

Sufficient conditions for stable equilibria

SRIHARI GOVINDAN

Department of Economics, University of Iowa

ROBERT WILSON

Graduate School of Business, Stanford University

A refinement of the set of Nash equilibria that satisfies two assumptions is shown to select a subset that is stable in the sense defined by Kohlberg and Mertens. One assumption requires that a selected set is invariant to adjoining redundant strategies and the other is a strong version of backward induction. Backward induction is interpreted as the requirement that each player's strategy is sequentially rational and conditionally admissible at every information set in an extensive-form game with perfect recall, implemented here by requiring that the equilibrium is quasi-perfect. The strong version requires 'truly' quasi-perfect in that each strategy perturbation refines the selection to a quasi-perfect equilibrium in the set. An exact characterization of stable sets is provided for two-player games.

KEYWORDS. Game theory, equilibrium selection, stability.

JEL CLASSIFICATION. C72.

1. INTRODUCTION

This article studies refinements of the equilibria of a non-cooperative game. As in other contributions to this subject, the aim is to sharpen Nash's (1950, 1951) original definition by imposing additional decision-theoretic criteria. We adopt the standard axiom of invariance to establish a connection between games in strategic (or 'normal') form and those in extensive form (with perfect recall, which we assume throughout). Our contribution is to show that a set selected by a refinement satisfying a strong form of the backward-induction criterion for an extensive-form game must be stable, as defined by Kohlberg and Mertens (1986) for the strategic form of the game but without their insistence on a minimal stable subset.

Thus in Section 3 we define formally the two criteria called Invariance and Strong Backward Induction, and then in Section 5 we prove:

THEOREM. *If a refinement of the Nash equilibria satisfies Invariance and Strong Backward Induction then each selected subset is stable.*

Srihari Govindan: srihari-govindan@uiowa.edu

Robert Wilson: rwilson@stanford.edu

This work was funded partly by a grant from the U.S. National Science Foundation.

The main concepts in the theorem are defined in [Section 3](#). Briefly:

Invariance Invariance requires that a refinement is immune to treating a mixed strategy as an additional pure strategy. Its role is to exclude some kinds of presentation effects. Its chief implication is that a refinement depends only on the *reduced form* of the game, i.e. only on the strategically equivalent game obtained by deleting redundant strategies from the strategic form.

Backward Induction For a game in extensive form, the criterion of backward induction or ‘sequential rationality’ is usually implemented by requiring that a selected subset includes a sequential equilibrium ([Kreps and Wilson 1982](#)). Here the criterion is strengthened by requiring conditional *admissibility*, i.e. by excluding a strategy that is weakly dominated in the continuation from some information set. This is implemented by requiring that the equilibrium is quasi-perfect ([van Damme 1984](#)).¹ Using quasi-perfection to represent backward induction brings the advantage that the generally accepted axioms of admissibility and conditional admissibility are included automatically.

Strong Backward Induction Sequential equilibrium outcomes can be sustained by many different conditional probability systems (‘beliefs’ in [Kreps and Wilson 1982](#)), and similarly, different quasi-perfect equilibria result from considering different perturbations of players’ strategies. Strong Backward Induction (SBI) requires that each perturbation of players’ strategies refines the selection further by identifying a quasi-perfect equilibrium within the selected subset.

Stability A subset of the Nash equilibria is stable if each nearby game, obtained by perturbing each player’s strategies by a ‘tremble,’ has a nearby equilibrium. This is the concept of stability defined by [Kohlberg and Mertens \(1986\)](#) but without their insistence on selecting a minimal stable subset.

SBI strengthens the criterion of ‘truly perfect’ (perfect with respect to all possible trembles) that originally motivated Kohlberg and Mertens’ (1986) definition of a stable set. In effect, SBI requires that a selected subset includes all the sequential equilibria in admissible strategies sustained by beliefs generated by perturbations of the game. Invariance and SBI together imply that a refinement selects a subset that includes a sequential equilibrium for *every* extensive form having the same *reduced* strategic form obtained by deleting redundant strategies ([Section 2.3](#) provides an explicit example). Due to Invariance, this implication is stronger than the property of a proper equilibrium of a strategic form; viz., a proper equilibrium induces a sequential equilibrium in each extensive form with the same (non-reduced) strategic form. In general a stable

¹A quasi-perfect equilibrium differs from a perfect equilibrium ([Selten 1975](#)) of the extensive form of a game by excluding a player’s anticipation of his own trembles. We slightly modify van Damme’s definition of an ε -quasi-perfect equilibrium.

subset might not contain a sequential equilibrium [Kohlberg and Mertens \(1986](#), example in Figure 11), so the subclass of stable subsets allowed by the theorem represents a refinement of stability.

For readers who are not familiar with the early literature on refinements from the 1970–80s, [Section 2](#) reviews informally the main antecedents of this article and presents some motivating examples. For a survey and critical examination of equilibrium refinements see [Hillas and Kohlberg \(2002\)](#). After the formulation in [Section 3](#), stability is characterized and the theorem is proved in [Section 4](#) for a game with two players, which is simpler than the general proof in [Section 5](#). Concluding remarks are in [Section 6](#). Appendices provide direct proofs for two cases of special interest.

2. BACKGROUND AND MOTIVATION

The central concept in the study of non-cooperative games is the definition of equilibrium proposed by [Nash \(1950, 1951\)](#). Nash interprets a player's strategy as a 'mixed' strategy, i.e. a randomization over pure strategies, each of which is a complete plan specifying the action to be taken in each contingency that might arise in the course of the game. Thus a game is specified by the strategic form that assigns to the players their utility payoffs from each profile of their pure strategies, and by extension, expected payoffs to each profile of their (mixed) strategies. Nash's definition of an equilibrium profile of strategies requires that each player's strategy is an optimal reply to the other players' strategies. Although Nash's definition can be applied to a game in the extensive form that describes explicitly the evolution of play, it depends only on the strategic form derived from the extensive form.

[Selten \(1965, 1975\)](#) initiated two lines of research aimed at refining Nash's definition of equilibrium. The first line invokes directly various decision-theoretic criteria that are stronger than the criteria Nash invokes. For example, admissibility and invariance are relevant criteria for a game in strategic form, and subgame perfection, sequential rationality (as in sequential equilibria), and quasi-perfection are relevant for a game in extensive form. The second line pursues a general method based on examining perturbations of the game. Its purpose is to obtain refinements that satisfy many decision-theoretic criteria simultaneously. For example, requiring that an equilibrium is affected slightly by perturbations excludes inadmissible equilibria, i.e. that use weakly dominated pure strategies. These two lines have basically the same goal although they use different methods. That goal is to characterize equilibria that are 'self-enforcing' according to a higher standard than Nash's definition requires. Perturbation methods have been remarkably successful, but the technique is often complicated, and for applications it often suffices to impose decision-theoretic criteria directly.

Both lines strengthen Nash's definition so as to exclude equilibria that are considered implausible. For example, in the context of the strategic form one wants to exclude an equilibrium that uses a weakly dominated pure strategy, or that depends on the existence of a pure strategy that is not an optimal reply at the equilibrium. In the context of the extensive form, one wants to exclude an equilibrium that is 'not credible' because it relies on a player's commitment to a strategy—when in fact no ability to commit is




FIGURE 1. A game with an equilibrium that is not credible if C cannot commit to her strategy a, d .

represented explicitly in the extensive form. For instance, consider the game in **Figure 1** in which player R (Row) chooses between T (top) and B (bottom) and then C (Col) responds. The equilibrium $(T; a, d)$ is considered not credible because it relies on C's threat to respond to B with d , whereas in the actual event C prefers c to d .

2.1 Refinements based on specific criteria

Selten (1965) began the first line of research. He argued that extensive-form considerations enable a selection among the Nash equilibria. He proposed selecting from among the equilibria one that induces an equilibrium in each subgame of the extensive form, i.e. one that is subgame-perfect. Subgame-perfection requires that each player's strategy is consistent with the procedure of backward induction used in the analysis of a decision tree with a single decision maker. Kreps and Wilson's (1982) definition of sequential equilibrium extends this approach to games with imperfect information. They require that the continuation from each contingency (a player's information set) is optimal with respect to a conditional probability system (for assessing the probabilities of prior histories) that is consistent with the structure of the game and other players' strategies. Van Damme's (1984) definition of quasi-perfect equilibrium imposes further restrictions described in **Section 2.2**.

The main deficiency of these refinements is that they depend sensitively on which extensive form is used, i.e. they are plagued with presentation effects. For instance, in **Figure 2** in the top presentation there is one sequential equilibrium $(r, r'; b)$ and others in which R chooses ℓ, ℓ' . But in the bottom presentation subgame-perfection yields only the outcome $(r, r'; b)$ in the subgame in strategic form. Some additional criterion like Invariance is needed to ensure that the refinement does not distinguish among strategically equivalent games. And even then deficiencies remain; e.g., a sequential equilibrium can use an inadmissible strategy, as in **Figure 4** below.

Typical of other work in this vein is Cho and Kreps' (1987) Intuitive Criterion for selecting among the sequential equilibria of signaling games. They require that there




FIGURE 2. A game with multiple sequential equilibrium outcomes in the top presentation, and only one in the bottom presentation (Hillas 1996).

cannot be some type of the sender that surely gains from deviating were the receiver to respond with a strategy that is optimal based on a belief that assigns zero probability to those types of the sender that cannot gain from the deviation. That is, an equilibrium fails the Intuitive Criterion if the receiver's belief fails to recognize that the sender's deviation is a credible signal about his type. For example, in Figure 3 the sequential equilibria with the outcome $(r, r'; b)$ are rejected because R would gain by choosing ℓ' in the bottom contingency if C were to choose b' , which is optimal for C if she recognizes the deviation as a credible signal that the bottom contingency has occurred—and indeed credibility is implied by the fact that in the top contingency R cannot gain by deviating to ℓ .

Stability essentially implies the Intuitive Criterion and its extensions by Banks and Sobel (1987). In particular, Kohlberg and Mertens (1986, Proposition 6) prove that a stable set S contains a stable set of the game obtained by deleting strategies that are inferior responses at all equilibria in S .

2.2 Refinements based on perturbations

Selten (1975) also opened the second line of research. He proposed selecting an equilibrium that is the limit of equilibria of perturbed games. The advantage of this computational method is that it assures that various decision-theoretic criteria are satisfied. Applied to the strategic form it assures admissibility, and applied to (the agent strategic form of) the extensive form it yields a sequential equilibrium. Selten defines a perfect




FIGURE 3. A signaling game with two sequential equilibrium outcomes, one of which is rejected by the Intuitive Criterion

equilibrium of the strategic form as the limit $\sigma = \lim_{\varepsilon \downarrow 0} \sigma^\varepsilon$ of a sequence of profiles of completely mixed strategies for which $\sigma_n^\varepsilon(s) \leq \varepsilon$ if the pure strategy s is an inferior reply for player n against σ^ε . Myerson (1978) defines a proper equilibrium analogously except that if s is inferior to s' then $\sigma_n^\varepsilon(s) \leq \varepsilon \sigma_n^\varepsilon(s')$. An advantage of a proper equilibrium of the strategic form is that it induces a quasi-perfect and hence sequential equilibrium in every extensive form with that strategic form van Damme (1984, Theorem 1), Kohlberg and Mertens (1986, Proposition 0).

Selten shows that an equivalent definition of a perfect equilibrium is that each player's strategy is an optimal reply to each profile of completely mixed strategies of other players in a sequence converging to their equilibrium profile. For a game in extensive form, van Damme's (1984) definition of a quasi-perfect equilibrium is similar: each player uses only actions at an information set that are part of an optimal continuation in reply to perturbations of other players' strategies converging to their equilibrium strategies. This ensures admissibility of continuation strategies in each contingency, but importantly, while taking account of small trembles by other players, the player ignores his own trembles both currently and also later in the game. Figure 4 shows van Damme's example in which both $(T, a; c)$ and $(B, a; c)$ are sequential equilibria,² but the second is not quasi-perfect—as is evident from the fact that R's strategy (T, a) weakly dominates (B, a) , which is therefore inadmissible in the strategic form. van Damme (1984) shows that a proper equilibrium induces a quasi-perfect equilibrium, and hence a sequential equilibrium, in every extensive form with that strategic form. A partial converse is that a quasi-perfect equilibrium induces a perfect equilibrium of the strategic form.

Subsequent development of refinements based on perturbations was influenced greatly by the work of Kohlberg and Mertens (1986) (KM hereafter). They envisioned characterizing an ideal refinement by decision-theoretic criteria adopted as axioms. To identify what the ideal refinement would be, they examined several that satisfy most of the criteria they considered.

²In fact each is perfect in the strategic form with three agents representing the two players.




FIGURE 4. A game with multiple sequential equilibria but only $(T, a; c)$ is quasi-perfect.

KM's analysis relies on a fundamental mathematical fact that we explain below using **Figure 5**. KM show that the graph of the equilibrium correspondence is homeomorphic to the (one point compactification of the) space of games obtained by varying players' payoffs, i.e. the graph is a deformed copy of the space of games. This 'structure theorem' is much more specific than the usual weak characterization of the equilibrium correspondence as upper-semi-continuous. The structure theorem is illustrated schematically in the figure as though the spaces of games and strategy profiles are each one-dimensional. Equilibrium components of game G , shown as vertical segments of the graph, are intrinsic to the study of refinements because (a) games in extensive form are nongeneric in the space of games, and (b) for an extensive-form game whose payoffs are generic in the subspace of games with the same game tree, the outcome of a sequential equilibrium is obtained by all equilibria in the same component (**Kreps and Wilson 1982**, **Govindan and Wilson 2001**), i.e. they agree along the equilibrium path.

KM's basic conclusion is that a refinement should select a subset of equilibria that is 'stable' against *all* perturbations of the strategic form of the game in a sufficiently rich class. To avoid confusions of terminology, below we use 'robust' rather than 'stable,' or say that the subset 'survives perturbations' in the sense that every perturbed game nearby has an equilibrium nearby. Their conclusion depends on several preliminary considerations.

1. A robust subset exists. The homeomorphism implies that every game has a component of its equilibria that survives all payoff perturbations in the sense that every nearby game has a nearby equilibrium. In the figure, the isolated equilibria #1 and #2 and the component #4 are robust in this sense. Not shown is the further important property that some robust subset satisfies Invariance, i.e. if game G is enlarged by treating some mixed strategy as a pure strategy then the strategically equivalent subset in the enlarged game is also robust.³

³In **Govindan and Wilson (2005)** we show that a component is essential (has nonzero index) if and only if it satisfies a slightly stronger criterion called uniform hyperstability.




FIGURE 5. Schematic diagram of the graph of equilibria over the space of games obtained by varying payoffs.

2. A refinement should consider a sufficiently rich class of perturbations. In the figure the component #3 of equilibria of game G has two endpoints that are each 'perfect' in the sense that nearby games to the right of G have nearby equilibria, but games to the left of G have no equilibria near this component.
3. A refinement should select a subset rather than a single equilibrium. In the figure the component #4 is robust as a set, but no single point is robust—games to the right of G have equilibria only near the top endpoint and games to the left have equilibria only near the bottom endpoint.
4. A refinement can consider all perturbations in any sufficiently rich class of perturbations. The preceding three considerations remain true for various classes of perturbations that are smaller than the class of all payoff perturbations. For example, the effect of perturbations of players' strategies by trembles (as in Selten's formulation) induces a (lower dimensional) subclass of payoff perturbations.

Based on these considerations, KM define three refinements based on successively smaller subclasses of perturbations of the strategic form of the game. In each case they include the auxiliary requirements that a selected subset is closed and minimal among those with the specified property. Further, Invariance is always assumed, so the property must persist for every enlargement of the game obtained by treating any finite set of mixed strategies as additional pure strategies—or equivalently, the refinement depends only on the reduced strategic form of the game.

- A *hyperstable* subset of the equilibria survives all payoff perturbations.
- A *fully-stable* subset survives all polyhedral perturbations of players' strategies. That is, each neighborhood of the subset contains an equilibrium of the perturbed game obtained by restricting each player to a closed convex polyhedron of completely mixed strategies, provided each of these polyhedra is sufficiently close (in Hausdorff distance) to the simplex of that player's mixed strategies.
- A *stable* subset survives all trembles of players' strategies. Specifically, for every $\varepsilon > 0$ there exists $\bar{\delta} > 0$ such that for each $\delta \in (0, \bar{\delta})^N$ and completely mixed profile η the perturbed game obtained by replacing each pure strategy s_n of each player n by the mixture $[1 - \delta_n]s_n + \delta_n\eta_n$ has an equilibrium within ε of the subset.

From (1) above, some component contains a hyperstable subset, and within that there is a fully-stable subset, which in turn contains a stable subset, since smaller subsets can survive smaller classes of perturbations. A fully-stable subset is useful because it necessarily contains a proper equilibrium of the strategic form that induces a sequential equilibrium in every extensive form with that strategic form. But a hyperstable or fully-stable subset can include equilibria that use inadmissible strategies, which is why KM focus on stable subsets. However, for an extensive-form game, a stable subset need not contain a sequential equilibrium.

KM's article ended perplexed that no one of their three refinements ensures both admissibility and backward induction. This conundrum was resolved later by **Mertens (1989, 1991)** who defined a stronger refinement (called here *Mertens-stability*) that satisfies all the criteria examined by KM, and more besides. Because Mertens-stability is couched in the apparatus of the theory of homology developed in algebraic topology, it is not widely accessible to non-specialists and we do not set forth its definition here.

Our purpose in this article is to show that KM's stability can be refined to select stable subsets that do indeed satisfy admissibility and backward induction. As described in the opening paragraphs of **Section 1**, we prove that a refinement that satisfies Invariance and Strong Backward Induction (SBI), defined formally in **Section 3**, contains a stable subset. Existence of stable subsets with these properties is assured because they are implied by Mertens-stability.

2.3 A simple example

In this subsection an example illustrates the interaction between Invariance and SBI. The example is sufficiently simple that it suffices to ignore admissibility and to represent backward induction by sequential equilibrium.

Figure 6 shows at the top an extensive-form game Γ in which players R and C alternate moves. In the subgame-perfect equilibrium each player chooses down at his first opportunity, which we represent by the pure strategy D , ignoring his subsequent choice were the player to err. With this convention the strategic form G of this game is shown in **Table 1**.

There is a single component of the Nash equilibria of G , in which R uses D and C uses any mixed strategy for which the probability of D is $\geq \frac{2}{3}$. The component of perfect




FIGURE 6. Top: A game Γ between players R and C. Bottom: The game modified so that player R can commit to the mixed strategy $x(\delta)$ after rejecting D .

equilibria requires further that C's probability of a is zero. The minimal stable subset consists of the two endpoints of the perfect-equilibrium component; viz., the subgame-perfect equilibrium (D, D) , and (D, y°) where $y^\circ = (\frac{2}{3}, \frac{1}{3}, 0)$ is the mixed strategy with probabilities $\frac{2}{3}$ and $\frac{1}{3}$ for D and d .

