

This is a repository copy of *Robust equilibria In generic extensive-form games*.

White Rose Research Online URL for this paper: <u>https://eprints.whiterose.ac.uk/id/eprint/225073/</u>

Version: Published Version

Monograph:

Pahl, L. orcid.org/0000-0002-0268-371X and Pimienta, C. (2025) Robust equilibria In generic extensive-form games. Working Paper. Sheffield Economic Research Paper Series (2025001). Department of Economics, University of Sheffield ISSN 1749-8368

© 2025 The Author(s). For reuse permissions, please contact the Author(s). The Sheffield Economic Research Paper Series (SERPS) offers a forum for the research output of the Department of Economics, University of Sheffield. Papers are reviewed for quality and presentation by two internal referees and a departmental editor. However, the contents and opinions expressed remain the responsibility of the author(s). Comments are welcomed and should be addressed to the individual author(s).

Reuse

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

Takedown

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





The University Of Sheffield.

School Of Economics

Robust Equilibria in Generic Extensive-Form Games

Lucas Pahl and Carlos Pimienta

Sheffield Economic Research Paper Series

SERPS no. 2025001

ISSN 1749-8368

24 Jan 2025

ROBUST EQUILIBRIA IN GENERIC EXTENSIVE-FORM GAMES: THE TWO-PLAYER CASE*

LUCAS PAHL[†] AND CARLOS PIMIENTA[‡]

ABSTRACT. We prove the 2-player, generic extensive-form case of the conjecture of Govindan and Wilson (1997a,b) and Hauk and Hurkens (2002) stating that an equilibrium component is essential in every equivalent game if and only if the index of the component is nonzero. This provides an index-theoretic characterization of the concept of hyperstable components of equilibria in generic extensive-form games, first formulated by Kohlberg and Mertens (1986). We also illustrate how to compute hyperstable equilibria in multiple economically relevant examples and show how the predictions of hyperstability compare with other solution concepts.

1. INTRODUCTION

Hyperstable equilibria, as formulated in Kohlberg and Mertens (1986), are those Nash equilibria that are robust to perturbations not only of the payoffs of the given game, but also to *equivalent games* of the given game (i.e., games that are obtained by adding duplicate strategies). In other words, it is a solution concept characterized by two simple properties: (1) robustness to payoff perturbations (*strategic essentiality*); (2) the principle that equivalent games should have equivalent solutions (also known as *invariance*).

Hyperstable equilibria constitute an important step in Kohlberg and Mertens' stability program as the first equilibrium concept to satisfy the basic properties of existence, sequential rationality, invariance and iterated dominance, making it an appealing refinement of Nash equilibria. Even if hyperstable equilibria do not always satisfy admissibility, in

Date: March 28, 2025.

^{*} We thank Srihari Govindan for drawing our attention to the conjecture we prove in this paper and for comments, and G V A Dharanan for proofreading the manuscript. We are grateful to Paulo Barelli, Joel Sobel, John Levy and Klaus Ritzberger for comments and suggestions. We thank audiences at the Paris Game Theory Seminar, BGSE Workshop in Bonn, CES-Sorbonne, LSE-Mathematics, U. of Surrey and Royal Holloway U. of London. Lucas thanks research support from the Hausdorff Center for Mathematics (DFG project no. 390685813). Carlos thanks research support from the Australian Research Council's Discovery Project DP190102629.

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

[‡] School of Economics, The University of New South Wales, Sydney, Australia.

JEL Classication code: C72; KEYWORDS: Game Theory, Robustness of Equilibria, Fixed Point Theory, Index Theory.

generic extensive-form games hyperstable equilibrium outcomes do satisfy all those properties.¹ However, hyperstability is a highly intractable concept. Not only does its definition require robustness to payoff perturbations in the base game—which is in itself hard to check—but also in every equivalent game to the base game.

In this paper we provide a characterization of hyperstable equilibria that makes it a more operational concept, as we intend to illustrate in the last section with a series of examples. This characterization dispenses with the need of considering equivalent games to the base game and, in several instances, makes the problem of checking for robustness to payoff perturbations almost trivial.

This gain in simplicity is achieved by formulating the characterization in terms of a concept of topological fixed point theory called the *index of fixed points*, which we now explain. In algebraic topology, tools such as the fixed point index are developed to study and distinguish among fixed points. In particular, when studying fixed points in an abstract setting, it is also of interest to identify robustness properties of the fixed points with respect to perturbations of their fixed-point maps. A topologically essential fixed point of a map is a point such that every nearby map has a close-by fixed point. Topologically essential fixed points do not always exist,² but considering the sets of fixed points of a given map, one can extend the definition of topological essentiality from points to *connected components of fixed points*:³ a connected component of the set of fixed points is topologically essential if for sufficiently small perturbations of the fixed-point map, there are fixed points close to the component. Topologically essential components can be identified by an integer number called their *index* (cf. O'Neill (1953)). The process through which this number is assigned is rather involved and will be explained later in more detail, but the important fact is that a component of fixed points is topologically essential if and only if the index is non-zero.

In a game-theoretic context, and for some function whose fixed points coincide with the Nash equilibria of a game, an index can therefore be assigned to every component of Nash equilibria in mixed strategies, and simple properties of the index of fixed points guarantee that at least one topologically essential component of equilibria exists. The conjecture that the strategically essential components coincide with those with nonzero index (Govindan and Wilson, 1997a,b), i.e. the topologically essential components, was disproved by Hauk

¹In generic extensive-form games, hyperstable equilibrium outcomes are always induced by a Mertens-stable set (Mertens, 1989) which, in turn, does satisfy all those properties.

