ON SUSTAINABLE EQUILIBRIA

SRIHARI GOVINDAN, RIDA LARAKI, AND LUCAS PAHL

Abstract. Following the ideas laid out in Myerson [10], Hofbauer [5] defined an equilibrium of a game as sustainable if it can be made the unique equilibrium of a game obtained by deleting a subset of the strategies that are inferior replies to it, and then adding others. Hofbauer also formalized Myerson's conjecture about the relationship between the sustainability of an equilibrium and its index: for generic games, an equilibrium is sustainable iff its index is +1. von Schemde and von Stengel [15] proved this conjecture for bimatrix games. This paper shows that the conjecture is true for all finite games.

1. Introduction

Myerson [10] proposes a refinement of Nash equilbria of finite games, which he calls *sustainable equilibria*, based on the hypothesis that most games, even if they are one-shot affairs, should be analyzed not as if they are played in isolation, but rather as particular instances of many plays of such games. Myerson agues that when, say, two members of a society play a Battle-of-Sexes game, if the game has a history in this society, it becomes a "culturally familiar" game for these two players, and the past history of plays, by other members of the society, should inform play in this game. An equilibrium is then a cultural norm, an institution, in this society, and the game is typically played according to this norm. Any Nash equilibrium of the underlying game that can emerge as a norm in some society is sustainable. From this perspective, Myerson reasons, the two pure-strategy equilibria in the Battle-of-Sexes game are sustainable while the mixed equilibrium is not.

In his search for a formal definition of a sustainable equilibrium—that uses as its only data the given game and not some extended model—Myerson considers, and then dismisses, on axiomatic grounds, existing refinements that yield the same prediction in the Battle-of-Sexes game as his heuristic argument does: for example, persistent equilibria fail invariance; and evolutionary stability fails existence. Myerson concludes his paper with a conjecture that the index of an equilibrium is a determinant of its sustainability.[1]

Date: February 2020.

[1]Interestingly, Myerson speculates that one could perhaps develop a theory of index of equilibria based on fixed-point theory, seemingly unaware, as Hofbauer [5] observes, of an extant theory in the literature, e.g. Ritzberger [13].

Hofbauer [5] distills the ideas in Myerson's paper to provide a definition of sustainable equilibria. Hofbauer posits that a minimum requirement of sustainability should be that if an equilibrium of a game is sustainable, it should remain sustainable in the game obtained by restricting players' strategies to the set of best replies to the equilibrium.[2] If one also accepts that an equilibrium that is unique is sustainable, then one is lead to the following definition. Say that a game-equilibrium pair is *equivalent* to another such pair if the restrictions of the two games to the set of best replies to their respective equilibria are the same game (modulo a relabelling of the players and their strategies) and the two equilibria coincide (under the same identification). An equilibrium of a game is *sustainable* if it has an equivalent pair where the equilibrium is unique. Call an equilibrium *regular* if locally the equilibrium is a differentiable function of the game. Following Myerson, Hofbauer conjectured that a regular equilibrium is sustainable iff its index is +1. von Stengel and von Schemde [15] proved this conjecture for bimatrix games. In this paper, we show that the conjecture holds for all N-person games.[3]

What are the key properties of the index of equilibria that drive this equivalence? To answer this question, let us see a sketch of our proof. In one direction, suppose that an equilibrium σ of a game G is sustainable, and that (G, σ) is equivalent to a pair $(\bar{G}, \bar{\sigma})$ where $\bar{\sigma}$ is the unique equilibrium of $\bar{G}$. Let G^* be the game obtained from G by deleting strategies that are inferior replies to σ. It follows from a property of index that the index of σ in G can be computed as the index of σ in G^*. The game G^* is also the game obtained from $\bar{G}$ by deleting inferior replies there. Therefore, the index of σ in G^* can also be computed as the index of $\bar{\sigma}$ in $\bar{G}$. As $\bar{\sigma}$ is the unique equilibrium of $\bar{G}$, its index is +1, which then gives us the result.

Going the other way, if we have a +1 equilibrium σ of a game G, then the sum of the indices of the other equilibria is zero. We can now take a map whose fixed points are the Nash equilibria of G and alter it outside a neighbourhood of σ so that the new map has no fixed points other than σ.[4] By a careful addition of strategies and specification of payoffs for these strategies, we obtain a game $\bar{G}$ where any equilibrium must translate into a fixed point of the modified map of G, making σ the unique equilibrium in $\bar{G}$.

[2]For generic games, this is equivalent to restricting players' strategies to the support of the equilibrium—see Section 4 for a discussion of this point.

[3]von Stengel and von Schemde, as well as our construction, are a little stronger than the conjecture: any isolated +1-index equilibrium can be made unique in an equivalent game without deleting strategies nor relabelling players or strategies.

[4]The possibility of such a construction follows from a deep result in algebraic topology, the Hopf Extension Theorem.

Equilibria with index $+1$ are also distinguished from their counterparts with index -1 in terms of their dynamic stability. It is well-known, both in general equilibrium and in game theory, that equilibria with index -1 are asymptotically unstable under any reasonable learning or adjustment process—cf. McLennan [9].[5] Even computational dynamics like those generated by homotopy algorithms (Lemke-Howson, linear-tracing procedure, etc.) converge to a $+1$ index equilibrium (Herings-Peeters [4]). While these results might be seen as eliminating -1 equilibria, they do not conclusively come down in favor of $+1$ equilibria. The main reason is that it is still an open question as to whether every regular game has an equilibrium (necessarily of index $+1$) that is stable with respect to some natural dynamical system—cf. Hofbauer [5] for more about this conjecture.

A word about our methodology is in order. In a bimatrix game, a player's payoff function is linear in his opponent's strategy. von Schemde and von Stengel [15] were able to exploit this feature in their proof and use tools from the theory of polytopes (see also von Schemde [16]). In the general case, their technique is inapplicable. What we do, instead, is start with a construction involving a fixed-point map and then convert it into a game-theoretically meaningful one. In this respect, our approach is similar in spirit to, but different in details from, that in Govindan and Wilson [3].

The rest of the paper is organized as follows. Section 2 sets up the problem and states our theorem. It also gives an informal summary of the theory of index of equilibria. Section 3 is devoted to proving the theorem. Section 4 provides a discussion of the role of regularity both in the definition and in the result. Finally, in the Appendix we review a construct from the theory of triangulations that we need in the proof.

2. Definitions and Statement of The Theorem

A finite game in normal form is a triple $(\mathcal{N}, (S_n)_{n\in\mathcal{N}}, G)$ where: $\mathcal{N} = \{ 1, \ldots, N \}$ is the set of players, with $N \geqslant 2$; for each $n \in \mathcal{N}$, S_n is a finite set of pure strategies; and, letting $S \equiv \prod_{n\in\mathcal{N}} S_n$ be the set of pure strategy profiles, $G : S \to \mathbb{R}^{\mathcal{N}}$ is the payoff function. By a slight abuse of notation, we will refer to a game by its payoff function G.

Given a game G, for each n, let Σ_n be the set of n's mixed strategies and let $\Sigma \equiv \prod_{n\in\mathcal{N}} \Sigma_n$. Also, for each n, $S_{-n} \equiv \prod_{m\neq n} S_n$, and $\Sigma_{-n} \equiv \prod_{m\neq n} \Sigma_n$. The payoff function G extends to Σ in the usual way and we will denote this extension by G as well.

Define an equivalence relation on game-equilibria pairs as follows. For $i = 1, 2$, let $((\mathcal{N}, (S^i_n)_{n\in\mathcal{N}}, G^i), \sigma^i)$ be a game-equilibrium pair, i.e., σ^i is an equilibrium of G^i. Say that $(G^1, \sigma^1) \sim (G^2, \sigma^2)$ if, up to a relabelling of strategies, the restriction of G^1 to the set of best

[5] In fact, this observation leads McLennan to articulate his *index* $+1$ *principle*, which selects $+1$ equilibria.

replies to σ^1 is the same game as the restriction of G^2 to the set of best replies to σ^2, and the equilibria coincide under this identification. It is easily checked that $\sim$ is an equivalence relation.

Definition 2.1. An equilibrium σ of a game G is *sustainable* if $(G, \sigma) \sim (\bar{G}, \bar{\sigma})$ for a game $\bar{G}$ where $\bar{\sigma}$ is the unique equilibrium.

Sustainability is a property of equivalence classes and we could say that the canonical representation of a sustainable equilibrium is the game-equilibrium pair where there are no inferior replies to the equilibrium.

2.1. **Index and Degree of Equilibria.** Both the index and the degree of equilibria are measures of the robustness of equilibria to perturbations. They differ in the space of perturbations they consider, but ultimately agree with one another.[6] We start with the degree of equilibria. For simplicity we give a definition of degree only for regular equilibria. This approach allows to bypass the use of algebraic topology, but more importantly it is germane to our problem, as we are concerned only with regular equilibria in this paper. For the general definition, see for e.g., Govindan and Wilson [3].

