path and earns $F_s^{(k)}$, or earns strictly less. Hence $\mathbf{e}_s$ is Nash. If y is a best response to $\mathbf{e}_s$, equality of aggregate payoffs and the absence of profitable component deviations imply that every strategy in every $S_y^{(k)}$ follows the resident path against s_{-k}. Induction on the history length shows that any profile combining these strategies and resident strategies follows that same infinite path: at each prescribed prefix all component strategies prescribe the next resident action, and none leaves. All partnerships relevant to $E(\mathbf{e}_s, y)$ and $E(y, y)$ therefore have the same role-specific payoffs $F_s^{(k)}$. Thus $E(\mathbf{e}_s, y) = E(y, y)$, proving neutral stability. *Q.E.D.*

The proof establishes the equality of payoffs separately in every role. Consequently the neutral-stability inequalities remain valid under any positive weighting of the roles. Section 5 derives the dynamic implications.

No purely periodic infinite path Φ^∞ can be protected. Indeed, a unilateral strategy that follows one complete period and then leaves generates a finite partnership with exactly the same normalized cycle payoff as the resident's infinite repetition. The strict inequality in (11) fails. This does not rule out *eventually* periodic protected paths with an initial transient, which are the basis of our construction.

3.5. *Deviation-leaving profiles and an incentive test*

Given an infinite path $\pi = (a^{[1]}, a^{[2]}, \ldots)$, its *deviation-leaving* profile s^π is defined, for every role k and every history $h \in A^t$, by

$$s_k^\pi(h) = \begin{cases} a^{[t+1],(k)}, & h = (a^{[1]}, \ldots, a^{[t]}), \\ \mathit{break}, & \text{otherwise}. \end{cases}$$

For $t = 0$ the first case applies. Thus players remain together while the prescribed path is followed and leave after any other history. This specifies a unique pure profile, including off-path behavior.

Write $F_\delta^{(k)} := F_{s^\pi}^{(k)}$ and define the continuation values

$$V_t^{(k)} := (1-\delta) \sum_{j=1}^{\infty} \delta^{j-1} U^{(k)}(a^{[t+j]}), \qquad t = 0, 1, \ldots.$$

In particular, $V_0^{(k)} = F_\delta^{(k)}$.

PROPOSITION 8—Incentive test: *The deviation-leaving profile s^π is Nash if and only if, for every k and $t \geq 1$,*

$$(1-\delta)b_k(a^{[t]}) + \delta F_\delta^{(k)} \leq V_{t-1}^{(k)}. \tag{13}$$

It is path-protecting if and only if all these inequalities are strict.

PROOF: For a fixed role k, consider the strategy that follows the path through stage $t-1$, plays a best response to $a_{-k}^{[t]}$ at stage t, and then leaves. Its cycle payoff is

$$D_t^{(k)} = \frac{(1-\delta)\left[\sum_{j=1}^{t-1} \delta^{j-1} U^{(k)}(a^{[j]}) + \delta^{t-1} b_k(a^{[t]})\right]}{1-\delta^t}.$$

Subtracting the resident payoff gives

$$(1-\delta^t)(D_t^{(k)} - F_\delta^{(k)}) = \delta^{t-1}\left[(1-\delta)b_k(a^{[t]}) + \delta F_\delta^{(k)} - V_{t-1}^{(k)}\right]. \tag{14}$$

Every departure from the resident path ends its partnership after a finite number τ of completed stages: an action deviation is followed by the residents' breakup, and a voluntary breakup ends the partnership directly. Its payoff is bounded above by $D_\tau^{(k)}$, since the best-response-and-leave strategy has the same earlier payoffs and maximizes the last-stage payoff. Conversely, for $t \geq 1$, each $D_t^{(k)}$ is attainable by such a unilateral strategy and has a finite path, distinct from π. The sign identity (14) therefore proves both assertions. *Q.E.D.*

Economically, the test asks whether the continuation of an established relationship dominates the gain from a best current action followed by restarting. The resident payoff $F_\delta^{(k)}$ is the benchmark in this comparison; (14) converts it into the payoff comparison for a strategy that repeatedly restarts, rather than a single isolated action change. After an action deviation, every other resident prescribes breakup, so no player can preserve that partnership by unilaterally changing the continuation decision. This explains the separation response in the construction without equating population Nash equilibrium with a sequential equilibrium concept.

LEMMA 9—Eventually periodic paths: *If $\pi = (\Phi_0, \Phi_p^\infty)$, where Φ_0 has length $\ell \geq 0$ and Φ_p has length $T \geq 1$, it is enough to check (13), weakly or strictly, for $1 \leq t \leq \ell + T$.*

PROOF: For $t > \ell$, both $b_k(a^{[t]})$ and $V_{t-1}^{(k)}$ repeat with period T; $F_\delta^{(k)}$ is fixed. Thus every later inequality coincides with one of those for $\ell+1,\dots,\ell+T$. The last value includes breakup after a complete period. *Q.E.D.*

EXAMPLE 1: In the Split the Pie game of Table I, perpetual play of (L,A) is Nash but cannot be protected, because it is purely periodic. Consider instead $\pi = ((L,R),(H,A)^\infty)$. Its payoffs are $F_\delta = (2\delta, 3\delta)$, and $V_1 = (2,3)$. The best-response vectors at the first two stages are $b(L,R) = (0,2)$ and $b(H,A) = (3,3)$. By Lemma 9, only those two stages need checking. The inequalities in (13) hold precisely when $\delta \geq 2/3$, and all are strict when $\delta > 2/3$. Thus the deviation-leaving profile is Nash for $\delta \geq 2/3$ and path-protecting for $\delta > 2/3$.

4. A FOLK THEOREM FOR n POPULATIONS

The feasible payoff set $\mathcal{F}$ is defined in Section 2.3, and the restarting thresholds $b_k(a^0)$ in Section 3.1. We now characterize the limiting payoffs of pure path-protecting profiles. Strict improvement over the thresholds for some common initial profile is necessary for path protection at any fixed δ and sufficient for payoff approximation as δ approaches one.

The economic constraint comes from the ability to start a new relationship. In the monomorphic conventions constructed here, each new partnership begins with the same action profile a^0; while the other roles play a^0_{-k}, player k can earn $b_k(a^0)$ by choosing a best response and leaving immediately. Continuing the prescribed relationship must improve on this repeatable alternative, which replaces the within-relationship continuation play used in fixed-partner folk theorems. The initial profile need not be Nash, and neither full dimensionality of $\mathcal{F}$ nor a nonequivalent-utilities condition is used.

4.1. *Construction and limiting payoffs*

Vector inequalities below are componentwise. Define the payoff region

$$\mathcal{R} := \bigcup_{a^0 \in A} \left\{ v \in \mathcal{F} : v_k > b_k(a^0) \text{ for every } k \right\}. \tag{15}$$

For a nonempty finite sequence $\Phi = (\Phi^{[t]})_{t=1}^{T}$ of action profiles, let

$$\bar{U}_{\Phi} := \frac{1}{T}\sum_{t=1}^{T} U(\Phi^{[t]}).$$

PROPOSITION 10—Construction from a common initial profile: *Fix $a^0 \in A$, an arbitrary finite sequence Φ_f of action profiles (possibly empty), and a nonempty finite sequence Φ_p satisfying*

$$\bar{U}^{(k)}_{\Phi_p} > b_k(a^0) \qquad (k = 1, \ldots, n).$$

There exist an integer $M \geq 1$ and $\bar{\delta} \in (0,1)$ such that the pure deviation-leaving profile s with path

$$\big((a^0)^M, \Phi_f, (\Phi_p)^{\infty}\big) \tag{16}$$

is path-protecting, and hence neutrally stable, for every $\delta \in (\bar{\delta}, 1)$. The profile s is independent of δ, and

$$\lim_{\delta \uparrow 1} F_{\delta}(s) = \bar{U}_{\Phi_p}.$$

The proof is given in Appendix A. The initial phase creates an entry cost in foregone relationship payoffs, not a fee or a waiting period outside the game. It is incurred after exogenous separations as well as voluntary departures. As exogenous separations become rare, its contribution to complying players' average payoffs vanishes, whereas repeated voluntary restarting continues to reintroduce it. The periodic tail determines the limiting payoff.

Any prescribed finite sequence can be inserted as the intermediate phase without changing that limit. Once the intermediate phase and periodic pattern are fixed, one chooses the initial duration M and then the continuation threshold $\bar{\delta}$. The result concerns sufficiently long-lived partnerships, not arbitrary payoff attainment at a fixed δ.

THEOREM 11—Folk theorem: *For every $v \in \mathcal{R}$ and every $\varepsilon > 0$, there exist a pure deviation-leaving strategy profile s and $\bar{\delta} \in (0,1)$ such that, for every $\delta \in (\bar{\delta}, 1)$, s is path-protecting and neutrally stable, and*

$$\max_{1 \leq k \leq n} |F^{(k)}_{\delta}(s) - v_k| < \varepsilon.$$

The same profile s works for all $\delta \in (\bar{\delta}, 1)$.