²Consider the identity map from the interval [0,1] to itself. Any point in [0,1] is a fixed point of the map, but if we consider uniform perturbations of this map to near-by maps, none of the points of [0,1] is robust.

³O'Neill (1953) shows that a minimal topologically essential fixed point set of an Euclidean space must be a connected component of the set of fixed points, thus justifying the formulation of topological essentiality to connected components of fixed points.

and Hurkens (2002) by means of a counterexample. They recast the conjecture and hypothesized that the hyperstable components are precisely those that have nonzero index. We show that this is indeed the case for generic two-player extensive-form games. Our notion of genericity implies the commonly used property of extensive-form games according to which each connected component of equilibria in mixed strategies induces a single outcome in an extensive-form game. Therefore, our characterization can be seen as identifying the Nash equilibrium outcomes of an extensive-form game which are hyperstable. While intuitively every topologically essential Nash equilibrium component should be hyperstable, this result implies that the space of perturbations considered in the definition of hyperstability i.e. payoff perturbations - is as rich as the one used by topological essentiality, i.e. the space of perturbations of a map whose fixed points are the equilibria of a game.

A similar result was proved by Govindan and Wilson (2005), henceforth GW, for arbitrary games. However, to obtain such a characterization they require a strengthening of hyperstability beyond the combination of invariance and robustness to payoff perturbations. GW call such a concept *uniform hyperstability*.⁴

As mentioned above, only information about the game under study is necessary to verify if a component is topologically essential. Indeed, Shapley (1974) shows that the index of equilibria in generic two-player games can be computed by a simple formula that only depends on parameters of the game. Balthasar (2009) generalized this formula to include non-generic two player games. Even if such formulas can be useful, for many relevant game-theoretical applications in the literature, it is enough to use the properties of the index to know if it is different from zero and, given our result, verify whether a component is hyperstable.⁵

From an axiomatic perspective, our result can be interpreted as providing an index theoretic characterization for the solution concept that, for each (generic two-player extensiveform) game selects all (closed and connected) subsets of Nash equilibria that satisfy (1) invariance and (2) robustness to payoff perturbations. A component of equilibria with nonzero index is known to be robust in every equivalent game, so a solution concept that selects the collection of equilibrium components with non-zero index is both invariant and robust to payoff perturbations. Conversely, if a solution concept selects all (closed and connected) sets of Nash equilibria that are invariant and robust to payoff perturbations then the selected solutions must be connected components of Nash equilibria.⁶ Furthermore, our main result

⁴ We define uniform hyperstability in Section 2.4.

⁵This is illustrated in Section 6.

⁶ Suppose that at least one solution is a proper subset of a connected component. Then our arguments imply that an equivalent game exists in which an equivalent solution is not robust to payoff perturbations.

LUCAS PAHL AND CARLOS PIMIENTA

implies that any zero index component is not robust in some equivalent game. Hence, all selected solutions are connected components with non-zero index.

In parallel to the developments in the strategic stability literature, the index of equilibria has also been used independently as a tool to select equilibria with interesting properties, especially with regards to their dynamic stability. Ritzberger (1994) formulated the index of equilibria for the replicator dynamics and presented some useful applications in game theory, including highlighting that non-zero index components of equilibria satisfy a number of desirable properties. Demichelis and Ritzberger (2003) obtained a necessary condition for the asymptotic stability of equilibria in dynamic adjustment processes that is formulated in terms of the index of equilibria. More recently, McLennan (2016) argued for the selection of equilibria with index +1 on both experimental and theoretical grounds, while Govindan et al. (2023) provided game-theoretic characterization for these. A natural question stemming from this strand of the literature is: what does the index precisely mean in game-theoretic terms? More precisely, while one can use non-zero index equilibria as a selection criterion to guarantee interesting properties for the selected equilibria, such a criterion is purely topological and may imply properties that are game-theoretically undesirable. Our main result provides an answer to this question by formulating the precise combination of game-theoretical properties that characterize non-zero index equilibria and define the classical concept of hyperstability.

This paper is organized as follows. Section 2 presents all the required notation and definitions. In particular, it introduces the necessary properties of the index that will be applied throughout the paper and states the main result. Section 3 defines one the main objects of the analysis—the *excluded game* associated to an equilibrium outcome. Section 4 lays down the precise genericity assumptions under which our main result holds. We prove our main result in Section 5. Such a proof is divided in three steps and is accompanied by a running example to illustrate the main ideas. To finalize, Section 6, contains an illustration on how to compute the index in many economically relevant examples, allowing to identify nonhyperstable equilibria. The Appendix and Online Appendix contain additional definitions and proofs that do not appear in the main text.

1.1. An example

To illustrate the main ideas of the result and the proof, consider the entry game in Figure 1. To analyze the example, we use two simple facts about the index of equilibria. The first one is that the index of a strict equilibrium is +1. The second is that the sum of the indexes is always +1. The game in Figure 1 has two components of Nash equilibria. In the



first component, player 1 moves In and then both players play (L, ℓ) in the subgame. Hence, this component is made of a single strategy profile that, furthermore, is a strict equilibrium as any deviation by either player leaves such player with a strictly lower payoff. (Note as well that (L, ℓ) is also a strict equilibrium in the subgame.) Since the index of a strict equilibrium is +1 and indexes must add up to one, the second component must have zero index. This second component is such that player 1 moves *Out* with probability one and player 2 plays r in the subgame with probability at least 1/3. We call the set of strategy profiles in the subgame such that player 2 plays in this way the *supporting polytope* of the component because it supports on-path play, i.e., at any point in this component player 2 cannot affect the outcome and player 1 is at most indifferent between *Out* and L when player 2 plays rwith probability exactly equal to 1/3.