Figure 6 shows at the bottom the expansion $\Gamma(\delta)$ of the extensive form in which player R can reject D and then choose either the mixed strategy $x(\delta) = (1 - \delta, \delta/4, 3\delta/4)$ or continue by choosing A and then later d or a if C chooses A . The two information sets of C indicate that C cannot know whether R chose $x(\delta)$. Thus, the expanded game has imperfect information in the sense of imperfect observability of R's choice. Even so, the reduced strategic form of the expanded game is the same as the original strategic form G in Table 1, since $x(\delta)$ is a redundant strategy.

Assume that $0 < \delta < 1$. One can easily verify that there is a unique sequential equilibrium in the expanded extensive form $\Gamma(\delta)$. In the strategic form this is the equilibrium in which R chooses D and C randomizes between D and d with probabilities $\alpha(\delta)$ and $1 - \alpha(\delta)$, where $\alpha(\delta) = [8 + \delta]/[12 - 3\delta]$. In the extensive form this is sustained by C's belief at her first information set that the conditional probability that R chose $x(\delta)$ given

		C		
		D	d	a
R	D	1, 0	1, 0	1, 0
	d	0, 2	3, 0	3, 0
	a	0, 2	0, 4	0, 0

TABLE 1. Strategic form G of the game Γ in [Figure 6](#).

that he rejected D is $\beta(\delta) = 2/[2 + \delta]$. By Bayes' Rule, the conditional probability that R chose $x(\delta)$ given that A occurred is $p = \frac{2}{3}$.

A refinement that includes the sequential equilibrium of each expanded extensive form $\Gamma(\delta)$ must therefore include every profile $(D; \alpha(\delta), 1 - \alpha(\delta), 0)$ as δ varies between zero and one. Since $\alpha(0) = \frac{2}{3}$ and $\alpha(1) = 1$ this requires the refinement to select the entire component of perfect equilibria. In fact, this is precisely the Mertens-stable subset. [Appendix A](#) extends the analysis of this example to general two-player games with perfect information.

As described in [Section 2.1](#), refinements like subgame-perfection and sequential equilibrium that focus on the extensive form aim to exclude equilibria that are not credible because they rely on an ability to commit to a strategy that is not modeled explicitly. In contrast, the strategic form seems to assume commitment. And seemingly worse, the above example illustrates that a refinement that satisfies Invariance allows a player to commit to a redundant strategy midway in the extensive form. The resolution of this conundrum lies in the additional assumption of backward induction. Together, Invariance and backward induction imply that a selected subset must include a sequential equilibrium of each expanded extensive form with the same reduced strategic form. In [Section 4](#) and [Section 5](#) we prove in general that the conjunction of Invariance and Strong Backward Induction implies that a selected subset must contain a stable subset of the reduced strategic form.

3. FORMULATION

We consider games with finite sets of players and pure strategies. The strategic form of a game is specified by a payoff function $G : \prod_{n \in N} S_n \rightarrow \mathbb{R}^N$ where N is the set of players and S_n is player n 's set of pure strategies. Interpret a pure strategy s_n as a vertex of player n 's simplex $\Sigma_n = \Delta(S_n)$ of mixed strategies. The sets of profiles of pure and mixed strategies are $S = \prod_n S_n$ and $\Sigma = \prod_n \Sigma_n$.

In a game G a pure strategy s_n of player n is *redundant* if n has in G a mixed strategy $\sigma_n \neq s_n$ that for every profile of mixed strategies of the other players yields for every player the same expected payoff as s_n yields. The strategic form is *reduced* if no pure strategy is redundant. Say that two games are *equivalent* if their reduced strategic forms are the same (except for labelling of pure strategies). We use the reduced strategic form of a game as the representative of its equivalence class. Each game in an equivalence class is an expansion of its reduced strategic form obtained by adjoining redundant pure strategies.

Say that two mixed strategies of a player in two equivalent games are *equivalent* if they induce the same probability distribution (called their reduced version) on his pure strategies in the reduced strategic form. Similarly, two profiles are equivalent if the players' strategies are equivalent, and two sets of profiles are equivalent if they induce the same sets of profiles in the reduced strategic form.

In general, a refinement is a correspondence that assigns to each game a collection of closed, nonempty subsets of its equilibria, called the selected subsets. However, each equilibrium induces a family of equivalent equilibria for each expansion of the game obtained by adding redundant strategies. Therefore, we assume:⁴

ASSUMPTION 1 (Invariance). Each selected subset is equivalent to a subset selected for an equivalent game. Specifically, if G and $\tilde{G}$ are equivalent games then a subset Σ° selected for G is equivalent to some subset $\tilde{\Sigma}^\circ$ selected for $\tilde{G}$.

In particular, every equivalent game has a selected subset whose reduced version is a selected subset of the reduced strategic form. This is slightly weaker than requiring that a refinement depends only on the equivalence classes of games and strategies; cf. [Mertens \(2003\)](#) for a detailed discussion of invariance, and more generally the concept of ordinality for games.

To each game in strategic form we associate those games in extensive form with perfect recall that have that strategic form. Each extensive form specifies a disjoint collection $H = \{H_n \mid n \in N\}$ of the players' information sets, and for each information set $h \in H_n$ it specifies a set $A_n(h)$ of possible actions by n at h . In its strategic form the set of pure strategies of player n is $S_n = \{s_n : H_n \rightarrow \cup_{h \in H_n} A_n(h) \mid s_n(h) \in A_n(h)\}$. The projection of S_n onto h and n 's information sets that follow h is denoted $S_{n|h}$; that is, $S_{n|h}$ is the set of n 's continuation strategies from h . Let $S_n(h)$ be the set of n 's pure strategies that choose all of n 's actions necessary to reach $h \in H_n$, and let $S_n(a|h)$ be the subset of strategies in $S_n(h)$ that choose $a \in A_n(h)$. Then a completely mixed strategy $\sigma_n \gg 0$ induces the conditional probability $\sigma_n(a|h) = \sum_{s_n \in S_n(a|h)} \sigma_n(s_n) / \sum_{s_n \in S_n(h)} \sigma_n(s_n)$ of choosing a at h . More generally, a behavior strategy $\beta_n \in \prod_{h \in H_n} \Delta(A_n(h))$ assigns to each information set h a probability $\beta_n(a|h)$ of action $a \in A_n(h)$ if h is reached. [Kuhn \(1953\)](#) shows that mixed and behavior strategies are payoff-equivalent in extensive-form games with perfect recall.

Given a game in extensive form, an action perturbation $\varepsilon : H \rightarrow (0, 1)^2$ assigns to each information set a pair $(\underline{\varepsilon}(h), \bar{\varepsilon}(h))$ of small positive numbers, where $0 < \underline{\varepsilon}(h) \leq \bar{\varepsilon}(h)$. Use $\{\varepsilon\}$ to denote a sequence of action perturbations that converges to 0.

DEFINITION 1 (Quasi-Perfect). A sequence $\{\sigma^\varepsilon\}$ of profiles is $\{\varepsilon\}$ -quasi-perfect⁵ if for

⁴The proof of the main theorem uses only a slightly weaker version: a selected subset is equivalent to a superset of one selected for an expanded game obtained by adding redundant strategies.

⁵This definition differs from [van Damme \(1984\)](#) in that the upper bound $\bar{\varepsilon}(\cdot)$ of the error probability can differ across information sets. However, it is easily shown that the set of quasi-perfect equilibria as defined by van Damme is the set of all profiles of behavioral strategies equivalent to limits of sequences of $\{\varepsilon\}$ -quasi-perfect equilibria as defined here in terms of mixed strategies. van Damme does not impose an explicit lower bound but because the strategies are completely mixed there is an implicit lower bound that

each $a \in A_n$, $h \in H_n$, $n \in N$ and each action perturbation ε :

- (i) $\sigma_n^\varepsilon(a|h) \geq \underline{\varepsilon}(h)$, and
- (ii) $\sigma_n^\varepsilon(a|h) > \bar{\varepsilon}(h)$ only if a is an optimal action at h in reply to σ^ε ; that is, only if $s_n(h) = a$ for some continuation strategy $s_n \in \arg \max_{s \in S_{n|h}} E[G_n | h, s, \sigma_{-n}^\varepsilon]$.

Suppose that $\sigma_n(\cdot|h) = \lim_{\varepsilon \downarrow 0} \sigma_n^\varepsilon(\cdot|h)$. Then this definition says that player n 's continuation strategy at h assigns a positive conditional probability $\sigma_n(a|h) > 0$ to action a only if a is chosen by a continuation strategy that is an optimal reply to sufficiently small perturbations $(\sigma_{n'}^\varepsilon)_{n' \neq n}$ of other players' strategies. Thus when solving his dynamic programming problem, player n takes account of vanishingly small trembles by other players but ignores his own trembles later in the game. In particular, this enforces admissibility of continuation strategies conditional on having reached h . **van Damme (1984)** shows that the pair $(\mu, \beta) = \lim_{\varepsilon \downarrow 0} (\mu^\varepsilon, \beta^\varepsilon)$ of belief and behavior profiles is a sequential equilibrium, where σ^ε induces at $h \in H_n$ the conditional probability $\mu_n^\varepsilon(t|h)$ of node $t \in h$ and the behavior $\beta_n^\varepsilon(a|h) = \sigma_n^\varepsilon(a|h)$ is player n 's conditional probability of choosing a at h .

Our second assumption requires that each sequence of action perturbations induces a further selection among the profiles in a selected set.

ASSUMPTION 2 (Strong Backward Induction). For a game in extensive form with perfect recall for which a refinement selects a subset Σ° of equilibria, for each sequence $\{\varepsilon\}$ of action perturbations there exists a profile $\sigma \in \Sigma^\circ$ that is the limit of a convergent subsequence $\{\sigma^\varepsilon\}$ of $\{\varepsilon\}$ -quasi-perfect profiles.

This assumption could as well be called ‘truly’ or Strong Quasi-Perfection. As the proofs in Sections 4 and 5 show, our **main theorem** remains true if **Assumption 2** is weakened by requiring action perturbations to satisfy the additional restriction that $\underline{\varepsilon}(h) = \bar{\varepsilon}(h)$ for all h . The reason we do not do so is conceptual. The lower bound $\underline{\varepsilon}(\cdot)$ reflects the requirement that every action of a player is chosen with positive probability, while $\bar{\varepsilon}(\cdot)$ provides the upper bound on the “error probability” of suboptimal actions at an information set.

We conclude this section by defining stability. In general, a closed subset of the equilibria of a game in strategic form is deemed stable if, for any neighborhood of the set, every game obtained from a sufficiently small perturbation of payoffs has an equilibrium in the neighborhood. However, to ensure admissibility, KM focus on sets that are stable only against those payoff perturbations induced by strategy perturbations.

For $0 \leq \delta \leq 1$, let $P_\delta = \{(\lambda_n \tau_n)_{n \in N} \mid (\forall n) 0 \leq \lambda_n \leq \delta, \tau_n \in \Sigma_n\}$ and let ∂P_δ be the topological boundary of P_δ . For each $\eta \in P_1$ and $n \in N$, let $\bar{\eta}_n = \sum_{s \in S_n} \eta_n(s)$. Given any $\eta \in P_1$, a perturbed game $G(\eta)$ is obtained by replacing each pure strategy s_n of player n with $\eta_n + (1 - \bar{\eta}_n)s_n$. Thus $G(\eta)$ is the perturbed game in which the strategy sets of the players are restricted so that the probability that n plays a strategy $s \in S_n$ must be

shrinks to zero as $\varepsilon \downarrow 0$.

at least $\eta_{n,s}$. For a vector $(\lambda, \tau) \in [0, 1]^N \times \Sigma$, we sometimes write $G(\lambda, \tau)$ to denote the perturbed game $G((\lambda_n \tau_n)_{n \in N})$.

DEFINITION 2 (Stability). A closed set Σ° of equilibria of the game G is stable if for each $\epsilon > 0$ there exists $\delta > 0$ such that for each $\eta \in P_\delta \setminus \partial P_\delta$ the perturbed game $G(\eta)$ has an equilibrium within ϵ of Σ° .

To avoid the trivially stable set of all equilibria, KM focus on minimal stable sets:

DEFINITION 3 (KM-Stability). A set of equilibria of the game G is KM-stable if it is a minimal stable set.

4. TWO-PLAYER GAMES

This section provides a direct proof of the main theorem for the special case of two players. It is simpler than the proof of the general case in [Section 5](#) because two-player games have a linear structure. This structure enables a generalization—(iii) in the following theorem—of the characterization of stability obtained by [Cho and Kreps \(1987\)](#) and [Banks and Sobel \(1987\)](#) for the special case of sender-receiver signaling games in extensive form with generic payoffs.

THEOREM 1 (Characterization of Stability). *Let G be a 2-player game, and let Σ° be a closed subset of equilibria of G . The following statements are equivalent.*

- (i) Σ° is a stable set of the game G .
- (ii) For each $\tau \in \Sigma \setminus \partial \Sigma$ there exists sequence σ^k in Σ converging to a point in Σ° and a corresponding sequence λ^k in $(0, 1)$ converging to 0, such that σ^k is an equilibrium of $G(\lambda^k \tau)$ for all k .
- (iii) For each $\tau \in \Sigma \setminus \partial \Sigma$ there exists $\sigma^\circ \in \Sigma^\circ$, a profile $\tilde{\sigma} \in \Sigma$, and $0 < \lambda \leq 1$ such that, for each player n , $\lambda \sigma_n^\circ(s) + [1 - \lambda] \tilde{\sigma}_n(s)$ is an optimal reply against both σ° and the profile $\sigma^* = \lambda \tau + [1 - \lambda] \tilde{\sigma}$.

PROOF. We prove first that (i) implies (ii). Suppose Σ° is a stable set. Fix $\tau \in \Sigma \setminus \partial \Sigma$. Then for each positive integer k one can choose $\lambda^k \in (0, 1/k)$ and an equilibrium σ^k of $G(\lambda^k \tau)$ whose distance from Σ° is less than $1/k$. Let σ° be the limit of a convergent subsequence of σ^k . Then $\sigma^\circ \in \Sigma^\circ$, which completes the proof.

Next we prove that (ii) implies (iii). Fix $\tau \in \Sigma \setminus \partial \Sigma$. (ii) assures us that there exists a sequence λ^k in $(0, 1)$ converging to zero and a sequence σ^k of equilibria of $G(\lambda^k \tau)$ converging to an equilibrium σ° in Σ° . By passing to a subsequence if necessary, we can assume that the set of optimal replies in G to σ^k is the same for all k . Define σ^* and λ to be the first elements of the sequences of σ^k and λ^k . And let $\tilde{\sigma} = [\sigma^* - \lambda \tau] / (1 - \lambda)$. Then, $\sigma^* = (1 - \lambda) \tilde{\sigma} + \lambda \tau$. Because σ^* is an equilibrium of $G(\lambda \tau)$, $\tilde{\sigma}$ is an optimal reply to σ^* . Because the best replies are constant along the sequence of σ^k , $\tilde{\sigma}$ is a best reply all along the sequence of σ^k and hence to the limit, σ° . Because the members of the sequence

of σ^k are equilibria of perturbed games converging to σ° , σ° is optimal against σ^k for large k and is therefore optimal against the entire sequence σ^k (in particular against σ^* and the limit σ° itself). Thus, $\tilde{\sigma}$ and σ° satisfy the optimality condition of (iii).

Lastly we prove that (iii) implies (i) by showing that Σ° satisfies the property in Definition 2 of a stable set. Fix an ϵ -neighborhood of Σ° . Take a sufficiently fine simplicial subdivision of Σ such that: (i) the union U of the simplices of this complex that intersect Σ° is contained in its ϵ -neighborhood; and (ii) the best-reply correspondence is constant over the interior of each simplex. Because G is a two-player game, this simplicial subdivision can be done such that each simplex is actually a convex polytope. Observe that U is itself a closed neighborhood of Σ° . Let Q be the set of all pairs $(\eta, \sigma) \in P_1 \times U$ such that σ is an equilibrium of $G(\eta)$; and let Q^0 be the set of $(0, \sigma) \in Q$, namely, the set of equilibria of the game G that are contained in U . By property (ii) of the triangulation and because the simplices are convex polytopes, Q and Q^0 are finite unions of polytopes. Triangulate Q such that Q^0 is a subcomplex, and take a barycentric subdivision so that Q^0 becomes a full subcomplex. Because Q is a union of polytopes, both the triangulation and the projection map $p : Q \rightarrow P_1$ can be made piecewise-linear. Let X be the union of simplices of Q that intersect Q^0 . Because Q^0 is a full subcomplex, the intersection of each simplex of Q with Q^0 is a face of the simplex. Let $X^0 = X \cap Q^0$ and let X^1 be the union of simplices of X that do not intersect Q^0 . Given $x \in X$, there exists a unique simplex K of X that contains x in its interior. Let K^0 be the face of K that is in X^0 , and let K^1 be the face of K spanned by the vertices of K that do not belong to K^0 . The face K^1 is then contained in X^1 . Therefore, x is expressible as a convex combination $[1 - \alpha]x^0 + \alpha x^1$, where $x^i \in K^i$ for $i = 0, 1$; moreover, this combination is unique if $x \notin X^0 \cup X^1$. Finally, $p(x) = [1 - \alpha]p(x^0) + \alpha p(x^1) = \alpha p(x^1)$ because the projection map p is piecewise affine.

