O'NEILL'S THEOREM FOR PL-APPROXIMATIONS

SRIHARI GOVINDAN AND LUCAS PAHL

Abstract. We present a version of O'Neill's Theorem (Theorem 5.2 in [6]) for piecewise linear approximations.

1. Introduction

Theorem 5.2 of [6] asserts that if f is a continuous function from a topological polyhedron to itself, C is a component of the set of fixed points of f, U is a Euclidean neighborhood of C containing no other fixed points of f, $r_1, \ldots, r_k$ are integers whose sum is the fixed point index of C, and $x_1, \ldots, x_k$ are distinct points of C, then there is a map arbitrarily close to f whose fixed points in U are $x_1, \ldots, x_k$, with the fixed point index of each x_i being r_i. This note establishes a version of this result in the PL category. Specifically: (i) we allow for the polyhedron to be a subset of a topological manifold, and not homeomorphic to an Euclidean neighborhood; (ii) we weaken the restriction that the component C be in the interior of the polyhedron and, consequently, have to allow for the x_i's to be arbitrarily close to it; (iii) we add the restriction that the manifold be the space of a simplicial complex and that the approximating function be piecewise linear; (iv) in order to obtain a regularity property for fixed points, we insist that they be interior points—barycenters, even—of full-dimensional simplices and that the displacement map of the approximating function be a homeomorphism locally around these fixed points, if the r_i's are ± 1.

Our interest in this problem was motivated by its intended use in game theory. Nash equilibria of games obtain as fixed points of self maps on strategy spaces. It is a frequent (and robust) feature of games that components of equilibria lie on the boundary of the strategy space, which prompts the weakening of O'Neill's condition sub (ii) above. Also, fixed point problems arising from games have a special structure, since the payoff functions of games are multilinear. Hence, perturbations of a given fixed point map associated with a game have to satisfy certain conditions if they are to be associated with fixed points of games, prompting us to investigate a multilinear version of O'Neill's theorem for games (See [2] for details.) This paper presents a linear version of the problem, where a stronger result is possible, and is possibly of wider interest as well.

Date: January 9, 2024.

We are grateful to Andy McLennan for suggesting this problem to us, and to Robert Brown for his comments. Lucas Pahl acknowledges financial support from the Hausdorff Center for Mathematics (DFG project no. 390685813).

2. Statement of The Theorem

We first set a few notational conventions and recall some definitions that will be required for the statement of the main theorem (Theorem 2.1) and its proof.

2.1. **(Notational) Conventions.** For $\zeta > 0$, define $B_\zeta(x)$ to be the ball around x with radius ζ. The symbol id_X denotes the identity map on the set X. Given $A \subseteq \mathbb{R}^n$ and a map $f : A \to \mathbb{R}^n$, $d_f(x) \equiv x - f(x)$. Let $X \subset \mathbb{R}^n$ be compact, and $f, g : X \to \mathbb{R}^n$ two continuous maps, we denote $\|f - g\| \equiv \sup_{x \in X} \|f(x) - g(x)\|_p$, where $\| \cdot \|_p$ denotes the ℓ_p-norm in $\mathbb{R}^n$. Unless explicitly stated otherwise, we will assume that $p = 2$ and will omit the subscript p for notational convenience. If $C \subseteq \mathbb{R}^n, x \in \mathbb{R}^n$, let $d(x, C) \equiv \inf_{y \in C} \|x - y\|$.

2.2. **Triangulations, Polyhedra and Pseudomanifolds.** Our terminology and notation for polyhedral complexes is mostly standard. In particular, we follow the convention of piecewise linear topology according to which a map from $X \subset \mathbb{R}^m$ to $\mathbb{R}^n$ is *linear* if it is the restriction to X of a map that is affine in the sense of linear algebra, i.e., the composition of a linear transformation and a translation.

As always, a *polytope* $P \subset \mathbb{R}^m$ is the convex hull of a finite set of points; an equivalent definition is that a polytope is an intersection of finitely many closed half-spaces that happens to be bounded, hence compact. The *dimension* of P is the dimension of its affine hull. The *faces* of P are P, the empty face, and the intersections of P with the boundaries of closed half-spaces that contain P; faces other than P are *proper*. A (finite, bounded) *polyhedral complex* Z in $\mathbb{R}^m$ is a finite collection of polytopes that contains each face of each of its elements, such that the intersection of any two of its elements is a face of both. If Y is a subset of Z that contains each of the faces of each of its elements, then Y is a *subcomplex* of Z. For $n = 0, \dots, m$, let Z^n be the set of n-dimensional elements of Z. Elements of Z^0 are *vertices* of Z. The *dimension* of Z is the largest n such that $Z^n \neq \emptyset$. The *mesh* of Z is the maximum of the diameters of the elements of Z. The *space* of Z is $|Z| = \bigcup_{P \in Z} P$. A set $P \subset \mathbb{R}^m$ is a *polyhedron* if it is the space of a polyhedral complex, and its *dimension* is the dimension of any such complex.