PROOF: Choose a^0 with $v_k > b_k(a^0)$ for all k. Since v is a convex combination of finitely many stage-payoff vectors, approximating its coefficients by rational probability weights gives a finite pattern Φ_p with mean r satisfying $r_k > b_k(a^0)$ and $\max_k |r_k - v_k| < \varepsilon/2$. Apply Proposition 10, for instance with Φ_f empty. Increase $\bar{\delta}$ if necessary so that $\max_k |F_\delta^{(k)}(s) - r_k| < \varepsilon/2$ whenever $\delta > \bar{\delta}$. *Q.E.D.*

PROPOSITION 12—Necessary restriction and asymptotic payoff region: *If a pure strategy profile s is path-protecting at a continuation probability $\delta \in (0,1)$ and begins with the action profile $a^{[1]}$, then*

$$F_\delta^{(k)}(s) > b_k(a^{[1]}) \qquad (k = 1, \ldots, n).$$

In particular, $F_\delta(s) \in \mathcal{R}$. Moreover, the set of vectors that can arise as limits $F_{\delta_m}(s_m) \to v$ along sequences $\delta_m \uparrow 1$, with each s_m a pure path-protecting profile at δ_m, is exactly $\overline{\mathcal{R}}$.

PROOF: In population k, consider a strategy that plays a best response to $a_{-k}^{[1]}$ at the first stage and then leaves. Against residents, every partnership lasts one stage, so its payoff is exactly $b_k(a^{[1]})$. Its finite partnership path differs from the resident's infinite path. Path protection therefore gives the strict inequality. The resident payoff is a convex combination of stage-payoff vectors, so it belongs to $\mathcal{F}$ and hence to $\mathcal{R}$.

Every limit in the statement must consequently belong to $\overline{\mathcal{R}}$. Conversely, given $v \in \overline{\mathcal{R}}$, choose $\{v_m \in \mathcal{R}\}_{m \in \mathbb{N}}$ tending to v. For each m, apply Theorem 11 with target v_m and tolerance $1/m$, obtaining a profile s_m and a threshold $\bar{\delta}_m$. Choose $\delta_m > \bar{\delta}_m$ so that $\delta_m \uparrow 1$. Then $F_{\delta_m}(s_m) \to v$. *Q.E.D.*

The last proposition concerns pure path-protecting profiles, not all Nash states or all neutrally stable population states. Its limit statement allows the profiles s_m to vary with m; it does not assert exact attainment of every payoff in $\mathcal{R}$ or its closure.

4.2. *Special cases and feasibility*

COROLLARY 13—Two-player games: *For $n = 2$, let m_1, m_2 be the pure minmax payoffs in (9). Then $\mathcal{R} = \{v \in \mathcal{F} : v_1 > m_1,\ v_2 > m_2\}$. Thus every feasible payoff vector strictly above the pure minmax payoffs admits the approximation in Theorem 11.*

PROOF: Choose $a^{0,(2)}$ attaining the minimum defining m_1 and $a^{0,(1)}$ attaining the minimum defining m_2. These choices are independent and give one profile a^0 with $b_1(a^0) = m_1$ and $b_2(a^0) = m_2$. For every other profile a, $b_k(a) \geq m_k$, which proves the claimed equality. *Q.E.D.*

COROLLARY 14—A pure Nash initial profile: *For any $n \geq 2$, if a^N is a pure Nash profile of the stage game, then every $v \in \mathcal{F}$ satisfying $v_k > U^{(k)}(a^N)$ for all k admits the approximation in Theorem 11.*

PROOF: At a pure Nash profile, $b_k(a^N) = U^{(k)}(a^N)$ for every k. *Q.E.D.*

A pure stage-game Nash profile therefore provides a ready-made initial phase whenever a feasible payoff vector strictly improves every role's payoff. Such a vector need not exist. The general theorem can also use non-Nash initial profiles.

With more than two roles, minimizing each restarting threshold separately need not produce a common initial profile. One role's initial action affects several other roles' opportunities, so reducing the gain from restarting for one role can increase it for another. The region $\mathcal{R}$ records whether these incentive constraints can be satisfied simultaneously.

REMARK 1—Checking nonemptiness: For each $a \in A$, consider the linear program

$$\begin{aligned}
\eta_a := \max_{\theta,\eta} \quad & \eta \\
\text{subject to} \quad & \theta_c \geq 0 \ (c \in A), \qquad \sum_{c \in A} \theta_c = 1, \\
& \sum_{c \in A} \theta_c U^{(k)}(c) \geq b_k(a) + \eta \qquad (k = 1, \ldots, n).
\end{aligned} \tag{17}$$

Here $\eta \in \mathbb{R}$. The largest common margin for a given θ is the minimum of finitely many continuous payoff differences, so the maximum exists on the compact probability simplex.

By (15), $\mathcal{R} \neq \varnothing$ if and only if $\max_{a\in A} \eta_a > 0$. Thus nonemptiness can be checked with $|A|$ linear programs.

4.3. *Polymorphic initial actions*

At a polymorphic state, a given role need not start every new partnership with the same action. Allowing a common profile of initial-action *distributions*, as in (10), can enlarge the first-stage feasible region by offering more ways to meet the constraints jointly. For a particular choice of distributions, it need not lower every role's threshold. Let $\Delta_{A^{(k)}}$ denote the probability simplex on $A^{(k)}$ and define

$$\widetilde{\mathcal{R}} := \bigcup_{\sigma \in \prod_k \Delta_{A^{(k)}}} \{v \in \mathcal{F} : v_k > \widetilde{b}_k(\sigma) \text{ for every } k\}, \tag{18}$$

and let $\widetilde{\mathcal{R}}^{\geq}$ be given by the same expression with weak inequalities. Degenerate distributions give $\mathcal{R} \subseteq \widetilde{\mathcal{R}} \subseteq \widetilde{\mathcal{R}}^{\geq}$, and Corollary 4 places the average payoff vector of every Nash state in $\widetilde{\mathcal{R}}^{\geq}$. Neither inclusion need be strict. These are first-stage feasibility regions only: membership of $\widetilde{\mathcal{R}}$ does not assert that any polymorphic equilibrium attains the corresponding payoffs, and the results of this section are proved for $\mathcal{R}$. For each fixed σ, the program (17) applies with $b(a)$ replaced by $\widetilde{b}(\sigma)$. When σ also varies, the expected deviation payoffs contain as variables the products $\prod_{\ell\neq k} \sigma_\ell(a_\ell)$ of the opponents' marginals. The formulation is therefore a linear program for $n = 2$ and generally nonlinear for $n \geq 3$.

REMARK 2—The relaxed region for two players: For $n = 2$, $\widetilde{b}_1(\sigma)$ depends only on σ_2 and $\widetilde{b}_2(\sigma)$ only on σ_1, so the two minimizations decouple exactly as in the proof of Corollary 13. Writing $\widetilde{m}_k := \min_{\sigma_{-k}} \widetilde{b}_k(\sigma)$ for the mixed minmax payoffs, which are attained because each $\widetilde{b}_k$ is a continuous function on a compact simplex,

$$\widetilde{\mathcal{R}} = \{v \in \mathcal{F} : v_1 > \widetilde{m}_1,\ v_2 > \widetilde{m}_2\}.$$

Since degenerate distributions are admissible, $\widetilde{m}_k \leq m_k$, so the difference between the pure and mixed minmax thresholds accounts for the gap between these two first-stage regions. For the three games in Table I the pure and mixed minmax vectors coincide, so there the two regions are the same. The set $\{v \in \mathcal{F} : v_k > \widetilde{m}_k,\ k = 1, 2\}$ is the strictly individually rational set appearing in the folk theorem for two-player repeated games with fixed partners and sufficiently patient players (Fudenberg and Maskin, 1986).

EXAMPLE 2—Conflicting restarting incentives: Consider three players with binary action sets and payoffs

$$\begin{aligned} U^{(1)}(a_1,a_2,a_3) &= 2\mathbf{1}_{\{a_2=a_3\}} + a_1, \\ U^{(2)}(a_1,a_2,a_3) &= 2\mathbf{1}_{\{a_1\neq a_3\}} + a_2, \qquad a_k \in \{0,1\}, \\ U^{(3)}(a_1,a_2,a_3) &= 2\mathbf{1}_{\{a_1\neq a_2\}} + a_3, \end{aligned}$$

where $\mathbf{1}$ denotes the indicator function. Each player's own action affects that player's payoff only through the last term. Thus action 1 strictly dominates action 0 for every player, and the unique stage-game Nash profile is $(1,1,1)$, with payoff vector $(3,1,1)$.

For any initial profile a, the best-response bounds are

$$b_1(a) = 2\mathbf{1}_{\{a_2=a_3\}} + 1, \qquad b_2(a) = 2\mathbf{1}_{\{a_1\neq a_3\}} + 1, \qquad b_3(a) = 2\mathbf{1}_{\{a_1\neq a_2\}} + 1.$$

The other two players can make the relevant indicator function vanish for each player separately, so the pure minmax payoff vector is $(1,1,1)$. Moreover, the average payoff vector of a pattern containing each of the eight action profiles once is

$$v = (3/2, 3/2, 3/2) \in \mathcal{F}.$$

Each player can guarantee 1, and the other roles can hold that player to 1 with deterministic actions, so mixed minmax payoffs are also $(1,1,1)$. Thus v is strictly individually rational under either convention. The feasible set is full dimensional: relative to $U(0,0,0)$, the payoff differences at $(0,0,1)$, $(0,1,0)$, and $(0,1,1)$ are

$$(-2,2,1), \qquad (-2,1,2), \qquad (0,3,3),$$

with determinant 12. Since v averages all eight payoff vectors with positive weights, it lies in the interior of $\mathcal{F}$. The example therefore satisfies the dimensionality hypothesis of the fixed-partner folk theorem.