Choose $\delta^* > 0$ such that for each $(\eta, \sigma) \in X^1$, $\max_n \bar{\eta}_n > \delta^*$. Such a choice is possible because X^1 is a compact subset of Q that is disjoint from Q^0 . Fix now $\delta_1, \delta_2 < \delta^*$ and $\tau \in \Sigma$. The proof is complete if we show that the game $G(\delta_1 \tau_1, \delta_2 \tau_2)$ has an equilibrium in U . By (iii), there exists $\sigma^\circ \in \Sigma^\circ$, $\tilde{\sigma} \in \Sigma$, and $0 < \mu \leq 1$ such that $\sigma(\gamma) = ((1 - \gamma \delta_n) \sigma_n^\circ + \gamma \delta_n ((1 - \mu) \tilde{\sigma}_n + \mu \tau_n))_{n=1,2}$ is an equilibrium of $G(\gamma \mu (\delta_1 \tau_1, \delta_2 \tau_2))$ for all $0 \leq \gamma \leq 1$. Because $\sigma(0) = \sigma^\circ \in \Sigma^\circ$, we can choose γ sufficiently small that the point $x = (\gamma \mu (\delta_1 \tau_1, \delta_2 \tau_2), \sigma(\gamma))$ belongs to $X \setminus (X^0 \cup X^1)$; hence there exists a unique $\alpha \in (0, 1)$ and $x^i \in X^i$ for $i = 0, 1$ such that x is an α -combination of x^0 and x^1 . As remarked before, $p(x) = \alpha p(x^1)$. Therefore, there exists $\sigma \in \Sigma$ such that $x^1 = (\gamma^* \mu (\delta_1 \tau_1, \delta_2 \tau_2), \sigma)$, where $\gamma^* = \gamma / \alpha$. Because points in X^1 project to $P_1 \setminus P_{\delta^*}$, $\gamma^* \mu \delta_n > \delta^*$ for some n ; that is, $\gamma^* \mu > 1$ since $\delta_n < \delta^*$ for each n by assumption. Therefore, the point $[1 - 1/\gamma^* \mu]x^0 + [1/\gamma^* \mu]x^1$ corresponds to an equilibrium of the game $G(\delta_1 \tau_1, \delta_2 \tau_2)$ that lies in U . This proves (i). $\square$

The characterization in (iii) of Theorem 1 can be stated equivalently in terms of a lexicographic probability system (LPS) as in Blume et al. (1991a,b). As a matter of terminology, given an LPS $(\sigma_m^0, \dots, \sigma_m^k)$ for player m , we say that for player $n \neq m$, a strategy σ_n is a *better reply* against the LPS than another strategy σ'_n if it is a lexicographic better reply. (Here, and throughout the paper, by a better reply we mean a *strictly* better reply as opposed to a *weakly* better reply.) And σ_n is a best reply against the LPS if there is no better reply.

COROLLARY 1 (Lexicographic Characterization). *A closed set Σ° of equilibria of G is a stable set if and only if for each $\tau \in \Sigma \setminus \partial\Sigma$ there exists $\sigma^0 \in \Sigma^\circ$, a profile $\tilde{\sigma} \in \Sigma$, and for each player n , an LPS $\mathcal{L}_n = (\sigma_n^0, \dots, \sigma_n^{K_n})$ where $K_n > 0$ and $\sigma_n^{K_n} = [1 - \lambda_n]\tilde{\sigma}_n + \lambda_n\tau_n$ for some $\lambda_n \in (0, 1]$, such that for each player n every strategy that is either: (i) in the support of σ^k with $k < K_n$ or (ii) in the support of $\tilde{\sigma}_n$ if $\lambda_n < 1$, is a best reply to the LPS of the other player.*

PROOF. The necessity of the condition follows from (iii). As for sufficiency, we show that the condition of the corollary implies (ii). Fix $\tau \in \Sigma \setminus \partial\Sigma$ and let $(\mathcal{L}_1, \mathcal{L}_2)$ be as in the corollary. Choose an integer K that is greater than K_n for each n . For each n , define a new LPS $\mathcal{L}'_n = (\hat{\sigma}_n^0, \dots, \hat{\sigma}_n^K)$ as follows: for $0 \leq k \leq K - K_n$, $\hat{\sigma}_n^k = \sigma_n^0$; for $K - K_n + 1 \leq k < K$, $\hat{\sigma}_n^k = \sigma_n^{k-K+K_n}$; $\hat{\sigma}_n^K = [\mu/\lambda_n]\sigma_n^{K_n} + [1 - \mu/\lambda_n]\sigma_n^0$, where $\mu = \min(\lambda_1, \lambda_2)$. Observe that for each n ,

$$\hat{\sigma}_n^K = \mu\tau_n + [(\mu(1 - \lambda_n)\tilde{\sigma}_n + (\lambda_n - \mu)\sigma_n^0)/\lambda_n].$$

Therefore, the LPS profile $(\mathcal{L}'_1, \mathcal{L}'_2)$ satisfies the condition of the corollary as well. For $\alpha > 0$, define $\sigma(\alpha)$ by $\sigma_n(\alpha) = (\sum_{k=0}^K \alpha^k)^{-1} (\sum_{k=0}^K \alpha^k \hat{\sigma}_n^k)$. For all small α , we now have that $\sigma(\alpha)$ is an equilibrium for the perturbed game $G(\lambda(\alpha)\tau)$ where $\lambda(\alpha) = (\sum_{k=0}^K \alpha^k)^{-1} \alpha^K \mu$. Because $\sigma(\alpha)$ converges to σ^0 as α goes to zero, the condition of the corollary implies (ii) of the **theorem**. $\square$

We show in **Appendix A** that for generic two-person extensive form games, the requirements for stability in the above lexicographic characterization can be weakened. As mentioned before, (a version of) the characterization of stability in (iii) of **Theorem 1** is obtained by **Cho and Kreps (1987)** and **Banks and Sobel (1987)** for the special case of sender-receiver signaling games in extensive form with generic payoffs—games like the one in **Figure 3**. In **Appendix C** we show directly that if a component of the equilibria violates this condition then a single redundant strategy can be adjoined to obtain an equivalent game that has no proper equilibrium yielding the same outcome.

We conclude this section by proving the main theorem for two-player games.⁶

PROPOSITION 1 (Sufficiency of the Assumptions). *If a refinement satisfies Invariance and Strong Backward Induction then for any two-player game a selected subset is a stable set of the equilibria of its strategic form.*

PROOF. Let G be the strategic form of a 2-player game. Suppose that $\Sigma^\circ \subset \Sigma$ is a set selected by a refinement that satisfies Invariance and Strong Backward Induction. Let $\tau = (\tau_1, \tau_2)$ be any profile in the interior of Σ . We show that Σ° satisfies the condition of **Corollary 1** for τ . Construct as follows the extensive-form game Γ with perfect recall that has a strategic form that is an expansion of G . In Γ each player n first chooses

⁶An anonymous referee has shown in his or her report that for 2-player games, in the definition of Strong Backward Induction, quasi-perfection in the extensive form can be replaced by perfection in the strategic form (we do not reproduce the referee's proof here). This reflects indirectly the fact that for 2-player games the sets of strategic-form perfect and admissible equilibria coincide.

whether or not to use the mixed strategy τ_n , and if not, then which pure strategy in S_n to use. Denote the two information sets at which n makes these choices by h'_n and h''_n . At neither of these does n have any information about the other player's analogous choices. In Γ the set of pure strategies for player n is $S_n^* = \{\tau_n\} \cup S_n$ (after identifying all strategies where n chooses to play τ_n at his first information set h'_n) and the corresponding simplex of mixed strategies is Σ_n^* . For each $\delta > 0$ in a sequence converging to zero, let $\{\varepsilon\}$ be a sequence of action perturbations that require the minimum probability of each action at h'_n to be $\underline{\varepsilon}(h'_n) = \delta$ and the maximum probability of suboptimal actions at h''_n to be $\bar{\varepsilon}(h''_n) = \delta^2$. By Invariance, the refinement selects a set $\tilde{\Sigma}^\circ$ for Γ that is a subset of those strategies equivalent to ones in Σ° . By Strong Backward Induction there exists a subsequence $\{\tilde{\sigma}^\varepsilon\}$ of $\{\varepsilon\}$ -quasi-perfect profiles converging to some point $\tilde{\sigma}^0 \in \tilde{\Sigma}^\circ$. If necessary by passing to a subsequence, by [Blume et al. \(1991b, Proposition 2\)](#) there exists for each player n : (i) an LPS $\tilde{\mathcal{L}}_n = (\tilde{\sigma}_n^0, \tilde{\sigma}_n^1, \dots, \tilde{\sigma}_n^{K_n})$, with members $\tilde{\sigma}_n^k \in \Sigma_n^*$, $\bigcup_{k=0}^{K_n} \text{supp } \tilde{\sigma}_n^k = S_n^*$; and (ii) for each $0 \leq k < K_n$ a sequence of positive numbers $\lambda_n^k(\varepsilon)$ converging to zero such that each $\tilde{\sigma}_n^k$ in the sequence is expressible as the nested combination $(1 - \lambda_n^0(\varepsilon))\tilde{\sigma}_n^0 + \lambda_n^0((1 - \lambda_n^1(\varepsilon))\tilde{\sigma}_n^1 + \lambda_n^1(\dots + \lambda_n^{K_n-1}(\varepsilon)\tilde{\sigma}_n^{K_n}))$. Let k_n^* be the smallest k for which $\tilde{\sigma}_n^k$ assigns positive probability to the “pure” strategy τ_n of the expanded game.

LEMMA 1. *Suppose $s_n \in S_n$ is not a best reply to the LPS of the other player. Then $s_n \in S_n$ is assigned zero probability by $\tilde{\sigma}_n^k$ for $k \leq k_n^*$, and $k_n^* > 0$.*

PROOF OF LEMMA. Since s_n is not a best reply to the LPS of the other player, sufficiently far along the sequence s_n is not a best reply against $\tilde{\sigma}^\varepsilon$. Quasi-perfection requires that

$$\tilde{\sigma}_n^\varepsilon(\tau_n|h'_n) \geq \underline{\varepsilon}(h'_n) = \delta \quad \text{and} \quad \tilde{\sigma}_n^\varepsilon(s_n|h''_n) \leq \bar{\varepsilon}(h''_n) = \delta^2.$$

Hence $\lim_{\varepsilon \downarrow 0} \tilde{\sigma}_n^\varepsilon(s_n|h''_n)/\tilde{\sigma}_n^\varepsilon(\tau_n|h'_n) = 0$. Therefore $\tilde{\sigma}_n^k(s_n) = 0$ for all $k \leq k_n^*$, which proves the first statement of the lemma. As for the second statement, observe that the pure strategy τ_n is not a best reply to the LPS of the other player, since τ_n , viewed as a mixed strategy in Σ_n , has full support. Therefore, far along the sequence, τ_n is not a best reply to $\tilde{\sigma}_n^\varepsilon$. Quasi-perfection now requires that τ_n is assigned the minimum probability δ by the sequence and hence its probability in the limit $\tilde{\sigma}_n^0$ of the sequence $\tilde{\sigma}_n^\varepsilon$ is zero, which means that $k_n^* > 0$. $\square$

From the LPS $\tilde{\mathcal{L}}$ construct for each player n an LPS $\mathcal{L}_n = (\sigma_n^0, \sigma_n^1, \dots, \sigma_n^{k_n^*})$ for the game G by letting σ_n^k be the mixed strategy in Σ_n that is equivalent to $\tilde{\sigma}_n^k$. Because $\tilde{\sigma}^0 \in \tilde{\Sigma}^0$, σ^0 belongs to Σ^0 . For each n , let λ_n be the probability assigned to the strategy τ_n in $\tilde{\sigma}_n^{k_n^*}$. By the definition of k_n^* , $\lambda_n > 0$. If $\lambda_n \neq 1$, let σ'_n be the mixed strategy in Σ_n that is given by the conditional distribution over S_n induced by $\tilde{\sigma}_n^{k_n^*}$, that is, the probability of a pure strategy $s \in S_n$ under σ'_n is $(1 - \lambda_n)^{-1}\tilde{\sigma}_n^{k_n^*}(s)$; if $\lambda_n = 1$, let σ'_n be an arbitrary strategy in Σ_n . By the definition of $\mathcal{L}_n$, $\sigma_n^{k_n^*} = \lambda_n\tau_n + [1 - \lambda_n]\sigma'_n$. For each n , define an LPS $\mathcal{L}'_n$ as follows: if $k_n^* > 0$, then $\mathcal{L}'_n = \mathcal{L}_n$; otherwise, $\mathcal{L}'_n = (\sigma_n^0, \sigma_n^0)$. We now show that the LPS profile $(\mathcal{L}'_1, \mathcal{L}'_2)$ satisfies the conditions of [Corollary 1](#) for τ . Each LPS $\mathcal{L}'_n$ has at

least two levels, where the last level is $\lambda_n \tau_n + (1 - \lambda_n) \sigma'_n$, $\lambda_n > 0$. We now show that the optimality property in [Corollary 1](#) holds for each player n . If $k_n^* = 0$ then by [Lemma 1](#) every strategy of player n is optimal. On the other hand, if $k_n^* > 0$ (and thus $\mathcal{L}'_n = \mathcal{L}_n$) then the first result of [Lemma 1](#) and the following two observations imply the optimality property. (1) By the definition of k_n^* and $\mathcal{L}_n$, for each $k < k_n^*$ the probability of each $s \in S_n$ is the same under σ_n^k and $\tilde{\sigma}_n^k$; (2) if $\lambda_n \neq 1$ then every strategy in the support of σ'_n is in the support of $\tilde{\sigma}_n^{k_n^*}$. Thus we have shown that Σ° is a stable set. $\square$

5. N-PLAYER GAMES

This section provides the proof of the main theorem for the general case with N players.⁷

We begin with some definitions. For a real-valued analytic function (or more generally a power series) $f(t) = \sum_{i=0}^{\infty} a_i t^i$ in a single variable t , the order of f , denoted $o(f)$, is the smallest integer i such that $a_i \neq 0$. The order of the zero function is $+\infty$. We say that a power series f is positive if $a_{o(f)} > 0$; thus if f is an analytic function then f is positive if and only if $f(t)$ is positive for all sufficiently small $t > 0$. Suppose f and g are two power series. We say that $f > g$ if $f - g$ is positive. We have the following relations for orders of power series: $o(fg) = o(f) + o(g)$; and $o(f + g) \geq \min(o(f), o(g))$, with equality if f and g are both nonnegative (or nonpositive). Suppose f and g are real-valued analytic functions defined on $(-\bar{t}, \bar{t})$ where $t > 0$ and $g \neq 0$. If $o(f) \geq o(g)$ then there exists an analytic function $h : (-\bar{t}, \bar{t}) \rightarrow \mathbb{R}$ such that for each $t \neq 0$, $h(t) = f(t)/g(t)$, i.e. dividing f by g yields an analytic function.

By a slight abuse of terminology, we call a function $F : [0, \bar{t}] \rightarrow X$, where X is a subset of a Euclidean space $\mathbb{R}^L$, analytic if there exists an analytic function $F' : (-\delta, \delta) \rightarrow \mathbb{R}^L$, $\delta > \bar{t}$, such that F' agrees with F on $[0, \bar{t}]$. For an analytic function $F : [0, \bar{t}] \rightarrow \mathbb{R}^k$, the order $o(F)$ of F is $\min_i o(F_i)$. If $\sigma : [0, \bar{t}] \rightarrow \Sigma$ is an analytic function then for each pure strategy s_n of player n his payoff $G_n(\sigma_{-n}(t), s_n)$ in the game G is an analytic function as well, since payoff functions are multilinear in mixed strategies. We say that s_n is a best reply of order k for player n against an analytic function $\sigma(t)$ if for all $s'_n \in S_n$, $G_n(\sigma_{-n}(t), s_n) - G_n(\sigma_{-n}(t), s'_n)$ is either nonnegative or has order at least $k + 1$; also, s_n is a best reply to σ if it is a best reply of order ∞ .

LEMMA 2. *Suppose $\sigma, \tau : [0, \bar{t}] \rightarrow \Sigma$ are two analytic functions such that $o(\sigma - \tau) > k$. If s_n is not a best reply of order k against σ then it is not a best reply of order k against τ .*

PROOF. Let s'_n be a pure strategy such that $G_n(\sigma_{-n}(t), s_n) - G_n(\sigma_{-n}(t), s'_n)$ is negative and has order, say, $\ell \leq k$. Let $\tau' = \tau - \sigma$. We can then write

$$G_n(\tau_{-n}(t), s_n) - G_n(\tau_{-n}(t), s'_n) = [G_n(\sigma_{-n}(t), s_n) - G_n(\sigma_{-n}(t), s'_n)] \\ + \sum_{s_{-n}} \sum_{N' \subsetneq N \setminus \{n\}} \left(\prod_{n' \in N'} \sigma_{n', s_{n'}}(t) \prod_{n'' \in N \setminus (N' \cup \{n\})} \tau'_{n'', s_{n''}}(t) \right) [G_n(s_{-n}, s_n) - G_n(s_{-n}, s'_n)].$$

⁷The proof suggests how to use the method in [Govindan and Klumpp \(2002\)](#) to extend to N -player games the characterization of stable sets for two-player games in [Theorem 1](#), but we do not include it here.

The first term on the right in the above expression is negative and has order ℓ by assumption. Therefore, to prove the result it is enough to show that the order of the double summation is at least $k + 1$: it then follows that the whole expression is negative and has order ℓ . To prove this last statement, using the above property of the order of sums of power series, it is sufficient to show that each of the summands in the second term has order at least $k + 1$. Consider now a summand for a fixed s_{-n} and $N' \subsetneq N \setminus \{n\}$. If both s_n and s'_n give the same payoff against s_{-n} then the order of this term is ∞ . Otherwise, using the property of the order of products of functions, the order of this term is

$$\sum_{n' \in N'} o(\sigma_{n', s_{n'}}) + \sum_{n'' \notin (N' \cup \{n\})} o(\tau'_{n'', s_{n''}}) > k,$$

where the inequality follows from the following two facts: (i) the order of each $\sigma_{n', s_{n'}}$ is at least zero; and (ii) there exists at least one $n'' \notin (N' \cup \{n\})$ and for any such n'' the order of $\tau'_{n'', s_{n''}}$ is greater than k by assumption. $\square$

We use the following version of a result of [Blume et al. \(1991b, Proposition 2\)](#).