A *simplicial complex* S in $\mathbb{R}^m$ is a polyhedral complex whose elements are all simplices. We say that S is a *triangulation* of $|S|$. The *carrier* $\Delta(x)$ of $x \in |S|$ in S is the smallest element of S that contains x, so it is the unique element of S whose interior contains x. If Z is a simplicial complex, we say that Y is a *subdivision* of Z if Y is a simplicial complex with $|Y| = |Z|$, and every simplex of Z is the union of simplices of Y.

When X is the space of a subcomplex of Z, we write $Z(X)$ to denote the subcomplex of Z composed by the simplices of Z which are contained in X.

If S, T are simplicial complexes, a function $f\colon |S| \to |T|$ is *simplicial* (relative to the triangulations S and T) if, for each $\sigma \in S$, there is a $\tau \in T$ such that f maps each vertex of σ to a vertex of τ and the restriction of f to σ is linear. If $P \subset \mathbb{R}^m$ and $Q \subset \mathbb{R}^\ell$ are polyhedra, a function $f\colon P \to Q$

is *piecewise linear* (PL) if there are simplicial subdivisions S of P and T of Q with respect to which f is simplicial. A sufficient condition for this (Theorem 2.14 of [7]) is that there is a simplicial subdivision S of P such that the restriction of f to each $\sigma \in S$ is linear.

A *polyhedron of homogeneous dimension* n is a polyhedron P that is the union of finitely many n-dimensional simplices, provided that the intersection of any two of the n-dimensional simplices is a (possibly empty) common face of both. The collection of the n-dimensional simplices together with all their faces then constitute a triangulation of P. If T is a triangulation of P, then ∂P is the union of those $\tau \in T^{n-1}$ that are a face of exactly one $\sigma \in T^n$; evidently ∂P is a polyhedron of homogeneous dimension $n-1$.

A polyhedron P of homogeneous dimension n is an *n-pseudomanifold*, provided the following hold for some triangulation T of P:

(1) Every element of T^{n-1} is a face of at most two elements of T^n;
(2) For any two n-simplices $\sigma, \sigma' \in T$ there is a finite chain $\sigma = \sigma_1, \ldots, \sigma_k = \sigma'$ of simplices in T^n such that $\sigma_i \cap \sigma_{i+1} \in T^{n-1}$.

2.3. **Statement of The Result.** Let $(Y, \partial Y)$ be a topological n-manifold with ∂Y denoting its boundary and assume $Y \subseteq \mathbb{R}^m$ for some finite $m > 0$. Let $(X, \partial X)$ be an n-pseudomanifold with boundary ∂X with $X \subseteq Y$. Suppose Y is a polyhedron of homogenous dimension n with triangulation T, and X is the space of a subcomplex of Y of homogenous dimension n, as well. We can assume without loss of generality that $m \geqslant n+1$, by embedding Y in a Euclidean space of dimension larger than n, when $m = n$. Let $S \equiv T(X)$. Let $f : X \to Y$ be a continuous function satisfying the following assumptions: A) either f has no fixed points on the boundary of X in Y, or $f(X) \subseteq X$; B) the map f has a unique connected component of fixed points (cf. Remark 2.3 for a generalization). Thanks to assumption A) about f, C has a well-defined index, call it c. Let U be a neighborhood of C in X with closure denoted $\bar{U}$.

Theorem 2.1. *For every $\varepsilon_0 > 0$, there exists $\delta_0 > 0$ such that for each $0 < \delta \leqslant \delta_0$ and each finite collection of points $x_1, \ldots, x_k$ and integers $r_1, \ldots, r_k$ such that: (a) for each $1 \leqslant i \leqslant k$, x_i belongs to the interior of an n-simplex of S, and $d(x_i, C) < \delta$, and (b) $\sum_i r_i = c$, there exist subdivisions S^* and T^* of S and T, resp., and a simplicial map $h^* : |S^*| \to |T^*|$ such that:*

(1) *$\|f - h^*\| < \varepsilon_0$;*
(2) *$h^*(X) \subseteq X$, if $f(X) \subseteq X$;*
(3) *the only fixed points of h^* in $\bar{U}$ are the x_i's, and the index of each x_i is r_i;*
(4) *for each i such that $r_i \in \{-1, +1\}$, there exist simplices $\sigma_i \in S^*$ and $\tau_i \in T^*$ such that:*
    (a) *$\sigma_i \subset \tau_i$ and x_i is the barycenter of both σ_i and τ_i;*
    (b) *h^* maps σ_i homeomorphically onto τ_i.*

Remark 2.2. The triangulation T^* can be chosen such that X is the space of a subcomplex of T^*. Also, S^* can be chosen such that outside of a neighborhood of the x_i's, it subdivides the triangulation $T^*(X)$. Apparently, we are unable to get the stronger condition that S^* subdivides the triangulation induced by T^*.

