MATHEMATICS OF OPERATIONS RESEARCH
Vol. 33, No. 4, November 2008, pp. 787–820
ISSN 0364-765X | EISSN 1526-5471 | 08 | 3304 | 0787

informs®

DOI 10.1287/moor.1080.0336

Metastable Equilibria

Srihari Govindan
Economics Department, University of Iowa, Iowa City, Iowa 52242,
srihari-govindan@uiowa.edu, http://www.biz.uiowa.edu/faculty/sgovindan

Robert Wilson
Stanford Business School, Stanford, California 94305,
rwilson@stanford.edu, http://faculty-gsb.stanford.edu/wilson/

Metastability is a refinement of the Nash equilibria of a game derived from two conditions: *embedding* combines behavioral axioms called *invariance and small-worlds*, and *continuity* requires games with nearby best replies to have nearby equilibria. These conditions imply that a connected set of Nash equilibria is metastable if it is arbitrarily close to an equilibrium of every sufficiently small perturbation of the best-reply correspondence of every game in which the given game is embedded as an independent subgame. Metastability satisfies the same decision-theoretic properties as Mertens' stronger refinement called stability. Metastability is characterized by a strong form of homotopic essentiality of the projection map from a neighborhood in the graph of equilibria over the space of strategy perturbations. Mertens' definition differs by imposing homological essentiality, which implies a version of small-worlds satisfying a stronger decomposition property. Mertens' stability and metastability select the same outcomes of generic extensive-form games.

Key words: game theory; equilibrium; refinement; stability
MSC2000 subject classification: Primary: 91A10; secondary: 91A18
OR/MS subject classification: Primary: games/group decisions: noncooperative; secondary: decision analysis: theory
History: Received June 25, 2007; revised: April 1, 2008. Published online in *Articles in Advance* November 3, 2008.

1. Introduction. This article contributes to the refinement of the Nash equilibria of a finite game. Previous contributions to this topic derive implications of proposed definitions.[1] Our method differs in that a refinement is derived from two primitive requirements. We obtain a new refinement called *metastability* from conditions called *embedding* and *continuity*. Embedding is a natural generalization of the invariance and small-worlds axioms.[2] Continuity requires that every game nearby has an equilibrium nearby.

Unlike previous contributions that impose an essentiality property in the definition of the refinement studied, we derive an essentiality property from the conjunction of embedding and continuity. However, it is essentiality in the sense of homotopy that is obtained, rather than homology as in Mertens' [21] definition of stability. Continuity is a weak requirement, but we show that in combination with embedding it is equivalent to a succinct mathematical test that is a strong form of homotopic essentiality called *stable essentiality*. Roughly, stable essentiality requires that the strategy set of the given game remains in the range of every homotopic deformation of every suspension of the local projection map from the graph of equilibria to the space of perturbed strategies (Appendix A's Definition A.1 provides the precise definition).

We prove that metastability satisfies all the standard decision-theoretic properties considered by Mertens [21, 22]. It also satisfies the projection property and a slightly weaker version of the decomposition property that are the two parts of Mertens' [23] version of the small-worlds axiom. Metastability is slightly weaker than stability as defined by Mertens. Technically, this weaker aspect appears in our characterization that a metastable set is one for which the local projection map (from a neighborhood in the graph of the equilibria correspondence to the space of strategy perturbations) satisfies a strong form of homotopic essentiality, which is slightly weaker than the homological essentiality that Mertens invokes in his definition of stability. However, in Appendix E we prove that metastability coincides with Mertens' stability for extensive-form games with perfect recall and generic payoffs.

1.1. Embedding and continuity. These two conditions are the following.

Embedding. We interpret any game G as a local version of any larger game $\widetilde{G}$ with additional features that do not affect optimal behavior in G. Imposed as an axiom, embedding requires that a selected set of equilibria of a game G is the projection of a selected set of any larger game $\widetilde{G}$ in which G is embedded. By "embedded" we mean trivially embedded in that optimal strategies of players in G are not affected by actions in $\widetilde{G}$ that are

[1] For a critical review of equilibrium refinements see Hillas and Kohlberg [10]. The initial sections of Hillas et al. [11] review further those refinements based on perturbations of a game's best-reply correspondence, which is the formulation adopted here.

[2] Kohlberg and Mertens [12] and Mertens [23] propose these axioms, and variants are studied by Hillas [9] and Hillas et al. [11].

not in G; that is, their best-reply correspondence is not affected. Embedding subsumes those axioms requiring that a refinement is not affected by extraneous features:

- *Small-worlds*. The additional features could be actions of players in $\tilde{G}$ who are not players in G, provided their actions have no effect on the optimal strategies of players in G.
- *Invariance*. The additional features could be redundant strategies of players in G, such as a pure strategy whose payoffs are replicated by some mixed strategy.
- *Rationality*. The additional features could be presentation effects, behavioral anomalies, or subjective beliefs that are not relevant for optimal play in G.

As Mertens [23, p. 555] remarks regarding the small-worlds axiom:

> "... such a property is essential if one wants to speak of a 'solution concept.' Indeed, otherwise one could never apply the 'solution concept' to a given game; one would first have to embed this game in some 'universal game' of everything going on in the world."

Continuity. Nearby games should have equilibria near selected equilibria of the game. Imposed as an axiom, continuity requires that each neighborhood of a set of equilibria selected for the game G should include an equilibrium of a sufficiently small perturbation of G. From a practical perspective, continuity is necessary for useful applied work and for experimental and empirical studies.

1.2. Metastability. Combining embedding and continuity yields the following brief characterization of metastability (a precise version is in §2.2):[3]

> A connected set of the equilibria of a game G is metastable if and only if every neighborhood contains the projection (into the strategy space of G) of an equilibrium of each sufficiently small perturbation of the best-reply correspondence of any larger game $\tilde{G}$ in which G is embedded.

1.3. Properties of metastable sets of equilibria. In §4 we establish the following decision-theoretic properties:

- *Admissibility, Perfection, and Backward Induction*. A metastable set includes only perfect (hence admissible) equilibria, and it includes a proper equilibrium, which induces a quasi-perfect (and hence a sequential) equilibrium in every extensive-form game with G as its strategic form.
- *Iterative Elimination of Weakly Dominated Strategies and Never Weak Best Replies, and Forward Induction*. A subset of a metastable set survives iterative elimination of weakly dominated strategies and strategies that are inferior replies at every equilibrium in the set.

Additional properties include Ordinality and Player-Splitting.

Our definition of metastability has two key aspects: (1) As in Govindan and Mertens [5], we use the best-reply correspondence of a game as the primitive, rather than a formulation in terms of payoffs; (2) As in Hillas [9] and Hillas et al. [11], we test for continuity using perturbations of the best-reply correspondence. The relevant mathematical tool is then homotopy theory rather than the stronger homology theory invoked by Mertens [21] for perturbations of strategies.

Metastability differs from Mertens' stability chiefly in the consequences of our reliance on perturbations of players' best-reply correspondence. As mentioned, this yields a characterization in terms of homotopic essentiality but precludes a tighter characterization in terms of homological essentiality—because the dimensions of the domain and range spaces of the projection map need not agree. This difference is manifest chiefly in the implications of the small-worlds axiom. Metastability satisfies the first part of Mertens' version of the small worlds property:

Projection Property. *A metastable set of a larger game projects to a metastable set of an embedded local game*.

That is, if G is embedded in $\tilde{G}$ and $\tilde{S}$ is a metastable set for $\tilde{G}$ then $S \equiv \text{proj}_G(\tilde{S})$ is a metastable set for G. Mertens' stability satisfies the projection property and also its converse in the following strong form:

(1*) Each stable set of G is the projection of a stable set of each game in which G is embedded.

[3] The requirement that a metastable set is connected excludes the set of all equilibria, which trivially satisfies (1) and (2). It also reflects the fact that all equilibria in a single connected component of the equilibria of a generic extensive-form game have the same outcome (Kreps and Wilson [16]), and more fundamentally, the fact that the uniformly hyperstable sets of Nash equilibria are necessarily connected because they are precisely the essential sets of fixed points of any map whose fixed points are the Nash equilibria (Govindan and Wilson [8]). Connectedness excludes the minimal stable sets studied by Kohlberg and Mertens [12]—Mertens [24] argues that minimality violates the ordinal properties of players' preferences.

(2^*) If S is a stable set of G, then for any stable set S' of any other game G' the product $S \times S'$ is a stable set of the product game $G \times G'$.

The projection property along with (1^*) is Mertens' small-worlds property. When the larger game is the product of two independent games, the projection property along with (2^*) is the *decomposition property*. Metastability satisfies the following weaker versions of (1^*) and (2^*).

(1) Each metastable set of G *contains* the projection of a metastable set of each game in which G is embedded.

(2) If S is a metastable set of G, then for any other game G' there *exists* a metastable set S' of G' such that $S \times S'$ is a metastable set of the product game $G \times G'$.

Thus (1) states that if S is a metastable set of G and G is embedded in $\tilde{G}$, then there is a metastable set $\tilde{S}$ of $\tilde{G}$ such that $S \supseteq \text{proj}_G(\tilde{S})$, where the projection property assures that $\text{proj}_G(\tilde{S})$ is itself a metastable set for G—and analogously in (2).

Our view is that these weaker versions are natural from a game-theoretic viewpoint. A metastable set is intended to be a collection of possible outcomes that can be refined further only with additional information. For instance, in an extensive-form game the play off the equilibrium path is typically indeterminate, but in specific contexts additional considerations might lead to selection of (say) minimal or maximal "punishments" for deviations from the equilibrium path. Analogously, embedding a game in a particular larger game provides such a context that can select a metastable set (the projection of a metastable set of the larger game) that is a strict subset of another metastable set. The decomposition property is thus seen as unduly strong when the product game $G \times G'$ allows correlated selections of metastable sets for the two embedded games; i.e., correlated selections in the product game destroy some of the presumed independence between play in the component games. Even so, these considerations are relevant only for nongeneric extensive-form games; we show in Appendix E that metastability agrees with Mertens' stability for generic extensive-form games.

The key step of our technical development is the demonstration that metastability (like Mertens' stability) of a connected set of equilibria is a property of its *germ*, i.e., its neighborhood in the graph of the equilibrium correspondence over the space of strategy perturbations. Theorem 3.2 characterizes metastability by the property that the local projection map from the graph of the equilibrium correspondence is *stably essential* in the sense of homotopy theory. That is, as defined in Appendix A's Definition A.1, the map remains essential when it is embedded in a space with extra dimensions—using the formal definition of a "suspension" from algebraic topology.[4] This characterization is especially relevant for Appendix E on extensive-form games because nearby points in the germ (i.e., equilibria of nearby games) induce beliefs that support sequential equilibria.

1.4. Synopsis. The formulation and definition of metastability are established in §2. In §3 metastability is characterized in terms of the stable essentiality of the local projection map. Then in §4 we verify that metastability satisfies the decision-theoretic axioms listed by Mertens [21]. To enable comparisons, §5 provides concise definitions of Mertens' stability and BR-stability of Hillas et al.[5] Section 6 provides concluding remarks.

Appendix A provides mathematical background regarding homotopic essentiality and defines stable essentiality of a map. Appendix B is a brief summary of the properties of multisimplicial and polyhedral complexes invoked in the proofs in Appendix C. Technical implications of the continuity axiom are established in Appendix C. In particular, §C.4 shows that the projection property is equivalent to stable essentiality of the projection map from the graph of the equilibrium correspondence to the space of strategy perturbations. Appendix D shows that the results of the paper are unaffected by alternative specifications of the space of strategy perturbations. Appendix E proves that for generic extensive-form games with perfect recall, the outcomes of metastable and Mertens' stable sets are the same.

2. Formulation and definition of metastability.

2.1. Preliminaries. Unless otherwise stated, all the sets we consider are subsets of Euclidean spaces equipped with the l_1-metric. Throughout, by a map or a function we mean a continuous function. By a correspondence we mean an upper-semicontinuous correspondence whose domain is compact, whose range is compact and convex, and whose values are nonempty compact convex sets. We rely on the following approximation theorem for correspondences—see McLennan [18, Proposition 2.25].

[4] For the applications here, essentiality—and hence also stable essentiality—fails to hold if the domain has a smaller dimension than the range; and stably essential is the same as essential when the dimensions of the domain and range are the same, but not when the domain has higher dimension than the range.

[5] These definitions are deferred to §5 at the suggestion of a referee, but readers unfamiliar with these refinements might want to examine their definitions before embarking on §§3 and 4.

Given a correspondence $\varphi\colon X \twoheadrightarrow Y$ and a scalar $\varepsilon > 0$, there exists a function $f\colon X \to Y$ whose graph is contained in the ε-neighborhood of the graph of φ.

For any two correspondences $\varphi\colon X \twoheadrightarrow Y$ and $\varphi'\colon X \twoheadrightarrow Y'$, their "product" $\varphi \otimes \varphi'\colon X \twoheadrightarrow Y \times Y'$ is defined by $(\varphi \otimes f')(x) = \varphi(x) \times \varphi'(x)$. For any two correspondences $\varphi, \varphi'\colon X \twoheadrightarrow Y$ their distance is $d(\varphi, \varphi') = \sup_{x\in X} d_H(\varphi(x), \varphi'(x))$, where d_H is the Hausdorff distance between compact sets. A correspondence φ' is a δ-perturbation of another φ if $d(\varphi, \varphi') \le \delta$.

We study a fixed finite game G in strategic form. The player set is $\mathcal{N} = \{1, \ldots, N\}$. Player n's set Σ_n° of pure strategies comprises the vertices of the simplex Σ_n of his mixed strategies. Let $\Sigma = \prod_{n\in\mathcal{N}} \Sigma_n$. Player n's payoff function is given by a multilinear function $G_n\colon \Sigma \to \mathbb{R}$. The best-reply correspondence of the game G is denoted by $R\colon \Sigma \twoheadrightarrow \Sigma$.

We do not use the l_1-metric on Σ. Instead, the distance between two points $\sigma, \sigma' \in \Sigma$ is given by $\|\sigma - \sigma'\| \equiv 0.5 \times \max_{n\in\mathcal{N}} \|\sigma_n - \sigma_n'\|_1$. The only reason for this convention is expositional convenience. Using this metric we have the following properties of perturbations of R. (1) For each $\tau \in \Sigma$, the correspondence $(1-\delta)R + \delta\tau$ is a δ-perturbation of R. (2) If φ is a δ-perturbation of R, then $\sum_{s_n\in\Sigma_n^\circ\setminus R(\sigma)} \tau_{n,s_n} \le \delta$ for all $\sigma \in \Sigma$, $\tau \in \varphi(\sigma)$, and $n \in \mathcal{N}$.

For each $\delta \in [0, 1]$, let Σ_δ be the set of $\sigma \in \Sigma$ such that $\sum_{s_n\in\Sigma_n^\circ\setminus R_n(\sigma)} \sigma_{n,s_n} \le \delta$ for each player n, i.e., the total probability of player n's strategies that are suboptimal against σ is at most δ. It follows from the previous paragraph that all the fixed points of a δ-perturbation of R are contained in Σ_δ.

2.2. Definition of metastability. We interpret the continuity axiom as saying that if $S \subseteq \Sigma$ is a solution of the game G, then for any small perturbation of G's best-reply correspondence R there exists a fixed point close to S. The embedding axiom says that if we embed G in a larger game $\widetilde{G}$, then for any solution S of G there exists a solution of $\widetilde{G}$ that projects to S. Combining these two axioms, we see that the minimal requirement on a solution concept is the following: For any solution S of G and any embedding of G in a larger game $\widetilde{G}$, there is a fixed point of each small perturbation of the best-reply correspondence of $\widetilde{G}$ whose projection is contained in S.

To formalize the above ideas in a definition, we need some more notation. For each nonnegative integer k let Δ^k be the k-dimensional unit simplex in $\mathbb{R}^{k+1}$. (Δ^0 is a singleton set.) Let $R^k\colon \Sigma \times \Delta^k \twoheadrightarrow \Sigma$ be the trivial extension of G's best-reply correspondence R, namely, $R^k(\sigma, \lambda) = R(\sigma)$. We interpret Δ^k as the strategy space of players other than those in $\mathcal{N}$ when we embed G in a larger game.[6]

Definition 2.1 (Premetastable Set). A closed set $S \subseteq \Sigma$ is premetastable if for every integer $k \ge 0$ and every $\varepsilon > 0$, there exists $\delta_{k\varepsilon} > 0$ such that every correspondence of the form $\varphi^k \otimes \pi^k\colon \Sigma \times \Delta^k \twoheadrightarrow \Sigma \times \Delta^k$, where φ^k is a $\delta_{k\varepsilon}$-perturbation of R^k has a fixed point whose projection to Σ is within ε of S.

Some remarks are in order concerning this preliminary definition. First, the perturbations we consider seem special in that they are products of two correspondences. Given our interpretation of Δ^k as a strategy space, it seems more natural to consider perturbations of $R^k \otimes \Pi$, where Π is viewed as the best-reply correspondence of the others. But the results in Appendix C show that our choice is without loss of generality. Our second remark concerns the seemingly stronger requirement that δ not depend on the correspondence π, just on k and ε. Again, Appendix C shows that the condition as stated in our definition is equivalent to the existence of a $\delta_{k\varepsilon}$ such that every $\delta_{k\varepsilon}$-perturbation of $R^k \otimes \mathrm{Id}$, where Id is the identity function on Δ^k, has a fixed point whose projection to Σ is within ε. Therefore, our definition is equivalent to one where we start with correspondences of the form $R^k \otimes \Pi$ and consider $\delta_{k\varepsilon,\Pi}$-perturbations of them.

The following equivalent characterization of premetastable sets provides the basis for our definition of metastable sets.

Theorem 2.1. *A closed set $S \subseteq \Sigma$ is premetastable if and only if there exist $\delta_* > 0$ and for each $0 < \delta \le \delta_*$ a closed subset V_δ of Σ_δ, such that $\lim_{\delta\to 0} V_\delta$ is the set of Nash equilibria in S; and for every k there exists $0 < \delta_k \le \delta_*$, such that for each $0 < \delta \le \delta_k$, V_δ contains the projection of a fixed point of the product $\varphi^k \otimes \pi^k$ of every δ-perturbation φ^k of R^k and every correspondence $\pi^k\colon \Sigma \times \Delta^k \twoheadrightarrow \Delta^k$.*

Proof. The sufficiency of the conditions is obvious. As for their necessity, suppose that S is a premetastable set. Let $\{\varepsilon_i\}_{i=0}^\infty$ be a decreasing sequence of positive numbers converging to zero. For each i, and $k \ge 0$, there exists $\delta_{k,\varepsilon_i} > 0$ satisfying the continuity condition in Definition 2.1. Without loss of generality we can assume

[6] If Δ^k is to be interpreted literally as the set of strategy profiles of a group of players, then it should be a product of simplices. But, as we show in Appendix C, Δ^k can be replaced by any convex set that is homeomorphic to it and all the results in the paper would hold. Thus, it is without loss of generality that we focus on simplices.

that for each i, the sequence $\{\delta_{k,\varepsilon_i}\}_{k=0}^{\infty}$ converges monotonically to zero. For each i choose a nonnegative integer $l(i) \geq i$ such that the sequence $\{\delta_{l(i),\varepsilon_i}\}_{i=0}^{\infty}$ decreases monotonically to zero.

Let δ_* be $\delta_{l(0),\varepsilon_0}$. For each $0 < \delta \leq \delta_*$, let V_δ be the intersection of Σ_δ with the ε_i-neighborhood of S, where i is the unique nonnegative integer such that $\delta_{l(i),\varepsilon_i} \geq \delta > \delta_{l(i+1),\varepsilon_{i+1}}$. Because the sequences of ε_i and $\delta_{l(i),\varepsilon_i}$ are monotonically decreasing, the V_δs are monotonically decreasing in δ. The limit of these sets is the set of Nash equilibria in S: Obviously, the latter set is contained in each V_δ and hence in the limit; on the other hand, any point that is not a Nash equilibrium fails to be in Σ_δ for small enough δ.

For each k, let δ_k be $\delta_{l(k),\varepsilon_k}$. To finish the proof, we prove that the V_δs satisfy the continuity condition of the theorem using these δ_ks. Fix an integer $k \geq 0$ and $0 < \delta \leq \delta_k$. There exists a unique $i \geq k$ such that $\delta_{l(i+1),\varepsilon_{i+1}} < \delta \leq \delta_{l(i),\varepsilon_i}$. By construction, every correspondence of the form $\varphi^k \otimes \pi^k$ where φ^k is δ_{k,ε_i}-perturbation of R^k has a fixed point whose projection to Σ is in the ε_i-neighborhood of Σ. Because $\delta_{k',\varepsilon_i}$ is a monotonically decreasing sequence in k' and $k \leq i \leq l(i)$, $\delta_{l(i),\varepsilon_i}$ and, hence, also δ, are smaller than δ_{k,ε_i}. Therefore, every correspondence of the form $\varphi^k \otimes \pi^k$, where φ^k is δ-perturbation of R^k has a fixed point whose projection σ to Σ is in the ε_i-neighborhood of Σ. Obviously, σ also belongs to Σ_δ and hence to V_δ. □

The V_δs in the above theorem can be viewed as encoding a set of beliefs supporting the equilibria in S as follows. For every sequence of δs converging to zero, and every corresponding sequence σ_δ of profiles in the sets $V_\delta \backslash \partial\Sigma$ converging to some point $\sigma \in S$, take the corresponding sequence of all conditional probability distributions over all possible subsets of pure strategy profiles; then going to the limit produces a set of limiting beliefs for σ. (If we did not insist on full support along the sequence, some of these conditional probabilities might not be defined.)

These systems of conditional probability distributions have a long history in decision theory, e.g., Rényi [26]. In game theory the basic formulation by Myerson [25] is extended by McLennan [20]; moreover, McLennan [19] shows that the space of such conditional systems is homeomorphic to a ball. The induced beliefs are especially important in studies of games in extensive form. Kohlberg and Reny [13] develop their role in justifying the role of "consistent assessments" in the formulation of sequential equilibrium by Kreps and Wilson [16], and more generally Blume et al. [1] develop their role in the general formulation of lexicographic equilibria and the special cases of equilibria that are sequential, perfect, or proper. In particular, when the game (as we define it here) arises from a nontrivial game in extensive form, these limiting beliefs induce posterior distributions at information sets that are excluded by σ. Thus, the sets V_δ we consider generate a set of beliefs that support the equilibria in S as in sequential equilibria.

Based on this interpretation, it is appropriate to impose tighter topological properties. First, we require that V_δ is the closure of $V_\delta \backslash \partial\Sigma$ and that S is the limit of the sets V_δ. Second, to ensure that not just S is connected but also that the induced beliefs are, we require that the sets $V_\delta \backslash \partial\Sigma$ are connected. This yields the following definition of the central solution concept in this paper.

Definition 2.2 (Metastable Set). A closed set $S \subseteq \Sigma$ is a metastable set if there exist $\delta_* > 0$ and for each $0 < \delta \leq \delta_*$ a closed subset V_δ of Σ_δ such that $\lim_{\delta\to 0} V_\delta = S$, and:

(i) Connectedness: $V_\delta \backslash \partial\Sigma$ is connected and dense in V_δ;

(ii) Continuity: For every k there exists $0 < \delta_k \leq \delta_*$ such that for each $0 < \delta \leq \delta_k$, V_δ contains the projection of a fixed point of the product $\varphi^k \otimes \pi^k$ of every δ-perturbation φ^k of R^k and every correspondence $\pi^k: \Sigma \times \Delta^k \twoheadrightarrow \Delta^k$.

Remark 2.1 One could strengthen the continuity requirement by requiring that δ_k is independent of k. Theorem 3.1 shows that this is equivalent to the condition as stated.

2.3. Stably essential sets. In §3 we characterize a metastable set by the property that the local projection map from the graph of the equilibrium correspondence to the space of strategy perturbations is stably essential in the sense of homotopy. Here we gather all the relevant notation and definitions concerning the graph of this equilibrium correspondence.

For each player n and each $0 < \delta \leq 1$, let $P_\delta = \{\varepsilon\tau \mid 0 \leq \varepsilon \leq \delta, \tau \in \Sigma\}$ and denote its topological boundary by ∂P_δ. For $\eta \in P_1$, $\bar{\eta}_n \equiv \sum_{s \in \Sigma_n^\circ} \eta_{n,s}$ is the minimum error probability; $\bar{\eta}_n$ is constant across players so we denote this number by $\bar{\eta}$. Given any $\eta \in P_1$ define the perturbed game $G(\eta)$ to be the game where the strategy set of each player n is the same as in G, but where the payoff from a strategy profile τ is the payoff in G from the profile $\sigma = (1-\bar{\eta})\tau + \eta$. Then we say that σ is a perturbed equilibrium of $G(\eta)$ if τ is an equilibrium of $G(\eta)$. Let $\mathscr{E}$ be the graph of the perturbed equilibrium correspondence over P_1, i.e.,

$$\mathscr{E} = \{(\eta, \sigma) \in P_1 \times \Sigma \mid \sigma \text{ is a perturbed equilibrium of } G(\eta)\}.$$

For $(\eta, \sigma) \in \mathscr{E}$ we use $\tau(\eta, \sigma) \equiv (1-\bar{\eta})^{-1}(\sigma - \eta)$ to denote the corresponding equilibrium of $G(\eta)$. Observe that a pure strategy s of player n is in the support of $\tau_n(\eta, \sigma)$ only if it is an optimal reply to σ in G.

Denote by p the natural projection from $\mathscr{E}$ to P_1. For $E \subseteq \mathscr{E}$, let $E_0 = \{(0, \sigma) \in E\}$, and for $0 < \delta \le 1$ let $(E_\delta, \partial E_\delta) = p^{-1}(P_\delta, \partial P_\delta) \cap E$. That is, E_δ is a germ for $\{\sigma \mid (0, \delta) \in E\}$ as described in §1.

Mertens [22] defined a concept of $*$-stability by requiring the projection map from subsets of E to be essential in cohomology (see §5 for Mertens' definition). In this definition if we were to require, instead, essentiality in homotopy, we would obtain the following, which is the basis for our notion of stably essential sets.

Definition 2.3 (Homotopy-Stable Set). $S \subseteq \Sigma$ is a homotopy-stable (h-stable) set if for some closed subset E of $\mathscr{E}$ with $E_0 = \{0\} \times S$:

(i) Connectedness: For every neighborhood V of E_0 in E, the set $V \backslash \partial E_1$ has a connected component whose closure is a neighborhood of E_0 in E;

(ii) Homotopic essentiality: p: $(E_\delta, \partial E_\delta) \to (P_\delta, \partial P_\delta)$ is essential in homotopy for some $\delta > 0$.

Remark 2.2. It follows from Lemma A.5 that, in the above definition, the homotopic essentiality of the projection for some $\delta > 0$ implies that for all $0 < \delta' < \delta$.

Appendix A, Definition A.1, defines the stronger property of stably essential in homotopy. The following is a version of this definition (cf. also Lemma A.8).

Definition 2.4 (Stably Essential Map). Given a map p: $(E, \partial E) \to (P, \partial P)$, where $(P, \partial P)$ is a ball pair with $p(E \backslash \partial E) \subseteq P \backslash \partial P$, and a k-dimensional simplex Δ^k, let p^k: $(E, \partial E) \times (\Delta^k, \partial \Delta^k) \to (P, \partial P) \times (\Delta^k, \partial \Delta^k)$ be the trivial extension for which $p^k(e, \lambda) = (p(e), \lambda)$. The map p is stably essential in homotopy if p^k is essential in homotopy for every k.

Thus a map is stably essential if it remains essential when its domain and range are extended trivially to higher dimensional spaces. Using this property, the following strengthens the definition of an h-stable set in Definition 2.3.

Definition 2.5 (Stably Essential Set). $S \subseteq \Sigma$ is a stably essential set if for some closed subset E of $\mathscr{E}$ with $E_0 = \{0\} \times S$:

(i) Connectedness: For every neighborhood V of E_0 in E, the set $V \backslash \partial E_1$ has a connected component whose closure is a neighborhood of E_0 in E;

(ii) Stable essentiality: the projection map p: $(E_\delta, \partial E_\delta) \to (P_\delta, \partial P_\delta)$ is stably essential in homotopy for some $\delta > 0$.

Remark 2.3. Using Definition 2.4 and Lemma A.5, if the projection map from $(E_\delta, \partial E_\delta)$ is stably essential for some δ, then it is essential for all smaller δ.

Instead of working with the graph of the perturbed equilibrium correspondence we could equivalently work with the graph of the equilibrium correspondence, i.e., the set of $(\eta, \tau) \in P_1 \times \Sigma$ such that τ is an equilibrium of $G(\eta)$ and thus $\sigma = (1 - \bar{\eta})\tau + \eta\tau$ is a perturbed equilibrium of $G(\eta)$. There is an obvious homeomorphism between the two spaces that commutes with the projections to P_1. Hence we obtain the same stable sets if we use the graph of equilibria. This observation is true for all "stability" definitions involving subsets of $\mathscr{E}$.

There is nothing canonical about using P as the space of perturbations. Indeed, we could have used $[0, 1] \times \Sigma$; or $[0, 1]^N \times \Sigma$, where the error probabilities are allowed to be different across players; or the actual games resulting from these perturbations. (Mertens [21, 22] shows that for his definition of stability one obtains the same collection if one uses any of these or other perturbation spaces derived from them.) In Appendix D, we show that a similar result holds for stably essential sets. (We suspect that a corresponding result is not true for h-stable sets.)

A subset E of $\mathscr{E}$ that satisfies the conditions of Definition 2.5 is called a germ for S. If E is a germ for S, then so is the closure $\widehat{E}$ of $E \backslash \partial E_1$. Indeed, the connectedness condition for E implies that $\widehat{E}_0 = \{0\} \times S$ and that $\widehat{E}$ satisfies the connectedness condition; $\widehat{E}$ also satisfies the essentiality condition because of Lemma A.5. Therefore, in our proofs, whenever we talk about a germ E for a stably essential set S, we implicitly assume that E is the closure of $E \backslash \partial E_1$. On a related point about our proofs, in proving the connectedness condition for a germ E that is the closure of $E \backslash \partial E_1$, we typically prove the following equivalent version: For a basis of neighborhoods V of E_0 in E, $V \backslash \partial E_1$ is connected.

Stable essentiality seems to be a reasonable solution concept in itself. All the game theoretic properties we prove in §4 for metastable sets are true also for stably essential sets. The one major defect might be a lack of compactness for the collection of these sets—although we do not know this for sure. As we show in §3, if we take the closure of these sets, then we obtain the collection of metastable sets.

3. Characterization of metastability. In this section we derive topological characterizations of metastable sets. Our first result Theorem 3.1 shows that the sets V_δ in the definition of metastability can be chosen to be a nested collection of semialgebraic sets that, moreover, satisfy the continuity property in the stronger form given by property (2) of Appendix C. Our second theorem shows that metastability can be defined just like stably essential sets, except that the stable essentiality requirement is weaker. Our next two theorems show that the implied difference between metastability and stably essential sets is minor, i.e., the collection of metastable sets is obtained by taking the Hausdorff closure of the collection of stably essential sets.

