



This is a repository copy of *O'Neill's Theorem for Games*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/227173/>

Version: Accepted Version

Article:

Govindan, S., Laraki, R. and Pahl, L. orcid.org/0000-0002-0268-371X (2025) O'Neill's Theorem for Games. *Mathematics of Operations Research*. ISSN 0364-765X

<https://doi.org/10.1287/moor.2022.0304>

© 2025 INFORMS. This is an author-produced version of a paper subsequently published in *Mathematics of Operations Research*. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

O'NEILL'S THEOREM FOR GAMES*

SRIHARI GOVINDAN[‡], RIDA LARAKI^{**}, AND LUCAS PAHL[†]

ABSTRACT. We present the following analog of O'Neill's Theorem (Theorem 5.2 in [17]) for finite games. Let C_1, \dots, C_k be the components of Nash equilibria of a finite normal-form game G . For each i , let c_i be the index of C_i . For each $\varepsilon > 0$, there exist pairwise disjoint neighborhoods V_1, \dots, V_k of the components such that for any choice of finitely many distinct completely mixed strategy profiles $\{\sigma^{ij}\}_{ij}, \bar{\sigma}^{ij} \in V_i$ for each $i = 1, \dots, k$ and numbers $r_{ij} \in \{-1, 1\}$ such that $\sum_j r_{ij} = c_i$, there exists a normal-form game \bar{G} obtained from G by adding duplicate strategies and an ε -perturbation \bar{G}^ε of \bar{G} such that the set of equilibria of \bar{G}^ε is $\{\bar{\sigma}^{ij}\}_{ij}$, where for each i, j : (1) $\bar{\sigma}^{ij}$ is equivalent to the profile σ^{ij} ; (2) the index $\bar{\sigma}^{ij}$ equals r_{ij} .

1. INTRODUCTION

Multiplicity of equilibria is a pervasive phenomenon in games, and refinement theory aims to reduce it by strengthening the Nash criterion through the imposition of additional requirements. For most well-known refinement concepts, their power derives from asking for robustness of the equilibria with respect to some type of perturbation. Perfect [21], Proper [16] and Stable equilibria (either the Kohlberg-Mertens [10] or the Mertens [14] variants) are examples of concepts that require robustness of Nash equilibria with respect to perturbations of the players' strategies. Essentiality [23] and Hyperstability [10] are examples of concepts that require robustness to payoff perturbations. Most strategy perturbations can be encoded in payoff perturbations, which implies that concepts that require robustness to payoff perturbations are usually more stringent than the strategic ones. In this paper we focus on payoff perturbations.

Every game induces an associated fixed-point problem, where the fixed-point map (usually called a Nash-map) or best-reply correspondence has as its fixed points the Nash equilibria of the game. Small perturbations of a game generate close-by perturbations of the fixed-point map, and therefore fixed-point theory can also inform us in the search for equilibria that are robust to payoff perturbations. The precise fixed-point theoretic tool that allows us to identify robustness of equilibria

Date: March 28, 2025.

* We thank Klaus Ritzberger and Mich Tvede for comments and suggestions. Lucas thanks research support from the Hausdorff Center for Mathematics (DFG project no. 390685813).

[†] SCHOOL OF ECONOMICS, UNIVERSITY OF SHEFFIELD, SHEFFIELD, UNITED KINGDOM.

[‡] DEPARTMENT OF ECONOMICS, UNIVERSITY OF ROCHESTER, ROCHESTER, USA.

^{**} MORROCAN CENTER FOR GAME THEORY, UM6P, RABAT, MOROCCO.

Email addresses: pahl.lucas@gmail.com, s.govindan@rochester.edu, rda.laraki@um6p.ma.

Corresponding author: Lucas Pahl.

JEL classification code: C72; KEYWORDS: Game Theory, Index Theory, Fixed Point Theory.

to perturbations of the fixed-point map is called the *fixed-point index*. The fixed-point index is an integer number associated to each connected component of fixed points of a map and provides a characterization of robustness: a component is robust to perturbations of its fixed-point map to nearby maps if and only if its index is nonzero.

The lesson for game theory obtained from this characterization is somewhat ambiguous: though nonzero index components are robust to payoff perturbations (cf. [19]), the converse is not true (see [7]). Moreover, a fundamental result in fixed point theory, O'Neill's Theorem (Theorem 5.2 in [17]) implies that given any fixed point in a component, there is a nearby map for which it is the unique fixed point in a neighborhood of the component. Though this result permits the interpretation that no fixed point in a connected component of fixed points is particularly distinguishable from the point of view of robustness to perturbations of the map, it yields no immediate prescription for game theory in terms of equilibrium selection: maps that approximate a Nash map may not be Nash maps of payoff-perturbed games, and therefore the same lesson cannot be translated immediately to selection of equilibria through payoff robustness.

Govindan and Wilson [4] proved that if we allow for duplicate strategies, then the class of fixed point maps generated by payoff perturbations is rich enough to capture the notion of essentiality in fixed point theory, i.e., of robustness to map-perturbations.¹ More precisely, they proved that a Nash component C of a finite game has index zero if and only if for any $\varepsilon > 0$, there exists an equivalent game and an ε -payoff perturbation of the equivalent game with no equilibria near the equivalent component to C . This result provides a clear picture of the relation between robustness to map perturbations and to payoff perturbations, but is still incomplete. What about equilibrium components of nonzero index? What is the structure of the equilibria generated by payoff perturbations around positive index components?

For general fixed-point problems, this question is answered by O'Neill's Theorem, which asserts that if $f : U \rightarrow X$ is a continuous function in a Euclidean neighborhood U of a fixed point component C of f , r_1, \dots, r_k are integers whose sum is the fixed-point index of C , and x_1, \dots, x_k are distinct points of C , then there is a map arbitrarily close to f whose fixed points are x_1, \dots, x_k , with the fixed point index of each x_i being r_i . The result of O'Neill provides us with the precise topological invariant to understand which fixed-point components can be robust, namely, the index of a component: for example, if a component has an index of $+2$, one can select any finite number of points inside the component and assign to them integers such that, as long as the sum of the integers is $+2$, an arbitrarily nearby map exists with those points as the only fixed points around the component, and the indices allocated to them are precisely the integers we have chosen.

¹The addition of duplicate strategies is not only of technical importance to bridge the gap between the space of maps and that of payoff perturbations in the result of [4], but is important in the theory of refinements as a decision-theoretic property: Kohlberg and Mertens list the property of *Invariance* as a requirement a solution concept should satisfy, in the sense that a solution concept should not depend on the addition/elimination of duplicate strategies to the game, i.e., the solution should depend only on the reduced normal form.

For game-theoretic problems, the immediate question is whether a similar addition of duplicate strategies and perturbation of payoffs of equivalent games would yield like in [4] an analog of O'Neill's Theorem for games. This is the main result we present in this paper. In essence, we prove that one can select in each equilibrium component a finite number of equilibria and associate to each one of them an integer equal to $+1$ or -1 , such that the sum of the numbers allocated to the selected points of the same component equals its index. For each $\varepsilon > 0$, one can then construct an equivalent game—obtained by adding duplicate strategies to the original game—and an ε -payoff perturbation of this equivalent game with a single Nash equilibrium close to each selected point, each of them having an index precisely equal to the pre-assigned number.

Let us illustrate this result in an example. Consider the following bi-matrix game (from Kohlberg and Mertens [10]). It has a unique Nash component of equilibria, homeomorphic to a circle.

	<i>L</i>	<i>M</i>	<i>R</i>
<i>t</i>	(1, 1)	(0, -1)	(-1, 1)
<i>m</i>	(-1, 0)	(0, 0)	(-1, 0)
<i>b</i>	(1, -1)	(0, -1)	(-2, -2)

Say we would like to make $(\frac{1}{2}t + \frac{1}{2}b, L)$ the unique equilibrium (and so necessarily with index $+1$) of a perturbed equivalent game, we duplicate strategy *L* and perturb the payoffs as follows.

	<i>L</i>	<i>L'</i>	<i>M</i>	<i>R</i>
<i>t</i>	$(1 + \varepsilon, 1)$	$(1, 1 + \varepsilon)$	$(0 + \varepsilon, -1)$	$(-1 + \varepsilon, 1)$
<i>m</i>	$(-1, 0)$	$(-1, +\varepsilon)$	$(0, 0)$	$(-1, 0)$
<i>b</i>	$(1, -1 + \varepsilon)$	$(1 + \varepsilon, -1)$	$(0, -1)$	$(-2, -2)$

As *m*, *M* and *R* are strictly dominated in the above game, eliminating them leads to a two-by-two game where the unique Nash equilibrium is $(\frac{1}{2}t + \frac{1}{2}b; \frac{1}{2}L + \frac{1}{2}L')$ (necessarily with index $+1$) which is equivalent in the original game to $\sigma = (\frac{1}{2}t + \frac{1}{2}b, L)$.

Consider now the three Nash equilibria (t, L) and (b, L) and $\sigma = (\frac{1}{2}t + \frac{1}{2}b, L)$ and associate to the first two the integer $+1$ and to the last the integer -1 . We will show how to add to the original game the same duplicate strategy *L'* as above and perturb the payoffs so as to generate exactly three equilibria in a perturbed game. Consider the following perturbation:

	<i>L</i>	<i>L'</i>	<i>M</i>	<i>R</i>
<i>t</i>	$(1 + \varepsilon, 1 + \varepsilon)$	$(1, 1)$	$(0 + \varepsilon, -1)$	$(-1 + \varepsilon, 1)$
<i>m</i>	$(-1, 0)$	$(-1, +\varepsilon)$	$(0, 0)$	$(-1, 0)$
<i>b</i>	$(1, -1)$	$(1 + \varepsilon, -1 + \varepsilon)$	$(0, -1)$	$(-2, -2)$

Strategies *m*, *M* and *R* are strictly dominated. Eliminating them leads to a battle of the sexes where only three Nash equilibria are present: the two strict equilibria (t, L) and (b, L') (hence with indices $+1$ each) and a completely mixed equilibrium $(\frac{1}{2}t + \frac{1}{2}b; \frac{1}{2}L + \frac{1}{2}L')$ with index -1 . But since *L'* was a duplicate of *L*, we obtain as Nash equilibria (t, L) and (b, L) and $(\frac{1}{2}t + \frac{1}{2}b, L)$, and their

indices correspond to the integers we fixed previously. Hence, perturbations to nearby equivalent games can single out σ as an isolated equilibrium with different indices: in the first perturbation we obtained it with index $+1$ and in the second with index -1 .

Our proof of the extension of O'Neill's Theorem to games is inspired by the one in Govindan Wilson [4]. Though the steps of the proof are similar, the details are very different and pose different technical challenges. It requires more preliminary work in the realm of fixed point theory in order to produce a suitable approximation of the best-reply correspondence (which is achieved with the help of a fundamental result in obstruction theory, namely, the Hopf Extension Theorem). Then we construct a game that has essentially that approximation as its best-reply correspondence (which will require tools from the theory of triangulations).

The paper is organized as follows. Section 2 sets up the problem and states the main result. Section 3 gives a detailed outline of the proof and includes a table of notation that is used in the proof. Section 4 contains the proof of the theorem, which is presented in a sequence of steps.

2. PRELIMINARY DEFINITIONS AND STATEMENT OF THE THEOREM




We start by introducing some notational conventions. A finite normal-form game is a tuple $G = (\mathcal{N}, (S_n)_{n \in \mathcal{N}}, (G_n)_{n \in \mathcal{N}})$, where $\mathcal{N} = \{1, \dots, N\}$ is the set of players, S_n is a finite set of pure strategies of player n , $S \equiv \times_{n \in \mathcal{N}} S_n$, and $G_n : S \rightarrow \mathbb{R}$ is the payoff function of player n . We denote the set of mixed strategies of player n by $\Sigma_n \equiv \Delta(S_n)$, and the set of mixed-strategy profiles is denoted $\Sigma \equiv \prod_n \Sigma_n$. For each n , we continue to use G_n to denote the extension of n 's payoff function to Σ . The best-reply correspondence of G is denoted BR.

For $V_n \subseteq \Sigma_n$, ∂V_n denotes the boundary of V_n in the affine space generated by Σ_n ; $\partial_{\Sigma_n} V_n$ refers to the boundary of V_n in Σ_n . Similarly, ∂V and $\partial_{\Sigma} V$ denote, resp., the boundaries of $V \subseteq \Sigma$ in the affine space generated by Σ and in Σ . Finally, $\|\cdot\|$ denotes the ℓ_{∞} -norm on Euclidean spaces.

A *polyhedral subdivision* or *polyhedral complex* \mathcal{P} of a polytope $X \subset \mathbb{R}^m$ is a collection of polytopes in \mathbb{R}^m that satisfies three conditions: first, the union of the polytopes in \mathcal{P} is equal to X ; second, every face of a polyhedron in the collection belongs to the collection; third, the intersection of any two polytopes in \mathcal{P} is a face of each (possibly empty). Given a polyhedral subdivision \mathcal{P} of X , a typical member of \mathcal{P} is called a *cell*. The cells of dimension 0 are called *vertices*. When the polyhedral subdivision is made of simplices, it is called a *triangulation*. The *closed star of a vertex v of a triangulation \mathcal{T}_n* is the union of simplices that have v as a vertex; and the *simplicial neighborhood* of the closed star of v is the union of all simplices that intersect the closed star.

Example 2.1. Figure 1 provides an example of a triangulation of a square. This triangulation has 16 vertices (labeled from a to p), 33 1-simplices, and 18 2-simplices. The space of this triangulation is the union of all the simplices, therefore, the square itself. The closed star of vertex f in the triangulation is depicted in red in Figure 2 and the simplicial neighborhood of this closed star is depicted in Figure 3.

FIGURE 1. Triangulation of a square



FIGURE 2. The closed star of vertex f FIGURE 3. The simplicial neighborhood of the closed star of f 

The *face complex* of a polytope $X \subset \mathbb{R}^m$ is the polyhedral subdivision \mathcal{P} of X given by the collection of its faces (X and the empty set are also considered faces, so they also belong to the face complex). If \mathcal{P} is any polyhedral subdivision of X , then a *subcomplex* of \mathcal{P} is a subset of \mathcal{P} that is itself a polyhedral subdivision. The *space of a complex* is the union of the cells that belong to it.

Example 2.2. In Figure 4 we have a 2-simplex ABC . The face complex of this 2-simplex is given by the empty face, vertices A, B, C , 1-simplices $[A, B], [B, C], [A, C]$, and the 2-simplex $[A, B, C]$. We highlighted in red the space of the subcomplex of the face complex composed by vertices A, B, C and 1-simplices $[A, B]$ and $[B, C]$.

Given a polyhedral subdivision \mathcal{T}_n of space X and $x \in X$, the *carrier of x in subdivision \mathcal{T}_n* is the unique cell of \mathcal{T}_n that contains x in its interior (relative to the affine space generated by the

FIGURE 4. Face complex and subcomplex

FIGURE 5. Carrier of point x in a triangulation

cell). Whenever the underlying polyhedral subdivision is understood, we shall omit reference to \mathcal{T}_n . Figure 5 illustrates the carrier of point x in a triangulation of a 2-simplex: the carrier of point x is the 2-simplex AGC .

For technical reasons, our proof will require considering the notion of a game defined over slightly more general strategy sets than simplices for each player (cf. Pahl [18]).

Definition 2.3. A *polytope-form game* is a tuple $G = (\mathcal{N}, (P_n)_{n \in \mathcal{N}}, (V_n)_{n \in \mathcal{N}})$, where $\mathcal{N} = \{1, \dots, N\}$ is the set of players, P_n is a polytope in \mathbb{R}^{m_n} (the set of strategies of player n) and $G_n : \prod_{n \in \mathcal{N}} P_n \rightarrow \mathbb{R}$ is the payoff function of player n , which is affine in each coordinate $p_j \in P_j, j = 1, \dots, N$.²

Every finite game in mixed strategies can be viewed as a game in polytope-form since strategy sets are simplices hence polytopes and payoff functions are multi-linear. Going the other way, given

²The original term to denote this class of games is *strategic-form games* and was coined in [15]. In order to avoid confusion with the current terminology of “strategic-form games”, we opted for the “polytope-form game” designation.

a polytope-form game, we can define a normal-form game where the set of pure strategies of each player is the set of vertices of P_n ; these two are “equivalent” in the sense to be defined now.

Definition 2.4. Two games $G = (\mathcal{N}, (P_n)_{n \in \mathcal{N}}, (G_n)_{n \in \mathcal{N}})$ and $G' = (\mathcal{N}, (P'_n)_{n \in \mathcal{N}}, (G'_n)_{n \in \mathcal{N}})$ in polytope form are *equivalent* if there exist a polytope-form game $\tilde{G} = (\mathcal{N}, (\tilde{P}_n)_{n \in \mathcal{N}}, (\tilde{G}_n)_{n \in \mathcal{N}})$, affine and surjective maps $\phi_n : P_n \rightarrow \tilde{P}_n$ and $\phi'_n : P'_n \rightarrow \tilde{P}_n$, such that $G_n(p) = \tilde{G}_n(\phi(p))$ and $G'_n(p') = \tilde{G}_n(\phi'(p'))$, with $\phi \equiv \times_n \phi_n, \phi' \equiv \times_n \phi'_n$.


Let G and \tilde{G} be as in Definition 2.4. We say $p_n \in P_n$ *projects* to $\tilde{p}_n \in \tilde{P}_n$, if $\phi_n(p_n) = \tilde{p}_n$; similarly, we say $p \in P$ projects to $\tilde{p} \in \tilde{P}$, if $\phi(p) = \tilde{p}$. We say that a set $A \subseteq P$ projects to $\tilde{A} \subseteq \tilde{P}$ if $\phi(A) = \tilde{A}$. Given a normal-form game G , a *duplicate strategy* s_n of player n is a pure strategy that projects to a (possibly mixed) strategy of player n that is distinct from s_n . A normal-form game \bar{G} obtained by adding duplicate strategies from G is a finite game where each pure strategy is either a pure strategy of G or is a duplicate of a strategy of G (pure or mixed).