Nevertheless, the three indicators cannot all vanish at any one profile: vanishing of the last two would give $a_1 = a_2 = a_3$, making the first indicator equal to one. Therefore every

a has $b_k(a) = 3$ for at least one player k. Since 3 is also that player's largest stage payoff, no feasible payoff vector can strictly exceed $b(a)$ in every coordinate. Consequently, $\mathcal{R} = \varnothing$. By Proposition 12, there is no pure path-protecting profile for any $\delta \in (0,1)$, even though the stage game has a unique pure Nash profile and feasible payoffs strictly above all pure minmax bounds. This example also shows why the closure in Proposition 12 cannot be obtained merely by replacing strict inequalities with weak ones in (15): $(3,1,1) = U(1,1,1) = b(1,1,1)$ satisfies the weak inequalities, while $\overline{\mathcal{R}}$ is empty.

PROPOSITION 15—Strict individual rationality does not suffice: *Some strictly individually rational payoff vectors lie at a positive distance from the set of Nash payoffs, uniformly in the continuation probability, even when the feasible set is full dimensional. More precisely, a payoff vector v in the interior of a full-dimensional feasible set $\mathcal{F}$ can strictly exceed every role's pure minmax payoff, and hence every role's mixed minmax payoff, and still satisfy, for some $\varepsilon > 0$,*

$$\left\| \left(E^{(k)}(x,x)\right)_{k=1}^{n} - v \right\|_{\infty} \geq \varepsilon$$

for every $\delta \in (0,1)$ and every Nash state $x \in \mathbb{D}$, including polymorphic states. In particular, v cannot be approximated by Nash payoffs as $\delta \uparrow 1$.

PROOF: It suffices to exhibit one stage game with these properties. Consider the three-player game of Example 2 and $v := (3/2, 3/2, 3/2)$. The example shows that $\mathcal{F}$ is full dimensional, that v lies in its interior, and that the pure minmax vector is $(1,1,1)$, so $v_k > m_k$ for every k. We prove the claim with $\varepsilon = 1/2$ by showing that, for every $\delta \in (0,1)$ and every Nash state $x \in \mathbb{D}$, including polymorphic states,

$$\max_k E^{(k)}(x,x) \geq 2, \qquad \left\| \left(E^{(k)}(x,x)\right)_{k=1}^{3} - v \right\|_{\infty} \geq 1/2.$$

Let σ be the common profile of initial-action distributions induced by the stationary pool of x, and put $t_k := 2\sigma_k(1) - 1$. Then

$$\widetilde{b}_1(\sigma) = 2 + t_2 t_3, \qquad \widetilde{b}_2(\sigma) = 2 - t_1 t_3, \qquad \widetilde{b}_3(\sigma) = 2 - t_1 t_2.$$

At least one bound is at least 2: otherwise $t_2 t_3 < 0$, $t_1 t_3 > 0$, and $t_1 t_2 > 0$, which is impossible. By Corollary 4, the corresponding role earns at least 2 at x, which proves the first inequality; the second follows because every coordinate of v equals $3/2$. Finally, if $\delta_m \uparrow 1$ and each x_m is a Nash state at δ_m, every average payoff vector $\left(E^{(k)}(x_m, x_m)\right)_{k=1}^{3}$ remains at sup-norm distance at least $1/2$ from v, so these vectors cannot converge to v. *Q.E.D.*

In Example 2, the obstruction is joint compatibility: each role can be held to its minmax payoff of 1 separately, but every common product distribution of initial actions gives some role a restarting payoff of at least 2. With two roles, the failure of approximation in Proposition 15 is impossible: any feasible vector strictly above both pure minmax payoffs is itself a limit of Nash payoffs, by Corollary 16 with w equal to that vector.

COROLLARY 16—Nash payoffs with two roles: *Let $n = 2$ and suppose some $w \in \mathcal{F}$ satisfies $w_k > m_k$ for $k = 1, 2$. Let $\mathcal{N}$ be the set of vectors that can arise as limits of average payoff vectors $\left(E^{(k)}(x_m, x_m)\right)_{k=1}^{2}$ along sequences $\delta_m \uparrow 1$, with each $x_m \in \mathbb{D}$ a Nash state at δ_m, including polymorphic states. Then*

$$\{v \in \mathcal{F} : v_k \geq m_k,\ k = 1, 2\} \subseteq \mathcal{N} \subseteq \{v \in \mathcal{F} : v_k \geq \widetilde{m}_k,\ k = 1, 2\}.$$

If the pure and mixed minmax payoffs coincide, both inclusions are equalities and $\mathcal{N} = \overline{\mathcal{R}}$: every limit of Nash payoffs is also a limit of payoffs of pure path-protecting profiles.

PROOF: By Corollary 13, $\mathcal{R} = \{v \in \mathcal{F} : v_k > m_k,\ k = 1, 2\}$, which contains w. The set $\{v \in \mathcal{F} : v_k \geq m_k,\ k = 1, 2\}$ is closed and contains $\mathcal{R}$; each of its points v is the limit, as $\alpha \downarrow 0$, of the points $(1 - \alpha)v + \alpha w \in \mathcal{R}$. Hence this set equals $\overline{\mathcal{R}}$. By Proposition 12, every point of $\overline{\mathcal{R}}$ is a limit of payoffs $F_{\delta_m}(s_m)$ of pure path-protecting profiles s_m at $\delta_m \uparrow 1$. By Proposition 7, each s_m is neutrally stable, hence Nash, and $F_{\delta_m}(s_m)$ is the average payoff vector of the Nash state $\mathbf{e}_{s_m}$. This proves the first inclusion.

For the second, let x be a Nash state at some $\delta \in (0, 1)$ and let σ be the profile of initial-action distributions induced by its pool. By Corollary 4, $\left(E^{(k)}(x, x)\right)_{k=1}^{2} \in \mathcal{F}$ and $E^{(k)}(x, x) \geq \widetilde{b}_k(\sigma) \geq \widetilde{m}_k$, the last inequality by the definition of the mixed minmax payoffs in Remark 2. Every Nash payoff vector, at every δ, therefore lies in the closed set $\{v \in \mathcal{F} : v_k \geq \widetilde{m}_k,\ k = 1, 2\}$, which also contains every limit of such vectors.

If $m_k = \widetilde{m}_k$ for $k = 1, 2$, the two bounds coincide and equal $\overline{\mathcal{R}}$. *Q.E.D.*

With two roles and under the nonemptiness hypothesis of Corollary 16, a limit of Nash payoffs that is not a limit of path-protecting payoffs must fall below at least one pure minmax bound while satisfying both mixed minmax bounds.

REMARK 3—First-stage feasibility does not establish equilibrium: In Example 2, the relaxed region (18) is nevertheless nonempty. Taking $\sigma_1(1) = 1/2$, $\sigma_2(1) = 0$ and $\sigma_3(1) = 1$ gives $\widetilde{b}(\sigma) = (1, 2, 2)$, while $U(0, 1, 1) = (2, 3, 3)$ exceeds it by one in every coordinate. Only the initial-action distribution of population 1 is nondegenerate in this witness. The witness establishes only strict first-stage feasibility, not existence of a polymorphic equilibrium: such existence would also require payoff equality among coexisting strategies and all continuation incentives, and neutral stability would impose additional conditions. The characterization in Proposition 12 concerns pure path-protecting profiles. A characterization of all Nash payoff limits is established here only for two roles when $\mathcal{R}$ is nonempty and the pure and mixed minmax payoffs coincide (Corollary 16).

5. EVOLUTIONARY STABILITY AND POPULATION STRUCTURE

The preceding results identify sustainable payoffs and construct neutrally stable profiles. We now show that these profiles are Lyapunov stable under strategy revision, and illustrate why allowing frequencies to vary independently across roles matters for that conclusion.

5.1. *Replicator dynamics and Lyapunov stability*

For finite sets S_k and $X = \prod_k \Delta_{S_k}$, the multipopulation replicator dynamics is

$$\dot{x}_{i_k}^{(k)} = x_{i_k}^{(k)} \left[F_{i_k}^{(k)}(x) - E^{(k)}(x, x) \right], \qquad k = 1, \ldots, n, \quad i_k \in S_k. \tag{19}$$

A strategy's share grows when its payoff exceeds the current average in its own role. Population sizes and roles remain fixed; it is the frequencies of contingent plans within each role that evolve. The vector field is Lipschitz by Section 2.3. Population masses remain one and zero coordinates remain zero, so X and its faces are forward invariant. Compactness then gives a unique forward solution for every initial state in X. This is the usual multipopulation replicator equation (Taylor and Jonker, 1978, Sandholm, 2010), with payoffs evaluated at the consistent stationary pool. Its use here retains Assumption 1; it does not model the transient adjustment of the pool after a change in population composition.

Lyapunov stability means that sufficiently small perturbations of a state remain small over time. The following result connects this dynamic property to neutral stability.

PROPOSITION 17: *Let $x \in \mathbb{D}$ be neutrally stable. For every choice of finite strategy sets $S_k \subseteq \Omega^{(k)}$ containing $\operatorname{supp} x^{(k)}$, $k = 1, \ldots, n$, the state x is Lyapunov stable under the multipopulation replicator dynamics on $X := \prod_{k=1}^{n} \Delta_{S_k}$.*

The proof in Appendix A adapts the relative-entropy Lyapunov argument of Thomas (1985) to the multipopulation setting; see also Sandholm (2010). It uses Lipschitz continuity and the two conditions in Definition 2, without requiring linear payoff functions (Bomze and Weibull, 1995), as payoffs are nonlinear in the population state. The proposition concerns every finite collection of strategies containing the residents. Stability holds for each collection and each $\delta < 1$ separately; it asserts neither attraction nor a neighborhood uniform across collections or as $\delta \uparrow 1$.