THEOREM 3.1. *S is metastable iff there exists $\delta_0 > 0$ and for each $0 < \delta \le \delta_0$ there exists a closed semialgebraic subset W_δ of Σ_δ such that*:

(i) *$W_\delta \backslash \partial\Sigma$ is connected and dense in W_δ*;

(ii) *For each k, each δ-perturbation φ of R^k, and each correspondence π: $\Sigma \times \Delta^k \to \Delta^k$, $\varphi \otimes \pi$ has a fixed point whose projection to Σ is contained in W_δ*;

(iii) *$W_{\delta'} \subseteq W_\delta$ if $0 < \delta' < \delta$*;

(iv) *$\bigcap_{0<\delta\le\delta_0} W_\delta = S$.*

PROOF. The sufficiency part of the proof is obvious. We prove the necessity of the conditions. Given $V_\delta \subseteq \Sigma_\delta$ for all small δ satisfying the conditions of Definition 2.2, choose a monotone sequence $\delta_k \downarrow 0$ from the δs such that for all k, k' with $k \ge k'$, $V_k \equiv V_{\delta_k}$ contains the projection of a fixed point of the product $\varphi^{k'} \otimes \pi^{k'}$ of every δ_k-perturbation $\varphi^{k'}$ of $R^{k'}$ and every correspondence $\pi^{k'}$: $\Sigma \times \Delta^{k'} \to \Delta^{k'}$. Take a triangulation of Σ. For each positive integer l, let Σ^l be the lth barycentric subdivision of this triangulation. Let P^l be the simplices of Σ^l that intersect S. For each l let X^l be the closure of $\{(\delta, \sigma) \in (0,1) \times P^l \mid \sigma \in \Sigma_\delta \backslash \partial\Sigma\}$ in $[0,1] \times P^l$, and let $\partial X^l = \{(\delta, \sigma) \in X^l \mid \delta = 0 \text{ or } \sigma \in \partial\Sigma\}$, and let g^l: $X^l \to [0,1]$ be the projection to the first coordinate. By Mertens [23, Lemma 2] there exists $0 < \bar{\delta}_l \le 1$, a finite number of closed semialgebraic subsets $X^{l,1}, \ldots, X^{l,j_l}$ of X^l such that for each $0 < \delta \le \bar{\delta}_l$ and each j, letting $X_\delta^{l,j} = ((g^l)^{-1}([0,\delta])) \cap X^{l,j}$, we have:

(a) $X_\delta^{l,j} \backslash (\partial X^l \cup (g^l)^{-1}(\delta))$ is connected and dense in $X_\delta^{l,j}$;

(b) $X_\delta^{l,j} \cap X_\delta^{l,j'} \subseteq \partial X^l$ for $j' \ne j$;

(c) $\bigcup_j X_\delta^{l,j} = (g^l)^{-1}([0,\delta])$.

Suppose for $0 < \delta, \delta' \le \bar{\delta}_l$ that (δ, σ) and (δ', σ) belong to $X^l \backslash \partial X^l$. Then $((\lambda\delta + (1-\lambda)\delta'), \sigma)$ belongs to $X^l \backslash \partial X^l$ for all $\lambda \in [0,1]$, and the above properties imply the property

(d) (δ, σ) belongs to $X_{\bar{\delta}_l}^{l,j}$ iff (δ', σ) does.

Because P^l is a neighborhood of S, it is a neighborhood of V_k for large k. Therefore, by the connectedness property for V_k, and also by the above properties of X^l, if k is also large enough such that $\delta_k \le \bar{\delta}_l$, then there exists $1 \le j_k \le j_l$ such that $\{\delta_k\} \times (V_k \backslash \partial\Sigma)$, and hence also its closure $\{\delta_k\} \times V_k$, is contained in $X_{\delta_k}^{l,j_k}$. Along a subsequence of V_ks now, j_k is constant, say 1. By replacing the sequence of V_k's with a subsequence that is obtained by the diagonalization process, we can assume that for each l and $k \ge l$, $V_k \subseteq P^l$, $\delta_k \le \bar{\delta}_l$ and $\{\delta_k\} \times V_k$ is contained in $X_{\delta_k}^{l,1}$. (The subsequence of V_ks that we use in place of the original sequence continues to satisfy the continuity property described before.)

For each l, let W^l be the projection of $X_{\bar{\delta}_l}^{l,1}$ to Σ. W^l is then a closed semialgebraic set that contains V_k for all $k \ge l$. By the continuity property for the V_ks, each W^l now satisfies property ($1k$) of Appendix C for each k. By Theorem C.9, it satisfies property (2) as well: There exists δ_l such that for each $0 < \delta \le \delta_l$ and each k, W^l contains the projection of a fixed point of the product $\varphi^k \otimes \pi^k$ of every δ_k-perturbation φ^k of R^k and every correspondence π^k: $\Sigma \times \Delta^k \to \Delta^k$. Without loss of generality, we can assume that $\delta_l \le \bar{\delta}_l$ for all l and that the sequence δ_l decreases monotonically to zero.

For each $0 < \delta \le \delta_1$ define W_δ to be the projection of $X_\delta^{l,1}$ to Σ, where l is the unique integer such that $\delta_{l+1} < \delta \le \delta_l$. The W_δs are closed semialgebraic sets. We show that they satisfy the four enumerated conditions of the theorem. For the first three properties, we fix $\delta_{l+1} < \delta \le \delta_l$.

PROPERTY 1. Property (a) implies that the set of $(\delta', \sigma) \in X_\delta^{l,1}$ such that $\sigma \notin \partial\Sigma$ is connected. Hence $W_\delta \backslash \partial\Sigma$, which is the image of this set under the projection to Σ, is connected. Suppose $\sigma \in W_\delta$. Then there exists $0 \le \delta' \le \delta$ such that $(\delta', \sigma) \in X_\delta^{l,1}$. By property (a) again, there exists a sequence (δ^i, σ^i) in $X_\delta^{l,1} \backslash \partial X^{l,1}$ converging to (δ', σ). Obviously the sequence σ^i belongs to $W_\delta \backslash \partial\Sigma$ and converges to σ. Hence $W_\delta \backslash \partial\Sigma$ is dense in W_δ.

PROPERTY 2. By Theorem C.1, it is sufficient to show that for each k, and each correspondence $\varphi^k \otimes \pi^k$: $\Sigma \times \Delta^k \to \Sigma \times \Delta^k$ where φ^k is a δ-perturbation of R^k whose values are contained in $\Sigma \backslash \partial\Sigma$, $W^l \backslash \partial\Sigma$ contains the projection σ of a fixed point of $\varphi^k \otimes \pi^k$. Fix such a correspondence $\varphi^k \otimes \pi^k$. Because W^l satisfies property (2) of Appendix C, we have that there exists a fixed point of $\varphi^k \otimes \pi^k$ whose projection to Σ, call it σ, is contained in W^l. We will prove now that, in fact, σ belongs to W_δ, which proves point (2) of the theorem. Observe that σ belongs to $W^l \backslash \partial\Sigma$. Therefore, there exists $0 \le \delta' \le \delta_l$ such that $(\delta', \sigma) \in X_{\bar{\delta}_l}^{l,1}$. If $\delta' \le \delta$, then $(\delta', \sigma) \in X_\delta^{l,1}$ and therefore $\sigma \in W_\delta$. Suppose now that $\delta' > \delta$. Because $\sigma \in W^l \cap \Sigma_\delta$, (δ, σ) belongs to X_δ^l, and, using property (d) above, (δ, σ) actually belongs to $X_\delta^{l,1}$; hence, σ belongs to W_δ in this case as well.

Property 3. It is sufficient to prove that if $\delta_{l+1} \le \delta' < \delta$ then $W_{\delta'} \subseteq W_\delta$. If $\delta' > \delta_{l+1}$, then the result follows from the fact that $X^{l,1}_{\delta'} \subseteq X^{l,1}_{\delta}$. Observe that if $\delta' = \delta_{l+1}$ then property (a) for $X^{l+1,1}_{\delta_{l+1}}$ implies that $X^{l+1,1}_{\delta_{l+1}} \backslash \partial X^{l+1}$ is a connected subset of $(g^l)^{-1}([0,\delta]) \backslash \partial X^l$; moreover, by construction, it contains $\{\delta_{l+1}\} \times (V_{l+1} \backslash \partial \Sigma)$. Since $X^{l,1}_{\delta} \backslash \partial X^l$ contains this latter set, by properties (a), (b), and (c) above $X^{l+1,1}_{\delta_{l+1}} \backslash \partial X^{l+1} \subseteq X^{l,1}_{\delta} \backslash \partial X^l$. Using property (a) again, we get that $X^{l+1,1}_{\delta_{l+1}}$ is contained in X^l_δ. Hence $W_{\delta_{l+1}} \subseteq W_\delta$.

Property 4. By property (3) it is sufficient to show that $\bigcap_l W_{\delta_l} = S$. For each l, since $\{\delta_k\} \times V_k$ is contained $X^{l,1}_{\delta_l}$ for all large k, W_{δ_l} contains V_k for all such large k. Because the V_k's converge to S, each W_{δ_l} contains W; thus, $\bigcap_l W_{\delta_l}$ contains S. On the other hand, for each l, W_{δ_l} is contained in P^l, and the P^l's form a basis of neighborhoods of S. Hence $\bigcap_l W_{\delta_l}$ is contained in S and thus we obtain (4). □

We now provide a characterization of metastability in terms of subsets of the graph $\mathscr{E}$ of the equilibria of perturbed games. For each $0 < \delta \le 1$ let $Q_\delta = \{\eta \in P_\delta \mid \bar{\eta} = \delta\}$ and let ∂Q_δ be its relative boundary. Let $(\mathscr{F}_\delta, \partial \mathscr{F}_\delta) = p^{-1}(Q_\delta, \partial Q_\delta)$. We have a well-defined correspondence ψ_δ: $\Sigma_\delta \twoheadrightarrow \mathscr{F}_\delta$ given by $\psi_\delta(\sigma) = \{(\eta, \sigma) \in \mathscr{F}_\delta\}$.

Theorem 3.2. *$S \subseteq \Sigma$ is metastable if and only if there exists a closed subset E of $\mathscr{E}$ with $E_0 = \{0\} \times S$ and*:

(i) *Connectedness*: *For every neighborhood V of E_0 in E, the set $V \backslash \partial E_1$ has a connected component whose closure is a neighborhood of E_0 in E*;

(ii) *Stable essentiality*: *There exists $\delta_0 > 0$ such that for each $0 < \delta \le \delta_0$, letting $(F_\delta, \partial F_\delta)$ be the pair $p^{-1}(Q_\delta, \partial Q_\delta) \cap E$, the natural projection q_δ: $(F_\delta, \partial F_\delta) \to (Q_\delta, \partial Q_\delta)$ is stably essential in homotopy.*

Proof. We prove here only the necessity of our conditions. The sufficiency part of this theorem follows from Theorems 3.3 and 3.4 below.

Given a metastable set S there exists $0 < \delta_0 < 1$, and a collection $\{W_\delta\}_{0<\delta\le\delta_0}$ satisfying the conditions in Theorem 3.1. Define E to be the closure of $\bigcup_{0<\delta\le\delta_0} \psi_\delta(W_\delta \backslash \partial \Sigma) \backslash \partial \mathscr{F}_\delta$. E is obviously a closed subset of $\mathscr{E}$. We prove that it satisfies the other conditions of the theorem.

We show first that $E_0 = \{0\} \times S$. Observe that $(0, \sigma) \in E_0$ iff there exists a sequence of δ_is converging to zero, and a corresponding sequence (η^i, σ^i) in $\psi_{\delta_i}(W_{\delta_i} \backslash \partial \Sigma) \backslash \partial \mathscr{F}_\delta$ converging to $(0, \sigma)$; this last condition is equivalent to the existence of a sequence of δ_is converging to zero and a corresponding sequence $\sigma^i \in W_{\delta_i} \backslash \partial \Sigma$ converging to σ. By properties 1 and 4 of Theorem 3.1, therefore, $(0, \sigma) \in E_0$ iff $\sigma \in S$.

We turn now to the stable essentiality condition. Fix $0 < \delta \le \delta_0$. By the continuity property for W_δ and Theorem C.8, the projection $\bar{q}_\delta$: $(\bar{F}_\delta, \partial \bar{F}_\delta) \to (Q_\delta, \partial Q_\delta)$ is stably essential, where $\bar{F}_\delta = \psi_\delta(W_\delta)$. Let $\hat{F}_\delta$ be the closure of $\bar{F}_\delta \backslash \partial \bar{F}_\delta$. Then by Lemma A.5, the projection $\hat{q}_\delta$: $(\hat{F}_\delta, \partial \hat{F}_\delta) \to (Q_\delta, \partial Q_\delta)$ is also stably essential. Observe that $\hat{F}_\delta$ is contained in F_δ. Indeed, if $(\eta, \sigma) \in \hat{F}_\delta \backslash \partial \hat{F}_\delta$, then $\sigma \in W_\delta \backslash \partial \Sigma$, $\eta \notin \partial Q_\delta$, and $(\eta, \sigma) \in \psi_\delta(\sigma)$. Therefore, $\hat{F}_\delta \backslash \partial \hat{F}_\delta$, and hence its closure $\hat{F}_\delta$ are contained in F_δ. The stable essentiality of $\hat{q}_\delta$ now implies that of q_δ by Remark A.1. Hence E satisfies the essentiality condition.

Finally, we prove the connectedness condition. Because $0 < \delta_0 < 1$, E is the closure of $E \backslash \partial E_1$. As the E_δs form a basis of neighborhoods of $\{0\} \times S$ in E, it is now sufficient to show that for each $0 < \delta \le \delta_0$, $E_\delta \backslash \partial E_\delta$ is a connected set. Therefore, fix $0 < \delta \le \delta_0$. For each $0 < \delta' < \delta$, because $W_\delta \backslash \partial \Sigma$ is connected and ψ_δ is a well-defined correspondence, $\psi_{\delta'}(W_{\delta'} \backslash \partial \Sigma) \backslash \partial \mathscr{F}_{\delta'}$ is connected. Also, for $0 < \delta' < \delta'' < \delta$, if $(\eta', \sigma) \in \psi_{\delta'}(W_{\delta'} \backslash \partial \Sigma) \backslash \partial \mathscr{F}_{\delta'}$, then $\sigma \in W_{\delta''}$ by property (3) of Theorem 3.1, and there exists $(\eta'', \sigma) \in \psi_{\delta''}(\sigma) \backslash \partial \mathscr{F}_{\delta''}$. Therefore, for each $\lambda \in [0, 1]$,

$$(\lambda \eta'' + (1-\lambda)\eta', \sigma) \in \psi_{\lambda\delta'' + (1-\lambda)\delta'}(W_{\lambda\delta''+(1-\lambda)\delta'} \backslash \partial \Sigma) \backslash \partial \mathscr{F}_{\lambda\delta' + (1-\lambda)\delta}.$$

Hence, $\bigcup_{0<\delta'<\delta} \psi_{\delta'}(W_{\delta'} \backslash \partial \Sigma) \backslash \partial \mathscr{F}_{\delta'}$ is connected. Because E is obtained by taking the closure of $\bigcup_{0<\delta'\le\delta_0} \psi_{\delta'}(W_{\delta'} \backslash \partial \Sigma) \backslash \partial \mathscr{F}_{\delta'}$, we have that $E_\delta \backslash \partial E_\delta$ is connected. Thus the connectedness condition holds. □

By the localization result for homotopic essentiality (Lemma A.5), the stable essentiality condition of the above theorem is weaker than the corresponding requirement in Definition 2.5. Therefore, metastability is weaker than stable essentiality. The next two theorems show that metastable sets are the limits of stably essential sets and thus that there is no substantial gap between these concepts.

Theorem 3.3. *Let E be a subset of $\mathscr{E}$ that satisfies the connectedness and essentiality conditions of Theorem 3.2 and let $S = \{\sigma \mid (0, \sigma) \in E\}$. Then S is the Hausdorff limit of a sequence S^l of semialgebraic stably essential sets. Moreover, the sequence can be chosen such that for each l, S^l has a germ E^l that is semialgebraic and satisfies the following stronger version of the connectedness requirement*: *There exists $\delta_l > 0$ such that for each $0 < \delta \le \delta_l$, $E^l_\delta \backslash \partial E^l_\delta$ is connected and dense in E^l_δ.*

Proof. Triangulate $\mathscr{E}$ such that $\mathscr{E}_0 = p^{-1}(0)$ and $\partial\mathscr{E}_1 = p^{-1}(\partial P_1)$ are full subcomplexes. Let $\widetilde{\mathscr{E}}$ be the union of the simplices of $\mathscr{E}$ that do not intersect $\mathscr{E}_0$. Denote by $\mathscr{E}_0^l$ the lth barycentric subdivision of $\mathscr{E}_0$. $\mathscr{E}_0^l$ and $\widetilde{\mathscr{E}}$ uniquely determine a triangulation $\mathscr{E}^l$ for $\mathscr{E}$.

By the connectedness condition there exists a decreasing sequence V^r of neighborhoods of S in E such that $V^r\backslash\partial E_1$ is connected and dense in V^r. Let $E^{l,r}$ be the union of the simplices of $\mathscr{E}^l$ whose interiors intersect $V^r\backslash\partial E_1$. Obviously V^r is contained in $E^{l,r}$ because $V^r\backslash\partial E_1$ is dense in V^r. For each l, the $E^{l,r}$s form a decreasing sequence in r. Because $\mathscr{E}^l$ is a finite complex there exists $r(l)$ such that for each $r \geq r(l)$, $E^{l,r}$ is constant, say E^l. If $l' \geq l$, then for each r, $E^{l,r} \supseteq E^{l',r}$ and hence $E^l \supseteq E^{l'}$. For each l, E_0^l contains E_0, because E^l contains V^r for large r. Hence, letting $S_0^l = \{\sigma \mid (0,\sigma) \in E^l\}$, we have $S \subseteq \bigcap_l S_0^l$. On the other hand, letting P^l be the set of simplices of $\mathscr{E}_0^l$ that intersect E_0, we have that $E_0^l \subseteq P^l$: indeed, each principal simplex of E^l intersects V^r for large r and hence intersects E_0. Because $\mathscr{E}_0^l$ is a full subcomplex of $\mathscr{E}^l$, the intersection of such a simplex with $\mathscr{E}_0^l$ is a face of the simplex and hence belongs to P^l; thus $E_0^l \subseteq P^l$. The fact that the P^ls form a decreasing sequence converging to E_0 therefore implies that $\bigcap_l E_0^l \subseteq E_0$. Consequently, the S_0^ls converge to S.

Fix l. Both E^l and S_0^l are obviously semialgebraic. To finish the proof, we show that E^l satisfies the stronger form of the connectedness condition in the statement of the theorem and also the essentiality condition in Definition 2.5, which then ensures that S_0^l is stably essential. To obtain the connectedness condition, we use a result in Mertens [22, §2, Theorem 1]: It is sufficient to show that (a) E^l is the closure of $E^l\backslash\partial E_1^l$; and (b) for each $0 < \alpha \leq 1$, the set W_α of points in $E^l\backslash\partial E_1^l$ whose simplicial distance from E_0^l is strictly smaller than α is connected. With regard to (a), since $\partial\mathscr{E}_1$ is a full subcomplex of $\mathscr{E}$, it is the space of a full subcomplex $\partial\mathscr{E}_1^l$ of $\mathscr{E}^l$. Therefore, the intersection of every principal simplex of E^l with $\partial\mathscr{E}_1^l$ is a face of the simplex; moreover, it cannot equal the simplex itself because the simplex intersects $V^r\backslash\partial E_1$ for all large r; hence E^l is the closure of $E^l\backslash\partial E_1^l$. Now we turn to (b). $\mathscr{E}_0^l$ and $\partial\mathscr{E}_1^l$ being full subcomplexes, the intersection of a principal simplex of E^l with $\mathscr{E}_0^l$ and $\partial\mathscr{E}_1^l$ are proper faces of it, hence its intersection with $W_\alpha\backslash\partial E_1^l$ is connected. But, for r large enough, the connected set $V^r\backslash\partial E_1$ is contained in W_α and intersects every principal simplex of E^l. Therefore, $W_\alpha\backslash\partial E_1^l$ is connected. Thus we have established the connectedness condition for E^l.

There remains to prove the essentiality condition for E^l. Because E^l contains V^r for large r, it contains E_δ for all small δ. Choose now a $\delta_0 > 0$ such that E^l contains E_{δ_0} and for each $0 < \delta \leq \delta_0$, F_δ satisfies the essentiality condition in Theorem 3.2. By Remark A.1, q_δ: $(F_\delta^l, \partial F_\delta^l) \to (Q_\delta, \partial Q_\delta)$ is stably essential for all small δ, where $(F_\delta^l, \partial F_\delta^l)$ is the inverse image of $(Q_\delta, \partial Q_\delta)$ in E^l under the natural projection. Because E^l is semialgebraic, by Lemma C.1, p_δ: $(E_\delta^l, \partial E_\delta^l) \to (P_\delta, \partial P_\delta)$ is stably essential for some small δ. Hence, E^l satisfies the essentiality condition in Definition 2.5 as well. □

Theorem 3.4. *The Hausdorff limit of a sequence of stably essential sets is metastable.*

Proof. Let S^l be a sequence of sets converging to a set S such that for each l there exists $E^l \subseteq \mathscr{E}$ such that $E_0^l = \{0\} \times S^l$, and E^l satisfies the essentiality and connectedness conditions of Definition 2.5. Using Lemma A.5, for each l the projection from $(F_\delta^l, \partial F_\delta^l)$ to $(Q_\delta^l, \partial Q_\delta^l)$ is stably essential for δ such that projection from $(E_\delta, \partial E_\delta) \to (P_\delta, \partial P_\delta)$ is stably essential. Therefore, we see that E^l satisfies the conditions of Theorem 3.3 and hence we can assume without loss of generality that E^l satisfies the stronger form of connectedness: for all small $\delta > 0$, $E_\delta^l\backslash\partial E_\delta$ is connected and dense in E_δ. Because the S^ls converge to S, the E_0^ls converge to $E_0 \equiv \{0\} \times S$, and we can now choose a sequence $\delta_l \downarrow 0$ such that the $E_{\delta_l}^l$s converge to E_0 and, for each l, $E_{\delta_l}^l\backslash\partial\mathscr{E}_1^l$ is connected and dense in $E_{\delta_l}^l$, and p_{δ_l}: $(E_{\delta_l}^l, \partial E_{\delta_l}^l) \to (P_{\delta_l}, \partial P_{\delta_l})$ is stably essential.

For each $0 < \delta \leq \delta_0$, let V_δ be the projection of E_δ^l to Σ where l is the unique integer such that $\delta_{l+1} < \delta \leq \delta_l$. Then $V_\delta\backslash\partial\Sigma$ is connected and dense in V_δ for each δ, as it is obtained from the projection of the set $E_\delta^l\backslash\partial E_\delta^l$, which is connected and dense in E_δ^l, by possibly adding some limit points. Also, the V_δs converge to S because in $\mathscr{E}$ the sets E_{δ_l} converge to E_0. Finally, we show that each V_δ satisfies the continuity condition of Definition 2.2. Because the projection from E_δ^l to P_δ is stably essential, by Lemma A.5, the projection from F_δ^l to Q_δ^l is stably essential. Because F_δ^l is contained in the set of $(\eta, \sigma) \in \mathscr{E}$ such that $\bar{\eta} = \delta$ and $\sigma \in V_\delta$, the projection from this latter set is stably essential by Remark A.1. By Theorem C.7, V_δ now satisfies the continuity condition for metastability. Thus S is metastable. □

It is natural to wonder if the Hausdorff limit of stably essential sets is itself stably essential, which would then imply equivalence between metastability and stable essentiality. In the two theorems above—whose proof techniques were borrowed from Mertens [22, §5B], where it is shown that the limit of a sequence of semialgebraic $*$-stable sets is itself $*$-stable—we could, like Mertens, use the approximations E^l to produce a germ E for S that satisfies the connectedness requirement of Definition 2.5. The problem is with the stable essentiality condition. In the case of $*$-stability the fact that Čech cohomology is weakly continuous is used to establish the essentiality condition for E. In our case there seems to be no analog of the following nature. Suppose $(X^l, \partial X^l)$

is a decreasing sequence of compact semialgebraic pairs converging to $(X, \partial X)$ and suppose there is a sequence of stably essential maps p^l: $(X^l, \partial X^l) \to (B, \partial B)$ where for $l > 1$, p^l is the restriction of the map p^1 to X^l. Is it then necessarily the case that the restriction of p^1 to X is also stably essential? While we do not have a counter example, the answer to this question appears to be no. In any event, the above three theorems readily imply the following compactness result for metastability.

Theorem 3.5. *The collection of metastable sets is the Hausdorff closure of the collection of stably essential sets.*

4. Properties of metastable sets. Kohlberg and Mertens [12] and Mertens [21, 22, 23] list a basic set of game-theoretic properties that they argue any reasonable solution concept should satisfy. In this section we show that metastability satisfies all their requirements, except that metastability satisfies a slightly weaker version of their decomposition property and small worlds axiom.

4.1. Basic properties. Because $*$-stable sets exist and are metastable (Theorem 5.1), we get existence for metastability. Also, by definition, metastable sets are connected sets of perfect equilibria. Metastable sets are BR-sets (Definition 5.1) so the proof in Hillas [9] shows that a metastable set contains a proper equilibrium. Because a proper equilibrium induces a quasi-perfect equilibrium, and hence a sequential equilibrium of every extensive-form game with perfect recall having G as its strategic form (Kohlberg and Mertens [12], van Damme [29]) metastability satisfies backward induction. Finally, by Theorem 3.5 the collection of metastable sets is compact in the Hausdorff topology.

4.2. Forward induction and iterated dominance. Kohlberg and Mertens [12] introduce the notion of forward induction by requiring that a solution to a game contain a solution to a game obtained by deleting a strategy that is not a best-reply against any equilibrium in the solution of the original game. Mertens [21] strengthens this property by requiring the solution to survive even under deletion of a strategy that, while possibly optimal against some equilibrium in the solution, is nonetheless inferior in any ε-perfect equilibrium close to the set. Here we prove this property for metastability.

Before stating and proving this property, we describe the setup. Let S be a metastable set of G. For each n, let T_n be a subset of pure strategies (such that T_n is nonempty for some n). Suppose there exists a neighborhood V of S and $\bar{\delta} > 0$ such that for each n, every pure strategy in T_n is an inferior reply against each $\sigma \in V \cap (\Sigma_{\bar{\delta}} \backslash \partial\Sigma)$. We claim that for each n, the strategies in T_n are used with zero probability at each $\sigma \in S$. Indeed, each equilibrium $\sigma \in S$ is perfect and hence is the limit of a sequence σ_ε of ε-perfect equilibria as ε converges to zero. For small enough ε, σ_ε belongs to V because V is a neighborhood of S, and it belongs to $\Sigma_{\bar{\delta}} \backslash \partial\Sigma$ as it completely mixed and ε goes to zero. Hence for all small ε, our assumptions imply that the strategies in T_n are used with probability at most ε for each n. Thus σ assigns zero probability to the strategies in T_n, as claimed.

Let $\bar{G}$ be the game obtained from G by deleting the strategies in T_n for each n. Let $\bar{\Sigma}$ be the face of Σ where for each n, all the strategies in T_n are used with zero probability. Then $\bar{\Sigma}$ can be viewed naturally as the mixed strategy space of the game $\bar{G}$ and by the arguments in the previous paragraph, therefore S can then be viewed as a set of equilibria in $\bar{G}$. The following theorem and its proof use this interpretation of $\bar{\Sigma}$ and S.

Theorem 4.1. *S contains a metastable set of the game $\bar{G}$.*

Proof. By Theorem 3.1 there exists $\delta_0 > 0$ and for each $0 < \delta \le \delta_0$ a closed semialgebraic subset W_δ of Σ satisfying the four properties listed there. Because V is a neighborhood of S and the W_δs converge to S, there is no loss of generality in assuming that W_{δ_0} is contained in V. Furthermore, we can also assume that $\delta_0 \le \bar{\delta}$.

The W_δs are a nested sequence converging to S; and S, as we saw above, is contained in $\bar{\Sigma}$. Therefore, for each $0 < \delta \le \delta_0$, $\overline{W}_\delta \equiv W_\delta \cap \bar{\Sigma}$ is a closed, nonempty, semialgebraic subset of $\bar{\Sigma}$, and the $\overline{W}_\delta$s converge to S. By Theorems C.8, C.10, and 3.4, it is sufficient to prove that for $0 < \delta \le \delta_0$, $\overline{W}_\delta$ satisfies property (2) of Appendix C.

Fix $0 < \delta \le \delta_0$. The set $\bar{P}_\delta$ of δ-perturbations for the game $\bar{G}$ can be viewed as the face of P_δ where the error probability for the strategies in T_n are zero for each n. Fix k, and let $\bar{f}$: $\bar{\Sigma} \times \Delta^k \to \Delta^k$ and $\bar{g}$: $\bar{\Sigma} \times \Delta^k \to \bar{P}_\delta$ be two maps. $\bar{g}$ defines a δ-perturbation $\bar{\varphi}^k_{\bar{g}}$ of the kth suspension $\bar{R}^k$ of the best-reply correspondence $\bar{R}$ in the game $\bar{G}$ as follows: $\bar{\varphi}^k_{\bar{g}}(\sigma, \lambda) = (1 - \bar{\eta})\bar{R}(\sigma) + \bar{g}(\sigma, \lambda)$, where $\bar{\eta}$ is the minimum error probability under $\bar{g}(\sigma, \lambda)$ for each player—see §C.3. By Theorem C.12 it is sufficient to show that $\bar{\varphi}^k_{\bar{g}} \otimes \bar{f}$ has a fixed point in $\overline{W}_\delta \times \Delta^k$. Because Δ^k and $\bar{P}_\delta$ are convex sets, and $\bar{\Sigma}$ is a closed subset of Σ, we can apply Tietze's extension theorem to obtain extensions f: $\Sigma \times \Delta^k \to \Delta^k$ and g: $\Sigma \times \Delta^k \to \bar{P}_\delta$ of $\bar{f}$ and $\bar{g}$, respectively. The map g induces a δ-perturbation φ^k_g of R^k as follows: $\varphi^k_g(\sigma, \lambda) = (1 - \bar{\eta})R(\sigma) + g(\sigma, \lambda)$, where $\bar{\eta}$ is the minimum error probability under $g(\sigma, \lambda)$

for each player. By the continuity property of W_δ, there exists a fixed point (σ, λ) of $\varphi_g^k \otimes f$ such that σ belongs to W_δ. If $\sigma \in \overline{W}_\delta$, then obviously it is also a fixed point of $\bar{\varphi}_g^k \otimes \bar{f}$. Hence to finish the proof we prove the existence of such a fixed point for $\varphi_g^k \otimes f$. Let δ^* be an interior point of P_δ, and for each positive integer l let g^l be the map $g^l(\sigma, \lambda) = (1 - l^{-1})g(\sigma, \lambda) + l^{-1}\delta^*$. The map g^l induces a δ-perturbation $\varphi_{g^l}^k$ of R^k exactly like g. By the continuity property of W_δ now, for each l there exists a fixed point (σ^l, λ^l) of $\varphi_{g^l}^k \otimes f$ such that $\sigma^l \in W_\delta$. Let $\eta^l = g^l(\sigma^l, \lambda^l)$. Because δ^* belongs to the interior of P_δ, η^l belongs to the interior of P_δ. Also, because $g(\sigma, \lambda)$ belongs to $\bar{P}_\delta$, $\eta^l_{n, s_n} = l^{-1}\delta^*_{n, s_n}$ for each n, $s_n \in T_n$. By the definition of $\varphi_{g^l}^k$, there exists $\tau^l \in R(\sigma^l)$ such that $\sigma^l = (1 - \bar{\eta}^l)\tau^l + \eta^l$. Therefore, $\sigma^l \in \Sigma_\delta \backslash \partial\Sigma$. Because W_δ is contained in V and $\delta \le \bar{\delta}$, our assumption on V implies that for each n, and $s_n \in T_n$, $\sigma^l_{n, s_n} = \eta^l_{n, s_n} = l^{-1}\delta^*_{n, s_n}$. By passing to a subsequence if necessary, the limit (σ, λ), which belongs to $W_\delta \times \Delta^k$, is a fixed point of $\varphi_g \otimes f$ where for each player n, the probability of each $s_n \in T_n$ is zero, i.e., $\sigma \in \overline{W}_\delta$. □

Remark 4.1. The proof actually implies a slightly stronger forward induction property. If for each player n the strategies in T_n are not best-replies against any point in W_δ for one of the sets W_δ described in Theorem 3.1—even if some of them are best-replies against points in $\Sigma_\delta \backslash \partial\Sigma$ that are arbitrarily close to S—then deleting these strategies yields a metastable set of the smaller game.