Example 2.5. Assume player n has three pure strategies $\{A, B, C\}$ in a finite normal-form game G . His mixed strategy set is therefore a 2-simplex as depicted in Figure 6. We now introduce the duplicate strategy b for player n , which corresponds to the barycenter of the 2-simplex. The newly defined pure strategy set of the player is then $\{A, B, C, b\}$, and the mixed strategy set is a tetrahedron. Payoffs for the player in the equivalent game \bar{G} where the strategy set of the player is a tetrahedron are defined by projecting b to the 2-simplex, that is: $\bar{G}_n(b, s_{-n}) = \frac{1}{3}G_n(A, s_{-n}) + \frac{1}{3}G_n(B, s_{-n}) + \frac{1}{3}G_n(C, s_{-n})$.

2.1. Fixed Point Index. We provide a brief review of the concept of fixed point index here. For a comprehensive and axiomatic treatment of the subject we refer the reader to [13]. Let X be a convex set and A a compact subset of X . Say that a correspondence $\varphi : A \rightrightarrows X$ is well-behaved if it is nonempty, compact, convex-valued and upper semicontinuous. Suppose φ is a well-behaved correspondence without fixed points on the boundary of A (in X). The index of φ (over A) is an integer that serves as an algebraic count of the number of fixed points of φ in A . In the case where A is the closure of an open subset of an Euclidean space, and φ is a function, the index of φ can be computed as the degree of the displacement d_φ of φ , which is defined as $d_\varphi \equiv \text{Id} - \varphi$.

We invoke three important properties of the fixed point index. The first property is continuity: there exists $\varepsilon > 0$ such that for all well-behaved correspondences $\varphi' : A \rightrightarrows X$ whose graph is in the ε -neighborhood of that of φ , the index of φ' is the same as that of φ . Since a well-behaved correspondence φ can be approximated by functions whose graphs are contained in arbitrarily small neighborhoods of the graph of φ , the index of φ can be computed by approximating it by a function.

The second property, which is related to continuity, is homotopy. Given two well-behaved correspondences $\varphi, \varphi' : A \rightrightarrows X$, if $\lambda\varphi + (1 - \lambda)\varphi'$ has no fixed points on the boundary of A for any $\lambda \in [0, 1]$, then the index of φ and φ' are the same.

FIGURE 6. Strategy set of player n and equivalent strategy set

The third property we need is the multiplicative property. If $\varphi_1 : A_1 \rightrightarrows X_1$ and $\varphi_2 : A_2 \rightrightarrows X_2$ are two well behaved correspondences without fixed points on the boundaries of A_1 and A_2 , the index of $\varphi_1 \times \varphi_2$ is the product of the indices of φ_1 and φ_2 .

Suppose $\varphi : X \rightrightarrows X$ is a well-behaved correspondence and X is compact and convex. Let C be a component of fixed points of φ . Let A be a closed neighborhood of C that contains no other fixed points. The index of C is defined to be the index of φ over A . The index of C is independent of the choice of the neighborhood, so long as the neighborhood does not contain any other fixed points.

Given a game G in polytope-form, let $\varphi : P \rightrightarrows P$ be the best-reply correspondence, i.e., $\varphi(p)$ is the set of $q \in P$ such that for each n , q_n is a best reply to p . φ is a well-behaved correspondence. The game G has finitely many components of Nash equilibria, obtained as the set of fixed points of φ . For each component C , we can then assign an index. The index of Nash equilibria can also be computed using any Nash map, i.e., a function f that is jointly continuous in payoffs and strategies and whose fixed points for any game are the Nash equilibria (cf. [3]). The sum of indices of each equilibrium component of any game is $+1$. (cf. [13]).

2.2. Statement of the Theorem. Fix a finite normal-form game G . Let C_1, \dots, C_k be the components of its equilibria. Let the index of each C_i be c_i .

Theorem 2.6. *For each $\varepsilon > 0$, there exist pairwise disjoint sets $V_1, \dots, V_k \subseteq \Sigma$, V_i a neighborhood of C_i for each i , such that for any choice of finitely many distinct points $\{\sigma^{ij}\}_{ij}, \sigma^{ij} \in V_i \setminus \partial V_i$ and*

numbers $r_{ij} \in \{-1, 1\}$ such that $\sum_j r_{ij} = c_i$, there exists a normal-form game \hat{G} that is equivalent to G and a ε -perturbation \hat{G}^ε of \hat{G} such that:

- (1) The set of equilibria of \hat{G}^ε equals $\{\hat{\sigma}^{ij}\}_{ij}$, where $\hat{\sigma}^{ij}$ projects to σ^{ij} , for each i, j .
- (2) The index of $\hat{\sigma}^{ij}$ equals r_{ij} .

The following corollary follows immediately from Theorem 2.6 and distinguishes components with a positive index from the others. Below, we say that a game is *generic* if it possesses finitely many equilibria, each with index $+1$ or -1 . It is known that the set of payoffs of a game for which this property is violated has a strictly lower dimension than the whole payoff space (cf. [19]).


Corollary 2.7. *A set $C \subseteq \Sigma$ of strategy profiles is a component of equilibria of G with a positive index iff it is a minimal set (under set inclusion) with respect to the following property: for each $\varepsilon > 0$, there exists $\delta > 0$ such that for each equivalent game \bar{G} and each generic \bar{G}^δ that is an δ -perturbation of \bar{G} , there is a $+1$ equilibrium whose projection to Σ is within ε of C .*

Corollary 2.7 shows that $+1$ equilibria always arise in generic perturbations of payoffs of a finite game, locally around a positive index component. In contrast, when negative index components are considered, there are always generic payoff perturbations for which all the equilibria arising locally have -1 index. The relevance of this result is that $+1$ equilibria possess a number of interesting properties their -1 counterparts do not: [8] shows that the $+1$ equilibria of a generic game are the only ones that are dynamically stable for at least one Nash dynamics. [12] argues for the selection of $+1$ index equilibria, on both theoretical as well as experimental grounds. [9] uses the $+1$ index selection criterion to refine equilibria in signaling games, and in [6], we show that they are the only equilibria that can be made unique in a larger game obtained by adding inferior replies to them.

3. A GUIDE TO THE PROOF OF THEOREM 2.6

This section provides a rough sketch of the proof of Theorem 2.6, which is presented formally in the next section. The main objective here is to highlight the key ideas involved in modifying O'Neill's theorem to construct a map that is game-theoretically meaningful, i.e., whose fixed points are Nash equilibria of a finite game.

We work with the best-reply correspondence, rather than a Nash map. The main reason is that for any approximation of the best-reply correspondence by a continuous function f , we have immediate information on the degree of suboptimality of $f(\sigma)$ against σ . More precisely, fix $\varepsilon > 0$ as in the statement of Theorem 2.6. Let BR_ε be the ε -best reply correspondence, i.e., $\text{BR}_\varepsilon(\sigma)$ is the set of all profiles τ such that for n , τ_n yields a payoff against σ that is strictly within ε of the best payoff achievable against σ . The graph of the ε -best-reply correspondence, denoted $\text{Graph}(\text{BR}_\varepsilon)$, is a neighborhood of the graph of BR . For any continuous function f whose graph is contained in $\text{Graph}(\text{BR}_\varepsilon)$, $f(\sigma)$ is an ε -best-reply to σ .

FIGURE 7. The displacement of point x under the extended map d' and map f' 

The sets V_i identified in the statement of the theorem are neighborhoods of the components that are pairwise disjoint and such that for any σ that belongs to some V_i , $(\sigma, \sigma) \in \text{Graph}(\text{BR}_\varepsilon)$. Now we fix points σ^{ij} as in Theorem 2.6. Our first task, carried out in Step 2, is to obtain a version of O'Neill's result by constructing an approximation of BR that has the σ^{ij} 's as its fixed points and with the assigned indices. To see why we cannot directly appeal to O'Neill's result, let us recapitulate what he does. O'Neill considers the case of a map $f : A \rightarrow X$ where the set of fixed points is a connected set C and A is a Euclidean neighborhood of C . Given finitely many points x_1, \dots, x_k and integers r_i adding up to the index of C , O'Neill constructs a nearby map f' with exactly the x_i 's as the fixed points and with the given indices r_i , by first perturbing the displacement map d of f to a map d' that is close by, and then he constructs f' as the map $x - d'(x)$. Thus, compared to our set up, there are two principal differences. Firstly, O'Neill is working with a function and not a correspondence. Secondly, and more importantly, it is crucial to his proof that A be a Euclidean neighborhood of C , as the following example illustrates, and this is a property that we cannot assume.

Example 3.1. Consider a map f from a 2-simplex to itself. The displacement map of f is defined as $d(x) = x - f(x)$. Figure 7 shows the case where a component of fixed points in red lies in the interior of the simplex, inside the Euclidean neighborhood V in pale blue. As in O'Neill's proof, we would like to obtain a map f' close to f , without any fixed points in V and such that it is equal to f outside of V . For that, the Hopf Extension Theorem is applied to the *displacement map* d in order to obtain a new map d' which is equal to d outside of the neighborhood, new and without fixed points inside of the neighborhood. After obtaining d' , the map defined by $x - d'(x)$ cannot be guaranteed to be inside the 2-simplex. However, it is possible, with the help of Urysohn's Lemma, to scale $d'(x)$ by a strictly positive function $\lambda(\cdot)$ so as to have $f'(x) = x - \lambda(x)d'(x)$ inside the simplex and close to f , as Figure 7 indicates.

However, if the component is in the boundary of the simplex like in Figure 8, the reasoning employed for the interior case does not apply, i.e., O'Neill's construction does not hold. As the

FIGURE 8. A component of fixed points in the boundary





figure suggests, in the case x is mapped to the point $x - d'(x)$ as indicated, no small positive scaling of $d'(x)$ by $\lambda > 0$ guarantees that $x - \lambda d'(x)$ is inside of the simplex.

In Step 2, we approximate BR by a function \hat{f} whose graph is in $\text{Graph}(\text{BR}_\varepsilon)$ and whose fixed points are in $\cup_i V_i \setminus (\partial\Sigma \cup \partial V_i)$. Now, we can proceed as in O'Neill to modify the function inside the V_i 's to obtain a function f whose graph is still in $\text{Graph}(\text{BR}_\varepsilon)$ but whose fixed points are the σ^{ij} 's with index r^{ij} .

Since we want to use f to mimic a best-reply correspondence, the analysis is made easier if the function f is locally affine around the σ^{ij} 's. This yields us the structure of a polymatrix game (one where the payoffs of a player are linear in the strategies of his opponents) at least locally around the equilibria of the game \hat{G}^ε , specified in the theorem. The job of Step 1 is to ensure that it is possible to make f locally affine. Local affineness, also, makes the task of constructing fixed points of f with indices $+1$ or -1 simple, as one can define f to be an affine homeomorphism locally with the corresponding σ^{ij} 's as fixed points.

We would like to view f as a proxy for a best-reply correspondence of a perturbed game, as that would indicate how to perturb payoffs. There are two hurdles to doing so, which we deal with sequentially in our proof. The first problem is that the function f maps into $\Sigma \setminus \partial\Sigma$ and therefore if $f_n(\sigma)$ is a “best-reply” to σ_{-n} , then every pure strategy of player n is also a best reply to σ_{-n} . The way around this problem is to say that $f(\sigma)$ is equivalent to a strategy in an equivalent game \tilde{G} , thus avoiding making everything a best-reply. Constructing a finite equivalent game, then, requires us to introduce finitely many points in each Σ_n as duplicate pure strategies and making $f(\sigma)$ mixtures over these points. Geometrically speaking, this leads us to a triangulation \mathcal{T}_n of Σ_n as carried out in Step 3. For each σ , player n 's “best replies” to σ are among the vertices of the simplex that contains $f_n(\sigma)$ in its interior. As Example 3.2 below demonstrates, resolving the first problem yields a new strategy space, call it $\tilde{\Sigma}_n$, for each player n , in which each pure strategy is labeled at the vertex of the triangulation.

FIGURE 9. Using f to construct a best-reply proxy in an equivalent game

Example 3.2. Figure 9 illustrates the procedure we employ to use f as a proxy for a best-reply correspondence in an equivalent game. The 2-simplex ABC (see (a)) represents the mixed strategy set of player 1 in a finite game. Note that $f_1(\sigma) \neq \sigma_1$, so $\sigma = (\sigma_1, \sigma_2)$ is not a fixed point of f . We introduce a duplicate strategy G in the mixed strategy set and from that produce a triangulation. In this triangulation, the strategy σ_1 has ABG as a carrier, and the strategy $f_1(\sigma)$ has ACG as a carrier. In (d), we declare each vertex of the triangulation as a pure strategy in an equivalent strategy set, that is, a tetrahedron (recall the construction of the equivalent game illustrated in Figure 6). Viewing $f_1(\sigma)$ and σ_1 as points in the tetrahedron, each of these strategies have distinct supports. If a perturbation of payoffs in the equivalent game can be made so that the support of $f_1(\sigma)$ can be made the set of pure best-replies to σ_1 (which is what is effectively meant by “proxy”), then σ_1 cannot be an equilibrium in this perturbed equivalent game, since its support contains sub-optimal replies.

The second problem is that, unlike a best-reply correspondence, the coordinate function f_n of player n could depend on his strategy, σ_n .³ In any game, for any profile σ of mixed strategies, the value of the best-reply correspondence of player n does not depend on player n ’s strategy σ_n . When generating the proxy for a best reply correspondence from the map f , recall that this map was obtained by approximating the best-reply correspondence of the original game by a continuous

³One could say that the problem is more acute around the σ^{ij} ’s where f_n is exclusively a function of σ_n !

map and then modifying it through the use of the Hopf Extension Theorem. The resulting map f unfortunately cannot be guaranteed to inherit that property of the best-reply correspondence. This impairs the possibility of converting $f_n(\cdot)$ into a best-reply correspondence of an equivalent strategy set of player n along the lines we illustrated in Example 3.2.

To deal with this problem, we consider an equivalent polytope-form game \tilde{G} in which each player n chooses a strategy in $\tilde{\Sigma}_n$ and one in $\tilde{\Sigma}_{n+1}$ (where $n+1 = 1$ if $n = N$), and the strategy space of player n is $\tilde{\Sigma}_n \times \tilde{\Sigma}_{n+1}$. The set $\tilde{\Sigma}_n \times \tilde{\Sigma}_{n+1}$ is a (non-simplicial) polytope. In order to define payoffs, we simply consider the first coordinates of each player and use the payoffs of the original game G : denote by $(\tilde{\sigma}_{n,0}, \tilde{\sigma}_{n,1}) \in \tilde{\Sigma}_n \times \tilde{\Sigma}_{n+1}$ a general strategy of player n in \tilde{G} . The payoff from a profile $(\tilde{\sigma}_{n,0}, \tilde{\sigma}_{n,1})_{n \in \mathcal{N}}$ in \tilde{G} to player n is defined as $\tilde{G}_n((\tilde{\sigma}_{n,0}, \tilde{\sigma}_{n,1})_{n \in \mathcal{N}}) = G_n((\tilde{\sigma}_{n,0})_{n \in \mathcal{N}})$. The construction is such that the choice by n of a strategy in $\tilde{\Sigma}_{n+1}$ is payoff-irrelevant for player n .

Now the function f can be “lifted” to this strategy space by making f_n depend on σ_{-n} and player $n-1$ ’s choice in Σ_n (where $n-1 = N$ if $n = 1$): since player n ’s coordinate function f_n might depend on Σ_n , we will use player $n-1$ ’s copy of those coordinates. We define a new map \tilde{f}_n over the cartesian product of strategy sets of the polytope-form game (with the exception of player n ’s) to $\tilde{\Sigma}_n \times \tilde{\Sigma}_{n+1}$ by $\tilde{f}_n((\tilde{\sigma}_{m,0}, \tilde{\sigma}_{m,1})_{m \neq n}) = f_n(\tilde{\sigma}_{-n,0}, \tilde{\sigma}_{n-1,1}) \times \tilde{\sigma}_{n+1,0}$.⁴

Combining this idea along with the one illustrated in Example 3.2, we surpass the two main hurdles to obtaining a function that can be used as a proxy for a best-reply correspondence: steps 1 to 3 are preparatory steps to ensure the existence of a game equivalent to the original one, where a proxy for a best-reply correspondence can be constructed so as to satisfy the requirements of our main result (with precisely the σ^{ij} ’s as fixed points and with the correct indices). Step 4 presents this equivalent game explicitly.

The remaining steps of the proof show the precise sense in which \tilde{f} is a proxy for a perturbed best-reply correspondence: they are responsible for constructing the payoff perturbation of the equivalent game whose best-reply correspondence will essentially match that proxy.

This construction has two main phases: in the first one conducted in steps 5 and 6, for each $\tilde{\sigma}$, we define a payoff perturbation of the polytope-form game \tilde{G} (constructed in step 4) such that for each n the best replies to $\tilde{\sigma}$ are among the vertices of the carrier of $\tilde{f}_n(\tilde{\sigma})$. The perturbation comes in the form of a bonus $\tilde{g}_n^0(\tilde{\sigma})$ for the first coordinate of n ’s strategy and $\tilde{g}_n^1(\tilde{\sigma})$ for the second coordinate. For each $\tilde{\sigma}$, the finite, perturbed game is denoted by $\tilde{G} \oplus \tilde{g}^0(\tilde{\sigma}) \oplus \tilde{g}^1(\tilde{\sigma})$. Step 5 details this construction. Step 6 analyses the relevant properties of the “best-reply” correspondence that assigns to each $\tilde{\sigma}$ the set of best replies to it in the perturbed game $\tilde{G} \oplus \tilde{g}^0(\tilde{\sigma}) \oplus \tilde{g}^1(\tilde{\sigma})$: note that the payoff perturbations in Step 5 depend on the strategy profile $\tilde{\sigma}$ and therefore may be different depending on which profile is played. Hence, though the correspondence that assigns to each $\tilde{\sigma}$ the

⁴The precise definition of this map is slightly more nuanced than what is discussed here, but the main idea for its definition is to avoid the dependency on player n ’s choices.

set of best replies to it in $\tilde{G} \oplus \tilde{g}^0(\tilde{\sigma}) \oplus \tilde{g}^1(\tilde{\sigma})$ matches our proxy \tilde{f} , the correspondence is not (yet) the best-reply correspondence of a *finite* game.

The final phase of the construction uses the maps \tilde{g}^0 and \tilde{g}^1 to finally construct a finite game with the required properties. This phase is the most technical part of the proof and providing a thorough intuition for it cannot be properly done without introducing an inordinate amount of notations and definitions. In spite of that, we would like to provide some insight into the final part of the argument.