Fix both the player set $\mathcal{N}$ and the strategy space S. The space of games with strategy space S is then the Euclidean space $\Gamma \equiv \mathbb{R}^{N|S|}$ of all payoff functions G. Let $\mathcal{E}$ be the graph of the Nash equilibrium correspondence over Γ, i.e., $\mathcal{E} = \{ (G, \sigma) \in \Gamma \times \Sigma \mid \sigma \text{ is a Nash equilibrium of } G \}$. Let $proj : \mathcal{E} \to \Gamma$ be the natural projection: $proj(G, \sigma) = G$. By the Kohlberg-Mertens Structure Theorem [6], there exists a homeomorphism $h : \mathcal{E} \to \Gamma$ such that h^{-1} is differentiable almost everywhere. Say that an equilibrium σ of a game G is *regular* if $proj \circ h^{-1}$ is differentiable and has a nonsingular Jacobian at $h(G, \sigma)$; and say that a game is *regular* if each equilibrium σ of G is regular. If an equilibrium σ is regular, then it is a *quasi-strict equilibrium*—that is, all unused strategies are inferior replies—and locally, the equilibrium is a smooth, even analytic, function of the game; moreover, it is also a regular equilibrium in the space of games obtained by deleting the unused strategies or, indeed, by adding strategies that are inferior replies to the equilibrium. The set of games that are not regular is a closed subset of lower dimension—actually codimension one—in Γ and thus generic games are regular.

If an equilibrium σ of a game G is regular, then we can assign a *degree* to it that is either $+1$ or -1 depending on whether the Jacobian of $proj \circ h^{-1}$ at $h(G, \sigma)$ has a positive or a negative determinant. An inspection of the formula for the Jacobian shows that the degree of

[6] The reason we are defining and reviewing both concepts, when apparently one would do, is that they are both useful in the exposition.

a regular equilibrium σ is the same as its degree computed in the space of games obtained by deleting the strategies that are inferior replies to σ. Therefore, if σ is a regular equilibrium of G and if $(G, \sigma) \sim (\bar{G}, \bar{\sigma})$ then $\bar{\sigma}$ is a regular equilibrium of $\bar{G}$ and it has the same degree as σ, making degree an invariant for an equivalence class.

As the Kohlberg-Mertens homeomorphism extends to the one-point compactifications of $\mathcal{E}$ and Γ, and $proj \circ h^{-1}$ is homotopic to the identity on this extension, the sum of the degrees of equilibria of a regular game is $+1$.[7]

In fixed-point theory, the index of fixed points contains information about their robustness when the map is perturbed. (See McLennan [8] for an account of index theory written primarily for economists.) Since Nash equilibria are obtainable as fixed points, index theory applies directly to them. For simplicity, suppose $f : U \to \Sigma$ is a differentiable map defined on a neighborhood U of Σ in $\mathbb{R}^{N|S|}$ and such that the fixed points of f are the Nash equilibria of a game G. Let d be the displacement of f, i.e., $d(\sigma) = \sigma - f(\sigma)$. Then the Nash equilibria of G are the zeros of d. Suppose now that the Jacobian of d at a Nash equilibrium σ of G is nonsingular. Then we can define the *index* of σ under f as ± 1 depending on whether the determinant of the Jacobian is positive or negative.

One potential problem with the definition of index is the dependence of the computation on the function f, as intuitively we would think of the index as depending only on the game G. But, under some regularity assumptions of f, we can show that the index is independent of f. Specifically, consider the class of continuous maps $F : \Gamma \times \Sigma \to \Sigma$ with the property that the fixed points of the restriction of F to $\{ G \} \times \Sigma$ are the Nash equilibria of G. Then the index of equilibria is independent of the particular map in this class that is used to compute it (Demichelis and Germano [2]); furthermore, the index of equilibria equals their degree (cf. Govindan and Wilson [3]).[8] Thus, for a regular equilibrium, we can talk unambiguously of its index and use the term degree interchangeably with it.

2.2. **Statement of The Theorem.** The following theorem settles the Myerson-Hofbauer conjecture in the affirmative.

Theorem 2.2. *A regular equilibrium is sustainable iff its index is* $+1$.

[7] If G is nongeneric, we can define the degree of a component of equilbria as the sum of the degrees of equilibria in a neighborhood of the component for a regular game that is in a neighborhood of G; this computation is independent of the neighborhoods chosen, as long as they are sufficiently small. The sum of the degrees of the components of equilibria of a game is $+1$.

[8] As with degree, we can define the index of a component of equilibria (first defined by Ritzberger [13]), and the index and degree of components agree as well.

As the sum of the indices of the equilibria of a regular game is $+1$, there are an odd number of equilibria—in particular at least one—with index $+1$. Thus, we have the following corollary.

Corollary 2.3. *Every regular game has at least one sustainable equilibrium.*

3. Proof of Theorem 2.2

We will present the proof in a sequence of steps, each of which will be carried out in a separate subsection.

3.1. **The Index of a Sustainable Equilibrium.** We begin with a proof of the necessity of the condition. Let σ^* be a regular equilibrium of a game G that is sustainable. Let $(G, \sigma^*) \sim (\bar{G}, \bar{\sigma})$, where $\bar{\sigma}$ is the unique equilibrium of $\bar{G}$. As we saw in the previous section, the index is constant on an equivalence class. As $\bar{\sigma}$ is the unique equilibrium of $\bar{G}$, its index is $+1$, and the result follows.

3.2. **Preliminaries.** The rest of the section is devoted to proving the sufficiency of the condition. In this subsection, we introduce some key ideas that we exploit in the proof.

First, we gather a list of notational conventions to be used. Throughout Section 3 (but not in the Appendix) we use the ℓ_∞-norm on Euclidean spaces. For any subset A of a topological space X, we let $\partial_X A$ be its topological boundary and $\mathrm{int}_X(A)$ its interior. If C is a convex set in a Euclidean space, then ∂C and $\mathrm{int}(C)$ refer to the boundary and the interior of C in the affine space generated by C. Given a payoff function G, and a vector $g \in \prod_{n \in \mathcal{N}} \mathbb{R}^{S_n}$, let $G \oplus g$ be the game where the payoff to player n from a profile $s \in S$ is $G_n(s) + g_{n,s_n}$.

For a game G, recall that Nash [11] obtains its equilibria as fixed points of a map on the strategy space. This function, which we denote by f, is defined as follows. For each n, s_n and σ, let $\phi_{n,s_n}(\sigma) \equiv \max\{\,0, G_n(s_n, \sigma_{-n}) - G_n(\sigma)\,\}$ and $\phi_n \equiv \prod_{s_n \in S_n} \phi_{n,s_n}$; then

$$f_{n,s_n}(\sigma) \equiv \frac{\sigma_{n,s_n} + \phi_{n,s_n}(\sigma)}{1 + \sum_{t_n \in S_n} \phi_{n,t_n}(\sigma)}.$$

If $f_n(\sigma) \neq 0$ for some n and σ, then letting

$$r_n(\sigma) = \left(\sum_{t_n \in S_n} \phi_{n,t_n}(\sigma) \right)^{-1} \phi_n(\sigma)$$

and

$$\lambda_n(\sigma) = \frac{1}{1 + \sum_{t_n \in S_n} \phi_{n,t_n}(\sigma)},$$

we have

$$f_n(\sigma) = \lambda_n(\sigma)\sigma_n + (1 - \lambda_n(\sigma))r_n(\sigma).$$

Thus $f_n(\sigma)$ is an average of σ and a mixed strategy $r_n(\sigma)$ that has the following properties: (1) it assigns a positive probability to a pure strategy iff it does strictly better than σ—in particular, it assigns zero probability to some strategy in the support of σ_n, as $f_n(\sigma) \neq \sigma_n$; (2) it assigns the highest probabilities to the best replies to σ.

A game $(\mathcal{N}, \bar{S}, \bar{G})$ *embeds* $(\mathcal{N}, S, G)$ if: (1) for each n: $S_n \subseteq \bar{S}_n$; and (2) the restriction of $\bar{G}$ to S equals G. Again for notational convenience, we will talk of a game $\bar{G}$ embedding G. When $\bar{G}$ embeds G, we view the set Σ of mixed strategies of G as a subset of the set $\bar{\Sigma}$ of mixed strategies in $\bar{G}$. Obviously, if $\bar{G}$ embeds G and σ is an equilibrium of $\bar{G}$ where for each n, the strategies that are not in S_n are inferior replies, then $(G, \sigma) \sim (\bar{G}, \sigma)$. Our proof technique is to show that for each regular +1 equilibrium σ^* of G, we can embed G in a game $\bar{G}$ where σ^* is the unique equilibrium and the newly added strategies are inferior replies to σ^*.

A game in *strategic form* is a triple $(\mathcal{N}, (P_n)_{n\in\mathcal{N}}, V)$ where: $\mathcal{N}$ is the player set; for each n, P_n is a polytope of strategies; letting $P \equiv \prod_n P_n$, $V : P \to \mathbb{R}^{\mathcal{N}}$ is a multilinear payoff function. Given such a game, we can define a normal-form game $(\mathcal{N}, (S_n)_{n\in\mathcal{N}}, G)$ where for each n, S_n is the set of vertices of P_n and for each s, $G(s) = V(s)$. We say that a strategic form game $\bar{V}$ embeds G if the associated normal form game $\bar{G}$ embeds G, or equivalently for each n, the strategies in S_n are vertices of $\bar{P}_n$. In our proof we construct embeddings of G in strategic form games $\bar{V}$ that have a simple structure: for each n, the strategies in S_n span a face of $\bar{P}_n$.