4.3. Ordinality and player-splitting. Kohlberg and Mertens [12] require that a solution be invariant under the addition or deletion of redundant strategies, i.e., a solution depends only on the reduced strategic form of the game obtained by deleting redundant strategies. Subsequently Mertens [24] provides a formal treatment of this notion, generalizing the idea to the concept of ordinality for solution concepts. Here we show that metastability is ordinal in the sense of Mertens. While Mertens considered the class of strategic-form games—where the strategy sets of the players are arbitrary polytopes and the payoff functions are multiaffine—we restrict ourselves here to games in normal form with finite pure strategy sets. Hence our treatment of ordinality is in the context of normal-form games (even though there is an obvious extension of metastability to this general class and ordinality obtains there as well).

Mertens [24, Theorem 2] gives two sufficient conditions for a solution to be ordinal. The following two theorems establish that metastability satisfies them.

A strategy τ_n is an admissible best-reply against a profile σ if there exists a sequence σ^k of completely mixed strategy profiles converging to σ such that τ_n is a best-reply against σ^k for all k. A profile τ is an admissible best-reply against σ if for each n, τ_n is an admissible best-reply against σ. One then obtains an admissible best-reply correspondence for the game that assigns to each σ the set of admissible best-replies.

Theorem 4.2. *Suppose G and $\widetilde{G}$ are two games with the same sets of players and strategies, and they have the same admissible best-reply correspondence. Then they have the same metastable sets.*

Proof. Let S be a metastable set of G and let E be a germ for S satisfying the conditions of Theorem 3.2. (Recall from §2 that we can assume without loss of generality that E is the closure of $E\backslash\partial E_1$.) Given $\delta > 0$ and a strategy profile $\tau \in \Sigma\backslash\partial\Sigma$, observe that $G(\delta\tau)$ and $\widetilde{G}(\delta\tau)$ have the same set of equilibria, since G and $\widetilde{G}$ have the same admissible best-reply correspondence. Therefore, E is also a subset of the graph of the perturbed equilibrium correspondence for the game $\widetilde{G}$. Hence, S is a metastable set of the game $\widetilde{G}$. The result follows from the symmetry between G and $\widetilde{G}$. □

We now state and prove a theorem that implies that metastability is invariant under addition of redundant strategies and also shows that the player-splitting property holds. Before discussing these properties, we present the theorem.

Suppose $\widetilde{G}$ and G are two strategic-form games with strategy spaces $\widetilde{\Sigma}$ and Σ, respectively. Suppose f is a surjective affine mapping from $\widetilde{\Sigma}$ to Σ such that for each $0 \le \delta \le 1$ and $\tilde{\tau} \in \widetilde{\Sigma}$, $\tilde{\sigma}$ is an equilibrium of the perturbed game $\widetilde{G}(\delta\tilde{\tau})$ iff $f(\tilde{\sigma})$ is an equilibrium of $G(\delta f(\tilde{\tau}))$.

Theorem 4.3. *If $\widetilde{S}$ is a metastable set of $\widetilde{G}$, then $f(\widetilde{S})$ is a metastable set of G. If S is a metastable set of G, then $f^{-1}(S)$ is a metastable set of $\widetilde{G}$.*

The proof uses the following lemma, which is a version of the generic local triviality theorem for the case of affine mappings between polyhedra—and because of the formulation it yields a "global triviality" result. Suppose $f\colon X \to Y$ is a surjective and affine map where X and Y are compact and convex polyhedra. Let $d = \dim(X) - \dim(Y)$. Because f is surjective, $d \ge 0$. Indeed, supposing the dimension of Y is k, we can choose vectors $y_0, \ldots, y_k$ in Y that are affinely independent. Then any set of vectors $x_0, \ldots, x_k$ in X, with $f(x_i) = y_i$ for $0 \le i \le k$, is affinely independent. Hence the dimension of X is at least k, i.e., $d \ge 0$. The same argument also shows that f is a homeomorphism if $d = 0$. The lemma now considers the case where $d > 0$.

Lemma 4.1. *There exists a surjective map* $h: Y \times [0,1]^d \to X$ *such that*:
(i) $h(\{y\} \times [0,1]^d) = f^{-1}(y)$ *for all* $y \in Y$ *(in particular, there exist continuous selections from* f^{-1}*)*;
(ii) h *maps* $(Y \backslash \partial Y) \times (0,1)^d$ *homeomorphically onto* $(X \backslash \partial X)$, *where* ∂X *and* ∂Y *are the relative boundaries of* X *and* Y, *respectively.*[7]

Proof of Lemma 4.1. Let k be the dimension of Y. There is no loss of generality in assuming that Y is a subset of $\mathbb{R}^k$. Indeed, let A be the affine space spanned by Y. A is also k-dimensional and there exists a homeomorphism $h: A \to \mathbb{R}^k$ that is affine. We can then replace f and Y with $h \circ f$ and $h(Y')$, respectively. Thus we can assume that Y is a subset of $\mathbb{R}^k$.

Because the projection map from the graph of f to X is an affine homeomorphism, it is sufficient to prove the lemma for the special case that f is a projection map onto, say, the first k coordinates.

We can further assume that X is a full-dimensional polyhedron in $\mathbb{R}^{k+d}$ and that for $0 \le l \le d-1$ the projection of X onto its first $k+l$ coordinates has dimension $k+l$. To see this, suppose X is a polyhedron in $\mathbb{R}^m$ with Y being the projection of X onto the first k coordinates. The affine space B spanned by X has dimension $k+d$. There exists a permutation of the last $m-k$ coordinates such that for $0 \le l \le d$, the projection of B onto its first $k+l$ coordinates is surjective and is actually a homeomorphism for $l = d$. Thus, replacing X with its projection onto its first $k+d$ coordinates, we can assume that X has the stated properties.

We have now reduced the problem to the case where X is a full-dimensional polyhedron in $\mathbb{R}^{k+d}$ and f is the projection of X onto the first k coordinates. Suppose first that $d = 1$. For each $y \in Y$, let $\bar{y}_{k+1}(y)$ and $\underline{y}_{k+1}(y)$ be the maximum and the minimum over $y_{k+1} \in \mathbb{R}$ such that $(y, y_{k+1}) \in X$. Because $f: X \to Y$ is a linear map, $f^{-1}: Y \twoheadrightarrow X$ is a continuous correspondence and, by the maximum theorem, $\bar{y}_{k+1}(y)$ and $\underline{y}_{k+1}(y)$ are continuous functions of y. Define now $h: Y \times [0,1] \to X$ by $h(y, \lambda) = (y, (1-\lambda)\,\underline{y}_{k+1}(y) + \lambda \bar{y}_{k+1}(y))$. h is a surjective map and $h(\{y\} \times [0,1]) = f^{-1}(y)$ for all $y \in Y$. There remains to prove point (2) of the lemma.

We first claim that h maps $(Y \backslash \partial Y) \times (0,1)$ homeomorphically onto its image. To prove this, it is obviously sufficient to prove that for each $y \in Y \backslash \partial Y$, $\underline{y}_{k+1}(y) < \bar{y}_{k+1}(y)$. The projection map from X to Y being linear, $f(X \backslash \partial X) = Y \backslash \partial Y$—see, for instance, Rockefellar [27, Theorem 6.6]. Therefore, for $y \in Y \backslash \partial Y$, $f^{-1}(y) \cap (X \backslash \partial X)$ is nonempty. Because $X \backslash \partial X$ is an open subset of $\mathbb{R}^{k+1}$, $\underline{y}_{k+1}(y) < \bar{y}_{k+1}(y)$, which proves our claim. To finish the proof of point (2) of the lemma we now show that the image of $(Y \backslash \partial Y) \times (0,1)$ under h is $X \backslash \partial X$. Because $\underline{y}_{k+1}(y) < \bar{y}_{k+1}(y)$ for all $y \in Y \backslash \partial Y$, by Rockefellar [27, Theorem 6.6], we have that $(y, y_{k+1}) \in X \backslash \partial X$ iff $y \in Y \backslash \partial Y$ and $\underline{y}_{k+1}(y) < y_{k+1} < \bar{y}_{k+1}$. Thus $h((Y \backslash \partial Y) \times (0,1)) = X \backslash \partial X$ and the proof is complete for the case $d = 1$.

Proceeding to the general case, for each $0 \le l \le d-1$, let Y_l be the projection of X onto its first $k+l$ coordinates, where $Y_0 = Y$. For each l, by construction, Y_l is a full-dimensional polyhedron in $\mathbb{R}^{k+l}$. Therefore, applying the above special case to the projection map from Y_{l+1} to Y_l—with the convention that $Y_d = X$—we obtain a homeomorphism $h^l: Y_l \times [0,1] \to Y_{l+1}$. For each l, define a function $g^l: Y \times [0,1]^{l+1} \to Y_{l+1}$ inductively as follows:

$$g^0 = h^0; \quad \text{for } 0 < l \le d-1, \qquad g^l(y, t_1, \ldots, t_{l+1}) = h^l(g^{l-1}(y, t_1, \ldots, t_l), t_{l+1}).$$

Then the map $h \equiv g^{d-1}$ is the required homeomorphism. □

Proof of Theorem 4.3. f is a surjective affine map and f^{-1} is a continuous correspondence. By Theorem 3.5, therefore, it is sufficient to prove the theorem assuming that $\tilde{S}$ and S are stably essential sets of $\tilde{G}$ and G, respectively. Also, by Theorem 3.3 we can further assume that the relevant sets have semialgebraic germs and satisfy the stronger connectedness condition given there.

For this proof we view $\mathscr{E}$ and $\tilde{\mathscr{E}}$ as graphs of equilibria (rather than perturbed equilibria) of perturbed games for G and $\tilde{G}$, respectively. Let $\xi: \tilde{\mathscr{E}} \to \mathscr{E}$ be the function $\xi(\delta\tilde{\tau}, \tilde{\sigma}) = (\delta f(\tilde{\tau}), f(\tilde{\sigma}))$.

Suppose $\tilde{S}$ is a stably essential set of $\tilde{G}$ with a semialgebraic germ $\tilde{E}$. Then we claim that $f^{-1}(f(\tilde{S}))$ is a stably essential set with $\xi^{-1}(\xi(\tilde{E}))$ as a semialgebraic germ. Indeed: (1) obviously, $\xi^{-1}(\xi(\tilde{E}))$ is a closed semialgebraic set with $f^{-1}(f(\tilde{S})) = \{\tilde{\sigma} \mid (0, \tilde{\sigma}) \in \xi^{-1}(\xi(\tilde{E}))\}$; (2) for each $(\eta, \sigma') \in \xi^{-1}(\xi(\tilde{E}))$, and $\lambda \in [0,1]$, $(\eta, \lambda\sigma' + (1-\lambda)\sigma) \in \xi^{-1}(\xi(\tilde{E}))$, where $(\eta, \sigma) \in \tilde{E}$ and $f(\sigma) = f(\sigma')$, and therefore, $\xi^{-1}(\xi(\tilde{E}))$ satisfies the connectedness condition of Theorem 3.3 because $\tilde{E}$ does; (3) finally, by Remark A.1, $\xi^{-1}(\xi(\tilde{E}))$ satisfies the essentiality condition because it contains $\tilde{E}$, which does. Thus $f^{-1}(f(\tilde{S}))$ is a stably essential set with a semialgebraic germ $\xi^{-1}(\xi(\tilde{E}))$.

Because ξ is a semialgebraic mapping, by our claim in the previous paragraph, to prove the theorem it is sufficient to show that S is a stably essential set of G with a semialgebraic germ E iff $f^{-1}(S)$ is a stably essential set with $\xi^{-1}(E)$ as a semialgebraic germ. To this end, fix a closed semialgebraic set E of $\mathscr{E}$ and denote $\xi^{-1}(E)$

[7] It can actually be shown that that there exist polyhedral subdivisions of X and $Y \times [0,1]^d$ such that the map h is piecewise-affine.

by $\widetilde{E}$. Observe first that $S = \{\sigma \mid (0, \sigma) \in E\}$ iff $f^{-1}(S) = \{\widetilde{\sigma} \mid (0, \widetilde{\sigma}) \in \widetilde{E}\}$. Second, E satisfies the connectedness condition of Theorem 3.3 iff $\widetilde{E}$ does, because of the linearity of f. Finally, we will show that E satisfies the essentiality condition of Definition 2.5 iff $\widetilde{E}$ does too.

Because E and $\widetilde{E}$ are semialgebraic, by Lemma C.1, it is sufficient to prove that E satisfies the essentiality condition of Theorem 3.2 iff $\widetilde{E}$ does. Fix $0 < \delta < 1$. We show that the projection q_δ from F_δ to Q_δ is essential iff $\widetilde{q}_\delta$ from $\widetilde{F}_\delta$ to $\widetilde{Q}_\delta$ is, where F_δ is the set of $(\delta\tau, \sigma) \in E$ and $\widetilde{F}_\delta = \xi^{-1}(F_\delta)$. Because Q_δ and $\widetilde{Q}_\delta$ are homeomorphic to Σ and $\widetilde{\Sigma}$, we view F_δ and $\widetilde{F}_\delta$ as subsets of $\Sigma \times \Sigma$ and $\widetilde{\Sigma} \times \widetilde{\Sigma}$, respectively. Thus F_δ is the set of $(\tau, \sigma) \in \Sigma \times \Sigma$ such that $(\delta\tau, \sigma) \in E_\delta$ and $\widetilde{F}_\delta$ is the set of $(\widetilde{\tau}, \widetilde{\sigma})$ such that $(\delta\widetilde{\tau}, \widetilde{\sigma})$ belongs to $\widetilde{E}_\delta$ (and therefore $(f(\widetilde{\tau}), f(\widetilde{\sigma}))$ belongs to F_δ). We view q_δ and $\widetilde{q}_\delta$ as the projection to the first factor.

Let $k = \dim(\widetilde{\Sigma}) - \dim(\Sigma)$. If $k = 0$, then f is homeomorphism between $\widetilde{\Sigma}$ and Σ and it induces a homeomorphism between $\widetilde{F}_\delta$ and F_δ with the property that these two homeomorphisms commute with the respective projections $\widetilde{q}_\delta$ and q_δ. Thus q_δ is essential iff $\widetilde{q}_\delta$. Assume therefore that $k > 0$. q_δ is stably essential iff the kth suspension of q_δ is. Using Theorem A.8, q_δ is stably essential iff q_δ^k: $(F_\delta, \partial F_\delta) \times (\Delta^k, \partial\Delta^k) \to (\Sigma, \partial\Sigma) \times (\Delta^k, \partial\Delta^k)$ given by $q_\delta^k(\tau, \sigma, \lambda) = (\tau, \lambda)$ is stably essential. Let $\bar{F}_\delta = (\widetilde{\tau}, \sigma) \in \widetilde{\Sigma} \times \Sigma$ be such that $(f(\widetilde{\tau}), \sigma) \in F_\delta$ and let $\bar{q}_\delta$ be the projection from $\bar{F}_\delta$ to $\widetilde{\Sigma}$. By Lemma 4.1 there exists a map h: $(\Sigma, \partial\Sigma) \times (\Delta^k, \partial\Delta^k) \to (\widetilde{\Sigma}, \partial\widetilde{\Sigma})$ whose restriction to $(\Sigma\backslash\partial\Sigma) \times (\Delta^k\backslash\partial\Delta^k)$ is a homeomorphism. The map induces a quotient map h_F from $F_\delta \times \Delta^k$ to $\bar{F}_\delta$ given by $h_F(\tau, \sigma, \lambda) = (h(\tau, \lambda), \sigma)$. And, $h \circ q_\delta^k = \bar{q}_\delta \circ h_F$. Therefore, using Lemma A.6, q_δ^k is stably essential iff $\bar{q}_\delta$ is stably essential.

Observe now that $\widetilde{q}_\delta = \bar{q}_\delta \circ (\mathrm{Id} \times f)$, where Id is the identity function on $\widetilde{\Sigma}$. Therefore, by Lemma A.2, if $\widetilde{q}_\delta$ is stably essential, then so is $\bar{q}_d$ and hence also q_δ. On the other hand, suppose q_δ is stably essential. Then $\bar{q}_\delta$ is stably essential as well. Letting g be a continuous selection from f^{-1}, we have $\bar{q}_\delta = \widetilde{q}_\delta \circ (\mathrm{Id} \times g)$; by Lemma A.2, we now have that $\widetilde{q}_\delta$ is stably essential. Thus q_δ is stably essential iff $\widetilde{q}_\delta$ is. □

This proof can be used to show that if g is a continuous selection from f^{-1}, then S is a metastable set of G iff $g(S)$ is. We are unable to ascertain the following stronger version of this property: $\widetilde{S}$ is metastable iff $f(\widetilde{S})$ is.

The above theorem applies to invariance and player-splitting as follows. Formally, suppose we have two games $\widetilde{G}$ and G with the same player set. Suppose for each player n there exists a surjective affine map f_n: $\widetilde{\Sigma}_n \to \Sigma_n$ such that if f: $\widetilde{\Sigma} \to \Sigma$ is the corresponding map between the spaces of strategy profiles then for each $\widetilde{\sigma} \in \Sigma$ the payoffs of the players in $\widetilde{G}$ are their payoffs in G from $f(\widetilde{\sigma})$. Because f is surjective, we can actually view Σ as a subspace of Σ by choosing, for each n and each pure strategy s_n in G a pure strategy $\tilde{s}_n$ in $\widetilde{G}$, such that $f_n(\tilde{s}_n) = s_n$. Thus $\widetilde{G}$ is obtained from G by adding redundant strategies. The above theorem now relates the solutions of G and $\widetilde{G}$ and yields the invariance property for metastability.

The player-splitting property states the following. In an extensive-form game, if we can partition some player's collection of information sets in such a way that no play of the game intersects more than one element of the partition, then the solution of the game should be the same if we consider the agent-normal form where this player has as many agents as there are elements in the partition. We now formally state this property for metastability.

Suppose one has an N-player extensive-form game in which one can partition some player n's collection $\mathscr{H}_n$ of information sets into two subcollections $\mathscr{H}_{n_1}$ and $\mathscr{H}_{n_2}$, such that no information set in one subcollection follows an information set in the other. Let $\widetilde{G}$ be the strategic form of the game. Consider now a new game G where we "split" player n into two players n_1 and n_2, i.e., the player set in G is $(N\backslash\{n\}) \cup \{n_1, n_2\}$. The strategy sets of the players other than n in $\widetilde{G}$ are the same as in the two games. Each pure strategy $\tilde{s}_n$ of player n in $\widetilde{G}$ prescribes actions at each information set in $\mathscr{H}_{n_i}$ for agent $i = 1, 2$ and thus gives a pure strategy for player n_i in G_n. Let S_{n_i} be player n_i's set of pure strategies in G, and let Σ_{n_i} be the corresponding set of mixed strategies. We now describe the payoff functions for the players. Observe that a pair (s_{n_1}, s_{n_2}) of pure strategies for the agents defines uniquely a pure strategy for player n in $\widetilde{G}$. Therefore, given a profile of pure strategies in G, the payoffs of the players other than the two agents are the payoffs they get from the corresponding profile in $\widetilde{G}$; for agent n_i, let it be n's payoff if the outcome induced by the profile follows an information set in $\mathscr{H}_{n_i}$, and let it be arbitrary otherwise.

For each i there is a well-defined affine function f_{n_i} from $\widetilde{\Sigma}_n$ to Σ_{n_i} that computes for each $\widetilde{\sigma}_n$ the corresponding marginal distribution over S_{n_i}. Let f: $\widetilde{\Sigma} \to \Sigma$ be the map $f(\sigma) = (\widetilde{\sigma}_1, \ldots, \widetilde{\sigma}_{n-1}, f_{n_1}(\widetilde{\sigma}_n), f_{n_2}(\widetilde{\sigma}_n), \widetilde{\sigma}_{n+1}, \ldots \sigma_N)$. Then f satisfies the conditions of Theorem 4.3, and we get the player-splitting property for metastability, in that it does not matter whether one treats the two agents as one player.

4.4. The small-worlds and decomposition properties. Suppose $\widetilde{G}$ is an $\mathcal{N}$-equivalent game. As specified by Mertens [23] the small-worlds axiom requires that solutions of G are precisely the projections of solutions of $\widetilde{G}$. Given Theorem C.13, one might expect metastability to satisfy the small-worlds axiom. As the following theorem shows, however, we obtain a slightly weaker version.

Theorem 4.4. *Let $\widetilde{G}$ be an $\mathcal{N}$-equivalent game. If $\tilde{S}$ is a metastable set of $\widetilde{G}$, then its projection to Σ is a metastable set of G. If S is a metastable set of G, then it contains a metastable set that is the projection of a metastable set of $\widetilde{G}$.*

Proof. Let $\tilde{S}$ be a metastable set of $\widetilde{G}$. Let $\{\widetilde{V}_\delta\}_{0<\delta\le\delta_0}$ be a collection of sets satisfying the conditions in Definition 2.2. For each δ let V_δ be the projection of $\widetilde{V}_\delta$ to Σ. Clearly the V_δs converge to the projection, call it S, of $\tilde{S}$. Also, the V_δs satisfy the connectedness condition because the $\widetilde{V}_\delta$s do. Finally, as for the continuity condition, fix k and $0<\delta\le\delta_k$, where δ_k is as in Definition 2.2. Given a correspondence $\varphi^k\colon \Sigma\times\Delta^k\to\Sigma$, where φ^k is a δ-perturbation of R^k, there is a correspondence $\tilde{\varphi}^k\otimes f\colon \widetilde{\Sigma}\times\Delta^k\to\widetilde{\Sigma}\times\Delta^k$ given by: $(\tilde{\varphi}^k\otimes f)(\sigma,\sigma_{-\mathcal{N}},\lambda)$ is the set of $(\sigma',\sigma'_{-\mathcal{N}},\lambda')$ such that $\sigma'\in\varphi^k(\sigma,\lambda)$, $\sigma'_{-\mathcal{N}}$ is a best-reply for the outsiders against $(\sigma,\sigma_{-\mathcal{N}})$, and $\lambda'\in f(\sigma,\lambda)$. By the continuity property for $\widetilde{V}_\delta$ there exists a fixed point $(\sigma,\sigma_{-\mathcal{N}},\lambda)$ of $\tilde{\varphi}^k\otimes f$ in $\widetilde{V}_\delta$. Then (σ,λ) is a fixed point of $\varphi^k\otimes f$ in V_δ, which shows that V_δ satisfies the continuity property and hence that S is metastable.

To prove the second statement, let S be a metastable set of G. Let $\{W_\delta\}_{0<\delta\le\delta_0}$ be a collection of sets converging to S and satisfying the conditions of Theorem 3.1. Let Σ^o be the mixed strategy space of the outsiders; and let $\widetilde{\Sigma}=\Sigma\times\Sigma^o$ be the mixed strategy space in $\widetilde{G}$. Let l be the dimension of Σ^o. Then for each k, Δ^{k+l} is homeomorphic to $\Sigma^o\times\Delta^k$. Hence, using Theorem C.5, Remark C.2, and Theorem C.2, the continuity condition for metastability of S gives us that for $0<\delta\le\delta_0$ and k, every correspondence $\tilde{\varphi}^k\otimes f\colon \widetilde{\Sigma}\times\Delta^k\to\widetilde{\Sigma}\times\Delta^k$, where $\tilde{\varphi}^k$ is a δ-perturbation of the best-reply correspondence in $\widetilde{G}$, has a fixed point in $W_\delta\times\Sigma^o$. $W_\delta\times\Sigma^0$ satisfies the continuity condition of Definition 2.2 in the game $\widetilde{G}$, i.e., property (2) as defined in Appendix C holds. Because $W_\delta\times\Sigma^0$ is semialgebraic, by Theorem C.8 and Theorem C.10, it contains a stably essential set $\widetilde{S}_\delta$, whose projection to Σ is contained in W_δ. By choosing an appropriate sequence of δs converging to zero for which $\widetilde{S}_\delta$ converges, the compactness property of metastability gives us that $S\times\Sigma^0$ contains a metastable set. By the first part of this theorem, its projection onto Σ is a metastable set of G, which is obviously contained in S. □

As shown in Mertens [24], the collection of q-stable sets (defined in Mertens [22]) satisfies the stronger form of the small-worlds property, namely that they are precisely the projections of q-stable sets of $\mathcal{N}$-equivalent games. Because q-stable sets are metastable as well, it would be interesting to know if the collection of metastable sets that satisfy the stronger property for metastability is exactly the collection of q-stable sets.

A property related to the small-worlds axiom is the decomposition property, which states the following. Suppose G^1 and G^2 are two games played by two sets of players in two different rooms. Suppose G is the composite game $G^1\times G^2$. Then,

(D1) The solutions of G project to solutions of G^1 and G^2;
(D2) The product of solutions to G^1 and G^2 are solutions to G.

Property (D1) is implied by Theorem 4.4, and hence metastability satisfies it. Metastability satisfies the following weaker form of (D2).

Theorem 4.5. *If S^1 is a metastable set of G^1, then there exists a metastable set S^2 of G^2 such that $S^1\times S^2$ is a metastable set of $G^1\times G^2$.*

Proof. By the compactness property for metastability, it is sufficient to prove the result when S^1 is stably essential. Moreover, by Theorem 3.3, we can assume that S^1 has a semialgebraic germ E^1; i.e., $S^1=\{\sigma^1\mid(0,\sigma^1)\in E^1\}$; and there exists $\delta_0>0$ such that for each $0<\delta\le\delta_0$ $E^1_\delta\backslash\partial E^1_\delta$ is connected and dense in E_δ, and the projection from E^1_δ is stably essential.

Let P_1 be the space of perturbations for the game $G^1\times G^2$, and let $\mathscr{E}$ be the graph of the perturbed equilibrium correspondence over it. Let $X=(\eta^1,\eta^2,\sigma^1,\sigma^2)\in\mathscr{E}$ such that $(\eta^1,\sigma^1)\in E^1$. X is a closed semialgebraic set. We claim that $p_{\delta_0}\colon (X_{\delta_0},\partial X_{\delta_0})\to(P_{\delta_0},\partial P_{\delta_0})$ is stably essential. To see this, let V_{δ_0} be the projection of X_{δ_0} onto Σ^1. V_{δ_0} equals the projection of $E^1_{\delta_0}$ onto Σ^1. Because the projection from $E^1_{\delta_0}$ is stably essential, Theorem C.7 tells us that V_{δ_0} satisfies the continuity condition in the definition of metastability for the game G^1. As in the proof of the second statement of Theorem 4.4, this implies—in conjunction with Theorem C.5, Remark C.2, and Theorem C.2—that for each k, each correspondence $\varphi\otimes f\colon (\Sigma^1\times\Sigma^2)\times\Delta^k\twoheadrightarrow(\Sigma^1\times\Sigma^2)\times\Delta^k$ where φ is a δ_0-perturbation of the best reply correspondence of G has a fixed point in $V_{\delta_0}\times\Sigma^2$. By Theorem C.8, $p_{\delta_0}\colon (X_{\delta_0},\partial X_{\delta_0})\to(P_{\delta_0},\partial P_{\delta_0})$ is stably essential as claimed.

By Theorem C.10, $\{\sigma\mid(0,\sigma)\in X\}$ contains a stably essential set, call it S. In fact, using Remark C.4, there exists $0<\delta_1\le\delta_0$ and a semialgebraic subset E of X such that: (i) $E_{\delta_1}\backslash\partial E_{\delta_1}$ is a connected component of $X_{\delta_1}\backslash\partial X_{\delta_1}$ whose closure equals E_{δ_1}; (ii) $p_{\delta_1}\colon (E_{\delta_1},\partial E_{\delta_1})\to(P_{\delta_1},\partial P_{\delta_1})$ is stably essential; and (iii) $S=\{\sigma\mid(0,\sigma)\in E\}$.

Let S_2 be the projection of S onto Σ^2. We show that $S = S^1 \times S^2$; the decomposition property (D1) then shows that S_2 is metastable as well, and we are done.