The perturbation $\tilde{g}^0(\cdot)$ maps $\tilde{g}^1(\cdot)$ are simply continuous maps defined over the strategy space of the polytope-form game \tilde{G} , namely, $\prod_{n \in \mathcal{N}} (\tilde{\Sigma}_n \times \tilde{\Sigma}_{n+1})$. If maps $\tilde{g}^0(\cdot)$ and $\tilde{g}^1(\cdot)$ were in addition multilinear maps of the strategy profiles of this game, then they would immediately define payoff perturbations allowing for the definition of a perturbed polytope-form game. Step 7 provides the preliminary work for generating approximations of the maps $\tilde{g}^0(\cdot)$ and $\tilde{g}^1(\cdot)$ so as to be able to define the desired polytope form. By triangulating the strategy sets of each player with sufficiently small diameter, we approximate the maps $\tilde{g}^0(\cdot)$ and $\tilde{g}^1(\cdot)$ by maps which are locally multilinear in each product of simplices of the triangulation. Each newly introduced vertex of the triangulation is declared to be a new strategy of an equivalent game \hat{G} to \tilde{G} —in the same fashion as exemplified in Figure 9—and this game can now be perturbed so as to achieve the required properties.

Step 8 defines the perturbation of the game \hat{G} , which has five components. The first two (\hat{G}^0 and \hat{G}^1) correspond to the perturbations that are induced by the multiaffine versions of the maps \tilde{g}^0 and \tilde{g}^1 . Two other perturbations (\hat{g}^0 and \hat{g}^1) are required to ensure that optimal strategies only mix over points that are vertices of a cell of the triangulation (this is illustrated in more detail in Example 4.18). The fifth and final component of payoff perturbations is required to ensure uniqueness of the equilibria (projecting to σ^{ij} 's). As explained, passing to equivalent games involves introducing duplicate strategies, which may result in sets of mixed strategies that are equivalent to a point σ^{ij} in the original game. One therefore needs to select one among the many profiles of strategies included in the equivalent set to be the isolated equilibrium of the equivalent game - and the payoff perturbations are constructed to achieve this goal. Step 9 completes the proof of the main theorem by showing that this perturbed game is the finite normal-form game we were after.

Since the proof involves a lot of notation, some of which are step-specific and others not, we end this section with a table of the notations that are used across multiple steps. The table should be used as a consultation tool throughout the reading of the proof, in case some notation that has already been encountered is forgotten.

TABLE 1. Table of Notations

Notation	Reference and/or Definition
Y_n^{ij}	Simplex contained in Σ_n , with σ_n^{ij} as barycenter (see Step 1)
X_n^{ij}	Simplex contained in Y_n^{ij} , with σ_n^{ij} as barycenter (see Step 1)
f_n^{ij}	$f_n^{ij} : X_n^{ij} \rightarrow Y_n^{ij}$; see Lemma 4.1
BR	Best-reply correspondence of game G
BR_ε	ε -Best-reply correspondence of game G (see Step 1)
f	$f : \Sigma \rightarrow \Sigma$ (See Lemma 4.3, Step 2)
\mathcal{T}_n	Triangulation of Σ_n (see Lemma 4.7, Step 3)
T_n^{ij}	Simplex of \mathcal{T}_n contained in X_n^{ij} , containing σ_n^{ij} as its barycenter (see Lemma 4.7, Step 3)
$\tilde{\Sigma}_{n,0}$	Vertices of the triangulation \mathcal{T}_n (see Step 4)
$\tilde{\Sigma}_{n,1}$	Vertices of the triangulation \mathcal{T}_{n+1} (see Step 4)
$\tilde{\Sigma}_{n,k}$	$\tilde{\Sigma}_{n,k} \equiv \Delta(\tilde{\Sigma}_{n,k}), k = 0, 1$ (See Step 4)
\tilde{G}	Polytope-form game where the strategy set of each player n is $\tilde{\Sigma}_n \equiv \tilde{\Sigma}_{n,0} \times \tilde{\Sigma}_{n,1}$ (see Step 4)
\tilde{BR}_ε	ε -Best-reply correspondence of game \tilde{G} (see Step 4)
$\tilde{\beta}_{n,k}$	$\tilde{\beta}_{n,k} : \Sigma_n \rightarrow \tilde{\Sigma}_{n,k}$ is the barycentric map: it associates to a point σ_n its barycentric coordinates in the triangulation \mathcal{T}_n , viewed as a face of $\tilde{\Sigma}_{n,k}, k = 0, 1$ (see Step 4)
$\phi_{n,k}$	$\phi_{n,k} : \tilde{\Sigma}_{n,k} \rightarrow \Sigma_{n+k}$ projects the point $\tilde{\sigma}_{n,k}$ to its location in $\Sigma_{n+k}, k = 0, 1$ (see Step 4)
$\tilde{\Sigma}_{n,k}^*$	$\tilde{\Sigma}_{n,k}^* \equiv \tilde{\beta}_{n,k}(\Sigma_{n+k}), k = 0, 1$ (see step 4)
$\mathcal{S}_{n,k}^*$	Subcomplex of the face complex of $\Delta(\tilde{\Sigma}_{n,k})$ whose space is $\tilde{\Sigma}_{n,k}^*, k = 0, 1$ (see Step 4)
$\tilde{T}_{n,k}^{ij}$	$\tilde{T}_{n,k}^{ij} \equiv \tilde{\beta}_{n,k}((f_{n+k}^{ij})^{-1}(T_{n+k}^{ij})), k = 0, 1$ (see Step 4)

Notation	Reference and/or Definition
$\tilde{f}_{n,0}$	$\tilde{f}_{n,0} : \tilde{\Sigma}_{-n} \rightarrow \tilde{\Sigma}_{n,0},$ $\tilde{f}_{n,0}(\tilde{\sigma}_{-n}) \equiv \tilde{\beta}_{n,0}(f_n(\phi_{-n,0}(\tilde{\sigma}_{-n,0}), \phi_{n-1,1}(\tilde{\sigma}_{n-1})))$ (see Step 4 and Lemma 4.11)
$\tilde{f}_{n,1}$	$\tilde{f}_{n,1}(\tilde{\sigma}_{n+1}) \equiv \tilde{\beta}_{n,1}(\phi_{n+1,0}(\tilde{\sigma}_{n+1}))$ (see Step 4 and Lemma 4.11)
$\tilde{g}^0(\cdot)$	$\tilde{g}^0 : \tilde{\Sigma} \rightarrow \prod_n \mathbb{R}^{\tilde{S}_{n,0}}$ (see Step 5)
$\tilde{g}^1(\cdot)$	$\tilde{g}^1 : \tilde{\Sigma} \rightarrow \prod_n \mathbb{R}^{\tilde{S}_{n,1}}$ (see Step 5)
$r_n(\cdot)$	$r_n : \Sigma_{-n} \rightarrow \mathbb{R}$ returns the best payoff player n obtains against σ_{-n} (see Step 5)
$\tilde{\mathcal{T}}_n$	Triangulation of $\Delta(\tilde{S}_{n,0}) = \Delta(\tilde{S}_{n-1,1})$ (see Step 7)
$\hat{S}_{n,0}$	$\hat{S}_{n,0}$ the set of vertices of $\tilde{\mathcal{T}}_n$, $\hat{S}_{n,0} = \hat{S}_{n-1,1}$ (see Step 8)
\hat{T}_n^{ij}	Carrier of $\tilde{\sigma}_n^{ij}$ in $\tilde{\mathcal{T}}_n$ (see Step 7 and Lemma 4.15)
$\hat{\phi}_{n,k}$	$\hat{\phi}_{n,k} : \hat{\Sigma}_{n,k} \rightarrow \tilde{\Sigma}_{n,k}$ projects strategies from $\hat{\Sigma}_{n,k}$ to $\tilde{\Sigma}_{n,k}$ (see Step 8)
$\hat{\mu}_{n,k}$	$\hat{\mu}_{n,k} : \hat{\Sigma}_n \rightarrow \hat{\Sigma}_{n,k}$ computes the marginal of $\hat{\sigma}_n$ over $\Sigma_{n,k}$ (see Step 8)
\hat{G}	Normal-form game with pure strategy set of player n equals to $\hat{S}_n \equiv \hat{S}_{n,0} \times \hat{S}_{n,1}$ and payoff $\hat{G}_n(\hat{\sigma}) = \hat{G}_n(\hat{\phi} \circ \hat{\mu}(\hat{\sigma}))$ (see Step 8)

4. PROOF OF THEOREM 2.6

As in the statement of the theorem, we have the following notation, used uniformly throughout the section. For each component C_i , $i = 1, \dots, k$, we have a finite (possibly empty) set J_i , and a finite set of integers $r^{ij} \in \{+1, -1\}$, $j \in J_i$, such that $\sum_{j \in J_i} r^{ij} = c_i$, with the sum being zero if J_i is an empty set.

Step 1. Fix a finite number of completely-mixed strategy profiles σ^{ij} , $i = 1, \dots, k$ and $j \in J_i$ such that for each n the σ_n^{ij} 's are distinct mixed strategies. Locally around each σ^{ij} , we construct an affine map that has σ^{ij} as its unique fixed point and assigns it the index r^{ij} . In Step 2, these maps are then extended to the whole of Σ (not affinely though) without introducing additional fixed points.

For each n , and i, j , let X_n^{ij} and Y_n^{ij} be two full-dimensional simplices in Σ_n with σ_n^{ij} as their barycenters and $X_n^{ij} \subseteq \text{int}(Y_n^{ij})$. Let $X^{ij} = \prod_n X_n^{ij}$ and $Y^{ij} = \prod_n Y_n^{ij}$. We will assume that the simplices Y_n^{ij} are small enough that the Y^{ij} 's are pairwise disjoint. For now, these are the only

restrictions we choose for Y_n^{ij} and X_n^{ij} in order to obtain Lemma 4.1. In Step 2, we will specify their choice further.

Lemma 4.1. *For each i, j and n , there is an affine homeomorphism $f_n^{ij} : X_n^{ij} \rightarrow Y_n^{ij}$ such that:*

- (1) σ_n^{ij} is the unique fixed point of f_n^{ij} ;
- (2) letting $f^{ij} = \prod_n f_n^{ij} : X^{ij} \rightarrow Y^{ij}$, the index of σ^{ij} under f^{ij} is r^{ij} ;
- (3) for all $\sigma_n \neq \sigma_n^{ij}$, $\|f_n^{ij}(\sigma_n) - \sigma_n^{ij}\| > 0$.

Proof. The result for the case $N = 1$ was proved in [5]. Use their result for each n , assigning index r^{ij} to player 1's map f_1^{ij} and +1 for the maps of all other players. The index of σ^{ij} under f^{ij} is r^{ij} by the multiplication property of the index. \square


Step 2. This step gives us a version of O'Neill's theorem with three main differences. (1) It applies to the BR correspondence. (2) The points chosen are in the interior of Σ and close to, but not necessarily in, C_i . (3) We insist that the map be locally affine around the fixed points.

Let $\text{Graph}(\text{BR})$ be the graph of the best-reply correspondence $\text{BR} : \Sigma \rightrightarrows \Sigma$ of the game G . For each $\varepsilon > 0$, let BR_ε be the ε -best-reply correspondence, i.e., for each σ , $\text{BR}_\varepsilon(\sigma)$ is the set of strategy profiles τ such that for each n , $G_n(\sigma_{-n}, \tau_n) > \max_{s_n} G_n(\sigma_{-n}, s_n) - \varepsilon$. The graph of BR_ε , denoted $\text{Graph}(\text{BR}_\varepsilon)$, is a neighborhood of $\text{Graph}(\text{BR})$. (In fact, the sets $\text{Graph}(\text{BR}_\varepsilon)$ form a basis of neighborhoods of $\text{Graph}(\text{BR})$.) From now on, we fix $\varepsilon > 0$ as required by the main theorem.

Choose $\delta > 0$ such that the δ -neighborhoods V_i of C_i satisfy the following properties: $V_i \cap V_j = \emptyset$ for all $i \neq j$; and for all i and $\sigma \in V_i$, (σ, σ) belongs to $\text{Graph}(\text{BR}_\varepsilon)$. These V_i 's are the ones specified in the statement of Theorem 2.6. These, too, are fixed from now on. Pick points σ^{ij} , $i = 1, \dots, k$, $j \in J_i$ as required by the same theorem and fix them from now on as well. If necessary by considering an equivalent game \bar{G} obtained by adding duplicate strategies, we can assume that the σ_n^{ij} 's are all distinct for each n and with disjoint supports. Indeed, consider an equivalent game \bar{G} where we duplicate every pure strategy of every player. Let $\bar{\Sigma}$ be the strategy space in \bar{G} and $\phi : \bar{\Sigma} \rightarrow \Sigma$ be the map that sends strategies in \bar{G} to equivalent profiles in G . Let $\bar{C}_i = \phi^{-1}(C_i)$ be the equilibrium component of \bar{G} for each i , and let $\bar{V}_i \equiv \phi^{-1}(V_i)$ be the corresponding neighborhood of \bar{C}_i . For each σ^{ij} , we can pick a point $\bar{\sigma}^{ij} \in \phi^{-1}(\sigma^{ij})$ such that for each n , the $\bar{\sigma}_n^{ij}$'s are distinct mixed strategies with disjoint supports. Because \bar{G} is equivalent to G , each \bar{C}_i has the same index as C_i (cf. Theorem 5 in [4]). Therefore, in what follows, we can replace G with \bar{G} . To simplify notation, we will refer to the game still as G and just assume that the chosen σ^{ij} 's involve distinct mixed strategies for each player. We illustrate this procedure in Example 4.2.

Example 4.2. In Figure 10, (a) represents a 1-simplex with vertices A and B where we singled-out point P , representing a mixed strategy. By duplicating the vertex A to A' and B to B' , we represent in (b) a tetrahedron, where the points P_1 and P_2 are singled-out duplicates of point P , with disjoint supports.

FIGURE 10. Duplicating pure strategies and choosing points with distinct supports



Lemma 4.3. *There exists a continuous function $f : \Sigma \rightarrow \Sigma$ with the following properties:*

- (1) *the graph of f is contained in $\text{Graph}(BR_\epsilon)$;*
- (2) *the fixed points of f are the points σ^{ij} and the index of each σ^{ij} is r^{ij} ;*
- (3) *for each n and i, j , there exist a polyhedral subdivision \mathcal{X}_n of Σ_n and full-dimensional simplices X_n^{ij} and Y_n^{ij} such that:*
 - (a) σ_n^{ij} *is the barycenter of both X_n^{ij} and Y_n^{ij} ;*
 - (b) $X_n^{ij} \subset \text{int}(Y_n^{ij}) \subset \Sigma_n \setminus \partial\Sigma_n$;
 - (c) $\prod_n Y_n^{ij} \times \prod_n Y_n^{ij} \subset \text{Graph}(BR_\epsilon)$;
 - (d) *the restriction f^{ij} of f to $\prod_n X_n^{ij}$ is a product map $\prod_n f_n^{ij}$ where f_n^{ij} is an affine homeomorphism of X_n^{ij} onto Y_n^{ij} ;*
 - (e) X_n^{ij} *is the space of a subcomplex of \mathcal{X}_n and the carrier of σ_n^{ij} in \mathcal{X}_n has dimension $\dim(\Sigma_n)$.*

Remark 4.4. As it will be clear from the proof of the lemma and illustrated in Example 4.6, the polyhedral subdivision \mathcal{X}_n is constructed by “extending” the faces of each X_n^{ij} , thus producing a polyhedral subdivision of Σ_n (see Figure 12). Point (e) of the Lemma implies that X_n^{ij} is a union of cells of the polyhedral subdivision. The simplices Y_n^{ij} play no role in the construction of \mathcal{X}_n and are neither a cell nor the space of a subcomplex of \mathcal{X}_n necessarily.

Proof. By Corollary 1 and Axiom 2 in [11], there exists a function $\hat{f} : \Sigma \rightarrow \Sigma$ such that

- (1) *the graph of \hat{f} is contained in $\text{Graph}(BR_\epsilon)$;*
- (2) *all its fixed points are contained in $\cup_i (V_i \setminus \partial_\Sigma V_i)$;*
- (3) *the index of \hat{f} over V_i is c_i for each i .*

If necessary by replacing \hat{f} with a suitable convex combination $(1 - \alpha)\hat{f} + \alpha\sigma^\circ$, where σ° is a completely mixed profile in Σ and $\alpha > 0$ is sufficiently small, we can assume furthermore that the fixed points of \hat{f} are contained in $\Sigma \setminus \partial\Sigma$.

Let $C_i(\hat{f})$ be the set of fixed points of \hat{f} in V_i . As V_i is the δ -neighborhood of the component C_i , it is semialgebraic (cf. Proposition 2.2.8 in [1]) and $V_i \setminus (\partial_{\Sigma_i} V_i \cup \partial\Sigma)$ is connected. Therefore for all $\eta > 0$ sufficiently small, the set of points in V_i whose distance from $\partial_{\Sigma_i} V_i \cup \partial\Sigma$ is strictly greater than η is a connected set. As the points σ^{ij} and the set $C_i(\hat{f})$ lie in $V_i \setminus (\partial_{\Sigma_i} V_i \cup \partial\Sigma)$, we can choose $\eta > 0$ small enough such that for each i , letting U_i be the set of $\sigma \in V_i$ such that $d(\sigma, \partial_{\Sigma_i} V_i \cup \partial\Sigma) > 2\eta$, we have that $U_i \setminus \partial U_i$ is connected and contains $C_i(\hat{f})$ as well as the points σ^{ij} . Furthermore, by choosing an even smaller η , if necessary, we can assume that $(\sigma, \tau) \in \text{Graph}(\text{BR}_\varepsilon)$ for $\sigma \in V_i$ and τ within 2η of σ .

Use Urysohn's lemma to construct a function $\gamma : \Sigma \rightarrow [\eta, 1]$ that equals 1 outside the interiors of V_i 's and η on U_i . Replace the function \hat{f} with the function $\hat{f}^1 \equiv \gamma\hat{f} + (1 - \gamma)\text{Id}$. The graph of \hat{f}^1 is also contained in $\text{Graph}(\text{BR}_\varepsilon)$ and its fixed points are still the points in the $C_i(\hat{f})$'s. Also, $\|\hat{f}^1(\sigma) - \sigma\| \leq 2\eta$ if $\sigma \in U_i$ for some i .