LEMMA 3. *If the map $\tau_n : [0, \bar{t}] \rightarrow \Sigma_n$ is analytic then there exists $0 < t_* \leq \bar{t}$ such that for each $t \leq t_*$, $\tau_n(t) = \sum_{k=0}^K f_n^k(t) \tau_n^k$, where $K \leq |S_n|$, each τ_n^k is in Σ_n , and each map $f_n^k : [0, t_*] \rightarrow \mathbb{R}_+$ is analytic.*

PROOF. There is nothing to prove if τ_n is a constant map. Therefore, assume that it is not. Let $\tau_n^0 = \tau_n(0)$ and let S_n^0 be the support of τ_n^0 . For each $0 \leq t \leq \bar{t}$, define $f_n^0(t)$ to be $\min_{s \in S_n^0} \tau_{n,s}(t) / \tau_{n,s^0}^0$. Remark that if for some t , $\tau_n(t) \neq \tau_n^0$, then $f_n^0(t) < 1$; indeed, if $f_n^0(t) \geq 1$ then for each pure strategy s (even if it is not in S_n^0) $\tau_{n,s}(t) \geq \tau_{n,s^0}^0$ and therefore $\tau_n(t) = \tau_n^0$. Since $\tau_n(t)$ is analytic, there now exist $0 < t_0 \leq \bar{t}$ and $s^0 \in S_n^0$ such that for all $t \leq t_0$, $f_n^0(t) = \tau_{n,s^0}(t) / \tau_{n,s^0}^0$, i.e. f_n^0 is analytic on $[0, t_0]$. Moreover, since $\tau_n(t)$ is not a constant function and since $\tau_{n,s^0}(0) > 0$, t_0 can be chosen such that for all $0 < t \leq t_0$: (i) $\tau_n(t) \neq \tau_n^0$; and (ii) $\tau_{n,s^0}(t) > 0$. Therefore, for $0 < t \leq t_0$, $0 < f_n^0(t) < 1$, where the fact that it is positive follows from the fact that $\tau_{n,s^0}(t)$ is positive while the other inequality follows from our earlier remark, since $\tau_n(t) \neq \tau_n^0$. We claim now that there is a well defined analytic function $\tau_n^1 : [0, t_0] \rightarrow \Sigma_n$ where for $t \neq 0$, $\tau_n^1(t) = [1 - f_n^0(t)]^{-1} [\tau_n(t) - f_n^0(t) \tau_n^0]$. Indeed to prove this claim it is sufficient to show that: (i) $1 - f_n^0(t)$ is positive; and (ii) for each $s \in S_n$, $\tau_{n,s}(t) - f_n^0(t) \tau_{n,s^0}^0$ is nonnegative with $o(\tau_{n,s}(t) - f_n^0(t) \tau_{n,s^0}^0) \geq o(1 - f_n^0(t))$. Point (i) follows from the fact that $f_n^0(t) < 1$. As for point (ii), for $s \notin S_n^0$, $\tau_{n,s}(t) - f_n^0(t) \tau_{n,s^0}^0 = \tau_{n,s}(t) \geq 0$, while if $s \in S_n^0$, $\tau_{n,s}(t) - f_n^0(t) \tau_{n,s^0}^0 \geq \tau_{n,s}(t) - (\tau_{n,s}(t) / \tau_{n,s^0}^0) \tau_{n,s^0}^0 = 0$. Thus $\tau_{n,s}(t) - f_n^0(t) \tau_{n,s^0}^0$ is nonnegative for each $s \in S_n$. And, as a result, we also have that for each $s \in S_n$, $o(\tau_{n,s}(t) - f_n^0(t) \tau_{n,s^0}^0) \geq \min_{s'} o(\tau_{n,s'}(t) - f_n^0(t) \tau_{n,s^0}^0) = o(\sum_{s'} (\tau_{n,s'}(t) - f_n^0(t) \tau_{n,s^0}^0)) = o(1 - f_n^0(t))$. Thus, $\tau_n^1(t)$ is a well-defined analytic function. We now have that $\tau_n(t) = f_n^0(t) \tau_n^0 + [1 - f_n^0(t)] \tau_n^1(t)$ for each $t \leq t_0$. For $0 < t \leq t_0$, the support of $\tau_n^1(t)$ is contained in that of $\tau_n(t)$, since $f_n(t) < 1$; also, $\tau_{n,s^0}^1(t)$ is zero, while $\tau_{n,s^0}(t)$ is obviously not. Thus, $\text{supp } \tau_n^1(t) \subsetneq \text{supp } \tau_n(t)$ for all $0 < t \leq t_0$.

If $\tau_n^1(t)$ is a constant function mapping to, say, $\tau_n^1 \in \Sigma_n$, let $f_n^1(t) = (1 - f_n^0(t))$ and then $\tau_n(t) = f_n^0(t)\tau_n^0 + f_n^1(t)\tau_n^1$ and we are done. So, assume that $\tau_n^1(t)$ is not a constant function. Now let $\tau_n^1 = \tau_n^1(0)$ and $S_n^1 = \text{supp } \tau_n^1$. We can repeat the above construction for the function $\tau_n^1(t)$ to obtain the following. There exists $0 < t_1 \leq t_0$, $s_n^1 \in S_n^1$, analytic functions $\hat{f}_n^1(t) : [0, t_1] \rightarrow \mathbb{R}$ and $\tau_n^2(t) : [0, t_1] \rightarrow \Sigma_n$ such that for each $0 < t \leq t_1$: $\hat{f}_n^1(t) = \min_{s \in S_n^1} \tau_{n,s}^1(t)/\tau_{n,s}^1 = \tau_{n,s^1}^1(t)/\tau_{n,s^1}^1$; $0 < \hat{f}_n^1(t) < 1$; $\tau_n^2(t) = [1 - \hat{f}_n^1(t)]^{-1}[\tau_n^1(t) - \hat{f}_n^1(t)\tau_n^1]$. Obviously, $\text{supp } \tau_n^2(t) \subseteq \text{supp } \tau_n^1(t)$. On the other hand, as before, $f_n^1(t) = \tau_{n,s^1}^1(t)/\tau_{n,s^1}^1$ and thus $\tau_{n,s}^2(t) = 0$ while $\tau_{n,s^1}^1(t)$ is not: in particular $\text{supp } \tau_n^2(t) \subsetneq \text{supp } \tau_n^1(t) \subsetneq \text{supp } \tau_n(t)$ for all $0 < t \leq t_1$. We now have $\tau_n(t) = f_n^0(t)\tau_n^0 + f_n^1(t)\tau_n^1 + f_n^2(t)\tau_n^2(t)$ where $f_n^1(t) = [1 - f_n^0(t)]\hat{f}_n^1(t)$ and $f_n^2(t) = [1 - f_n^0(t)][1 - \hat{f}_n^1(t)]$. We can continue this process to obtain a sequence of analytic functions $\tau_n^k : [0, t_k] \rightarrow \Sigma_n$, $k = 0, 1, \dots$, with $0 < t_k \leq t_{k-1}$, $\tau_n^k \equiv \tau_n^k(0)$, and a corresponding sequence of analytic functions $f_n^k : [0, t_k] \rightarrow [0, 1]$ such that for each k and $0 < t \leq t_k$, $\text{supp } \tau_n(t) \supsetneq \text{supp } \tau_n^1(t) \supsetneq \text{supp } \tau_n^2(t) \cdots \supsetneq \text{supp } \tau_n^k(t)$ and $\tau_n(t) = \sum_{i=0}^{k-1} f_n^i(t)\tau_n^i + f_n^k(t)\tau_n^k(t)$. This process must terminate in a finite number of steps, in the sense that there exists $K \leq |S_n|$ such that $\tau_n^K(t)$ is a constant function. $\square$

MAIN THEOREM. *If a refinement satisfies Invariance and Strong Backward Induction then for any game a selected subset is a stable set of the equilibria of its strategic form.*

PROOF. We show that if a refinement selects a subset $\Sigma^\circ \subset \Sigma$ of profiles that is not a stable set for the strategic-form game G then it satisfies Invariance only if it violates Strong Backward Induction.

Suppose Σ° is not a stable set. Then there exists $\epsilon > 0$ such that for each $\delta \in (0, 1)$ there exists $\eta \in P_\delta \setminus \partial P_\delta$ such that the perturbed game $G(\eta)$ does not have an equilibrium in the ϵ -neighborhood U of Σ° . Take a closed semi-algebraic neighborhood X of Σ° that is contained in U . Let $A = \{(\lambda, \tau) \in (0, 1)^N \times (\Sigma \setminus \partial \Sigma) \mid G(\lambda, \tau) \text{ has no equilibrium in } X\}$; then A is nonempty and there exists $\tau^\circ \in \Sigma$ such that $(0, \tau^\circ)$ is in the closure of A . Further, because X is semi-algebraic, A too is semi-algebraic. Therefore, by the Nash Curve Selection Lemma (cf. [Bochnak et al. 1998](#), Proposition 8.1.13) there exists $\bar{t} > 0$ and a semialgebraic analytic map $t \mapsto (\lambda(t), \tau(t))$ from $[0, \bar{t}]$ to $[0, 1]^N \times \Sigma$ such that $(\lambda(0), \tau(0)) = (0, \tau^\circ)$ and $(\lambda(t), \tau(t)) \in A$ for all $t \in (0, \bar{t}]$. Define the compact semi-algebraic set

$$Y = \{(t, \sigma) \in [0, \bar{t}] \times X \mid (\forall s_n \in S_n) \sigma_{n,s_n} \geq \lambda_n(t)\tau_{n,s_n}(t)\}.$$

Observe that if $(t, \sigma) \in Y$ with $t \neq 0$ then σ is not an equilibrium of $G(\lambda(t), \tau(t))$.

LEMMA 4. *There exists a positive integer p such that for every analytic function $z \mapsto (t(z), \sigma(z))$ from an interval $[0, \bar{z}]$ to Y , where $t(z)$ is a positive function, there exists a player n and a pure strategy $s_n \in S_n$ such that $\sigma(z) > \lambda_n(t(z))\tau_{n,s_n}(t(z))$ and s_n is not a best reply of order $o(t(z))p$ against $\sigma(z)$.*

PROOF OF LEMMA. Define the maps $\alpha, \beta : Y \rightarrow \mathbb{R}$ via

$$\alpha(t, \sigma) = \max_{n, s_n \in S_n} \{[\sigma_{n,s_n} - \lambda_n(t)\tau_{n,s_n}(t)] \times \max_{s'_n \in S_n} [G_n(s'_n, \sigma_{-n}) - G_n(s_n, \sigma_{-n})]\}$$

and $\beta(t, \sigma) = t$. By construction, $\alpha, \beta \geq 0$ and $\alpha^{-1}(0) \subseteq \beta^{-1}(0)$. By Lojasiewicz's inequality (Bochnak et al. 1998, Corollary 2.6.7) there exist a positive scalar c and a positive integer p such that $c\alpha \geq \beta^p$. Given an analytic map $z \mapsto (t(z), \sigma(z))$ as in the statement of the lemma, observe for each n, s_n, s'_n , that $\sigma_{n, s_n}(z) - \lambda(t(z))\tau_{n, s_n}(t(z))$ and $G_n(s'_n, \sigma_{-n}(z)) - G_n(s_n, \sigma_{-n}(z))$ are also analytic in z . Therefore there exists a pair n, s_n that achieves the maximum in the definition of α for all small z . Then

$$\max_{s'_n} [G_n(s'_n, \sigma_{-n}(z)) - G_n(s_n, \sigma_{-n}(z))] \geq \alpha(t(z), \sigma(z)) \geq (t(z))^p / c,$$

where the first inequality follows from the fact that $\sigma_{n, s_n}(z) - \lambda_n(t(z))\tau_{n, s_n}(t(z)) \leq 1$. By assumption, $t(z)$ is positive. Therefore, $\max_{n, s_n, s'_n} [G_n(s'_n, \sigma_{-n}(z)) - G_n(s_n, \sigma_{-n}(z))]$ is also a positive analytic function and, being greater than $c^{-1}(t(z))^p$, has order at most $o(t(z))p$. $\square$

Using Lemma 3 express each $\tau_n(t)$ as the sum $\sum_{k=0}^{K_n} f_n^k(t)\tau_n^k$, where each τ_n^k is a mixed strategy in Σ_n and $f_n^k: [0, \bar{t}] \rightarrow \mathbb{R}_+$ is analytic. Construct the game Γ in extensive form in which each player n chooses among the following, while remaining uninformed of the others' choices. Player n first chooses whether to commit to the mixed strategy τ_n^0 or not; if not then n chooses between τ_n^1 or not, and so on for $k = 2, \dots, K_n$; and if n does not commit to any strategy τ_n^k then n chooses among the pure strategies in S_n . Let $\tilde{S}$ and $\tilde{\Sigma}$ be the sets of pure and mixed-strategy profiles in Γ . (As in the two-person case, for each player n and each $0 \leq k \leq K_n$ we identify all strategies of n that choose, at the relevant information set, to play the strategy τ_n^k .) Because the strategic form of Γ is an expansion of G , Invariance implies that for the game Γ the refinement selects a subset of those strategies equivalent to Σ° . We now show that the refinement does not satisfy Strong Backwards Induction in the game Γ . The argument is by contradiction. Suppose it does satisfy SBI. For perturbations of the game Γ use the following action perturbation: for the information set where n chooses between τ_n^k or not, use $\underline{\varepsilon}_n^k(t) = \bar{\varepsilon}_n^k(t) = \lambda_n(t)f_n^k(t)$; and at the information set where n chooses among the strategies in S_n , use $\underline{\varepsilon}_n^{K_n+1}(t) = \bar{\varepsilon}_n^{K_n+1}(t) = t^{p+1}$, where p is as in Lemma 4.

Let E be the set of $(t, \sigma) \in (0, \bar{t}] \times \tilde{\Sigma}$ such that σ is an $\{\varepsilon(t)\}$ -quasi-perfect equilibrium of Γ (i.e. satisfying conditions 1 and 2 of Definition 2.2) whose reduced-form strategy profile in Σ lies in X . Because the minimum error probabilities are analytic functions of t , E is a semi-analytic set.⁸ Strong Backward Induction requires that there exists $\tilde{\sigma}^0 \in \Sigma^*$ such that the reduced form of $\tilde{\sigma}^0$ belongs to Σ^0 and $(0, \tilde{\sigma}^0)$ belongs to the closure of E . By the Curve Selection Lemma (cf. Lojasiewicz 1993, II.3) there exists an analytic function $z \mapsto (t(z), \tilde{\sigma}(z))$ from $[0, \bar{z}]$ to $[0, \bar{t}] \times \tilde{\Sigma}$ such that $(t(z), \tilde{\sigma}(z)) \in E$ for all $z > 0$ and $(t(0), \tilde{\sigma}(0)) = (0, \tilde{\sigma}^0)$. By construction, $0 < o(t(z)) < \infty$.

From $\tilde{\sigma}(z)$ construct the analytic function $\hat{\sigma}(z)$ as follows: for each player n , choose a strategy s_n^* in Γ such that $o(\tilde{\sigma}_{n, s_n^*})$ is zero; that is, a strategy in the support of $\tilde{\sigma}_n(0)$. Let S'_n be the set of all pure strategies s_n of the original game G that are chosen with the

⁸ $A \subseteq \mathbb{R}^k$ is semi-analytic if for all $x \in \mathbb{R}^k$ there exists a neighborhood U of x such that $A \cap U$ is a finite union of sets of the form $\{y \in U \mid f_1(y) = \dots = f_m(y) = 0, g_1(y) > 0, \dots, g_n(y) > 0\}$ where $f_1, \dots, f_m, g_1, \dots, g_n$ are analytic on U .

minimum probability in $\tilde{\sigma}(t)$ (that is, with probability $(t(z))^{p+1}$); let $\hat{\sigma}_{n,s_n}(z) = 0$ for each $s_n \in S'_n$; define $\hat{\sigma}_{n,s_n^*}(z) = \tilde{\sigma}_{n,s_n^*}(z) + |S'_n|(t(z))^{p+1}$; and finally, let the probabilities of the other strategies in $\hat{\sigma}$ be the same as in $\tilde{\sigma}$. Obviously, $o(\hat{\sigma}(z) - \tilde{\sigma}(z)) \geq o(t(z))(p+1) > o(t(z))p$, where the second inequality follows from the fact that $0 < o(t(z)) < \infty$.

If $\hat{\sigma}_{n,s_n}(z) > 0$ for some $s_n \in S_n$ then s_n is a best reply against $\tilde{\sigma}(z)$; hence by **Lemma 2** s_n is a best reply of order $o(t(z))p$ against $\hat{\sigma}(z)$. Likewise, for each k the strategy s_n that plays τ_n^k at the appropriate information set is optimal of order $o(t(z))p$ against $\hat{\sigma}_n(z)$ if $\hat{\sigma}_{n,s_n}(z) > \lambda_n(t(z))f_n^k(t(z))$.

Let $\sigma(z)$ be the reduced form of $\hat{\sigma}(z)$ in the game G . Because $\hat{\sigma}(0) = \tilde{\sigma}(0)$, there exists $z^* > 0$ such that $\sigma(z) \in X$ for all $z \leq z^*$. We claim now that we have a well-defined analytic function $\varphi : [0, z^*] \rightarrow Y$ given by $\varphi(z) = (t(z), \sigma(z))$: indeed, t and σ are analytic functions and, as we remarked, $\sigma(z)$ is contained in X ; also, for each n and $s_n \in S_n$, $\sigma_{n,s_n}(z) \geq \lambda_n(t(z))\tau_{n,s_n}(t(z))$, since in $\tilde{\sigma}(z)$ (and therefore in $\hat{\sigma}(z)$) the “pure” strategy τ_n^k is chosen with probability at least $\lambda_n(t(z))f_n^k(t(z))$. Therefore, φ is a well-defined map, and by **Lemma 4**, there exist n, s_n such that $\sigma_n(z)$ assigns s_n more than the minimum probability even though it is not a best reply of order $o(t(z))p$ against $\sigma_n(z)$ (and $\hat{\sigma}(z)$). By the definition of $\sigma(z)$ and $\hat{\sigma}(z)$, either (i) s_n is assigned a positive probability by $\hat{\sigma}(z)$ or (ii) a strategy τ_n^k (containing s_n in its support, when viewed as a mixed strategy in Σ_n) is assigned a probability greater than $\lambda_n(\tau(z))f_n^k(t(z))$, even though it is not a best reply of order $o(t(z))p$ against $\hat{\sigma}(z)$, which contradicts the conclusion of the previous paragraph. In the game Γ , therefore, for any sequence of sufficiently small t there cannot be a sequence of $\{\varepsilon(t)\}$ -quasi-perfect profiles whose reduced forms are in X . Thus Strong Backward Induction is violated. $\square$

6. CONCLUDING REMARKS

The contribution of the **Main Theorem** is the demonstration that a ‘truly perfect’ form of backward induction, namely Strong Backward Induction, in an extensive-form game implies stability in the strategic form, provided one links the two forms by requiring Invariance.

We accept the arguments for Invariance and admissibility adduced by Kohlberg and Mertens as entirely convincing—to do otherwise would reject a cornerstone of decision theory. Our assumptions differ primarily in using quasi-perfection to specify a form of backward induction that ensures admissibility. Our result differs in that we obtain a refinement in which a stable set must include a quasi-perfect and hence a sequential equilibrium of every extensive form with the same reduced strategic form.⁹

In spite of its awkward name, quasi-perfection seems to be an appropriate refinement of weaker forms of backward induction such as sequential equilibrium. Some

⁹Kohlberg and Mertens (1986, Appendix D) establish a comparable result in the special case of an isolated equilibrium that assigns positive probability to every optimal strategy, and that is perfect with respect to every perturbation of behavior strategies in every extensive form with the same reduced strategic form. But such an ‘essential’ equilibrium need not exist; indeed, this seems to be the original motivation for their definition of stable sets of equilibria. In Kohlberg and Mertens (1986, Appendix E) they argue that implementing backward induction by sequential equilibrium cannot suffice to imply stability. They cite van Damme’s (1984) work on quasi-perfection but unfortunately they do not consider whether it is an appropriate implementation of backward induction.

strengthening is evidently necessary since a sequential equilibrium can use inadmissible strategies and strategies that are dominated in the continuation from an information set. Strong Backward Induction is used in the proofs mainly to establish existence of lexicographic probability systems that ‘respect preferences’ as defined by [Blume et al. \(1991b\)](#). Thus [Assumption 2](#) might state directly that each sequence of perturbations of an extensive form should refine the selected set by selecting a lexicographic equilibrium that respects preferences, analogously to (iii) of [Theorem 1](#) and [Corollary 1](#) for two-player games. It seems plausible that quasi-perfection can be characterized in terms of a lexicographic equilibrium with the requisite properties.