Choose $\sigma^2 \in S^2$. There exists $\sigma^1 \in S^1$ such that $\sigma \equiv (\sigma^1, \sigma^2) \in S$. Because $E_{\delta_1} \backslash \partial E_{\delta_1}$ is a semialgebraic set that is connected and dense in E_{δ_1}, the Nash curve selection lemma (Bochnak et al. [2, Theorem 8.1.13]) implies $(0, 0, \sigma^1, \sigma^2) = \lim_{\delta \to 0}(\eta^1(\delta), \eta^2(\delta), \sigma^1(\delta), \sigma^2(\delta))$ for some analytic path $(\eta^1(\delta), \eta^2(\delta), \sigma^1(\delta), \sigma^2(\delta))$ in $E_{\delta_1} \backslash \partial E_{\delta_1}$. Because the path is analytic, we can assume without loss of generality that $\bar{\eta}^1(\delta)$ (and hence also $\bar{\eta}^2(\delta)$) equals δ. If $\tilde{\sigma}^1$ is another strategy in S^1, then using the connectedness for property for the semi-algebraic set E^1, $(0, \tilde{\sigma}^1) = \lim_{\delta \to 0}(\tilde{\eta}^1(\delta), \tilde{\sigma}^1(\delta))$ for a path in $E^1_{\delta_1} \backslash \partial E^1_{\delta_1}$. As with the path in E_{δ_1}, we can assume that $\bar{\tilde{\eta}}^1(\delta) = \delta$ along this path as well. Therefore, $(\tilde{\eta}^1(\delta), \eta^2(\delta), \tilde{\sigma}^1(\delta), \sigma^2(\delta))$ is a path in $X_{\delta_1} \backslash \partial X_{\delta_1}$ whose limit is $(0, \tilde{\sigma}^1, \sigma^2)$. Again using the connectedness property for E^1, we have that the two paths $(\eta^1(\delta), \eta^2(\delta), \sigma^1(\delta), \sigma^2(\delta))$ and $(\tilde{\eta}^1(\delta), \eta^2(\delta), \tilde{\sigma}^1(\delta), \sigma^2(\delta))$ belong to the same connected component of $X_{\delta_1} \backslash \partial X_{\delta_1}$. $E_{\delta_1} \backslash \partial E_{\delta_1}$ is a connected component of $X_{\delta_1} \backslash \partial X_{\delta_1}$ and contains the former path; therefore, it contains the latter path, too. Hence $(\tilde{\sigma}^1, \sigma^2)$ belongs to S. Because $\tilde{\sigma}^1$ and σ^2 were arbitrary strategies in S_1 and S_2, respectively, $S = S^1 \times S^2$, and the proof is complete. □

5. Related refinements. For completeness and to provide comparisons, in this section we provide concise definitions of the two refinements most closely related to metastability. The first is the weaker refinement BR-stability defined by Hillas et al. [11]. Like metastability, it is defined in terms of the players' best-reply correspondence. The second is the stronger refinement called stability by Mertens [21]. Although his definition is cast in terms of a game defined by its payoffs, the analog in terms of the best-reply correspondence is established by Govindan and Mertens [5].

The following definition is due to Hillas et al. [11].

Definition 5.1 (BR-set, BR-stable set). A closed subset S of Σ is a best-response set (BR-set) if for every $\varepsilon > 0$ there exists $\delta > 0$, such that each δ-perturbation of R has a fixed point within ε of S. A BR-set is BR-stable if it is a connected set of perfect equilibria.

Hillas et al. [11] show that a BR-stable set satisfies several of the properties listed by Mertens [21]. BR-stability is a weaker refinement than metastability in that the latter invokes the embedding principle described in §1 and thus entails homotopic essentiality. In effect, metastability requires BR-stability for any larger game in which the given game might be embedded.

We turn now to the definition of stability provided by Mertens. Let $\check{H}$ refer to Čech cohomology with integer coefficients.

Definition 5.2 ($*$-Stable Set). $S \subseteq \Sigma$ is an $*$-stable set if for some closed subset E of $\mathscr{E}$ with $E_0 = \{0\} \times S$:

(i) Connectedness: For every neighborhood V of E_0 in E, the set $V \backslash \partial E_1$ has a connected component whose closure is a neighborhood of E_0 in E;

(ii) Cohomological essentiality: $p^*\colon \check{H}^*(P_\delta, \partial P_\delta) \to \check{H}^*(E_\delta, \partial E_\delta)$ is nonzero for some $\delta > 0$.

The following theorem shows that $*$-stability is stronger than metastability.

Theorem 5.1. *$*$-stable sets are metastable.*

Proof. As is shown in Mertens [22, §4E] S is $*$-stable iff there exists a sequence of closed d-dimensional semialgebraic subsets E^l of $\mathscr{E}$, where d is the dimension of P_1, such that for each l, E^l satisfies the essentiality and continuity conditions in Definition 2.3, and the sequence $S^l_0 \equiv \{\sigma \mid (0, \sigma) \in E^l\}$ converges to S. Cohomological essentiality being a stronger requirement than homotopic essentiality, S^l_0 is h-stable with E^l as a germ. Because E^l is d-dimensional, using Lemma A.7 we therefore have that S^l_0 is stably essential and hence metastable. By Theorem 3.5, S is metastable as well. □

As an example, Mertens [22, §4F] shows that requiring essentiality in cohomology is stronger than requiring stable essentiality if the domain has a larger dimension than the range. Thus $*$-stability is indeed a strict refinement of metastability.

Mertens [21, 22] proposes several definitions of stability in which the essentiality requirement is cast in terms of singular homology with coefficients in an Abelian group M. Then he shows that $*$-stability is the union over M of all these refinement concepts and thus is the most inclusive solution concept.

Among the various definitions considered by Mertens, the most natural in some respects is the one where the coefficient group is the group of integers. This concept, called 0-stability, can equivalently be defined exactly like $*$-stability, but with one small difference: The cohomological essentiality is with respect to Čech cohomology with coefficients in the field of rational numbers. We know of no game-theoretic property that suggests a distinguished role for 0-stability to the exclusion of these other concepts. However, in Appendix E we show that 0-stability, $*$-stability, metastability, and h-stability all select exactly the same set of outcomes for generic extensive-form games; see Corollary E.1.

6. Concluding remarks. The refinements defined in §2 and §5 differ chiefly in the formulation of the corresponding version of continuity. As Hillas et al. [11] show, homotopy stability is more restrictive than BR-stability because homotopic essentiality invokes a richer class of perturbations. Stable essentiality is an even stronger requirement because it invokes the embedding principle, including invariance and the projection property of small worlds. (Co)homological essentiality is evidently the strongest criterion—and importantly, unlike homotopy criteria, it ensures that essential maps are surjective. It is this difference that accounts for the slightly weaker form (compared to Mertens' refinement) of the small-worlds property established in Theorem 4.4, and the possible failure of metastability to satisfy (D2) of the decomposition property. However, we show in Appendix E that this difference occurs only for a game whose extensive form has nongeneric payoffs.

For the foundations of game theory, the development of a canonical refinement of Nash equilibria requires one to choose among these topological criteria. This choice must ultimately be guided by decision-theoretic criteria. The results in this paper suggest that the weakest topological criterion that preserves the standard list of decision-theoretic axioms is stable essentiality. Our exposition is cast differently in that we begin with the definition of metastability and its motivation in terms of embedding and continuity, and then establish that this definition is equivalent to stable essentiality of the projection map from the equilibrium graph. But this is the crux of the matter technically.

Our view is that metastability is a viable substitute for Mertens' refinements based on (co)homological essentiality of the projection map, because metastability yields basically the same decision-theoretic properties. From a computational viewpoint, the test for metastability (stable essentiality) is more difficult to apply. Its advantage in applications might therefore lie in its conceptual justification and its agreement with Mertens' stability in generic extensive-form games.

Appendix A. Mathematical background. In this appendix we prove several useful results concerning essential and inessential maps. The main results are Lemmas A.3 and A.4, which show that essentiality of f is equivalent to a strong fixed-point property of f. Also, Lemma A.5 provides a localization result: The restrictions of f to certain subsets of X are also essential. In §A.1 we extend the notion of essentiality to suspensions of f. While the notion of suspension is defined in a traditional manner, the main result of the subsection, Lemma A.8, can be taken to be the definition of the essentiality of the kth suspension of f; in particular, it provides a definition of stable essentiality of f, which is the essentiality of all suspensions of f.

Let $f\colon (X, \partial X) \to (B, \partial B)$ be a map where $(X, \partial X)$ is an arbitrary compact Hausdorff pair and $(B, \partial B)$ is homeomorphic to a ball pair. (We are not assuming here that ∂X is the boundary of X.) f is *inessential in homotopy*—or simply *inessential*, for short—if it is homotopic relative to ∂X to a map to ∂B, i.e., if there exists a map $F\colon (X, \partial X) \times [0, 1] \to (B, \partial B)$ such that for all $x \in X$, $F(x, 0) = f(x)$, $F(x, 1) \in \partial B$, and $F(x, t) = f(x)$ for all $t \in [0, 1]$ if $x \in \partial X$. We say that f is *essential in homotopy*—or simply *essential*, for short—if it is is not inessential.

Mertens [22, §4E, Lemma 2] proves the following equivalent characterizations of inessentiality. We use these characterizations in the paper without referencing this lemma.

Lemma A.1. *The following statements are equivalent.*

- *f is inessential in homotopy.*
- *f is freely homotopic to a map $g\colon X \to \partial B$ as maps between pairs, i.e. there exists $F\colon (X, \partial X) \times [0, 1] \to (B, \partial B)$ such that for each $x \in X$, $F(x, 0) = f(x)$, and $F(x, 1) = g(x)$.*
- *There exists a map $g\colon X \to \partial B$ that agrees with f on ∂X.*

It follows from the second statement of Lemma A.1 that essentiality is a property of the homotopy class of f, i.e., if f is homotopic to another map g as maps between pairs, then f is essential iff g is.

The next lemma gives us an important property of essentiality, which follows easily from the third statement of Lemma A.1.

Lemma A.2. *If f is the composite of two maps $h\colon (X, \partial X) \to (Y, \partial Y)$ and $g\colon (Y, \partial Y) \to (B, \partial B)$, then g is essential if f is.*

Remark A.1. A useful particular case of Lemma A.2 obtains when h is an inclusion map: The mapping from a set is essential if the restriction of the mapping to a subset is already essential.

The next lemma shows that a map that is essential in homotopy has strong fixed-point properties. It is a slight generalization of a lemma in Mertens [22, §3, lemma of Theorem 6].

Lemma A.3. *If f is essential, then every map $g\colon X \to B$ has a point of coincidence with f, i.e., there exists $x \in X$ such that $f(x) = g(x)$. Moreover, if X is metrizable and B is convex, then every correspondence $\varphi\colon X \twoheadrightarrow B$ has a point of coincidence with f.*

Proof. Suppose there exists $g: X \to B$ that has no point of coincidence with f. We claim that f is inessential. Indeed, viewing B as a ball, define a map $h: X \to \partial B$ as follows: For each $x \in X$, $h(x)$ is the unique point in ∂B that is closer to $f(x)$ than $g(x)$ on the line from $g(x)$ through $f(x)$. Clearly h coincides with f on ∂X, and hence f is inessential as claimed.

Assume now the additional hypotheses of the second statement. Using McLennan [18, Proposition 2.25], for each positive integer n there exists a map $g_n: X \to B$ whose graph is within n^{-1} of the graph of φ. By what we have proved, for each n there exists x_n such that $f(x_n) = g_n(x_n)$. Let x be the limit of a convergent subsequence of x_n. Then x is a point of coincidence between f and φ. □

As the following lemma shows, the above coincidence property completely characterizes the essentiality of f in some cases.

Lemma A.4. *Suppose $f(X\backslash\partial X) \subseteq B\backslash\partial B$. If f is inessential, then there exists a map $g: X \to B$ with no point of coincidence with f. Furthermore, the map g can be taken to be such that its image is contained in $B\backslash\partial B$.*

Proof. There is no loss of generality in assuming that B is the unit ball in a Euclidean space. Suppose f is inessential. Then there exists a map $h: X \to \partial B$ that agrees with f on ∂X. Define $g: X \to B$ by letting $g(x)$ be the "antipode" of $h(x)$ in B, i.e. $g(x) = -h(x)$. Clearly f has no point of coincidence with g.

The map g that we have constructed here has all its values in ∂B. But, for a fixed $b \in B\backslash\partial B$, the map sending x to $(1-\delta)g(x) + \delta b$ has no point of coincidence with f for sufficiently small $\delta > 0$ and has all its values in $B\backslash\partial B$, which proves the second statement. □

In Lemma A.4, the property that g can be constructed so as to have its image contained in the interior of B is quite handy—cf., for example, the proof of the next lemma.

We now turn to a localization result that shows the implications of the essentiality of f for the essentiality of restrictions of f to subsets of X. Let $(A, \partial A)$ be a compact pair that is homeomorphic to a ball pair and such that $A \subseteq B$. Let $(W, \partial W)$ be a compact pair such that $W\backslash\partial W = f^{-1}(A\backslash\partial A)$ and $f(\partial W) \subseteq \partial A$. Then the restriction of f to W is a map f_W from $(W, \partial W)$ to $(A, \partial A)$.

Lemma A.5. *If f is essential, then f_W is essential.*

Proof. Suppose f_W is inessential. Because $f(W\backslash\partial W) \subseteq A\backslash\partial A$, Lemma A.4 applies to f_W, and there exists a map g from W to A that has no point of coincidence with f_W and such that $g(W) \subseteq A'$ for some ball A' that is contained in $A\backslash\partial A$. Extend g to a map from X in such a way that all of X gets mapped into A'. Because $f^{-1}(A\backslash\partial A) \subseteq W\backslash\partial W$, this extension does not have a point of coincidence with f, and hence f is inessential by Lemma A.3. □

An important implication of Lemma A.5 is the following result, which we rely on quite heavily in the paper: Suppose $(W, \partial W)$ is a compact subpair of $(X, \partial X)$ such that $X\backslash\partial X = W\backslash\partial W$ and $f(X\backslash\partial X) \subseteq B\backslash\partial B$. Then if f is essential, so is the restriction of f to W. (Of course, the converse is also true, by Lemma A.1.)

We are often interested in quotient spaces X' and B' of X and B obtained by identifying some points in ∂X and ∂B, respectively, in such a way that the map f induces a map f' from X' to B'. Under some conditions, the essentiality of f is equivalent to the essentiality of f'. We present here a lemma that gives such a result. Suppose $(X', \partial X')$ and $(B', \partial B')$ are compact Hausdorff pairs, with the latter pair being homeomorphic to a ball pair. For $Y = X, B$, suppose $q_Y: (Y, \partial Y) \to (Y', \partial Y')$ is a surjective map that sends $Y\backslash\partial Y$ homeomorphically onto $Y'\backslash\partial Y'$. Because Y is compact and Y' is Hausdorff, q_Y is actually a closed map and, in particular, it is a quotient map—under it, Y' is obtained from Y by collapsing some points in ∂Y. Let $f: (X, \partial X) \to (B, \partial B)$ and $f': (X', \partial X') \to (B', \partial B')$ be two maps such that: $f' \circ q_X = q_B \circ f$ and $f'(X'\backslash\partial X') \subseteq B'\backslash\partial B'$.

Lemma A.6. *f is essential iff f' is.*

Proof. Suppose f is inessential. Let $g: X \to \partial B$ be a map that agrees with f on ∂X. Define $g': X' \to \partial B'$ by $g'(x') = q_B(g(q_X^{-1}(x')))$. We will show that g' is a well-defined map that coincides with f' on X', which proves that f' is inessential. First, observe that $g'(x')$ is a unique point in $\partial B'$ for $x' \in X'\backslash\partial X'$ because q_X maps $X\backslash\partial X$ homeomorphically onto $X'\backslash\partial X'$. Next, for the same reason, and also because q_X is surjective, for $x' \in \partial X'$, $\varnothing \neq q_X^{-1}(x') \subseteq \partial X$; therefore, $q_B(g(q_X^{-1}(x'))) = q_B(f(q_X^{-1}(x'))) = f'(q_X(q_X^{-1}(x'))) = f'(x')$. Thus, g' is single-valued and coincides with f' on $\partial X'$. Finally, continuity of g' follows from the fact that q_X is a closed map and from the continuity of q_B and g. Thus g' is a well-defined map and, consequently, f' is inessential.

Suppose f' is inessential. By Lemma A.4 there exists a map $g': X' \to B'$ that does not have a point of coincidence with f' and whose image is contained in $B'\backslash\partial B'$. Then, the map $q_B^{-1} \circ g' \circ q_X$ is well defined and has no point of coincidence with f. Thus f is inessential. □

Remark A.2. The proof that f' is essential if f is does not use the fact that q_X maps $X\backslash\partial X$ homeomorphically onto $X'\backslash\partial X'$, something that is crucial for the proof of the converse.

Remark A.3. One special case of the lemma that we use in this paper occurs when q_X is the identity map. Thus, if a map f' from $(X, \partial X)$ to $(B', \partial B')$ can be written as the composition $q_B \circ f$, then f' is essential iff f is.

A.1. Extension of maps to suspensions. The (unreduced) suspension SX of X is defined as the quotient space of $X \times [0, 1]$ obtained by identifying $X \times \{0\}$ to a point and $X \times \{1\}$ to another point. For each nonnegative integer k, one then defines the kth suspension S^kX of X inductively as follows: $S^0X = X$; and $S^kX = SS^{k-1}X$ for each $k > 0$.

$S\partial X$ can be viewed as a subset of SX if for $i = 0, 1$, we identify the "point" $\partial X \times \{i\}$ with $X \times \{i\}$—thus $S\partial X$ is the set $\partial X \times (0, 1)$ along with the points $X \times \{i\}$ for $i = 0, 1$. Therefore, for each nonnegative integer k, we have the kth suspension of the pair $(X, \partial X)$ given by $S^k(X, \partial X) \equiv (S^kX, S^k\partial X)$. The suspensions of $(B, \partial B)$ are defined analogously. If B is n-dimensional, then SB is an $(n+1)$-ball and $S\partial B$ is an n-sphere. Thus, $S(B, \partial B)$ is an $(n+1)$-ball pair.

One defines the suspension of f as the function Sf: $S(X, \partial X) \to S(B, \partial B)$ given by $Sf(x, t) = (f(x), t)$ for $(x, t) \in X \times (0, 1)$, and $Sf(X \times \{i\}) = B \times \{i\}$ for $i = 0, 1$. Then one defines inductively the map S^kf: $S^k(X, \partial X) \to S^k(B, \partial B)$ as follows: $S^0f = f$; and $S^kf = SS^{k-1}f$ for each $k > 0$.

Definition A.1 (Stably Essential). f is stably essential if, for each k, S^kf is essential in homotopy.

Remark A.4. Suppose S^kf is inessential in homotopy for some k. Let g: $S^kX \to S^k\partial B$ be a map that agrees with f on $S^k\partial X$. It is easily checked that the suspension of g, which is a map from $S^{k+1}X$ to $S^{k+1}\partial B$, agrees with $S^{k+1}f$ on $S^{k+1}\partial X$; and thus $S^{k+1}f$ is also inessential in homotopy. The converse is not true in general—see for instance Mertens [22, §4F]—but the following lemma gives a sufficient condition.

Lemma A.7. *Suppose $(X, \partial X)$ is a CW complex that has the same dimension as $(B, \partial B)$. If f is essential, then it is stably essential.*

Proof. For $k \geq 0$, suppose the kth suspension S^kf of f is essential. We show that $S^{k+1}f$ is essential. (Recall that $S^0f = f$ is essential, so that this would indeed prove the lemma.) $S^k(X, \partial X)$ is obviously a CW complex. Also, if n is the dimension of B then $k + n$ is the dimension of S^kX and S^kB. Let (Y, y_0) be the space obtained from S^kX by collapsing $S^k\partial X$ to a point y_0. Likewise, let (C, c_0) be the corresponding space obtained by collapsing $S^k\partial B$ to a point c_0. Let g: $(Y, y_0) \to (C, c_0)$ be the map induced by S^kf. By Mertens [22, §4.E, theorem], because S^kf is essential, g is not homotopic to the constant map that sends every $y \in Y$ to c_0.

Let (Y_1, y_1) be the quotient space of $S(Y, y_0)$ obtained by collapsing Sy_0 to a point, i.e., the quotient space of $Y \times [0, 1]$ obtained by collapsing $(Y \times 0) \cup (Y \times 1) \cup (y_0 \times [0, 1])$ to a point y_1. Let (C_1, c_1) be the corresponding space obtained from (C, c_0). Let g_1: $(Y_1, y_1) \to (C_1, c_1)$ be the map induced by the suspension Sg of g. By Spanier [28, suspension Theorem 8.5.11], because g is not homotopic to the constant map sending points to c_0, g_1 is also not homotopic to the constant map that sends every $y \in Y_1$ to c_1. Obviously, (Y_1, y_1) is the quotient space of the $(k + n + 1)$-dimensional CW complex $S^{k+1}(X, \partial X)$ obtained by collapsing $S^{k+1}\partial X$ to a point y_1. The same is true of (C_1, c_1). Hence we can again apply Mertens [22, §4.E, theorem] to conclude that $S^{k+1}f$ is essential. ☐

If the dimension of $(X, \partial X)$ is smaller than the dimension of $(B, \partial B)$, then the map f is not even surjective, so it is inessential in homotopy. It is when X has a higher dimension than B that stable essentiality is possibly stronger than essentiality.

When $f(X\backslash\partial X) \subseteq (B\backslash\partial B)$ the following lemma enables us to express essentiality of suspensions without having to go to quotient spaces. It is a very useful lemma to prove the stable essentiality of maps.

Lemma A.8. *Suppose $f(X\backslash\partial X) \subseteq B\backslash\partial B$. For $k > 0$, let $(B^k, \partial B^k)$ be a pair that is homeomorphic to a k-ball pair and let f^k: $(X, \partial X) \times (B^k, \partial B^k) \to (B, \partial B) \times (B^k, \partial B^k)$ be the map $f^k(x, b^k) = (f(x), b^k)$. Then S^kf is essential iff f^k is.*

Proof. It is sufficient to prove the lemma for the case $(B^k, \partial B^k) = ([0, 1]^k, \partial[0, 1]^k)$. Because $f(X\backslash\partial X) \subseteq B\backslash\partial B$, we have that $S^kf(S^kX\backslash S^k\partial X) \subseteq S^kB\backslash S^k\partial B$. For $Y = X, B$, $S^k(Y, \partial Y)$ is a quotient space of $S^{k-1}(Y, \partial Y) \times ([0, 1], \{0, 1\})$ under a quotient map, call it q_Y^k; let $(S^{k-1}f)^1$: $S^{k-1}(X, \partial X) \times ([0, 1], \{0, 1\}) \to S^{k-1}(B, \partial B) \times ([0, 1], \{0, 1\})$ be given by $(S^{k-1}f)^1(x, t) = (S^{k-1}f(x), t)$. Then, $q_B^k \circ (S^{k-1}f)^1 = S^kf \circ q_X^k$, and the conditions of Lemma A.6 hold. Therefore, S^kf is essential iff $(S^{k-1}f)^1$ is essential. Again, using the same lemma, $(S^{k-1}f)^1$ is essential iff the map

$$(S^{k-2}f)^2: S^{k-2}(X, \partial X) \times ([0, 1]^2, \partial[0, 1]^2) \to S^{k-2}(B, \partial B) \times ([0, 1]^2, \partial[0, 1]^2)$$

given by $(S^{k-2}f)^2(x, t_1, t_2) = (S^{k-1}f(x), t_1, t_2)$ is essential. Continuing this downward induction yields the result because $(S^0f)^k$ is the map f^k. ☐

Appendix B. Multisimplicial and polyhedral complexes. The purpose of this appendix is to state the multisimplicial approximation theorem and the relevant properties of polyhedral complexes invoked in Appendix C.

B.1. Multisimplicial complexes. The material of this subsection is based on Govindan and Wilson [8, Appendix B]. We refer the reader to that article for a proof of the multisimplicial approximation theorem stated below as Theorem B.1.

A set of points $\{v_0, \ldots, v_n\}$ in $\mathbb{R}^N$ is affinely independent if the equations $\sum_{i=0}^n \lambda_i v_i = 0$ and $\sum_i \lambda_i = 0$ imply that $\lambda_0 = \cdots = \lambda_n = 0$. An n-simplex K in $\mathbb{R}^N$ is the convex hull of an affinely independent set $\{v_0, \ldots, v_n\}$. Each v_i is a vertex of K and the collection of vertices is called the vertex set of K. Each $\sigma \in K$ is expressible as a unique convex combination $\sum_i \lambda_i v_i$; and for each i, $\sigma(v_i) \equiv \lambda_i$ is the v_ith barycentric coordinate of σ. The interior of K is the set of σ such that $\sigma(v_i) > 0$ for all i. A face of K is the convex hull of a nonempty subset of the vertex set of K.

A (finite) simplicial complex $\mathcal{K}$ in $\mathbb{R}^N$ is a finite collection of simplices in $\mathbb{R}^N$ such that the face of each simplex in $\mathcal{K}$ belongs to $\mathcal{K}$, and the intersection of two simplices is either empty or a face of each. The set V of 0-dimensional simplices is called the vertex set of $\mathcal{K}$. The set given by the union of the simplices in $\mathcal{K}$ is called the space of the simplicial complex and is denoted $|\mathcal{K}|$. For each $\sigma \in |\mathcal{K}|$ there exists a unique simplex K of $\mathcal{K}$ containing σ in its interior; define the barycentric coordinate function σ: $V \to [0, 1]$ by letting $\sigma(v) = 0$ if v is not a vertex of K and otherwise by letting $\sigma(v)$ be the corresponding barycentric coordinate of σ in the simplex K.

A subdivision of a simplicial complex $\mathcal{K}$ is a simplicial complex $\mathcal{K}^*$ such that each simplex of $\mathcal{K}^*$ is contained in a simplex of $\mathcal{K}$ and each simplex of $\mathcal{K}$ is the union of simplices in $\mathcal{K}^*$. Obviously $|\mathcal{K}| = |\mathcal{K}^*|$.

A multisimplex is a set of the form $K_1 \times \cdots \times K_m$, where for each i, K_i is a simplex. A multisimplicial complex $\mathcal{K}$ is a product $\mathcal{K}_1 \times \cdots \times \mathcal{K}_m$, where for each i, $\mathcal{K}_i$ is a simplicial complex. (The vertex set V of a multisimplicial complex $\mathcal{K}$ is the set of all $(v_1, \ldots, v_m)$ such that for each i, v_i is a vertex of $\mathcal{K}_i$. The space of the multisimplicial complex is $\prod_i |\mathcal{K}_i|$ and is denoted $|\mathcal{K}|$. A subdivision of a multisimplicial complex $\mathcal{K}$ is a multisimplicial complex $\mathcal{K}^* = \prod_i \mathcal{K}_i^*$ where for each i, $\mathcal{K}_i^*$ is a subdivision of $\mathcal{K}_i$. In the following, $\mathcal{K} = \mathcal{K}_1 \times \cdots \times \mathcal{K}_m$ is a fixed multisimplicial complex and $\mathcal{L}$ is a fixed simplicial complex.

Definition B.1 (Multisimplicial Map). A map f: $|\mathcal{K}| \to |\mathcal{L}|$ is called multisimplicial if for each multisimplex K of $\mathcal{K}$ there exists a simplex L in $\mathcal{L}$ such that:

(i) f maps each vertex of K to a vertex of L;

(ii) f is multilinear on K; i.e., for each $\sigma \in K$, $f(\sigma) = \sum_{v\in V} f(v) \times \prod_i \sigma_i(v_i)$.

By property (i) of the definition, vertices of K are mapped to vertices of L. Therefore, for each $\sigma \in K$, $f(\sigma)$ is an average of the values at the vertices of K. Because the simplex L is a convex set, the image of the multisimplex K is contained in L. If $\mathcal{K}$ is a simplicial complex, then Definition B.1 coincides with the usual definition of a simplicial map. In this case the image of a (multi)simplex K under f is a simplex of $\mathcal{L}$, but it need not be so in general.

Definition B.2 (Multisimplicial Approximation). Let g: $|\mathcal{K}| \to |\mathcal{L}|$ be a map. A multisimplicial map f: $|\mathcal{K}| \to |\mathcal{L}|$ is a multisimplicial approximation to g if for each $\sigma \in |\mathcal{K}|$, $f(\sigma)$ belongs to the simplex that contains $g(\sigma)$ in its interior.

We could equivalently define a multisimplicial approximation by requiring that for each σ, and each simplex L of $\mathcal{L}$, $g(\sigma) \in L$ implies $f(\sigma) \in L$. The following theorem is the multisimplicial version of the simplicial approximation theorem.

Theorem B.1. *Let g: $|\mathcal{K}| \to |\mathcal{L}|$ be a map. There exists $\eta > 0$ such that for each subdivision $\mathcal{K}^*$ of $\mathcal{K}$ with the property that the diameter of each multisimplex is at most η, there exists a multisimplicial approximation f: $|\mathcal{K}^*| \to |\mathcal{L}|$ of g.*

B.2. Polyhedral complexes. A polyhedral complex $\mathcal{P}$ is a finite collection of polyhedra such that: (i) each face of a polyhedron in $\mathcal{P}$ belongs to $\mathcal{P}$; and (ii) the intersection of two polyhedra in $\mathcal{P}$ is either empty or a face of each of them. The union of the polyhedra in $\mathcal{P}$ is the space of the polyhedral complex and denoted $|\mathcal{P}|$. Every multisimplicial complex, for example, is a polyhedral complex where the polyhedra are the multisimplices.

A polyhedral complex $\mathcal{P}'$ is a polyhedral subdivision of $\mathcal{P}$ if each polyhedron in $\mathcal{P}'$ is contained in a polyhedron of $\mathcal{P}$ and each polyhedron in $\mathcal{P}$ is the union of polyhedra in $\mathcal{P}'$. The following lemma is the basis for defining Player 0's payoff function in Step 2 of the proof of Theorem C.13.

Theorem B.2. *Let $\mathcal{P}$ be a polyhedral complex such that $|\mathcal{P}|$ is a d-dimensional polyhedron in $\mathbb{R}^n$. There exists a polyhedral subdivision $\mathcal{P}'$ of $\mathcal{P}$ and a convex, piecewise-affine function γ: $|\mathcal{P}| \to \mathbb{R}$ such that the maximal convex domains on which γ is affine are the d-dimensional polyhedra in $\mathcal{P}'$.*

Proof. The polyhedral complex $\mathscr{P}'$ is derived from $\mathscr{P}$ as follows Eaves and Lemke [3]. Let $\mathscr{P}_1$ be the set of all $(d-1)$-dimensional polyhedra in $\mathscr{P}$. For each polyhedron $P \in \mathscr{P}_1$, let $H_P = \{z \in \mathbb{R}^n \mid a'_P z = b_P\}$ be the hyperplane that includes P, and if $d < n$ is orthogonal to $|\mathscr{P}|$. Let $\mathscr{P}'_0$ be the set of all polyhedra of the form $|\mathscr{P}| \cap [\bigcap_{P \in \mathscr{P}_1} H_P^i]$, where each $i \in \{+, -\}$ and H_P^+ and H_P^- are the two closed half spaces whose intersection is H_P. P'_0 is a collection of d-dimensional polyhedra whose union is $|\mathscr{P}|$. Let $\mathscr{P}'$ be the polyhedral complex consisting of all the polyhedra that are faces of some polyhedron in $\mathscr{P}'_0$. By construction, $\mathscr{P}'$ is a polyhedral subdivision of $\mathscr{P}$. Associate with $\mathscr{P}'$ the map γ: $|\mathscr{P}| \to \mathbb{R}_+$ for which $\gamma(\sigma) = \sum_{P \in \mathscr{P}_1} |a'_P \sigma - b_P|$. Then γ is convex and piecewise affine. Moreover, the maximal convex domains on which γ is affine are the polyhedra in P'_0, which are the d-dimensional polyhedra of $\mathscr{P}'$. □

Appendix C. Suspensions of the best-reply correspondence. The purpose of this appendix is to obtain results about the continuity condition in the definition of metastability. Theorem C.3, shows that it is sufficient to consider correspondences $\varphi \otimes f$ where f is a function. In §C.1, the continuity condition is related to the stable essentiality conditions of Definition 2.5 and Theorem 3.2 via Theorems C.8–10; and, under an additional assumption, it is shown to be equivalent to a stronger condition in Theorem C.12. In §C.3, so-called CKM perturbations (defined by Hillas et al. [11]) are shown to be sufficient in generating perturbations of the best-reply correspondences for the continuity condition. Finally, in §C.4 we obtain a characterization of the small-worlds axiom in terms of a stably essential projection from the equilibrium graph.