For each i, j and n choose simplices X_n^{ij} and Y_n^{ij} such that:

- (4) X_n^{ij} is contained in the interior of Y_n^{ij} , and Y_n^{ij} is contained in the interior of U_i ;
- (5) σ_n^{ij} is the barycenter of both X_n^{ij} and Y_n^{ij} ;
- (6) Y_n^{ij} is contained in the η -radius ball around σ^{ij} ;
- (7) $\{Y_n^{ij}\}_{i,j}$ are pairwise disjoint for each n .

Each dimension $\dim(\Sigma_n) - 1$ face F_n^{ij} of X_n^{ij} is contained in a unique hyperplane $H_{F_n^{ij}}$ in \mathbb{R}^{S_n} orthogonal to Σ_n . Denote by $H_{F_n^{ij}}^*$, $*$ $\in \{+, -\}$, the positive and negative half-spaces defined by this hyperplane. For a fixed X_n^{ij} , if a hyperplane $H_{F_n^{ij}}$ intersects $\sigma_n^{\hat{i}\hat{j}}$ for $\hat{i} \neq i$ or $\hat{j} \neq j$, then the vertices of F_n^{ij} are not in general position in Σ_n , i.e., they belong to a subspace of Σ_n of dimension strictly lower than $\dim(\Sigma_n)$. Therefore, choosing the vertices of X_n^{ij} conveniently in the interior of Y_n^{ij} , we can assume that $H_{F_n^{ij}}$ does not intersect $\sigma_n^{\hat{i}\hat{j}}$ for $\hat{i} \neq i$ or $\hat{j} \neq j$. We assume this holds for each i, j, n . We can now define \mathcal{X}_n . Consider the intersection of the hyperplane $H_{F_n^{ij}}$ with Σ_n , for each i, j . This defines a polyhedral subdivision where a cell of the subdivision is given by the intersection of finitely many half-spaces $H_{F_n^{ij}}^*$.⁵ Due to our assumption, the carrier of σ_n^{ij} in this subdivision is a cell of dimension $\dim(\Sigma_n)$. By construction of this subdivision, X_n^{ij} is the space of a subcomplex of \mathcal{X}_n . This proves 3(e).

Use Step 1 to construct for each n and i, j an affine homeomorphism $f_n^{ij} : X_n^{ij} \rightarrow Y_n^{ij}$ satisfying property (1) of Lemma 4.3 and such that $f^{ij} = \prod_n f_n^{ij}$ satisfies property (2) of Lemma 4.3. Letting

⁵In the Appendix, we elaborate on polyhedral subdivisions constructed through this procedure.

A be the affine space generated by $\Sigma - \Sigma$, define $d : (\Sigma \setminus \cup_i (U_i \setminus \partial U_i)) \cup (\cup_{i,j} X^{ij}) \rightarrow A$ by $d(\sigma) = \sigma - \hat{f}^1(\sigma)$ if $\sigma \notin \cup_i (U_i \setminus \partial U_i)$; $d(\sigma) = \sigma - f^{ij}(\sigma)$ if $\sigma \in X^{ij}$. The only zeros of d are the σ^{ij} 's. The set U_i is semi-algebraic and $U_i \setminus \partial U_i$ is an open and connected set. Therefore, $(U_i, \partial U_i)$ is a pseudomanifold with boundary. Hence, $U_i \setminus \cup_j (X^{ij} \setminus \partial X^{ij})$ is a pseudo-manifold with boundary $(\cup_j \partial X^{ij}) \cup \partial U_i$ and the restriction of d to this boundary has degree zero. As there are at least two players with at least two pure strategies each, the dimension of Σ is at least two. Therefore, by the Hopf Extension Theorem, d extends to the sets U_i in such a way that its norm over $\cup_i U_i$ is still no more than 2η and it has no additional zeros.⁶ Thus, we have a displacement map defined over the whole of Σ , denoted still by d . Define f as $\text{Id} - d$.

Outside $\cup_i U_i$, d is the displacement of \hat{f}^1 , whose graph is contained in $\text{Graph}(\text{BR}_\varepsilon)$; on the U_i 's d has norm 2η or less, which by the choice of η , therefore, means the graph of f is contained in $\text{Graph}(\text{BR}_\varepsilon)$, proving point (1) of the Lemma. Point 3(c) of the lemma follows from the fact that the diameter of Y^{ij} is at most 2η ; all other points of the lemma follow from the construction of the function f . \square


Example 4.5. We explain and illustrate our claim on the degree of d over $(\cup_j \partial X^{ij}) \cup \partial U_i$ being zero. One can apply Theorem 5.1 in [17] to obtain this result: the degree of d over ∂U_i is determined by the index of \hat{f}^1 in U_i and is c_i . The degree of d over ∂X^{ij} is determined by the index of f^{ij} over X^{ij} which is r^{ij} by construction. This degree is measured with respect to the orientation induced on ∂X^{ij} by X^{ij} . But the orientation induced by $U_i \setminus \cup_j (X^{ij} \setminus \partial X^{ij})$ over ∂X^{ij} is opposite to the one with respect to which f^{ij} is measured, by the definition of orientability. According to this orientation, the index of f^{ij} is opposite, hence, $-r^{ij}$. Therefore, since $\sum_j r^{ij} = c_i$, the restriction of d to $(\cup_j \partial X^{ij}) \cup \partial U_i$ has degree zero.

In (a) Figure 11, we illustrate a counterclockwise orientation induced on the boundary of the disc, and the counterclockwise orientations on each of the boundaries of the smaller red discs centered in points P_1 and P_2 . These orientations agree, since the ones on the boundary of the smaller discs are induced (by excision) from the one on the boundary of the blue disc. But when orienting the boundary of the blue disc with the two smaller red discs removed, note that the inner boundaries are now clockwise oriented, by the definition of orientability. So if a map has degree $+2$ measured according to the orientation on the boundary of the blue disc (or according to the sum of the degrees measured according to the orientation on the boundary of the red discs, say, each being $+1$), then, when the red discs are removed, the orientation on the inner boundary is now clockwise (therefore implying the degree of the map, with respect to this orientation, is -1). So the degree

⁶The version of the Hopf Extension Theorem we use here is Corollary 18, Chapter 8 in [22]. We note this Corollary is stated for singular cohomology (with coefficients in \mathbb{Z}), and provides a condition for the extension of $d : (\cup_j \partial X^{ij}) \cup \partial U_i \rightarrow A - \{0\}$ to $U_i \setminus \cup_j (X^{ij} \setminus \partial X^{ij})$. The condition is that $\delta(d^*(s^*)) = 0$, where s^* is a generator of $H^m(A - \{0\})$ ($A - \{0\}$ being a homotopy sphere), where $m = \dim(A) - 1$ and δ is the coboundary cohomology morphism of the long exact sequence of the pair $(U_i \setminus \cup_j (X^{ij} \setminus \partial X^{ij}), (\cup_j \partial X^{ij}) \cup \partial U_i)$. If the degree of $d : (\cup_j \partial X^{ij}) \cup \partial U_i \rightarrow A - \{0\}$ is zero (because it is the displacement of \hat{f}^1), then by definition of the singular cohomology functor, $d^*(s^*) = 0$, which implies $\delta(d^*(s^*)) = 0$.

on the outer-boundary is $+2$, and on the inner-boundaries is -1 each, resulting in the degree of zero in the union of all boundaries.

FIGURE 11. Orientation



Example 4.6. Figure 12 illustrates how the polyhedral subdivision of Lemma 4.3 is obtained. The two points in red in Figure (a) can be viewed as two of the σ_n^{ij} 's, first picked in the interior of the simplex. We omit any reference to Y_n^{ij} 's in this construction because they do not play a role in obtaining the subdivision. In (b) we choose the vertices for each of the X_n^{ij} 's that should have the fixed red points as their barycenters. In (c), the 2-simplices $A1 - B1 - C1$ and $A2 - B2 - C2$ should be viewed as two distinct X_n^{ij} 's of player n . When the vertices $A1, B1, C1$ and $A2, B2, C2$ are well-chosen, for previously fixed points $BC1$ and $BC2$, then they define 2-simplices such that, when their 1-dimensional faces are extended as in the figure, they do not intersect any of the points $BC1$ and $BC2$. The result of the process is a polyhedral subdivision as depicted in (d), where $BC1$ and $BC2$ have 2-dimensional carriers and where simplices $A1 - B1 - C1$ and $A2 - B2 - C2$ are unions of cells of the subdivision, i.e., both are spaces of a subcomplex of the the polyhedral subdivision: in particular, $A1 - B1 - C1$ is a cell of the resulting polyhedral subdivision, whereas $A2 - B2 - C2$ is not.

Step 3. From here on, we fix the map f given in Step 2, along with the multisimplices $X^{ij} \equiv \prod_n X_n^{ij}$ and $Y^{ij} \equiv \prod_n Y_n^{ij}$. We now describe a triangulation \mathcal{T}_n of Σ_n for each n .

Lemma 4.7. *For each n , there exists a triangulation \mathcal{T}_n of Σ_n such that:*

- (1) *For each i, j , there exists a simplex $T_n^{ij} \subset X_n^{ij}$ of \mathcal{T}_n with σ_n^{ij} as its barycenter;*
- (2) *For each i, j , letting $T^{ij} \equiv \prod_n T_n^{ij}$, if $\sigma \notin \cup_{i,j} (T^{ij} \setminus \partial T^{ij})$, there exists n such that for each τ_n that belongs to the carrier of σ_n in \mathcal{T}_n , σ_n does not belong to the carrier of $f_n(\sigma_{-n}, \tau_n)$ in \mathcal{T}_n .*

- (3) $(\sigma, \tau) \in \text{Graph}(\text{BR}_\varepsilon)$ if for each n , τ_n is a vertex of \mathcal{T}_n and $f_n(\sigma)$ belongs to the simplicial neighborhood of the closed star of τ_n ;
- (4) if $(\sigma, \tau) \in \text{Graph}(\text{BR}_\varepsilon)$ where $\sigma \in T^{ij}$ for some ij , then for each $\sigma' \in T^{ij}$ and each n , $G_n(\sigma_{-n}, \tau_n) > \max_{s_n} G_n(\sigma'_{-n}, s_n) - \varepsilon$.

Example 4.8. In Figure 13, we illustrate condition (2) of Lemma 4.7. In the figure, the carrier of τ_n and σ_n is the 2-simplex $A - G - C$, and the carrier of $f_n(\sigma_{-n}, \tau_n)$ is the 2-simplex $A - G - B$.

Proof. There exists $\tilde{\eta} > 0$ such that:

- (5) if $\sigma \notin \cup_{i,j} (X^{ij} \setminus \partial X^{ij})$, then there exists n such that $\|f_n(\sigma_{-n}, \tau_n) - \sigma_n\| > \tilde{\eta}$, for all τ within $\tilde{\eta}$ of σ ;
- (6) $(\sigma, \tau) \in \text{Graph}(\text{BR}_\varepsilon)$, if τ is within $2\tilde{\eta}$ of $f(\sigma)$.

FIGURE 12. An illustration for the polyhedral subdivision of Lemma 4.3





FIGURE 13. Illustration of condition (2) of Lemma 4.7



Let T_n^{ij} be the simplex that has as its vertices $\{(1 - .5\tilde{\eta})\sigma_n^{ij} + .5\tilde{\eta}v \mid v \text{ a vertex of } X_n^{ij}\}$. Taking $\tilde{\eta} > 0$ smaller if necessary, we can assume that T_n^{ij} is contained in the interior of the carrier of σ_n^{ij} in \mathcal{X}_n (cf. Lemma 4.3) and such that

- (7) for each $(\sigma, \tau) \in \text{Graph}(\text{BR}_\varepsilon)$ where $\sigma \in T^{ij}$, it follows that, for all n , $G_n(\sigma, \tau_n) > \max_{s_n} G_n(\sigma', s_n) - \varepsilon$ for all $\sigma' \in T^{ij}$.

By construction, T_n^{ij} is contained in X_n^{ij} for each n and i, j . Take a triangulation \mathcal{T}'_n of Σ_n which refines \mathcal{X}_n such that:

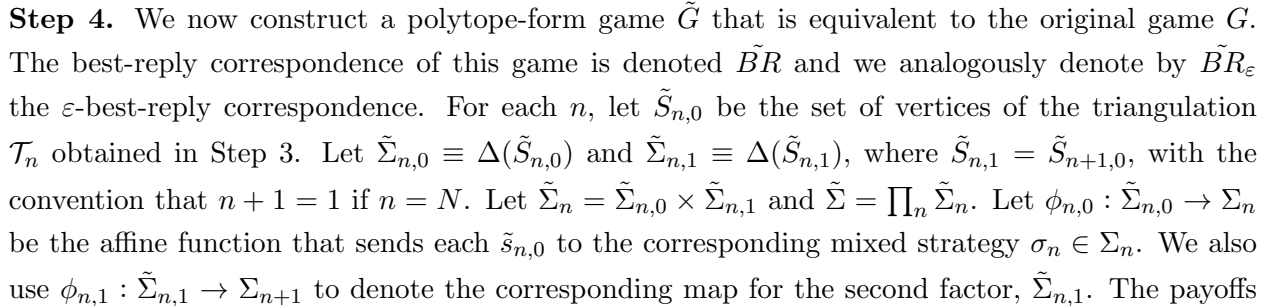
- (8) T_n^{ij} is a simplex of \mathcal{T}'_n for each i, j ;

- (9) The diameter of each simplex in \mathcal{T}'_n is no more than $\tilde{\eta}$.

Then, properties (1) (3) and (4) hold for the triangulation \mathcal{T}'_n (from (7), (8) and (9)); for $\sigma \notin \cup_{i,j}(X^{ij} \setminus \partial X^{ij})$, property (2) also holds for this triangulation because of (4). Since f_n^{ij} is an affine homeomorphism of X_n^{ij} onto Y_n^{ij} , both of which have σ_n^{ij} as their barycenters, we have $f_n^{ij}(\partial T_n^{ij}) \cap T_n^{ij} = \emptyset$. Therefore, there exists a sufficiently fine subdivision \mathcal{T}_n of \mathcal{T}'_n modulo $\cup_{i,j} T_n^{ij}$ —i.e., \mathcal{T}_n subdivides all simplices which are not faces of any T^{ij} (cf. [24])—such that property (2) holds even for $\sigma \in \cup_{i,j}(X^{ij} \setminus (T^{ij} \setminus \partial T^{ij}))$. \square

Example 4.9. The example illustrates how to refine the triangulation shown in Figure 14 modulo the 2-simplex $A1 - C1 - B1$. The point p in the figure is the barycenter of $A1 - B1 - C1$ and q is another point that does not belong to this 2-simplex. One can see then in (b) how we can insert vertices away from the 2-simplex $A1 - B1 - C1$ so as to obtain another triangulation that refines the initial one in (a). The carriers of p and q in the refined triangulation are, in particular, disjoint. Note also that, for the simplices of the refined triangulation that do not share an edge

FIGURE 14. Refining a triangulation modulo a 2-simplex



are now defined as $\tilde{G}_n(\tilde{\sigma}) \equiv G_n((\phi_{n,0}(\tilde{\sigma}_{n,0})_{n \in \mathcal{N}})$. In particular, the choice of the strategy $\tilde{\sigma}_{n,1}$ of any player n is payoff-irrelevant. Therefore, the game \tilde{G} is equivalent to G .

Let $\text{Graph}(\tilde{\text{BR}}_\varepsilon)$ be the graph of the ε -best-reply correspondence of \tilde{G} . We now use the map f to construct a map $\tilde{f} : \tilde{\Sigma} \rightarrow \tilde{\Sigma}$ such that for each n , \tilde{f}_n is independent of n 's coordinate $\tilde{\sigma}_n$. Let $\tilde{\beta}_{n,0} : \Sigma_n \rightarrow \tilde{\Sigma}_{n,0}$ be the map that assigns to each σ_n , the barycentric coordinates of σ_n in the triangulation \mathcal{T}_n : that is, σ_n can be written uniquely as a convex combination of the vertices of its carrier in \mathcal{T}_n and $\tilde{\beta}_{n,0}$ assigns σ_n to the same convex combination in the face of $\tilde{\Sigma}_{n,0}$ generated by the vertices of that carrier. The map $\tilde{\beta}_{n,1} : \Sigma_{n+1} \rightarrow \tilde{\Sigma}_{n,1}$ is defined similarly. We denote

$$\tilde{\beta}_{n,k} \equiv \times_{\tilde{s}_{n,k}} \beta_{n,k}^{\tilde{s}_{n,k}},$$

where $k = 0, 1$, $\beta_{n,0}^{\tilde{s}_{n,0}} : \Sigma_n \rightarrow \mathbb{R}$ assigns the value of the coordinate $\tilde{s}_{n,0}$ and $\beta_{n,1}^{\tilde{s}_{n,1}} : \Sigma_{n+1} \rightarrow \mathbb{R}$ assigns the value of the coordinate $\tilde{s}_{n,1}$.

Example 4.10. Figure 15 illustrates the definition of the barycentric coordinate map $\tilde{\beta}_{n,k}$. Figure (a) represents a triangulation of a 1-simplex AB , in three vertices $\{A, B, C\}$ and two 1-simplices (AC and BC). We singled-out a point in blue which has BC as a carrier. We define the barycentric coordinate map b from the 1-simplex AB to a 2-simplex. First, the vertices of the 2-simplex are images under b of vertices of the triangulation, as indicated in figure (b). In each 1-simplex of the triangulation represented in (a), the barycentric coordinate map is affine over the same colored face in (b). The barycentric coordinates of the point in blue depicted in (a) are its coordinates with respect to the simplex that carries it in the triangulation, so its coordinates are $(1/2, 1/2)$ —where the first entry corresponds to vertex C and the second to B . These coordinates are preserved in figure (b), since b is affine from the red simplex on (a) over the red simplex on (b). Note in addition that the image of b is the space of a subcomplex of the face complex (cf. Example 2.2) of (b): the subcomplex is defined by the green and red 1-simplices and their vertices, and the space is their union.