APPENDIX

A. GENERIC TWO-PERSON EXTENSIVE-FORM GAMES

This appendix proves the analog of [Corollary 1](#) for generic extensive-form games.

We consider a fixed finite game tree with perfect recall and two players having the sets S_n and Σ_n of pure and mixed strategies for each player $n \in N = \{1, 2\}$, and the spaces $S = \prod_n S_n$ and $\Sigma = \prod_n \Sigma_n$ of pure and mixed strategy profiles. Given this game tree, let $\mathcal{G}$ be the Euclidean space of players’ payoffs in extensive-form games with this tree.

Define $\mathcal{E}$ to be the following subset of the graph of the perturbed equilibrium correspondence over the space of games: $\mathcal{E}$ is the set of those $(G, \lambda, \tau, \sigma) \in \mathcal{G} \times [0, 1] \times \Sigma \times \Sigma$ such that (a) σ is an equilibrium of $G(\lambda\tau)$ and (b) if $\lambda = 0$ then there exists a sequence $\lambda^k \in (0, 1)$ converging to zero and a sequence σ^k of equilibria of $G(\lambda^k\tau)$ converging to σ . Denote by p the natural projection from $\mathcal{E}$ to $\mathcal{G}$.

LEMMA 5. *For each game G , $p^{-1}(G)$ is compact.*

PROOF. Fix a game $G \in \mathcal{G}$. Let Q be the set of those $(\eta, \sigma) \in P_1 \times \Sigma$ such that σ is an equilibrium of $G(\eta)$. And, let Q^0 be the subset of $(0, \sigma) \in Q$. Because G is a two-player game, Q and Q^0 are unions of polytopes. Triangulate Q such that Q^0 is a full subcomplex and each simplex of Q is convex. Let X be the union of the simplices of Q that intersect Q^0 but are not contained in Q^0 . Because Q^0 is a full subcomplex of Q , the intersection of each simplex of X with Q^0 is a proper face of the simplex. Let $X^0 = X \cap Q^0$ and let X^1 be the union of simplices of X that do not intersect Q^0 . Given a point $x \in X$, let K be the simplex that contains x in its interior. Then $K \cap X^i$ is a nonempty face of K for $i = 0, 1$ and every vertex of K belongs to either X^0 or X^1 . Therefore, x is expressible as a convex combination $\alpha(x)x^1 + [1 - \alpha(x)]x^0$, where $x^i \in K \cap X^i$ for $i = 0, 1$.

We are now ready to prove the lemma. Obviously it is sufficient to prove that $p^{-1}(G)$ is closed. Accordingly, consider a sequence $(G, \lambda^k, \tau^k, \sigma^k)$ in $\mathcal{E}$ converging to a point $(G, \lambda, \tau, \sigma)$. We show that $(G, \lambda, \tau, \sigma)$ belongs to $\mathcal{E}$. The result is clear if $\lambda \neq 0$. Therefore, assume that $\lambda = 0$. By the definition of $\mathcal{E}$, we can assume without loss of generality that $\lambda^k \in (0, 1)$ for all k .

Because $\lambda^k > 0$, $(\eta^k, \sigma^k) \in Q \setminus Q^0$, where $\eta^k = \lambda^k \tau^k$. If necessary by passing to a subsequence, we can assume that the entire sequence (η^k, σ^k) is contained in the interior of a simplex K of Q . The limit $(0, \sigma)$ of the sequence then belongs to K and,

therefore, K is contained in X . For each k there exists $(0, \sigma^{0k}) \in X^0 \cap K$, $(\eta^{1k}, \sigma^{1k}) \in X^1 \cap K$ and a number $\alpha^k > 0$ such that $(\eta^k, \sigma^k) = [1 - \alpha^k](0, \sigma^{0k}) + \alpha^k(\eta^{1k}, \sigma^{1k})$. Clearly, for each k , $\eta^{1k} = \mu^k \tau^k$, where $\mu^k = \lambda^k / \alpha^k$. Since $X^1 \cap K$ is compact, by passing to an appropriate subsequence, we can assume that (η^{1k}, σ^{1k}) converges to some $(\eta^1, \sigma^1) \in X^1 \cap K$, where obviously $\eta^1 \neq 0$. Because $\eta^{1k} = \mu^k \tau^k$, and τ^k converges to τ , this implies that μ^k converges to some $\mu \neq 0$ and $\eta^1 = \mu \tau$.

As $(0, \sigma)$ and (η^1, σ^1) belong to K , so does $(v\eta^1, v\sigma^1 + (1-v)\sigma)$ for all $v \in [0, 1]$. For all $v \in (0, 1]$ now, $(G, v\mu, \tau, v\sigma^1 + (1-v)\sigma)$ belongs to $\mathcal{E}$, which shows that $(G, 0, \tau, \sigma)$ belongs to $\mathcal{E}$. $\square$

Define $\mathcal{E}_o(G) = \{(\tau, \sigma) \mid (G, 0, \tau, \sigma) \in \mathcal{E}\}$. By (iii) of **Theorem 1**, a closed set Σ° of equilibria of G is a stable subset if and only if for each $\tau \in \Sigma \setminus \partial\Sigma$ there exists $\sigma \in \Sigma^\circ$ such that $(\tau, \sigma) \in \mathcal{E}_o(G)$. We invoke a slightly weaker concept of stability that turns out to be equivalent to stability for generic games in $\mathcal{G}$.

For $G \in \mathcal{G}$, define

$$\tilde{\mathcal{E}}_o(G) = \{(\tau, \sigma) \mid \exists (G^k, \lambda^k, \tau, \sigma^k) \rightarrow (G, 0, \tau, \sigma) \text{ s.t. } \forall k, \lambda^k > 0, (G^k, \lambda^k, \tau, \sigma^k) \in \mathcal{E}\}.$$

Call a closed set Σ° of equilibria of G a pseudo-stable set of G if $(\forall \tau \in \Sigma \setminus \partial\Sigma)(\exists \sigma \in \Sigma^\circ)$ such that $(\tau, \sigma) \in \tilde{\mathcal{E}}_o(G)$. Since $\mathcal{E}_o(G) \subseteq \tilde{\mathcal{E}}_o(G)$, if Σ° is a stable set then it is a pseudo-stable set as well. The following proposition shows that for generic games the two concepts coincide.

PROPOSITION 2. *There exists a closed lower-dimensional semialgebraic subset $\mathcal{G}^\dagger$ of $\mathcal{G}$ such that for each $G \in \mathcal{G} \setminus \mathcal{G}^\dagger$ every pseudo-stable set of G is a stable set.*

PROOF. As in **Blume and Zame (1994, §2 Lemma)** one applies the Generic Local Triviality Theorem to the projection map p from $\mathcal{E}$ to $\mathcal{G}$ to establish existence of a closed lower-dimensional subset $\mathcal{G}^\dagger \subset \mathcal{G}$ such that for each connected component $\mathcal{G}^i$ of $\mathcal{G} \setminus \mathcal{G}^\dagger$ there exists a semi-algebraic fiber C^i and a homeomorphism $h^i : \mathcal{G}^i \times C^i \rightarrow p^{-1}(\mathcal{G}^i)$ such that $[p \circ h^i](G, c) = G$ for all $(G, c) \in \mathcal{G}^i \times C^i$. By **Lemma 5**, p^{-1} is compact valued and thus C^i is compact.

To prove the proposition it is sufficient to show that for each $G \in \mathcal{G} \setminus \mathcal{G}^\dagger$, $\mathcal{E}_o(G) \supseteq \tilde{\mathcal{E}}_o(G)$. Let G belong to a component $\mathcal{G}^i$ of $\mathcal{G} \setminus \mathcal{G}^\dagger$. Pick $(\tau, \sigma) \in \tilde{\mathcal{E}}_o(G)$. We now show that $(G, 0, \tau, \sigma) \in \mathcal{E}$. Since $(\tau, \sigma) \in \tilde{\mathcal{E}}_o(G)$, by definition, there exists a convergent sequence $(G^k, \lambda^k, \tau, \sigma^k) \rightarrow (G, 0, \tau, \sigma)$ with each $(G^k, \lambda^k, \tau, \sigma^k) \in \mathcal{E}$. Because $\mathcal{G}^i$ is a component of the open set $\mathcal{G} \setminus \mathcal{G}^\dagger$, the sequence $(G^k, \lambda^k, \tau, \sigma^k)$ is eventually in $p^{-1}(\mathcal{G}^i)$. Therefore, for large k there exists c^k in C^i such that $h^i(G^k, c^k) = (G^k, \lambda^k, \tau, \sigma^k)$. Because C^i is compact, there exists $c \in C^i$ that is a limit of a convergent subsequence of c^k . Obviously $(G, 0, \tau, \sigma)$ is the image of (G, c) under h^i and is therefore in $\mathcal{E}$. $\square$

We now prove the analog of **Corollary 1**.

THEOREM 2. *For each $G \in \mathcal{G} \setminus \mathcal{G}^\dagger$ a closed subset of equilibria Σ° is stable if and only if*

(*) for each $\tau \in \Sigma \setminus \partial \Sigma$ there exists an equilibrium σ^0 in Σ° , a profile $\tilde{\sigma} \in \Sigma$, and for each player n a lexicographic probability system $\mathcal{L}_n = (\sigma_n^0, \dots, \sigma_n^{K_n})$ where $K_n > 0$ and $\sigma_n^{K_n} = [1 - \lambda_n]\tilde{\sigma}_n + \lambda_n\tau_n$ for some $\lambda_n \in (0, 1]$, such that for each strategy $s \in S_n$ with $\sum_{k=0}^{K_n-1} \sigma_n^k(s) + [1 - \lambda_n]\tilde{\sigma}_n(s) > 0$ and each information set h of n that s does not exclude, s is a conditionally optimal reply at h for player n against $\sigma_m^{k(h)}$, where $m \neq n$ and $k(h)$ is the first $k \in \{0, \dots, K_m\}$ such that h is reached with positive probability when n plays s and m plays σ_m^k .

PROOF. The necessity follows by applying (iii) of [Theorem 1](#). To prove sufficiency we invoke [Proposition 2](#): it is sufficient to show that Σ° is a pseudo-stable set and thus to show that for each $\tau \in \Sigma \setminus \partial \Sigma$ there exists $\sigma \in \Sigma^\circ$ with $(\tau, \sigma) \in \tilde{\mathcal{E}}_\circ(G)$. Fix $\tau \in \Sigma \setminus \partial \Sigma$. Let $(\mathcal{L}_1, \mathcal{L}_2)$ be an LPS profile as in the statement of the theorem. The proof uses the construction and notation from the (sufficiency) proof of [Corollary 1](#). As there, obtain for each player n , the LPS $\mathcal{L}'_n = (\tilde{\sigma}_n^0, \dots, \tilde{\sigma}_n^K)$. For each n ,

$$\tilde{\sigma}_n^K = \mu\tau_n + [(\mu(1 - \lambda_n)\tilde{\sigma}_n + (\lambda_n - \mu)\sigma^0)]/\lambda_n$$

and is thus expressible as an average $\mu\tau_n + (1 - \mu)\tilde{\sigma}_n$, where $\tilde{\sigma}_n$ is a convex combination of $\tilde{\sigma}_n$ and σ_n° that equals σ_n° when $\lambda_n = 1$ and $\mu < 1$. The profile $(\mathcal{L}'_1, \mathcal{L}'_2)$ satisfies the optimality property in (*). Specifically, for each player n , each strategy s that is in the support of $\tilde{\sigma}_n^k$ for $k < K$ or in the support of $\tilde{\sigma}_n$ when $\mu < 1$ is optimal against $\mathcal{L}'_m$ in the following sense: at each information set h that s does not exclude, s is conditionally optimal against $\tilde{\sigma}_m^{k(h)}$, where $k(h)$ is the first level of $\mathcal{L}'_m$ that does not exclude h . Define $\sigma_n(\alpha)$ as in the proof of [Corollary 1](#). The optimality property for $(\mathcal{L}'_1, \mathcal{L}'_2)$ implies that for each sufficiently small α there exists a perturbed game $G(\alpha)$ such that: (a) $G = \lim_{\alpha \downarrow 0} G(\alpha)$; and (b) each strategy in the support of any σ^k with $k < K$, or in the support of $\tilde{\sigma}$ if $\mu < 1$, is an optimal reply to $\sigma(\alpha)$ in the game $G(\alpha)$.¹⁰ Therefore, $(G(\alpha), \lambda(\alpha), \tau, \sigma(\alpha))$ is in $\tilde{\mathcal{E}}$, where $\lambda(\alpha) = (\sum_{k=0}^K \alpha^k)^{-1} \alpha^K \mu$. Thus, (τ, σ^0) is in $\tilde{\mathcal{E}}_\circ(G)$ as claimed. $\square$

B. GAMES WITH PERFECT INFORMATION

This appendix generalizes the analysis of the example in [Section 2.3](#) using the formulation in [Section 3](#). We consider an extensive-form game Γ with perfect information and, for simplicity, generic payoffs and two players. As mentioned in [Section 2.2](#), a proper equilibrium induces a quasi-perfect and hence sequential equilibrium in every extensive-form game with that strategic form. Therefore, in this section we represent Strong Backward Induction simply by the assumption that a selected subset contains a proper equilibrium of the strategic form. Together with Invariance, this implies that a subset Σ° is selected for the strategic form G of Γ only if, for every game G' equivalent to G , there exists a proper equilibrium of G' that is equivalent to some equilibrium in Σ° .

¹⁰The proof of this fact follows the construction in the appendix of [Kreps and Wilson \(1982\)](#) by working backwards through the tree.

To show that this property implies that a selected subset is stable, for each strategy perturbation η of G we consider an equivalent game G' obtained by adjoining redundant strategic-form strategies that induce redundant behavioral strategies in the extensive form Γ .

The interesting aspect of the following **Proposition 3** is that, even though the proper equilibrium is merely the unique subgame-perfect equilibrium of the extensive-form game Γ , Invariance requires a refinement to select a subset that includes other equilibria in a stable set in the same component as the subgame-perfect equilibrium. This illustrates vividly that even for perfect-information games, which are usually considered trivial and easily solved by applying subgame-perfection, if one assumes Invariance then the selected subset must survive all trembles of the players' strategies.

Let $\Gamma \in \mathcal{G} \setminus \mathcal{G}^\dagger$ be a two-person perfect-information game with generic payoffs, and denote by G its strategic form. (The sets $\mathcal{G}$ and $\mathcal{G}^\dagger$ are as in **Appendix A** and we further assume that Γ has finitely many equilibrium outcomes, which is also a property of generic extensive-form games.) Because Γ is a perfect-information game, each node x of the game tree of Γ defines a subgame $\Gamma(x)$ that has x as its root. When referring to player n , we use m to denote his opponent. We say that a strategy $\sigma_n \in \Sigma_n$ *enables* (or, alternatively, *does not exclude*) a node x in the game Γ if there exists $\sigma_m \in \Sigma_m$ such that x is reached with positive probability under the profile (σ_n, σ_m) . If a strategy σ_n enables a node x then it induces a continuation strategy in the subgame $\Gamma(x)$ in the obvious way.

Let Σ° be a closed set of equilibria, all of which induce the same probability distribution P° over outcomes. For a node x in the extensive form of Γ , $P^\circ(x)$ denotes the probability of reaching x under the distribution P° . Because Γ is a generic game with perfect information, for each x such that $P^\circ(x) > 0$ and x is a node of one of the players, all equilibria in Σ° enable x and prescribe the same (pure) action there.

PROPOSITION 3. *Suppose a refinement satisfies Invariance and selects only subsets containing proper equilibria. If this refinement selects Σ° in the game G then Σ° is a stable set of G .*

The basic idea of the proof is to obtain the LPS induced by a proper equilibrium of an expanded game, and then truncate the LPS to obtain one that satisfies condition (*) in **Theorem 2**. Unlike the example in **Section 2.3**, here we add two types of redundant strategies to cover all possibilities in the general case. First we establish notation and then, since the proof is long, we break it into a series of claims.

Fix $\tau \in \Sigma \setminus \partial \Sigma$. We show that Σ° satisfies the condition (*) for the profile τ in **Theorem 2**. Denote by b the behavioral strategy profile that is equivalent to τ . Since τ is completely mixed, every node of Γ is reached with positive probability under b .

For each player n , let X_n be the set of nodes where he moves.

- Let $V_n \subset X_n$ comprise those nodes x for which $P^\circ(x) = 0$ and the last node y preceding x such that $P^\circ(y) > 0$ belongs to player m .

Each node $x \in V_n$ is excluded by player m 's equilibrium action at an earlier node $y \in X_m$ that is on the equilibrium path.

Let $W_n = X_n \setminus V_n$ and let W_n^0 be the set of $x \in X_n$ such that $P^\circ(x) > 0$. Obviously, $W_n^0 \subseteq W_n$.

- Let $Q_n \subseteq S_n$ be the set of pure strategies s of player n such that for each $x \in W_n^0$, s plays the unique action associated with Σ° at x .

Let $R_n = S_n \setminus Q_n$.

The set Q_n is nonempty, while R_n is nonempty if and only if there exists a node in W_n^0 where player n has at least two actions available.

The following claim follows readily from our definitions.

CLAIM 1. *Suppose $x \notin W_n^0$.*

- (1) *If x belongs to V_n then every node y that belongs to X_m (resp. X_n) and that precedes or succeeds x belongs to W_m (resp. $V_n \cup W_n^0$). Moreover, x is enabled only by strategies in R_m and it is enabled by some strategy in Q_n .*
- (2) *If x belongs to W_n then any node y in X_m (resp. X_n) and that precedes or succeeds x belongs to $V_m \cup W_m^0$ (resp. W_n). Moreover, x is enabled by some strategy in Q_m and is enabled only by strategies in R_n .*

The support of every equilibrium in Σ° is contained in $Q_m \times Q_n$. Also, for each n , every strategy in Q_n is a best reply against every equilibrium in Σ° , since the strategies in Q_n differ only at nodes that are in V_n (which are excluded by the equilibria in Σ° , by point (1) of Claim 1) or in $W_n \setminus W_n^0$ (which are excluded by all strategies in Q_n by point (2) of Claim 1).

For each player n , fix some pure strategy $\bar{s}_n \in Q_n$.

- Let $\bar{q}_n$ be a mixed strategy equivalent to the following behavioral strategy: at each node $x \in V_n$ player n plays according to b_n , and at each node $x \in W_n$ player n plays according to $\bar{s}_n$. Obviously, $\bar{q}_n$ is a mixture over strategies in Q_n .¹¹
- In case R_n is nonempty, let $\bar{r}_n$ be a mixed strategy equivalent to the following behavioral strategy: at each node $x \in V_n$ player n plays according to $\bar{s}_n$, and at each node $x \in W_n$ player n plays according to b_n . Then $\bar{r}_n$ is a mixture over strategies in R_n and Q_n that assigns a positive probability to some strategy in R_n .