Throughout this section, V is a closed subset of Σ. Consider the following properties.

$(1k)$ There exists $0 < \delta_k \leq 1$ such that for each δ_k-perturbation φ^k of R^k and each correspondence π^k: $\Sigma \times \Delta^k \twoheadrightarrow \Delta^k$, $\varphi^k \otimes \pi^k$ has a fixed point in $V \times \Delta^k$.

(2) There exists $0 < \delta \leq 1$ such that for each k, each δ-perturbation φ^k of R^k, and each correspondence π^k: $\Sigma \times \Delta^k \twoheadrightarrow \Delta^k$, $\varphi^k \otimes \pi^k$ has a fixed point in $V \times \Delta^k$.

Asking for property (2) to hold is obviously stronger than requiring property $(1k)$ to hold for each k, since it is uniform in k.

Remark C.1. Suppose that property $(1k)$ does not hold. Then for $l > k$ property $(1l)$ does not hold. Indeed, fix $\delta > 0$ and suppose φ^k is a δ-perturbation of R^k such that for some π^k, $\varphi^k \otimes \pi^k$ does not have a fixed point in $V \times \Delta^k$. For $l > k$ let ρ^l: $\Delta^l \to \Delta^k$ be a retraction (viewing Δ^k as a face of Δ^l). Define the correspondence φ^l (resp. π^l) from $\Sigma \times \Delta^l$ into Σ (resp. Δ^l) by $\varphi^l(\sigma, \lambda) = \varphi^k(\sigma, \rho^l(\lambda))$ (resp. $\pi^l(\sigma, \lambda) = \pi^k(\sigma, \rho^l(\lambda))$). Then φ^l is a δ-perturbation of R^l and $\varphi^l \otimes \pi^l$ does not have a fixed point in $V \times \Delta^l$.

We begin with some preliminary results about property $(1k)$. The analogous results for property (2) should be clear, and therefore we omit them. Also, while these theorems can be stated so as to give equivalent characterizations of property $(1k)$ we present only one direction of most of these results because their converses are obvious.

Theorem C.1. *Property $(1k)$ holds if there exists $\delta_k > 0$ such that for every δ_k-perturbation φ^k of R^k with $\varphi(\sigma) \subset \Sigma \backslash \partial\Sigma$ for all $\sigma \in \Sigma$, and every correspondence π^k: $\Sigma \times \Delta^k \to \Delta^k$, $\varphi^k \otimes \pi^k$ has a fixed point whose projection to Σ is contained in V.*

Proof. Assume that there exists a $\delta_k > 0$ satisfying the conditions of the theorem. Fix a correspondence $\varphi^k \otimes \pi^k$: $\Sigma \times \Delta^k \to \Sigma \times \Delta^k$ where φ^k is a δ_k-perturbation of R^k. Define $\bar{\varphi}^k$: $\Sigma \times \Delta^k \to \Sigma$ by $\bar{\varphi}^k(\sigma, \lambda) = (1-\delta_k)R(\sigma) + \delta_k \sigma^*$ where σ^* is the profile of uniform mixed strategies. For each $t \in [0,1]$, let φ_t^k be the correspondence $t\bar{\varphi}^k + (1-t)\varphi^k$. φ_t^k is then a δ_k perturbation of φ^k for each t. Moreover for $t > 0$, $\varphi_t^k(\sigma) \subset \Sigma \backslash \partial\Sigma$ for all $\sigma \in \Sigma$ and, therefore, by assumption there exists $\sigma^t \in V$ that is the projection of a fixed point of $\varphi_t^k \otimes \pi^k$. A limit point of σ^t as t goes to zero is then the projection of a fixed point of $\varphi^k \otimes \pi^k$ that belongs to V. □

Theorem C.2. *If property $(1k)$ holds then every correspondence ψ^k: $\Sigma \times \Delta^k \twoheadrightarrow \Sigma \times \Delta^k$, where $\varphi^k \equiv proj_\Sigma \circ \psi^k$ is a δ_k-perturbation of R^k, has a fixed point in $V \times \Delta^k$.*

Proof. Let ψ^k be a correspondence from $\Sigma \times \Delta^k$ to itself such that $\varphi^k \equiv \mathrm{proj}_\Sigma \circ \psi^k$ is a δ_k-perturbation of R^k. Let W be the set of $(\sigma, \lambda) \in V \times \Delta^k$ such that $\sigma \in \varphi^k(\sigma, \lambda)$. W is a closed subset of $V \times \Delta^k$. Moreover, it is nonempty because otherwise the correspondence $\varphi^k \otimes (\mathrm{proj}_{\Delta^k} \circ \psi^k)$ would not have a fixed point in V, contrary to the hypothesis. For $(\sigma, \lambda) \in W$ let $\pi^k(\sigma, \lambda)$ be the set of $\lambda' \in \Delta^k$ such that $(\sigma, \lambda') \in \psi^k(\sigma, \lambda)$. By the definition of W, π^k is nonempty valued. Also, since $\pi^k(\sigma, \lambda)$ equals the projection onto Δ^k of $(\{\sigma\} \times \Delta^k) \cap \psi^k(\sigma, \lambda)$, π^k is compact and convex valued. Finally, upper-semicontinuity of π^k follows from that of ψ^k. Therefore, π^k: $W \to \Delta^k$ is a well-behaved correspondence. For each positive integer m choose a map f_m^k: $W \to \Delta^k$ such that the graph of f_m^k is contained in the m^{-1}-neighborhood of the graph of π^k. Extend f_m^k to a map over $\Sigma \times \Delta^k$, denoting it still by f_m^k. By hypothesis, for each m, $\varphi^k \otimes f_m^k$ has a fixed point in $(\sigma_m, \lambda_m) \in V \times \Delta^k$. The fact that

$\sigma_m \in \varphi^k(\sigma_m, \lambda_m)$ now means that $(\sigma_m, \lambda_m) \in W$. By construction, therefore, $(\sigma_m, \lambda_m, \lambda_m)$ is within $1/m$ of some point $(\sigma'_m, \lambda'_m, \lambda''_m)$ in the graph of π^k. A subsequence of (σ_m, λ_m) converges to, say, (σ, λ). We now have that $(\sigma, \lambda) \in W$ and $\lambda \in \pi^k(\sigma, \lambda)$. Therefore (σ, λ) is a fixed point of ψ^k in $V \times \Delta^k$. $\square$

THEOREM C.3. *Property $(1k)$ holds if there exists $\delta_k > 0$ such that for every δ_k-perturbation φ^k of R^k and every map $f^k: \Sigma \times \Delta^k \to \Delta^k$, $V \times \Delta^k$ contains a fixed point of $\varphi^k \otimes f^k$.*

PROOF. Consider a correspondence $\pi^k: \Sigma \times \Delta^k \twoheadrightarrow \Delta^k$ and a δ_k-perturbation φ^k of R. We show that $\varphi^k \otimes \pi^k$ has a fixed point in $V \times \Delta^k$. For each positive integer m choose a map $f_m: \Sigma \times \Delta^k \to \Delta^k$ whose graph is contained in the m^{-1}-neighborhood of the graph of π^k. By assumption there exists a fixed point $(\sigma_m, \lambda_m) \in V \times \Delta^k$ of $\varphi^k \otimes f_m$. If necessary by passing to a subsequence, the limit of the sequence (σ_m, λ_m) is a fixed point of $\varphi^k \otimes \pi$ that belongs to $V \times \Delta^k$. $\square$

Our next theorem says that the properties we are studying are related to suspensions of R. The metric on $\Sigma \times \Delta^k$ that is used in it is the following: the distance between two points (σ, λ) and (σ', λ') is $\max\{\|\sigma - \sigma'\|, \|\lambda - \lambda'\|_1\}$. (Recall that the metric on Σ is given by $\|\sigma - \tau\| = 0.5 \max_n \|\sigma_n - \tau_n\|_1$.)

THEOREM C.4. *Property $(1k)$ holds if and only if there exists $\delta_k > 0$ such that each δ_k-perturbation of $R^k \otimes Id_{\Delta^k}$ has a fixed point in $V \times \Delta^k$.*

PROOF. The necessity follows from Theorem C.2. As for sufficiency, suppose property $(1k)$ does not hold. Then by Theorem C.3, for each $\delta_k > 0$ there exists a δ_k-perturbation φ^k of R^k and a map $f: \Sigma \times \Delta^k \to \Sigma \times \Delta^k$ such that $\varphi^k \otimes f$ does not have a fixed point in $V \times \Delta^k$. Let $g: \Sigma \times \Delta^k \to \Delta^k$ be the map $g(\sigma, \lambda) = (1 - 0.5\delta_k)\lambda + 0.5\delta_k f(\sigma, \lambda)$. Then $\varphi^k \otimes g$ is a δ_k-perturbation of $R^k \otimes \mathrm{Id}_{\Delta^k}$ that does not have a fixed point in $V \times \Delta^k$. $\square$

REMARK C.2. For each k let $\tilde{\Delta}^k$ be a convex set that is homeomorphic to Δ^k. We could study property $(1\tilde{k})$ obtained by replacing Δ^k with $\tilde{\Delta}^k$ in property $(1k)$. The above theorems are still valid for property $(1\tilde{k})$, and with exactly the same proofs, when we substitute $\tilde{\Delta}^k$ for Δ^k. Furthermore, the following theorem links these two properties.

THEOREM C.5. *Let $\tilde{\Delta}^k$ be a convex set that is homeomorphic to Δ^k and let $\tilde{R}^k: \Sigma \times \tilde{\Delta}^k \twoheadrightarrow \Sigma$ be the correspondence $\tilde{R}^k(\sigma, \tilde{\lambda}) = R(\sigma)$. If property $(1k)$ holds then for each δ_k-perturbation $\tilde{\varphi}^k$ of $\tilde{R}^k$ and each correspondence $\tilde{\pi}^k: \Sigma \times \tilde{\Delta}^k \twoheadrightarrow \Delta^k$ there exists a fixed point of $\tilde{\varphi}^k \otimes \tilde{\pi}^k$ whose projection to Σ is contained in V.*

PROOF. Because Δ^k is homeomorphic to $\tilde{\Delta}^k$, property $(1k)$ implies that for every δ_k-perturbation $\tilde{\varphi}^k$ of $\tilde{R}^k$ and every map $\tilde{f}^k: \Sigma \times \tilde{\Delta}^k \to \tilde{\Delta}^k$ there exists a fixed point whose projection to Σ is contained in V. The theorem now follows by Remark C.2 and the version of Theorem C.3 for $\tilde{\Delta}^k$. $\square$

C.1. Essentiality of projections. In this subsection we show the connection among properties $(1k)$ and (2) and the essentiality of suspensions of the projection map from $\mathscr{E}$, the graph of the perturbed equilibrium correspondence. Throughout this subsection, let $E = \{(\eta, \sigma) \in \mathscr{E} \mid \sigma \in V\}$. To make the domain clear, we write p_δ to denote the projection from $(E_\delta, \partial E_\delta)$ to $(P_\delta, \partial P_\delta)$. $S^k p_\delta$ refers to the kth suspension of p_δ, and $p_\delta^k: (E_\delta, \partial E_\delta) \times (\Delta^k, \partial \Delta^k) \to (P_\delta, \partial P_\delta) \times (\Delta^k, \partial \Delta^k)$ is the map $p_\delta^k((\eta, \sigma), \lambda) = (\eta, \lambda)$. By Lemma A.8, $S^k p_\delta$ is essential in homotopy if and only if p_δ^k is.

For each $0 < \delta \leq 1$, let $Q_\delta = \{\eta \in P \mid \overline{\eta} = \delta\}$ and denote its boundary by ∂Q_δ. Let $(F_\delta, \partial F_\delta) = p^{-1}(Q_\delta, \partial Q_\delta) \cap E_\delta$. Denote the natural projection by $q_\delta: (F_\delta, \partial F_\delta) \to (Q_\delta, \partial Q_\delta)$. Define $S^k q_\delta$ and q_δ^k like their counterparts above. As before, $S^k q_\delta$ is essential in homotopy if and only if q_δ^k is.

Before relating properties $(1k)$ and (2) to the essentiality of suspensions of p_δ and q_δ, we begin with a preliminary result about these projections, whose proof follows directly from Lemma A.5.

THEOREM C.6. *If p_δ^k is stably essential, then so is q_δ^k.*

THEOREM C.7. *Property $(1k)$ holds if $q_{\delta_k}^k$ is essential in homotopy for some $\delta_k > 0$.*

PROOF. Suppose $q_{\delta_k}^k$ is essential in homotopy for $\delta_k > 0$. Let φ^k be a δ_k-perturbation of $R^k: \Sigma \times \Delta^k \twoheadrightarrow \Sigma$ and let $f: \Sigma \times \Delta^k \to \Delta^k$ be a map. By Theorem C.3, it is sufficient to show that $\varphi^k \otimes f$ has fixed point in $V \times \Delta^k$. Define $\psi: F_{\delta_k} \times \Delta^k \twoheadrightarrow Q_{\delta_k} \times \Delta^k$ as follows: $\psi((\eta, \sigma), \lambda)$ is the set of all $(\eta', f(\sigma, \lambda)) \in Q_{\delta_k} \times \Delta^k$ such that there exists $\sigma' \in \Sigma$ such that $\sigma' \in \varphi^k(\sigma, \lambda)$ and $\eta' = \sigma' - (1 - \delta_k)\tau(\eta, \sigma)$. (Recall that $\tau(\eta, \sigma)$ is the equilibrium of $G(\eta)$ that corresponds to the perturbed equilibrium σ, i.e. $\sigma = (1 - \overline{\eta})\tau(\eta, \sigma) + \eta$.) It follows from its definition that ψ is compact and convex valued and upper-semicontinuous. Thus, to show that ψ is a correspondence there remains to check that ψ is nonempty valued. Fix $((\eta, \sigma), \lambda)$. Let $s^1, \ldots, s^{k_1}$ be the set of all pure strategy profiles that are best replies against σ. Because φ^k is a δ_k-perturbation of R^k, there exists for each $1 \le i \le k_1$ a mixed strategy profile $\sigma^i \in \varphi^k(\sigma, \lambda)$ that is within δ_k of s^i (when viewing s^i as a point in Σ). Therefore, $\eta^i \equiv \sigma^i - (1-\delta_k)s^i$ belongs to Q_{δ_k} for all i. Because $\tau(\eta, \sigma)$ is the equilibrium of $G(\eta)$ that corresponds to the perturbed equilibrium σ, it is a best reply against σ in the game G and can thus be expressed as a convex combination $\sum_{i=1}^{k_1} \alpha^i s^i$ of the pure best replies. Because φ^k is convex valued, we now have that $\sum_i \alpha^i \sigma^i$ belongs to $\varphi^k(\sigma, \lambda)$. Moreover,

$$\sum_i \alpha^i \sigma^i - (1-\delta_k)\tau(\eta, \sigma) = \sum_i \alpha^i(\sigma^i - (1-\delta_k)s^i) = \sum_i \alpha^i \eta^i \in Q_{\delta_k}.$$

Hence $(\sum_i \alpha^i \eta^i, f(\sigma, \lambda))$ belongs to $\psi((\eta, \sigma), \lambda)$ and ψ is nonempty valued. Our assumption and Lemma A.3 now imply that ψ has a point of coincidence with $q^k_{\delta_k}$: There exists (η, σ, λ) such that $(\eta, \lambda) \in \psi((\eta, \sigma, \lambda))$ and $\bar{\eta} = \delta_k$. By the definition of ψ, there exists $\sigma' \in \Sigma$ such that $\sigma' \in \varphi^k(\sigma, \lambda)$ and $\eta = \sigma' - (1-\delta_k)\tau(\eta, \sigma) \in Q_{\delta_k}$. Therefore, using the definition of $\tau(\eta, \sigma)$ we have $\eta = \sigma' - (1-\delta_k)(1-\bar{\eta})^{-1}(\sigma - \eta) = \sigma' - \sigma + \eta$. Thus, $\sigma = \sigma' \in \varphi^k(\sigma, \lambda)$. Also, because (σ, η, λ) is a point of coincidence between ψ and $q^k_{\delta_k}$, we obtain $\lambda = f(\sigma, \lambda)$. Then (σ, λ) is a fixed point of $\varphi^k \otimes f$. Because $(\eta, \sigma) \in E$, $\sigma \in V$ by definition, and the proof is complete. □

Using Theorem C.6, the above theorem shows that the essentiality of $p^k_{\delta_k}$ implies property $(1k)$. The following theorem gives a partial converse.

Theorem C.8. *Let d be the dimension of P_1. If property $(1k)$ holds for some $k \ge d$, then $p^{k-d}_{\delta_k}$ is essential in homotopy. If it holds for some $k \ge d-1$, then $q^{k-d+1}_{\delta_k}$ is essential in homotopy.*

Proof. We prove that $p^{k-d}_{\delta_k}$ is essential in homotopy if property $(1k)$ holds for $k \ge d$. The proof for the other case is analogous. By Lemma A.4 it is sufficient to show that every map $f\colon E_{\delta_k} \times \Delta^{k-d} \to P_{\delta_k} \times \Delta^{k-d}$ has a point of coincidence with the map $p^{k-d}_{\delta_k}$. Accordingly, fix such a map f. Extend f to a map from $P_{\delta_k} \times \Sigma \times \Delta^{k-d} \to P_{\delta_k} \times \Delta^{k-d}$, denoting it still by f.

Let $\varphi^k\colon P_{\delta_k} \times \Sigma \times \Delta^{k-d} \twoheadrightarrow \Sigma$ be defined as follows. For each (σ, η, λ), letting η' be the projection of $f(\eta, \sigma, \lambda)$ to P_{δ_k}, $\varphi^k(\sigma, \eta, \lambda)$ is the set of $\tau \ge \eta'$ such that for each $n \in \mathcal{N}$ and $s \in \Sigma^0_n$, $\tau_{n,s} = \eta'_{n,s}$ if strategy s is not a best-reply for player n against σ. Then φ^k is a δ-perturbation of the correspondence $\tilde{R}^k\colon P_{\delta_k} \times \Sigma \times \Delta^{k-d} \twoheadrightarrow \Sigma$ that is defined by $\tilde{R}^k(\eta, \sigma, \lambda) = R(\sigma)$. Because $P_\delta \times \Delta^{k-d}$ is homeomorphic to Δ^k, Theorem C.5 and our assumption imply the existence of a point (η, σ, λ) such that $\sigma \in V$, $\varphi^k(\eta, \sigma, \lambda) = \sigma$, and $(\eta, \lambda) = f(\eta, \sigma, \lambda)$. By the construction of φ^k we have that (η, σ) belongs to E_{δ_k} and hence that (η, σ, λ) is a point of coincidence between $p^{k-d}_{\delta_k}$ and f, which proves that $p^{k-d}_{\delta_k}$ is essential in homotopy. □

Remark C.3. There is a certain asymmetry in the preceding two theorems. While the essentiality of $S^k p$ or $S^k q$ implies the continuity property for R^k, our proof of Theorem C.8 requires the continuity property for R^{k+d-1} (resp. R^{k+d})—not just that of R^k—to obtain the essentiality of $S^k q$ (resp. $S^k p$). It is not obvious to us if continuity of R^k suffices, though we conjecture that it does not.

If property $(1k)$ holds for each k, and when k grows large if δ_k goes to zero, property (2) fails to obtain. We do not have an example exhibiting this phenomenon. Our next theorem gives sufficient conditions when this does not happen.

Theorem C.9. *If V is a semialgebraic set and property $(1k)$ holds for each k, then property (2) holds.*

The proof of this theorem uses the following lemma, which is stated in a slightly more general form here because it is used in §3. Let X be a closed semialgebraic subset of $\mathcal{E}$. For each $0 < \delta \le 1$, let $(Y_\delta, \partial Y_\delta)$ be the inverse image of $(Q_\delta, \partial Q_\delta)$ under the projection map p_δ from X_δ to P_δ, and let q_δ be the projection from Y_δ to Q_δ. For each k, q^k_δ, and p^k_δ are defined exactly as we defined them for the sets F_δ and E_δ, respectively.

Lemma C.1. *There exists $\delta_0 > 0$ such that for each k and $0 < \delta \le \delta_0$, q^k_δ is essential if and only if $q^k_{\delta_0}$ is essential. Moreover, $p^k_{\delta_0}$ is essential if and only if $q^{k+1}_{\delta_0}$ is.*

Proof of Lemma. For player n and each $s_n \in \Sigma^\circ_n$, let F^{n,s_n} be the face of P_1 consisting of the set of η for which $\eta_{n,s_n} = 0$. Because X is semialgebraic, for each F^{n,s_n}, $X^{n,s_n} \equiv X \cap p_1^{-1}(F^{n,s_n})$ is semialgebraic as well. Let $\varepsilon\colon X \to [0,1]$ be the map $\varepsilon(\eta, \sigma) = \bar{\eta}$. By the generic local triviality theorem (Bochnak et al. [2, Proposition 9.3.2]) there exists $\delta_0 > 0$, a semialgebraic fiber C, with for each n, s_n, a closed semialgebraic subset C^{n,s_n} of C, and a homeomorphism $h\colon (0, \delta_0] \times C \to \varepsilon^{-1}(0, \delta_0]$, such that: (i) for each n, s_n, h maps $(0, \delta_0] \times C^{n,s_n}$ into X^{n,s_n}; (ii) h maps $\{\delta\} \times C$ onto Y_δ for $0 < \delta \le \delta_0$.

Let $\partial C = \bigcup_{n,s_n} C^{n,s_n}$. For each $0 < \delta \le \delta_0$, define $h_\delta\colon (C, \partial C) \to (Y_\delta, \partial Y_\delta)$ by $h_\delta(c) = h(\delta, c)$. Then h_δ is a homeomorphism. Obviously, q_δ is essential iff $q_\delta \circ h_\delta$ is. For each k, we now have a homeomorphism $h^k_\delta\colon (C, \partial C) \times (\Delta^k, \partial \Delta^k) \to (Y_\delta, \partial Y_\delta) \times (\Delta^k, \partial \Delta^k)$ given by $h^k_\delta(c, \lambda) = (h_\delta(c), \lambda)$. q^k_δ is essential if and only if $q^k_\delta \circ h^k_\delta$ is.

For $0<\delta<\delta_0$, define f_δ: $(C,\partial C)\to(Q_\delta,\partial Q_\delta)$ by $f_\delta=r_\delta\circ q_{\delta_0}\circ h_{\delta_0}$, where r_δ: $(Q_{\delta_0},\partial Q_{\delta_0})\to(Q_\delta,\partial Q_\delta)$ is the map $r_\delta(\eta)=\delta\eta/\delta_0$. We claim that f_δ is linearly homotopic to $q_\delta\circ h_\delta$ as maps between pairs; indeed, this claim follows from the fact that if $c\in C^{n,s_n}$ for some n, s_n, then its images under $q_\delta\circ h_\delta$ and $q_{\delta_0}\circ h_{\delta_0}$ (and hence also f_δ) belong to F^{n,s_n}, because its images under h_δ and h_{δ_0} belong to X^{n,s_n}. Define f^k_δ: $(C,\partial C)\times(\Delta^k,\partial\Delta^k)\to(Q_\delta,\partial Q_\delta)\times(\Delta^k,\partial\Delta^k)$ by $f^k_\delta(c,\lambda)=(f_\delta(c),\lambda)$. Obviously, f^k_δ is linearly homotopic to $q^k_\delta\circ h^k_\delta$ as maps between pairs. Hence, q^k_δ is essential iff f^k_δ is. Because r_δ is a homeomorphism, it follows readily from the definition of essentiality that f^k_δ (and, therefore, q^k_δ) is essential iff $q^k_{\delta_0}\circ h^k_{\delta_0}$ is essential. Hence, q^k_δ is essential iff $q^k_{\delta_0}$ is.

We now turn to the second statement. As we saw above, $q^{k+1}_{\delta_0}$ is essential iff $q^{k+1}_{\delta_0}\circ h^{k+1}_{\delta_0}$ is. Let $(Z,\partial Z)=([0,\delta_0],\{0,\delta_0\})\times(C,\partial C)$. For each k, let

$$g^k\colon (Z,\partial Z)\times(\Delta^k,\partial\Delta^k)\to([0,\delta_0],\{0,\delta_0\})\times(Q_{\delta_0},\partial Q_{\delta_0})\times(\Delta^k,\partial\Delta^k)$$

be the map given by $g^k(\delta,c,\lambda)=(\delta,q_{\delta_0}(h_{\delta_0}(c)),\lambda)$. By Lemma A.8, $q^{k+1}_{\delta_0}\circ h^{k+1}_{\delta_0}$ is essential iff g^k is. Therefore, $q^{k+1}_{\delta_0}$ is essential iff g^k is essential.

Define $\bar Z$ to be the quotient space of $Z\equiv[0,\delta_0]\times C$ obtained by collapsing $\{0\}\times C$ to a point, call it c_0. And let $\partial\bar Z$ be the image of $((([0,\delta_0])\times\partial C))\cup(\{0,\delta_0\}\times C)$ under the quotient map. Likewise, there is a quotient map from $([0,\delta_0],\{0,\delta_0\})\times(Q_{\delta_0},\partial Q_{\delta_0})$ to $(P_{\delta_0},\partial P_{\delta_0})$ that sends (δ,η) to $(\delta/\delta_0)\eta$. Define $\bar g$: $(\bar Z,\partial\bar Z)\to(P_{\delta_0},\partial P_{\delta_0})$ as follows: $\bar g^k(c_0)=0$; for $0<\delta\le\delta_0$, $c\in C$, $\bar g^k(\delta,c)=g(\delta,c)$. The map $\bar g^k$: $(\bar Z,\partial Z)\times(\Delta^k,\partial\Delta^k)\to(P_{\delta_0},\partial P_{\delta_0})\times(\Delta^k,\partial\Delta^k)$ is the obvious trivial extension. By Lemma A.6, g^k is essential iff $\bar g^k$ is. Thus $q^{k+1}_{\delta_0}$ is essential iff $\bar g^k$ is.

Let $\bar X$ be the quotient space of X_{δ_0} obtained by collapsing $p^{-1}_{\delta_0}(0)$ to a point, x_0, and let $\partial\bar X$ be the image of ∂X_{δ_0} under the quotient map. The map h extends to a homeomorphism $\bar h$ between $(\bar Z,\partial\bar Z)$ and $(\bar X,\partial\bar X)$ by letting $\bar h(c_0)=x_0$. Let $\bar p_{\delta_0}$ be the projection map from $\bar X$ to P_{δ_0}: it maps x_0 to 0 and each other point (η,σ) to η. Then $\bar g$ is linearly homotopic to $\bar p_{\delta_0}\circ\bar h$ as maps between pairs for the same reason that f_δ was linearly homotopic to $q_\delta\circ h_\delta$ earlier on in the proof: the image of a point under one map belongs to a face F^{n,s_n} iff it does so under the other. For each k, define $\bar h^k$: $(\bar Z,\partial\bar Z)\times(\Delta^k,\partial\Delta^k)\to(\bar X,\partial\bar X)\times(\Delta^k,\partial\Delta^k)$ and $\bar p^k_{\delta_0}$: $(\bar X,\partial\bar X)\times(\Delta^k,\partial\Delta^k)\to(P_{\delta_0},\partial P_{\delta_0})\times(\Delta^k,\partial\Delta^k)$ in the obvious way. $\bar p^k_{\delta_0}\circ\bar h^k$ is now linearly homotopic to $\bar g^k$ as maps between pairs. Hence, and also because $\bar h^k$ is a homeomorphism, $q^{k+1}_{\delta_0}$ is essential in homotopy iff $\bar p^k_{\delta_0}$ is. By Lemma A.6, $\bar p^k_{\delta_0}$ is essential iff $p^k_{\delta_0}$ is essential. Thus $q^{k+1}_{\delta_0}$ is essential iff $p^k_{\delta_0}$ is. □

Proof of Theorem C.9. Because V is semialgebraic, E is semialgebraic and the above lemma applies. In particular, there exists $\delta_0>0$ satisfying the conditions given there. By assumption, for each k, there exists δ_k such that property $(1k)$ holds. Without loss of generality we can assume that $\delta_k\le\delta_0$. Using Theorem C.8 and the above lemma, $q^k_{\delta_0}$ is essential for each k. By Theorem C.7, property (2) holds. □

Semialgebraicity of V has an important consequence. If property (2) holds then V contains a stably essential set, as the following theorem in conjunction with Theorem C.8 shows.

Theorem C.10. *Let E' be a semialgebraic subset of E. If p_δ: $(E'_\delta,\partial E'_\delta)\to(P_\delta,\partial P_\delta)$ is stably essential for some $\delta>0$ then $S\equiv\{\sigma\mid(0,\sigma)\in E'\}$ contains a stably essential set.*

Proof. Let X be the closure of $E'\backslash\partial E'_1$. Clearly, X is a compact semialgebraic set and $\{\sigma\mid(0,\sigma)\in X\}\subseteq S$. Moreover, by Lemma A.5, the projection from X_δ is also stably essential for some, and then all, small $\delta>0$. Let ε: $X\to\mathbb{R}$ be the map $\varepsilon(\eta,\sigma)=\bar\eta$. By definition, we have that for each $\delta\ge0$, $\varepsilon^{-1}([0,\delta])=X_\delta$. By Mertens [23, Lemma 2] there exist a positive integer l, a real number $\delta_1>0$, and semialgebraic subsets $X^1,\dots,X^l$ of X such that for each $0<\delta\le\delta_1$: (i) for $i=1,\dots,l$, $X^i_\delta\backslash\partial X^i_\delta$ is connected and dense in X^i_δ; (ii) $i\ne j$ implies $X^i_\delta\cap X^j_\delta\subseteq\partial X_1$; and (iii) $\bigcup_i X^i_\delta=X_\delta$. By property (i), each of the X^is satisfies the connectedness requirement of Definition 2.5. Choose $0<\delta\le\delta_1$ such that the projection from X_δ is stably essential. We claim now that the projection from $(X^i_\delta,\partial X^i_\delta)$ to $(P_\delta,\partial P_\delta)$ is stably essential for some i. Indeed, otherwise, there exists for each i a nonnegative integer k_i such that the k_ith suspension of the projection from $(X^i_\delta,\partial X^i_\delta)$ to $(P_\delta,\partial P_\delta)$ is inessential in homotopy. Let $k=\max_i k_i$. By Remark A.4, the kth suspension of the projection from $(X^i_\delta,\partial X^i_\delta)$ to $(P_\delta,\partial P_\delta)$ is inessential for each i. Hence, for each i there exists a map g^i: $X^i_\delta\times\Delta^k\to(\partial P_\delta\times\Delta^k)\cup(P_\delta\times\partial\Delta^k)$ that coincides with the the projection map over $(\partial X^i_\delta\times\Delta^k)\cup(X^i_\delta\times\partial\Delta^k)$. By properties (ii) and (iii) of the X^is described above, the g^is define a map from X_δ to $(\partial P_\delta\times\Delta^k)\cup(P_\delta\times\partial\Delta^k)$ that coincides with the projection map over $(\partial X_\delta\times\Delta^k)\cup(X_\delta\times\partial\Delta^k)$. That would imply that the projection from X_δ is inessential, which is impossible. Therefore, there exists an i such that X^i satisfies the stable essentiality property of Definition 2.5. Because it also satisfies the connectedness condition, $\{\sigma\mid(0,\sigma)\in X^i\}$ is a stably essential set. □

Remark C.4. As our proof shows, the stably essential set in Theorem C.10 has a semialgebraic germ X^i with the property that, for all small δ, $X^i_\delta\backslash\partial X^i_\delta$ is a connected component of $E'_\delta\backslash\partial E'_\delta$ whose closure is X^i_δ.