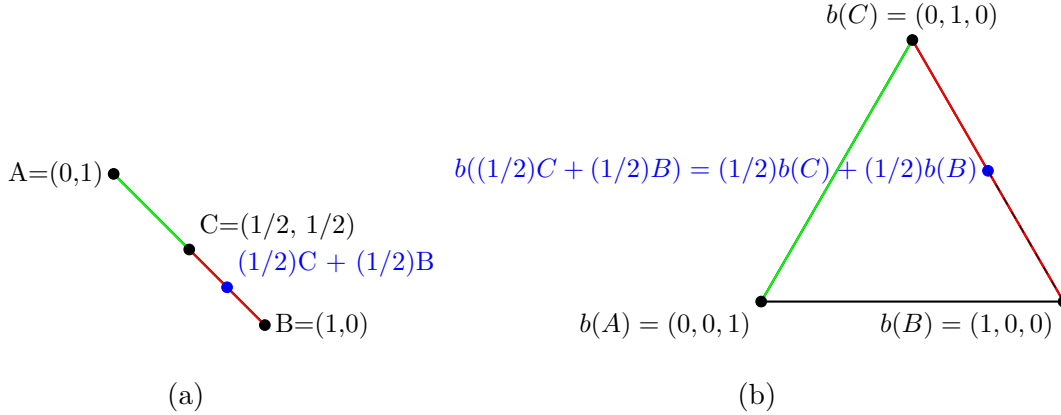
Define \tilde{f} by $\tilde{f}_n(\tilde{\sigma}) = (\tilde{f}_{n,0}(\tilde{\sigma}), \tilde{f}_{n,1}(\tilde{\sigma}))$, where:

$$\tilde{f}_{n,0}(\tilde{\sigma}) \equiv \tilde{\beta}_{n,0}(f_n(\phi_{-n,0}(\tilde{\sigma}_{-n,0}), \phi_{n-1,1}(\tilde{\sigma}_{n-1,1})))$$

with the convention that $n - 1 = N$ if $n = 1$, and

$$\tilde{f}_{n,1}(\tilde{\sigma}) \equiv \tilde{\beta}_{n,1}(\phi_{n+1,0}(\tilde{\sigma}_{n+1,0})).$$

Let $\tilde{\Sigma}_{n,k}^* \equiv \tilde{\beta}_{n,k}(\Sigma_{n+k})$, for $k = 0, 1$. The set $\tilde{\Sigma}_{n,k}^*$ is then the space of a subcomplex $\mathcal{S}_{n,k}^*$ of the face complex of $\Delta(\tilde{\Sigma}_{n,k})$. For each n and $k = 0, 1$, let $\tilde{T}_{n,k}^{ij} \equiv \tilde{\beta}_{n,k}((f_{n+k}^{ij})^{-1}(T_{n,k}^{ij}))$. Recall that $f_n^{ij} : X_n^{ij} \rightarrow Y_n^{ij}$, $X_n^{ij} \subset \text{int}(Y_n^{ij})$ is an affine homeomorphism. Since $T_{n,k}^{ij} \subseteq \text{int}(X_n^{ij})$, then, for $k \in \{0, 1\}$, $(f_n^{ij})^{-1}(T_{n,k}^{ij}) \subseteq T_{n,k}^{ij}$. Hence, $\tilde{\beta}_{n,k}(T_{n,k}^{ij})$ is a face of the simplex $\tilde{\Sigma}_{n,k}$ and $\tilde{T}_{n,k}^{ij}$ is contained in that face.

FIGURE 15. Illustration of $\tilde{\beta}_{n,k}$ 

Let

$$\tilde{\Sigma}_n^* \equiv \tilde{\Sigma}_{n,0}^* \times \tilde{\Sigma}_{n,1}^* \quad ; \quad \tilde{T}_n^{ij} \equiv \tilde{T}_{n,0}^{ij} \times \tilde{T}_{n,1}^{ij} \quad ; \quad \tilde{\Sigma}^* \equiv \prod_n \tilde{\Sigma}_n^* \quad ; \quad \tilde{T}^{ij} \equiv \prod_n \tilde{T}_n^{ij}.$$

For each i, j, n and k , let $\tilde{\sigma}_{n,k}^{ij} \equiv \tilde{\beta}_{n,k}(\sigma_{n+k}^{ij})$, $\tilde{\sigma}_n^{ij} \equiv (\tilde{\sigma}_{n,0}^{ij}, \tilde{\sigma}_{n,1}^{ij})$ and $\tilde{\sigma}^{ij} \equiv (\tilde{\sigma}_n^{ij})_{n \in \mathcal{N}}$. By Property (1) of Lemma 4.7, for each n, i, j and $k = 0, 1$, $\tilde{\sigma}_{n,k}^{ij}$ is the barycenter of both \tilde{T}_n^{ij} and $\tilde{\beta}_{n,k}(T_{n,k}^{ij})$.

The following lemma summarizes key properties of \tilde{f} .

Lemma 4.11. *The function \tilde{f} maps $\tilde{\Sigma}$ into $\tilde{\Sigma}^*$ and player n 's coordinate functions $\tilde{f}_{n,0}$ and $\tilde{f}_{n,1}$ are independent of the coordinate $\tilde{\sigma}_n$. Moreover,*

- (1) $\tilde{f}_{n,1}$ is a function of $\tilde{\sigma}_{n+1,0}$ and the restriction of the function to $\tilde{\Sigma}_{n+1,0}^*$ is affine on each simplex of $\mathcal{S}_{n+1,0}^*$;
- (2) the restriction of \tilde{f}_n to each \tilde{T}_n^{ij} is the cartesian product of affine homeomorphisms $\tilde{f}_{n,0}^{ij} : \tilde{T}_{n-1,1}^{ij} \rightarrow \tilde{\beta}_{n,0}(T_n^{ij})$ and $\tilde{f}_{n,1}^{ij} : \tilde{T}_{n+1,0}^{ij} \rightarrow \tilde{\beta}_{n,1}(T_{n+1}^{ij})$;
- (3) if for each $n \in \mathcal{N}, k = 0, 1$, $\tilde{f}_{n,k}(\tilde{\sigma})$ belongs to the simplicial neighborhood of the closed star of a pure strategy $\tilde{s}_{n,k} \in \tilde{\Sigma}_{n,k}^*$ in the triangulation $\mathcal{S}_{n,k}^*$, then $(\tilde{\sigma}, \tilde{s}) \in \text{Graph}(\tilde{B}R_\varepsilon)$;
- (4) if $\tilde{\sigma} \notin \tilde{T}^{ij} \setminus \partial \tilde{T}^{ij}$ for any i, j and if for all n , $\phi_{n-1,1}(\tilde{\sigma}_{n-1,1})$ belongs to the carrier of $\phi_{n,0}(\tilde{\sigma}_{n,0})$ in \mathcal{T}_n , then there exists n such that $\tilde{\sigma}_{n,0}$ does not belong to the carrier of $\tilde{f}_{n,0}(\tilde{\sigma})$ in $\mathcal{S}_{n,0}^*$;
- (5) the only fixed points of \tilde{f} are the profiles $\tilde{\sigma}^{ij}$ and each $\tilde{\sigma}^{ij}$ has index r^{ij} .

Proof. It follows from the definition of \tilde{f} that each player's coordinate function \tilde{f}_n is independent of his strategies. By definition of \tilde{f} it maps into $\tilde{\Sigma}^*$. The numbered items (1), and (2) also follow by definition of \tilde{f} and $\tilde{\beta}_{n,k}$. As for (3), it follows from property (3) of Lemma 4.7. Property (4) follows from property (2) of Lemma 4.7 in the case where $(\phi_{n,0}(\tilde{\sigma}_{n,0}))_{n \in \mathcal{N}} \notin \cup_{i,j} (T^{ij} \setminus \partial T^{ij})$. Consider now the case where $(\phi_{n,0}(\tilde{\sigma}_{n,0}))_{n \in \mathcal{N}}$ belongs to the set $T^{ij} \setminus (\partial T^{ij} \cup ((f^{ij})^{-1}(T^{ij} \setminus \partial T^{ij})))$ for some i, j .

There exists then n such that $\phi_{n,0}(\tilde{\sigma}_{n,0})$ belongs to the set $T_n^{ij} \setminus (\partial T_n^{ij} \cup ((f_n^{ij})^{-1}(T_n^{ij} \setminus \partial T_n^{ij}))$. As $\phi_{n-1,1}(\tilde{\sigma}_{n-1,1})$ belongs to the same simplex as $\phi_{n,0}(\tilde{\sigma}_{n,0})$, it also lies in $T_n^{ij} \setminus ((f_n^{ij})^{-1}(T_n^{ij} \setminus \partial T_n^{ij}))$. Its image under f_n^{ij} (recall it depends only on $\phi_{n-1,1}(\tilde{\sigma}_{n-1,1})$) lies outside $T_n^{ij} \setminus \partial T_n^{ij}$. This completes the proof of property (4).

There remains to prove point (5). The index of $\tilde{\sigma}^{ij}$ can be computed using the restriction of \tilde{f} to \tilde{T}^{ij} . Call this restriction \hat{f} and note that each coordinate function \hat{f}_n is a cartesian product of affine homeomorphisms: $\hat{f}_n = \tilde{f}_{n,0} \times \tilde{f}_{n-1,1} : \tilde{T}_{n,0}^{ij} \times \tilde{T}_{n-1,1}^{ij} \rightarrow \tilde{\beta}_{n,0}(T_n^{ij}) \times \tilde{\beta}_{n-1,1}(T_n^{ij})$, where $\hat{f}_n(\tilde{\sigma}_{n,0}, \tilde{\sigma}_{n-1,1}) = (\tilde{\beta}_{n,0}(f_n^{ij}(\phi_{n-1,1}(\tilde{\sigma}_{n-1,1}))), \tilde{\sigma}_{n,0})$. For $\lambda \in [0, 1]$, let \hat{f}^λ be the map sending $(\tilde{\sigma}_{n,0}, \tilde{\sigma}_{n-1,1})$ to $(\tilde{\beta}_{n,0}(f_n^{ij}((1-\lambda)\phi_{n,0}(\tilde{\sigma}_{n,0}) + \lambda\phi_{n-1,1}(\tilde{\sigma}_{n-1,1}))), (1-\lambda)\tilde{\sigma}_{n-1,1} + \lambda\tilde{\sigma}_{n,0})$. When $\lambda = 1$, we get the map \hat{f} . For $\lambda > 0$, the unique fixed point is $(\tilde{\sigma}_n^{ij}, \tilde{\sigma}_{n-1,1}^{ij})$ and for $\lambda = 0$, we have a component of fixed points $\{\tilde{\sigma}_{n,0}^{ij}\} \times \tilde{T}_{n-1,1}^{ij}$. By the multiplication property of the index, the index of the component of fixed points for $\lambda = 0$ is the index of σ_n^{ij} under the map f_n^{ij} . Therefore, the index of $\tilde{\sigma}^{ij}$ is r^{ij} , as claimed. \square

Step 5. We now construct a family of ε -perturbed games, indexed by $\tilde{\sigma} \in \tilde{\Sigma}$. First we need to introduce some notation. Given vectors $\tilde{g}^0 \in \prod_n \mathbb{R}^{\tilde{s}_{n,0}}$ and $\tilde{g}^1 \in \prod_n \mathbb{R}^{\tilde{s}_{n,1}}$, we define the game $\tilde{G} \oplus \tilde{g}^0 \oplus \tilde{g}^1$ as a polytope-form game where for each $\tilde{\sigma}$ and each n , the payoff is $\tilde{G}_n(\tilde{\sigma}) + \tilde{g}_n^0 \cdot \tilde{\sigma}_{n,0} + \tilde{g}_n^1 \cdot \tilde{\sigma}_{n,1}$ (where $a \cdot b$ denotes the scalar product of the vectors a and b). Thus $\tilde{g}_{n,\tilde{s}_{n,0}}^0$ represents a “bonus” for playing $\tilde{s}_{n,0}$ and similarly $\tilde{g}_{n,\tilde{s}_{n,1}}^1$, for each $\tilde{s}_{n,1}$. Given functions $\tilde{g}^0 : \tilde{\Sigma} \rightarrow \prod_n \mathbb{R}_+^{\tilde{s}_{n,0}}$, $\tilde{g}^1 : \tilde{\Sigma} \rightarrow \prod_n \mathbb{R}_+^{\tilde{s}_{n,1}}$ and profile $\tilde{\sigma}$, we have a finite polytope-form game $\tilde{G} \oplus \tilde{g}^0(\tilde{\sigma}) \oplus \tilde{g}^1(\tilde{\sigma})$.

For the time being fix $\varepsilon_0 > 0$. The exact choice will be determined later. We first define $\tilde{g}^1(\cdot)$. For each n and $\tilde{s}_{n,1}$, it depends only on $\tilde{\sigma}_{n+1,0}$ and it is given by

$$\tilde{g}_{n,\tilde{s}_{n,1}}^1(\tilde{\sigma}_{-n}) = \varepsilon_0 \tilde{\beta}_{n,1}^{\tilde{s}_{n,1}}(\phi_{n+1,0}(\tilde{\sigma}_{n+1,0})).$$

where $\tilde{\beta}_{n,1}^{\tilde{s}_{n,1}}$ gives the $\tilde{s}_{n,1}$ coordinate of the barycentric coordinates w.r.t. the triangulation \mathcal{T}_{n+1} . To define \tilde{g}^0 , we first construct for each n , a function $r_n : \Sigma_{-n} \rightarrow \mathbb{R}$ satisfying:

- (1) For all $\sigma_{-n} \in \Sigma_{-n}$, $r_n(\sigma_{-n}) \geq \max_{s_n} G_n(s_n, \sigma_{-n})$;
- (2) The map r_n is constant over T_{-n}^{ij} .
- (3) If τ_n is an ε -best reply against σ_{-n} , then $r_n(\sigma_{-n}) - G_n(\tau_n, \sigma_{-n}) < \varepsilon$.

To construct this function, for each i, j and $\sigma \in T^{ij}$, define $r_n(\sigma) = \max_{s_n \in S_n, \sigma' \in T^{ij}} G_n(\sigma', s_n)$. On the sets T^{ij} the function satisfies the three desired properties (thanks to point (4) of Lemma 4.7). There exists an open neighborhood W^{ij} of T^{ij} such that (constantly) extending this function to W^{ij} still satisfies the three properties. By Urysohn's Lemma, there is a continuous function $\xi : \Sigma \rightarrow [0, 1]$ that is one on $\cup_{ij} T^{ij}$ and zero outside $\cup_{ij} W^{ij}$. The function r_n is then obtained

as follows: $r_n(\sigma) = \max_{s_n} G_n(s_n, \sigma_{-n})$ if $\sigma \notin W^{ij}$; $r_n(\sigma) = \xi(\sigma) \max_{s_n \in S_n, \sigma' \in T^{ij}} G_n(\sigma', s_n) + (1 - \xi(\sigma)) \max_{s_n} G_n(s_n, \sigma_{-n})$ if $\sigma \in W^{ij}$.

Recall that $\tilde{\Sigma}_{n,0}^*$ is the space of a subcomplex of the face complex of $\tilde{\Sigma}_{n,0}$. Let $\tilde{\gamma}_{n,s_{n,0}} : \tilde{\Sigma}_{n,0}^* \rightarrow [0, 1]$ be a Urysohn function that is one on the closed star of $\tilde{s}_{n,0}$ in the complex $\mathcal{S}_{n,0}^*$ and zero outside the simplicial neighborhood of the closed star of $\tilde{s}_{n,0}$ in this complex. Then,

$$\tilde{g}_{n,\tilde{s}_{n,0}}^0(\tilde{\sigma}_{-n}) \equiv \tilde{\gamma}_{n,\tilde{s}_{n,0}}(\tau_n)[r_n(\phi_{-n,0}(\tilde{\sigma}_{-n,0})) - G_n(\tilde{\sigma}_{-n,0}, \phi_{n,0}(\tilde{s}_{n,0}))] + \varepsilon_0 \tilde{f}_{n,\tilde{s}_{n,0}}(\tilde{\sigma}_{-n}),$$

where

$$\tau_n = \tilde{\beta}_{n,0}(f_n(\phi_{-n,0}(\tilde{\sigma}_{-n}), \phi_{n-1,1}(\tilde{\sigma}_{n-1,1}))),$$

and $r_n(\cdot)$ is the function we just defined. The following lemma sets out key properties of the perturbations.