Remark 2.3. If the map f has finitely many connected components of fixed points $C_1, ..., C_k$ (for example, if f is semialgebraic), the proof of Theorem 2.1 applies with insignificant modifications in order to obtain a simplicial approximation g of f where the result stated in Theorem 2.1 holds for each C_i.

Remark 2.4. When comparing Theorem 2.1 with Theoerem 5.2 in [6], our statement, ignoring the PL structure and applying it to triangulable manifolds, provides a couple of generalizations. First, Theorem 2.1 allows for fixed-point components to intersect the boundary of X in Y, whereas in O'Neill, a fixed point component is located in the interior of the pseudomanifold X. Second, we allow for a pseudomanifold X that is the subset of a topological manifold of the same dimension as X and contained in a Euclidean space, while O'Neill requires X to be homeomorphic to a Euclidean neighborhood. When the first case occurs, then $f(X) \subseteq X$, by our assumption on f, and the index is well-defined (explicitly, by the trace formula of O'Neill).

3. Auxiliary Results

Lemma 3.1. *Let τ be a n-simplex in $\mathbb{R}^m$ with barycenter x and let c be an integer. There exists an n-simplex $\sigma \subset \tau$ with x as a barycenter and a PL map $h : \sigma \to \tau$ such that x is the unique fixed point of h and its index is c. Furthermore, if $c \in \{-1, +1\}$, h can be chosen to be an affine homeomorphism.*

Proof. Consider first the case where $|c| \neq 1$. We can assume without loss of generality that $m = n$ and $x = 0$. Take $\delta > 0$ such that ℓ_1-distance between 0 and $\partial\tau$ is greater than 2δ. Letting $B \subset \tau$ be the ℓ_1-ball of radius δ around 0, it is sufficient to construct a PL function $h : B \to \tau$ such that 0 is the unique fixed point of h and its index is c. We can further reduce the problem to the case $n = 2$: intersect τ (and B) with the linear subspace H of $\mathbb{R}^n$ consisting of points where the last $n - 2$ coordinates are zero. If we have a PL function $h : H \cap B \to H \cap \tau$ where the index of 0 is c, we can extend it to B by composing it with the projection from B to $H \cap B$. The point 0 still has index c under the extension.

By the choice of δ, the problem is solved if we can find a PL function $d : B \to B$—to serve as the displacement of h— such that 0 is the only zero of d and has degree c. The case $c = 0$ is obvious: map 0 to 0, the boundary of B to some constant on ∂B and all other points by linear interpolation. Fix now c such that $|c| > 1$. The ℓ_1-ball B can be triangulated as the union of four triangles (one in each orthant). Subdivide each of the triangles into $|c|$ triangles all of which having 0 as a vertex. There now exists a PL map from B to itself that sends each of the $4c$ triangles of the subdivision to one of the triangles of B and that has degree c.

For the case $|c| = 1$, the lemma requires h to be an affine homeomorphism, so we approach the problem slightly differently. Let $w_0, \dots, w_n$ be the vertices of τ. Take a simplex $\sigma \subset \tau$ of diameter less than δ, that has x as the barycenter, and is such that, letting $v_0, \ldots, v_n$ be the vertex set of σ, there is $\lambda > 1$ for which $w_i = x + \lambda(v_i - x)$ for all i.

For any permutation $\pi : \{0, \ldots, n\} \to \{0, \ldots, n\}$ we can define an affine homeomorphism $f^\pi : \sigma \to \tau$ that sends v_i to $w_{\pi(i)}$. Obviously x is the only fixed point of f^π. By virtue of the assumptions on σ, there is a retraction $r : \tau \to \sigma$ that sends w_i to v_i for each i, and that is affine on each face of τ. For a permutation π, x is also an isolated fixed point under $f^\pi \circ r$ and its index is the same under f^π and $f^\pi \circ r$.

Suppose π is a cyclic permutation where the only cycle involves all $n+1$ elements. Then the index of x under f^π is $+1$ as under $f^\pi \circ r$ it is the unique fixed point. To obtain a fixed point of index -1, consider a permutation π that leaves, say, 0 fixed, and is cyclic on the others. Under the map $f^\pi \circ r$, there are three fixed points, w_0, x, and the barycenter of the face opposite w_0. The index of the first and the last fixed points is $+1$, assigning x an index of -1. □

Lemma 3.2. *Let $\hat{T}$ be a triangulation of Y. Let $\{x_i\}_{i=1}^k$ be a subset of Y, with each x_i contained in the interior of a simplex $\tau_i \in \hat{T}^n$. For each $\delta > 0$, there exists a triangulation $\tilde{T}$ of Y that subdivides $\hat{T}$ and satisfies the following:*

1. *The mesh of $\tilde{T}$ is less than δ;*
2. *For each $i = 1, ..., k$, there exist n-simplices $\sigma_i \in \tilde{T}$ and $\tau_i \in \hat{T}$ with $\sigma_i \subset \tau_i$, x_i the barycenter of σ_i.*