3.3. **A Simple Consequence of Regularity.** From now on fix a game G and let σ^* be a regular equilibrium with index +1. For each n, let S^*_n be the support of σ^*_n. Our objective in this subsection is to record the following simple, and yet consequential, property of σ^*. There exists $\bar{\varepsilon} > 0$ such that: if $\sigma \neq \sigma^*$ is an equilibrium of G, then there exist two different players n_1, n_2, such that for $i = 1, 2$, there exists $s_{n_i} \in S^*_{n_i}$ with $\sigma_{n_i,s_{n_i}} < \sigma^*_{n_i,s_{n_i}} - \bar{\varepsilon}$. Indeed if this property is not true, there exist a sequence $k \to \infty$, a corresponding sequence σ^k of equilibria converging to some σ, and a player n such that: (1) $\sigma^k \neq \sigma^*$ for all k; and (2) for all $m \neq n$ and $s_m \in S^*_m$, $\sigma^k_{m,s_m} \geqslant \sigma^*_{m,s_m} - k^{-1}$. Therefore, $\sigma_m = \sigma^*_m$ for all $m \neq n$. But $\sigma_n \neq \sigma^*_n$ as σ^* is regular and, hence, isolated. This implies that $\lambda\sigma + (1 - \lambda)\sigma^*$ is an equilibrium for all $\lambda \in [0, 1]$, again contradicting the fact that σ^* is isolated. Thus, there exists $\bar{\varepsilon}$ with the stated property.[9]

[9]Note that this proof only uses the fact that σ^* is isolated.

For $0 < \varepsilon \leqslant \bar{\varepsilon}$, and each n, let B_n^ε be the set of $\sigma_n \in \Sigma_n$ such that $\sigma_{n,s_n} \geqslant \sigma^*_{n,s_n} - \varepsilon$ for all $s_n \in S_n^*$; and let B^ε be the set of σ such that σ_n is not in B_n^ε for at most one n. (N.B. $\prod_{n\in\mathcal{N}} B_n^\varepsilon \subsetneq B^\varepsilon$). Since $B^\varepsilon = \bigcup_n \Sigma_n \times \prod_{m\neq n} B_m^\varepsilon$ is the union of finitely many closed sets, it is closed.

3.4. **Killing All Fixed Points of f Other Than σ^*.** From the viewpoint of fixed-point theory, our problem amounts to embedding Σ as a proper face of a polytope $\bar{\Sigma}$, extending f to a function $\bar{f}$ on it, and then modifying $\bar{f}$ such that its only fixed point is σ^*. From a game-theoretic viewpoint, there is an additional problem introduced by the caveat that $\bar{f}$ should, in a sense, be realizable as a fixed-point map of a game $\bar{G}$ that embeds G—i.e., a map whose fixed points are the equilibria of game $\bar{G}$. In this subsection, we solve the first problem partially, by constructing a map f^0 that coincides with f on some B^ε and that has no fixed points outside it. We will later use this map to construct the embedding $\bar{G}$.

If necessary by adding a strictly dominated strategy for each player, we can assume that σ_n^* belongs to $\partial\Sigma_n$ for each n. (Recall that sustainability and index are properties of equivalence classes of regular equilibria, so that the addition of such strategies is harmless.) Let $V \equiv B^{\bar{\varepsilon}}$; $X \equiv \Sigma \setminus \mathrm{int}_\Sigma(V)$. The boundary of X is relative to the affine space generated by Σ, i.e., $\partial X \equiv (\partial\Sigma \setminus V) \cup \partial_\Sigma V$. We claim now that $(X, \partial X)$ is homeomorphic to a ball with boundary. Indeed, the desired homeomorphism can be constructed as follows. Pick a completely-mixed strategy-profile σ^0 such that $\sigma^0_{n,s_n} < \sigma^*_{n,s_n} - \bar{\varepsilon}$ for all n and $s_n \in S_n^*$. (Such a choice is possible since σ^* belongs to the boundary of Σ.) The set X is star-convex at σ^0: for each $\sigma \in X$, $\lambda\sigma + (1-\lambda)\sigma^0 \in X$ for all $\lambda \in [0,1]$. Therefore, there is now a ball B around σ^0 in $\Sigma \setminus \partial\Sigma$ that is contained in $X \setminus \partial X$ that is homeomorphic to X using radial projections from σ^0.

Define $\tilde{f} : \Sigma \to \Sigma$ as follows. First let $\alpha : \Sigma \to [0,1]$ be a continuous function that is zero on V and positive everywhere else. Then, letting τ_n^0 be the barycenter of Σ_n for each n, define $\tilde{f}(\sigma) = (1-\alpha(\sigma))f(\sigma) + \alpha(\sigma)\tau^0$. The function $\tilde{f}$ equals f on V and therefore σ^* is the unique fixed point of $\tilde{f}$ in V. The set $\Sigma \setminus V$ is mapped by $\tilde{f}$ to $\Sigma \setminus \partial\Sigma$. Hence all the other fixed points of $\tilde{f}$ belong to $X \setminus \partial X$.

Let $\tilde{d}$ be the displacement of $\tilde{f}$: $\tilde{d}(\sigma) \equiv \sigma - \tilde{f}(\sigma)$. Letting A be the hyperplane in $\mathbb{R}^{N|S|}$ through the origin and with normal $(1,\ldots,1)$, $\tilde{d}$ maps Σ into A. As the index of σ^* is $+1$, the sum of the indices of the other components of fixed points of $\tilde{f}$, which are contained in $X \setminus \partial X$, is zero. Therefore, $\tilde{d} : (X, \partial X) \to (A, A - 0)$ has degree zero. By the Hopf Extension Theorem (cf. Corollary 8.1.18, [14]) there exists a map $\hat{d}^0$ from X to $A - 0$ such that its restriction to ∂X coincides with $\tilde{d}$. Extend $\hat{d}^0$ to the whole of Σ by letting it be $\tilde{d}$ outside X, i.e., on $V \setminus \partial_\Sigma V$.

For each $\sigma \in \Sigma$, there exists $\lambda \in (0,1]$ such that $\sigma - \lambda\hat{d}^0(\sigma) \in \Sigma$: indeed, this is obvious for $\sigma \in X \setminus \partial X$, since σ belongs to the interior of Σ; for $\sigma \in \partial X \cup V$, we can take $\lambda = 1$ as $\hat{d}^0 = \tilde{d}$. Now for each σ, let $\lambda(\sigma)$ be the largest $\lambda \in [0,1]$ such that $\sigma - \lambda(\sigma)\hat{d}^0(\sigma) \in \Sigma$. Note that the map $\sigma \mapsto \lambda(\sigma)$ is continuous. Define $f^0 : \Sigma \to \Sigma$ by: $f^0(\sigma) = \sigma - \lambda(\sigma)\hat{d}^0(\sigma)$. The map f^0 is continuous, coincides with f on V, and has σ^* as its unique fixed point.

3.5. **A Parametrized Family of Perturbed Games.** Ideally, we would like a game G^0 such that f^0 is the Nash map of G^0. This seems to be too strong a property to hold. However, f^0 does contain enough information for us to construct a function $g : \Sigma \to \prod_{n\in\mathcal{N}} \mathbb{R}^{S_n}$ such that: (1) $g(\cdot)$ is zero on B^ε for some sufficiently small ε; (2) σ is an equilibrium of $G \oplus g(\sigma)$ iff $\sigma = \sigma^*$.

Choose $0 < \varepsilon < \bar{\varepsilon}$ and let $U \equiv B^\varepsilon$; note that $U \subsetneq V$. For each n, let Z_n be $(d_n^0)^{-1}(0) \cap (\Sigma \setminus \mathrm{int}_\Sigma(U))$, where d^0 is the displacement of f^0. Let Z_n^1 be the complement of Z_n in $\Sigma \setminus \mathrm{int}_\Sigma(U)$. Define $r_n^0 : Z_n^1 \to \partial\Sigma_n$ as follows. For each $\sigma \in Z_n^1$, let $r_n^0(\sigma)$ be the unique point in $\partial\Sigma_n$ on the ray from σ_n through $f_n^0(\sigma)$, i.e., it is the unique point of the form $(1-\alpha)\sigma_n + \alpha f_n^0(\sigma)$ for $\alpha \geqslant 1$ that belongs to $\partial\Sigma_n$. If $\sigma \in Z_n^1$, then there exists some t_n that is in the support of σ_n but not $r_n^0(\sigma)$. For each n, s_n, let Z_{n,s_n}^+ be the closure of the set of $\sigma \in Z_n^1$ for which $r_{n,s_n}^0(\sigma) \geqslant r_{n,t_n}^0(\sigma)$ for all $t_n \in S_n$. If $f^0(\sigma) = f(\sigma)$ and $\sigma \in Z_n^1$, then $r_n^0(\sigma)$ equals $r_n(\sigma)$ as defined in subsection 3.2; therefore, $\sigma \in Z_{n,s_n}^+$ iff s_n is a best reply to σ in G.