COROLLARY 18—Dynamic stability of protected profiles: *The monomorphic state of every path-protecting profile is Lyapunov stable under (19) on every finite strategy space containing it.*

PROOF: Combine Propositions 7 and 17. *Q.E.D.*

REMARK 4—Positive role weights: For fixed weights $\lambda_k > 0$, put

$$E_\lambda(y, x) := \sum_{k=1}^{n} \lambda_k E^{(k)}(y, x),$$

and call x λ-neutrally stable if it satisfies Definition 2 with E_λ in place of E. Since each role's distribution can vary independently, Nash states and best-response sets are unchanged by these weights. At a Nash state, $E^{(k)}(x, x) - E^{(k)}(y, x) \geq 0$ for every k; hence equality in the weighted aggregate is equivalent to equality in every role.

For each fixed $\lambda > 0$, λ-neutral stability implies the conclusion of Proposition 17 for the unchanged dynamics (19). On a fixed finite X, apply the local payoff-dominance argument in its proof to the Lipschitz map with components $\lambda_k F_i^{(k)}$. This gives $E_\lambda(y, y) - E_\lambda(x, y) \leq$

0 near x. On a relative neighborhood where $y_i^{(k)}>0$ whenever $x_i^{(k)}>0$, use

$$H_{x,\lambda}(y):=\sum_{k=1}^{n}\lambda_k\sum_{i:x_i^{(k)}>0}x_i^{(k)}\log\frac{x_i^{(k)}}{y_i^{(k)}}.$$

This continuous function is nonnegative and vanishes only at x; along (19), $\dot H_{x,\lambda}=E_\lambda(y,y)-E_\lambda(x,y)$. The same sublevel-set argument proves Lyapunov stability.

The main definition retains equal weights; its second condition can depend on that choice for other states.

REMARK 5—Payoff-equivalent strategies and the stationary pool: Let s be path-protecting with path π. Fix finite sets S_k containing s_k, and let $N_k\subseteq S_k$ collect the strategies that follow π against s_{-k}. The induction in the proof of Proposition 7 shows that every profile with components in $N_1,\dots,N_n$ has path π, hence $T_I=\infty$. On the face $\prod_k\Delta_{N_k}$, the duration tensor is therefore constant, the pool–population map is the identity by Section 2.2, and every state is a rest point of (19) at which role k earns $F_s^{(k)}$.

Thus a stable outcome need not select a unique contingent plan: strategies that agree along the protected path can coexist while prescribing different responses to histories that never occur. Every strategy in $S_k\setminus N_k$ earns strictly less than s_k against the resident state $\mathbf{e}_s$; this strict comparison need not hold uniformly over the entire neutral face.

5.2. *Why the population structure matters*

Independent variation across roles can change stability even when the stage game is symmetric. Consider a voluntarily repeated Prisoner's Dilemma (VRPD) with symmetric payoffs $(R,S,T,P)=(3,0,4,1)$, where $R=U_{CC}$, $S=U_{CD}$, $T=U_{DC}$, and $P=U_{DD}$, and let $\delta=0.9$. This is a different stage game from the asymmetric Prisoner's Dilemma in Table I.

There are two strategies, *CSL* and *DSL* (Izquierdo, Izquierdo, and Vega-Redondo, 2014). CSL always cooperates, whereas DSL always defects; both stay if their partner cooperates and leave otherwise. Panel (a) of Figure 1 shows the single-population dynamics. Besides the all-DSL sink, there is a partially cooperative stable equilibrium whose basin covers more than 90% of the interval of population shares.

In panel (b), the same game, strategies, and continuation probability are used in two populations. The single-population dynamics coincide with the dynamics along the antidiagonal, where the two populations have identical strategy frequencies. The partially cooperative equilibrium becomes a saddle: its Jacobian has a negative eigenvalue along the diagonal and a positive transverse eigenvalue. The all-DSL equilibrium remains asymptotically stable. The two off-diagonal interior equilibria have purely imaginary Jacobian eigenvalues; this linear classification does not establish nonlinear stability or periodic orbits.

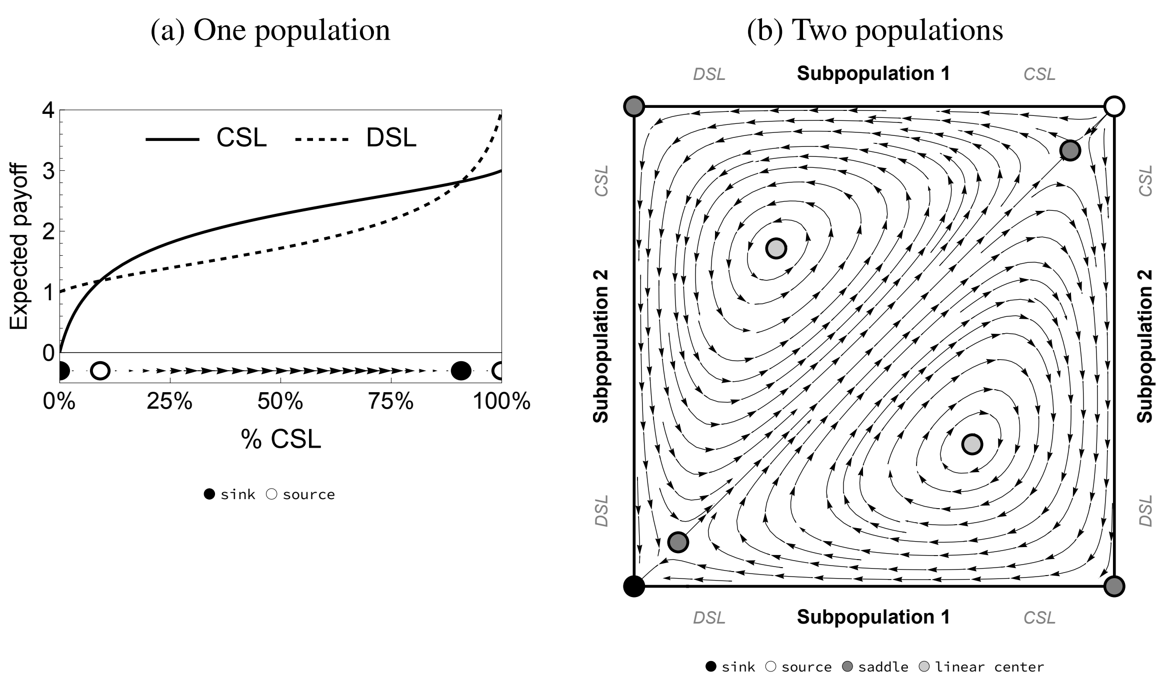

FIGURE 1.—Population structure and stability in the symmetric VRPD with $(R,S,T,P) = (3,0,4,1)$, $\delta = 0.9$, and strategies CSL and DSL. Coordinates are population shares of CSL. Panel (a) has two sinks and two sources. In panel (b), the all-DSL corner is a sink, the all-CSL corner is a source, the other two corners and the two diagonal interior equilibria are saddles, and the two off-diagonal interior equilibria are linear centers. The partially cooperative sink in panel (a) is a saddle in panel (b). Interactive versions: https://doi.org/10.5281/zenodo.22872413 and https://doi.org/10.5281/zenodo.22872689.

5.3. *A protected convention*

For the symmetric Prisoner's Dilemma just introduced, consider a strategy that we call *Leaver* (Izquierdo, Izquierdo, and Boyd, 2026). This strategy defects at the first stage of a new partnership and reacts to the action profile played in a partnership (own action, mate's action) using the mapping $CC \mapsto C$, $CD \mapsto \mathit{break}$, $DC \mapsto \mathit{break}$, and $DD \mapsto C$. A *Leaver–Leaver* partnership generates the sequence of action profiles $(DD, CC, CC, \ldots)$.

EXAMPLE 3—The Leaver profile: For the symmetric Prisoner's Dilemma payoffs $(R, S, T, P) = (3, 0, 4, 1)$, the path $(DD, CC, CC, \ldots)$ gives each player $F_\delta^{(k)} = 1 + 2\delta$, with $b_k(DD) = 1$, $b_k(CC) = 4$, and $V_t^{(k)} = 3$ for $t \geq 1$. Against Leaver, every first unilateral action departure from this path produces CD or DC, after which the resident leaves; a voluntary breakup also ends the partnership. Thus the first-departure argument in the proof of Proposition 8 applies to the actual Leaver strategy, even though its choices after all possible off-path histories need not coincide with those of the canonical deviation-leaving profile. The first-stage inequality is strict for every $\delta \in (0, 1)$, and all later inequalities reduce to

$$4(1-\delta) + \delta(1+2\delta) < 3 \quad \iff \quad (1-\delta)(1-2\delta) < 0.$$

Hence Leaver–Leaver is path-protecting whenever $\delta > 1/2$, including $\delta = 0.6$ in Figure 2. Its neutral and Lyapunov stability follow from Proposition 7 and Corollary 18.

Figure 2 shows the two-population replicator dynamics for these symmetric payoffs and $\delta = 0.6$, with strategy set $\{\mathit{Leaver}, \mathit{Defector}\}$, where *Defector* always plays D and never leaves. The interior saddle occurs at a Leaver share of approximately 0.247 in each population. This restricted phase portrait illustrates the result; Lyapunov stability of the all-Leaver state holds on every finite strategy space containing Leaver.