Proof. For each $i = 1, ..., k$, consider an n-simplex $\sigma_i \subset \mathrm{int}(\tau_i)$ with diameter less than δ that has x_i as a barycenter. For each i, take a polyhedral subdivision P_i of τ_i that has σ_i as an n-dimensional polyhedron, without introducing new vertices in τ_i beyond those of σ_i and τ_i. There exists a triangulation $\hat{T}_i'$ of τ_i which subdivides P_i, without introducing new vertices (cf. Proposition 2.9 in [7]). The simplices of the triangulation $\hat{T}_i'$, for each i, together with the other simplices of the triangulation $\hat{T}$, form a triangulation $\hat{T}'$ of Y. Now iterating sufficiently many times the barycentric subdivision of $\hat{T}'$ modulo $\cup_i \sigma_i$, (cf. [11]), we obtain a triangulation $\tilde{T}$ that subdivides $\hat{T}'$ and has mesh less than δ as well. The triangulation $\tilde{T}$ satisfies both requirements of the lemma. □

Lemma 3.3. *Let $\hat{T}$ be a triangulation of Y and let $\sigma \in \hat{T}^n$. Let $\hat{S}$ be any triangulation of σ. There exists a triangulation $\tilde{T}$ of Y that subdivides $\hat{T}$ such that $\tilde{T}(\sigma) = \hat{S}$ and the simplices of $\hat{T}$ that are disjoint from σ are simplices of $\tilde{T}$.*

Proof. Let $\hat{\mathcal{S}}$ be the collection of simplices in $\hat{S}$ that are contained in maximal proper faces of σ. Let $\hat{\mathcal{T}}$ be the collection of simplices in $\hat{T}$ that intersect σ but are not contained in σ. Let $\hat{\mathcal{T}}_0 = \{\, \tau \in \hat{T} \mid \tau \cap \sigma = \emptyset, \tau \subset \varrho \in \hat{\mathcal{T}} \,\}$. Let $f \in \hat{\mathcal{S}}$ and assume $\varrho \in \hat{\mathcal{T}}$ contains f. The convex closure of f with any simplex in $\mathcal{T}_0 \cap \varrho$ is a simplex. Taking the convex closure of simplices in $\varrho \cap \hat{S}$ and in $\hat{\mathcal{T}}_0 \cap \varrho$ produces a triangulation of ϱ, which adds no vertices to the faces ϱ that are not contained in σ. The simplices of $\hat{S}$, the simplices obtained by the triangulation just defined in the simplices of $\hat{\mathcal{T}}$ and simplices of the triangulation $\hat{T}$ which do not intersect σ, define the triangulation $\tilde{T}$ of the statement. $\square$

We say the triangulation $\tilde{T}$ from Lemma 3.3 *extends* the triangulation $\hat{S}$ from σ to Y.

Definition 3.4. A *fiber bundle (with fiber F)* is a triple (E, B, F, p) where:

(1) $p : E \to B$ is a continuous surjective map from the *total space* E to the *base space* B;
(2) For each $x \in B$, there exists a neighborhood $U \subseteq B$ of x such that $h_x : p^{-1}(U) \to U \times F$ is a homeomoprhism that satisfies $p = \mathrm{p}_1 \circ h_x$, where p_1 is the projection over the first coordinate.

Two fiber bundles $(\bar{E}, \bar{B}, \bar{F}, \bar{p})$ and (E, B, F, p) are *isomorphic* if there exist homeomorphisms $\bar{h} : \bar{E} \to E$ and $h : \bar{B} \to B$ such that $h \circ \bar{p} = p \circ \bar{h}$. The fiber bundle (E, B, F, p) is *trivial* if $E = B \times F$ and p is the projection over the first coordinate. For notational convenience, we will say that a fiber bundle is *trivial* if it is isomorphic to a trivial bundle.

Definition 3.5. A *n-microbundle* over the base space B is a triple (E, B, e, p) where $e : B \to E$ and $p : E \to B$ are continuous maps such that:

(1) $p \circ e = id_B$;
(2) For every $b \in B$, there are a neighborhood $U \subseteq B$ of b and a neighborhood $V \subseteq E$ of $e(b)$ such that $e(U) \subseteq V$, $p(V) \subseteq U$ and $h_V : V \to U \times B_1^n(0)$ a homeomorphism satisfying: (i) $p_1 \circ h_V = p|_V$, and (ii) $h \circ e|_U = i$, where $i : B \to B \times B_1^n(0)$, $i(b) \equiv (b, 0)$ and p_1 is the projection over the first coordinate.