C.2. Sufficiency of essential projections. By the results of the previous subsection, checking whether property (2) holds is equivalent to verifying whether p_δ^k or q_δ^k is stably essential for each k, which involves checking the essentiality of an infinity of maps. There is hence the question of whether there exists a k such that the essentiality of the kth suspension of p_δ implies that p_δ is stably essential. We do not know the answer to this question in general. However, we know from Lemma A.7 that there are conditions when the essentiality of p_δ implies its stable essentiality. This result, therefore, yields the following theorem, whose proof is obvious.

THEOREM C.11. *Suppose V is semialgebraic and let d be the dimension of P_1. If the dimension of E is less than d, then property* (2) *fails to hold. If the dimension of E is d, then property* (2) *holds if and only if p_δ is essential for some $\delta > 0$.*

C.3. CKM perturbations. Hillas et al. [11] introduce the notion of *continuous Kohlberg-Mertens perturbations* (CKM perturbations). A CKM perturbation is a map $g\colon \Sigma \to P_1$. Such a map g produces a perturbation φ_g of R defined as follows: $\varphi_g(\sigma)$ is the set of $(1-\overline{g(\sigma)})\tau + g(\sigma)$ such that $\tau \in R(\sigma)$. Analogous to BR-sets, S is a CKM set if for every $\varepsilon > 0$ there exists $\delta > 0$ such that for each $g\colon \Sigma \to P_\delta$, φ_g has a fixed point within ε of S. Hillas et al. [11] show that the CKM sets are exactly the BR-sets.

The results of this section show that one obtains such an equivalence between the two notions in our context. Observe that in the proof of Theorem C.8, we need only a specific type of perturbation of R^k, which are "suspensions" of CKM perturbations. Specifically, given a map g from $\Sigma \times \Delta^k$ to P_δ (or Q_δ), we can generate a perturbation φ_g^k of R^k by letting $\varphi_g^k(\sigma, \lambda)$ be the set of $(1-\overline{g(\sigma,\lambda)})\tau + g(\sigma,\lambda)$, where τ is a best reply to σ. For q_δ to be stably essential, it sufficient that for each such g, and $f\colon \Sigma \times \Delta^k \to \Delta^k$, $\varphi_g^k \otimes f$ has a fixed point in $V \times \Delta^k$. Using Theorem C.7 as well, we therefore have the following theorem.

THEOREM C.12. *Property* (2) *holds if and only if for each k and maps $f\colon \Sigma \times \Delta^k \to \Delta^k$, $g\colon \Sigma \times \Delta^k \to P_\delta$, the correspondence $\varphi_g^k \otimes f$ has a fixed point in $V \times \Delta^k$.*

C.4. The small-worlds projection property. We show here that property (2) is equivalent to a projection property as in the small-worlds axiom. We begin with a definition.

DEFINITION C.1 ($\mathcal{N}$-EQUIVALENT GAME). A finite game $\widetilde{G}$ in strategic form is $\mathcal{N}$-equivalent to G if: (i) the player set of $\widetilde{G}$ includes $\mathcal{N}$; (ii) for each $n \in \mathcal{N}$, his set of pure strategies is Σ_n°; and (iii) the payoffs of each $n \in \mathcal{N}$ depend only on the strategy choices of the players in $\mathcal{N}$ and coincide with his payoffs in G.

In an $\mathcal{N}$-equivalent game $\widetilde{G}$ we refer to the players in $\mathcal{N}$ as the "insiders" and those not in $\mathcal{N}$ as the "outsiders."

THEOREM C.13. *Property* (2) *holds if and only if there exists $\delta > 0$ such that for each game $\widetilde{G}$ that is $\mathcal{N}$-equivalent to G every correspondence of the form $\tilde{\varphi} \otimes \tilde{\pi}$, where $\tilde{\varphi}$ is a δ-perturbation of the best-reply correspondence of the insiders and $\tilde{\pi}$ is the best-reply correspondence of the outsiders, has a fixed point whose projection to Σ is contained in V.*[8]

PROOF. The necessity of the condition follows from Theorem C.5. We now prove the sufficiency part by contradiction. Suppose there exist k, a δ-perturbation φ of R^k, and a map $f\colon \Sigma \times \Delta^k \to \Delta^k$, such that $\varphi \otimes f$ does not have a fixed point in $V \times \Delta^k$. We now construct an equivalent game and a δ-perturbation of its best-reply correspondence that leaves the best-reply correspondence of the outsiders unchanged and that does not have a fixed point projecting to a point in V. This construction is quite involved and is therefore broken into five steps. It uses definitions and results in Appendix B about multisimplices (i.e., products of simplices) and polyhedral complexes.

Step 1 (*Simplicial preliminaries*). Let C be the closed set consisting of those points $(\sigma, \lambda) \in V \times \Delta^k$ such that $\sigma \in \varphi(\sigma, \lambda)$. For each $(\sigma, \lambda) \in C$, $\lambda \neq f(\lambda)$. Therefore, we can choose a number $\alpha > 0$ that is strictly smaller than $\|f(\sigma, \lambda) - \lambda\|$ for all $(\sigma, \lambda) \in C$. Let $\mathcal{K}$ be the simplicial complex obtained by taking a sufficiently fine simplicial subdivision of Δ^k so that each simplex of $\mathcal{K}$ has diameter at most $\alpha/2$. By the multisimplicial approximation theorem (see Theorem B.1), there exists a barycentric subdivision $\mathcal{L}$ of $\mathcal{K}$ and, for each n, a barycentric subdivision $\mathcal{T}_n$ of Σ_n such that the map $f\colon \Sigma \times \Delta^k \to \Delta^k$ has a multisimplicial approximation g from the complex $\mathcal{T} \times \mathcal{L}$ to $\mathcal{K}$, where $\mathcal{T} \equiv \prod_n \mathcal{T}_n$. If $(\sigma, \lambda) \in C$, then $\|f(\sigma, \lambda) - \lambda\| > \alpha$, while $\|f(\sigma, \lambda) - g(\sigma, \lambda)\| \leq \alpha/2$ because the diameter of each simplex of $\mathcal{K}$ is at most $\alpha/2$; therefore, $\|g(\sigma, \lambda) - \lambda\| > \alpha/2$; i.e., there does not exist a simplex of $\mathcal{K}$ that contains both λ and $g(\sigma, \lambda)$. Finally, let $\mathcal{Q}$ be a polyhedral complex that is a

[8] As in our analysis of essential sets of equilibria in Govindan and Wilson [8], there is an open question of whether there is an analog of this theorem in which the quantifiers are reversed, say "for each game $\widetilde{G}$... there exists $\delta > 0$ such that... ." However, as mentioned in Remark 2.1, the order of quantifiers is immaterial to the definition of metastability.

refinement of $\mathcal{T} \times \mathcal{L}$ (viewed as a polyhedral complex) such that there exists a convex, piecewise-affine map $\gamma: \Sigma \times \Delta^k \to \mathbb{R}$ with the property that the maximal convex domains on which γ is affine are the full-dimensional polyhedra of $\mathcal{Q}$—see Theorem B.2.

Step 2 (*Constructing an equivalent game* $\widetilde{G}$). We are now ready to define a game $\widetilde{G}$ that is $\mathcal{N}$-equivalent to G. The set $\widetilde{\mathcal{N}}$ of players in $\widetilde{G}$ is $\mathcal{N} \cup O$, where the set O of outsiders comprises three players, denoted o_1, o_2, o_3.

(a) *The strategy sets.* The strategy sets of players in $\mathcal{N}$ are as before. The mixed strategy sets Σ_{o_1} and Σ_{o_2} of players o_1 and o_2 have as their vertex sets $\Sigma^\circ_{o_1}$ and $\Sigma^\circ_{o_1}$ (these players' pure strategies) the sets of vertices of $\mathcal{L}$ and $\mathcal{K}$, respectively. The mixed strategy set Σ_{o_3} of player o_3 has as its vertex set the class of full-dimensional polyhedra of $\mathcal{Q}$. Observe that each pure strategy $s \in \Sigma^\circ_{o_1}$ corresponds to a unique point $\lambda(s)$ in Δ^k; since Δ^k is a convex set, there is now an affine map h from the set Σ_{o_1} of his mixed strategies to the set Δ^k that sends each mixed strategy to the corresponding average of the vertices, i.e., for each $\sigma_{o_1} \in \Sigma_{o_1}$, $h(\sigma_{o_1}) = \sum_{s\in\Sigma^\circ_{o_1}} \sigma_{o_1,s}\lambda(s) \in \Delta^k$, where $\sigma_{o_1,s}$ is the probability of pure strategy s in σ_{o_1}.

(b) *The payoff functions.* To complete the description of the game we describe the payoff functions. For players in $\mathcal{N}$, their payoffs depend only on the strategy choices of the players in $\mathcal{N}$ and coincide with their payoffs in G.

Player o_1's payoffs depend only on his choice and that of player o_2, and they are defined as follows. For pure strategies $s \in \Sigma^\circ_{o_1}$ and $t \in \Sigma^\circ_{o_2}$, o_1's payoff $\widetilde{G}_{o_1}(s, t)$ is the t-th barycentric coordinate of s in the complex $\mathcal{K}$. (Thus player o_1 tries to mimic o_2's strategy.)

Player o_2's payoffs depend on the strategy choices of all players. They are defined as follows. Suppose player o_3 picks a pure strategy s_{o_3}. Let $T \times L$ be the unique multisimplex of $\mathcal{T} \times \mathcal{L}$ that contains the full-dimensional polyhedron corresponding to s_{o_3}. For each pure strategy $s \in \Sigma^\circ_{o_2}$, define a multilinear function $\widetilde{G}_{o_2,s,s_{o_3}}: T \times L \to \mathbb{R}$ as follows: for each vertex v of $T \times L$, $\widetilde{G}_{o_2,s,s_{o_3}}(v) = 1$ if $g(v) = s$ and zero otherwise. Because $T \times L$ is a full-dimensional multisimplex of $\Sigma \times \Delta^k$, this function extends uniquely to a multilinear function over the whole of $\Sigma \times \Delta^k$, denoted still as $\widetilde{G}_{o_2,s,s_{o_3}}$. Given now a mixed strategy profile $\sigma \in \Sigma$ for players in $\mathcal{N}$, a mixed strategy τ_{o_1} for player o_1, and the pure strategy s_{o_3} for player o_3, player o_2's payoff if he plays a pure strategy s is $\widetilde{G}_{o_2,s,s_{o_3}}(\sigma, h(\tau_{o_1}))$. Obviously, player o_2's payoff function is multilinear in the strategies of his opponents.

Player o_3's payoffs are a linear function of the strategies of players in $\widetilde{\mathcal{N}}\backslash\{o_2\}$ and are defined as follows. Let s be a pure strategy of player o_3. Let $Q \subset \Sigma \times \Delta^k$ be the full-dimensional polyhedron corresponding to s. The restriction of γ to Q is affine, and because Q is full-dimensional, admits a unique affine extension—call it $\widetilde{G}_{o_3,s}$—to the whole of $\Sigma \times \Delta^k$. Player o_3's payoff function is given by $\widetilde{G}_{o_3}(s_{o_3}, \tilde{\sigma}_{-o_3}) = \widetilde{G}_{o_3,s_{o_3}}(\sigma, \lambda)$, where $\lambda = h(\tilde{\sigma}_{o_1})$ and σ is the projection of $\tilde{\sigma}$ to Σ.

The description of $\widetilde{G}$ is now complete. By construction $\widetilde{G}$ is $\mathcal{N}$-equivalent to G. Let $\tilde{R}$ and $\tilde{\pi}$ be the best-reply correspondences of the insiders and the outsiders in the game $\widetilde{G}$.

Step 3 (*Analyzing* $\tilde{\pi}$). Let σ^* be a mixed strategy profile and let σ be the projection of σ^* to Σ. Let $\lambda = h(\sigma^*_{o_1})$. The following lemma summarizes the relevant aspects of $\tilde{\pi}(\sigma^*)$.

Lemma C.2. *For each outsider m, let Σ^*_m be the support of σ^*_m. Then*:

(i) *Suppose that the vertices in $\Sigma^*_{o_2}$ span a simplex K^* in $\mathcal{K}$. Then $h(s)$ belongs to K^* for each pure optimal reply s of player o_1.*

(ii) *Suppose that each $s \in \Sigma^*_{o_3}$ contains (σ, λ). Each pure optimal reply for o_2 against σ^* is a vertex of the unique simplex of $\mathcal{K}$ that contains $g(\sigma, \lambda)$ in its interior.*

(iii) *Player o_3's set of pure optimal replies to σ^* is the class of all polyhedra that contain (σ, λ).*

Proof of Lemma C.2.

1. Player o_1's payoff, if he plays $s_{o_1} \in \Sigma^\circ_{o_1}$ against σ^* is $\sum_{s\in\Sigma^*_{o_2}} \sigma^*_{o_2,s} \times s_{o_1}(s)$, where $\sigma_{o_2,s}$ is the probability of s in $\sigma^*_{o_2}$ and $s_{o_1}(s)$ is the sth barycentric coordinate of s_{o_1} in the complex $\mathcal{K}$. By assumption, the support of σ_{o_2} spans a simplex in $\mathcal{K}$, so it follows that a pure strategy s_{o_1} is an optimal reply if and only if $s_{o_1}(s) = 0$ for all $s \notin \Sigma^{**}_{o_2}$, where $\Sigma^{**}_{o_2} \subseteq \Sigma^*_{o_2}$ is the subset of pure strategies of o_2 that are assigned the highest probability under σ_{o_2}. Because the vertices in $\Sigma^{**}_{o_2}$ span a simplex that is a face of K^*, for each pure optimal reply s for o_1, $h(s)$ belongs to K^*.

2. Let $T^* \times L^*$ be the unique multisimplex of $\mathcal{T} \times \mathcal{L}$ that contains (σ, λ) in its interior. Let K' be the unique simplex of $\mathcal{K}$ that contains $g(\sigma, \lambda)$ in its interior. By the construction of player o_2's payoff function, for each polyhedron $s_{o_3} \in \Sigma^\circ_{o_3}$ that contains (σ, λ) the payoff to player o_2 from playing $s_{o_2} \in \Sigma^\circ_{o_2}$, if player o_3 plays s_{o_3} and all others play according to σ^*, is positive if s_{o_2} is a vertex of K' and is zero otherwise. Because, by assumption, each $s_{o_3} \in \Sigma^*_{o_3}$ contains (σ, λ), each optimal reply for o_2 is a vertex of K'.

3. By construction, for each pure strategy s of player o_3, $G^*_{o_3}(s, \sigma^*_{-o_3}) \le \gamma(\sigma, \lambda)$ where the inequality is strict unless the polyhedron s contains (σ, λ). Thus, player o_3's set of pure optimal replies is the class of polyhedra that contain (σ, λ). □

Step 4 (*Perturbing* $\tilde R$). Define a δ-perturbation $\tilde\varphi$ of $\tilde R$ as follows. For each mixed strategy profile σ^* in the game $\widetilde G$, $\tilde\varphi(\sigma^*) = \varphi(\sigma, \lambda)$, where σ is the projection of σ^* to Σ and $\lambda = h(\sigma^*_{o_1})$. Because φ is a δ-perturbation of R^k, $\tilde\varphi$ is indeed a δ-perturbation of $\tilde R$.

Step 5 (*Fixed points of* $\tilde\varphi \otimes \tilde\pi$). To finish the proof we show that $\tilde\varphi \otimes \tilde\pi$ has no fixed point whose projection to Σ is contained in V. Let σ^* be a fixed point of $\tilde\varphi \otimes \tilde\pi$, and let σ be the projection of σ^* into Σ. Also, let $\lambda = h(\sigma^*_{o_1})$. Because σ^* is a fixed point of $\tilde\varphi \otimes \tilde\pi$, $\sigma \in \varphi(\sigma, \lambda)$. Using the definition of C, it is sufficient to show that $(\sigma, \lambda) \notin C$.

For each outsider m, let Σ^*_m be the support of σ^*_m. Also, let K be the unique simplex of $\mathscr{K}$ that contains $g(\sigma, \lambda)$ in its interior. By property 3 of Lemma C.2, (σ, λ) belongs to each $s \in \Sigma^*_{o_3}$. Hence, by property 2, the vertices in $\Sigma^*_{o_2}$ span a simplex K^* that is a face of K. By property 1 now, for each $s \in \Sigma^*_{o_1}$, $h(s)$ belongs to K^*. Hence, $\lambda \in K^* \subseteq K$. Because $g(\sigma, \lambda)$ belongs to K as well, $(\sigma, \lambda) \notin C$. (Recall from step 1 that for a point $(\sigma, \lambda) \in C$, λ and $g(\sigma, \lambda)$ belong to different simplices.) Thus we have shown that $\tilde\varphi \otimes \tilde\pi$ has no fixed point projecting to a point in V. □

Appendix D. Spaces of perturbations for stably essential sets. As promised in §2, we show in this appendix that alternative specifications of the space of strategy perturbations yield equivalent definitions of stably essential sets. There are two results here. The first concerns perturbations that are specified as error probabilities, albeit slightly different from what is used in the main text. The second uses the games induced by these perturbations as the relevant space.

We begin with three spaces of perturbations that are closely related to P. For $0 < \delta \le 1$,

$$\text{let } \hat P_\delta = [0, \delta] \times \Sigma; \qquad P'_\delta = \{(\delta_n \tau_n)_{n \in \mathcal{N}} \mid \forall\, n, 0 \le \delta_n \le \delta, \tau_n \in \Sigma_n\}; \qquad \hat P'_\delta = [0, \delta]^N \times \Sigma.$$

We suppress the subscript δ for $\delta = 1$. Each perturbation in these spaces generates a perturbed game in the obvious way. For $(\delta, \tau) \in \hat P_\delta$ the perturbed game $G(\delta, \tau)$ is the game $G(\delta\tau)$; for $\eta' \in P'$, the perturbed game is the game—call it $G'(\eta')$—where for each player the payoff from the strategy profile $\sigma \in \Sigma$ is the payoff in G from the profile $((1 - \bar\eta'_n)\sigma_n + \eta'_n)_{n \in \mathcal{N}}$, with $\bar\eta'_n \equiv \sum_{s_n \in S_n} \eta'_{n, s_n}$; for $(\delta, \tau) \in \hat P'$, the perturbed game $G'(\delta, \tau)$ is the game $G'((\delta_n \tau_n)_{n \in \mathcal{N}})$. The perturbed equilibria for these games are defined analogously. Let $\hat{\mathscr{E}}$, $\mathscr{E}'$, and $\hat{\mathscr{E}}'$ be the graphs of the perturbed equilibrium correspondence over $\hat P$, P', and $\hat P'$, respectively. The projection maps to the spaces of perturbations are denoted $\hat p$, p', and $\hat p'$. For each δ and each subset $\hat E$ of $\hat{\mathscr{E}}$, $(\hat E_\delta, \partial \hat E_\delta) = \hat p^{-1}(\hat P_\delta, \partial P_\delta) \cap \hat E$. We define such pairs similarly for the other two spaces, $\mathscr{E}'$ and $\hat{\mathscr{E}}'$.

With these notations in place, we can define stably essential sets precisely as in Definition 2.5 except that we use one of these three other spaces of perturbations and consider subsets of the graph over them.

Theorem D.1. *The collection of stably essential sets remains the same if we use $\hat P$ or P' or $\hat P'$ as the space of perturbations.*

Proof. We first consider the space P' and begin with some definitions. Let $\tilde\Delta = \{\tilde\lambda \in \mathbb{R}^N \mid 0 \le \tilde\lambda_n \le 1\, \forall\, n, \min_n \tilde\lambda_n = 0\}$ and denote by $\partial\tilde\Delta$ the set of $\tilde\lambda \in \tilde\Delta$ such that $\tilde\lambda_n = 1$ for some n. $(\tilde\Delta, \partial\tilde\Delta)$ is homeomorphic to the $(N-1)$-dimensional simplicial pair $(\Delta^{N-1}, \partial\Delta^{N-1})$. Fix $0 < \tilde\delta < 1$ and define π: $P_{\tilde\delta} \times \Sigma \times \tilde\Delta \to P'_{\tilde\delta}$ as follows: For each player n his perturbation vector under π is given by $\pi_n(\eta, \sigma, \tilde\lambda) = \eta_n + \tilde\lambda_n(\eta_n - \sigma_n)\beta_n(\eta, \sigma)$, where $\beta_n(\eta_n, \sigma_n) = \min\{(\sigma_{n, s_n} - \eta_{n, s_n})^{-1}\eta_{n, s_n} \mid s_n \in S_n, \sigma_{n, s_n} > \eta_{n, s_n}\}$. Because $\bar\eta_n \le \tilde\delta < 1$, $\beta_n(\eta_n, \sigma_n)$ is well defined and continuous at (η_n, σ_n) if either $\sigma_n \gg 0$ or $\eta_n \gg 0$. Therefore, π is continuous at all points $(\eta, \sigma, \tilde\lambda)$ such that either $\eta \gg 0$ or $\sigma \gg 0$. π defines a map h: $(\mathscr{E}_{\tilde\delta} \backslash \partial\mathscr{E}_{\tilde\delta}) \times (\tilde\Delta \backslash \partial\tilde\Delta) \to \mathscr{E}'_{\tilde\delta} \backslash \partial\mathscr{E}'_{\tilde\delta}$, given by $h(\eta, \sigma, \lambda) = (\pi(\eta, \sigma, \lambda), \sigma)$. It is easily checked that h is a homeomorphism. Indeed, to compute the inverse image $(\eta, \sigma, \tilde\lambda)$ of $(\eta', \sigma, \lambda) \in \mathscr{E}'_{\tilde\delta} \backslash \partial\mathscr{E}'_{\tilde\delta}$, first compute each η_n as follows: $\eta_n = \eta'_n$ if $\bar\eta_n = \max_m \bar\eta'_m$; otherwise, $\eta_n = (1 - \bar\eta'_n)^{-1}((1 - \bar\eta'_m)\eta'_n + (\bar\eta'_m - \bar\eta'_n)\sigma)$, where m is some player for whom $\bar\eta_m = \max_{m'} \bar\eta'_{m'}$. Then compute $\tilde\lambda_n = (1 - \bar\eta_n)^{-1}(\bar\eta_n - \bar\eta'_n)(\beta_n(\eta, \sigma))^{-1}$. An important property of h is that for each $0 < \delta \le \tilde\delta$, it maps $(\mathscr{E}_\delta \backslash \partial\mathscr{E}_\delta) \times (\tilde\Delta \backslash \partial\tilde\Delta)$ onto $\mathscr{E}'_\delta \backslash \partial\mathscr{E}'_\delta$.

Suppose for some $S \subseteq \Sigma$ there exists a germ E' in $\mathscr{E}'$ satisfying the essentiality and connectedness conditions. We will prove that S is stably essential as in Definition 2.5. There is no loss of generality in assuming that E' is contained in $\mathscr{E}'_{\tilde\delta}$ and that it is the closure of $E' \backslash \partial E'_1$. Let E be the closure of $\mathrm{proj}_{\mathscr{E}} \circ h^{-1}(E' \backslash \partial E'_1)$, where $\mathrm{proj}_{\mathscr{E}}$ is the natural projection from $\mathscr{E} \times \tilde\Delta \to \mathscr{E}$. For each $0 < \delta \le \tilde\delta$, $\mathrm{proj}_{\mathscr{E}} \circ h^{-1}$ maps $(E'_\delta \backslash \partial E'_\delta)$ onto $E_\delta \backslash \partial E_\delta$. Because S is the set of σ such that $(0, \sigma) \in E'$, we therefore have that $S = \{\sigma \mid (0, \sigma) \in E\}$. As for the connectedness condition, let $\mathscr{W}'$ be a basis of neighborhoods W' of E'_0 in E' such that $W'\backslash\partial E'_1$ is connected. For each $W' \in \mathscr{W}$, let W be the closure of $\text{proj}_{\mathscr{E}} \circ h^{-1}(W'\backslash\partial E'_1)$. Then $W\backslash\partial E_1$ is connected because it is obtained from $\text{proj}_{\mathscr{E}} \circ h^{-1}(W'\backslash\partial E'_1)$ by adding some limit points. Also, the Ws form a basis of neighborhoods of E_0 in E. Indeed, because the sets in $\mathscr{W}'$ are neighborhoods of E'_0 in E', for each $W' \in \mathscr{W}'$ there exists $\delta > 0$ such that E'_δ is contained in the interior of W' in E; E_δ is then contained in W, making W a neighborhood of E_0 in E. Moreover, because the W's form a basis of neighborhoods of E'_0 in E', for each $\delta > 0$ there exists W' such that W' is contained in E'_δ; obviously W is then contained in E_δ. Thus, the Ws form a basis of neighborhoods of E_0 in E, which then satisfies the connectedness condition. Finally, as for stable essentiality, fix $\delta > 0$ such that the projection map p'_δ from $(E'_\delta, \partial E'_\delta)$ to $(P'_\delta, \partial P'_\delta)$ is stably essential. Let $\widetilde{E}'_\delta$ be the closure of the set of $(\eta, \sigma) \in E'_\delta$ such that $\eta \notin \partial P'_1$ and $\bar{\eta}_n = \bar{\eta}_m$ for all $n, m \in \mathscr{N}$. P_δ being a subset of P'_δ, an application of Lemma A.5 implies that the projection from $(\widetilde{E}'_\delta, \partial\widetilde{E}'_\delta)$ to $(P_\delta, \partial P_\delta)$ is stably essential, where $\partial\widetilde{E}'_\delta$ is the inverse image of ∂P_δ under p'_δ in E'. Obviously $\widetilde{E}'_\delta\backslash\partial\widetilde{E}'_\delta$, and hence also its closure $\widetilde{E}'_\delta$ is a subset of E_δ. Therefore, the projection p_δ from $(E_\delta, \partial E_\delta)$ to $(P_\delta, \partial P_\delta)$ is stably essential. Thus we have shown that S is stably essential using our original definition.

To prove the converse, suppose S is stably essential with a germ E. We can assume that E is the closure of $E\backslash\partial E_1$. Let E' be the closure of $h((E\backslash\partial E_1) \times (\tilde{\Delta}\backslash\partial\tilde{\Delta}))$. Obviously, we have $S = \{\sigma \mid (0, \sigma) \in E'\}$. The proof of the connectedness condition parallels that given for the other direction. For each W in basis of neighborhoods $\mathscr{W}$ of E_0 in E such that $W\backslash\partial E_1$ is connected, let W' be the closure of $h((W\backslash\partial E_1) \times (\tilde{\Delta}\backslash\partial\tilde{\Delta}))$. Then $W'\backslash\partial E'_1$ is connected since it equals $h((W\backslash\partial E_1) \times (\tilde{\Delta}\backslash\partial\tilde{\Delta}))$. Also the fact that $h((E_\delta\backslash\partial E_\delta) \times (\tilde{\Delta}\backslash\partial\tilde{\Delta})) = E'_\delta\backslash\partial E'_\delta$ for all $\delta \leq \tilde{\delta}$ implies that the collections of these sets W' form a basis of neighborhoods by an argument similar to that in the previous paragraph. There remains to prove that E'_δ satisfies the essentiality condition, which we do below.

Fix now δ such that projection p_δ from E_δ is stably essential. Let σ^* be the profile of uniform mixed strategies and define the function F from $(E_\delta\backslash\partial E_\delta \times (\tilde{\Delta}\backslash\partial\tilde{\Delta}) \times \{0\}) \cup (E_\delta \times \tilde{\Delta} \times (0, 1])$ to P'_δ by $F(\eta, \sigma, \tilde{\lambda}, t) = \pi(\eta, (1-t)\sigma_n + t\sigma^*, \tilde{\lambda})$. The corresponding continuity properties of π imply that F is continuous. Moreover, it is readily checked that F is a proper map, i.e., the inverse image of a compact set is compact. For each $t \in [0, 1]$, let F^t be the corresponding cross section of F obtained by fixing t. In particular, F^0 is the map from $(E_\delta\backslash\partial E_\delta) \times (\tilde{\Delta}\backslash\partial\tilde{\Delta})$ to P'_δ and is given as the restriction of h to this domain; and, for $t > 0$, F^t gives a map from $(E_\delta, \partial E_\delta) \times (\tilde{\Delta}, \partial\tilde{\Delta}) \to (P'_\delta, \partial P'_\delta)$. Obviously, F defines a homotopy between F^t and $F^{t'}$ for $t, t' \in (0, 1]$. We claim that F^1, and hence also each F^t, for $t > 0$, is stably essential. To see this, observe that the map F^1 can be written as the composite map $\pi^* \circ p^{N-1}_\delta$ where p^{N-1}_δ: $E_\delta \times \tilde{\Delta}^k \to P_\delta \times \tilde{\Delta}$ is the map $p^{N-1}_\delta(\eta, \sigma, \tilde{\lambda}) = (\eta, \tilde{\lambda})$ and π^*: $(P_\delta, \partial P_\delta) \times (\tilde{\Delta}, \partial\tilde{\Delta}) \to (P'_\delta, \partial P'_\delta)$ is the map $\pi^*(\eta, \tilde{\lambda}) = \pi(\eta, \sigma^*, \tilde{\lambda})$. The stable essentiality of p_δ implies that of p^{N-1}_δ. Because π^* is a quotient map that sends $(P_\delta\backslash\partial P_\delta) \times (\tilde{\Delta}\backslash\partial\tilde{\Delta})$ homeomorphically onto $P'_\delta\backslash\partial P'_\delta$, by Remark A.3, F^1 is stably essential. Therefore, F^t is stably essential for each $t > 0$.

We show that if p'_δ from E'_δ is not stably essential then F^t is not stably essential for small but positive t, which contradicts what we proved above and thus establishes the result. If p'_δ is not stably essential, then by Lemma A.4 there exist k and a map f' from $E'_\delta \times \Delta^k$ to $P'_\delta \times \Delta^k$ that does not have a point of coincidence with $(p'_\delta)^k$ and whose image is a subset of $B' \times \Delta^k$, where B' is a ball that is contained in the interior of P'_δ.