We can think of the "Out"-component as excluding the subgame, and of the supporting polytope as containing behavior in the subgame that supports that exclusion. The supporting polytope contains two Nash equilibria of the subgame, (R,r) and the mixed strategy profile $(\frac{3}{4}L + \frac{1}{4}R, \frac{1}{4}\ell + \frac{3}{4}r)$. When considering the subgame in isolation, the former equilibrium has index +1 because it is strict, and the latter has index -1 because the sum of the indexes in the subgame must be +1 and the other two equilibria are strict ((R,r) as just mentioned, and (L,ℓ)). If we extend the notion of index of a component of equilibria to the index of a (suitably chosen) neighborhood by adding the indexes of the components included in the neighborhood, we note that the index of any sufficiently small neighborhood of the supporting polytope must be zero.

Given a game, we show in Section 5.2 that if the index of a neighborhood is zero, then there exists a payoff perturbation of an equivalent game with no equilibrium inside the neighborhood and such that, payoffs *outside* such a neighborhood are as close as we want to those in the original game. Figure 2 illustrates. The interval (x_1, x_2) contains two fixed points in which the function cuts the 45° line from below (index -1) and then from above





(index +1).⁷ We can pull the graph of f so that the new function whose graph coincides with the dashed line when $x \in (x_1, x_2)$ has no fixed point in (x_1, x_2) . Thus, if G is the subgame in Figure 1, there exists an equivalent game \bar{G} (i.e., a game which is obtained by adding finitely many pure strategies that are duplicates of mixed strategies of G) and a perturbation of it \bar{G}^{δ} with no equilibria in the (subset of the strategy profile equivalent to the) supporting polytope and such that payoffs outside that polytope are at least δ -close to those in \bar{G} .⁸

With this information, consider the extensive-form game in Figure 3. Nature moves first. With probability $1-\varepsilon$ players play the original entry game but with the subgame G replaced by the equivalent subgame \overline{G} . With probability ε , players play the perturbed game \overline{G}^{δ} that does not have an equilibrium in the supporting polytope. Player 1 observes the move of Nature, but player 2 does not know if she is playing because Nature chose the perturbed game \overline{G}^{δ} or because player 1 chose In. This is captured by player 2's information set connecting nodes representing games \overline{G} and \overline{G}^{δ} . Note that when $\varepsilon = 0$ this extensive-form game is equivalent to the game in Figure 1.⁹ It can be proved that when ε is small enough this equivalent game does not have any Nash equilibrium component in which player 1 plays *Out*. Intuitively, if there were such outcome, player 2 would have to be playing in the (set equivalent to the) supporting polytope - so that *Out* is a best reply for player 1. Knowing

 $[\]overline{{}^{7}}$ In this one-dimensional case, if x^* is a fixed point, the index of x^* is $\operatorname{Ind}_f(x^*) = \operatorname{sign}\left(\frac{\partial (Id-f)(x)}{\partial x}\right|_{x=x^*}$

⁸The proof of existence of \bar{G}^{δ} relies on the techniques developed in GW, but it is not an exact application of their result, as Online Appendix B demonstrates.

⁹More precisely, the game with $\varepsilon = 0$ is equivalent to the original game in the sense that both have the same reduced normal forms.



FIGURE 3. A game equivalent to that in Figure 1 when $\varepsilon = 0$

that, in such an equilibrium, if player 2 is called to move, then she knows it is because Nature chose \bar{G}^{δ} . But such a game does not have an equilibrium in the supporting polytope. We can similarly prove that there is no sequence of equilibria in which player 1's strategy converges to playing *Out* with probability one as ε converges to zero. Hence, when ε is small enough, Figure 3 gives us a perturbation of a game equivalent to the entry game of Figure 1 with no equilibrium close to the zero-index "*Out*"-component.

In the rest of the paper we show that under some generic conditions every key step in this example generalizes. In particular, given a two-player game and a component that induces a unique outcome, we can always define an "excluded game" (which is not necessarily a subgame) whose supporting polytope (capturing behavior in the excluded game that prevents it from being reached) has zero index in such an excluded game. And with that, we can construct a game equivalent to the original and a perturbation of its payoffs analogous to Figure 3 with the property that the resulting game has no equilibrium close to the zero index component.

2. Preliminaries

2.1. Equivalent strategies, equivalent normal-form games

A finite two-player normal-form game $\mathbb{G} \equiv (S_1, S_2, \mathbb{G}_1, \mathbb{G}_2)$ is a four-tuple where, for each n = 1, 2, player *n*'s finite set of pure strategies is S_n and $\mathbb{G}_n : S_1 \times S_2 \to \mathbb{R}$ is player *n*'s payoff function. As usual, $\Sigma_n \equiv \Delta(S_n)$ is player *n*'s set of mixed strategies and $S \equiv S_1 \times S_2$ and $\Sigma \equiv \Sigma_1 \times \Sigma_2$ are, respectively, the sets of pure and mixed strategy profiles. We also denote by \mathbb{G}_n the multilinear extension of player *n*'s payoff function to the set of mixed strategy profiles Σ .

Two strategies $\sigma_n, \sigma'_n \in \Sigma_n$ are *equivalent* if for both m = 1, 2 we have $\mathbb{G}_m(\sigma_n, s_{-n}) = \mathbb{G}_m(\sigma'_n, s_{-n})$ for all $s_{-n} \in S_{-n}$. Given $\mathbb{G} = (S_1, S_2, \mathbb{G}_1, \mathbb{G}_2)$, the *reduced normal-form* of \mathbb{G} is a normal-form game $\mathbb{G}' = (S'_1, S'_2, \mathbb{G}_1, \mathbb{G}_2)$ that is obtained from \mathbb{G} by eliminating pure strategies that are equivalent to some existing mixed strategy. That is, $S'_n \subseteq S_n$ and if $s_n \in S_n \setminus S'_n$ then

there is $\sigma_n \in \Sigma'_n \equiv \Delta(S'_n)$ that is equivalent to s_n . Note that, up to relabeling of strategies, the reduced normal-form of a finite game is unique.