We are now ready to define the function $g(\sigma)$. First, let $v_n(\sigma) = \max_{s_n} G_n(s_n, \sigma_{-n})$. Second, let $\beta_n^1 : \Sigma \to [0,1]$ be a continuous function that is zero on Z_n and positive everywhere else. Third, for each n, s_n, let $\beta_{n,s_n}^2 : \Sigma \to [0,1]$ be a continuous function that is one on Z_{n,s_n}^+ and strictly smaller than one elsewhere. Finally, let $\beta^3 : \Sigma \to [0,1]$ be a continuous function that is one on $\Sigma \setminus \mathrm{int}_\Sigma(V)$, zero on U and strictly positive everywhere else. For each n, s_n and σ, define:

$$g_{n,s_n}(\sigma) = \beta^3(\sigma)\beta_{n,s_n}^2(\sigma)[v_n(\sigma) - G_n(s_n, \sigma_{-n}) + \beta_n^1(\sigma)].$$

If $\sigma \in U$, then $g(\sigma) = 0$ as $\beta^3(\sigma) = 0$; and σ is an equilibrium of $G \oplus g(\sigma)$ iff $\sigma = \sigma^*$. Suppose $\sigma \notin U$. Since σ^* is the only fixed point of f^0, there exists some n such that $f_n^0(\sigma) \neq \sigma_n$. For this n, there exist s_n such that: $\sigma \in Z_{n,s_n}^+$ (take s_n s.t. $r_{n,s_n}^0(\sigma) \geqslant r_{n,t_n}^0(\sigma)$ for all $t_n \in S_n$); and there is t_n in the support of σ_n but not in the support of $r_n^0(\sigma)$. This implies $\beta_{n,s_n}^2(\sigma) = 1$, while $\beta_{n,t_n}^2(\sigma) < 1$.

If $\sigma \notin V$, then $\beta^3(\sigma) = 1$ and so,

$$G_n(s_n, \sigma_{-n}) + g_{n,s_n}(\sigma) = v_n(\sigma) + \beta_n^1(\sigma) > v_n(\sigma) + \beta_{n,t_n}^2\beta_n^1(\sigma) \geqslant G_n(t_n, \sigma_{-n}) + g_{n,t_n}(\sigma),$$

showing that σ is not an equilibrium of $G \oplus g(\sigma)$.

If $\sigma \in V \setminus U$, then as f^0 coincides with f, s_n is a best reply against σ while t_n is not. Thus $G_n(s_n, \sigma_{-n}) = v_n(\sigma)$ and $G_n(t_n, \sigma_{-n}) < v_n(\sigma)$. Since $\beta^3(\sigma) > 0$, we obtain that

$$G_n(s_n, \sigma_{-n})+g_{n,s_n}(\sigma) = v_n(\sigma)+\beta^3(\sigma)\beta^1_n(\sigma) > v_n(\sigma)+\beta^3(\sigma)\beta^2_{n,t_n}(\sigma)\beta^1_n(\sigma) > G(t_n, \sigma_{-n})+g_{n,t_n}(\sigma),$$

and again σ is not an equilibrium of $G \oplus g(\sigma)$. Thus the function g has the desired properties.

3.6. **Ridding the Bonus $g_n(\sigma)$ of its Dependence on σ_n.** An important feature of a best-reply correspondence is the independence of a player's coordinate correspondence from his own strategy. In our context, this would translate to having g_n be independent of σ_n, since as we saw in the previous subsection, we used the information about the function r^0_n—which, at least partially, encodes the preference orderings over strategies—to define g. Given that f^0, from which r was derived, was constructed by an appeal to a theorem from algebraic topology, there is no guarantee that the function g has these features. In this subsection, we show that by an embedding of G in a game $\hat{G}$, we can "lift" g to a higher-dimensional space where we get the required independence.

For simplicity, the game $\hat{G}$ will be represented in strategic form as follows. The strategy set $\hat{\Sigma}_n$ of player n is a product $\Sigma_n \times \Sigma_{n,n+1}$, where $\Sigma_{n,n+1} \equiv \Sigma_{n+1}$, with the convention that $n+1 = 1$ if $n = N$. A strategy $\hat{\sigma}_n$ is typically denoted as a pair $(\sigma_n, \sigma_{n,n+1})$. Let $\hat{\Sigma} = \prod_n \hat{\Sigma}_n$. For a strategy profile $\hat{\sigma}$, we write σ for the profile of coordinates of the first factor. The coordinates in $\Sigma_{n,n+1}$ are payoff irrelevant, i.e., for each $\hat{\sigma} \in \hat{\Sigma}$, $\hat{G}(\hat{\sigma}) \equiv G(\sigma)$. For each n, if we let s^0_{n+1} be a fixed pure strategy in S_{n+1}, then $\hat{G}$ embeds G as Σ_n can be identified with the face $\Sigma_n \times \{\, s^0_{n+1} \,\}$. For each n, s_n, we define $\hat{g}_{n,s_n} : \hat{\Sigma}_{-n} \to \mathbb{R}$ by:

$$\hat{g}_{n,s_n}(\hat{\sigma}_{-n}) = g_{n,s_n}(\sigma_{n-1,n}, \sigma_{-n}).$$

Obviously $\hat{g}_n$ depends only on $\hat{\sigma}_{-n}$.

3.7. **Isolating σ^*.** Before we can use the perturbation $\hat{g}_n$, we need to first embed G in a game $\tilde{G}$ where σ^* is the only equilibrium in the face Σ of $\tilde{G}$ and in fact the only equilibrium in which the strategy of even one of the players is in Σ_n. (The perturbation g is then used on the face opposite to Σ.) Since we will ultimately have to add the payoff-irrelevant factor Σ_{n+1} for each n, we will actually embed the game $\hat{G}$ defined in subsection 3.6 in a game $\tilde{G}$. As previously, the embedding is in a strategic form game.

Choose $0 < \varepsilon^* < \varepsilon$ such that $\sigma^*_{n,s_n} > \varepsilon^*$ for each n and $s_n \in S^*_n$. Let $U^* \equiv B^{\varepsilon^*}$ and $U^*_n \equiv B^{\varepsilon^*}_n$, for each $n \in \mathcal{N}$; the set U^* is a strict subset of U. For each n, choose an arbitrary object 0^*_n (not in S^*_n). Let Θ_n be the set of distributions over $\prod_{m \neq n}(S^*_m \cup \{\, 0^*_m \,\})$. As with Σ in previous section, it is convenient to equip n with a copy of Θ_{n+1}. Thus, let $\Theta_{n,n+1} \equiv \Theta_{n+1}$.

For each player n, his strategy set $\tilde{\Sigma}_n$ is a product $\Sigma_n \times \Theta_n \times \Sigma_{n,n+1} \times \Theta_{n,n+1}$. A typical element $\tilde{\sigma}_n \in \tilde{\Sigma}_n$ has coordinates $(\sigma_n, \theta_n, \sigma_{n,n+1}, \theta_{n,n+1})$.

We will now describe the payoff functions. For each $\theta_n \in \Theta_n$ and $m \neq n$, we let $\theta_{n,m}$ be the marginal distribution of θ_n over $S^*_m \cup \{ 0^*_m \}$; let $\Theta_{n,m}$ be the set of all marginal distributions over $S^*_m \cup \{ 0^*_m \}$. For each $m \neq n$, let $\gamma_{n,m} : \Theta_{n,m} \times \Sigma_m \to \mathbb{R}$ be a bilinear function defined as follows. For all σ_m: $\gamma_{n,m}(0^*_m, \sigma_m) = 0$, while for $s_m \in S^*_m$, $\gamma_{n,m}(s_m, \sigma_m) = 1 - (\sigma^*_{m,s_m} - \varepsilon^*)^{-1}\sigma_{m,s_m}$. For each $\tilde{\sigma}$, player n's payoff in $\tilde{G}$ is:

$$\tilde{G}_n(\tilde{\sigma}) = G_n(\sigma) + \gamma_n(\theta_n, \sigma_{-n}),$$

where

$$\gamma_n(\theta_n, \sigma_{-n}) \equiv \sum_{m \neq n} \gamma_{n,m}(\theta_{n,m}, \sigma_m).$$

Notice that the payoff function of each player n is affine over each strategy set $\tilde{\Sigma}_m$, $m = 1, ..., N$, so $\tilde{G}$ is indeed a well-defined game in strategic form. For each n, let $\theta^0_n = (0^*_m)_{m \neq n}$. Then G is embeddable in $\tilde{G}$ as the face $\Sigma_n \times \{ \theta^0_n \} \times \{ (s^0_{n+1}, \theta^0_{n+1}) \}$ (recall that s^0_{n+1} was a fixed strategy used in defining $\hat{G}$ in subsection 3.6).

Suppose $\tilde{\sigma}$ is an equilibrium of $\tilde{G}$, then σ is an equilibrium of G, as the functions γ_n of each player n do not depend on σ_n. If $\sigma = \sigma^*$, then the unique $\theta_{n,m}$ that is optimal for each $n \neq m$ is 0^*_m and thus the equilibrium uses θ^0_n for each n. On the other hand if $\sigma \neq \sigma^*$, then there are at least two players m for whom $\sigma_m \notin U^*_m$. Therefore, for each n, there is at least one $m \neq n$ such that $0^*_{n,m}$ is not optimal. Thus for each n, the support of θ_n does not include θ^0_n.