Define g': $(E_\delta\backslash\partial E_\delta) \times (\tilde{\Delta}\backslash\partial\tilde{\Delta}) \times \Delta^k \to P'_\delta \times \Delta^k$ by $g'(\eta, \sigma, \tilde{\lambda}, \lambda) = f'(h(\eta, \sigma, \tilde{\lambda}), \lambda)$. Then, obviously, g' has no point of coincidence with $F^0 \times \text{Id}$: $(E_\delta\backslash\partial E_\delta) \times (\tilde{\Delta}\backslash\partial\tilde{\Delta}) \times \Delta^k \to P'_\delta \times \Delta^k$, where Id is the identity map on Δ^k. Because F is a proper map, and F^t maps $(E_\delta \times \partial\tilde{\Delta}) \cup (\partial E_\delta \times \tilde{\Delta})$ into ∂P_δ for each $t > 0$, there exists a compact subset X of $(E_\delta\backslash\partial E_\delta) \times (\tilde{\Delta}\backslash\partial\tilde{\Delta})$ such that $F^{-1}(B') \subseteq X \times [0, 1]$. The restrictions of $F^0 \times \text{Id}$ and g' to $X \times \Delta^k$ have no point of coincidence, either. Because $F^{-1}(B) \subseteq X \times [0, 1]$, using the continuity of F over $X \times [0, 1]$, we therefore have that for t small enough the restriction of $F^t \times \text{Id}$ to $X \times \Delta^k$ has no point of coincidence with the restriction of g' to $X \times \Delta^k$. In particular, for such t, the restriction of $F^t \times \text{Id}$ to $(F^t)^{-1}(B) \times \Delta^k$ (which is contained in $X \times \Delta^k$) has no point of coincidence with the corresponding restriction of g'. But by Lemma A.5 that would be impossible because F^t is stably essential. Therefore, the projection from E'_δ is stably essential and thus the definitions involving P and P' are equivalent.

We now prove the equivalence between the definitions using P and $\hat{P}$. The corresponding proof for P' and $\hat{P}'$ is analogous and, therefore, we omit it. The function f: $(\hat{P}, \partial\hat{P}) \to (P, \partial P)$ given by $f(\delta, \tau) = \delta\tau$ is a quotient map that collapses $\{0\} \times \Sigma$ to a single point 0, and it induces a homeomorphism between $\hat{P}\backslash\partial\hat{P}$ and $P\backslash\partial P$. The function g: $(\hat{\mathscr{E}}, \partial\hat{\mathscr{E}}) \to (\mathscr{E}, \partial\mathscr{E})$ given by $g(\delta, \tau, \sigma) = (\delta\tau, \sigma)$ is also a quotient map with similar properties. Moreover, these maps commute with the respective projections, i.e., $p \circ g = f \circ \hat{p}$, where p and $\hat{p}$ are the projections from $\mathscr{E}$ and $\hat{E}$, respectively, to P and $\hat{P}$. Therefore, given a germ E in $\mathscr{E}$ for a stably essential set S, the closure $\hat{E}$ of $g^{-1}(E\backslash\partial E_1)$ is a germ for S in $\hat{E}$. Indeed, $S = \{\sigma \mid (0, \tau, \sigma) \in \hat{E}\}$; the connectedness condition for $\hat{E}$ is immediate from the fact that g maps $\hat{E}\backslash\partial\hat{E}_1$ homeomorphically onto $E\backslash\partial E_1$, and that E satisfies the connectedness condition; and the essentiality condition follows from Lemma A.6. Going the other way, if $\hat{E}$ is a germ in $\hat{E}$ for a set S, then $g(\hat{E})$ is a germ for E for the same reasons. □

We turn now to perturbation spaces obtained from P. (The analogs for P' are obvious so we do not discuss them.) Following Mertens [21, Appendix I], a normalization is an equivalence relation on P such that two perturbations are equivalent only if the games they induce have the same equilibria. Formally, a normalization is a quotient mapping f from P_{δ_0}, for some $\delta_0 > 0$, to some set $\tilde{Q}$ such that if η and η' belong to $P_{\delta_0} \backslash \partial P_1$ and $f(\eta) = f(\eta')$, then $G(\eta)$ and $G(\eta')$ have the same set of Nash equilibria. For example, the map sending η to $G(\eta)$ is a normalization, because it identifies all perturbations that yield the same game. Mertens [21, §6, Appendix I] provides other examples of normalizations, which yield spaces of much smaller dimension. He also shows that for any normalization f, there exists a basis $\mathscr{W}$ of neighborhoods W of $0 \in P_{\delta_0}$ that are homeomorphic to simplices, and that are mapped by f to a basis $\widetilde{\mathscr{W}}$ of neighborhoods of $f(0)$ in $\tilde{Q}$ that are also homeomorphic to simplices. Also, for each $W \in \mathscr{W}$, there is a homeomorphism h_W of W with $f(W) \times W_k$ for $W_k \subset \mathbb{R}^k$ for some k such that the restriction of f to W is the composite of h_W with the projection from $f(W) \times W_k$ to $f(W)$. This set W_k is homeomorphic to a quotient space of a cube $[0,1]^k$ obtained by identifying some points on the boundary (Mertens [21, Appendix I, §4, Theorems 1 and 2]).

Given a quotient mapping $f: P_{\delta_0} \to \tilde{Q}$, we can now define stably essential sets, using $\tilde{Q}$ as the space of perturbations as follows. Choose some set W_0 in the basis $\mathscr{W}$ of neighborhoods of $0 \in P$. Viewing $\mathscr{E}$ as the graph of equilibria over P, let $\mathscr{E}_{W_0}$ be the inverse image in $\mathscr{E}$ of W_0 under the projection p to P and define $g: \mathscr{E}_{W_0} \to \tilde{Q} \times \Sigma$ by $g(\eta, \sigma) = (f(\eta), \sigma)$. Let $\tilde{p}$ be the projection from $\widetilde{\mathscr{E}}$ to $\tilde{Q}$. Recall that $\widetilde{W}_0 \equiv f(W_0)$ is a neighborhood of $f(0)$ in $\tilde{Q}$ that is homeomorphic to a simplex. Denote its boundary by $\tilde{\partial} W_0$. Let $\partial\widetilde{\mathscr{E}}$ be the set of $\tilde{p}^{-1}(\partial \widetilde{W}_0)$.

For a subset $\widetilde{E}$ of $\widetilde{\mathscr{E}}$ and a neighborhood $\widetilde{W}$ of $f(0) \in \tilde{Q}$ that is homeomorphic to a simplex, let $(\widetilde{E}_{\widetilde{W}}, \partial \widetilde{E}_{\widetilde{W}})$ be the inverse image of $(\widetilde{W}, \partial \widetilde{W})$ in $\widetilde{E}$ under p, where $\partial \widetilde{W}$ is the boundary of $\widetilde{W}$. Also let $\widetilde{E}_0$ be the set of $\{(f(0), \sigma) \in \widetilde{E}\}$.

Definition D.1 ($\tilde{Q}$-Stably Essential Set). $S \subseteq \Sigma$ is a $\tilde{Q}$-stably essential set if for some closed subset $\widetilde{E}$ of $\widetilde{\mathscr{E}}$ with $\widetilde{E}_0 = \{f(0)\} \times S$:

(i) Connectedness: For every neighborhood $\widetilde{V}$ of $\widetilde{E}_0$ in $\widetilde{E}$, the set $\widetilde{V} \backslash \partial \widetilde{\mathscr{E}}$ has a connected component whose closure is a neighborhood of $\widetilde{E}_0$ in $\widetilde{E}$.

(ii) Stable essentiality: the projection map $\tilde{p}: (\widetilde{E}_{\widetilde{W}}, \partial \widetilde{E}_{\widetilde{W}}) \to (P_{\widetilde{W}}, \partial P_{\widetilde{W}})$ is stably essential in homotopy for some neighborhood $\widetilde{W}$ of $f(0)$ that is homeomorphic to a simplex.

With these definitions, we have the following theorem.

Theorem D.2. *S is stably essential if and only if it is $\tilde{Q}$-stably essential.*

Proof. Let S be a stably essential set with a germ E. We can assume without loss of generality that E is contained in $\mathscr{E}_{W_0}$ and that it is the closure of $E \backslash \partial E_1$. Let $\widetilde{E} = g(E)$. Then $\widetilde{E}$ is a closed subset of $\widetilde{\mathscr{E}}$ with $\widetilde{E}_0 = \{f(0)\} \times S$. We will show that $\widetilde{E}$ satisfies the connectedness and essentiality conditions as well. By the connectedness condition for E, there exists a basis $\mathscr{V}$ of neighborhoods V of E_0 in E such that $V \backslash \partial E_1$ is connected. For each $V \in \mathscr{V}$, let $\widetilde{V} = g(V)$. Then $\widetilde{V} \backslash \partial \widetilde{\mathscr{E}}$ is connected because it is obtained from the connected set $g(V \backslash \partial E_1)$ by adding some limit points. Moreover, the $\widetilde{V}$s form a basis of neighborhoods of $\widetilde{E}_0$ in $\widetilde{E}$ because the sets $p^{-1}(W) \cap E$ for $W \in \mathscr{W}$ form a basis of neighborhoods of E_0 in E whose images under g are the sets $\widetilde{E}_{\widetilde{W}}$, for $\widetilde{W} \in \widetilde{W}$, which form a basis of neighborhoods of $\widetilde{E}_0$ in $\widetilde{E}$. Thus, $\widetilde{E}$ satisfies the connectedness condition. We turn now to the essentiality property. Because E satisfies the essentiality condition, by Lemma A.5, there exists some $W \in \mathscr{W}$ such that the projection $p: (E_W, \partial E_W) \to (W, \partial W)$ is essential. By the properties of f we saw before, there exists a quotient map π from $\widetilde{W} \times [0,1]^k$ to W that collapses some points in $\widetilde{W} \times \partial [0,1]^k$ such that $f \circ \pi = \mathrm{proj}_{\widetilde{W}}$. π induces a map ψ from $\widetilde{E}_{\widetilde{W}} \times [0,1]^k$ to E_W given by $\psi((\tilde{\eta}, \sigma), t) = (\pi(\tilde{\eta}, t), \sigma)$. And $p \circ \psi = \pi \circ \tilde{p}^k$, where $\tilde{p}^k: (\widetilde{E}_{\widetilde{W}}, \partial \widetilde{E}_{\widetilde{W}}) \times ([0,1]^k, \partial [0,1]^k) \to (\widetilde{W}, \partial \widetilde{W}) \times ([0,1]^k, \partial [0,1]^k)$ is given by $\tilde{p}^k(\tilde{\eta}, \sigma, t) = (\tilde{\eta}, t)$. By Lemma A.6, the stable essentiality of p implies that of $\tilde{p}^k$. The stable essentiality of $\tilde{p}^k$ now clearly implies that of $\tilde{p}$ itself. Thus, S is $\tilde{Q}$-stably essential.

We turn now to the converse. Suppose S is a $\tilde{Q}$-stably essential set with a germ $\widetilde{E}$. Let $E = g^{-1}(\widetilde{E})$. Then E is a closed subset of $\mathscr{E}$ with $E_0 = \{0\} \times S$. Given a basis $\widetilde{\mathscr{V}}$ of neighborhoods $\widetilde{V}$ of $\widetilde{E}_0$ in $\widetilde{E}$ such that $\widetilde{V} \backslash \partial \widetilde{E}$ is connected, let $\mathscr{V}$ be the collection of $g^{-1}(\widetilde{V})$ for $\widetilde{V} \in \widetilde{\mathscr{V}}$. Then $g^{-1}(\widetilde{V}) \backslash \partial E_1$ is a connected set because it is the image of $\widetilde{V} \backslash \partial \widetilde{E}_1 \times (0,1)^k$ under the map ψ described in the last paragraph. Also $\mathscr{V}$ is a basis of neighborhoods of E_0 in E because the sets $\widetilde{E}_{\widetilde{W}}$ for $W \in \mathscr{W}$ form a basis of neighborhoods of $\widetilde{E}_0$ in $\widetilde{E}$ whose inverse images under g are the sets $p^{-1}(W) \cap E$, $W \in \mathscr{W}$, which form a basis of neighborhoods of E_0 in E. Hence, E satisfies the connectedness condition. Finally, the essentiality condition for E follows as above: The essentiality of $\tilde{p}$ for some $\widetilde{W} \in \widetilde{W}$ implies that of $\tilde{p}^k$ defined above. Lemma A.6 now shows that the restriction of p to E_W is stably essential. Because W is a neighborhood of 0, it contains P_δ for some small $\delta > 0$ and Lemma A.5 yields the result. □

Reworking the two proofs in this section for the case of homotopy stability reveals the following relationships among the four concepts that obtain from using these four spaces of perturbations. Definition 2.3 is equivalent to the definition using $\hat{P}$. Likewise the definitions using P' and $\hat{P}'$ are equivalent to one another and imply Definition 2.3. Also Definition 2.3 implies $\tilde{Q}$-homotopy stability. It seems likely that Definition 2.3 does not imply P'-homotopy stable and that it is not implied by $\tilde{Q}$-homotopy stability. Indeed, given that $\mathscr{E}'$ and $\mathscr{E}$ are "suspensions" of $\mathscr{E}$ and $\widetilde{\mathscr{E}}$, respectively, one would need essentiality of at least a few higher order suspensions of projection maps to be essential for these two results to hold.

Appendix E. Generic extensive-form games. In this appendix we show that metastability is equivalent to Mertens' stability for games in extensive form with perfect recall and generic payoffs. Specifically, the result obtains as a corollary to the main result of the appendix, which is that for such games the outcomes of Mertens' stable sets remain the same if homological essentiality is replaced by homotopic essentiality.

E.1. Formulation. We consider a fixed finite N-player game tree $\mathscr{T} = (X, \prec; \mathscr{N}, U; \sigma_\circ)$ with perfect recall. The set of players is $\mathscr{N}$ and the set of nodes is X. The precedence relation $\prec$ on X is acyclic, totally orders the predecessors of each node, and has a single root. The set $Z \subset X$ of terminal nodes comprises those with no successors. The partition U of $X \backslash Z$ divides the set of non-terminal nodes into information sets that are the contingencies in which players and nature choose actions; $\sigma_\circ$ is the mixed strategy of nature. For each player $n \in \mathscr{N}$, $U_n \subset U$ is the (nonempty) collection of his information sets. For each n and each $u \in U_n$, let $A_n(u)$ be the set of player n's actions at u. Let $A_n = \bigcup_{u \in U_n} A_n(u)$ be the entire set of n's actions, all labeled differently. For $x \in X$ and $u \in U$, write $u \prec x$ if there is a node in u that precedes x. Likewise, $(u, i) \prec x$, or simply $i \prec x$, if x succeeds a node in u after action i at u. Because $\mathscr{T}$ has perfect recall, if a node of an information set $u' \in U_n$ of player n succeeds another $u \in U_n$ by a choice of action i then every node of u' succeeds u by the choice of i. In this case, we write $(u, i) \prec u'$, or simply $i \prec u'$. Player n's set of pure strategies is $\{s_n: U_n \to A_n \mid s_n(u) \in A_n(u)\ \forall\, u \in U_n\}$. His set Σ_n of mixed strategies is the set of probability distributions on his pure strategies.

Given the game tree $\mathscr{T}$, an extensive-form game G is specified by assigning for each $n \in \mathscr{N}$ and $z \in Z$ a payoff G_{nz} to player n at the terminal node z. Thus the space of extensive-form games generated by $\mathscr{T}$ is $\mathscr{G} \equiv \mathbb{R}^{\mathscr{N} \times Z}$. It is well known (Kreps and Wilson [16], Govindan and Wilson [6]) that a generic extensive-form game with perfect recall has finitely many outcomes from Nash equilibria. That is, there exists a closed lower dimensional semialgebraic subset $\mathscr{G}_0$ of $\mathscr{G}$ such that each game $G \in \mathscr{G} \backslash \mathscr{G}_0$ has finitely many equilibrium outcomes; in particular, all equilibria in each connected set of equilibria have the same outcome.

Because $\mathscr{T}$ has perfect recall, by Kuhn's [17] theorem every mixed strategy of a player can be implemented by a behavioral strategy. In Govindan and Wilson [7] we introduced the set of enabling strategies, essentially equivalent to sequence-form strategies, which has the dimensional advantage of behavioral strategies and the linearity advantage of mixed strategies. Enabling strategies slightly refine the sequence-form strategies introduced by Koller and Megiddo [14], Koller et al. [15], von Stengel [30] for two-player games.

Enabling strategies are defined as follows. First, for each $u \in U_n$ define $E_n(u)$ to be the set of n's pure strategies that enable the contingency u to occur, i.e., $E_n(u)$ is the set of n's pure strategies s_n such that $s_n(u') = i$ for each $(u', i') \prec u$ and thus do not exclude u. Also for each action $i \in A_n(u)$, let $E_n(u, i) \subset E_n(u)$ be the subset that enable u and choose i at u. Say that an action $i \in A_n$ is a *last* action for n preceding $z \in Z$ if: $i \prec z$, and $u \not\prec z$ for all $u \in U_n$ such that $i \prec u$. For each $z \in Z$, let $\ell_n(z)$ be the set of n's last actions preceding z, which is a singleton set when nonempty. Let $L_n \equiv \bigcup_z \ell_n(z)$ be the set of player n's last actions. Define $\pi_n: \Sigma_n \to [0, 1]^{L_n}$ as follows: For each $\sigma_n \in \Sigma_n$, and each last action $i \in L_n$, $\pi_{ni}(\sigma_n) = \sum_{s_n \in E_n(u, i)} \sigma_n(s_n)$, where u is the unique information set in U_n with $i \in A_n(u)$. The set $\widetilde{\Sigma}_n \equiv \pi_n(\Sigma_n)$ is the set of player n's enabling strategies. Because π_n is linear and Σ_n is a compact and convex polyhedron, $\widetilde{\Sigma}_n$ is also a compact convex polyhedron. Its vertices are the "pure" enabling strategies, each identified by the terminal nodes it does not exclude. Define $\pi \equiv (\pi_n)_{n \in \mathscr{N}}$ and $\widetilde{\Sigma} \equiv \prod_{n \in \mathscr{N}} \widetilde{\Sigma}_n$.

Given an extensive-form game $G \in \mathscr{G}$, we use $\widetilde{G}$ to denote the corresponding enabling-form game in which each player n's strategy space is $\widetilde{\Sigma}_n$ and his expected payoff from $\tilde{\sigma} \in \widetilde{\Sigma}$ is

$$\widetilde{G}_n(\tilde{\sigma}) = \sum_{z \in Z} G_{nz} \sigma_\circ(z) \prod_{m:\, \ell_m(z) \neq \varnothing} \tilde{\sigma}_m(\ell_m(z)),$$

where $\sigma_\circ(z)$ is the product of the probabilities of the actions of nature that precede z, with the convention that $\sigma_\circ(z) = 1$ if there is no such action. The strategic and enabling forms of the extensive-form game G are equivalent in that the payoffs from a mixed strategy profile σ in the strategic form of G are the payoffs from the profile $\tilde{\sigma} = \pi(\sigma)$ of enabling strategies in its enabling form $\tilde{G}$. Thus, $\sigma \in \Sigma$ is an equilibrium of G if and only if $\tilde{\sigma}$ is an equilibrium of $\tilde{G}$. In particular, for each $G \in \mathscr{G} \backslash \mathscr{G}_0$, all equilibria in a connected set of equilibria of its enabling-form $\tilde{G}$ induce the same outcome.

E.2. Stable sets of the enabling-form game. For each $\varepsilon \in (0, 1]$ let $\tilde{P}_\varepsilon = [0, \varepsilon] \times \tilde{\Sigma}$ and let $\partial \tilde{P}_\varepsilon$ be its relative boundary. Given an extensive-form game $G \in \mathscr{G}$ and $(\delta, \tilde{\tau}) \in \tilde{P}_1$, define analogously the perturbed enabling-form game $\tilde{G}(\delta, \tilde{\tau})$ to be the enabling-form game where the strategy set is the set $\tilde{\Sigma}$ of enabling strategies in $\tilde{G}$, but the payoffs from a profile $\tilde{\sigma}$ are the payoffs from the profile $(1-\delta)\tilde{\sigma} + \delta\tilde{\tau}$ in $\tilde{G}$. Let $\tilde{\mathscr{E}}(G)$ be the graph of the equilibrium correspondence over the space $\tilde{P}$ of perturbations; viz.,

$$\tilde{\mathscr{E}}(G) = \{(\delta, \tilde{\tau}, \tilde{\sigma}) \in \tilde{P}_1 \times \tilde{\Sigma} \mid \tilde{\sigma} \text{ is an equilibrium of } \tilde{G}(\delta, \tilde{\tau})\}.$$

Due to the equivalence between G and $\tilde{G}$, for each $(\delta, \tau) \in P_1$, σ is a Nash equilibrium of $G(\delta, \tau)$ if and only if $\pi(\sigma)$ is a Nash equilibrium of $\tilde{G}(\delta, \pi(\tau))$. Therefore, there is a well-defined map $\phi\colon E(G) \to \tilde{\mathscr{E}}(G)$ given by $\phi(\delta, \tau, \sigma) = (\delta, \pi(t), \pi(\sigma))$. Using the equilibrium graph $\tilde{\mathscr{E}}(G)$ and the projection map $\tilde{p}$ from $\tilde{\mathscr{E}}(G)$ to $\tilde{P}_1$ (instead of $E(G)$ and p) in Definitions 2.3 and 5.2, one obtains analogous definitions of the h-stable and $*$-stable sets of $\tilde{G}$ represented as enabling strategies. Recall from §5 that 0-stable sets are defined exactly like $*$-stability except that the cohomological essentiality uses the field of rational numbers. We can now define 0-stable sets of $\tilde{G}$ just like $*$-stable sets. To avoid confusion, we will specify whether stability is with respect to G using perturbations of mixed strategies, or $\tilde{G}$ using perturbations of enabling strategies as here. Lemmas E.1 and E.2 establish the relationship between stability for G and stability for $\tilde{G}$.

E.3. Statement of the theorem. Consider now an 0-stable or h-stable set S_0 of a generic extensive-form game $G \in \mathscr{G} \backslash \mathscr{G}_0$. Because G has finitely many equilibrium outcomes, and S_0 is a connected set due to the connectedness property, all the strategy profiles in S_0 induce the same outcome, which is then called a 0-stable or h-stable outcome of G, respectively. The 0-stable or h-stable outcomes of G's enabling-form $\tilde{G}$ are defined analogously. Our main result in this appendix is the following theorem.

Theorem E.1. *There exists a lower dimensional semialgebraic subset $\mathscr{G}_1$ of $\mathscr{G}$ such that for each game $G \in \mathscr{G} \backslash (\mathscr{G}_0 \cup \mathscr{G}_1)$ an equilibrium outcome is h-stable for G if and only if it is* 0-*stable for G.*

That is, the class of extensive-form games with perfect recall for which the 0-stable and h-stable outcomes coincide is generic. The plan of the proof is the following. The "if" part of the theorem is obvious because 0-stability is a stronger refinement than h-stability. The proof of the "only if" part involves three steps. Consider a generic game $G \in \mathscr{G} \backslash \mathscr{G}_0$ whose enabling form is $\tilde{G}$. The first step shows that an outcome is 0-stable for G if and only if it is 0-stable for $\tilde{G}$, and that it is h-stable for G only if it is h-stable for $\tilde{G}$. This is accomplished via two lemmas that are phrased in slightly more general terms in that they deal with the stable sets of any game in $\mathscr{G}$. Thanks to this step, it is sufficient to show that there exists a lower dimensional semialgebraic subset $\mathscr{G}_1$ of $\mathscr{G}$ such that if $G \in \mathscr{G} \backslash (\mathscr{G}_0 \cup \mathscr{G}_1)$ and an outcome is h-stable for $\tilde{G}$, then it is 0-stable for $\tilde{G}$.

A structure theorem for extensive-form games (Govindan and Wilson [7, Theorem 5.2]) implies that there exists a lower dimensional semialgebraic subset $\mathscr{G}_1$ of $\mathscr{G}$ such that if $G \notin \mathscr{G}_1$ then the graph $\tilde{\mathscr{E}}(G)$ of its equilibrium correspondence over $\tilde{P}_1$ has the same dimension as $\tilde{P}_1$. This already shows that an h-stable outcome of $\tilde{G}$ is $*$-stable. However, the fact that $\tilde{\mathscr{E}}(G)$ and $\tilde{P}_1$ have the same dimension does not suffice. The potential problem is that a germ that renders an outcome h-stable might, for instance, be a nonorientable manifold, and a cycle (mod 2) in the germ could project to a cycle (mod 2) in $\tilde{P}_1$. To show that this is *not* the case, the second step shows that for all small $\varepsilon > 0$, $\tilde{\mathscr{E}}_\varepsilon(G)$ (actually, the closure of $\tilde{\mathscr{E}}_\varepsilon(G) \backslash \partial \tilde{\mathscr{E}}_\varepsilon(G)$) is a finite union of orientable pseudomanifolds. These pseudomanifolds are shown to satisfy the connectedness condition for stability, and the intersection of any two of them lies in $\partial \tilde{\mathscr{E}}_1(G)$. Equipped with this result, the third step establishes the 0-stability of h-stable outcomes of $\tilde{G}$ as follows. The third step shows that a germ that makes an outcome h-stable for $\tilde{G}$ must be contained in one of these pseudomanifolds, say $\tilde{T}$. The projection from $\tilde{T}$ satisfies the essentiality condition for h-stability because it contains the germ that does. The fact that $\tilde{T}$ is orientable then implies that the projection map is essential in cohomology with coefficients in $\mathbb{Q}$.

E.4. Relation between stable sets of G and $\tilde{G}$. Because the equilibria of G and $\tilde{G}$ are equivalent, the ordinality property for 0-stability (Mertens [24, Proposition 11]) provides the following result.

Lemma E.1. *If S is an 0-stable set of G, then $\pi(S)$ is an 0-stable set of $\widetilde{G}$. And if $\widetilde{S}$ is an 0-stable set of $\widetilde{G}$, then $\pi^{-1}(\widetilde{S})$ is the union of those 0-stable sets of G whose image under π is $\widetilde{S}$.*

Regarding h-stability, we obtain the following partial result for an arbitrary extensive-form game G, although Lemma E.1 and later Lemma E.4 show that if G is generic, then the h-stable outcomes of G and $\widetilde{G}$ are the same.

Lemma E.2. *If S is an h-stable set of G then $\pi(S)$ is an h-stable set of $\widetilde{G}$.*

Proof. Let S be an h-stable set of G and let E be a germ as in Definition 2.3. Let $\widetilde{E}=\phi(E)$. Obviously $\widetilde{E}$ is a closed subset of $\widetilde{\mathscr{E}}(\widetilde{G})$ and $\pi(S)=\{\widetilde{\sigma}\mid \exists(0,\tilde{\tau},\widetilde{\sigma})\in\widetilde{E}\}$. We prove that $\widetilde{E}$ satisfies the essentiality and connectedness conditions for h-stability in $\widetilde{G}$ as in Definition 2.3.

Fix ε such that the projection p from E_ε is essential in homotopy. We claim that the projection $\tilde{p}$: $(\widetilde{E}_\varepsilon,\partial\widetilde{E}_\varepsilon)\to(\widetilde{P}_\varepsilon,\partial\widetilde{P}_\varepsilon)$ is essential in homotopy. By Lemma A.4, to prove this claim it is sufficient to show that for each continuous map $\tilde{f}$: $\widetilde{E}_\varepsilon\to\widetilde{P}_\varepsilon$ there is a point of coincidence between $\tilde{f}$ and $\tilde{p}$. Accordingly, fix such a map $\tilde{f}$. The map from P_ε to $\widetilde{P}_\varepsilon$ that sends (δ,τ) to $(\delta,\pi(\tau))$ is a surjective affine map between polyhedra; hence its inverse has a continuous selection h by Lemma 4.1. Then $h\circ\tilde{f}\circ\phi$ is a well-defined continuous map from E_ε to P_ε. Because p is essential in homotopy, by Lemma A.3 it has a point of coincidence with $h\circ\tilde{f}\circ\phi$; i.e., there exists $(\delta,\tau,\sigma)\in E_\varepsilon$ such that $(\delta,\tau)=h\circ\tilde{f}\circ\phi(\delta,\tau,\sigma)$. Then $(\delta,\pi(\tau))=h^{-1}(\delta,\tau)=\tilde{f}\circ\phi(\delta,\tau,\sigma)=\tilde{f}(\delta,\pi(\tau),\pi(\sigma))$, so $(\delta,\pi(\tau),\pi(\sigma))$ is a point of coincidence between $\tilde{p}_\varepsilon$ and f. Thus $\tilde{p}$ is essential in homotopy as claimed.

We turn now to the connectedness condition. Let $\widetilde{U}$ be a neighborhood of $\pi(S)$ in $\widetilde{E}$, i.e., $\widetilde{U}$ is a neighborhood of $\{(0,\tilde{\tau},\widetilde{\sigma})\in E\mid\widetilde{\sigma}\in\pi(S)\}$ in $\widetilde{E}$. Pick a closed neighborhood $\widetilde{V}\subseteq\widetilde{U}$ of $\pi(S)$ in $\widetilde{E}$. Then $V\equiv E\cap\phi^{-1}(\widetilde{V})$ is a closed neighborhood of S in E that is contained in $\phi^{-1}(\widetilde{U})$. By the connectedness condition for E there exists a connected component of $V\backslash\partial E_1$ whose closure—call it W—is a neighborhood of S in E. Obviously, W is contained $\phi^{-1}(\widetilde{U})$. Let $\widetilde{W}\equiv\phi(W)$, then $\widetilde{W}$ is contained in $\widetilde{U}$, and it is a neighborhood of $\pi(S)$ in $\widetilde{E}$—indeed, because W is a neighborhood of S in E, there exists ε such that W contains E_ε, $\widetilde{W}$ then contains the neighborhood $\widetilde{E}_\varepsilon$ of $\pi(S)$ in $\widetilde{E}$. Finally, we show that $\widetilde{W}\backslash\partial\widetilde{E}_1$ is connected and dense in $\widetilde{W}$, which proves the connectedness condition because the closure of the connected component of $\widetilde{U}\backslash\partial\widetilde{E}_1$ that contains $\widetilde{W}\backslash\partial\widetilde{E}_1$ contains the neighborhood $\widetilde{W}$ of $\pi(S)$ in $\widetilde{E}$. Let $W'=W\cap\phi^{-1}(\widetilde{W}\backslash\partial\widetilde{E}_1)$. Because $W\backslash\partial E_1$ is contained in W' and is connected and dense in W, W' is connected and dense in W. Hence $\widetilde{W}\backslash\partial\widetilde{E}_1$ is connected and dense in $\widetilde{W}$. Thus $\widetilde{E}$ satisfies the connectedness condition as well and $\pi(S)$ is h-stable in $\widetilde{G}$. □

E.5. The structure of $\widetilde{\mathscr{E}}(G)$. The results of the previous subsection reduce the problem to analyzing the relationship between 0-stability and h-stability in $\widetilde{G}$. We prove here a useful lemma concerning the structure of $\widetilde{\mathscr{E}}(G)$ for a generic game G that will help show the equivalence between the two stability concepts in $\widetilde{G}$.