Definition 2.1. Two games $\overline{\mathbb{G}}$ and \mathbb{G} are *equivalent* if they have the same reduced normal-form.

Given two equivalent games $\mathbb{G} = (S_1, S_2, \mathbb{G}_1, \mathbb{G}_2)$ and $\overline{\mathbb{G}} = (\overline{S}_1, \overline{S}_2, \overline{\mathbb{G}}_1, \overline{\mathbb{G}}_2)$ we extend the notion of equivalence between strategy profiles in the same game to equivalence of strategy profiles in equivalent games. Say that $\overline{\sigma}_n \in \overline{\Sigma}_n$ and $\sigma_n \in \Sigma_n$ are *equivalent* if there exists a strategy σ'_n in their (common) reduced normal-form game such that $\overline{\sigma}_n$ is equivalent to σ'_n (both viewed as strategy profiles in $\overline{\mathbb{G}}$) and σ_n is equivalent to σ'_n (both viewed as strategy profiles in $\overline{\mathbb{G}}$) and σ_n is equivalent to σ'_n (both viewed as strategy profiles in $\overline{\mathbb{G}}$). If \mathbb{G} and $\overline{\mathbb{G}}$ are two equivalent games we say profile $\sigma \in \Sigma$ is equivalent to profile $\overline{\sigma} \in \overline{\Sigma}$ if for each player n, the strategy σ_n is equivalent to $\overline{\sigma}_n$. A subset $T \subset \Sigma$ is equivalent to $\overline{T} \subset \overline{\Sigma}$, if for each $\sigma \in T$ there exists $\overline{\sigma} \in \overline{T}$ such that σ is equivalent to $\overline{\sigma}$ and for each $\overline{\sigma} \in \overline{T}$ such that $\overline{\sigma}$ is equivalent to $\overline{\sigma}$.

2.2. Extensive-form games

We introduce notation and basic definitions. For a formal definition of extensive-form game with perfect recall see, e.g., Osborne and Rubinstein (1994). Consider a *two-player finite game tree with perfect recall* $\Gamma \equiv (T, <, U, A, \rho)$. The set of nodes is T and < denotes the precedence relation in the tree. The set of terminal nodes is $Z \subset T$. The collection Uis a partition of $T \setminus Z$ into information sets of players and Nature. The set $U_n \subset U$ is the collection of information sets for player n = 0, 1, 2, where player 0 represents Nature. (Every element in U_0 is a singleton.) The set of all actions in the game is A and $A_n(u)$ represents player n's set of actions available at her information set $u \in U_n$. Let $A_n \equiv \bigcup_{u \in U_n} A_n(u)$ be the entire set of player n's actions. Moves of Nature are given by the function ρ that, to every $u \in U_0$, assigns a completely mixed probability distribution on $A_0(u)$.

For n = 1, 2, player *n*'s set of pure strategies is $S_n \equiv \{s : U_n \to A_n \mid s(u) \in A_n(u)\}$ and her mixed strategy set is $\Sigma_n \equiv \Delta(S_n)$. Hence, we have the sets of strategy profiles $S \equiv S_1 \times S_2$ and $\Sigma = \Sigma_1 \times \Sigma_2$. Since we only consider games with perfect recall, Kuhn's theorem implies that we can work with behavior strategies whenever convenient. Given player n = 1, 2, a behavior strategy $b_n = (b_n(i))_{i \in A_n}$ satisfies $(b_n(i))_{i \in A_n(u)} \in \Delta(A_n(u))$ for every $u \in U_n$. We let B_n be player *n*'s set of behavioral strategies.

For a fixed two-player game tree Γ , the space of games is $\mathscr{G} \equiv \mathbb{R}^{2|Z|}$. A game $G \in \mathscr{G}$ assigns payoff $G_n(z)$ to player n at final node z. We refer to the extensive-form game defined by Γ and terminal payoffs G simply by G. The space of outcomes is $\Delta(Z)$, where an outcome $Q \in \Delta(Z)$ assigns probability Q(z) to z. Given action $a \in A$ the set of terminal nodes that come after action a is Z(a). Similarly, given an information set $u \in U$ we let Z(u) be the set of terminal nodes that come after some action available at u. We denote expressions Q(Z(a)) and Q(Z(u)) simply as Q(a) and Q(u). Given a profile of mixed strategies $\sigma \in \Sigma$ the induced outcome when players play according to σ is $\mathbb{P}(\cdot | \sigma) \in \Delta(Z)$. With abuse of notation we also let $G : \Delta(Z) \to \mathbb{R}$ denote the expected utility function associated with game $G \in \mathcal{G}$. The normal-form representation of Gis $\mathbb{G} = (S_1, S_2, \mathbb{G}_1, \mathbb{G}_2)$ where, for each n = 1, 2, we have $\mathbb{G}_n(s) = G_n(\mathbb{P}(\cdot | s))$ for all $s \in S$.

2.3. Index Theory

The fixed point index contains information about the robustness of fixed points of a map when such a map is perturbed to a nearby map. Since Nash equilibria are fixed points, we can apply index theory to them (cf. Ritzberger, 1994). The classical introduction to index theory usually requires some concepts from algebraic topology. For the purposes of this paper, this can be bypassed without much hindrance. Results in this section can be found in McLennan (2018, pp 245-265) and Pahl (2023).