The following claim also follows readily from Claim 1 and from our definitions.

CLAIM 2. *The mixed strategy $\bar{r}_n$ enables each node in V_m and W_n , and $\bar{q}_n$ enables each node in W_m and V_n .*

For each $\delta \in [0, 1]$ and $s_n \in Q_n$, define:

¹¹If one picks a different $\bar{s}_n$ to define $\bar{q}_n$ then one obtains an equivalent mixed strategy, since all strategies in Q_n agree at nodes in W_n^0 and, by Claim 1, exclude those in $W_n \setminus W_n^0$.

- $q_n(s_n, \delta)$ is the mixed strategy that plays s_n with probability $1 - \delta$ and $\bar{q}_n$ with probability δ . The mixed strategy $q_n(s_n, \delta)$, like $\bar{q}_n$, is a mixture over strategies in Q_n .

Let $Q_n(\delta)$ be the collection of these mixed strategies.

- In case R_n is nonempty, $r_n(s_n, \delta)$ is the mixed strategy that plays s_n with probability $1 - \delta$ and $\bar{r}_n$ with probability δ . Like $\bar{r}_n$, $r_n(s_n, \delta)$ assigns a positive probability to some strategy in R_n . And since $s_n \in Q_n$, the conditional distribution over R_n that is induced by $r_n(s_n, \delta)$ is the same as that induced by $\bar{r}_n$.

Let $R_n(\delta)$ be the collection of these mixed strategies.

In sum,

$$q_n(s_n, \delta) = [1 - \delta]s_n + \delta\bar{q}_n \quad \text{and} \quad r_n(s_n, \delta) = [1 - \delta]s_n + \delta\bar{r}_n$$

where the contingency table for the choices by $\bar{q}_n$ and $\bar{r}_n$ at nodes $x \in X_n$ is:

	$x \in V_n$	$x \in W_n$
$\bar{q}_n[x] =$	$b_n[x]$	$\bar{s}_n[x]$
$\bar{r}_n[x] =$	$\bar{s}_n[x]$	$b_n[x]$

Let $G(\delta)$ be the equivalent game obtained by adding for each player n all strategies in $Q_n(\delta) \cup R_n(\delta)$ as redundant pure strategies. Let $S(\delta)$ and $\Sigma(\delta)$ be the spaces of pure and mixed strategy profiles for the game $G(\delta)$, and let $I_n = |S_n(\delta)|$. If R_n is nonempty, $I_n = |S_n| + 2 \times |Q_n|$; otherwise, $I_n = 2 \times |S_n|$.

By assumption, for each δ there exists a proper equilibrium $\bar{\sigma}^1(\delta)$ of $G(\delta)$ whose equivalent profile in Σ , call it $\sigma^1(\delta)$, is in Σ° . By Blume et al. (1991b, Proposition 5) there exists for each player n an LPS $\mathcal{K}_n(\delta) \equiv (\bar{\sigma}_n^1(\delta), \dots, \bar{\sigma}_n^{I_n(\delta)}(\delta))$ over $S(\delta)$ with full support, where $I_n(\delta) \leq I_n$, such that the LPS profile $(\mathcal{K}_1(\delta), \mathcal{K}_2(\delta))$ respects preferences; viz., if against $\mathcal{K}_m(\delta)$ a pure strategy s_n for player n in $G(\delta)$ is a better reply than another t_n , then in $\mathcal{K}_n(\delta)$, s_n is infinitely more likely than t_n .¹² (In particular, $\bar{\sigma}_n^1(\delta)$ is a best reply against $\mathcal{K}_m(\delta)$.) Since I_n is independent of δ , by replacing $\mathcal{K}_n(\delta)$ with the LPS $(\bar{\sigma}_n^1(\delta), \bar{\sigma}_n^1(\delta), \dots, \bar{\sigma}_n^1(\delta), \bar{\sigma}_n^2(\delta), \dots, \bar{\sigma}_n^{I_n(\delta)}(\delta))$ where we insert $I_n - I_n(\delta)$ copies of $\bar{\sigma}_n^1(\delta)$ before $\bar{\sigma}_n^2(\delta)$, we can assume that $\mathcal{K}_n(\delta)$ has I_n levels, i.e. $I_n(\delta) = I_n$. Let $\mathcal{L}_n(\delta) = (\sigma_n^1(\delta), \dots, \sigma_n^{I_n}(\delta))$ be the corresponding LPS over S_n induced by $\mathcal{K}_n(\delta)$; viz., $\sigma_n^i(\delta)$ is the strategy in Σ_n equivalent to $\bar{\sigma}_n^i(\delta)$. Obviously for $\sigma_n, \sigma'_n \in \Sigma_n$, σ_n is a better reply against $\mathcal{K}_m(\delta)$ than σ'_n if and only if it is a better reply against $\mathcal{L}_m(\delta)$ than σ'_n . In particular, $\sigma_n^1(\delta)$ is a best reply against $\mathcal{K}_m(\delta)$ and hence also against $\mathcal{L}_m(\delta)$.

CLAIM 3. In $\mathcal{L}_n(\delta)$ every strategy $q_n \in Q_n$ is infinitely more likely than every strategy $r_n \in R_n$.

PROOF OF CLAIM. Fix a strategy $q_n \in Q_n$. Because $\sigma_m^1(\delta)$ belongs to Σ° , q_n is a best reply to $\sigma_m^1(\delta)$. Let i be the first level of $\mathcal{L}_n(\delta)$ that assigns q_n a positive probability.

¹²Recall from Section 4 that by a better reply we mean a lexicographically *strict* better reply.

The first level i' of $\mathcal{K}_n(\delta)$ that assigns a positive probability to q_n is such that $i' \geq i$. (It is possible that $i' > i$, since a duplicate strategy $q_n(s_n, \delta)$ that has q_n in its support might be assigned a positive probability by level i of $\mathcal{K}_n(\delta)$, thus accounting for the presence of q_n in the support of $\sigma_n^i(\delta)$.) Then, since $\mathcal{K}_n(\delta)$ respects preferences, every strategy in the support of $\sigma_n^j(\delta)$ (and hence of $\sigma_n^i(\delta)$) for $j \leq i$ is a best reply to $\sigma_m^1(\delta)$. As remarked above, $\sigma_m^1(\delta)$ is a best reply against $\mathcal{L}_n(\delta)$. Therefore, for sufficiently small $\varepsilon > 0$ the strategy profile $\sigma(\delta, \varepsilon) = (\sigma_m^1(\delta), \sigma_n(\delta, \varepsilon))$ is an equilibrium of G , where $\sigma_n(\delta, \varepsilon) = (\varepsilon + \dots + \varepsilon^i)^{-1} \sum_{j=1}^i \varepsilon^j \sigma_n^j(\delta)$. Because Γ is a generic extensive-form game, the outcome induced by $\sigma(\delta, \varepsilon)$ is the same for all ε , viz., they all induce P° . Therefore, all strategies in the support of $\sigma_n^j(\delta)$ for $j \leq i$ choose the unique equilibrium action at each node $x \in W_n^0$. Thus these strategies are all in Q_n , and those in $R_n = S_n \setminus Q_n$ are only in the supports of strategies in levels $j > i$ of the LPS. $\square$

Now choose a sequence $\delta(k)$ converging to zero as k increases and let $(\mathcal{K}_1(\delta(k)), \mathcal{K}_2(\delta(k)))$ and $(\mathcal{L}_1(\delta(k)), \mathcal{L}_2(\delta(k)))$ be the corresponding sequences of LPS profiles, where the former profile respects preferences and $(\sigma_1^1(\delta(k)), \sigma_2^1(\delta(k))) \in \Sigma^\circ$. Passing to a subsequence of the $\delta(k)$'s if necessary, [Blume et al. \(1991b, Proposition 2\)](#) implies the following. There exists for each n and each $1 \leq i \leq I_n$, a positive integer J_n^i , an LPS $\mathcal{L}_n^i \equiv (\sigma_n^{i1}, \dots, \sigma_n^{iJ_n^i})$ over S_n , and a sequence $(\mu_n^1(k), \dots, \mu_n^{J_n^i-1}(k)) \in (0, 1)^{J_n^i-1}$ converging to zero such that, for each k , $\sigma_n^i(\delta(k))$ is expressible as the nested combination

$$\sigma_n^i(\delta(k)) = (1 - \mu_n^1(k))\sigma_n^{i1} + \mu_n^1(k)[(1 - \mu_n^2(k))\sigma_n^{i2} + \mu_n^2(k)[\dots + \mu_n^{J_n^i-1}(k)\sigma_n^{iJ_n^i}]\dots].$$

Thus we obtain an LPS $\mathcal{L}_n = (\sigma_n^{ij})$, $1 \leq i \leq I_n$, $1 \leq j \leq J_n^i$, where the ordering is the following: $ij < i'j'$ (that is, level ij is infinitely more likely than level $i'j'$) if $i < i'$ or $i = i'$ and $j < j'$.

The strategy profile where each player n plays σ_n^{11} , the first level of his LPS, is in Σ° as it is the limit of the sequence of profiles where each player n plays $\sigma_n^1(\delta(k))$. Because each level i of $\mathcal{L}_n(\delta)$ is a nested combination of the LPS $\mathcal{L}_n^i$, a strategy σ_n is a better reply against $\mathcal{L}_m$ than another strategy σ'_n if and only if for all large k it is a better reply against $\mathcal{L}_m(\delta(k))$ than σ'_n is. In particular, strategies in the support of σ_n^{11} , which are obviously in the support of $\sigma_n^1(\delta(k))$ for all k , are best replies against $\mathcal{L}_m$. Also, the nested property for $\mathcal{L}_n$ along with [Claim 3](#) immediately implies:

CLAIM 4. *In $\mathcal{L}_n$ every strategy in Q_n is infinitely more likely than every strategy in R_n .*

In case R_n is nonempty, let $i_n^1 j_n^1$ be the first level of $\mathcal{L}_n$ that assigns a positive probability to some strategy in R_n and let $\sigma_n^{i_n^1 j_n^1}(R_n)$ be the conditional distribution over R_n under $\sigma_n^{i_n^1 j_n^1}$. For each k , since each level i of $\mathcal{L}_n(\delta(k))$ is a nested combination of levels of $\mathcal{L}_n^i$, i_n^1 is the first level of $\mathcal{L}_n(\delta(k))$ that assigns positive probability to a strategy in R_n . Moreover, letting $\sigma_n^{i_n^1}(\delta(k), R_n)$ be the conditional distribution over R_n induced by $\sigma_n^{i_n^1}$, we have that $\sigma_n^{i_n^1 j_n^1}(R_n)$ is limit of $\sigma_n^{i_n^1}(\delta(k), R_n)$. We will now compute $\sigma_n^{i_n^1 j_n^1}(R_n)$ using the sequence $\mathcal{K}_n(\delta(k))$.

Clearly, i_n^1 is the first level of $\mathcal{X}_n(\delta(k))$ that assigns positive probability to a strategy in either R_n or $R_n(\delta)$. (Recall that each strategy in $R_n(\delta)$ assigns a positive probability to some strategy in R_n .) Let α_{nk}^1 (resp. β_{nk}^1) be the sum of the probabilities of the strategies in $R_n(\delta(k))$ (resp. R_n) under $\tilde{\sigma}_n^{i_n^1}(\delta(k))$. Then $\alpha^1(\delta(k)) + \beta^1(\delta(k)) > 0$ since otherwise the strategies in R_n would be assigned zero probability by $\sigma_n^{i_n^1}(\delta(k))$. Observe that the probability $\sigma_{n,s}^{i_n^1}(\delta(k))$ of a strategy $s \in R_n$ under $\sigma_n^{i_n^1}(\delta(k))$ equals $\tilde{\sigma}_{n,s}^{i_n^1}(\delta(k)) + \delta(k)\alpha_{nk}^1 \bar{r}_{n,s}$. Define a strategy $\sigma_{nk}^*(R_n) \in \Sigma_n$ as follows. If $\beta_{nk}^1 > 0$, then let $\sigma_{nk}^*(R_n)$ be the mixed strategy in Σ_n that is obtained by taking the conditional distribution over R_n that is induced by $\tilde{\sigma}_n^{i_n^1}(\delta(k))$: that is, the probabilities of strategies in Q_n under $\sigma_{nk}^*(R_n)$ are zero, while that of a strategy $r_n \in R_n$ is $(\beta_{nk}^1)^{-1} \tilde{\sigma}_{n,r_n}^{i_n^1}(\delta(k))$. If $\beta_{nk}^1 = 0$, let $\sigma_{nk}^*(R_n)$ be an arbitrary mixed strategy in Σ_n whose support is contained in R_n . Let R_{nk}^* be the support of $\sigma_{nk}^*(R_n)$. By construction, R_{nk}^* is contained in R_n , and in case $\beta_{nk}^1 \neq 0$ it equals the set of strategies that are assigned a positive probability by level i_n^1 of $\mathcal{X}_n(\delta(k))$. Let $\mu_{nk}^* = \delta(k)\alpha_{nk}^1 / [\delta(k)\alpha_{nk}^1 + \beta_{nk}^1]$. We now have that $\sigma_n^{i_n^1}(\delta(k), R_n)$ is given by the conditional distribution over R_n that is induced by $[1 - \mu_{nk}^*]\sigma_{nk}^*(R_n) + \mu_{nk}^* \bar{r}_n$.

By going to an appropriate subsequence of $\delta(k)$, μ_{nk}^* converges to some μ_n^* and $\sigma_{nk}^*(R_n)$ converges to some strategy $\sigma_n^*(R_n)$ with support, say, R_n^* . Because the support of $\sigma_n^*(R_n)$ and $\bar{r}_n$ contain strategies in R_n , $\sigma_n^{i_n^1 j_n^1}(R_n)$ is now the conditional distribution over R_n that is induced by $(1 - \mu_n^*)\sigma_n^*(R_n) + \mu_n^* \bar{r}_n$.

Say that a pure strategy s_n of player n is *conditionally optimal* against $\mathcal{L}_m$ if at every node $x \in X_n$ that s_n does not exclude the action prescribed by s_n at x is optimal against the first level of $\mathcal{L}_m$ that does not exclude x .

CLAIM 5. Suppose R_n is nonempty. (1) If $\mu_n^* \neq 1$ then every strategy in R_n^* is conditionally optimal against $\mathcal{L}_m$. (2) $\mu_n^* \neq 0$.

PROOF OF CLAIM. Suppose $\mu_n^* \neq 1$. Fix a pure strategy $r_n \in R_n$ that is not conditionally optimal against $\mathcal{L}_m$. We show that $r_n \notin R_n^*$. The set R_n^* is the support of $\sigma_n^*(R_n)$, which is defined as the limit of $\sigma_{nk}^*(R_n)$. Therefore, it is sufficient to show that r_n does not belong to the support R_{nk}^* of $\sigma_{nk}^*(R_n)$ for all large k . Since $\mu_n^* \neq 1$, $\beta_{nk}^1 \neq 0$ for all large k . As remarked above, this implies that for each such k , R_{nk}^* is the set of strategies that are assigned a positive probability by level i_n^1 of $\mathcal{X}_n(\delta(k))$. To show that for large k , $r_n \notin R_{nk}^*$, i.e. that it is assigned probability zero by $\tilde{\sigma}_n^{i_n^1}(\delta(k))$, it is sufficient to show that for large enough k there is a strategy $r'_n \in R_n \cup R_n(\delta(k))$ that is a better reply against $\mathcal{L}_m(\delta(k))$. Indeed, the result then follows from the fact that $\mathcal{X}_n(\delta(k))$ respects preferences and that i_n^1 is the first level of $\mathcal{X}_n(\delta(k))$ that assigns positive probability to some strategy in $R_n \cup R_n(\delta(k))$.

Let x be a node of player n that r_n does not exclude and where it is not conditionally optimal. Consider first the case $P^\circ(x) > 0$. Then σ_m^{11} is the first level of $\mathcal{L}_m$ that does not exclude x . It cannot be the case that r_n chooses the equilibrium action at x since that choice is optimal. Thus r_n chooses a non-equilibrium action at x that is not optimal against σ_m^{11} at x . Pick $s_n \in Q_n$ that is in the support of σ_n^{11} . The strategy s_n is a best reply

against $\mathcal{L}_m$, while r_n is not even a best reply against σ_m^{11} . Therefore, there exists $\delta_0 > 0$ such that $r_n(s_n, \delta_0)$ is a better reply against $\mathcal{L}_m$ than r_n . Hence, there exists k_0 such that $r_n(s_n, \delta_0)$ is a better reply against $\mathcal{L}_m(\delta(k))$ for $k \geq k_0$. Because s_n is a best reply against $\mathcal{L}_m$, there exists $k_1 \geq k_0$ such that s_n is a best reply against $\mathcal{L}_m(\delta(k))$ for $k \geq k_1$. Consequently, for each $\delta < \delta_0$, and $k \geq k_1$, $r_n(s_n, \delta)$, which is a convex combination of s_n and $r_n(s_n, \delta_0)$, is a weakly better reply against $\mathcal{L}_m(\delta(k))$ than $r_n(s_n, \delta_0)$. In particular, for k such that $k \geq k_1$ and $\delta(k) < \delta_0$, $r_n(s_n, \delta(k))$ is a weakly better reply against $\mathcal{L}_m(\delta(k))$ than $r_n(s_n, \delta_0)$ and therefore a better reply than r_n . Thus $r_n \notin R_{nk}^*$ for large k .

Now suppose $P^\circ(x) = 0$. Let r'_n be a strategy that differs from r_n only in that it prescribes an optimal continuation at x . Obviously r'_n is in R_n as well. Moreover, against $\mathcal{L}_m$ it is now a better reply than r_n . Hence for all large k , r'_n is a better reply against $\mathcal{L}_m(\delta(k))$ than r_n . Again r_n is not in R_{nk}^* for large k .