Lemma 4.12. *Recall that $\varepsilon > 0$ is fixed from the beginning of the proof (cf. step 2). Each player n 's coordinate functions \tilde{g}_n^0 and \tilde{g}_n^1 are independent of $\tilde{\sigma}_n$. Moreover:*

- (1) *the restriction of $\tilde{g}_{n,\tilde{s}_{n,1}}^1$ to $\Sigma_{n+1,0}^*$ is affine on each simplex of $\mathcal{S}_{n+1,0}^*$;*
- (2) *the restriction of \tilde{g}_n^0 to each \tilde{T}^{ij} is a cartesian product $\prod_{\tilde{s}_{n,0}} \tilde{g}_{n,\tilde{s}_{n,0}}^{ij,0}$, where for each $\tilde{s}_{n,0} \in \tilde{T}_n^{ij}$, $\tilde{g}_{n,\tilde{s}_{n,0}}^{ij,0} : \tilde{T}_{n,0}^{ij} \times \tilde{T}_{n-1,1}^{ij} \rightarrow \mathbb{R}_+$ and $\tilde{G}_n(\tilde{s}_{n,0}, \tilde{\sigma}_{-n}) + \tilde{g}_{n,\tilde{s}_{n,0}}^{ij,0}(\tilde{\sigma}_{-n}) = r_n(\phi_{-n,0}(\tilde{\sigma}_{-n,0})) + \varepsilon_0 \tilde{f}_{n,\tilde{s}_{n,0}}(\tilde{\sigma}_{n-1,1})$, which is an affine function.*
- (3) *$\|\tilde{g}^0\| + \|\tilde{g}^1\| < \varepsilon$, if ε_0 is small.*

Proof. Item (1) follows from the fact that $\tilde{\beta}_{n,1}^{\tilde{s}_{n,1}}$ is affine in each cell of $\mathcal{S}_{n+1,0}^*$ and $\phi_{n+1,0}$ is affine by definition. For (2), the map \tilde{g}_n^0 is a cartesian product $\prod_{\tilde{s}_{n,0}} \tilde{g}_{n,\tilde{s}_{n,0}}^{ij,0}$ by definition. Fix $\tilde{\sigma}_{-n} \in \tilde{T}_{-n}^{ij}$. Note that any vertex $\tilde{s}_{n,0}$ of \tilde{T}_n^{ij} is such that $\tilde{\gamma}_{n,\tilde{s}_{n,0}}(\tau_n) = 1$, where $\tau_n = \tilde{\beta}_{n,0}(f_n(\phi_{n-1,1}(\tilde{\sigma}_{n-1,1})))$ (recall f_n depends exclusively on $\sigma_{n-1,1}$ in this region). Therefore, $\tilde{G}_n(\tilde{s}_{n,0}, \tilde{\sigma}_{-n}) + \tilde{g}_{n,\tilde{s}_{n,0}}^{ij,0}(\tilde{\sigma}_{-n}) = r_n(\phi_{-n,0}(\tilde{\sigma}_{-n,0})) + \varepsilon_0 \tilde{f}_{n,\tilde{s}_{n,0}}(\tilde{\sigma}_{n-1,1})$. By definition $r_n(\phi_{-n,0}(\tilde{\sigma}_{-n,0}))$ is constant on \tilde{T}^{ij} , which makes $r(\phi_{-n,0}(\tilde{\sigma}_{-n,0})) + \varepsilon_0 \tilde{f}_{n,\tilde{s}_{n,0}}(\tilde{\sigma}_{n-1,1})$ affine. For (3), it is sufficient to prove that the difference

$$\tilde{\gamma}_{n,\tilde{s}_{n,0}}(\tau_n)[r_n(\phi_{-n,0}(\sigma_{-n,0})) - G_n(\tilde{\sigma}_{-n,0}, \phi_{n,0}(\tilde{s}_{n,0}))]$$

is strictly less than $\varepsilon > 0$. First, if τ_n is outside the simplicial neighborhood of the closed star of $\tilde{s}_{n,0}$, then it follows that $\tilde{\gamma}_{n,\tilde{s}_{n,0}}(\tau_n) = 0$ and the result follows. If τ_n is inside that simplicial neighborhood, then item (3) of Lemma 4.11 implies $\tilde{s}_{n,0}$ is an ε -best reply to $\tilde{\sigma}$ in \tilde{G} . Hence $\varepsilon > r_n(\phi_{-n,0}(\sigma_{-n,0})) - G_n(\tilde{\sigma}_{-n,0}, \phi_{n,0}(\tilde{s}_{n,0})) \geq 0$, which implies the desired result. \square

We conclude this step with a lemma concerning best replies in these perturbed games. Its proof, too, follows from the construction of the perturbed games.

Lemma 4.13. *The profile $(\tilde{s}_{n,0}, \tilde{s}_{n,1})$ is a best reply to $\tilde{\sigma}$ in the game $\tilde{G} \oplus \tilde{g}^0(\tilde{\sigma}) \oplus \tilde{g}^1(\tilde{\sigma})$ only if:*

- (1) $\tilde{s}_{n,0}$ is a vertex of the carrier of $\tilde{f}_{n,0}(\tilde{\sigma}_{-n,0}, \tilde{\sigma}_{n-1,1})$ in $\mathcal{S}_{n,0}^*$;
- (2) $\tilde{s}_{n,1}$ is a vertex of the carrier of $\tilde{f}_{n,1}(\tilde{\sigma}_{n+1,0})$ in $\mathcal{S}_{n+1,0}^*$.

Proof. To prove (1), take $(\tilde{t}_{n,0}, \tilde{s}_{n,1})$ such that $\tilde{t}_{n,0}$ belongs to the carrier of $\tilde{f}_{n,0}(\tilde{\sigma}_{-n,0}, \tilde{\sigma}_{n-1,1})$ in $\mathcal{S}_{n,0}^*$. Then the payoff it obtains in $\tilde{G} \oplus \tilde{g}^0(\tilde{\sigma}) \oplus \tilde{g}^1(\tilde{\sigma})$ is $r_n(\phi_{-n,0}(\tilde{\sigma}_{-n,0})) + \varepsilon_0 \tilde{f}_{n,\tilde{t}_{n,0}}(\tilde{\sigma}_{-n}) + \tilde{g}_{n,\tilde{s}_{n,1}}^1(\tilde{\sigma}_{-n})$; this quantity is strictly larger than the payoff from $(\tilde{s}_{n,0}, \tilde{s}_{n,1})$ where $\tilde{s}_{n,0}$ is not in that carrier, because $\gamma_{n,\tilde{s}_{n,0}}(\tau_n) \leq 1$ and $\tilde{f}_{n,\tilde{s}_{n,0}}(\tilde{\sigma}_{-n}) = 0$. The proof of (2) follows a similar reasoning: since $\tilde{g}_{n,\tilde{s}_{n,1}}^1(\tilde{\sigma}_{-n}) = \varepsilon_0 \tilde{\beta}_{n,1}^{\tilde{s}_{n,1}}(\phi_{n+1,0}(\tilde{\sigma}_{n+1,0})) = \varepsilon_0 \tilde{f}_{\tilde{s}_{n,1}}(\tilde{\sigma}_{n+1,0})$, a best reply $\tilde{s}_{n,1}$ must be a vertex of the of the carrier $\tilde{f}_{n,1}(\tilde{\sigma}_{n+1,0})$. \square

Step 6. Let $\tilde{\varphi} : \tilde{\Sigma} \rightrightarrows \tilde{\Sigma}$ be the correspondence that assigns to each $\tilde{\sigma}$ the set of best replies to it in the polytope-form game $\tilde{G} \oplus \tilde{g}^0(\tilde{\sigma}) \oplus \tilde{g}^1(\tilde{\sigma})$. The correspondence $\tilde{\varphi}$ is non-empty, compact, and convex-valued. It is also upper semi-continuous. Hence it has a fixed point and a well-defined index for its fixed points. We now characterize the fixed points.

Lemma 4.14. *The fixed points of $\tilde{\varphi}$ are the profiles $\tilde{\sigma}^{ij}$. The index of each $\tilde{\sigma}^{ij}$ is r^{ij} .*

Proof. To prove this step we will show that for all $\lambda \in [0, 1]$ the only fixed points of $\lambda \tilde{f} + (1 - \lambda) \tilde{\varphi}$ are the $\tilde{\sigma}^{ij}$'s. The proof then follows from item (5) of Lemma 4.11.


Take some $\tilde{\sigma}$ that is a fixed point of $\lambda \tilde{f} + (1 - \lambda) \tilde{\varphi}$ for some $\lambda \in [0, 1]$. It follows from (1) of Lemma 4.13 that $\tilde{\sigma}_{n,0}$ belongs to the simplex of $\mathcal{S}_{n,0}^*$ that contains $\tilde{f}_{n,0}(\tilde{\sigma}_{-n,0}, \tilde{\sigma}_{n-1,1})$ in its interior. An optimal $\tilde{s}_{n-1,1}$ has to be a vertex of the simplex containing $\tilde{\sigma}_{n,0}$ (which is also the carrier of the image of $\tilde{\sigma}_{n,0}$ under $\tilde{f}_{n-1,1}$). Hence $\tilde{\sigma}_{n-1,1}$ belongs to the simplex containing $\tilde{\sigma}_{n,0}$. By property (4) of Lemma 4.11, $\tilde{\sigma}$ must be in \tilde{T}^{ij} for some i, j . It then follows in this case $\tilde{\sigma} = \tilde{\sigma}^{ij}$. \square

Recall that for each $\tilde{\sigma} \in \tilde{\Sigma}$, we can define a polytope-form game $\tilde{G} \oplus \tilde{g}^0(\tilde{\sigma}) \oplus \tilde{g}^1(\tilde{\sigma})$. Hence we have a continuous family of polytope-form games indexed by $\tilde{\sigma} \in \tilde{\Sigma}$: $\{\tilde{G} \oplus \tilde{g}^0(\tilde{\sigma}) \oplus \tilde{g}^1(\tilde{\sigma})\}_{\tilde{\sigma} \in \tilde{\Sigma}}$. We will now use this family to construct a game \hat{G} that is equivalent to \tilde{G} and identify a finite payoff perturbation that attains the objectives of our main theorem. This construction is performed in the final three steps.

Step 7. In this step, we describe a triangulation of $\tilde{\Sigma}_n$ for each n that allows us to convert the continuous family of perturbed games to a single perturbed game. We first observe that we can obtain a $\xi > 0$ such that:

- (A) for $\tilde{\sigma} \notin \cup_{i,j}(\tilde{T}^{ij} \setminus \partial \tilde{T}^{ij})$, if $\tilde{g}^0 \in \mathbb{R}^{\sum_n |\tilde{S}_{n,0}|}$ and $\tilde{g}^1 \in \mathbb{R}^{\sum_n |\tilde{S}_{n,1}|}$ satisfy $\|\tilde{g}^0 - \tilde{g}^0(\tilde{\sigma})\| + \|\tilde{g}^1 - \tilde{g}^1(\tilde{\sigma})\| < 2\xi$, then $\tilde{\sigma}$ is not an equilibrium of $\tilde{G} \oplus \tilde{g}^0 \oplus \tilde{g}^1$.

FIGURE 16. Illustration of property (2) of Lemma 4.15



Claim (A) follows from the fact that if $\tilde{\sigma} \notin \cup_{i,j}(\tilde{T}^{ij} \setminus \partial \tilde{T}^{ij})$, then $\tilde{\sigma}$ is not an equilibrium of $\tilde{G} \oplus \tilde{g}^0(\tilde{\sigma}) \oplus \tilde{g}^1(\tilde{\sigma})$ (cf. Lemma 4.14). Therefore, for sufficiently small perturbations \tilde{g}^0 of $\tilde{g}^0(\tilde{\sigma})$ and \tilde{g}^1 of $\tilde{g}^1(\tilde{\sigma})$, $\tilde{\sigma}$ is also not an equilibrium of the finite game $\tilde{G} \oplus \tilde{g}^0 \oplus \tilde{g}^1$.

Since $g(\cdot)$ is uniformly continuous, there exists $\zeta > 0$ such that $\|g(\tilde{\sigma}) - g(\tilde{\sigma}')\| < \xi$ if $\|\tilde{\sigma} - \tilde{\sigma}'\| < \zeta$. Recall that $\tilde{S}_{n,0}$ is the set of vertices of \mathcal{T}_n and $\tilde{\Sigma}_{n,0} = \tilde{\Sigma}_{n-1,1} = \Delta(\tilde{S}_{n,0})$. Therefore, any triangulation of $\Delta(\tilde{S}_{n,0})$ induces a corresponding triangulation of $\tilde{\Sigma}_{n,0}$ and $\tilde{\Sigma}_{n-1,1}$.


Lemma 4.15. *For each n , there exists a triangulation $\tilde{\mathcal{T}}_n$ of $\Delta(\tilde{S}_{n,0})$ with the following properties.*

- (1) *the diameter of each simplex is less than ζ ;*
- (2) *for each i, j , $\tilde{T}_{n,0}^{ij}$ is the space of a subcomplex of \mathcal{T}_n , $\tilde{\sigma}_{n,k}^{ij}$ lies in the interior of a simplex $\hat{T}_{n,k}^{ij} \in \mathcal{T}_n$ for $k = 0, 1$, with $\dim(\hat{T}_{n,k}^{ij}) = \dim(\Sigma_{n+k})$, and $(\tilde{\sigma}_{n,0}^{ij}, \tilde{\sigma}_{n,1}^{ij})$ is not in the convex hull of $(\dim(\Sigma_n) + \dim(\Sigma_{n+1}))$ vertices of $\hat{T}_{n,0}^{ij} \times \hat{T}_{n,1}^{ij}$;*
- (3) *there is a convex piecewise linear function $\gamma_n : \Delta(\tilde{S}_{n,0}) \rightarrow \mathbb{R}_+$ such that:*
 - (a) *the map is linear precisely on the simplices of $\tilde{\mathcal{T}}_n$;*
 - (b) *the function is constant over each $\hat{T}_{n,0}^{ij}$.*

Proof. See Appendix. □

Example 4.16. Figure 16 illustrates condition (2) of Lemma 4.15. The idea is to view AD as the 1-simplex $\hat{T}_{n,0}^{ij}$ and AB as the 1-simplex $\hat{T}_{n,1}^{ij}$. The condition says the point σ_n^{ij} should avoid both diagonals AC and DB (as figure (b) shows). The purpose of this requirement will be clear in step 8 and it is illustrated in Example 4.17.

Step 8. We now define a game \hat{G} in normal form. Let $\hat{S}_{n,0}$ be the set of vertices of $\tilde{\mathcal{T}}_n$. For each n , let $\hat{S}_{n,1} \equiv \hat{S}_{n+1,0}$ and $\hat{S}_n \equiv \hat{S}_{n,0} \times \hat{S}_{n,1}$. The pure strategy set of player n in \hat{G} is \hat{S}_n . Let

FIGURE 17. Illustration of \hat{C}_n^{ij} 

$\hat{\Sigma}_{n,0} = \Delta(\hat{S}_{n,0})$, $\hat{\Sigma}_{n,1} = \Delta(\hat{S}_{n,1})$, and $\hat{\Sigma} = \hat{\Sigma}_{n,0} \times \hat{\Sigma}_{n,1}$. For $k = 0, 1$, let $\hat{\mu}_{n,k} : \hat{\Sigma}_n \rightarrow \hat{\Sigma}_{n,k}$ be the map that sends each $\hat{\sigma}_n$ to the marginal distribution over $\hat{S}_{n,k}$. Each $\hat{s}_{n,0} \in \hat{S}_{n,0}$ corresponds to a mixed strategy in $\hat{\Sigma}_{n,0}$. Therefore, there exists for $k = 0, 1$, an affine map $\hat{\phi}_{n,k} : \hat{\Sigma}_{n,k} \rightarrow \hat{\Sigma}_{n,k}$. Let $\hat{\phi}_n \equiv \hat{\phi}_{n,0} \times \hat{\phi}_{n,1}$ and $\hat{\mu}_n \equiv \hat{\mu}_{n,0} \times \hat{\mu}_{n,1}$. The payoffs in \hat{G} are defined, for each player n , by $\hat{G}_n(\hat{\sigma}) = \tilde{G}_n(\hat{\phi}_0 \circ \hat{\mu}_0(\hat{\sigma}))$, where $\hat{\phi}_0 \equiv \times_n \hat{\phi}_{n,0}$ and $\hat{\mu}_0 \equiv \times_n \hat{\mu}_{n,0}$. The game \hat{G} just defined is equivalent to \tilde{G} and hence to G .

For each n, i, j and $k = 0, 1$, let $\hat{S}_{n,k}^{ij}$ be the set of vertices of $\hat{T}_{n,k}^{ij}$ obtained in (2) of Lemma 4.15 and let $\hat{\Sigma}_{n,k}^{ij}$ be the face of $\hat{\Sigma}_{n,k}$ consisting of strategies with support contained in $\hat{S}_{n,k}^{ij}$. Let $\hat{\Sigma}_n^{ij}$ be the set of mixed strategies $\hat{\sigma}_n$ such that $\hat{\mu}_{n,k}(\hat{\sigma}_n) \in \hat{\Sigma}_{n,k}^{ij}$ for each k . Let $\hat{\Sigma}^{ij} \equiv \prod_n \hat{\Sigma}_n^{ij}$. Let $\hat{C}^{ij} = \{\hat{\sigma} \in \hat{\Sigma}^{ij} \mid \hat{\phi}_{n,k}(\hat{\mu}_{n,k}(\hat{\sigma}_n)) = \tilde{\sigma}_{n,k}^{ij}, \text{ for each } n, k\}$. The set \hat{C}^{ij} is then a polytope, since $\hat{\phi}_{n,k} \circ \hat{\mu}_{n,k}$ is affine, for $k = 0, 1$. Observe that $\hat{C}^{ij} = \prod_n \hat{C}_n^{ij}$, where $\hat{C}_n^{ij} = \{\hat{\sigma}_n \in \hat{\Sigma}_n^{ij} \mid \hat{\phi}_{n,k}(\hat{\mu}_{n,k}(\hat{\sigma}_n)) = \tilde{\sigma}_{n,k}^{ij}, \text{ for each } k = 0, 1\}$.

By Property (2) of Lemma 4.15, each extreme point of \hat{C}_n^{ij} belongs to a face of $\hat{\Sigma}_n$ of dimension $\dim(\Sigma_n) + \dim(\Sigma_{n+1})$ (see Example 4.17). Let $\hat{\sigma}^{ij}$ be one of the extreme points of the polytope. For each n , let $\hat{S}_n^{ij,*}$ be the support of $\hat{\sigma}^{ij}$.

Example 4.17. Figure 17 illustrates the set \hat{C}_n^{ij} and the effects of property (2) of Lemma 4.15. The vertices of the square in (b) of Figure 16 are declared to be pure strategies for player n in \hat{G} , as we described above, yielding a tetrahedron as a mixed strategy set, like in Figure 17. Because of property (2) of Lemma 4.15, the end-points of the resulting component of equilibria (one illustrated in red, and the other in blue), are contained in the interior of faces of dimension equal to $\dim(\Sigma_n)$. Without property (2) of Lemma 4.15, it could be the case that such end points would fall in faces of lower dimension, which could restrict the index number we would like to assign to equilibria. This point will be addressed in the last step.

We will now construct an ε -perturbation of \hat{G} . In the next and final step, we show that the equilibria of this perturbed game are the points $\hat{\sigma}^{ij}$ and that the index of each $\hat{\sigma}^{ij}$ is r^{ij} .