Let \mathbb{G} be a normal-form game with mixed strategy Σ . Given a neighborhood O of Σ suppose $f: O \to \Sigma$ is a differentiable map. Let d_f be the displacement of f, i.e., $d_f(\sigma) = \sigma - f(\sigma)$. Then the fixed points of f are the zeros of d_f . Suppose now that the Jacobian of d_f at a zero σ of f is nonsingular. We assign the index +1 to σ if the determinant of the Jacobian of d_f is positive or -1 if such a determinant is negative. For the next definition, given a subset \mathcal{O} of Σ , we denote by $cl_{\Sigma}(\mathcal{O})$ the closure of \mathcal{O} with respect to Σ .

Definition 2.2. An open neighborhood $\mathcal{O} \subset \Sigma$ of a component of equilibria (in mixed strategies) *K* of normal-form game \mathbb{G} is *admissible* if every equilibrium of \mathbb{G} in $cl_{\Sigma}(\mathcal{O})$ belongs to *K*.

When a finite game \mathbb{G} has a component of equilibria K that consists of more than a single strategy profile, we extend the definition of the index as follows. Take a continuous map $g_{\mathbb{G}}$: $\Sigma \to \Sigma$ such that the fixed points of $g_{\mathbb{G}}$ are the Nash equilibria of game \mathbb{G} and $g_{\mathbb{G}}$ continuously depends on the payoffs of \mathbb{G} . An example of such a map is the map Nash (1951) constructed to prove existence of equilibria in mixed strategies. Consider a neighborhood O of Σ and $r: O \to \Sigma$ a retraction to the closest point in Σ . Let $\mathcal{O} \subset O$ be an open neighborhood of K in the affine space generated by Σ , whose closure contains no other fixed point of g. We approximate ($g_{\mathbb{G}} \circ r$) uniformly by a differentiable function $f: O \to \Sigma$ without fixed points on the boundary of \mathcal{O} and such that the displacement of f at any fixed point has nonsingular Jacobian. For any sufficiently close approximation, the sum of the indexes of the (isolated) fixed points of f in \mathcal{O} is constant and can be taken as the definition of the index of the component K.¹⁰

We can now define the index with respect to the best-reply correspondence of game \mathbb{G} following the same procedure as in GW. Consider now $\mathcal{O} \cap \Sigma$. For notational convenience, we will denote this intersection from now on as \mathcal{O} . Let W be an open neighborhood of Graph(BR^G) such that $W \cap \{(\sigma, \sigma) \in \Sigma \times \Sigma \mid \sigma \in cl_{\Sigma}(\mathcal{O}) - \mathcal{O}\} = \emptyset$. There exists $\tilde{W} \subset W$ a neighborhood of Graph(BR^G) such that any two continuous functions f_0 and f_1 from Σ to Σ whose graphs are in the neighborhood \tilde{W} are homotopic by a homotopy $H: [0,1] \times \Sigma \to \Sigma$ with $\operatorname{Graph}(H(t,\cdot)) \subset \widetilde{W}$ for all $t \in [0,1]$ (cf. McLennan, 1989). Take a continuous map $f: \Sigma \to \Sigma$ with $\operatorname{Graph}(f) \subset \tilde{W}$. We define the *best-reply index* of component *K*, denoted $\operatorname{Ind}_{\operatorname{BR}^{\mathbb{G}}}(K)$, as the fixed point index of the continuous map $f|_{\mathcal{O}}: \mathcal{O} \to \Sigma$. The choice of the neighborhood W and the homotopy property of the index (see Dold, 1972, Chapter VII, 5.15) imply that the index of the component is the same for any continuous map with graph in the neighborhood \tilde{W} . We note that defining the index of a component from, say, the map that Nash used in (Nash, 1951) or from the best-reply correspondence are two distinct ways of defining the index, which can be shown to be equivalent (cf. GW). In addition, this process to define the index applies with insignificant changes if the correspondence between the simplices of strategies is contractible valued instead of convex valued (cf. McLennan, 1989).¹¹ This fact will play an important role in Section 5.1, as we need to consider selections of a best-reply correspondence which are not necessarily convex-valued, but are contractible-valued.

One can generalize the definition of the best-reply index of a component of equilibria to the *best-reply index of an admissible neighborhood*, by using the exact same procedure as in the previous paragraph. This yields the index of a neighborhood by summing the indexes of the components of equilibria which are contained in it. For convenience, whenever we refer to the index of a component or an admissible neighborhood, it will be implicit that we refer to the best-reply index.

We are now ready to recall a few known properties of the index of equilibria which we will use in the proof of our main result. The proofs that the index satisfies such properties can be found in McLennan (2018) or GW.

I.1 The index of an admissible neighborhood is locally constant with respect to payoff perturbations of the game. Formally, fix an admissible neighborhood O in the mixed

¹⁰Under the continuous dependence of the payoffs of G, this definition of index is independent of the particular map $g_{\mathbb{G}}$ used, it only depends on the game G (cf. DeMichelis and Germano (2000)).

¹¹ A topological space X is contractible if there exists a continuous map $T : [0,1] \times X \to X$ and $x_0 \in X$ such that $T(0,\cdot) = \operatorname{id}_X$ and $T(1,\cdot) = x_0$.

strategy set of a finite game G. Then there exist $\overline{\delta} > 0$, such that for any $0 \le \delta \le \overline{\delta}$ and a δ -payoff-perturbation \mathbb{G}^{δ} of game G, the index of O (with respect to \mathbb{G}^{δ}) is constant.