Let $Y^* = Y \sqcup_{\partial Y} Y$ be the compact, connected, n-dimensional, boundaryless topological manifold containing Y, obtained by attaching Y with itself along its boundary. Let p_1 be the natural projection from $Y^* \times Y^*$ to its first factor. Let $\Delta = \{ (y, y) \in Y^* \times Y^* \}$. Let $D : Y^* \to Y^* \times Y^*$ be the diagonal map, which sends $x \in Y^*$ to $(x, x) \in \Delta$. For each $\delta > 0$, let $B_\delta(\Delta)$ be the the set of $(x, y) \in Y^* \times Y^*$ such that $\|x - y\| \leqslant \delta$. Let $B_1^n(0)$ be the unit ball of $\mathbb{R}^n$. Given open sets V in $Y^* \times Y^*$ and U in Y we say that a homeomorphism $h : V \to U \times B_1^n(0)$ is *trivializing* for D if $h \circ D(x) = (x, 0)$. We say h is *trivializing* for p_1 if $p_1 = q_1 \circ h$, where $q_1 : Y^* \times B_1^n(0) \to Y^*$ is the projection over the first coordinate.

The n-microbundle $(Y^* \times Y^*, Y^*, D, p_1)$ is called the *tangent microbundle of Y^** (see Example (iii) in Chapter 14 of [9]).

Lemma 3.6. *For each $\delta > 0$ there exists a neighborhood Z_δ of Δ in $B_\delta(\Delta)$ such that the restriction of p_1 to Z_δ is a fiber bundle $(Z_\delta, Y^*, B_1^n(0), p_1|_{Z_\delta})$.*

Proof. We start by constructing a microbundle (O_δ, Y^*, D, p_1) where $O_\delta \subset B_\delta(\Delta)$. Consider the tangent microbundle of Y^*. For each $x \in Y^*$, there exist then an open neighborhood $U_x \subset Y^*$ of x, an open neighborhood $V_x \subset Y^* \times Y^*$ of (x, x) and a trivializing homeomorphism h_x :

$V_x \to U_x \times B_1^n(0)$ for both the diagonal map D and the projection p_1. By compactness of Δ, there exist finitely many $x_1, ..., x_k$ such that $\bigcup_{i=1}^k V_{x_i}$ is a neighborhood of the diagonal Δ. For each x_i, there exists $\lambda_i > 0$, such that $h_{x_i}^{-1}(U_{x_i} \times B_{\lambda_i}^n(0)) \subset B_\delta(\Delta)$. Take $\lambda = \min_i\{\lambda_i\}$ and let $W_i \equiv h_{x_i}^{-1}(U_{x_i} \times B_\lambda^n(0)) \subset B_\delta(\Delta)$. The union $O_\delta \equiv \bigcup_i W_i$ is therefore a microbundle such that $O_\delta \subset B_\delta(\Delta)$. Applying the Kister-Mazur Theorem (Theorem 2 in [4]), we obtain a neighborhood $Z_\delta \subset O_\delta$ of the diagonal Δ such that $(Z_\delta, Y^*, B_1^n(0), p_1|_{Z_\delta})$ is a fiber bundle. $\square$

We now present the final auxiliary result which will be used in the proof of Theorem 2.1. The result is known, but we have not found a complete proof of it anywhere, so we include one here for completeness.

Lemma 3.7. *Let (E, B, F, p) be a fiber bundle over a paracompact and contractible space B. Then (E, B, F, p) is trivial.*

Proof. Since B is contractible, let $f : B \to \{*\}$ and $g : \{*\} \to B$ be two continuous maps that are homotopy-inverses of each other. Let $(g \circ f)^*(E) \equiv \{(b, e) \in B \times E \mid p(e) = (g \circ f)(b)\}$ be the pull-back bundle induced by $g \circ f$. Then $((g \circ f)^*(E), B, F, \mathrm{proj}_1)$ is a fiber bundle, and it is immediately checked it is trivial, since $g \circ f$ is constant. Since $g \circ f$ is homotopic to id_B, from Theorem 2.1 in [1], it follows that $((g \circ f)^*(E), B, F, \mathrm{proj}_1)$ is isomorphic to (E, B, F, p).[1] Hence, (E, B, F, p) is trivial. $\square$

4. Proof of Theorem 2.1

With preparations complete, we proceed to the proof of Theorem 2.1 per se. Let $W \subset Y^*$ be a neighborhood of Y for which there exists a retraction $r_Y : W \to Y$. There exists $\tilde{\delta} > 0$ such that $\tilde{\delta}$-neighborhood $Y(\tilde{\delta})$ around Y in Y^* is contained in W and the $\tilde{\delta}$-neighborhood $X(\tilde{\delta})$ around X in Y^* retracts to X. We denote this retraction also by r_X for notational convenience. Define $\ell_X : [0, \tilde{\delta}] \to \mathbb{R}_+$ by the maximum of $\|x - r_X(x)\|$ over all $x \in Y^*$ such that $d(x, X) \leqslant \delta$. If else, define $\ell_Y : [0, \tilde{\delta}] \to \mathbb{R}_+$ by the maximum of $\|x - r_Y(x)\|$ over all $x \in W$ such that $d(x, Y) \leqslant \delta$. Observe that for $* \in \{X, Y\}$, ℓ_* is continuous and $\ell_*(0) = 0$. For $\delta > 0$, denote by $B_\delta(C)$ the δ-neighborhood around C in $\mathbb{R}^m$.