To finish the proof of the claim, it remains to show that $\mu_n^* > 0$. Suppose to the contrary that $\mu_n^* = 0$. Then $\sigma_n^{i_n^1 j_n^1}(R_n) = \sigma_n^*(R_n)$ and thus by (1) all strategies that belong to R_n and also to the support of $\sigma_n^{i_n^1 j_n^1}$ are optimal against σ_m^{11} . Any other strategy that appears in the support of a level ij that equals or precedes $i_n^1 j_n^1$ in player n 's LPS belongs to Q_n and is also optimal against σ_m^{11} . Also, σ_m^{11} is a best reply to $\mathcal{L}_n$. Hence, for all small ε , the strategy profile $(\sigma_m^{11}, \sigma_n(\varepsilon))$ is an equilibrium of G , where

$$\sigma_n(\varepsilon) = \left(\sum_{ij \leq i_n^1 j_n^1} \varepsilon^{d(ij)} \right)^{-1} \sum_{ij \leq i_n^1 j_n^1} \varepsilon^{d(ij)} \sigma_n^{ij},$$

with $d(ij) = \sum_{i'=1}^{i-1} J_n^{i'} + j$. The strategy $\sigma_n(\varepsilon)$ assigns positive probability to strategies in R_n , which by definition choose a non-equilibrium action at some node on the equilibrium path that they do not exclude. Hence for $\varepsilon > 0$, $(\sigma_m^{11}, \sigma_n(\varepsilon))$ induces an outcome different from P° , which is impossible. Therefore, $\mu_n^* > 0$. $\square$

CLAIM 6. For $x \in V_n$, the first level of $\mathcal{L}_m(\delta)$ that enables it is $i_m^1 j_m^1$. Moreover, the continuation strategy in $\Gamma(x)$ induced by $\sigma_m^{i_m^1 j_m^1}$ is the same as that induced by $(1 - \mu_m^*)\sigma_m^*(R_m^*) + \mu_m^* \bar{r}_m$.

PROOF OF CLAIM. By **Claim 1**, any strategy of m that enables x belongs to R_m —in particular, R_m is nonempty and $i_m^1 j_m^1$ is well defined. By the definition of $i_m^1 j_m^1$, therefore, x is not enabled by $ij < i_m^1 j_m^1$. Again because x is enabled only by strategies in R_m , and also because $\sigma_m^{i_m^1 j_m^1}(R_m)$ is the conditional distribution induced by $(1 - \mu_m^*)\sigma_m^*(R_m^*) + \mu_m^* \bar{r}_m$, the proof is complete if we show that x is enabled by $(1 - \mu_m^*)\sigma_m^*(R_m^*) + \mu_m^* \bar{r}_m$. This last point follows from **Claim 2** and point (2) of **Claim 5**. $\square$

Let Q_n^+ be the subset of strategies in Q_n that are conditionally optimal against $\mathcal{L}_m$. Let $i_n^0 j_n^0$ be the first level ij with the following property: if $Q_n = Q_n^+$, then every strategy in Q_n is assigned a positive probability by some level $i'j' \leq ij$; and if $Q_n \neq Q_n^+$, then level ij assigns a positive probability to some $s \in Q_n \setminus Q_n^+$. When R_n is nonempty, $i_n^0 j_n^0 < i_n^1 j_n^1$ by **Claim 4**. Moreover, if $Q_n \neq Q_n^+$, $i_n^0 j_n^0 > 11$: σ_n^{11} assigns probability only to best replies against $\mathcal{L}_m$, and any strategy in $Q_n \setminus Q_n^+$ is obviously not a best reply against $\mathcal{L}_m$.

CLAIM 7. Suppose $Q_n \neq Q_n^+$. Then $\bar{q}_n$ assigns positive probability to some strategy in $Q_n \setminus Q_n^+$.

PROOF OF CLAIM. Let $s_n \in Q_n \setminus Q_n^+$. Let $x \in X_n$ be a node that s_n enables but where it is not conditionally optimal. Obviously $P^c(x) = 0$, since s_n prescribes only the equilibrium action at nodes on the equilibrium path. By Claim 1, we therefore have $x \in V_n$. The result now follows from the fact that $\bar{q}_n$ enables each node in V_n (Claim 2) and is equivalent to a behavioral strategy that mixes over all the actions at each such node. $\square$

Suppose $Q_n \neq Q_n^+$. Let Q_n^0 be the set of $q_n \in Q_n$ assigned zero probability by levels $i_j < i_n^{j_0}$. By definition there exists q_n that belongs to $Q_n^0 \setminus Q_n^+$ and that is assigned positive probability by level $i_n^{j_0}$. Therefore, $\sigma_n^{i_n^{j_0}}$ induces a well-defined conditional distribution over Q_n^0 , call it $\sigma_n^{i_n^{j_0}}(Q_n^0)$. To compute this conditional distribution, we mimic what we did for R_n above. (The only difference between what we present in the next paragraph and their counterparts from before is that we use Q_n^0 instead of R_n and $\bar{q}_n$ instead of $\bar{r}_n$.)

For each k , level i_n^0 is the first level of $\mathcal{L}_n(\delta(k))$ that assigns a positive probability to some strategy in Q_n^0 . Let $\sigma_n^{i_n^0}(\delta(k), Q_n^0)$ be the conditional distribution over Q_n^0 that is induced by $\sigma_n^{i_n^0}(\delta(k))$. Using the above claim, i_n^0 is also the first level of $\mathcal{K}_n(\delta(k))$ to assign a positive probability to a strategy in Q_n^0 or in $Q_n(\delta(k))$. As with R_n , there exist $\lambda_{nk}^* \in [0, 1]$ and a mixed strategy $\sigma_{nk}^*(Q_n^0)$ with support, say, $Q_{nk}^* \subseteq Q_n^0$ such that $\sigma_n^{i_n^0}(\delta(k), Q_n^0)$ is the conditional distribution over Q_n^0 induced by $(1 - \lambda_{nk}^*)\sigma_{nk}^*(Q_n^0) + \lambda_{nk}^*\bar{q}_n$. If $\lambda_{nk}^* \neq 1$ then Q_{nk}^* consists of strategies $s_n \in Q_n^0$ assigned a positive probability by $\bar{\sigma}_n^{i_n^{j_1}}$. Going to the limit produces the strategy $[1 - \lambda_n^*]\sigma_n^*(Q_n^0) + \lambda_n^*\bar{q}_n$, whose conditional distribution over Q_n^0 is $\sigma_n^{i_n^{j_0}}(Q_n^0)$. Denote by Q_n^* the support of $\sigma_n^*(Q_n^0)$.

We want to establish an analogue of the optimality property of R_n^* in Claim 5 for Q_n^* . Before doing so, we need a preliminary claim.

CLAIM 8. All strategies in Q_n are equally good replies against any strategy for player m whose support is contained in Q_m .

PROOF OF CLAIM. The result follows from the following three observations. All strategies in Q_n agree at each $x \in W_n^0$; they exclude nodes in $W_n \setminus W_n^0$, by point 2 of Claim 1; finally a node $x \in V_n$ is excluded by every strategy in Q_m by point 1 of Claim 1. $\square$

CLAIM 9. Suppose $Q_n \neq Q_n^+$. (1) If $\lambda_n^* \neq 1$, each strategy in Q_n^* is conditionally optimal against $\mathcal{L}_m$, i.e. it belongs to Q_n^+ . (2) $\lambda_n^* \neq 0$.

PROOF OF CLAIM. Suppose $q_n \in Q_n \setminus Q_n^+$. We show that $q_n \notin Q_n^*$. As in the proof of statement (1) of Claim 5, it sufficient to show that $q_n \notin Q_{nk}^*$, i.e. it is assigned zero probability by $\bar{\sigma}_n^{i_n^0}(\delta(k))$, for large k . We claim that we are done if we can show that there exists $s_n \in Q_n$ such that $q_n(s_n, \delta(k))$ is a better reply against $\mathcal{K}_m(\delta(k))$ for large k . Indeed, if

there does exist such an s_n , then letting $i(k)$ and $i(q_n, k)$ be, resp., the first levels where $q_n(s_n, \delta(k))$ and q_n are assigned positive probability by $\mathcal{K}_n(\delta(k))$, we have $i(k) < i(q_n, k)$ for large k , because $\mathcal{K}_n(\delta(k))$ respects preferences. By **Claim 7**, level $i(k)$ of $\mathcal{L}_n(\delta(k))$ assigns a positive probability to some $q'_n \in Q_n \setminus Q_n^+$. Hence, $i_n^0 \leq i(k) < i(q_n, \delta(k))$ and q_n is assigned zero probability by $\bar{\sigma}_n^{i_n^0}$. Thus, it is sufficient to show the existence of $s_n \in Q_n$ such that $q_n(s_n, \delta(k))$ is a better reply against $\mathcal{K}_m(\delta(k))$ for large k .

Since q_n is not conditionally optimal against $\mathcal{L}_m$, there exists $x \in X_n$ that q_n enables and where it is not optimal against the first level of $\mathcal{L}_m$ that enables it. Clearly, $x \notin W_n^0$, since q_n chooses the equilibrium action at each node on the equilibrium path. As q_n excludes each node in $W_n \setminus W_n^0$, $x \in V_n$ and by **Claim 6**, x is enabled by level $i_m^1 j_m^1$. Let q'_n be a strategy that agrees with q_n everywhere except that in the subgame at x , it prescribes a continuation strategy that is optimal against $\sigma_m^{i_m^1 j_m^1}$. Obviously, q'_n belongs to Q_n and is a better reply to level $i_m^1 j_m^1$ of $\mathcal{L}_m$ than q_n . Pick s_n in the support of $\sigma_n^{i_n^1}$. Because $s_n \in Q_n$, s_n and q'_n are equally good replies against every level $ij < i_m^1 j_m^1$ of $\mathcal{L}_m$, by **Claim 8**. Because s_n is a best reply to $\mathcal{L}_m$, it is at least as good a reply as q'_n against level $i_m^1 j_m^1$. Hence it is a better reply to $\sigma_m^{i_m^1 j_m^1}$ than q_n . There exists $\delta_0 > 0$ such that $q_n(s_n, \delta_0)$ is a better reply to $\sigma_m^{i_m^1 j_m^1}$ than q_n . Since the support of $q_n(s_n, \delta_0)$ is contained in Q_n , again using **Claim 8**, $q_n(s_n, \delta_0)$ is a better reply to $\mathcal{L}_m(\delta)$ than q_n . Hence there exists k_0 such that for all $k \geq k_0$, $q_n(s_n, \delta_0)$ is a better reply to $\mathcal{L}_m(\delta(k))$ than q_n . Because σ_n is a best reply to $\mathcal{L}_m$, there exists $k_1 > k_0$ such that s_n is a weakly better reply to $\mathcal{L}_m(\delta(k))$ than $q_n(s_n, \delta_0)$ for all $k \geq k_1$. Hence for all k such that $\delta(k) \leq \delta_0$ and $k \geq k_1$, $q_n(s_n, \delta(k))$ is a better reply to $\mathcal{L}_k(\delta(k))$ than q_n , which completes the proof of (1).

We now prove point (2). Suppose $\lambda_n^* \neq 1$. Then by point (1) every strategy in the support of $\sigma_n^*(Q_n^0)$ is in Q_n^+ . Since $\sigma_n^{i_n^0 j_n^0}$ assigns positive probability to some strategy in $Q_n^0 \setminus Q_n^+$, it must be that $\lambda_n^* \neq 0$. $\square$

CLAIM 10. *Let $x \in W_n$. Then x is enabled by σ_m^{ij} for some $ij \leq i_m^0 j_m^0$. Moreover, if $Q_m \neq Q_m^+$ and x is not enabled by level $ij < i_m^0 j_m^0$, then the continuation strategy that $\sigma_m^{i_m^1 j_m^1}$ induces in $\Gamma(x)$ is the same as that induced by $(1 - \lambda_m^*)\sigma_m^*(Q_m^0) + \lambda_m^* \bar{q}_m$.*

PROOF OF CLAIM. By **Claim 1**, there is a strategy s_m in Q_m that enables x . Suppose $Q_m = Q_m^+$. By the definition of $i_m^0 j_m^0$, s_m is assigned positive probability by some level $ij \leq i_m^0 j_m^0$ of $\mathcal{L}_m$, which then enables x .

Suppose now that $Q_m \neq Q_m^+$ and x is not enabled by any level $ij < i_m^0 j_m^0$. Then by the definition of Q_m^0 , the only strategies in Q_m that enable x belong to Q_m^0 . The conditional distribution $\sigma_m^{i_m^0 j_m^0}(Q_m^0)$ over Q_m^0 induced by level $i_m^0 j_m^0$ is that induced by $\lambda_m^* \bar{q}_m + (1 - \lambda_m^*)\sigma_m^*(Q_m^0)$. Thus, to complete the proof it is sufficient to show that $\lambda_m^* \bar{q}_m + (1 - \lambda_m^*)\sigma_m^*(Q_m^0)$ enables x . This last follows from point (2) of **Claim 9** and **Claim 2**. $\square$

CLAIM 11. *For each n there exist $v_n^* \in (0, 1]$ and $\sigma_n^* \in \Sigma_n$ such that we have the following properties, where $\bar{\sigma}_n = (1 - v_n^*)\sigma_n^* + v_n^*\tau_n$.*

- (1) *If $v_n \neq 1$ then σ_n^* is conditionally optimal against $\mathcal{L}_m$.*
- (2) *For each $x \in V_m$, $\bar{\sigma}_n$ enables x and prescribes the same continuation strategy in $\Gamma(x)$ as level $i_n^1 j_n^1$ of $\mathcal{L}_n$.*
- (3) *If $Q_n \neq Q_n^+$ then for each $x \in W_m$ that is not enabled by $ij < i_n^0 j_n^0$, $\bar{\sigma}_n$ enables x and prescribes the same continuation strategy in $\Gamma(x)$ as level $i_n^0 j_n^0$.*

PROOF OF CLAIM. Suppose first that $Q_n = Q_n^+$ and R_n is empty, then let $\bar{\sigma}_n = \tau_n$. (We have $v_n^* = 1$ and the choice of σ_n^* is irrelevant.) Points (1) and (3) of the claim hold trivially. As for point (2), if R_n is empty, V_m is empty, by Claim 1, and hence it too holds trivially.

Suppose now that $Q_n = Q_n^+$ and R_n is nonempty. Then, let $\sigma_n^* = \sigma_n^*(R_n)$ and $v_n^* = \mu_n^*$. Point (1) of this claim follows from point (1) of Claim 5 and point (3) holds trivially. As for point (2) of this claim, remark first that by Claim 6, the continuation strategy for player n $\Gamma(x)$ that is prescribed by $\sigma_n^{i_n^1 j_n^1}$ is that given by $(1 - \mu_n^*)\sigma_n^*(R_n) + \mu_n^*\bar{r}_n$. Since $x \in V_m$, by Claim 1, every node of n that precedes or succeeds x belongs to W_n . The strategies $\bar{r}_n$ and τ_n are equivalent to behavioral strategies that choose the same mixture at each node in W_n . Therefore, $\bar{\sigma}_n$ prescribes the continuation strategy in the subgame at x as $(1 - \mu_n^*)\sigma_n^*(R_n) + \mu_n^*\bar{r}_n$.

Finally, suppose now that $Q_n \neq Q_n^+$. Let V_n^0 be the set of initial vertices of V_n , viz., V_n^0 is the set of nodes in V_n that are not preceded by any other vertex in V_n . By the definition of V_n^0 and Claim 1, nodes in X_n that precede a node in V_n^0 belong to W_n^0 . Also, V_n^0 is nonempty. Indeed, if it is empty then V_n is empty; but then $Q_n = Q_n^+$ because all strategies in Q_n prescribe the equilibrium action at each node in W_n^0 and exclude all nodes in $W_n \setminus W_n^0$. Thus, V_n^0 is nonempty. For $v \in V_n^0$, define v_{nv}^* to be the probability that τ_n enables v , i.e. v_{nv}^* is the total probability under τ_n of the set of pure strategies that enable v . For each $v \in V_n^0$, define σ_{nv}^* to be the mixed strategy that is equivalent to the following behavioral strategy: in the subgame $\Gamma(v)$ play according to $\sigma_n^*(Q_n)$; elsewhere play according to σ_n^{11} . The strategy σ_{nv}^* is a mixture over strategies in Q_n and if $\lambda_n^* \neq 1$ then σ_{nv}^* is conditionally optimal against $\mathcal{L}_m$, since strategies $\sigma_n^*(Q_n)$ and σ_n^{11} are.

In case R_n is empty, define $\hat{\sigma}_n$ to be σ_n^{11} and let $\hat{\mu}_n^* = 1$. In case R_n is nonempty, let $\hat{\mu}_n^* = \mu_n^*$ and define $\hat{\sigma}_n^*$ to be the mixed strategy that is equivalent to the following behavioral strategy: at each node $x \in W_n$, play according to $\sigma_n^*(R_n)$; at each node $x \in V_n$, play according to σ_n^{11} . Obviously, when $\hat{\mu}_n^* \neq 1$, $\hat{\sigma}_n$ is also conditionally optimal against $\mathcal{L}_m$.

Now let

$$\bar{\sigma}_n = \left(\lambda_n^* + \sum_{v \in V_n^0} (1 - \lambda_n^*) \hat{\mu}_n^* v_{nv}^* \right)^{-1} \left((1 - \hat{\mu}_n^*) \lambda_n^* \hat{\sigma}_n + \sum_{v \in V_n^0} (1 - \lambda_n^*) \hat{\mu}_n^* v_{nv}^* \sigma_{nv}^* + \lambda_n^* \hat{\mu}_n^* \tau_n \right)$$

and let

$$v_n^* = \frac{\lambda_n^* \hat{\mu}_n^*}{\lambda_n^* + \sum_{v \in V_n^0} (1 - \lambda_n^*) \hat{\mu}_n^* v_{nv}^*}.$$

By point (2) of Claims 5 and 9, $\lambda_n^* \neq 0 \neq \hat{\mu}_n^*$. Thus, $\bar{\sigma}_n$ is a well defined strategy and $0 < v_n^* \leq 1$. If $v_n^* \neq 1$, let

$$\sigma_n^* = \left(\lambda_n^* (1 - \hat{\mu}_n^*) + \sum_{v \in V_n^0} (1 - \lambda_n^*) \hat{\mu}_n^* v_{nv}^* \right)^{-1} \left((1 - \hat{\mu}_n^*) \lambda_n^* \hat{\sigma}_n + \sum_{v \in V_n^0} (1 - \lambda_n^*) \hat{\mu}_n^* v_{nv}^* \sigma_{nv}^* \right)$$

and otherwise let it be an arbitrary strategy. Therefore $\bar{\sigma}_n = (1 - v_n^*) \sigma_n^* + v_n^* \tau_n$. Suppose $v_n^* \neq 1$. Then either $\lambda_n^* \neq 1$ or $\hat{\mu}_n^* \neq 1$. If $\lambda_n^* \neq 1$ then each σ_{nv}^* is optimal against $\mathcal{L}_m$. And if $\hat{\mu}_n^* \neq 1$ then $\hat{\sigma}_n$ is optimal against $\mathcal{L}_m$. Therefore, if $v_n^* \neq 1$ then σ_n^* is optimal against $\mathcal{L}_m$. To prove point (2) we argue as follows. Suppose $x \in V_n$. By Claim 1, x is enabled only by strategies in R_n . (In particular if R_n is nonempty, $\hat{\sigma}_n = \sigma_n^*(R_n)$ and $\hat{\mu}_n^* = \mu_n^*$.) The only “components” of $\bar{\sigma}_n$ that enable x are τ_n and $\hat{\sigma}_n$. Therefore, the continuation strategy in $\Gamma(x)$ under $\bar{\sigma}_n$ is that given $(1 - \hat{\mu}_n^*) \hat{\sigma}_n + \hat{\mu}_n^* \tau_n$. As before, any node of n that precedes or succeeds x belongs to W_m . The strategy $\hat{\sigma}_n$ agrees with $\sigma_n^*(R_n)$ at each node in W_n while τ_n agrees with $\bar{r}_n$ at each such node. Thus the continuation strategy in $\Gamma(x)$ that is induced by $\bar{\sigma}_n$ is the same as that given by $(1 - \mu_n^*) \sigma_n^*(R_n) + \mu_n^* \bar{r}_n$. The result now follows from Claim 6.