- **1.2** The index of an equilibrium component is invariant to equivalent presentations of a game. If K is a component of equilibria of \mathbb{G} with index c, then for any equivalent game $\overline{\mathbb{G}}$, the index of the equivalent component \overline{K} is also c.
- **I.3** The index of a component is invariant to the deletion of strictly inferior replies to the component. If K is a component of equilibria with index c of game G, then deleting from the normal-form of game G the pure strategies of player n which are strictly inferior replies to every profile in the component K yields a new game G' with the same component K as an equilibrium component in G' and with the same index c.

We also need three well-known properties of the index. For our purposes, we particularize their statements as follows:

- **I.4** *Multiplication*: Let $BR^{\mathbb{G}} : \Sigma \Rightarrow \Sigma$ and $BR^{\mathbb{G}'} : \Sigma' \Rightarrow \Sigma'$ be the best-reply correspondences of (respectively) games \mathbb{G} and \mathbb{G}' . Let $BR^{\mathbb{G}} \times BR^{\mathbb{G}'}$ be the correspondence taking $(\sigma, \tau) \in$ $\Sigma \times \Sigma'$ to $BR^{\mathbb{G}}(\sigma) \times BR^{\mathbb{G}'}(\tau)$. If $O \times O' \subset \Sigma \times \Sigma'$ contains no fixed points of $BR^{\mathbb{G}} \times BR^{\mathbb{G}'}$ in its boundary, then O (respectively O') has no fixed points of $BR^{\mathbb{G}}$ (respectively $BR^{\mathbb{G}'}$) in its boundary, and the index of $O \times O'$ with respect to $BR^{\mathbb{G}} \times BR^{\mathbb{G}'}$ is the multiplication of the indexes of O (with respect to $BR^{\mathbb{G}}$) and of O' (with respect to $BR^{\mathbb{G}'}$).
- **I.5** *Commutativity*: Let $F : \Sigma \Rightarrow \Sigma$ be an upper hemicontinuous correspondence that is nonempty, compact and convex valued. Let $e : \Sigma \to \Sigma'$ be a continuous map with left-inverse $q : \Sigma' \to \Sigma$. If X is a component of fixed points of $e \circ F \circ q : \Sigma' \Rightarrow \Sigma'$ then q(X) is a component of fixed points of F and their indexes agree.
- **I.6** *Excision*: Let $F : \Sigma \Rightarrow \Sigma$ be an upper hemicontinuous correspondence that is nonempty, compact and convex valued. Suppose $\tilde{\mathcal{O}}$ and \mathcal{O} are both admissible neighborhoods in Σ whose closures in Σ contain the same fixed points of F. Then the index of \mathcal{O} and that of $\tilde{\mathcal{O}}$ with respect to F are identical.

We conclude this section with the following result. In particular, it implies that if a component has zero index, then some ending node is reached with probability zero.

Proposition 2.3. If an equilibrium outcome Q induced by a component of equilibria in mixed strategies K has full support, then K has non-zero index.

Proof. See Appendix A.

2.4. Main Result

We recall the definition of Hyperstability.

Definition 2.4. A component *K* of equilibria in mixed strategies of a normal-form game \mathbb{G} is *hyperstable* if for each equivalent game $\overline{\mathbb{G}}$ and every $\varepsilon > 0$, there exists $\delta > 0$ such that any δ -payoff-perturbation $\overline{\mathbb{G}}^{\delta}$ of $\overline{\mathbb{G}}$ has an equilibrium which is ε -close to \overline{K} (the equivalent component in $\overline{\mathbb{G}}$ to *K*).

Hyperstability is a property first formulated in Kohlberg and Mertens (1986) to refine Nash equilibria. It implies several desirable properties.¹² In Section 5 we prove the following result.

Theorem 2.5. Fix a two-player game tree with perfect recall. Apart from a lower-dimensional, semi-algebraic set of payoffs, an equilibrium component is hyperstable if and only if it has non-zero index.

GW also offer a characterization of non-zero index Nash equilibrium components. However, their characterization requires a strengthening of hyperstability. Namely, GW show that a component has non-zero index if and only if it is *uniformly hyperstable*, that is, if for each $\varepsilon > 0$, there exists $\delta > 0$ such that for any equivalent game $\overline{\mathbb{G}}$ and any δ -perturbation $\overline{\mathbb{G}}^{\delta}$, there exists an equilibrium of $\overline{\mathbb{G}}^{\delta}$ which is within ε of \overline{K} (the equivalent component of $\overline{\mathbb{G}}$). Note that δ is fixed across all equivalent games, making uniform hyperstability a stronger concept than the combination of robustness to payoff perturbations and invariance. Nonetheless, our result implies that for generic two-player extensive-form games, a component is uniformly hyperstable if and only if it is hyperstable.

3. EXCLUDED GAMES

The "Out"-component in the game in Figure 1 has player 1 excluding the subgame. In turn, behavior in the subgame as prescribed by such a component prevents player 1 from profiting by deviating and choosing In. In general, observed behavior crucially depends on the robustness of unobserved behavior. As we discussed in the Introduction, the reason the "Out"-component has zero index is that it specifies a set of strategy profiles in the subgame that, when analyzed relative to such a subgame, has itself zero index. In this section, we generalize this insight along with the notion of "excluded game" so that, given an equilibrium component, we can examine the interaction between the two players that occurs both on-path and off-path.

Take a two-player finite game tree with perfect recall Γ and let $\mathscr{G} = \mathbb{R}^{2|Z|}$ be the space of payoffs over terminal nodes. For payoffs $G \in \mathscr{G}$, let \mathbb{G} be the induced normal form. From

 $^{^{12}}$ Hyperstable components of equilibria always contain proper (and therefore sequential) equilibria, satisfy invariance, and are robust to payoff perturbations.