Let $\varepsilon_0 > 0$. By continuity of $\ell_*(\cdot), * \in \{X, Y\}$, choose $\bar{\delta} > 0$ sufficiently small such that $\ell_*(\bar{\delta}) + \bar{\delta} < \varepsilon_0$. Fix $\delta_0 > 0$ such that

$$\mathrm{Graph}(f) \cap (B_{\delta_0}(C) \times Y) \subset Z_{\bar{\delta}}. \tag{0}$$

Fix any $\delta \in (0, \delta_0)$ and choose points $x_1, \ldots, x_k$ in the interior of n-simplices of S with $d(x_i, C) < \delta$. Let $r_1, \ldots, r_k$ be integers such that $\sum r_i = c$.

[1]We note that that the proof of Theorem 2.1 in [1] applied above relies on the Covering Homotopy Theorem, which requires (see [3]) the base space B to be paracompact.

Apply now the Hopf Approximation Theorem (Theorem 2.5, Appendix C in [10]) to obtain two subdivisions T_0 and S_0 of T, with S_0 a subdivision of T_0, and a simplicial map $g : |S_0(X)| \to |T_0|$ such that:

(1) $\forall x \in X, d(f(x), g(x)) < \bar{\delta}$;
(2) $g(X) \subseteq X$ if $f(X) \subseteq X$;
(3) $\mathrm{Graph}(g) \cap (B_\delta(C) \times Y) \subset Z_{\bar{\delta}}$;
(4) g has finitely many fixed points, each of which is contained in the interior of an n-simplex in $S_0(X)$;
(5) The boundary of $B_\delta(C) \cap X$ in X has no fixed points of g and the index of g over $B_\delta(C)$ is c;
(6) All fixed points of g are contained in $B_\delta(C)$.

Let $F(g)$ be the set of fixed points of g in $B_\delta(C)$. Consider an open neighborhood $V \subset X \setminus \partial X$ of $F(g) \cup \bigcup_{i=1}^k \{x_i\}$ that is contractible and contained in $B_\delta(C) \cap (X \setminus \partial X)$. Using the fact that V is contractible, Lemmas 3.6 and 3.7 imply that the restriction of $p_1|_{Z_{\bar{\delta}}}$ to $Z_{\bar{\delta}}|_V \equiv (p_1|_{Z_{\bar{\delta}}})^{-1}(V)$ defines the trivial bundle $(Z_{\bar{\delta}}|_V, V, B_1^n(0), p_1)$. Therefore, letting $q_1 : V \times B_1^n(0) \to V$ be the natural projection on the first factor, there exists a homeomorphism $\varphi : Z_{\bar{\delta}}|_V \to V \times B_1^n(0)$ such that $p_1|_{Z_{\bar{\delta}}|_V} = q_1 \circ \varphi$. We note that $\mathrm{Graph}(g|_V) \subset Z_{\bar{\delta}}|_V$ (from (3) above). The restriction of φ to the x-section $(Z_{\bar{\delta}}|_V)_x = \{(x, y) \in Z_{\bar{\delta}}|_V\}$ is a homeomorphism with $\{x\} \times B_1^n(0)$. Let now $(h_x)_{x \in B_1^n(0)}$ be a continuous family of homeomorphisms from $B_1^n(0)$ to itself, such that h_x sends x to 0; let φ_2 be the coordinate map of φ mapping to $B_1^n(0)$. We can now define $\psi : Z_{\bar{\delta}}|_V \to V \times B_1^n(0)$ by $(x, y) \mapsto (x, h_{\varphi_2(x,y)} \circ \varphi_2(x, y))$; this is a homeomorphism that sends (y, y) to $y \times \{0\}$. Letting $Z^*_{\bar{\delta}}|_V \equiv Z_{\bar{\delta}}|_V - \{\Delta\}$, it follows that $\psi|_{Z^*_{\bar{\delta}}|_V}$ is a homeomorphism $Z^*_{\bar{\delta}}|_V \to V \times (B_1^n(0) - \{0\})$.

Let $\delta_1 > 0$ be such that the set-distance $d(V, \partial X) > \delta_1$ and $\min_{i \neq j} d(x_i, x_j) \geqslant 3\delta_1$. Let $\delta_2 > 0$ be such that for each $i = 1, ..., k$, any n-simplex τ_i with barycenter at x_i and diameter less than δ_2 is contained in V and is such that $\tau_i \times \tau_i \subset Z_{\bar{\delta}}$. Consider now a closed connected neighborhood B of $F(g) \cup \bigcup_i \{x_i\}$ that is contained in the interior of V. Let $d(\partial V, B) \geqslant \delta_3 > 0$. Fix $\eta \equiv \min\{\delta_1, \delta_2, \delta_3\}$. Lemma 3.2 now gives a simplicial subdivision T_1 of T_0 with mesh less than η such that each x_i is the barycenter of an n-simplex $\tau_i \in T_1$. Our choice of η implies that the collection of simplices τ_i is pairwise disjoint. Let P be the closed star of B with respect to T_1. The set P is a orientable, connected n-pseudomanifold with boundary ∂P (with associated triangulation $T_1(P)$) contained in V.