To conclude, if $\tilde{\sigma} = (\sigma_n, \theta_n, \sigma_{n,n+1}, \theta_{n,n+1})_{n \in \mathcal{N}}$ is an equilibrium of $\tilde{G}$, then σ is an equilibrium of G and either: (1) $\sigma = \sigma^*$ and $\theta_n = \theta^0_n$, for each $n \in \mathcal{N}$; or (2) $\sigma \neq \sigma^*$ and the support θ_n does not contain θ^0_n for any $n \in \mathcal{N}$.

3.8. **The Embedding $\bar{G}^\delta$.** The embedding that allows us to obtain σ^* as the unique equilibrium (and regular as well) will be built from $\tilde{G}$ by adding a finite number of mixed strategies as pure strategies and by modifying the payoffs to eliminate all other equilibria.[10]

The set $\Sigma \backslash \text{int}_\Sigma(U^*)$ is compact, $g(\cdot)$ is continuous, and no $\sigma \in \Sigma \backslash \text{int}_\Sigma(U^*)$ is an equilibrium of $G \oplus g(\sigma)$, as shown in subsection 3.5. Hence, there exists $\eta > 0$ such that no $\sigma \in \Sigma \backslash \text{int}_\Sigma(U^*)$ is an equilibrium of $G \oplus g$ for any g with $\|g - g(\sigma)\| \leqslant \eta$. Also, since g is uniformly continuous,

[10]von Schemde and von Stengel have an explicit bound on the number of strategies they add, which is three times the number of pure strategies in G. In our construction, we have no way of obtaining such a bound: the construction depends on the fixed-point map f^0 obtained in subsection 3.4, which in turn relies on an existence result, the Hopf Extension Theorem.

there exists $0 < \zeta < 1/3$ such that $\|g(\sigma) - g(\sigma')\| \leqslant \eta$, if $\|\sigma - \sigma'\| \leqslant \zeta$. Reduce ζ to ensure that it is also smaller than the distance between U_n^* and $\partial_{\Sigma_n} U_n$ for each n.

For each n, take a triangulation $\mathcal{T}_n$ of $\Sigma_n \times \Theta_n$ with the following properties. (See the Appendix for the details.) (1) The only vertices in $\Sigma_n \times \{\theta_n^0\}$ of $\mathcal{T}_n$ are pure strategies (s_n, θ_n^0), $s_n \in S_n$; (2) letting Θ_n^1 be the face of Θ_n where θ_n^0 has zero probability, if $T_n \in \mathcal{T}_n$ is a simplex either with a face in $\Sigma_n \times \Theta_n^1$, or shares a face with such a simplex, then the diameter of T_n is less than ζ; (3) there exists a convex function $\varrho_n : \Sigma_n \times \Theta_n \to \mathbb{R}_+$ such that: (a) $\varrho_n(\lambda x + (1-\lambda)y) = \lambda\varrho(x) + (1-\lambda)\varrho(y)$ iff x and y belong to a simplex T_n of $\mathcal{T}_n$; (b) ϱ_n is zero on $\Sigma_n \times \{\theta_n^0\}$.

Let $\bar{S}_n^0$ be the set of vertices of the triangulation $\mathcal{T}_n$. Let $\bar{S}_n^1 \equiv \bar{S}_{n+1}^0$. For $i = 0, 1$, let $\bar{\Sigma}_n^i$ be the set of mixtures over $\bar{S}_n^i$. The pure strategy set of player n in the game $\bar{G}^\delta$ in normal form is $\bar{S}_n \equiv \bar{S}_n^0 \times \bar{S}_n^1$. The set of mixed strategies is denoted $\bar{\Sigma}_n$. For each mixed strategy $\bar{\sigma}_n$, and $i = 0, 1$, we let $\bar{\sigma}_n^i$ be the marginals over $\bar{S}_n^i$. Define $\bar{S} \equiv \prod_n \bar{S}_n$ and $\bar{\Sigma} \equiv \prod_n \bar{\Sigma}_n$. Also, let $\bar{S}^i \equiv \prod_n \bar{S}_n^i$ and $\bar{\Sigma}^i \equiv \prod_n \bar{\Sigma}_n^i$ for $i = 0, 1$.

Fix $\delta > 0$. We will now define the payoff function $\bar{G}^\delta$. For each n, let $\mathcal{T}_n^1$ be the collection of simplices of $\mathcal{T}_n$ that have nonempty intersection with $\Sigma_n \times \Theta_n^1$. Given a pure strategy profile $\bar{s} \in \bar{S}$ with $\bar{s}_n = (\sigma_n, \theta_n, \sigma_{n,n+1}, \theta_{n,n+1})$ for each n, the payoff $\bar{G}_n^\delta(\bar{s})$ has five distinct components:

$$\bar{G}_n^\delta(\bar{s}) = G_n(\sigma) + \sum_{s_n \in S_n} g_{n,s_n}^1(\bar{s}_{-n})\sigma_{n,s_n} + \gamma_n(\theta_n, \sigma_{-n}) + \pi_n(\bar{s}_n^1, \bar{s}_{n+1}^0) - \delta\varrho_n(\bar{s}_n^0).$$

The first and the third terms have been defined before. The function ϱ_n in the last term is the convex function defined above. We will specify the other two terms.

$$g_{n,s_n}^1(\bar{s}_{-n}) = \xi_n(\bar{s}_{-n}^0) g_{n,s_n}(\sigma_1, \ldots, \sigma_{n-1}, \sigma_{n-1,n}, \sigma_{n+1}, \ldots, \sigma_N),$$

where $\xi_n(\bar{s}_{-n}^0)$ is one if for each $m \neq n$, $\bar{s}_m^0$ is a vertex of some simplex in $\mathcal{T}_m^1$; otherwise it is zero. The function π_n is 0 if either: (1) $\bar{s}_{n+1}^0$ is a vertex of some simplex in $\mathcal{T}_{n+1}^1$ and $\bar{s}_n^1 = \bar{s}_{n+1}^0$; or (2) $\bar{s}_{n+1}^0$ is not a vertex of such a simplex, but $\bar{s}_n^1 = (s_{n+1}^0, \theta_{n+1}^0)$ (the chosen pure strategy when we defined $\tilde{G}$); elsewhere it is -1. The definition of $\bar{G}^\delta$ clearly implies that it embeds G.

We want to make a couple of remarks about the payoffs. First, the function π_n incentivizes player n to mimic player $n+1$ whenever the latter is choosing a strategy close to $\Sigma_{n+1} \times \Theta_{n+1}^1$: if $n+1$ randomizes over the vertices of a simplex $T_{n+1} \in \mathcal{T}_{n+1}^1$, then player n's best replies must be among the vertices of the simplex. This will be a crucial property, since the choices in $\Sigma_{n-1,n}$ play a role in the evaluation of the bonus function g_n^1. The idea is that whenever the bonus function g_n^1 is active, meaning that all players $m \neq n$ randomize over the vertices of a simplex T_m in $\mathcal{T}_m^1$, then each pure best-reply for player n must choose a vertex of T_{n+1}. On the other hand, if player $n+1$ is randomizing over $S_{n+1} \times \{\, \theta_{n+1}^0 \,\}$ then it follows that the unique best-reply for player n is to choose the previously fixed strategies $\bar{s}_n^1 = (s_{n+1}^0, \theta_{n+1}^0)$.

Our second remark concerns the nature of the payoffs for mixed strategies. For n and each $i = 0, 1$, there is a linear map $p_n^i : \bar{\Sigma}_n :\to \Sigma_{n+i} \times \Theta_{n+i}$ that sends each pure $\bar{s}_n$ to the corresponding mixed strategy in $\Sigma_{n+i} \times \Theta_{n+i}$. For each n, the first and the third terms of the payoffs depend on $\bar{\sigma}$ only through their images under $p^0 = \prod_{n \in \mathcal{N}} p_n^0$; the fourth term depends on all the information in $\bar{\sigma}_n^1$ and $\bar{\sigma}_{n+1}^0$. The second term depends on $\bar{\sigma}_n$ only through p_n^0, but requires the entire information in $\bar{\sigma}_{-n}$, while the last term requires the information in $\bar{\sigma}_n^0$.

3.9. **Wrapping up the Proof.** Let $\bar{\sigma}^*$ be the profile where for each player n, the marginal on $\bar{\Sigma}_n^0$ is (σ^*, θ_n^0) and the marginal on $\bar{\Sigma}_n^1$ is $(s_{n+1}^0, \theta_{n+1}^0)$. For $\delta \geqslant 0$, $\bar{\sigma}^*$ is an equilibrium of $\bar{G}^\delta$, and $(G, \sigma^*) \sim (\bar{G}^\delta, \bar{\sigma}^*)$. We will now show that for δ sufficiently small, this is the only equilibrium of the game, which completes the proof.

Say that a strategy $\bar{\sigma}_n$ is *admissible* for player n if the support of its marginal $\bar{\sigma}_n^0 \in \bar{\Sigma}_n^0$ is the set of vertices of a simplex T_n in $\mathcal{T}_n$. Observe that for any $\delta > 0$, every best reply for player n is admissible. Indeed, the first three components of n's payoff function depend on n's strategy only through its projection to $\Sigma_n \times \Theta_n$ and the fourth is independent of these choices. Therefore, any two strategies for n that project under p_n^0 to the same point in $\Sigma_n \times \Theta_n$ yield the same payoffs for these four terms, leaving the fifth to decide which one is better. But the map ϱ_n is convex, and it is linear precisely on the simplices of $\mathcal{T}_n$, which then forces each best reply to be a mixture over the vertices of a simplex of $\mathcal{T}_n$.