13

now on, fix an equilibrium component K that induces a unique outcome $Q \in \Delta(Z)$. Given the equilibrium component K, an information set u is said to be *on-path* if Q(u) > 0. Player n's collection of information sets that are on-path is $U_n^+ \equiv \{u_n \in U_n \mid Q(u) > 0\}$. Similarly, an information set is *off-path* if Q(u) = 0. Player n's collection of information sets that are off-path of Q is $U_n^0 \equiv U_n \setminus U_n^+$.

Let Z^0 be the set of terminal nodes which have probability 0 under Q. We define

$$S_n^n \equiv \{s_n \in S_n \mid \text{ there exists } \sigma \in K \text{ with } \mathbb{P}(Z^0 \mid s_n, \sigma_{-n}) > 0\}$$

and call it player *n*'s set of *observable deviations*. It consists of those player *n*'s pure strategies in which, at some information set on-path according to *K*, player *n* deviates and plays an action that has zero probability under *Q*. As usual, $\sum_{n=1}^{n} \equiv \Delta(S_{n}^{n})$.

Let

(1)
$$S_n^+ \equiv \left\{ s_n^+ : U_n^+ \to A_n \mid s_n^+(u) \in A_n(u) \text{ for every } u \in U_n^+ \right\}$$

and $\Sigma_n^+ \equiv \Delta(S_n^+)$. Note that these are not pure strategies because they assign choices only to on-path (of *K*) information sets.

The subset $S_n^{\otimes} \subset S_n^+$ is the collection of assignments of actions at on-path information sets that only assign actions that are taken in the equilibrium component K, that is

$$S_n^{\otimes} \equiv \left\{ s_n^{\otimes} \in S_n^+ \mid Q(s_n^{\otimes}(u)) > 0 \text{ for every } u \in U_n^+ \right\},\$$

where the symbol $Q(s_n^{\otimes}(u))$ follows the notational convention established in section 2.2. Define $\Sigma_n^{\otimes} \equiv \Delta(S_n^{\otimes})$ and $\Sigma^{\otimes} \equiv \Sigma_1^{\otimes} \times \Sigma_2^{\otimes}$. Analogously to (1), define also

(2)
$$S_n^0 \equiv \left\{ s_n^0 : U_n^0 \to A_n \mid s_n^0(u) \in A_n(u) \text{ for every } u \in U_n^0 \right\}$$

together with $\Sigma_n^0 \equiv \Delta(S_n^0)$.

We will now split the mixed strategies of a player between an on- and off-path part. Given some player *n*'s pure strategy $s_n \in S_n$ the restriction of s_n to information sets in, respectively, U_n^+ and U_n^0 is $q_n^+(s_n)$ and $q_n^0(s_n)$. These define maps $q_n^+: S_n \to S_n^+$ and $q_n^0: S_n \to S_n^0$. Linear interpolation extends q_n^+ to a map from Σ_n to Σ_n^+ . Analogously, q_n^0 is extended to a map from Σ_n to Σ_n^0 . Given a mixed strategy $\sigma_n \in \Sigma_n$, $q_n^+(\sigma_n)$ is, therefore, the marginal of σ_n over S_n^+ ; and $q_n^0(\sigma_n)$ is the marginal of σ_n over S_n^0 . In particular, if $s_n \in S_n \setminus S_n^n$, then s_n can be written as a pair $(s_n^{\otimes}, s_n^0) \in S_n^{\otimes} \times S_n^0$ and we let $q_n^{\otimes}(s_n) = s_n^{\otimes}$. Similarly to above, $q_n^{\otimes} : \Delta(S_n \setminus S_n^n) \to \Sigma_n^{\otimes}$ is the affine map for which $q_n^{\otimes}(\sigma_n^{\otimes})$ is the marginal of $\sigma_n \in \Delta(S_n \setminus S_n^n)$ over Σ_n^{\otimes} . Let $q^+ \equiv q_1^+ \times q_2^+$, $q^{\otimes} \equiv q_1^{\otimes} \times q_2^{\otimes}$, and $q^0 \equiv q_1^0 \times q_2^0$.

Given any pair $(\sigma_n^+, \sigma_n^0) \in \Sigma_n^+ \times \Sigma_n^0$ let $\sigma_n^+ * \sigma_n^0$ represent the product strategy σ_n that satisfies $\sigma_n(s_n^+, s_n^0) = \sigma_n^+(s_n^+) \cdot \sigma_n^0(s_n^0)$ for every $s_n = (s_n^+, s_n^0)$.

If σ_n is a product strategy then, by definition, there is a pair $(\sigma_n^+, \sigma_n^0) \in \Sigma_n^+ \times \Sigma_n^0$ such that $\sigma_n = \sigma_n^+ * \sigma_n^0$.

Obviously, not every mixed strategy is a product strategy over $S_n^+ \times S_n^0$. Nonetheless, for any mixed strategy of a player there exists an outcome-equivalent product strategy. This observation is formalized below and follows directly from Kuhn's Theorem.

Lemma 3.1. For every $\sigma_n \in \Sigma_n$ there exists a pair $(\hat{\sigma}_n^+, \hat{\sigma}_n^0) \in \Sigma_n^+ \times \Sigma_n^0$ such that, given any $\sigma_{-n} \in \Sigma_{-n}$, the strategy $\hat{\sigma}_n^+ * \hat{\sigma}_n^0$ induces the same outcome as σ_n .