The perturbation of the payoffs of \hat{G} for each n has five components: $\hat{G}_n^0 : \hat{\Sigma} \rightarrow \mathbb{R}_{++}$; $\hat{G}_n^1 : \hat{\Sigma} \rightarrow \mathbb{R}_+$; $\hat{G}_n^* : \hat{\Sigma} \rightarrow \mathbb{R}_-$; $\hat{g}_n^0 : \hat{S}_{n,0} \rightarrow \mathbb{R}$; and $\hat{g}_n^1 : \hat{S}_{n,1} \rightarrow \mathbb{R}$. For $\hat{s} \in \hat{S}$, with $\hat{s}_m = (\hat{s}_{m,0}, \hat{s}_{m,1})$ for each m :


$$\begin{aligned}\hat{G}_n^0(\hat{s}) &= \hat{\phi}_{n,0}(\hat{s}_{n,0}) \cdot \tilde{g}_n^0(\hat{\phi}_{-n,0}(\hat{s}_{-n,0}), \hat{\phi}_{n-1,1}(\hat{s}_{n-1,1})) \\ \hat{G}_n^1(\hat{s}) &= \hat{\phi}_{n,1}(\hat{s}_{n,1}) \cdot \tilde{g}_n^1(\hat{\phi}_{n+1,0}(\hat{s}_{n+1,0})) \\ \hat{G}_n^*(\hat{s}) &= -\mathbb{1}_{[\cup_{i,j}(\hat{S}_n^{ij} \setminus \hat{S}_n^{ij,*}) \times \hat{S}_{-n}^{ij}]} \\ \hat{g}_n^0(\hat{s}_n) &= -\gamma_n(\hat{\phi}_{n,0}(\hat{s}_{n,0})) \\ \hat{g}_n^1(\hat{s}_n) &= -\gamma_{n+1}(\hat{\phi}_{n,1}(\hat{s}_{n,1})).\end{aligned}$$

The function \hat{G}_n^0 depends on $\hat{\sigma}_{-n,0}$ and $\hat{\sigma}_{n,0}$ (the marginal of strategies $\hat{\sigma}_n$ on $\hat{\Sigma}_{n,0}$); the function \hat{G}_n^1 depends on $\hat{\sigma}_{n+1,0}$ and $\hat{\sigma}_{n,1}$. The function \hat{G}_n^1 is affine on subsets of $\hat{\Sigma}_{n+1,0}$ that project under $\hat{\phi}_{n+1,0}$ to simplices of $\tilde{\mathcal{T}}_{n+1}$: recall from item (1) in Lemma 4.12 that this is the case in the subcomplex $\mathcal{S}_{n+1,0}^*$, and everywhere else $\tilde{g}_{n,\hat{s}_{n,1}}^1$ is identically 0. These two functions provide a multilinear approximation of the continuous perturbations \tilde{g}_n^0 and \tilde{g}_n^1 : they are linear in the strategies of each player m . Observe that two strategies $\hat{\sigma}_n$ of player n that have the same marginals $\hat{\sigma}_{n,0}$ and $\hat{\sigma}_{n,1}$ are equally good replies if we use only these perturbations \hat{G}_n^0 and \hat{G}_n^1 : in particular there are a lot of equilibria of the game that are equivalent to the equilibria σ^{ij} . In order to eliminate this multiplicity, we introduce three other components in the perturbation.

The function \hat{G}^* penalizes n 's use of a strategy in $\cup_{i,j} \hat{S}_n^{ij}$ that does not belong to the support of $\hat{\sigma}_n^{ij}$ (i.e., $\hat{S}_n^{ij,*}$) when n 's opponents are playing a strategy in \hat{S}_{-n}^{ij} . Even with all other perturbations, there is a multiplicity of equilibria equivalent to σ^{ij} as those arguments explicitly depend on the marginals on Σ_n^0 and Σ_n^1 , rather than the normal-form strategies $\hat{\Sigma}_n$. This penalty helps to isolate a unique profile among those that are equivalent to σ^{ij} .

The functions \hat{g}_n^0, \hat{g}_n^1 are bonus functions like we constructed for the game \tilde{G} . Their purpose is to eliminate equivalent equilibria that arise when we introduce duplicate strategies as the vertices of triangulation. We illustrate this in Example 4.18 below.

Example 4.18. Recall that γ_n is a piecewise linear convex function that is linear precisely in the cells of the triangulation $\tilde{\mathcal{T}}_n$ (cf. property (3) of Lemma 4.15). Suppose player n has a 1-simplex as mixed strategy set, represented in the x -axis in Figure 18. This simplex is triangulated as indicated, i.e., partitioned in intervals, and the vertices introduced in this triangulation are duplicates of mixed strategies in an equivalent game. Once the vertices become duplicates, player n has multiple ways of playing the same mixed strategy: for example, in order to play the strategy corresponding to 0.5 in the figure, player n can either play 0 with probability 1/2 and 1 with probability 1/2, or 0.25 with probability 1/2 and 0.75 with probability 1/2, or even 0.5 with probability 1. All these strategies induce the same payoffs against any strategy of the opponents in the equivalent game. However, when we add to the payoffs of player n a piecewise linear convex penalty, as illustrated

FIGURE 18. The role of $\hat{g}_n^k, k = 0, 1$ 

by the graph of the function in red, then the strategy with the lowest penalty associated is 0.5. In general, for any strategy σ_n of the 1-simplex, among the many convex combinations that can induce this strategy in the equivalent game, the vertices of the carrier of σ_n in the triangulation are the strategies with lowest associated penalties in the equivalent game. So, if σ_n is a best reply to some opponents' profile in the original game, adding \hat{g}_n^k to the payoffs in the equivalent game allows us to select the support of the best reply (in the equivalent game) among all strategies equivalent to σ_n .⁷

For positive constants α and α^* , define the game $\hat{G}^{\alpha, \alpha^*} \equiv \hat{G} \oplus \hat{G}^0 \oplus \hat{G}^1 \oplus \alpha^* \hat{G}^* \oplus \alpha \hat{g}^0 \oplus \alpha \hat{g}^1$. Fix $\alpha > 0$ sufficiently small such that $\hat{G}^{\alpha, 0}$ is an ε -perturbation. If α^* is chosen small as well then $\hat{G}^{\alpha, \alpha^*}$ is also an ε -perturbation of \hat{G} .

Step 9. Let us first analyze the game $\hat{G}^{\alpha, 0}$, i.e, the perturbation with $\alpha^* = 0$. Suppose $\hat{\phi}(\mu(\hat{\sigma})) = \tilde{\sigma}^{ij}$ for some ij . We claim that $\hat{\sigma}$ is an equilibrium iff $\hat{\sigma} \in \hat{C}^{ij}$. If $\hat{\sigma}_n \notin \hat{C}_n^{ij}$ for some n , then a direct computation shows that $\hat{\sigma}_n^{ij}$ gives: (1) the same payoff as $\hat{\sigma}_n$ in the components $\hat{G} \oplus \hat{G}^0 \oplus \hat{G}^1$; and (2) a strictly higher payoff in $\hat{g}^0 \oplus \hat{g}^1$. Thus, if $\hat{\sigma}$ is an equilibrium then $\hat{\sigma} \in \hat{C}^{ij}$. Conversely, a similar computation shows that if $\hat{\sigma} \in \hat{C}^{ij}$, then any pure strategy that is not in \hat{S}_n^{ij} is a strictly inferior reply against $\hat{\sigma}$, and $\hat{\sigma}$ is an equilibrium of $\hat{G}^{\alpha, 0}$.

We will now show that the \hat{C}^{ij} 's are the only equilibria of the game $\hat{G}^{\alpha, 0}$ (note that in the previous paragraph we showed that they are the only equilibria of $\hat{G}^{\alpha, 0}$ among those profiles $\hat{\sigma}$ for which $\hat{\phi}(\mu(\hat{\sigma})) = \tilde{\sigma}^{ij}$). Let $\hat{\sigma}$ be an equilibrium of $\hat{G}^{\alpha, 0}$. For each n , let $\tilde{\sigma}_{n,k} = \hat{\phi}_{n,k}(\hat{\mu}_{n,k}(\hat{\sigma}_n))$ for $k = 0, 1$. It is sufficient to show that $\tilde{\sigma} = \tilde{\sigma}^{ij}$ for some i, j .

Any strategy $\hat{\tau}_n$ of player n in $\hat{G}^{\alpha, 0}$ such that $\tilde{\sigma}_{n,k} = \hat{\phi}_{n,k}(\hat{\mu}_{n,k}(\hat{\tau}_n))$ for each k yields the same payoff in \hat{G} , \hat{G}^0 , and \hat{G}^1 . The choice of which among these strategies is optimal depends on the payoffs they obtain under \hat{g}_n^k for $k = 0, 1$. The optimality of $\hat{\sigma}_n$ implies that $\hat{\sigma}_{n,0}$ is a mixture

⁷The same technique has been used to select supports of best replies in [4] and [6].

over the vertices of the simplex of $\tilde{\mathcal{T}}_n$ that contains $\tilde{\sigma}_{n-1,1}$ in its interior and $\hat{\sigma}_{n,1}$ is a mixture over the vertices of the simplex of $\tilde{\mathcal{T}}_{n+1}$ that contains $\tilde{\sigma}_{n+1,0}$ in its interior (cf. with our discussion in Example 4.18). Thus $\tilde{\sigma}$ is an equilibrium of the game $\tilde{G} \oplus g^0 \oplus g^1$ where:

$$\begin{aligned} g_{n,\tilde{s}_{n,0}}^0 &= \sum_{\tilde{s}_{-n}} \prod_{m \neq n} \tilde{\sigma}_m(\tilde{s}_m) \tilde{g}_{n,\tilde{s}_{n,0}}^0(\tilde{s}_{-n,0}, \tilde{s}_{n-1,1}) \\ g_{n,\tilde{s}_{n,1}}^1 &= \sum_{\tilde{s}_{n+1}} \tilde{\sigma}_{n+1}(\tilde{s}_{n+1}) \tilde{g}_{n,\tilde{s}_{n,1}}^1(\tilde{s}_{n+1,0}). \end{aligned}$$

Since for each $k = 0, 1$, $\|g^k - \tilde{g}^k(\tilde{\sigma})\| < \xi$, (A) of step 7 implies that $\tilde{\sigma}$ is in $\tilde{T}^{ij} \setminus \partial \tilde{T}^{ij}$ for some i, j . If $\tilde{\sigma} \neq \tilde{\sigma}^{ij}$, there is some n for which either (a) $\tilde{\sigma}_{n,0} \neq \tilde{\sigma}_{n,0}^{ij}$ or (b) $\tilde{\sigma}_{n,1} \neq \tilde{\sigma}_{n,1}^{ij}$. Case (a) implies that $\tilde{\sigma}_{n-1,1}$ belongs to the boundary of $T_{n-1,1}^{ij}$ by definition of \tilde{g}_{n-1}^1 and the fact it is affine over $\tilde{T}_{n,0}^{ij}$, which is a contradiction with our assumption that $\tilde{\sigma}$ is in $\tilde{T}^{ij} \setminus \partial \tilde{T}^{ij}$. Hence, (a) does not occur, and it must be (b). The payoffs from $\tilde{\sigma}$ to player $n+1$ in $\tilde{G} \oplus g^0 \oplus g^1$ are then $\sum_{\tilde{s}_{-n}} \prod_{m \neq n} \tilde{\sigma}_m(\tilde{s}_m) r_n(\phi_{n,0}(\tilde{s}_{-n,0})) + \tilde{\sigma}_{n+1,0} \cdot \varepsilon_0 \tilde{f}_n(\tilde{\sigma}_{n,1}) + \tilde{\sigma}_{n+1,1} \cdot g_n^1$. By construction, r_n is constant on \tilde{T}_{-n}^{ij} (cf. step 5), which implies it is $\tilde{f}_{n+1,0}(\tilde{\sigma}_{n,1})$ that decides which $\tilde{\sigma}_{n+1,0}$ is a best reply. But since $\tilde{f}_{n+1,0}(\tilde{\sigma}_{n,1}) \neq \tilde{\sigma}_{n+1}^{ij}$ (cf. (3) of Lemma 4.1), it follows that $\tilde{\sigma}_{n+1,0}$ belongs to the boundary of $\tilde{\beta}_{n+1,0}(T_{n+1}^{ij})$, which is again a contradiction. Therefore, $\tilde{\sigma} = \tilde{\sigma}^{ij}$. From the claim proved at the beginning of this step, it follows that the \hat{C}^{ij} 's are the only equilibria of the game $\hat{G}^{\alpha,0}$, as we wanted to prove.

Now we turn to the equilibria of $\hat{G}^{\alpha,\alpha^*}$. For each ij , $\hat{\sigma}^{ij}$ is an isolated equilibrium of $\hat{G}^{\alpha,\alpha^*}$ because of the penalty function \hat{G}^* . We will show that the index of $\hat{\sigma}^{ij}$ is r^{ij} . Since the strategies not in $\hat{\Sigma}^{ij*}$ are inferior replies to $\hat{\sigma}^{ij}$, the index can be computed in the game obtained by deleting them. Recall also that \hat{g}_n^0 and \hat{g}_n^1 are constant over $\hat{\Sigma}_n^{ij}$. Therefore, the payoff of player n in $\hat{G} \oplus \hat{G}^0 \oplus \hat{G}^1$ restricted to $\hat{\Sigma}_n^{ij}$ from any strategy $(\hat{s}_{n,0}, \hat{s}_{n,1})$ is an affine function of $\hat{\sigma}_{-n}$ equal to $D + \hat{\phi}_{n,0}(\hat{s}_{n,0}) \cdot \tilde{f}_{n,0}(\hat{\phi}_{n-1,1}(\hat{\mu}_{n-1,1}(\hat{\sigma}_{n-1}))) + \hat{\phi}_{n,1}(\hat{s}_{n,1}) \cdot \tilde{f}_{n,1}(\hat{\phi}_{n+1,0}(\hat{\mu}_{n+1,0}(\hat{\sigma}_{n+1})))$, where the constant D comes from the fact that r_n is constant in $\hat{\Sigma}^{ij}$. Therefore, the best reply correspondence of this game assigns to $\hat{\sigma}^{ij}$ the same index as $\tilde{\varphi}$ assigns to $\tilde{\sigma}^{ij}$.

To finish the proof, it remains to be shown that for small $\alpha^* > 0$, the only equilibria of $\hat{G}^{\alpha,\alpha^*}$ are the $\hat{\sigma}^{ij}$'s. Let $\hat{\sigma}$ be an equilibrium of $\hat{G}^{\alpha,\alpha^*}$ where $\alpha^* > 0$ and let $\tilde{\sigma} = \hat{\phi} \circ \hat{\mu}(\hat{\sigma})$. If α^* is sufficiently small, then $\tilde{\sigma}$ is close to $\tilde{\sigma}^{ij}$ for some ij and there exists $\hat{\sigma}^*$ with support S^{ij*} such that $\hat{\phi} \circ \hat{\mu}(\hat{\sigma}^*) = \tilde{\sigma}$. Also, any strategy that is not in S_n^{ij} for player n is a inferior reply to $\hat{\sigma}$, as this property holds when $\alpha^* = 0$ and the equilibrium is in \hat{C}^{ij} . This implies that $\hat{\sigma}_n$ has support in \hat{S}_n^{ij} , for each n . Because of the penalty function \hat{G}_n^* , the support of $\hat{\sigma}_n$ must be in particular in \hat{S}_n^{ij*} . Therefore, $\hat{\sigma}^* = \hat{\sigma}$. The restriction of the game to S^{ij*} is a polymatrix game and as such the set of its equilibria with support S^{ij*} is a convex set. As $\hat{\sigma}^{ij}$ is an isolated equilibrium with support S^{ij*} , $\hat{\sigma}^{ij} = \hat{\sigma}$ and the proof is complete.

APPENDIX: PROOF OF LEMMA 4.15

Let $m \equiv \dim(\Delta(\tilde{S}_{n,0}))$, so the number of vertices of $\Delta(\tilde{S}_{n,0})$ is $m + 1$. Recall the definition of \mathcal{T}_n in Lemma 4.7. Note that \mathcal{T}_n can be assumed to have any finite number of vertices. Letting $K_n \equiv \#\{\tilde{\sigma}_n^{ij}\}_{i,j}$, we will therefore assume that $K_n[\dim(\Sigma_n) + 1] < m + 1$.

A polyhedral subdivision \mathcal{P}_n of $\Delta(\tilde{S}_{n,0})$ is said to be *regular* (cf. [2]) if there exists a *height function* $h : P_n \rightarrow \mathbb{R}$ from the vertex set P_n of \mathcal{P}_n , where the vertex set $\{p_t\}_t$ of a maximal-dimensional cell of the subdivision is mapped to points in \mathbb{R} s.t. $\{(p_t, h(p_t))\}_t$ spans a (non-vertical) hyperplane in $\Delta(\tilde{S}_{n,0}) \times \mathbb{R}$ and all other points $p \in P_n$ are such that $(p, h(p))$ is strictly above the hyperplane. In this case, we say that the set $\{p_t\}_t$ is *lifted* by h to a *hyperplane*, with all other vertices *lifted above* the hyperplane.

Let $\tilde{\mathcal{T}}_n$ be a regular triangulation with vertex set T_n . For each $p \in T_n$, let $v_p \in \Delta(\tilde{S}_{n,0})$ be any point with the same carrier as p in the face complex of $\Delta(\tilde{S}_{n,0})$. For $\varepsilon \geq 0$, let $p(\varepsilon) \equiv (1 - \varepsilon)p + \varepsilon v_p$. For $\varepsilon > 0$ sufficiently small, define the triangulation $\mathcal{T}_n^\varepsilon$ as follows: $B(\varepsilon) \equiv [p_0(\varepsilon), \dots, p_m(\varepsilon)]$ is a maximal dimensional cell of $\tilde{\mathcal{T}}_n^\varepsilon$, given by the convex hull of $\{p_0(\varepsilon), \dots, p_m(\varepsilon)\}$ iff $B \equiv [p_0, \dots, p_m]$ is a maximal dimensional cell in $\tilde{\mathcal{T}}_n$ given by the convex hull of $\{p_0, \dots, p_m\}$.