We claim that if $\delta = 0$ and the only admissible equilibrium of $\bar{G}^\delta$ is $\bar{\sigma}^*$, then for sufficiently small $\delta > 0$, $\bar{\sigma}^*$ is the only equilibrium of $\bar{G}^\delta$. To prove this claim, suppose that we have a sequence $\bar{\sigma}^\delta$ of equilibria of $\bar{G}^\delta$ converging to some equilibrium $\bar{\sigma}^0$ of $\bar{G}^0$, then as we saw above $\bar{\sigma}^\delta$ must be admissible, and hence also its limit $\bar{\sigma}^0$. As we have assumed that $\bar{\sigma}^*$ is the unique admissible equilibrium of $\bar{G}^0$, $\bar{\sigma}^0 = \bar{\sigma}^*$. Observe now that for each n, every pure best reply in $\bar{G}^0$ to $\bar{\sigma}^*$ is of the form $(s_n, \theta_n^0, s_{n+1}^0, \theta_{n+1}^0)$ where $s_n \in S_n$ is a best reply to σ^*; and this property holds for best replies to $\bar{\sigma}^\delta$, for small δ. Thus for each such δ, and for each n, $\bar{\sigma}_n^\delta$ is of the form $(\sigma_n^\delta, \theta_n^0, s_{n+1}^0, \theta_{n+1}^0)$, where σ_n^δ is a best reply to σ^* in G. In other words, σ^δ is an equilibrium of G. As σ^δ converges to σ^* and as σ^* is an isolated equilibrium of G, $\sigma^\delta = \sigma^*$ for all small δ. Thus the claim follows and it is sufficient to show that $\bar{\sigma}^*$ is the only admissible equilibrium for $\delta = 0$.

To prove this last point, fix now an admissible equilibrium $\bar{\sigma}$ with marginals $(\bar{\sigma}^0, \bar{\sigma}^1) \in \bar{\Sigma}^0 \times \bar{\Sigma}^1$ of the game $\bar{G}^0$. For each n, let (σ_n, θ_n) and $(\sigma_{n,n+1}, \theta_{n,n+1})$ be the image of $\bar{\sigma}_n$ under p_n^0 and p_n^1, resp. Also, let T_n be the simplex of $\mathcal{T}_n$ generated by the support of $\bar{\sigma}_n^0$ for each n.

Suppose first for each n, T_n belongs to $\mathcal{T}_n^1$. For each n, θ_n assigns probability less than ζ, which is smaller than one, to θ_n^0. Also, for at least two n, $\sigma_n \notin \mathrm{int}_{\Sigma_n}(U_n^*)$: indeed, otherwise there is one player n all of whose opponents m are choosing in $\mathrm{int}_{\Sigma_n}(U_m^*)$, making θ_n^0 the unique optimal choice, which is impossible. Thus, $\sigma_n \notin \mathrm{int}_{\Sigma_n}(U_n^*)$ for at least two n, i.e., $\sigma \notin \mathrm{int}_{\Sigma_n}(U^*)$ and, hence, σ is not an equilibrium of $G \oplus g(\sigma)$. For each n, and each $\bar{s}_{-n}$ in the support of $\bar{\sigma}_{-n}$, $\xi_n(\bar{s}_{-n}^0) = 1$ as $\bar{s}_m^0$ is a vertex of the simplex T_m, which is in $\mathcal{T}_m^1$, for each m; because of the function π_n, the optimality of $\bar{\sigma}_{n-1}^1$ implies that each $\bar{s}_{n-1}^1$ in the support of $\bar{\sigma}_{n-1}^1$ is a vertex of T_n. Therefore, for each $\bar{s}_{-n}$ in the support of $\bar{\sigma}_{-n}$, $\|g^1(\bar{s}_{-n}) - g(\sigma)\| \leqslant \eta$ and then $\|g^1(\bar{\sigma}_{-n}) - g(\sigma)\| \leqslant \eta$. As σ is not an equilibrium of $G \oplus g(\sigma)$, by the choice of η in subsection 3.8, it is not an equilibrium of $G \oplus g^1(\bar{\sigma})$, which contradicts the fact that $\bar{\sigma}$ is an equilibrium of $\bar{G}^0$.

Now suppose that for exactly one n, say $n = 1$, T_n does not belong to $\mathcal{T}_n^1$. Then, θ_1^0 has positive probability under θ_1. Therefore, because of the definition of γ_n, $\sigma_n \in U_n^*$ for $n > 1$, i.e., $\sigma \in U^*$. For $n > 1$, the fact that $\sigma_n \in U_n^*$ and T_n belongs to $\mathcal{T}_n^1$ imply that each s_n in the support of σ_n belongs to U_n (as the diameter of T_n is less than ζ, which is smaller than the distance between U^* and $\partial_\Sigma U$.) Thus, $g_1^1(\sigma) = 0$. We will now show that $g_n^1(\sigma) = 0$ for $n > 1$. The payoff function π_n for each $n \neq N$ forces each strategy $\bar{s}_n^1 = (\sigma_{n,n+1}, \theta_{n,n+1})$ in the support of $\bar{\sigma}_n^1$ to be a vertex of T_{n+1} and hence $\sigma_{n,n+1}$ is in U_{n+1}. Therefore, for $n > 1$, $\sigma_{n-1,n} \in U_n$. Recall that $g_n(\cdot)$ was constructed to be 0 on U. Consequently, for each $n > 1$, $g_n^1(\bar{s}_{-n}) = 0$ for each $\bar{s}_{-n}$ in the support of $\bar{\sigma}_{-n}$, i.e., $g_n^1(\bar{\sigma}_{-n}) = 0$.

The fact that $g^1(\sigma) = 0$, implies that $\sigma = \sigma^*$. Optimality of θ_n for $n > 1$ now requires that it assign probability one to θ_n^0. This is a contradiction: since $T_n \in \mathcal{T}_n^1$, its diameter is smaller than ζ (and hence one), putting it at positive distance from $\Sigma_n \times \{\, \theta_n^0 \,\}$.

Finally, suppose that for at least two players n, T_n does not belong to $\mathcal{T}_n^1$. Then, again because of γ_n, for each n, $\sigma_n \in U_n^*$. We claim that for each n, $g_n^1(\bar{s}_{-n}) = 0$ for each $\bar{s}_{-n}$ in the support of $\bar{\sigma}_{-n}$. Indeed, if for some $m \neq n$, T_m does not have a vertex that belongs to a simplex in $\mathcal{T}_m^1$, then ξ_n is zero by construction at each $\bar{s}_{-n}$ in the support and we are done. Otherwise, again by construction, the diameter of each T_m is less than ζ for each $m \neq n$, which would put each vertex of each T_m in U_m. Therefore, $g_n^1(\cdot)$ is again zero on the support of $\bar{\sigma}_{-n}$.

It follows from the previous paragraph that σ is an equilibrium of G, i.e., $\sigma = \sigma^*$, making $\theta_n = \theta_n^0$. Finally, optimality of $\bar{\sigma}_n^1$ implies that it is $(s_{n,n+1}^0, \theta_{n,n+1}^0)$, as it yields zero with others yielding -1. Thus, $\bar{\sigma} = \bar{\sigma}^*$, which concludes the proof.

4. Discussion

There are two aspects of genericity—one concerning the definition of sustainability and the other the statement of the theorem—that we would like to highlight. First, in the definition of equivalence between game-equilibrium pairs, we require that the games obtained by restricting the strategies to the best replies to the equilibria, rather than to the support of the equilibria, be the same. Of course, if both games are generic, then we could have used either requirement.[11] Also, if we are dealing with two-player games, then, too, we could use just the support due to the fact that a unique equilibrium of a bimatrix game is quasi-strict (Norde [12]). However, in N-person games, uniqueness does not guarantee quasi-strictness (Brandt and Fischer [1]). We now construct an example that exploits this feature of N-person games to show that our theorem would fail if we merely restrict strategies to the support of equilibria. Consider the following three-player game G, where player 3 is a dummy player, whose unique action is W.

	l	r
t	$(6,6,1)$	$(0,0,1)$
b	$(0,0,1)$	$(6,6,1)$

This game has three equilibria: two strict equilibria, (t,l) and (b,r), which have index $+1$, and a mixed equilibrium σ^* where players 1 and 2 mix uniformly, which has index -1.

Define now a game $\bar{G}$ where each player has three choices. Player 1's strategy set is $\{ (T,t), (T,b), B \}$; 2's strategy set is $\{ (L,l), (L,r), R \}$; 3's strategy set is $\{ W, E^w, E^\varepsilon \}$. The payoffs are:

<table>
<tr><td></td><td></td><td></td><td colspan="2">L</td><td>R</td></tr>
<tr><td></td><td></td><td></td><td>l</td><td>r</td><td></td></tr>
<tr><td rowspan="2">$W:$</td><td rowspan="2">T</td><td>t</td><td>$(6,6,1)$</td><td>$(0,0,1)$</td><td rowspan="2">$(3,3,0)$</td></tr>
<tr><td>b</td><td>$(0,0,1)$</td><td>$(6,6,1)$</td></tr>
<tr><td></td><td>B</td><td colspan="3">$(3,0,1)$</td><td>$(0,3,1)$</td></tr>
</table>

[11] Hofbauer states his definition in the body of his paper using the supports, but he adds a footnote that we need to include best replies for nongeneric games.