6. DISCUSSION AND CONCLUSIONS

The opportunity to restart a relationship links the incentive constraints of its members. Under anonymous rematching and strategies based on the current partnership's history, every new match must expose all roles to one common profile of initial actions. A low restarting payoff for each role separately need not be compatible with such a profile. With two roles, the pure minmax bounds can be attained simultaneously. If some feasible payoff

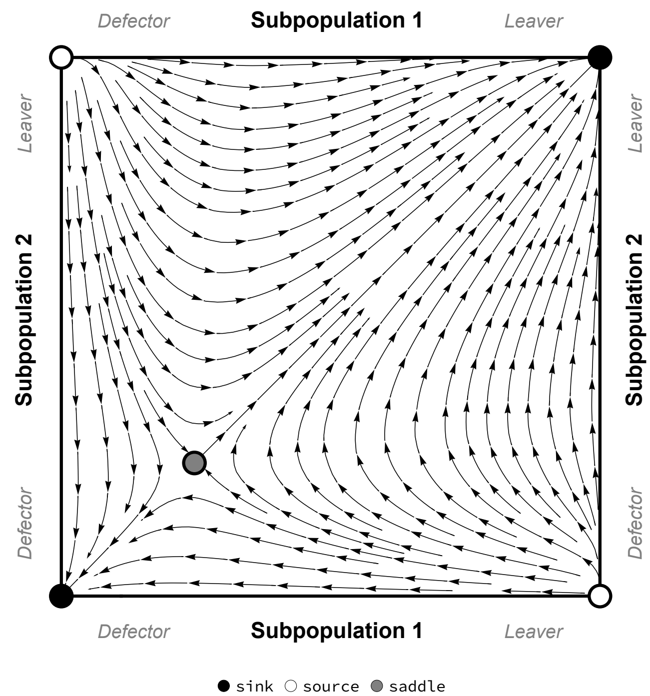

FIGURE 2.—A protected convention in the symmetric VRPD with $(R,S,T,P)=(3,0,4,1)$, $\delta=0.6$, and strategies Leaver and Defector. Coordinates are population shares of Leaver. Leaver starts with D, cooperates after DD or CC, and leaves after CD or DC; Defector always plays D and never leaves. The all-Defector and all-Leaver corners are sinks, the other two corners are sources, and the interior equilibrium at approximately $(0.247, 0.247)$ is a saddle.

vector strictly exceeds both bounds and the pure and mixed minmax payoffs coincide, the limits of all Nash payoffs coincide with those of pure path-protecting profiles. With three roles, even an interior, strictly individually rational payoff vector can remain a fixed distance from every Nash payoff.

For pure path-protecting profiles, the limiting payoff set is exactly $\overline{\mathcal{R}}$: the closure of feasible vectors strictly exceeding all one-stage best-response payoffs for some common initial profile. An initial phase deters voluntary restarting while its cost to complying players vanishes as exogenous separations become rare. This gives payoff approximation without a pure Nash initial profile, a full-dimensional feasible set, or a nonequivalent-utilities condition. The result characterizes this class of profiles; first-stage feasibility alone does not establish a polymorphic equilibrium.

Stationary matching also makes evolutionary analysis possible. The duration tensor determines a unique pool and Lipschitz payoffs on each fixed finite strategy space. Path-protecting profiles are neutrally stable for every positive weighting of roles and Lyapunov stable under replicator dynamics. Different contingent plans can coexist along the same protected path, so stability need not select a unique behavioral rule.

An explicit joint dynamics of matching and strategy revision would allow the stationary approximation to be tested. Behavioral errors could also affect whether separation continues to deter deviations. Separation costs would change the restarting comparison itself; related two-player analyses include Enquist and Leimar (1993), Gutiérrez-Mielgo, Izquierdo, and Izquierdo (2026), Zhang and Li (2026). Extending those analyses requires specifying who bears the costs and how they enter the stationary payoff. In each of these directions, the comparison between continuing an established relationship and restarting elsewhere remains the mechanism to be evaluated.

APPENDIX A: PROOFS

A.1. *The pool–population map*

Fix the finite strategy sets and the notation of Section 2.2. The following scaling equations characterize the inverse map.

PROPOSITION 19: *For every $x \in X$, the tensor-scaling equations*

$$d^{(k)}_{i_k} \sum_{i_{-k}} B_I \prod_{\ell \neq k} d^{(\ell)}_{i_\ell} = x^{(k)}_{i_k} \qquad (k = 1, \ldots, n,\ i_k \in S_k) \tag{20}$$

admit nonnegative solutions $d^{(k)} \in \mathbb{R}^{S_k}$. Every such solution has $\|d^{(k)}\|_1 > 0$ and characterizes the unique pool state $p = f^{-1}(x)$ by

$$p^{(k)} = \frac{d^{(k)}}{\|d^{(k)}\|_1}, \qquad k = 1, \ldots, n.$$

The scaling factors need not be unique: all solutions are related by $d^{(k)} \mapsto c_k d^{(k)}$, where $c_k > 0$ and $\prod_k c_k = 1$. Indeed, every solution has the same normalization $p = f^{-1}(x)$, and summing (20) gives

$$1 = Z(d) = Z(p)\prod_{k=1}^{n}\|d^{(k)}\|_1.$$

The ratios of these strictly positive norms for any two solutions give the stated rescaling; conversely, such rescalings preserve (20).

The proofs use positive tensor scaling (Friedland, 2011), followed by a derivative and compactness argument. The tensor-scaling literature also includes Bapat (1982), Franklin and Lorenz (1989); see Idel (2016) for a survey. For $n = 2$, (20) is the matrix-scaling problem

$$d^{(1)} \circ (Bd^{(2)}) = x^{(1)}, \qquad d^{(2)} \circ (B^T d^{(1)}) = x^{(2)},$$

where $\circ$ denotes entrywise multiplication. Classical matrix-scaling results include Sinkhorn and Knopp (1967); see also Idel (2016), Rothblum and Schneider (1989). No separate matrix proof is needed.

BIJECTIVITY IN PROPOSITION 1 AND PROOF OF PROPOSITION 19. For a nonnegative vector $z = (z_i)_{i\in S}$, write $\operatorname{supp} z := \{i \in S : z_i > 0\}$.

The polynomials in (1) satisfy

$$d_{i_k}^{(k)} M_{i_k}^{(k)}(d) = \sum_{i_{-k}} q_I(d), \qquad \sum_{i_k\in S_k} d_{i_k}^{(k)} M_{i_k}^{(k)}(d) = Z(d). \tag{21}$$

For $p \in X$, both $Z(p)$ and each $M_{i_k}^{(k)}(p)$ are convex combinations of entries of B. As $B_I \geq 1$,

$$Z(p) \geq 1 > 0, \qquad M_{i_k}^{(k)}(p) \geq 1 > 0 \qquad (p \in X). \tag{22}$$

The *matching-pool map* is

$$f_{i_k}^{(k)}(p) := \frac{p_{i_k}^{(k)} M_{i_k}^{(k)}(p)}{Z(p)} = \frac{\sum_{i_{-k}} q_I(p)}{Z(p)}. \tag{23}$$

Equations (21)–(22) show that $f : X \to X$ is well defined and that

$$\operatorname{supp} f^{(k)}(p) = \operatorname{supp} p^{(k)} \qquad (p \in X,\ k = 1, \dots, n). \tag{24}$$

The (k, i_k)-*slice sum* of a tensor C is $\sum_{i_{-k}} C_I$. A tensor C is *positively diagonally equivalent* to a tensor B on the same index set if $C_I = B_I \prod_k d^{(k)}_{i_k}$ for strictly positive factors $d^{(k)}_{i_k}$.

Fix $x \in X$ and put

$$\underline{S}_k := \operatorname{supp} x^{(k)}, \qquad \mathcal{S} := \underline{S}_1 \times \cdots \times \underline{S}_n, \qquad \underline{S}_{-k} := \prod_{\ell \neq k} \underline{S}_\ell.$$

Each $\underline{S}_k$ is nonempty.

Surjectivity, and existence of scaling solutions. The restricted tensor $B|_{\mathcal{S}}$ is strictly positive and therefore has no zero slice. The tensor $Y_I := \prod_k x^{(k)}_{i_k}$, $I \in \mathcal{S}$, is also strictly positive and has slice sums

$$\sum_{i_{-k} \in \underline{S}_{-k}} Y_I = x^{(k)}_{i_k} \prod_{\ell \neq k} \Big(\sum_{i_\ell \in \underline{S}_\ell} x^{(\ell)}_{i_\ell} \Big) = x^{(k)}_{i_k}.$$

Thus Y has the same zero pattern as $B|_{\mathcal{S}}$ and realizes the prescribed strictly positive slice-sum vectors $x^{(k)}|_{\underline{S}_k}$. By Friedland (2011, Corollary 1.2), there are strictly positive factors $d^{(k)}_{i_k}$, $i_k \in \underline{S}_k$, such that

$$C_I := B_I \prod_{k=1}^{n} d^{(k)}_{i_k} \quad (I \in \mathcal{S}), \qquad \sum_{i_{-k} \in \underline{S}_{-k}} C_I = x^{(k)}_{i_k}. \tag{25}$$

In particular, C has total mass one. Extend each $d^{(k)}$ by zero outside $\underline{S}_k$ and set

$$p^{(k)} := d^{(k)} / \|d^{(k)}\|_1, \qquad \lambda := \prod_{k=1}^{n} \|d^{(k)}\|_1.$$

Then $p \in X$, $C_I = \lambda q_I(p)$ on $\mathcal{S}$, and $q_I(p) = 0$ outside $\mathcal{S}$. Summing gives $1 = \lambda Z(p)$, so the tensor $q(p)/Z(p)$ has slice sums x by (25). Equation (23) therefore gives $f(p) = x$, proving surjectivity. The zero-extended factors $d^{(k)}$ also solve (20), by (25).