Lemma 4.19. *Suppose $\tilde{\mathcal{T}}_n$ is a regular triangulation. There exists $\varepsilon > 0$ sufficiently small such that any triangulation $\tilde{\mathcal{T}}_n^\varepsilon$ is regular.*

Proof. Take $\{p_0, \dots, p_m\}$ to be the vertex set of a maximal dimensional cell B of $\tilde{\mathcal{T}}_n$ and let $p \in T_n \setminus B$. Since $\tilde{\mathcal{T}}_n$ is regular, let its associated height function be denoted by h . Then $(p, h(p))$ lies above the hyperplane defined by $(p_t, h(p_t))_{t=0}^m$. Therefore, substituting p_t for $p_t(\varepsilon)$ and letting $\varepsilon > 0$ be sufficiently small, we have that $(p_t(\varepsilon), h(p_t(\varepsilon)))_{t=0}^m$ defines a non-vertical hyperplane in \mathbb{R}^{m+2} with $(p(\varepsilon), h(p(\varepsilon)))$ above this hyperplane, for any other vertex $p(\varepsilon)$. From this, we can define the graph of a convex piecewise linear function $h^\varepsilon : \Delta(\tilde{S}_{n,0}) \rightarrow \mathbb{R}$, by letting $h^\varepsilon(p) \equiv h(p)$ for every vertex p of $\mathcal{T}_n^\varepsilon$, and, for another arbitrary point p , we consider the carrier of p in $\tilde{\mathcal{T}}_n^\varepsilon$ and define h^ε by linear interpolation. \square

We now define a *generalized barycentric subdivision* of the simplex $\Delta(\tilde{S}_{n,0})$, which is a minor generalization of the classic barycentric subdivision. Let $\tilde{\mathcal{T}}_n^0$ be the face complex of $\Delta(\tilde{S}_{n,0})$. Assume $\tilde{\mathcal{T}}_n^{k-1}$ is defined and let us define $\tilde{\mathcal{T}}_n^k$: the 0-dimensional *generalized barycenters* are the vertices of $\tilde{\mathcal{T}}_n^{k-1}$. For each 1-dimensional cell $\tilde{\tau}$ of $\tilde{\mathcal{T}}_n^{k-1}$, chose exactly one point b^1 in its interior. The point b^1 is referred to as a 1-dimensional *generalized barycenter* of $\tilde{\tau}$. One then proceeds by choosing points in the interior of cells of increasing dimension, in the exact same fashion as with the classical barycentric subdivision.

The full-dimensional cells of $\tilde{\mathcal{T}}_n^k$ are also defined in the exact same fashion as in the k -th-iterate of the classical barycentric subdivision: an m -dimensional cell of $\tilde{\mathcal{T}}_n^k$ is denoted by its vertices $(b_k^0, b_k^1, \dots, b_k^m)$, where b_k^0 is a vertex of $\tilde{\mathcal{T}}_n^{k-1}$, b_k^1 a generalized barycenter of a 1-dimensional cell of

$\tilde{\mathcal{T}}_n^{k-1}$, b_k^2 a generalized barycenter of a 2-dimensional cell of $\tilde{\mathcal{T}}_n^{k-1}$, etc. When generalized barycenters are chosen sufficiently near the actual barycenters, the diameter of a generalized barycentric subdivision shrinks, as k increases, just like with the classical barycentric subdivision.

Letting $\zeta > 0$ be as in the statement of Lemma 4.15, there exists $\delta > 0$ sufficiently small and $k_0 \in \mathbb{N}$ such that for any $k \geq k_0$, any choice of generalized barycenters within δ of the actual barycenters yields a generalized barycentric subdivision $\tilde{\mathcal{T}}_n^k$ with diameter less than ζ . Moreover, we can assume without loss of generality that for each i, j , the carrier of $\tilde{\sigma}_n^{ij}$ in $\tilde{\mathcal{T}}_n^k$ has all its vertices in the interior of \tilde{T}_n^{ij} . We fix from now on such a $\delta > 0$ and k_0 .

We now define a polyhedral refinement of $\tilde{\mathcal{T}}_n^k$, which we call the *Eaves-Lemke (EL)-refinement* \mathcal{P}_n of $\tilde{\mathcal{T}}_n^k$. This is precisely the same polyhedral refinement as the one used in [4], we recall its construction for completeness. For each $(m-1)$ -dimensional cell τ of $\tilde{\mathcal{T}}_n^k$, let $H_\tau \equiv \{x \in \mathbb{R}^{m+1} \mid a_\tau \cdot x = b_\tau\}$ be the hyperplane that includes τ and is orthogonal to $\Delta(\tilde{S}_{n,0})$. Each full-dimensional cell of \mathcal{P}_n is the intersection of $\Delta(\tilde{S}_{n,0})$ with a polyhedron $\cap_\tau H_\tau^i$, where $i \in \{+, -\}$ and H_τ^i is one of the two half-spaces defined by the hyperplane H_τ .

Lemma 4.20. *There exists a polyhedral subdivision \mathcal{P}_n of $\Delta(\tilde{S}_{n,0})$ satisfying the following properties:*

- (1) *for each i, j , $\tilde{\sigma}_n^{ij}$ has a carrier ϱ^{ij} in \mathcal{P}_n of dimension $\dim(\Sigma_n)$;*
- (2) *there is a convex piecewise linear function $\gamma_n : \Delta(\tilde{S}_{n,0}) \rightarrow [0, 1]$ which is linear precisely in each cell of \mathcal{P}_n ;*
- (3) *the diameter of \mathcal{P}_n is less than ζ .*

Proof. We first focus on constructing a polyhedral subdivision that satisfies (1). In $\tilde{\mathcal{T}}_n^0$, choose generalized barycenters which are distinct from $\tilde{\sigma}_n^{ij}$, for each i, j , defining $\tilde{\mathcal{T}}_n^1$. Consider now the derived (EL)-polyhedral subdivision \mathcal{P}_n^1 of $\tilde{\mathcal{T}}_n^1$. If there exists i, j such that the carrier of $\tilde{\sigma}_n^{ij}$ in \mathcal{P}_n^1 has dimension less than $\dim(\Sigma_n)$, then there exists a hyperplane H_τ which contains $\tilde{\sigma}_n^{ij}$, where τ is a face containing an m -dimensional generalized barycenter b_1^m . Let $\{p_0, p_1, \dots, p_{m-2}, b_1^m\}$ be affinely independent vertices of the triangulation $\tilde{\mathcal{T}}_n^1$ defining $H_\tau \cap \Delta(\tilde{S}_{n,0})$. If H_τ contains in addition $\tilde{\sigma}_n^{ij}$, then $\{p_0, p_1, \dots, p_{m-2}, b_1^m, \tilde{\sigma}_n^{ij}\}$ is an affinely dependent set. Choosing therefore b_1^m outside an affine set with dimension strictly lower than m in $\Delta(\tilde{S}_{n,0})$ implies H_τ does not contain $\tilde{\sigma}_n^{ij}$. This procedure can be iteratively applied for each point $\tilde{\sigma}_n^{ij}$, ensuring all of them have a carrier of dimension $\dim(\Sigma_n)$ in \mathcal{P}_n^1 . Suppose now we have obtained $\mathcal{P}_n^{k_0-1}$ from $\tilde{\mathcal{T}}_n^{k_0-1}$ satisfying the property that all $\tilde{\sigma}_n^{ij}$ have carriers with dimension equal to $\dim(\Sigma_n)$. Choose generalized barycenters in $\tilde{\mathcal{T}}_n^{k_0-1}$ such that each generalized barycenter is distinct from each $\tilde{\sigma}_n^{ij}$ and consider the (EL)-polyhedral subdivision $\mathcal{P}_n^{k_0}$. If there exists i, j , such that the carrier of $\tilde{\sigma}_n^{ij}$ in $\mathcal{P}_n^{k_0}$ has dimension less than $\dim(\Sigma_n)$, then it follows that there exists τ containing an m -dimensional generalized barycenter $b_{k_0}^m$ such that H_τ contains $\tilde{\sigma}_n^{ij}$. Similarly as before, let $\{p_0, p_1, \dots, p_{m-2}, b_{k_0}^m\}$

be affinely independent vertices of the triangulation $\tilde{\mathcal{T}}_n^{k_0}$ defining $H_\tau \cap \Delta(\tilde{S}_{n,0})$. If H_τ contains in addition $\tilde{\sigma}_n^{ij}$, then $\{p_0, p_1, \dots, p_{m-2}, b_{k_0}^m, \tilde{\sigma}_n^{ij}\}$ is an affinely dependent set. Choosing therefore $b_{k_0}^m$ outside an affine set with dimension strictly lower than m in $\Delta(\tilde{S}_{n,0})$ implies H_τ does not contain $\tilde{\sigma}_n^{ij}$. This procedure can be iteratively applied for each point $\tilde{\sigma}_n^{ij}$, ensuring all of them have a carrier of dimension $\dim(\Sigma_n)$ in $\mathcal{P}_n^{k_0}$. This shows $\mathcal{P}_n^{k_0}$ satisfies (1). Since $\mathcal{P}_n^{k_0}$ refines $\tilde{\mathcal{T}}_n^{k_0}$, its diameter is also less than ζ , implying (3). To show (2), consider the following function: $\gamma_n : \Delta(\tilde{S}_{n,0}) \rightarrow [0, 1]$, $\gamma_n(\sigma_n) \equiv \alpha \sum_\tau |a_\tau \cdot \sigma_n - b_\tau|$, where the scaling factor $\alpha > 0$ is chosen so as to let $\gamma_n(\Delta(\tilde{S}_{n,0})) \subseteq [0, 1]$. The function γ_n then satisfies (2). \square

We now derive a triangulation from \mathcal{P}_n in Lemma 4.20, by triangulating each full-dimensional cell of \mathcal{P}_n . This triangulation will add no vertices beyond those of \mathcal{P}_n , i.e., it will simply subdivide each m -dimensional cell of \mathcal{P}_n into m -dimensional simplices, without adding new vertices to the triangulation. This triangulation, up to an arbitrarily small perturbation of its vertices, achieves the objectives of Lemma 4.15.

Lemma 4.21. *There exists a triangulation $\tilde{\mathcal{T}}_n$, which refines \mathcal{P}_n without adding new vertices, satisfying the following conditions:*

- (1) *The diameter of $\tilde{\mathcal{T}}_n$ is less than ζ ;*
- (2) *For i, j , \tilde{T}_n^{ij} is the space of a subcomplex of $\tilde{\mathcal{T}}_n$;*
- (3) *For i, j , the carrier \hat{T}_n^{ij} of $\tilde{\sigma}_n^{ij}$ in $\tilde{\mathcal{T}}_n$ has dimension $\dim(\Sigma_n)$;*
- (4) *There is a convex piecewise linear function $\gamma_n : \Delta(\tilde{S}_{n,0}) \rightarrow \mathbb{R}_+$ such that:*
 - (4.1) *γ_n is linear precisely on the simplices of $\tilde{\mathcal{T}}_n$;*
 - (4.2) *γ_n is constant in each \hat{T}_n^{ij} ;*
- (5) *For each i, j , $(\tilde{\sigma}_n^{ij}, \tilde{\sigma}_{n+1}^{ij})$ is not in the convex hull of $\dim(\Sigma_n) + \dim(\Sigma_{n+1})$ (or less) vertices of $\hat{T}_n^{ij} \times \hat{T}_{n+1}^{ij}$.*

Proof. The polyhedral complex \mathcal{P}_n refines the generalized barycentric subdivision $\tilde{\mathcal{T}}_n^{k_0}$. Therefore it also has diameter less than ζ and has \tilde{T}_n^{ij} as the space of a subcomplex. There exists a piecewise-linear, convex function $\hat{\gamma}_n : \Delta(\tilde{S}_{n,0}) \rightarrow \mathbb{R}_+$, which is linear precisely in the cells of \mathcal{P}_n (cf. Lemma 4.20). The map $\hat{\gamma}_n(\cdot)$ is in particular a height function in the set of vertices of $\tilde{\mathcal{T}}_n^k$. By Lemma 2.3.15 in [2], any sufficiently small (generic) perturbation γ'_n of this height function gives a refinement $\tilde{\mathcal{T}}_n$ of \mathcal{P}_n (without adding vertices to \mathcal{P}_n), where $\tilde{\mathcal{T}}_n$ is a triangulation, yielding a convex piecewise linear function $\gamma'_n : \Delta(\tilde{S}_{n,0}) \rightarrow \mathbb{R}_+$ which is linear precisely in the cells of $\tilde{\mathcal{T}}_n$.

We show that it is without loss of generality to assume that the carrier of $\tilde{\sigma}_n^{ij}$ in $\tilde{\mathcal{T}}_n$ is a full-dimensional cell of $\tilde{\mathcal{T}}_n$. If there exists i, j such that $\tilde{\sigma}_n^{ij}$ is contained in an $(\dim(\Sigma_n) - 1)$ -dimensional cell τ of $\tilde{\mathcal{T}}_n$, using Lemma 4.19, we can consider a sufficiently small $\varepsilon > 0$ and a $\tilde{\mathcal{T}}_n^\varepsilon$ in which the carrier of $\tilde{\sigma}_n^{ij}$ in $\tilde{\mathcal{T}}_n^\varepsilon$ has dimension $\dim(\Sigma_n)$: this is done by choosing suitable perturbations $p(\varepsilon)$ of vertices p of the carrier of $\tilde{\sigma}_n^{ij}$ in $\tilde{\mathcal{T}}_n$, yielding the desired $\tilde{\mathcal{T}}_n^\varepsilon$. Moreover, for a sufficiently small ε ,

$\tilde{\mathcal{T}}_n^\varepsilon$ satisfies (1), (2), (3) and (4.1). Therefore, if necessary by considering a small perturbation $\tilde{\mathcal{T}}_n^\varepsilon$, we can assume that $\tilde{\mathcal{T}}_n$ satisfies (1), (2), (3) and (4.1).

Note now that the convex hull of any collection of $\dim(\Sigma_n) + \dim(\Sigma_{n+1})$ (or less) vertices of $\hat{T}_n^{ij} \times \hat{T}_{n+1}^{ij}$ is a union of convex sets whose dimensions are strictly less $\dim(\Sigma_n) + \dim(\Sigma_{n+1})$. This implies that if the vertices of $\hat{T}_n^{ij} \times \hat{T}_{n+1}^{ij}$ are positioned in \tilde{T}_n^{ij} outside a set of dimension strictly less than $\dim(\Sigma_n) + \dim(\Sigma_{n+1})$, then $\tilde{\sigma}_n^{ij}$ satisfies property (5). By using Lemma 4.19 just like in the previous paragraph, we can then assume $\tilde{\mathcal{T}}_n$ satisfy (5).

To finish our construction it remains to prove that we can construct a piecewise linear convex function $\gamma_n : \Delta(\tilde{S}_{n,0}) \rightarrow \mathbb{R}_+$, linear precisely on each cell of $\tilde{\mathcal{T}}_n$ and constant over each \hat{T}_n^{ij} . Since $K_n[\dim(\Sigma_n) + 1] < m + 1$, the union of the set of vertices of \hat{T}_n^{ij} over i, j spans an affine set with dimension lower than m . This implies we can define an affine function $\phi_n : \Delta(\tilde{S}_{n,0}) \rightarrow \mathbb{R}$ satisfying the following conditions: (a) $\phi_n(\hat{p}_n^{ij}) = \hat{\gamma}_n(\hat{p}_n^{ij})$, for all vertices \hat{p}_n^{ij} of \hat{T}_n^{ij} ; (b) $\phi_n(p_n) < \hat{\gamma}_n(p_n)$ for all remaining vertices p_n of $\tilde{\mathcal{T}}_n$. Letting now $\gamma_n \equiv \gamma'_n - \phi_n$ yields the desired function satisfying all the required properties in the statement of the Lemma. \square

REFERENCES

- [1] Bochnak, J., M. Coste, and M.-F. Roy (1998): *Real Algebraic Geometry*. Berlin: Springer Verlag.
- [2] De Loera, J.A., J. Rambau, and F. Santos (2010): *Triangulations: Structures for Algorithms and Applications*. Berlin: Springer Verlag.
- [3] DeMichelis, S. and F. Germano (2020): “On the Indices of Zeros of Nash Fields,” *Journal of Economic Theory*, 94, 192-217.
- [4] Govindan, S., and R. Wilson (2005): “Essential Equilibria,” *Proceedings of the National Academy of Sciences*, 102, 15706–11.
- [5] Govindan, S., and L. Pahl (2024): “O’Neill’s Theorem for PL-Approximations,” *J. Fixed Point Theory Appl.* 26, 31.
- [6] Govindan, S., R. Laraki and L. Pahl (2022): “On Sustainable Equilibria” *Journal of Economic Theory*, 213, 105736.
- [7] Hauk, E., and S. Hurkens (2002): “On Forward Induction and Evolutionary and Strategic Stability”, *Journal of Economic Theory*, 106, 66-90.
- [8] Hofbauer, J. (2003): “Some Thoughts on Sustainable/Learnable Equilibria,” in *Conference Book of the XV Italian Meeting of Game Theory and Applications*. Urbino, IT: Università Degli Studi di Urbino.
- [9] Hofbauer, J., C. Pawlowitsch. (2023): “The evolutionary dynamics of costly signaling,” *Mimeo*.
- [10] Kohlberg, E., J.-F. Mertens, (1986): “On the Strategic Stability of Equilibria,” *Econometrica*, 54, 5, 1003–7.
- [11] McLennan, A. (1989): “Fixed Points of Contractible Valued Correspondences,” *International Journal of Game Theory*, 18, 175-184.
- [12] McLennan, A. (2016): “The +1 Index Principle,” *Mimeo*.
- [13] McLennan, A. (2018): *Advanced Fixed Point Theory for Economics*. Berlin: Springer Verlag.
- [14] Mertens, J.-F. (1989): “Stable Equilibria—A Reformulation, Part I, Definition and Basic Properties”, *Mathematics of Operations Research*, 14, pp. 575-625.
- [15] Mertens, J.-F. (2004): “Ordinality in non cooperative games” , *International Journal of Game Theory*, 32, pp. 387-430.
- [16] Myerson, R. (1978): “Refinements of the Nash equilibrium concept”, *International Journal of Game Theory*, 7, 73-83.

- [17] O'Neill, B. (1953): "Essential sets and fixed points," *American Journal of Mathematics*, 75, 497–509.
- [18] Pahl, L. (2022): "Polytope-form Games and Index/Degree theories for Extensive-form Games," *Games and Economic Behavior*, 141, pp. 444-471.
- [19] Ritzberger, K. (1994): "The theory of normal form games from the differentiable viewpoint," *International Journal of Game Theory*, 23, pp. 207-236.
- [20] Rourke, C.P., Sanderson, B.J. (2021): *Piecewise-Linear Topology*. Berlin: Springer Verlag.
- [21] Selten, R. (1975): "Re-examination of the Perfectness Concept for Equilibrium Points in Extensive Games", *International Journal of Game Theory*, 4, pp. 25-55.
- [22] Spanier, E. H., (1966): *Algebraic Topology*, Springer-Verlag, Berlin.
- [23] Wu, W-T., and J-H. Jiang, (1962): "Essential equilibrium points of n -person non-cooperative games", *Sci. Sinica*, 11, pp. 1307–22.
- [24] Zeeman, E.C., "Relative Simplicial Approximation," *Proc. Camb. Phil. Soc.*, 60, 39.