It remains to prove point (3). Let $x \in W_m$ be a node such that the first level of $\mathcal{L}_n$ that enables it is $i_n^0 j_n^0$. Since $Q_n \neq Q_n^+$, $i_n^0 j_n^0 > 11$. There must exist $v \in V_n^0$ that precedes it: otherwise, by Claim 1, the only nodes preceding x belong to W_n^0 and x is enabled by σ_n^{11} . By construction, the strategies $\sigma_{nv'}^*$ for $v' \neq v$ and $\hat{\sigma}_n^*(R_n)$ choose the continuation strategy prescribed by σ_n^{11} in the subgame $\Gamma(v)$. Because σ_n^{11} does not enable x , clearly these strategies do not either. Therefore, the only “components” of $\bar{\sigma}_n$ that enable x are σ_{nv}^* and τ_n . Thus, the continuation action prescribed by $\bar{\sigma}_n$ in $\Gamma(x)$ is that prescribed by the conditional distribution $[(1 - \lambda_n^*) \hat{\mu}_n^* v_{nv}^* + \lambda_n^* \hat{\mu}_n^*]^{-1} ((1 - \lambda_n^*) \hat{\mu}_n^* v_{nv}^* \sigma_{nv}^* + \lambda_n^* \hat{\mu}_n^* \tau_n)$, where the factor $\hat{\mu}_n^*$ cancels out. In the subgame $\Gamma(v)$, $\bar{q}_n$ prescribes the same continuation as τ_n while σ_{nv}^* prescribes the same continuation as $\sigma_n^*(Q_n^0)$. The probability of enabling v under $\bar{q}_n$ or $\sigma_n^*(Q_n)$ or σ_{nv}^* is 1, whereas the corresponding probability for τ_n is v_{nv}^* . Hence it is clear that $\bar{\sigma}_n$ prescribes the same continuation strategy in $\Gamma(v)$ —and, therefore, in $\Gamma(x)$, as x succeeds v —as does the strategy $(1 - \lambda_n^*) \sigma_n^*(Q_n) + \lambda_n^* \bar{q}_n$. Point (3) follows by applying Claim 10. $\square$

We construct a new LPS $\bar{\mathcal{L}}_n$ as follows. Suppose first that $Q_n = Q_n^+$. The LPS $\bar{\mathcal{L}}_n$ is obtained by deleting all levels succeeding $i_n^0 j_n^0$ and adding the mixed strategy $\bar{\sigma}_n$ defined above as the last level. If $Q_n \neq Q_n^+$ then $i_n^0 j_n^0 > 11$ so $\bar{\mathcal{L}}_n$ has more than one level. Delete all levels succeeding $i_n^0 j_n^0$ and replace level $i_n^0 j_n^0$ with a mixed strategy $\bar{\sigma}_n$ as defined above.

To finish the proof of the proposition it remains to show that every strategy in the support of σ_n^* (as defined in Claim 11) or of a level of $\bar{\mathcal{L}}_n$ before the last is conditionally optimal against $\bar{\mathcal{L}}_n$. Point (1) of Claim 11 in conjunction with the following claim establishes this optimality property.

CLAIM 12. *A strategy of player n is conditionally optimal against $\overline{\mathcal{L}}_m$ if and only if it is conditionally optimal against $\mathcal{L}_m$.*

PROOF OF CLAIM. We show the following. For each node $x \in X_n$ the first level of $\overline{\mathcal{L}}_m$ that does not exclude x induces the same continuation strategy for m in $\Gamma(x)$ as the first level of $\mathcal{L}_m$ that does not exclude x .

If $x \in W_n^0$ then the first level of $\mathcal{L}_m$ enables it and coincides with the first level of $\overline{\mathcal{L}}_m$. If $x \in V_n$ then by **Claim 6** the first level of $\mathcal{L}_m$ that enables it is $i_m^1 j_m^1$ and the result follows in this case by point (2) of the previous claim. If $x \in W_n \setminus W_n^0$ then, by **Claim 10**, x is enabled by some level $ij \leq i_m^0 j_m^0$ of $\mathcal{L}_m$. If $ij < i_m^0 j_m^0$ or if $Q_m = Q_m^+$ then x is enabled by the same level of $\overline{\mathcal{L}}_m$ as that of $\mathcal{L}_m$. Finally, if $Q_m \neq Q_m^+$ and x is enabled only by level $i_m^0 j_m^0$ of $\mathcal{L}_m$ then the result follows from point (3) of the previous claim. $\square$

C. SIGNALING GAMES

This appendix provides a simple proof of a variant of **Theorem 1** for two-player two-stage signaling games of the kind depicted in **Figure 3** in **Section 2.1**. For these games, the strategic-form and extensive-form perfect equilibria coincide, and if payoffs are generic then they are the same as the sequential equilibria and all equilibria in each component yield the same outcome.

The extensive form is described by three nonempty finite sets (T, M, A) and the scenario in which player 1 (the sender) observes some type $t \in T$ and then sends a message $m \in M$; next, player 2 (the responder) observes only the message m that was sent and based on this observation chooses an action $a \in A$. The set of pure strategies for the sender is $S = M^T$, and for the responder, $R = A^M$. A particular game G is obtained by specifying the players' payoffs and nature's probabilities: let the payoff to player n be $u_n(t, m, a)$ and let $\pi_t > 0$ be the prior probability of type t . A generic signaling game has finitely many Nash equilibrium outcomes. Therefore, in each component of equilibria, the sender has a unique equilibrium strategy and the indeterminacy in equilibria arises only from multiple actions for the receiver following out-of-equilibrium messages from the sender. **Cho and Kreps (1987)** cite and **Banks and Sobel (1987, Theorem 3)** prove the following characterization.¹³

PROPOSITION 4 (Cho and Kreps, Banks and Sobel). *For a generic class of signaling games, an equilibrium component is stable if and only if for each unsent message m and each probability distribution $\theta \in \Delta(T)$ there exists in the component a sequential equilibrium sustained by the sender's belief $\mu \in \Delta(T)^M$ such that $\mu(\cdot|m)$ is in the convex hull of θ and $\Delta(T_m)$, where T_m is the subset of types (if any) indifferent between sending m and using an equilibrium strategy. [That is, $T_m = \{t \mid \sum_{a \in A} u_1(t, m, a)\gamma(a|m) = u_1^\circ(t)\}$ where $u_1^\circ(t)$ is the equilibrium conditional expected payoff of the sender given her type t and the responder's strategy γ .]*

¹³This proposition can be proved also using the characterization result in **Theorem 2**.

Such a sequential equilibrium satisfies the sufficient condition for stability in **Theorem 1** even if it is not induced by a proper equilibrium. The following shows that this anomaly disappears when Invariance is invoked.

PROPOSITION 5. *Fix an equilibrium component Σ° of a generic signaling game. If there exists an unsent message m and a probability distribution $\theta \in \Delta(T)$ for which the condition in **Proposition 4** is violated then there exists an equivalent game, obtained by adding a single mixed strategy of the original game as an additional pure strategy, with no proper equilibrium that is equivalent to some equilibrium in Σ° .*

PROOF. Fix a message m that is an out-of-equilibrium message in the component Σ° . Pick $\theta \in \Delta(T)$. Choose a pure strategy $s \in M^T$ in the support of the sender's equilibrium strategy in Σ° . For each type t let s_{tm} be the pure strategy that agrees with s except that $s_{tm}(t) = m$ —in s_{tm} type t sends the message m while all other types send the message prescribed by s , which is different from m , since m is an out-of-equilibrium message. Let $\hat{\sigma}(\delta)$ be the mixed strategy that assigns probability $1 - \delta$ to s and probability $\delta \eta_{tm}$ to s_{tm} , where $0 < \delta < 1$ and $\eta_{tm} = [\theta_t / \pi_t] / \sum_{t \in T} [\theta_t / \pi_t]$. Thus, δ is the probability of a tremble in favor of m that with probability η_{tm} deviates in the event t from $s(t)$ to m . Observe that conditional on the strategy $\hat{\sigma}(\delta)$ and a tremble in favor of the unsent message m , the posterior probability that the message m came from type t is

$$\mu(t|m, \hat{\sigma}(\delta)) = \pi_t \eta_{tm} / p(m|\hat{\sigma}(\delta)) \quad \text{where} \quad p(m|\hat{\sigma}(\delta)) = \sum_{t \in T} \pi_t \eta_{tm}.$$

Consequently, $\mu(t|m, \hat{\sigma}(\delta)) = \theta_t$. Now append $\hat{\sigma}(\delta)$ as a new pure strategy in the strategic form. This expanded game has the same reduced strategic form as the original game. Assume now that for every small $\delta > 0$, there is a proper equilibrium of the expanded game that induces the same outcome. We show that (m, θ) must satisfy the condition in **Proposition 4**.

Fix δ and let $(\sigma^\delta, \rho^\delta)$ be a proper equilibrium whose reduced form belongs to Σ° . By definition, there exists a sequence $(\sigma^{\delta, \varepsilon}, \rho^{\delta, \varepsilon})$ of ε -proper equilibria converging to $(\sigma^\delta, \rho^\delta)$. Let $\mu^{\delta, \varepsilon}(\cdot|m)$ be the induced sequence of posteriors on the types conditional on m , and let $\gamma^{\delta, \varepsilon} \in \Delta(A)$ be the induced mixed action of the responder after receiving message m . Let μ^δ and γ^δ be the corresponding limits. Let $u_1^\circ(t)$ be the equilibrium payoff to type t of the sender in the component Σ° . Define $f^\delta(t) = \sum_{t' \neq t} \pi_{t'} u_1^\circ(t') + \pi_t u_1(t, m, \gamma^\delta)$, which is the expected payoff to the sender when type t deviates by sending m and the responder replies with the behavioral strategy γ^δ . Let $T^\delta = \arg \max_t f^\delta(t)$, the set of those types with the smallest disincentive to sending m when the responder uses γ^δ .

We claim that μ^δ belongs to the convex hull of θ and $\Delta(T^\delta)$. To see this, consider any pure strategy s' of the original game that sends message m for a subset $T' \not\subseteq T^\delta$ of types. Fix some $t' \in T' \setminus T^\delta$. The strategy s' does no better against ρ^δ than the strategy $s_{t'm}$. (Recall that the strategy $s_{t'm}$ is one in which type t' sends message m while all other types send the message prescribed by strategy s , which is in the support of the sender's equilibrium strategy in Σ° .) For $t^* \in T^\delta$, $f(t^*) > f(t')$, since $t' \notin T^\delta$. Therefore,

the strategy s_{t^*m} does strictly better than $s_{t'm}$ (and s') against ρ^δ and hence against $\rho^{\delta,\varepsilon}$ for all small ε . Consequently, the limit belief μ^δ is not determined by any pure strategy of the original game where a type not in T^δ sends m . In other words, μ^δ is determined by the relative probabilities of $\hat{\sigma}(\delta)$ and those pure strategies for which only types in T^δ send message m . Hence, μ^δ belongs to the convex hull of θ and $\Delta(T^\delta)$ as asserted.

Let $S(\delta)$ be the set of s_{tm} such that $t \in T^\delta$. Observe that by definition all strategies in $S(\delta)$ yield the same payoff against ρ^δ . Now suppose for some δ that the new pure strategy $\hat{\sigma}(\delta)$ does strictly better than the pure strategies in $S(\delta)$ against ρ^δ . Then we claim that the limit belief is θ and the condition in **Proposition 4** holds. By the conclusion from the last paragraph, to prove this claim it is sufficient to show that for any strategy s' where a nonempty subset $\hat{T}^\delta$ of T^δ of types send message m , $\hat{\sigma}(\delta)$ is a better reply against ρ^δ than s' . To prove this last point, consider such a pure strategy s' : the pure strategy s_{tm} , where $t \in \hat{T}^\delta$, is at least as good a reply as s' against ρ^δ ; in turn s_{tm} does strictly worse than $\hat{\sigma}(\delta)$ as it belongs to $S(\delta)$. Thus μ^δ equals θ if $\hat{\sigma}(\delta)$ does strictly better than the pure strategies in $S(\delta)$.

Finally, if for each δ sufficiently small the strategies in $S(\delta)$ do at least as well as $\hat{\sigma}(\delta)$ then their payoffs against ρ^δ must be arbitrarily close to $\sum_t \pi_t u_1^\circ(t)$, since the payoff to $\hat{\sigma}(\delta)$ is at least $(1-\delta)\sum_t u_1^\circ(t) - \delta K$ for a positive constant K . Consider now a sequence of δ 's going to zero such that T^δ is constant, call it T^* , and $(\mu^\delta, \gamma^\delta)$ converges to, say, (μ^*, γ^*) . Then μ^* belongs to the convex hull of θ and $\Delta(T^*)$. Moreover, for each type $t \in T^*$, $u_1(t, m, \gamma^*) = u_1^\circ(t)$. Thus again the condition in **Proposition 4** holds. $\square$

For instance, the result of applying the Intuitive Criterion to the examples in Figures 2 and 3 in **Section 2.1** can be obtained by invoking stability, or simply by observing that the preferred component is the one containing a proper equilibrium.

Cho and Kreps (1987, p. 220) conclude that, “if there is an intuitive story to go with the full strength of stability, it is beyond our powers to offer it here.” **Proposition 5** shows that if one recognizes that “the full strength of stability” entails Invariance, and that a proper equilibrium induces a sequential equilibrium in every extensive-form game with the same strategic form, then the “intuitive story” is less mysterious.

REFERENCES

- Banks, Jeffrey S. and Joel Sobel (1987), “Equilibrium selection in signaling games.” *Econometrica*, 55, 647–661. [171, 180, 182, 202]
- Blume, Lawrence E., Adam Brandenberger, and Eddie Dekel (1991a), “Lexicographic probabilities and choice under uncertainty.” *Econometrica*, 59, 61–79. [181]
- Blume, Lawrence E., Adam Brandenberger, and Eddie Dekel (1991b), “Lexicographic probabilities and equilibrium refinements.” *Econometrica*, 59, 81–98. [181, 183, 185, 189, 194, 195]
- Blume, Lawrence E. and William Zame (1994), “The algebraic geometry of perfect and sequential equilibrium.” *Econometrica*, 62, 783–794. [190]

- Bochnak, Jacek, Michel Coste, and Marie-Françoise Roy (1998), *Real Algebraic Geometry*. Springer-Verlag, Berlin. [186, 187]
- Cho, In-Koo and David M. Kreps (1987), “Signaling games and stable equilibria.” *Quarterly Journal of Economics*, 102, 179–221. [170, 180, 182, 202, 204]
- Govindan, Srihari and Tilman Klumpp (2002), “Perfect equilibrium and lexicographic beliefs.” *International Journal of Game Theory*, 31, 229–243. [184]
- Govindan, Srihari and Robert Wilson (2001), “Direct proofs of generic finiteness of Nash equilibrium outcomes.” *Econometrica*, 69, 765–769. [173]
- Govindan, Srihari and Robert Wilson (2005), “Essential equilibria.” *Proceedings of the National Academy of Sciences USA*, 102, 15706–15711. [173]
- Hillas, John (1996), “How much of forward induction is implied by backward induction and ordinality.” Technical report, University of Auckland, New Zealand. [171]
- Hillas, John and Elon Kohlberg (2002), “Foundations of strategic equilibrium.” In *Handbook of Game Theory* (Robert J. Aumann and Sergiu Hart, eds.), volume 3, chapter 42, 1597–1663, Elsevier, Amsterdam. [169]
- Kohlberg, Elon and Jean-François Mertens (1986), “On the strategic stability of equilibria.” *Econometrica*, 54, 1003–1038. [167, 168, 169, 171, 172, 188]
- Kreps, David and Robert Wilson (1982), “Sequential equilibria.” *Econometrica*, 50, 863–894. [168, 170, 173, 191]
- Kuhn, Harold W. (1953), “Extensive games and the problem of information.” In *Contributions to the Theory of Games* (Harold W. Kuhn and Albert W. Tucker, eds.), volume II, 193–216, Princeton University Press, Princeton, New Jersey. Reprinted in Harold W. Kuhn (ed.), *Classics in Game Theory*, Princeton University Press, Princeton, New Jersey, 1997. [178]
- Lojasiewicz, Stanislaw (1993), “Sur la géométrie semi- et sous-analytique.” *Annales de l’institut Fourier*, 43, 1575–1595. [187]
- Mertens, Jean-François (1989), “Stable equilibria—a reformulation. Part I. Definition and basic properties.” *Mathematics of Operations Research*, 14, 575–625. [175]
- Mertens, Jean-François (1991), “Stable equilibria—a reformulation. Part II. Discussion of the definition, and further results.” *Mathematics of Operations Research*, 16, 694–753. [175]
- Mertens, Jean-François (2003), “Ordinality in non cooperative games.” *International Journal of Game Theory*, 32, 387–430. [178]
- Myerson, Roger B. (1978), “Refinements of the Nash equilibrium concept.” *International Journal of Game Theory*, 7, 73–80. [172]

Nash, John F (1950), "Equilibrium points in N-person games." *Proceedings of the National Academy of Sciences USA*, 36, 48–49. Reprinted in Harold W. Kuhn (ed.), *Classics in Game Theory*, Princeton University Press, Princeton, New Jersey, 1997. [167, 169]

Nash, John F (1951), "Non-cooperative games." *Annals of Mathematics*, 54, 286–295. [167, 169]

Selten, Reinhard (1965), "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit." *Zeitschrift für die gesamte Staatswissenschaft*, 121, 301–324, 667–689. [169, 170]

Selten, Reinhard (1975), "Reexamination of the perfectness concept for equilibrium points in extensive games." *International Journal of Game Theory*, 4, 25–55. Reprinted in Harold W. Kuhn (ed.), *Classics in Game Theory*, Princeton University Press, Princeton, New Jersey, 1997. [168, 169, 171]

van Damme, Eric (1984), "A relation between perfect equilibria in extensive form games and proper equilibria in normal form games." *International Journal of Game Theory*, 13, 1–13. [168, 170, 172, 178, 179, 188]

Submitted 2005-6-2. Final version accepted 2005-11-3. Available online 2005-12-2.