Before stating and proving the lemma, we introduce some additional notation. Let $\check{\mathscr{E}}=\{(G,(\delta,\tilde{\tau},\widetilde{\sigma}))\mid G\in\mathscr{G},(\delta,\tilde{\tau},\widetilde{\sigma})\in\widetilde{\mathscr{E}}(G)\}$. For each G, its equilibrium graph $\widetilde{\mathscr{E}}(G)$ corresponds to the cross-section of $\check{\mathscr{E}}$ over the game G. For each $\varepsilon\in(0,1]$ let $(\check{\mathscr{E}}_\varepsilon,\partial\check{\mathscr{E}}_\varepsilon)$ be the inverse image of $(\widetilde{P}_\varepsilon,\partial\widetilde{P}_\varepsilon)$ in $\check{\mathscr{E}}$ under the natural projection to $\widetilde{P}_1$. Finally, let k be the dimension of $\widetilde{P}_1$.

The sets we construct later to show 0-stability of h-stable outcomes are all compact and semialgebraic; in particular, they are compact polyhedra. Hence all cohomology theories on such sets coincide. Given this, even though the definition of stability requires Čech cohomology, we can use the singular cohomology theory. Accordingly, in what follows we use the singular cohomology and homology theories with integer coefficients.

Lemma E.3. *There exists a lower dimensional semialgebraic subset $\mathscr{G}_1$ of $\mathscr{G}$ such that for each $G\in\mathscr{G}\backslash\mathscr{G}_1$ there exist $\varepsilon_0>0$, a positive integer J, and for each $i=1,\ldots,J$, a k-dimensional semialgebraic subset $\widetilde{T}^i$ of $\widetilde{\mathscr{E}}(G)$ such that for each $0<\varepsilon\le\varepsilon_0$:*
(i) *$\widetilde{T}^i_\varepsilon\backslash\partial\widetilde{T}^i_\varepsilon$ is connected and dense in $\widetilde{T}^i_\varepsilon$ for each i.*
(ii) *$\widetilde{T}^i_\varepsilon\cap\widetilde{T}^j_\varepsilon\subseteq\partial\widetilde{\mathscr{E}}_1(G)$ for $i\neq j$.*
(iii) *$\bigcup_i(\widetilde{T}^i_\varepsilon\backslash\partial\widetilde{T}^i_1)=\widetilde{\mathscr{E}}_\varepsilon(G)\backslash\partial\widetilde{\mathscr{E}}_1(G)$.*
(iv) *$H^k(\widetilde{T}^i_\varepsilon,\partial\widetilde{T}^i_\varepsilon)\approx\mathbb{Z}$ for each i.*

Proof. If $\widetilde{\sigma}$ is an equilibrium of $\widetilde{G}(\delta,\tilde{\tau})$, then say that $\hat{\sigma}\equiv(1-\delta)\widetilde{\sigma}+\delta\tilde{\tau}$ is a perturbed equilibrium of $\widetilde{G}(\delta,\tilde{\tau})$. For each $\varepsilon>0$, let $\hat{\mathscr{E}}_\varepsilon$ (resp. $\partial\hat{\mathscr{E}}_\varepsilon$) be the set of $(G,\delta,\tilde{\tau},\hat{\sigma})$ such that $(\delta,\tilde{\tau})$ belongs to $\widetilde{P}_\varepsilon$ (resp. $\partial\widetilde{P}_\varepsilon$), and $\hat{\sigma}$ is a perturbed equilibrium of $\widetilde{G}(\delta,\tilde{\tau})$. Obviously, there is a homeomorphism between $(\check{\mathscr{E}}_\varepsilon,\partial\check{\mathscr{E}}_\varepsilon)$ and $(\hat{\mathscr{E}}_\varepsilon,\partial\hat{\mathscr{E}}_\varepsilon)$.

For a fixed $(\delta,\tilde{\tau})\in\widetilde{P}_1\backslash\partial\widetilde{P}_1$, we show in Govindan and Wilson [7, Theorem 5.2] that there is a homeomorphism $\hat{\psi}_{\delta,\tilde{\tau}}$ between $\{(G,\delta,\tilde{\tau},\hat{\sigma})\in\hat{\mathscr{E}}_1\}$ and $\mathscr{G}$. Moreover, it follows from the construction there that the homeomorphism is itself a continuous function of the parameters $(\delta, \tilde{\tau})$. Thus there exists a homeomorphism $\hat{\psi}$: $\hat{\mathscr{E}}_1\backslash\partial\hat{\mathscr{E}}_1 \to \mathscr{G}\times(\tilde{P}_1\backslash\partial\tilde{P}_1)$ given by $\hat{\psi}(G,\delta,\tilde{\tau},\hat{\sigma}) = (\hat{\psi}_{\delta,\tilde{\tau}}(G,\delta,\tilde{\tau},\hat{\sigma}),(\delta,\tilde{\tau}))$. Because $\hat{\mathscr{E}}_1$ is homeomorphic to $\check{\mathscr{E}}_1$, there is a homeomorphism $\tilde{\psi}$ between $\check{\mathscr{E}}_1\backslash\partial\check{\mathscr{E}}_1$ and $\mathscr{G}\times(\tilde{P}_1\backslash\partial\tilde{P}_1)$.

Let $(\bar{\mathscr{E}}, \partial\bar{\mathscr{E}}) = \bigcup_{\varepsilon\in(0,1)}(\{\varepsilon\}\times(\check{\mathscr{E}}_\varepsilon,\partial\check{\mathscr{E}}_\varepsilon))$. Then $\bar{\mathscr{E}}$ is a semialgebraic set. Clearly, $\bar{\mathscr{E}}\backslash\partial\bar{\mathscr{E}}$ is open in $\bar{\mathscr{E}}$, and it is homeomorphic to $(0,1)\times\mathscr{G}\times(\tilde{P}_1\backslash\partial\tilde{P}_1)$ using the map that sends $(\varepsilon, G,\delta,\tilde{\tau},\tilde{\sigma})$ to $(\varepsilon',\tilde{\psi}(G,\delta,\tilde{\tau},\tilde{\sigma}))$, where $\varepsilon' = (1-\delta)^{-1}(\varepsilon-\delta)$. Therefore, $\bar{\mathscr{E}}\backslash\partial\bar{\mathscr{E}}$ is homeomorphic to $\mathbb{R}^{NZ+k+1}$.

Let $\mathscr{H} = (0,1)\times\mathscr{G}$ and define the projection q: $\bar{\mathscr{E}}\to\mathscr{H}$ by $q(\varepsilon,(G,\delta,\tilde{\tau},\tilde{\sigma})) = (\varepsilon, G)$. By the generic local triviality theorem (Bochnak et al. [2, Theorem 9.3.2]) there exists a closed lower dimensional semialgebraic subset $\mathscr{H}_1$ of $\mathscr{H}$ such that for each component C of $\mathscr{H}\backslash\mathscr{H}_1$, there exists a closed semialgebraic pair $(F,\partial F)$ and a semialgebraic homeomorphism φ: $C\times(F,\partial F)\to(q^{-1}(C), q^{-1}(C)\cap\partial\bar{\mathscr{E}})$, such that $q\circ\varphi = \mathrm{proj}_C$. For each G, the set of ε such that $(\varepsilon, G)\in\mathscr{H}_1$ is a semialgebraic set and has finitely many connected components; therefore, there exists $0<\bar{\varepsilon}<1$ such that $(0,\bar{\varepsilon}]\times\{G\}$ belongs to $\mathscr{H}_1$ or $\mathscr{H}\backslash\mathscr{H}_1$. Let $\mathscr{G}_1$ be the set of games G such that there exists $\bar{\varepsilon}>0$ such that $(0,\bar{\varepsilon}]\times\{G\}$ is contained in $\mathscr{H}_1$. (We then have that for each $G\notin\mathscr{G}_1$, there exists $\bar{\varepsilon}$ such that $(0,\bar{\varepsilon}]\times\{G\}$ belongs to $\mathscr{H}\backslash\mathscr{H}_1$.) Obviously $\mathscr{G}_1$ is a semialgebraic set. Because the dimension of $\mathscr{H}_1$ is at most NZ, an application of Govindan and McLennan [4, Proposition 4.8] to the projection map from $\mathscr{H}_1\cap((0,1)\times\mathscr{G}_1)$ to $\mathscr{G}_1$ shows that the dimension of $\mathscr{G}_1$ is at most $NZ-1$.

Now fix $G\in\mathscr{G}\backslash\mathscr{G}_1$. As we saw above, $(0,\bar{\varepsilon}]\times\{G\}$ is contained in $\mathscr{H}\backslash\mathscr{H}_1$ for some $\bar{\varepsilon}>0$. There now exists a component C of $\mathscr{H}\backslash\mathscr{H}_1$ such that $(0,\bar{\varepsilon}]\times\{G\}\subset C$. Let $(F,\partial F)$ and φ be as in the previous paragraph. As C is open in $\mathscr{H}$, $q^{-1}(C)$ is open in $\bar{\mathscr{E}}$. Because $\bar{\mathscr{E}}\backslash\partial\bar{\mathscr{E}}$ is an open subset of $\bar{\mathscr{E}}$ that has dimension $NZ+k+1$, we now have that $q^{-1}(C)\backslash\partial\bar{\mathscr{E}}$ is an open subset of $\bar{\mathscr{E}}$ and, evidently being nonempty, has dimension $NZ+k+1$. Because φ is a homeomorphism, $C\times(F\backslash\partial F)$ is an open $(NZ+k+1)$-dimensional subset of $C\times F$. Therefore, $F\backslash\partial F$ is open in F; C, being a nonempty open subset of $\mathscr{H}\backslash\mathscr{H}_1$, has dimension $NZ+1$, which implies that $F\backslash\partial F$ has dimension k. To summarize, $F\backslash\partial F$ is an open k-dimensional semialgebraic subset of F.

Proposition E.1. *There exists a positive integer J and for each $1\le i\le J$ a compact and semialgebraic pair $(F^i,\partial F^i)\subseteq(F,\partial F)$, with $\dim(F^i)=k$, such that*:
(i) *$F^i\backslash\partial F^i$ is connected and dense in F^i for each i.*
(ii) *$F^i\cap F^j\subseteq\partial F$ for $j\ne i$.*
(iii) *$\bigcup_i(F^i\backslash\partial F^i) = F\backslash\partial F$*
(iv) *$H^k(F^i,\partial F^i)\approx\mathbb{Z}$ for each i.*

Proof of Proposition. Because $F\backslash\partial F$ is semialgebraic, it has finitely many connected components, say F^{0i}, $1\le i\le J$. For each i, denote the closure of F^{0i} by F^i. F^i is then a compact semialgebraic subset of F. Because $F\backslash\partial F$ is open in F, F^{0i} is open in F, and $\varphi(C\times F^{0i})$ is an open subset of $\bar{\mathscr{E}}\backslash\partial\bar{\mathscr{E}}$, which is $(NZ+k+1)$-dimensional. As was the case with F, this implies that F^{0i}, and hence also its closure F^i, is k-dimensional. Let $\partial F^i = F^i\backslash F^{0i}$. ∂F^i is a closed semialgebraic subset of ∂F. The first three enumerated points of the claim follow readily from the definition of the pairs $(F^i,\partial F^i)$.

Finally, we show that $H^k(F^i,\partial F^i)\approx\mathbb{Z}$ for each i. Let B be a closed ball in C and let ∂B be its boundary. (Such a ball exists because C is open in $\mathscr{H}$.) Then $(B\backslash\partial B)\times(F^i\backslash\partial F^i)$ is homeomorphic to an open connected subset of $\mathbb{R}^{NZ+k+1}$, because it is connected and its image under φ is an open subset of $\bar{\mathscr{E}}\backslash\partial\bar{\mathscr{E}}$. Let $(M^i,\partial M^i) = (B,\partial B)\times(F^i,\partial F^i)$. Because $(M^i,\partial M^i)$ and $(B,\partial B)$ are polyhedral pairs, the Künneth formula for cohomology (Spanier [28, Theorem V.5.11]) yields the following short sequence:

$$0\to\bigoplus_{p+q=NZ+k+1}H^p(B,\partial B)\otimes H^q(F^i,\partial F^i)\to H^{NZ+k+1}(M^i,\partial M^i)\to\bigoplus_{p+q=NZ+k+2}H^i(B,\partial B)*H^j(F^i,\partial F^i)\to 0.$$

Because $(B,\partial B)$ is an $(NZ+1)$-ball, $H^n(B,\partial B)$ is $\mathbb{Z}$ if $n=NZ+1$ and zero otherwise. Also, $H^{k+1}(F^i,\partial F^i)$ is zero, because F^i is k-dimensional. Therefore, the above sequence implies that $H^{NZ+k+1}(M^i,\partial M^i)\approx H^k(F^i,\partial F^i)$. It is sufficient to show that $H^{NZ+k+1}(M^i,\partial M^i)\approx\mathbb{Z}$. Because it is homeomorphic to an open subset of $\mathbb{R}^{NZ+k+1}$, $M^i\backslash\partial M^i$ is an orientable manifold. Therefore, by Spanier [28, Theorem 6.2.19 (the Lefschetz duality theorem) and Theorem 6.1.11]), $H^{NZ+k+1}(M^i,\partial M^i)$ is isomorphic to $H_0(M^i\backslash\partial M^i)$. The latter group is isomorphic to $\mathbb{Z}$ because $M^i\backslash\partial M^i$ is a connected set, and hence path connected because it is semialgebraic. Therefore, $H^{NZ+k+1}(M^i,\partial M^i)$ $\approx\mathbb{Z}$, which completes the proof of the proposition. □

Continuing the proof of the lemma, let φ_G: $(0,\bar{\varepsilon}]\times F\to(0,\bar{\varepsilon}]\times\tilde{\mathscr{E}}_{\bar{\varepsilon}}(G)$ be the map defined by $\varphi_G(\varepsilon, f) = (\varepsilon, t)$, where $t = \mathrm{proj}_{\tilde{P}_1\times\tilde{\Sigma}}\varphi(\varepsilon, G, f)$. For each $0<\varepsilon\le\bar{\varepsilon}$, let $\varphi_{G,\varepsilon}$: $F\to\tilde{\mathscr{E}}_{\bar{\varepsilon}}(G)$ be the map $\varphi_{G,\varepsilon}(f) = \varphi_G(\varepsilon, f)$. Clearly, for each ε, and $1\le i\le J$, $\varphi_{G,\varepsilon}$ maps $(F^i,\partial F^i)$ homeomorphically onto its image in $(\tilde{\mathscr{E}}_\varepsilon(G),\partial\tilde{\mathscr{E}}_\varepsilon(G))$; moreover, $\varphi_{G,\varepsilon}(F\backslash\partial F) = \tilde{\mathscr{E}}_\varepsilon(G)\backslash\partial\tilde{\mathscr{E}}_\varepsilon(G)$. The image under $\varphi_{G,\varepsilon}$ of the sets $F^i\backslash\partial F^i$, which are the connected components of $F\backslash\partial F$, are therefore the connected components of $\tilde{\mathscr{E}}_\varepsilon(G)\backslash\partial\tilde{\mathscr{E}}_\varepsilon(G)$.

For each i, let $(\widetilde{T}^i, \partial\widetilde{T}^i)$ be the image of $(F^i, \partial F^i)$ under $\varphi_{G,\bar{\varepsilon}}$. By the corresponding properties for the F^is, each $\widetilde{T}^i$ is a closed k-dimensional semialgebraic set such that $\widetilde{T}^i\backslash\partial\widetilde{T}^i$ is connected and dense in $\widetilde{T}^i$, $\widetilde{T}^i \cap \widetilde{T}^j \subseteq \partial\widetilde{\mathscr{E}}_{\bar{\varepsilon}}(G)$ for $j \neq i$, and $\bigcup_i(\widetilde{T}^i\backslash\partial\widetilde{T}^i) = \widetilde{\mathscr{E}}_{\bar{\varepsilon}}(G)\backslash\partial\widetilde{\mathscr{E}}_{\bar{\varepsilon}}(G)$.

Choose $0 < \varepsilon_0 < \bar{\varepsilon}$. For each $0 < \varepsilon \leq \varepsilon_0$, we will prove that $\widetilde{T}^i_\varepsilon$ satisfies the four enumerated properties. Property (2) follows from the fact that $\widetilde{T}^i \cap \widetilde{T}^j \subseteq \partial\widetilde{\mathscr{E}}_{\bar{\varepsilon}}(G)$ for $j \neq i$, and property (3) follows from the fact that $\bigcup_i(\widetilde{T}^i\backslash\partial\widetilde{T}^i) = \widetilde{\mathscr{E}}_{\bar{\varepsilon}}(G)\backslash\partial\widetilde{\mathscr{E}}_{\bar{\varepsilon}}(G)$. The other two properties follow from the analogous properties for the F^is if we show that for each i and $0 < \varepsilon < \bar{\varepsilon}$, $\varphi_{G,\varepsilon}(F^i) = \widetilde{T}_\varepsilon$.

For each i and $0 < \varepsilon < \bar{\varepsilon}$, to show that $\varphi_{G,\varepsilon}(F^i) = \widetilde{T}_\varepsilon$ observe first that $\varphi_{G,\varepsilon}(F^i\backslash\partial F^i)$ is contained in $\widetilde{T}_\varepsilon\backslash\partial\widetilde{T}_\varepsilon$ because $\mathrm{proj}_{\widetilde{P}_1\times\widetilde{\Sigma}} \circ \varphi_G([\varepsilon, \bar{\varepsilon}] \times F^i\backslash\partial F^i)$ is a connected subset of $\widetilde{\mathscr{E}}_{\bar{\varepsilon}}(G)$ that intersects the connected component $\widetilde{T}^i\backslash\partial\widetilde{T}^i_{\bar{\varepsilon}}$ of $\widetilde{\mathscr{E}}_{\bar{\varepsilon}}(G)\backslash\partial\widetilde{\mathscr{E}}_{\varepsilon(G)}(G)$ by the definition of $\widetilde{T}^i$. Because $F^i\backslash\partial F^i$ is dense in F^i and $\widetilde{T}^i$ is closed, we have that $\varphi_{G,\varepsilon}(F^i) \subseteq \widetilde{T}^i_\varepsilon$. To show the reverse inclusion, fix $t \in \widetilde{T}^i_\varepsilon$. Because $\widetilde{T}^i\backslash\partial\widetilde{T}^i$ is dense in $\widetilde{T}^i$, there exists a sequence $t^n \in \widetilde{T}^i\backslash\partial\widetilde{T}^i$ converging to t. There now exists a sequence ε^n converging to ε such that $t^n \in \widetilde{T}^i_{\varepsilon^n}\backslash\partial\widetilde{T}^i_{\varepsilon^n}$ all along the sequence. For $j \neq i$, and each n, $t^n \notin \varphi_{G,\varepsilon^n}(F^j)$ since $\widetilde{T}^i \cap \widetilde{T}^i \subseteq \partial\widetilde{T}^i$ and $\varphi_{G,\varepsilon^n}(F^j)$, as we saw above, is contained in $\widetilde{T}^j$. Because $\bigcup_j \varphi_{G,\varepsilon^n}(F^j\backslash\partial F^j) = \widetilde{\mathscr{E}}_{\varepsilon_n}(G)\backslash\partial\widetilde{\mathscr{E}}_{\varepsilon_n}(G)$, there must exist $f^n \in F^i$ such that $t^n = \varphi_G(\varepsilon^n, f^n)$. Going to a subsequence if necessary, let f be the limit of the sequence f^n. Then $t = \varphi_{G,\varepsilon}(f)$, and thus $\widetilde{T}^i_\varepsilon \subseteq \varphi_{G,\varepsilon}(F^i)$. □

One can actually show for each i and $0 < \varepsilon \leq \varepsilon_0$ that $(\widetilde{T}^i_\varepsilon, \partial\widetilde{T}^i_\varepsilon)$ is an orientable k-dimensional pseudomanifold with boundary. To show this it is sufficient to show that the pair $(F^i, \partial F^i)$ constructed in Proposition E.1 is an orientable k-dimensional pseudomanifold with boundary. Let B be the ball defined in the last paragraph of the proof of Proposition E.1. As we saw there, $(B\backslash\partial B) \times (F^i\backslash\partial F^i)$ is a manifold of dimension $NZ + k + 1$. Therefore, for each $b \in B\backslash\partial B$ and $f \in F^i\backslash\partial F^i$, $H^{NZ+K+1}((B, B-b) \times (F^i, F^i - f)) \approx \mathbb{Z}$. Using the Künneth formula we then have that $H^k(F^i, F^i - f) \approx \mathbb{Z}$ for each $f \in F^i\backslash\partial F^i$. Thus, $F^i\backslash\partial F^i$ is a homology k-manifold. Because $(F^i, \partial F^i)$ is a polyhedral pair, it is a k-dimensional pseudomanifold. Also, it is orientable because $H^k(F^i, \partial F^i) \approx \mathbb{Z}$, hence the result. It is not clear if $\widetilde{T}^i_\varepsilon$ is actually a manifold with boundary.

E.6. Proof of the Theorem. The "if" part of Theorem E.1 follows from the fact that essentiality in cohomology implies that in homotopy. The "only if" part of the theorem follows from Lemma E.4 along with Lemmas E.1 and E.2.

Lemma E.4. *For a game $G \in \mathscr{G}\backslash(\mathscr{G}_0 \cup \mathscr{G}_1)$, if an outcome is h-stable in $\widetilde{G}$, then it is 0-stable in $\widetilde{G}$.*

Proof. Fix a game $G \in \mathscr{G}\backslash(\mathscr{G}_0 \cup \mathscr{G}_1)$ and let $\widetilde{S}$ be an h-stable set of equilibria in $\widetilde{G}$. We show that there exists an 0-stable set of $\widetilde{G}$ that contains $\widetilde{S}$ and is therefore outcome equivalent to $\widetilde{S}$.

Because $\widetilde{S}$ is h-stable in $\widetilde{G}$, there exists a subset $\widetilde{E}$ of $\widetilde{\mathscr{E}}(G)$ that satisfies the conditions for h-stability and $\widetilde{S} = \{\widetilde{\sigma} \mid \exists(0, \tilde{\tau}, \widetilde{\sigma})\} \in \widetilde{E}$. By Lemma E.3, there exists $\varepsilon_0 > 0$ and semialgebraic sets $\widetilde{T}^i$ with the properties described there. By the connectedness condition for $\widetilde{E}$ and properties 2 and 3 of Lemma E.3, there exist a unique i and $0 < \varepsilon_1 \leq \varepsilon_0$ such that $\widetilde{E}_{\varepsilon_1}$ is contained in $\widetilde{T}^i_{\varepsilon_1}$. For simplicity in notation, we refer to $\widetilde{T}^i$ simply as $\widetilde{T}$. We prove that $\widetilde{T}_0 \equiv \{\widetilde{\sigma} \mid \exists(0, \tilde{\tau}, \widetilde{\sigma}) \in \widetilde{T}\}$ is an 0-stable set of $\widetilde{G}$: because it contains $\widetilde{S}$, that completes the proof of the lemma.

Observe that $\widetilde{T}^i$ satisfies the connectedness condition for 0-stability in $\widetilde{G}$ by property 1 of Lemma E.3. Hence it is sufficient to prove that it satisfies the essentiality condition. As mentioned after Definition 2.3, the set $\widetilde{E}$ that renders $\widetilde{S}$ h-stable in $\widetilde{G}$ also satisfies the condition for essentiality in homotopy for all sufficiently small ε. Fix $\varepsilon \leq \varepsilon_1$, where ε_1 is as described above, such that the projection from $(\widetilde{E}_\varepsilon, \partial\widetilde{E}_\varepsilon)$ is essential in homotopy. Because $\widetilde{T}$ contains $\widetilde{E}_\varepsilon$, $\tilde{p}$: $(\widetilde{T}_\varepsilon, \partial\widetilde{T}_\varepsilon) \to (\widetilde{P}_\varepsilon, \partial\widetilde{P}_\varepsilon)$ is essential in homotopy; see Remark A.1. Observe that $\widetilde{T}_\varepsilon$ is a semialgebraic subset of the k-dimensional set $\widetilde{T}$; by property 4 of Lemma E.3, it is therefore k-dimensional, just like the set $\widetilde{P}_\varepsilon$. Therefore, using the theorem in Mertens [22, §4E], the fact that the projection from $\widetilde{T}_\varepsilon$ is essential in homotopy now implies that $\tilde{p}^*$: $H^k(\widetilde{P}_\varepsilon, \partial\widetilde{P}_\varepsilon) \to H^k(\widetilde{T}_\varepsilon, \partial\widetilde{T}_\varepsilon)$ is nonzero. Again because $\widetilde{T}_\varepsilon$ is k-dimensional, the universal coefficient theorem (Spanier [28, V.5.10]) shows that the following diagram commutes:

$$\begin{array}{ccc} H^k(\widetilde{P}_\varepsilon, \partial\widetilde{P}_\varepsilon) \otimes \mathbb{Q} & \xrightarrow{\approx} & H^k(\widetilde{P}_\varepsilon, \partial\widetilde{P}_\varepsilon; \mathbb{Q}) \\ \tilde{p}^* \otimes 1 \big\downarrow & & \tilde{p}^* \big\downarrow \\ H^k(\widetilde{T}_\varepsilon, \partial\widetilde{T}_\varepsilon) \otimes \mathbb{Q} & \xrightarrow{\approx} & H^k(\widetilde{T}_\varepsilon, \partial\widetilde{T}_\varepsilon; \mathbb{Q}). \end{array}$$

Because $\tilde{p}^*$ is nonzero in cohomology with integer coefficients and $H^k(\widetilde{T}_\varepsilon, \partial\widetilde{T}_\varepsilon) \approx \mathbb{Z}$, $\tilde{p}^* \otimes 1$ above is nonzero. Hence $\tilde{p}^*$ is nonzero in cohomology with coefficients in $\mathbb{Q}$. Thus we have shown that $\widetilde{T}$ satisfies the conditions for 0-stability in $\widetilde{G}$. □

Observe that for generic extensive-form games, metastability and stable essentiality induce the same set of outcomes. Indeed, this follows from the fact that metastable sets are obtained by taking Hausdorff limits of stably essential sets, and that these sets are connected sets of equilibria. Also, stable essentiality is a stronger refinement than h-stability. Finally, 0-stability is a stronger refinement than $*$-stability, which in turn is stronger than metastability. Therefore, we obtain the following corollary.

Corollary E.1. *For a generic extensive-form game with perfect recall, an outcome is* 0*-stable if and only if it is* $*$*-stable, and the sets of metastable outcomes and stably essential outcomes coincide with the set of* $*$*-stable outcomes.*

Acknowledgment. This work was funded in part by a grant from the National Science Foundation.

References

[1] Blume, L., A. Brandenburger, E. Dekel. 1991. Lexicographic probabilities and equilibrium refinements. *Econometrica* **59** 81–98.
[2] Bochnak, J., M. Coste, M.-F. Roy. 1998. *Real Algebraic Geometry*. Springer-Verlag, Berlin.
[3] Eaves, B., C. Lemke. 1981. Equivalence of LCP and PLS. *Math. Oper. Res.* **6** 475–484.
[4] Govindan, S., A. McLennan. 2001. On the generic finiteness of equilibrium outcomes in game forms. *Econometrica* **69** 455–471.
[5] Govindan, S., J.-F. Mertens. 2003. An equivalent definition of stable equilibria. *Internat. J. Game Theory* **32** 339–357.
[6] Govindan, S., R. Wilson. 2001. Direct proofs of generic finiteness of Nash equilibrium outcomes. *Econometrica* **69** 765–769.
[7] Govindan, S., R. Wilson. 2002. Structure theorems for game trees. *Proc. Natl. Acad. Sci.* **99** 9077–9080, http://www.pnas.org/cgi/reprint/99/13/9077.pdf.
[8] Govindan, S., R. Wilson. 2005. Essential equilibria. *Proc. Natl. Acad. Sci.* **102** 15706–15711, http://www.pnas.org/cgi/reprint/102/43/15706.pdf.
[9] Hillas, J. 1990. On the definition of the strategic stability of equilibria. *Econometrica* **58** 1365–1390.
[10] Hillas, J., E. Kohlberg. 2002. Foundations of strategic equilibrium. R. Aumann, S. Hart, eds. *Handbook of Game Theory*, Vol. III, Chapter 42. Elsevier, Amsterdam, 1597–1663.
[11] Hillas, J., M. Jansen, J. Potters, D. Vermeulen. 2001. On the relation among some definitions of strategic stability. *Math. Oper. Res.* **26** 611–635.
[12] Kohlberg, E., J.-F. Mertens. 1986. On the strategic stability of equilibria. *Econometrica* **54** 1003–1038.
[13] Kohlberg, E., P. Reny. 1997. Independence on relative probability spaces and consistent assessments in game trees. *J. Econom. Theory* **75** 280–313.
[14] Koller, D., N. Megiddo. 1992. The complexity of two-person zero-sum games in extensive form. *Games Econom. Behav.* **4** 528–552.
[15] Koller, D., N. Megiddo, B. von Stengel. 1996. Efficient computation of equilibria for extensive two-person games. *Games Econom. Behav.* **14** 247–259.
[16] Kreps, D., R. Wilson. 1982. Sequential equilibria. *Econometrica* **50** 863–894.
[17] Kuhn, H. 1953. Extensive games and the problem of information. *Contributions to the Theory of Games*, Vol. II. 193–216. Reprinted in H. Kuhn, ed. *Classics in Game Theory*, Princeton University Press, Princeton, NJ, 1997.
[18] McLennan, A. 1988. *Selected Topics in the Theory of Fixed Points*. University of Minnesota, Minneapolis.
[19] McLennan, A. 1989a. The space of conditional systems is a ball. *Internat. J. Game Theory* **18** 125–139.
[20] McLennan, A. 1989b. Consistent conditional systems in noncooperative game theory. *Internat. J. Game Theory* **18** 141–174.
[21] Mertens, J.-F. 1989. Stable equilibria—A reformulation, Part I: Definition and basic properties. *Math. Oper. Res.* **14** 575–625.
[22] Mertens, J.-F. 1991. Stable equilibria—A reformulation, Part II: Discussion of the definition and further results. *Math. Oper. Res.* **16** 694–753.
[23] Mertens, J.-F. 1992. The small worlds axiom for stable equilibria. *Games Econom. Behav.* **4** 553–564.
[24] Mertens, J.-F. 2004. Ordinality in noncooperative games. *Internat. J. Game Theory* **32** 387–430.
[25] Myerson, R. 1986. Multistage games with communication. *Econometrica* **54** 323–358.
[26] Rényi, A. 1956. On conditional probability spaces generated by a dimensionally ordered set of measures. *Theory Prob. Appl.* **1** 61–71.
[27] Rockefellar, R. 1970. *Convex Analysis*. Princeton University Press, Princeton, NJ.
[28] Spanier, E. 1966. *Algebraic Topology*. McGraw-Hill, New York, Reprinted by Springer-Verlag, New York, 1989.
[29] van Damme, E. 1984. A relation between perfect equilibria in extensive form games and proper equilibria in normal form games. *Internat. J. Game Theory* **13** 1–13.
[30] von Stengel, B. 1996. Efficient computation of behavior strategies. *Games Econom. Behav.* **14** 220—246.