Injectivity. Suppose $f(p) = f(\widehat{p}) = x$. By (24), both $p^{(k)}$ and $\widehat{p}^{(k)}$ have support $\underline{S}_k$. The tensors on $\mathcal{S}$ given by

$$G_I := \frac{q_I(p)}{Z(p)}, \qquad \widehat{G}_I := \frac{q_I(\widehat{p})}{Z(\widehat{p})}$$

are positively diagonally equivalent to $B|_{\mathcal{S}}$, with factors $Z(p)^{-1/n} p^{(k)}_{i_k}$ and $Z(\widehat{p})^{-1/n} \widehat{p}^{(k)}_{i_k}$, respectively. They have the same strictly positive slice-sum vectors $x^{(k)}|_{\underline{S}_k}$, each of total mass one. Since $B|_{\mathcal{S}}$ has no zero slice, the uniqueness assertion in Friedland (2011, Theorem 1.1) gives $G = \widehat{G}$. Dividing entrywise by $B_I > 0$ yields

$$\frac{\prod_k p^{(k)}_{i_k}}{Z(p)} = \frac{\prod_k \widehat{p}^{(k)}_{i_k}}{Z(\widehat{p})} \qquad (I \in \mathcal{S}). \tag{26}$$

Summing over $\mathcal{S}$ gives $Z(p) = Z(\widehat{p})$, because each product tensor has total mass one. Cancelling this common denominator in (26) and summing over all coordinates except i_k gives $p^{(k)}_{i_k} = \widehat{p}^{(k)}_{i_k}$ on $\underline{S}_k$. Both vectors vanish outside $\underline{S}_k$, so $p = \widehat{p}$.

The inverse formula. Let d be any nonnegative solution of (20). Summing its equations over i_k for any fixed k gives $Z(d) = \sum_{i_k} x^{(k)}_{i_k} = 1$ by (21). Hence no $d^{(k)}$ is $\mathbf{0}$, since otherwise every $q_I(d)$ would vanish. Thus $\|d^{(k)}\|_1 > 0$. Set $\widetilde{p}^{(k)} := d^{(k)}/\|d^{(k)}\|_1$ and $\nu := \prod_k \|d^{(k)}\|_1$, so $\widetilde{p} \in X$ and $q_I(d) = \nu\, q_I(\widetilde{p})$, $Z(d) = \nu\, Z(\widetilde{p})$. Using (20) and $Z(d) = 1$ yields

$$f^{(k)}_{i_k}(\widetilde{p}) = \frac{\sum_{i_{-k}} q_I(d)}{Z(d)} = x^{(k)}_{i_k}.$$

By bijectivity, $\widetilde{p} = f^{-1}(x)$, proving the result.

Q.E.D.

Derivatives are taken along the affine hull of X. Define

$$\mathcal{A} := \operatorname{aff}(X) = \Big\{ p \in \prod_{k=1}^{n} \mathbb{R}^{S_k} : \sum_{i_k \in S_k} p^{(k)}_{i_k} = 1 \text{ for every } k \Big\},$$

$$\mathcal{E} := \Big\{ w \in \prod_{k=1}^{n} \mathbb{R}^{S_k} : \sum_{i_k \in S_k} w^{(k)}_{i_k} = 0 \text{ for every } k \Big\}.$$

Then $\mathcal{E}$ is the translation space of $\mathcal{A}$, and X is compact and convex. By (22), the set $\mathcal{O} := \{p \in \mathcal{A} : Z(p) > 0\}$ is an open neighborhood of X in $\mathcal{A}$.

LEMMA 20: *Formula (23) defines a C^1 map $\widetilde{f} : \mathcal{O} \to \mathcal{A}$ extending f. For every $p \in X$, its derivative $D\widetilde{f}(p) : \mathcal{E} \to \mathcal{E}$ is injective.*

PROOF: The components of $\widetilde{f}$ are quotients of polynomials with nonvanishing denominator on $\mathcal{O}$, so $\widetilde{f}$ is C^1. The polynomial identity (21) gives $\sum_{i_k} \widetilde{f}^{(k)}_{i_k} = 1$ for every k; thus $\widetilde{f}(\mathcal{O}) \subset \mathcal{A}$ and $D\widetilde{f}(p) : \mathcal{E} \to \mathcal{E}$.

Fix $p \in X$. To prove injectivity, let $w \in \mathcal{E}$ satisfy $D\widetilde{f}(p)w = \mathbf{0}$. We show that $w = \mathbf{0}$. Set

$$\underline{S}_k := \operatorname{supp} p^{(k)}, \qquad \mathcal{S} := \prod_{k=1}^{n} \underline{S}_k, \qquad \underline{S}_{-k} := \prod_{\ell \neq k} \underline{S}_\ell, \qquad \mu := \frac{DZ(p)w}{Z(p)}.$$

By the quotient rule applied to (23), the equations $D\widetilde{f}(p)w = \mathbf{0}$ are equivalent to the linearized slice-sum equations

$$\sum_{i_{-k}} Dq_I(p)w = \mu \sum_{i_{-k}} q_I(p) \qquad (k = 1, \ldots, n,\ i_k \in S_k), \tag{27}$$

where the product rule gives

$$Dq_I(p)w = B_I \sum_{r=1}^{n} w^{(r)}_{i_r} \prod_{\ell \neq r} p^{(\ell)}_{i_\ell}. \tag{28}$$

If $i_k \notin \underline{S}_k$, then (27) reduces, using (28), to $M^{(k)}_{i_k}(p) w^{(k)}_{i_k} = 0$. Positivity in (22) implies

$$w^{(k)}_{i_k} = 0 \qquad (i_k \notin \underline{S}_k). \tag{29}$$

Consequently, if $I \notin \mathcal{S}$, both $q_I(p)$ and $Dq_I(p)w$ vanish: choose an inactive coordinate i_k; in (28), the term $r = k$ contains $w^{(k)}_{i_k} = 0$ and every other term contains $p^{(k)}_{i_k} = 0$. Thus the sums in (27) for $i_k \in \underline{S}_k$ may be restricted to $\underline{S}_{-k}$.

For the active coordinates and profiles, put

$$v^{(k)}_{i_k} := \frac{w^{(k)}_{i_k}}{p^{(k)}_{i_k}} \quad (i_k \in \underline{S}_k), \qquad s_I := \sum_{k=1}^{n} v^{(k)}_{i_k} - \mu \quad (I \in \mathcal{S}).$$

On $\mathcal{S}$, we have $q_I(p) > 0$ and $Dq_I(p)w = q_I(p) \sum_k v^{(k)}_{i_k}$. Hence the active equations are

$$\sum_{i_{-k} \in \underline{S}_{-k}} q_I(p) s_I = 0 \qquad (k = 1, \ldots, n,\ i_k \in \underline{S}_k). \tag{30}$$

Summing (30) over $i_k \in \underline{S}_k$ for any fixed k gives

$$\sum_{I \in \mathcal{S}} q_I(p) s_I = 0. \tag{31}$$

Now multiply each equation in (30) by $v^{(k)}_{i_k}$ and sum over k and i_k. Using (31), we obtain

$$0 = \sum_{I \in \mathcal{S}} q_I(p) s_I \sum_{k=1}^{n} v^{(k)}_{i_k} = \sum_{I \in \mathcal{S}} q_I(p) s_I (s_I + \mu) = \sum_{I \in \mathcal{S}} q_I(p) s_I^2.$$

All the weights are positive, so $s_I = 0$ for every $I \in \mathcal{S}$. Comparing two profiles that differ only in their k-th coordinate shows that $v^{(k)}$ is constant on $\underline{S}_k$; denote its value by α_k. Together with (29), this gives $w^{(k)} = \alpha_k p^{(k)}$. Since $w \in \mathcal{E}$ and $p \in X$,

$$0 = \sum_{i_k \in S_k} w^{(k)}_{i_k} = \alpha_k \sum_{i_k \in S_k} p^{(k)}_{i_k} = \alpha_k.$$

Thus $w = \mathbf{0}$, as required. *Q.E.D.*

REGULARITY IN PROPOSITION 1. Use Euclidean norms and their induced operator norms. Since all norms in finite dimension are equivalent, this choice does not affect the bi-Lipschitz property.

If X is a singleton, the conclusion is immediate. Otherwise, let $\widetilde{f}$ be the extension in Lemma 20. We prove that there are constants $c > 0$ and $L < \infty$ such that

$$c\|p - p'\| \le \|f(p) - f(p')\| \le L\|p - p'\| \qquad (p, p' \in X). \tag{32}$$

Continuity of $D\widetilde{f}$ and compactness of X give

$$L := \max_{p \in X} \|D\widetilde{f}(p)\|_{\mathrm{op}} < \infty.$$

For $p, p' \in X$, convexity and the fundamental theorem of calculus yield

$$f(p) - f(p') = \int_0^1 D\widetilde{f}\big(p' + t(p - p')\big)(p - p')\, dt. \tag{33}$$

Taking norms proves the upper bound in (32).