E^w :		L		R
		l	r	
T	t	$(-3,0,4)$	$(1,4,0)$	$(1,0,1)$
	b	$(1,4,0)$	$(1,4,0)$	
B		$(3,0,0)$		$(0,3,0)$

E^e :		L		R
		l	r	
T	t	$(1,4,0)$	$(1,4,0)$	$(3,0,1)$
	b	$(1,4,0)$	$(-3,0,4)$	
B		$(3,0,0)$		$(0,3,0)$

If we delete B, R, E^w, E^e, we get the game G and so $\bar{G}$ embeds G. The unique equilibrium of this game is the mixed equilibrium σ^* of G. The pair $(\bar{G}, \sigma^*)$ is not equivalent to (G, σ^*) under the definition of the paper since, for e.g., strategy B is a best reply to σ^* but it is not a strategy in G. However, if we merely ask for equivalence using supports, then the two pairs are equivalent and we would make a regular -1 equilibrium sustainable.

The second point about genericity is that we focus on regular equilibria in our theorem in order to align our paper with the Hofbauer-Myerson conjecture. But what really matters, as a careful reading of the proof shows, is one particular implication of regularity, namely that regular equilibria are isolated. Thus, we could obtain a slightly stronger result: an isolated equilibrium is sustainable iff its index is $+1$.

When the game is truly non-generic, i.e., it has nontrivial components of equilibria, then there is an obvious extension of the definition of sustainability to components, which replaces equilibria with components. However, such a refinement may fail to exist: for example, we could have two components of equilibria in a game, one with index $+2$ and the other with index -1; it is impossible to obtain either as the unique solution of a larger game.

Even when a given game has a unique component of equilibria, which has index $+1$, then the indeterminacy of equilibria may persist in any equivalent game with an expanded set of strategies, as we will now show. Consider the following game G:

	L	R
T	$(1,1)$	$(0,1)$

This game has a unique component of equilibria $\{(T, yL + (1-y)R),\ y \in [0,1]\}$ and two pure equilibria (T, L) and (T, R). We claim that no equilibrium of G can be made unique by adding strategies. To see this, let $\widehat{G}$ be obtained from G by adding rows $r \in R$ and columns $l \in L$ with associated payoffs $g(r,l)$ and suppose that an equilibrium σ of G can be made unique in $\widehat{G}$. Observe that σ cannot be pure because otherwise it will be non quasi-strict in G, and so non-quasi strict in $\widehat{G}$, a contradiction with Norde's result. Suppose now that $\sigma = (T, yL + (1-y)R)$, $y \in]0,1[$. Let $\tau = (T, L)$ be a pure equilibrium. Then τ is the unique equilibrium of the following perturbation G_ε.

	L	R
T	$(1, 1+\varepsilon)$	$(0, 1)$

Now, we add to G_ε the extension of G which are the rows $r \in R$ and the columns $l \in L$ with associated payoffs $g(r, l)$. This defines a perturbed game $\widehat{G}_\varepsilon$ of $\widehat{G}$. Since $\widehat{G}$ has a unique equilibrium σ, by upper-hemi-continuity of the equilibrium correspondence, $\widehat{G}_\varepsilon$ has an equilibrium σ_ε that converges to σ as ε goes to zero. Because τ is the unique equilibrium of G_ε, necessarily there is $r \in R$ (or $l \in L$) in the support of σ_ε (otherwise, σ_ε would be an equilibrium of G_ε and so $\sigma_\varepsilon = \tau$, a contradiction). By continuity, r (or l) is still a best response to σ, showing that σ is not a quasi-strict equilibrium in $\widehat{G}$, establishing a contradiction.

Interestingly, if we can add also players then any equilibrium of the last game can be made unique.[12] For example, the pure and non quasi-strict equilibrium (T, L) of the game above is the projection of the unique equilibrium (T, L, W) of the following 3-player game.[13]

$W:$

	L	R
T	$(1, 1, 1)$	$(1, 1, 0)$
B	$(1, 0, 1)$	$(0, 1, 1)$

$E:$

	L	R
T	$(0, 1, 1)$	$(1, 0, 1)$
B	$(1, 0, 0)$	$(0, 1, 0)$

The above considerations suggest that an extension of the definition of sustainability to nongeneric games—and index-theoretic characterizations of it—is a nontrivial question, which we hope future research will address.

References

[1] Brandt, F., and F. Fischer (2008): "On the Hardness and Existence of Quasi-Strict Equilibria," SAGT 2008, LNCS 4997, 291-302.

[2] Demichelis, S., and F. Germano (2000): "On the Indices of Zeros of Vector Fields," *Journal of Economic Theory*, 94, 192–218.

[3] Govindan, S., and R. Wilson (2005): "Essential Equilibria," *Proceedings of the National Academy of Sciences*, 102, 15706–11.

[4] Herings, P.J-J. and R. Peeters (2010). Homotopy methods to compute equilibria in game theory, *Economic Theory*, 42, 119–156.

[5] Hofbauer, J. (2000): "Some Thoughts on Sustainable/Learnable Equilibria," Mimeo.

[6] Kohlberg, E., and J.-F. Mertens (1986): "On the Strategic Stability of Equilibria," *Econometrica*, 54, 1003–37.

[7] Loera, J.A., J. Rambau, and F. Santos (2010): *Triangulations: Structures for Algorithms and Applications.* Berlin: Springer-Verlag.

[8] McLennan, A. (2018): *Advanced Fixed Point Theory for Economics*. Singapore: Springer-Verlag.

[12]We thank Joseph Hofbauer for raising this issue.

[13]This is a mild strengthening of the example in Brandt and Fischer [1], where the unique equilibrium –not quasi-strict– was in mixed strategies.

[9] McLennan, A. (2016): *The Index +1 principle*. Mimeo.

[10] Myerson, R.B. (1996): "Sustainable Equilibria in Culturally Familiar Games," in *Understanding Strategic Interaction: Essays in Honor of Reinhard Selten*, edited by W. Albers et al, Springer, 111–21.

[11] Nash, J.F. (1951): "Noncooperative Games," *Annals of Mathematics*, 54, 286–95.

[12] Norde, H. (1999): "Bimatrix Games Have a Quasi-Strict Equilibria," *Mathematical Programming*, 85, 35–49.

[13] Ritzberger, K. (1994): "The Theory of Normal Form Games from the Differentiable Viewpoint," *International Journal of Game Theory*, 23: 207-236.

[14] Spanier, E.H. (1966): *Algebraic Topology.* New York: McGraw Hill: New York, Reprinted in New York: Springer-Verlag, 1989.

[15] von Schemde, A., and B. von Stengel (2008): "Strategic Characterization of the Index of an Equilibrium," SAGT 2008, LNCS 4997, 242–54.

[16] von Schemde, A. (2005): *Index and Stability in Bimatrix Games*, Lecture Notes in Economics and Mathematical Systems, Springer.

Appendix

Construction of the Triangulation. Here we construct a triangulation $\mathcal{T}_n$ of $\Sigma_n \times \Theta_n$ for each n with the properties stated in subsection 3.8. We start with some definitions.

A *simplex* T in $\mathbb{R}^d$ is the convex hull of affinely independent points $x_0, x_1, \ldots, x_k$ $(k \leqslant d)$; a *face* of T is the convex hull of a subset of the points x_i. A *triangulation* $\mathcal{T}$ of a polytope $C \subset \mathbb{R}^d$ is a finite collection of simplices T in $\mathbb{R}^d$ such that: (1) if $T \in \mathcal{T}$, so is every face of T; (2) the intersection of two simplices in $\mathcal{T}$ is a face of both (possibly empty); (3) the union of the simplices in $\mathcal{T}$ equals C.

Throughout this Appendix, we use the ℓ_2 norm, unless we specify differently. Suppose that the convex hull C of the points $x_0, \ldots x_k$ is d-dimensional. Given a finite collection $\{ x_0, x_1, \ldots, x_k \}$ of points, (where k is now an arbitrary positive integer) let C be its convex hull. Suppose that the x_i's are in *general position* in $\mathbb{R}^d$—i.e., no $d+2$ points lie in any $(d-1)$-sphere (centered at any point and of any radius) in $\mathbb{R}^d$. We can construct a triangulation of the convex hull C, called the *Delaunay triangulation*, as follows. (Cf. Loera et al [7] for details). Let D be the convex hull of the set of points $(x_i, \|x_i\|^2) \in \mathbb{R}^{d+1}$, $i = 1, \ldots k$. Let D_0 be the lower envelope of D. The natural projection $(x, y) \mapsto x$ from D_0 to $\mathbb{R}^d$ is C and D_0 is the graph of a piece-wise linear convex function $\varrho : C \to \mathbb{R}$ with the property that the subsets on which ϱ is linear are simplices, whose projections then yield the cells of a triangulation of C.