For each i, using Lemma 3.1 in each τ_i, we obtain a n-simplex $\sigma_i \subseteq \tau_i$ and a PL map $h_i : \sigma_i \to \tau_i$ such that x_i is the barycenter of both σ_i and τ_i, and the only fixed point of h_i, with index r_i. Take now a subdivision T_2 of T_1 that has each σ_i as an n-simplex of T_2 if $|r_i| = 1$. Using Theorem 2.14 in [7], there exist for each i for which $|r_i| \neq 1$, simplicial subdivisions $\hat{S}(\sigma_i)$ and $\hat{T}(\tau_i)$ of σ_i and τ_i, such that $h_i : |\hat{S}(\sigma_i)| \to |\hat{T}(\tau_i)|$ is simplicial. Using Lemma 3.3, there exist subdivisions $\hat{S}$ of T_2 and $\hat{T}$ of T_1 that extend $\hat{S}(\sigma_1)$ and $\hat{T}(\tau_1)$. Since σ_2 and τ_2 are disjoint from σ_1 and τ_1, respectively, the same lemma guarantees that σ_2 is an n-simplex of $\hat{S}$, and τ_2 an n-simplex of $\hat{T}$. This observation applied iteratively together with Lemma 3.3 implies there exists a subdivision $\hat{S}_2$ of T_2 and $\hat{T}_2$ of T_1 such that for each $i = 1, ..., k$, $\hat{S}_2$ extends the triangulation $\hat{S}(\sigma_i)$ and $\hat{T}_2$ extends the triangulation $\hat{T}(\sigma_i)$. For notational convenience we drop the subscripts of $\hat{T}_2$ and $\hat{S}_2$ and refer to these triangulations only as $\hat{T}$ and $\hat{S}$. Note that if $|r_i| = 1$, then we can assume that $\sigma_i \in \hat{S}^n$ and $\tau_i \in \hat{T}^n$.

Define $q : \partial P \cup \bigcup_i \sigma_i \to Y^*$ by $q|_{\partial P} \equiv g|_{\partial P}$ and for each $i = 1, ..., k$, $q|_{\sigma_i} \equiv h_i$. Let $Q = P \setminus \bigcup_{x_i \in V}(\sigma_i \setminus \partial\sigma_i)$. The set Q is a connected, orientable, n-pseudomanifold with boundary $\partial Q = \partial P \cup \bigcup_{x_i \in V} \partial\sigma_i$. Define now a map $d_q : \partial Q \to B_1^n(0) - \{0\}$ by $d(x) = q_2(\psi(x, q(x)))$, where q_2 is the projection on the second factor. Clearly, the degree of d is zero. By the Hopf Extension Theorem (Corollary 18, Chapter 8 in [8]), d_q extends to a map over Q, still denoted d_q. This defines a map $h : Q \to Y^*$ by letting $h(x) = p_2(\psi^{-1}(x, d_q(x)))$, where p_2 is the projection on the second factor.

The graph of h is guaranteed to be in $B_{\bar{\delta}}(\Delta)$ but not in $Q \times Y$, so, from h we now construct another map whose graph is in $Q \times Y$. Since $\mathrm{Graph}(h) \subset Z^*_{\bar{\delta}}|_V \subset B_{\bar{\delta}}(\Delta)$, if $h(x) \in Y^* \setminus Y$, then it follows that $h(x) \in Y(\bar{\delta})$; if $f(X) \subset X$, then we have that $h(x) \in X(\bar{\delta})$. In the latter case, define $\hat{h}_X : Q \to Y$ by $\hat{h} = r_X \circ h$; in the former case, let $\hat{h}_Y = r_Y \circ h$. Therefore we have that for each $x \in Q \subset V \subset X \setminus \partial X$, if $f(X) \subseteq X$, then $\hat{h}_X(Q) \subseteq X$ and $\|x - \hat{h}_X(x)\| \leqslant \ell_X(\bar{\delta}) + \bar{\delta}$; if else, $\|x - \hat{h}_Y(x)\| \leqslant \ell_Y(\bar{\delta}) + \bar{\delta}$. In either case, we can extend the map $\hat{h}_*, * \in \{X, Y\}$ to a map over X by letting it be equal to g everywhere on $X \setminus P$, denoting the extension still by $\hat{h}_*$.

For notational convenience, because the proofs in the two cases ($f(X) \subseteq X$ and $f(X) \not\subset X$) are equal, we will omit the subscripts X and Y from ℓ_X and ℓ_Y, as well as from $\hat{h}_X$ and $\hat{h}_Y$, writing only ℓ and $\hat{h}$.