If the lower bound failed, there would be distinct $p_m, p'_m \in X$ with

$$\frac{\|f(p_m) - f(p'_m)\|}{\|p_m - p'_m\|} \longrightarrow 0. \tag{34}$$

Set $u_m := (p_m - p'_m)/\|p_m - p'_m\| \in \mathcal{E}$. By compactness of X and of the unit sphere of $\mathcal{E}$, pass to a subsequence such that $p_m \to p$, $p'_m \to p'$ and $u_m \to u$, where $p, p' \in X$ and $\|u\| = 1$. Since X is bounded, (34) implies $f(p_m) - f(p'_m) \to \mathbf{0}$. Continuity and Proposition 1 then give $f(p) = f(p')$ and hence $p = p'$.

Dividing (33) by $\|p_m - p'_m\|$, we obtain

$$\frac{f(p_m) - f(p'_m)}{\|p_m - p'_m\|} = \int_0^1 D\widetilde{f}\big(p'_m + t(p_m - p'_m)\big)u_m\,dt \longrightarrow D\widetilde{f}(p)u.$$

Indeed, the segment points converge to p uniformly in $t \in [0,1]$, $D\widetilde{f}$ is continuous at p, and $u_m \to u$. By (34), the left-hand side converges to $\mathbf{0}$. Thus $D\widetilde{f}(p)u = \mathbf{0}$ with $u \in \mathcal{E}$ and $\|u\| = 1$, contradicting Lemma 20.

This proves (32). By Proposition 1, f is bijective, and the lower bound gives

$$\|f^{-1}(x) - f^{-1}(y)\| \le c^{-1}\|x - y\| \qquad (x, y \in X).$$

Hence f and f^{-1} are Lipschitz, and in particular continuous. The constants may depend on B and X. *Q.E.D.*

A.2. *Construction of path-protecting profiles*

PROOF OF PROPOSITION 10. Put $r := \bar{U}_{\Phi_p}$, $b_k := b_k(a^0)$ and $U_k^{\max} := \max_{a \in A} U^{(k)}(a)$. Let $T_f \ge 0$ and $T_p \ge 1$ be the lengths of Φ_f and Φ_p, respectively, and set $H := T_f + T_p$. Since $r_k > b_k$, choose one integer $M \ge 1$ such that, for every k,

$$\alpha_k := \frac{Mb_k + HU_k^{\max}}{M + H} < r_k. \tag{35}$$

For example, it suffices that $M(r_k - b_k) > H(U_k^{\max} - r_k)$ for every k. Fix this M and the deviation-leaving profile s defined by (16).

Consider a unilateral deviation in population k. Until the first departure from the prescribed path, the other players use their prescribed actions. An action deviation ends the partnership after that stage because the resident players then leave; a decision to break ends it after its last completed stage. Thus a deviation either follows the resident path forever or generates a finite partnership cycle of some length $t \geq 1$, which restarts with every new partnership.

First consider cycles with $t \leq M + H$. If $t \leq M$, every stage payoff to the deviator is at most b_k. If $M < t \leq M + H$, the first M payoffs are at most b_k and the remaining $t - M$ payoffs are at most $U_k^{\max}$. Consequently, the undiscounted mean of such a cycle is at most

$$\frac{Mb_k + (t - M)U_k^{\max}}{t} \leq \alpha_k,$$

and the same bound holds when $t \leq M$, since $b_k \leq \alpha_k$. For any fixed finite cycle, its normalized discounted value extends continuously to $\delta = 1$, with value equal to its undiscounted mean. There are only finitely many relevant cycles with $t \leq M + H$: their pre-deviation histories are fixed, and their possible terminal actions belong to finite action sets; compliance followed by a voluntary breakup is the case in which the terminal action is the prescribed one.

The resident path has a prefix of length $\ell := M + T_f$ followed by a periodic tail. Explicitly,

$$F_\delta^{(k)}(s) = (1 - \delta) \sum_{t=1}^{\ell} \delta^{t-1} U^{(k)}(a_s^{[t]}) + \delta^\ell \frac{\sum_{j=1}^{T_p} \delta^{j-1} U^{(k)}(\Phi_p^{[j]})}{\sum_{j=1}^{T_p} \delta^{j-1}} \longrightarrow r_k.$$

In view of (35), continuity and the finiteness of the cycle list give a common $\bar{\delta} \in (0, 1)$ such that every deviation ending by stage $M + H$ gives player k strictly less than $F_\delta^{(k)}(s)$, for every k and every $\delta > \bar{\delta}$.

The resident path has a prefix of length $M + T_f$ and a period of length T_p. By Proposition 8 and Lemma 9, strict unprofitability of the best-response-and-leave cycles ending by stage $M + H$ implies strict unprofitability at every later stage as well. This includes breakups at boundaries between successive repetitions.

Thus every unilateral departure is strictly unprofitable, proving path protection by Definition 3. Neutral stability follows from Proposition 7. *Q.E.D.*

A.3. *Neutral stability and Lyapunov stability*

PROOF OF PROPOSITION 17. Fix the sets S_k and write states as concatenated population vectors. We first show that

$$E(y,y) - E(x,y) \le 0 \quad \text{for every } y \text{ in a relative neighborhood of } x \text{ in } X. \tag{36}$$

This argument uses only Lipschitz continuity of the payoff map F on X, not linearity of the payoffs.

For each population, define

$$\beta_k := \max_{i \in S_k} F_i^{(k)}(x), \qquad \mathcal{B}_k := \{i \in S_k : F_i^{(k)}(x) = \beta_k\}.$$

Since x is Nash, $\operatorname{supp} x^{(k)} \subseteq \mathcal{B}_k$ and $\sum_{i \in S_k} x_i^{(k)} F_i^{(k)}(x) = \beta_k$. For $y \in X$, let

$$r_k := \sum_{i \in S_k \setminus \mathcal{B}_k} y_i^{(k)}, \qquad r := \sum_{k=1}^{n} r_k, \qquad \widetilde{y}_i^{(k)} := \begin{cases} y_i^{(k)} + r_k x_i^{(k)}, & i \in \mathcal{B}_k, \\ 0, & i \notin \mathcal{B}_k. \end{cases}$$

Then $\widetilde{y} \in X$ is a best response to x. Consequently, neutral stability gives

$$(\widetilde{y} - x)^T F(\widetilde{y}) \le 0. \tag{37}$$

With $\|\cdot\|_1$ denoting the sum of absolute values over all population coordinates, the construction also gives

$$\|y - \widetilde{y}\|_1 = 2r, \qquad \|\widetilde{y} - x\|_1 \le \|y - x\|_1. \tag{38}$$

For the inequality, in each population use the triangle inequality on $\mathcal{B}_k$ and observe that the mass of $y^{(k)}$ outside $\mathcal{B}_k$ is r_k, whereas $x^{(k)}$ has no mass there.

If $\mathcal{B}_k = S_k$ for every k, every $y \in X$ is a best response to x, so (36) follows directly from neutral stability. Otherwise, the finite set of non-best-response strategies has a strictly positive payoff gap

$$\gamma := \min_{\substack{k=1,\dots,n \\ i \in S_k \setminus \mathcal{B}_k}} \big(\beta_k - F_i^{(k)}(x)\big) > 0.$$

By continuity, in a relative neighborhood U of x,

$$F_i^{(k)}(y) - \sum_{j \in S_k} x_j^{(k)} F_j^{(k)}(y) \le -\gamma/2 \qquad (i \in S_k \setminus \mathcal{B}_k).$$

Hence, for $y \in U$,

$$(y - \widetilde{y})^T F(y) = \sum_{k=1}^{n} \sum_{i \in S_k \setminus \mathcal{B}_k} y_i^{(k)} \left(F_i^{(k)}(y) - \sum_{j \in S_k} x_j^{(k)} F_j^{(k)}(y) \right) \le -\gamma r/2.$$

Choose $K \ge 1$ such that $\|F(y) - F(z)\|_\infty \le K\|y - z\|_1$ for all $y, z \in X$. Using (37) and (38), we obtain

$$\begin{aligned} E(y,y) - E(x,y) &= (y - \widetilde{y})^T F(y) + (\widetilde{y} - x)^T \big(F(y) - F(\widetilde{y})\big) + (\widetilde{y} - x)^T F(\widetilde{y}) \\ &\le -\gamma r/2 + 2Kr\|y - x\|_1. \end{aligned}$$

After shrinking U so that $\|y - x\|_1 < \gamma/(4K)$, this proves (36).

We now apply the relative-entropy Lyapunov argument (Thomas, 1985, Sandholm, 2010). On a relative neighborhood of x where $y_i^{(k)} > 0$ whenever $x_i^{(k)} > 0$, define

$$H_x(y) := \sum_{k=1}^{n} \sum_{i : x_i^{(k)} > 0} x_i^{(k)} \log \frac{x_i^{(k)}}{y_i^{(k)}}.$$

This function is continuous, nonnegative, and vanishes exactly at x. Along the multipopulation replicator dynamics,

$$\dot{H}_x(y) = \sum_{k=1}^{n} \left(\sum_{i \in S_k} y_i^{(k)} F_i^{(k)}(y) - \sum_{i \in S_k} x_i^{(k)} F_i^{(k)}(y) \right) = E(y,y) - E(x,y).$$

Thus $\dot{H}_x \le 0$ near x by (36). The positive minimum of H_x on the boundary of a sufficiently small relative ball about x prevents trajectories starting in a smaller sublevel set from leaving that ball. This proves Lyapunov stability. If X is a singleton, the conclusion is immediate. *Q.E.D.*

A.4. *Finite breakup periods*

PROOF OF LEMMA 5. Let $T := T_s < \infty$. If $a_s^{[T]}$ is not a stage-game Nash profile, there is a role k with $b_k(a_s^{[T]}) > U^{(k)}(a_s^{[T]})$. Consider a mutant that follows the resident through stage $T-1$, plays a best response at stage T, and then leaves. Its partnership length remains T, so its payoff improvement is

$$\frac{(1-\delta)\delta^{T-1}}{1-\delta^T}\left[b_k(a_s^{[T]}) - U^{(k)}(a_s^{[T]})\right] > 0.$$

This contradicts the Nash property of $\mathbf{e}_s$. *Q.E.D.*

PROOF OF LEMMA 6. Let $T := T_s < \infty$, suppose all residents choose to leave after stage T, and put $Q := E(\mathbf{e}_s, \mathbf{e}_s)$. By (3), Q is a convex combination, with strictly positive weights, of the joint stage payoffs along the finite resident path. If some such payoff is less than J, then $Q < J$.