There is a dual representation of the Delaunay triangulation, known as the *Voronoi Diagram*, which works as follows. For each i, let P_i be the polyhedron in $\mathbb{R}^d$ consisting of points y in $\mathbb{R}^d$ such that $\|y - x_i\| \leqslant \|y - x_j\|$ for all $j \neq i$. We then have a polyhedral complex (which is exactly like a simplicial complex but with polyhedra rather than simplices) where the maximal polyhedra are the P_i. There is an edge between two vertices x_i and x_j in the Delaunay triangulation iff the polyhedra P_i and P_j have a nonempty intersection. Also, the intersection of $d+1$ of these polyhedra when nonempty is a single point (because of general position), that is then the center of a ball that contains $d+1$ points of the collection on its boundary and no other point in the ball itself—these $d+1$ points span a d-dimensional simplex in the Delaunay triangulation.

For our purposes, we need a triangulation with the diameter of certain simplices to be smaller than ζ. Therefore, we will now show how to obtain such triangulations with arbitrarily small diameters. Let C be a d-dimensional polytope; let B be the subset of C given by the set of $x \in C$ s.t. $a \cdot x \leqslant b$, for some $a \in \mathbb{R}^d$ and $b \in \mathbb{R}$ such that B is a proper subset of C with dimension d. Assume that the set of vertices of B and C are in general position for a Delaunay triangulation of C. Let $\delta > 0$ be such that $\|x - y\| \geqslant \delta/2$, for all $x \in B$ and vertices $y \in C \setminus B$.

We will now introduce additional vertices to obtain our desired triangulation. Let X_δ be a collection of points in C such that: (1) X_δ contains the vertices of C but no other point in $C \setminus B$; (2) for $x \in B$, there is a point $x_\delta \in B \cap X_\delta$ such that $\|x - x_\delta\| < \delta/2$ and x_δ belongs to the face of B that contains x in its interior; (3) every point in $\text{int}_C(B) \cap X_\delta$ is at least $\delta/2$ from $\partial_C B$; (4) the points in X_δ are in general position for a Delaunay triangulation of C.

Let $\mathcal{T}_\delta$ be the triangulation of C using the point-set X_δ. This triangulation achieves two properties: (i) every simplex with vertices in B has diameter at most δ; (ii) every simplex of $\mathcal{T}_\delta$ that has a vertex outside B does not intersect $\text{int}_C(B)$. We start by proving that (i) is indeed satisfied. Define $r : \mathbb{R}^d \to B$ by letting $r(x)$ be the point in B that is closest to x. If $r(x) \neq x$, $r(x)$ belongs to a proper face of B, and then we can write $r(x)$ as $x - p$ where p is a normal for a supporting hyperplane at $r(x)$ with $p \cdot r(x) \geqslant p \cdot y$ for all $y \in B$. If $r(x) \in \text{int}_C(B)$, then $r(x)$ is at the boundary of C and so $p \cdot r(x) \geqslant p \cdot y$ for all $y \in C$ as well. Suppose $r(x) \neq x$ and let $r(x) = x - p$. Let y be a point such that $p \cdot y \leqslant p \cdot r(x)$. Let z be the nearest-point projection of y onto the line from x through $r(x)$. Then

$$\|x - y\|^2 = (\|x - r(x)\| + \|r(x) - z\|)^2 + \|z - y\|^2 \geqslant \|r(x) - x\|^2 + \|r(x) - y\|^2,$$

with the inequality being an equality iff $z = r(x)$, i.e., $p \cdot y = p \cdot r(x)$.

We are now ready to prove that $\mathcal{T}_\delta$ has the requisite property. Let x_δ be a point in $X_\delta \cap B$ and let x be a point in $\mathbb{R}^d$ that belongs to the Voronoi polyhedron $P(x_\delta)$ of x_δ. We claim that $\|r(x) - x_\delta\| < \delta/2$. If $r(x) = x$, this follows directly from Property (2) of X_δ. Suppose that $r(x) \neq x$. Then $r(x)$ belongs to the interior of a proper face B' of B and as we saw in the last paragraph, $r(x)$ can be written as $x - p$. By definition of $r(x)$, $p \cdot x_\delta \leqslant p \cdot r(x)$ and thus: $\|x - x_\delta\|^2 \geqslant \|r(x) - x\|^2 + \|r(x) - x_\delta\|^2$. By Property (2), there exists y_δ in $B' \cap X_\delta$ such that $\|r(x) - y_\delta\| < \delta/2$. Obviously $p \cdot y_\delta = p \cdot r(x)$ and since $x \in P(x_\delta)$, it follows that $\|x - x_\delta\|^2 \leq \|x - y_\delta\|^2 < \|r(x) - x\|^2 + \delta^2/4$; therefore, $\|r(x) - x_\delta\| < \delta/2$, as claimed. Observe that this also proves that $\|x - x_\delta\|^2 < \|r(x) - x\|^2 + \delta^2/4$, a fact we will use below.

From the above paragraph, for each $x_\delta \in X_\delta \cap B$ and each $x \in P(x_\delta)$, the distance between $r(x)$ and x_δ is less than $\delta/2$; therefore, the diameter of each simplex in B is less than δ. Actually, letting x_δ and y_δ be two vertices of a simplex in B, then their Voronoi cells intersect, so we can take x in the intersection. Since $r(x)$ is of distant $\delta/2$ from x_δ and $\delta/2$ from y_δ, x_δ and y_δ are of distant less than δ. This concludes the proof that $\mathcal{T}_\delta$ satisfies (i).

We now prove that $\mathcal{T}_\delta$ satisfies (ii): for this, it is sufficient to show that the intersection of $P(x_\delta)$ and $P(y_\delta)$ is empty for all $x_\delta \in \text{int}_C(B) \cap X_\delta$ and y_δ in $X_\delta \setminus B$. Take such a pair x_δ, y_δ. Fix $x \in P(x_\delta)$. If $r(x) = x$, then $\|x - x_\delta\| < \delta/2$, while by the definition of δ, $\|x - y_\delta\| \geqslant \delta/2$ and thus $x \notin P(y_\delta)$. Suppose $r(x) \neq x$. Since $x_\delta \in \text{int}_C(B)$, by Property (3) of the set X_δ, $r(x)$ cannot belong to the face of B described by the equation $a \cdot y = b$, as the distance between x_δ and $r(x)$ is greater than $\delta/2$. Writing $r(x)$ as $x - p$, we then infer that p is a normal to a hyperplane containing one of the faces of C and thus $p \cdot y_\delta \leqslant p \cdot r(x)$. Hence, $\|x - y_\delta\|^2 \geqslant \|r(x) - x\|^2 + \|r(x) - y_\delta\|^2 \geqslant \|r(x) - x\|^2 + \delta^2/4$ by the definition of δ, while as we saw in the previous paragraph, $\|x - x_\delta\|^2 < \|r(x) - x\|^2 + \delta^2/4$; thus again $x \notin P(y_\delta)$ and we are done.

For our problem of triangulating $\Sigma_n \times \Theta_n$, letting d_n be its dimension, take an affine function $F_n : \Sigma_n \times \Theta_n \to \mathbb{R}^{d_n}$ with the following properties: (1) F_n maps $\Sigma_n \times \Theta_n$ homeomorphically onto its image C; (2) it maps the vertices of the form (s_n, θ_n^0), for $s_n \in S_n$, to unit vectors in $\mathbb{R}^{d_n}$ and all other vertices of $\Sigma_n \times \Theta_n$ to vectors with ℓ_2-norm strictly greater than one; (3) the polytope $D_n \equiv \{(\sigma_n, \theta_n) \in \Sigma_n \times \Theta_n | \theta_n$ assigns probability at most 2ζ to θ_n^0 $\}$ is such that its vertices and the vertices of C are in general position for the Delaunay triangulation of C. The Delaunay triangulation of C induces (through F_n^{-1}) a triangulation of $\Sigma_n \times \Theta_n$, but that is not enough for our purposes, as it does not satisfy property (2) of Subsection 3.8. To get this additional property as well, we have to add more vertices.

Since F_n is an affine homeomorphism, $\|x - y\|_\infty \leqslant M\|F_n(x) - F_n(x)\|$ for some $M > 0$. Let $B = F_n(D_n)$. Let δ be smaller than $M\zeta$. Using the construction above, we now have a triangulation $\mathcal{T}_\delta$ of C where each point in $\text{int}_C(B)$ belongs to a simpex with diameter less than δ, giving us properties (1) and (2) of Subsection 3.8. As for property (3), if ϱ_n is the convex function associated to the Delaunay triangulation $\mathcal{T}_\delta$, the composition $\varrho_n \circ F_n$ is convex and linear precisely on each cell of the triangulation. By Property (2) of the function

F_n, our convex function takes value one on $\Sigma_n \times \{ \theta_n^0 \}$ and is strictly above one elsewhere. Subtracting now 1 from $\varrho_n \circ F_n$ we have a convex function satisfying property (3).

Department of Economics, University of Rochester, NY 14627, USA.
Email address: s.govindan@rochester.edu

CNRS, Lamsade, University of Paris Dauphine-PSL, 75016 Paris, France, and Department of Computer Science, University of Liverpool, Liverpool L69 3BX, UK
Email address: rida.laraki@gmail.com

Dept. of Economics, University of New South Wales, Sydney NSW 2052, Australia.
Email address: pahl.lucas@gmail.com