Recall that: (i) $P \subset V \subset B_\delta(C) \cap (X \setminus \partial X)$, so, from (0), $\mathrm{Graph}(f|_P) \subset Z_{\bar{\delta}} \subset B_{\bar{\delta}}(\Delta)$, which implies that $\|f|_P - id_P\| \leq \bar{\delta}$; (ii) $\|id_Q - \hat{h}|_Q\| \leqslant \ell(\bar{\delta}) + \bar{\delta}$; (iii) for each i, since $\tau_i \times \tau_i \subset Z_{\bar{\delta}} \subset B_{\bar{\delta}}(\Delta)$, then $\|id_{\sigma_i} - \hat{h}|_{\sigma_i}\| \leq \bar{\delta}$. Since $P = Q \cup \bigcup_i \sigma_i$, (i) - (iii) imply $\|f|_P - \hat{h}|_P\| \leq \ell(\bar{\delta}) + 2\bar{\delta}$. In $X \setminus P$, the map $\hat{h}$ equals g, and therefore, from (1), $\|f|_{X \setminus P} - \hat{h}|_{X \setminus P}\| \leq \bar{\delta}$. Hence, we have $\|f - \hat{h}\| \leq \ell(\bar{\delta}) + 2\bar{\delta}$.

Note now that by construction $\hat{h}$ has no fixed points in $X \setminus \bigcup_i(\sigma_i \setminus \partial\sigma_i)$. Since this is a compact set, let $0 < \alpha < \bar{\delta}$ be such that $\|x - \hat{h}(x)\| > 3\alpha$ for all $x \in X \setminus \bigcup_i(\sigma_i \setminus \partial\sigma_i)$. By Lemma 3.2, we can take a subdivision T^* of $\hat{T}$ such that:

(1) The diameter of each simplex is less than α;
(2) for each i, τ_i is the space of a subcomplex $T^*(\tau_i)$ of T^*;
(3) For each i for which $|r_i| = 1$, there is a full-dimensional simplex τ_i^* of T^* that has x_i as its barycenter.

Recall that, for each i, the map $\hat{h}|_{\sigma_i} = h_i : \sigma_i \to \tau_i$ is simplicial by construction w.r.t. to triangulations $\hat{S}(\sigma_i)$ of σ_i and $\hat{T}(\tau_i)$ of τ_i. Since $T^*(\tau_i)$ is a subdivision of $\hat{T}(\tau_i)$, by Lemma 2.16 in [7], there exists, for each i, a subdivision $S^*(\sigma_i)$ of $\hat{S}(\sigma_i)$ such that $\hat{h}|_{\sigma_i} : |S^*(\sigma_i)| \to |T^*|$ is simplicial for each i. When $r_i = \pm 1$, as $\hat{h}|_{\sigma_i}$ is an affine homeomorphism to τ_i^*, the simplex σ_i^* of $S^*(\sigma_i)$ that maps to τ_i^* has x_i as the barycenter for each such i. As before, applying Lemma 3.3 recursively for $i = 1, \ldots, k$, we can extend the triangulation $S^*(\cup_i \sigma_i)$ to a triangulation $\hat{S}^*$ of X. Using now the Theorem and Addendum in [11], we can consider a sufficiently fine barycentric subdivision S^* of $\hat{S}^*$ modulo $S^*(\cup_i \sigma_i)$ and a simplicial map $h^* : |S^*| \to |T^*|$ such that the restriction of h^* to $\cup_i \sigma_i$ equals $\hat{h}$ and $\|h^* - \hat{h}\| < 2\alpha$. It is easily verified that h^* has all the stated properties of the theorem.

References

[1] Cohen, R. (1998): *The Topology of Fiber Bundles*. Lecture Notes. Available here.

[2] Govindan, S., R. Laraki, and L. Pahl (2022): "O'Neill's Theorem for Games," *ArXiv*: *arXiv:2312.03392*.

[3] Huebsch, W. (1955): "On the Covering Homotopy Theorem," *Annals of Mathematics*, vol. 61, No. 3, pp. 555-563.

[4] Kister, J.M. (1964): "Microbundles are Fiber Bundles," *Annals of Mathematics*, vol. 8, No. 1, pp. 190-199.

[5] Milnor, J. (1964): "Microbundles. I," *Topology*, vol. 3, pp. 53-80.

[6] O'Neill, B. (1953): "Essential Sets and Fixed Points," *American Journal of Mathematics*, 75, 3, 497-509.

[7] Rourke, C.P., B.J. Sanderson (1982): *Introduction to Piecewise-Linear Topology*, Springer-Verlag, Berlin.

[8] Spanier, E. H. (1966): *Algebraic Topology*, Springer-Verlag, Berlin.

[9] Switzer, R.M. (1973): *Algebraic Topology — Homology and Homotopy*, Springer-Verlag, Berlin.

[10] Kiang, T.-H. (1987): *The Theory of Fixed Point Classes*, Springer-Verlag, Berlin.

[11] Zeeman, E.C. (1964): "Relative Simplicial Approximation," *Proc. Camb. Phil. Soc.*, 60, 39.

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

Hausdorff Center for Mathematics and Institute for Microeconomics, University of Bonn, Adenaueralle 24-42, 53113 Bonn, Germany.
Email address: `pahl.lucas@gmail.com`