Choose $a^* \in A$ with $\sum_k U^{(k)}(a^*) = J$. For each role k, construct s'_k to prescribe the resident's action at every on-path history of length less than T and to prescribe $a^{*,(k)}$ after every history of the form $(a_s^{[1,T]}, (a^*)^j)$, $j \geq 0$. Specify arbitrary choices after all remaining histories and put $s' := (s'_1, \ldots, s'_n)$. A unilateral s'_k mutant matched with residents in all other roles still separates after stage T, since $n \geq 2$ and those residents choose to leave. Likewise, a resident of role k matched with s'_{-k} ends that partnership after stage T. Thus

$$E(\mathbf{e}_{s'}, \mathbf{e}_s) = E(\mathbf{e}_s, \mathbf{e}_s) = E(\mathbf{e}_s, \mathbf{e}_{s'}) = Q.$$

In contrast, all-mutant partnerships have no endogenous breakup and yield

$$E(\mathbf{e}_{s'}, \mathbf{e}_{s'}) = (1-\delta^T)Q + \delta^T J > Q.$$

Since neutral stability implies Nash, the equality $E(\mathbf{e}_{s'}, \mathbf{e}_s) = E(\mathbf{e}_s, \mathbf{e}_s)$ makes $\mathbf{e}_{s'}$ a best response to $\mathbf{e}_s$; the strict inequality contradicts its neutral stability. Hence $Q = J$, and positivity of the weights implies that every stage on the resident path attains J. *Q.E.D.*

REFERENCES

Abreu, Dilip, Prajit K. Dutta, and Lones Smith (1994), "The folk theorem for repeated games: A NEU condition." *Econometrica*, 62 (4), 939–948. [3]

Bapat, Ravindra (1982), "D_1AD_2 theorems for multidimensional matrices." *Linear Algebra and its Applications*, 48 (C), 437 – 442. [33]

Barclay, Pat and Nichola Raihani (2016), "Partner choice versus punishment in human prisoner's dilemmas." *Evolution and Human Behavior*, 37 (4), 263–271. [4]

Bomze, Immanuel M. and Jörgen W. Weibull (1995), "Does neutral stability imply Lyapunov stability?" *Games and Economic Behavior*, 11, 173–192. [27]

Carmichael, H. Lorne and W. Bentley MacLeod (1997), "Gift giving and the evolution of cooperation." *International Economic Review*, 38 (3), 485–509. [4]

Datta, Saikat (1996), "Building trust." *STICERD - Theoretical Economics Paper Series*, TE/1996/305. [2, 4]

Deb, Joyee, Takuo Sugaya, and Alexander Wolitzky (2020), "The folk theorem in repeated games with anonymous random matching." *Econometrica*, 88 (3), 917–964. [4]

Ellison, Glenn (1994), "Cooperation in the Prisoner's Dilemma with anonymous random matching." *The Review of Economic Studies*, 61 (3), 567–588. [4]

Enquist, Magnus and Olof Leimar (1993), "The evolution of cooperation in mobile organisms." *Animal Behaviour*, 45, 747–757. [32]

Franklin, Joel and Jens Lorenz (1989), "On the scaling of multidimensional matrices." *Linear Algebra and its Applications*, 114-115, 717–735. [33]

Friedland, Shmuel (2011), "Positive diagonal scaling of a nonnegative tensor to one with prescribed slice sums." *Linear Algebra and its Applications*, 434 (7), 1615–1619. [8, 33, 34, 35]

Fudenberg, Drew, David K. Levine, and Eric Maskin (1994), "The folk theorem with imperfect public information." *Econometrica*, 62 (5), 997–1039. [4]

Fudenberg, Drew and Eric Maskin (1986), "The folk theorem in repeated games with discounting or with incomplete information." *Econometrica*, 54 (3), 533–554. [3, 23]

Fujiwara-Greve, Takako and Masahiro Okuno-Fujiwara (2009), "Voluntarily separable repeated prisoner's dilemma." *Review of Economic Studies*, 76 (3), 993–1021. [4]

Ghosh, Parikshit and Debraj Ray (1996), "Cooperation in community interaction without information flows." *The Review of Economic Studies*, 63, 491–519. [4, 5]

Graser, Christopher, Takako Fujiwara-Greve, Julián García, and Matthijs van Veelen (2025), "Repeated games with partner choice." *PLOS Computational Biology*, 21, e1012810. [4]

Gutiérrez-Mielgo, Alejandro, Luis R. Izquierdo, and Segismundo S. Izquierdo (2026), "Games with costly endogenous separation." SSRN working paper, URL https://doi.org/10.2139/ssrn.7110598. [32]

Hayashi, Nahoko and Toshio Yamagishi (1998), "Selective play: Choosing partners in an uncertain world." *Personality and Social Psychology Review*, 2, 276–289. [4]

Idel, Martin (2016), "A review of matrix scaling and Sinkhorn's normal form for matrices and positive maps." arXiv:1609.06349, URL https://arxiv.org/abs/1609.06349. [33]

Izquierdo, Luis R., Segismundo S. Izquierdo, and Robert Boyd (2026), "Successful strategies in the voluntarily repeated prisoner's dilemma." *Journal of Theoretical Biology*, URL https://doi.org/10.1016/j.jtbi.2026.112598. [4, 30]

Izquierdo, Luis R., Segismundo S. Izquierdo, and Fernando Vega-Redondo (2014), "Leave and let leave: A sufficient condition to explain the evolutionary emergence of cooperation." *Journal of Economic Dynamics and Control*, 46, 91–113. [4, 28]

Izquierdo, Segismundo S. and Luis R. Izquierdo (2026), "Stable strategies in repeated games with endogenous separation." SSRN working paper, URL https://doi.org/10.2139/ssrn.7110538. [4, 15]

Izquierdo, Segismundo S., Luis R. Izquierdo, and Fernando Vega-Redondo (2010), "The option to leave: Conditional dissociation in the evolution of cooperation." *Journal of Theoretical Biology*, 267 (1), 76–84. [4]

Kandori, Michihiro (1992), "Social norms and community enforcement." *The Review of Economic Studies*, 59 (1), 63–80. [4]

Kranton, Rachel E. (1996), "The formation of cooperative relationships." *Journal of Law, Economics, and Organization*, 12, 214–233. [4]

Mailath, George J. and Larry Samuelson (2006), *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press, New York. [2]

Maynard Smith, John (1982), *Evolution and the Theory of Games*, Cambridge University Press, Cambridge. [13]

Rand, David G., Samuel Arbesman, and Nicholas A. Christakis (2011), "Dynamic social networks promote cooperation in experiments with humans." *Proceedings of the National Academy of Sciences*, 108, 19193–19198. [4]

Rob, Rafael and Huanxing Yang (2010), "Long-term relationships as safeguards." *Economic Theory*, 43 (2), 143–166. [2]

Rothblum, Uriel G. and Hans Schneider (1989), "Scalings of matrices which have prespecified row sums and column sums via optimization." *Linear Algebra and its Applications*, 114-115, 737–764. [33]

Sandholm, William H. (2010), *Population Games and Evolutionary Dynamics*, The MIT Press, Cambridge, MA. [26, 27, 41]

Schuessler, Rudolf (1989), "Exit threats and cooperation under anonymity." *Journal of Conflict Resolution*, 33, 728–749. [4]

Selten, Reinhard and Peter Hammerstein (1984), "Gaps in Harley's argument on evolutionarily stable learning rules and in the logic of 'tit for tat'." *Behavioral and Brain Sciences*, 7 (1), 115–116. [14]

Sinkhorn, Richard and Paul Knopp (1967), "Concerning nonnegative matrices and doubly stochastic matrices." *Pacific Journal of Mathematics*, 21 (2), 343–348. [33]

Taylor, Peter D. and Leo B. Jonker (1978), "Evolutionary stable strategies and game dynamics." *Mathematical Biosciences*, 40, 145–156. [26]

Thomas, Bernhard (1985), "On evolutionarily stable sets." *Journal of Mathematical Biology*, 22 (1), 105–115. [27, 41]

Vanberg, Viktor J. and Roger D. Congleton (1992), "Rationality, morality, and exit." *American Political Science Review*, 86, 418–431. [4]

Watson, Joel (2002), "Starting small and commitment." *Games and Economic Behavior*, 38 (1), 176–199. [4]

Zhang, Minmin and Cong Li (2026), "The effect of exit cost on the evolution of cooperation in the separable repeated prisoner's dilemma game." *Chaos, Solitons & Fractals*, 209, 118514. [32]