Definition 3.2 (Excluded game). Fix $\hat{\tau} \in K$ such that $\hat{\tau}_n = \hat{\tau}_n^{\otimes} * \hat{\tau}_n^0$ with $\hat{\tau}_n^{\otimes} \in int(\Sigma_n^{\otimes})$, where the interior is relative to the affine space generated by Σ_n^{\otimes} . Player *n*'s *excluded game* is the normal-form game in which player *n*'s strategy set is S_n^n , player *-n*'s strategy set is S_{-n}^0 and the payoff function for player m = 1, 2 is defined for each profile $(s_n^n, s_{-n}^0) \in S_n^n \times S_{-n}^0$ as

(3)
$$\mathbb{G}_m^n(s_n^n, s_{-n}^0) \equiv \sum_{\substack{s_{-n}^{\otimes} \in S_{-n}^{\otimes}}} \hat{\tau}_{-n}^{\otimes}(s_{-n}^{\otimes}) \mathbb{G}_m(s_n^n, s_{-n}^{\otimes}, s_{-n}^0).$$

Since the equilibrium component K induces a unique probability distribution Q, the payoff function (3) does not depend on the particular choice of $\hat{\tau} \in K$, on the fact that $\hat{\tau}_n$ is a product (as $q_n^{\otimes}(\hat{\tau}_n)$ can be used instead), or on the fact that $\hat{\tau}_n^{\otimes}$ is chosen (for convenience) to be in the interior of Σ_n^{\otimes} .

Remark 3.3. We note two particular cases. First, when no information set of player -n is off-path in component K then player n's excluded game is just a decision problem. Second, if K is such that player n does not have any observable deviation, as player 2 in the "Out"-component of Figure 1, then we say that player n does not have an excluded game. Also note that while in Figure 1 player 1's excluded game coincides with the proper subgame, in general, an excluded game is not necessarily a subgame.

Let K_{-n}^n be the subset of player -n's mixed strategies in player n's excluded game such that player n obtains a payoff that is no larger than what she obtains under the equilibrium component K:

$$K_{-n}^n \equiv \left\{ \sigma_{-n}^0 \in \Sigma_{-n}^0 \mid \text{ for all } s_n^n \in S_n^n \text{ we have } \mathbb{G}_n^n(s_n^n, \sigma_{-n}^0) \le G_n(Q) \right\}.$$

And let ∂K_{-n}^n be the subset of K_{-n}^n where player *n* has at least one deviation for which she is indifferent:

$$\tilde{\partial} K_{-n}^n \equiv \Big\{ \sigma_{-n}^0 \in K_{-n}^n \mid \text{ for some } s_n^n \in S_n^n \text{ we have } \mathbb{G}_n^n(s_n^n, \sigma_{-n}^0) = G_n(Q) \Big\}.$$

The set $\tilde{\partial} K_{-n}^n$ contains the boundary of K_{-n}^n relative to Σ_{-n}^0 but, in general, it is not equal to it. Player *n*'s supporting polytope is $K^n \equiv \Sigma_n^n \times K_{-n}^n$. We also define $\tilde{\partial} K^n \equiv \Sigma_n^n \times \tilde{\partial} K_{-n}^n$.

Later we show that under a generic choice of terminal payoffs for the game tree K^n is a full dimensional neighborhood in the strategy set of the excluded without equilibria in its boundary realtive to this strategy set, thus admiting an index.

Example 3.4. In Figure 1, the supporting polytope K^1 to the "Out" component is $\Delta(\{L,R\}) \times \{\alpha \ell + (1-\alpha)r \mid 0 \le \alpha \le 2/3\}.$

Remark 3.5. Proposition 2.3 implies that if a zero-index component of equilibria in mixed strategies K induces a unique outcome then such an outcome must not have full support. This result does not require any genericity assumption and it is a consequence of the twoplayer environment. Furthermore, assume that the component K is such that for both n = 1, 2 we have $\partial K^n = \emptyset$. Then every strategy $s_n^n \in S_n^n$ pays strictly less than $G_n(Q)$ against any equilibrium strategy of player -n. Hence, we can eliminate the deviations s_n^n from \mathbb{G} for both players leaving the index of component K invariant (cf. Property I.3). The resulting game obtained after this elimination has an extensive-form where the outcome Q induced by K is completely mixed. Proposition 2.3 then implies that K has non-zero index. Therefore, any zero-index equilibrium component K is such that $\partial K^n \neq \emptyset$ for some n = 1, 2.

Definition 3.6 (Included game). The *included game* \mathbb{G}^{\otimes} associated with component K is the normal-form game in which player n's strategy set is S_n^{\otimes} and payoff functions for player m = 1, 2 are given by:

$$\mathbb{G}_m^{\otimes}(s_n^{\otimes}, s_{-n}^{\otimes}) \equiv \mathbb{G}_m(s_n, s_{-n}),$$

where $(s_n, s_{-n}) \in (q_n^{\otimes} \times q_{-n}^{\otimes})^{-1}(s_n^{\otimes}, s_{-n}^{\otimes})$.

Note that the payoffs in \mathbb{G}^{\otimes} are well-defined since they do not depend on the particular choice of $(s_n, s_{-n}) \in (q_n^{\otimes} \times q_{-n}^{\otimes})^{-1}(s_n^{\otimes}, s_{-n}^{\otimes})$.

4. GENERICITY ASSUMPTIONS

In Section 5 we fix a game G with game tree Γ and a zero-index component K and prove that K is not hyperstable provided it satisfies the following two assumptions:

- A.1 The outcome associated to K is unique. Moreover, after eliminating all branches and nodes from Γ which have zero probability under the equilibrium outcome induced by K, the set $q^{\otimes}(K)$ is a component of equilibria of G^{\otimes} .
- **A.2** If K induces a unique outcome in which player $n \in \{1,2\}$ has an excluded game, K^n is a full-dimensional polytope in Σ^n and every equilibrium payoff of \mathbb{G}^n to player n in K^n is strictly lower than player n's equilibrium payoff induced by K.





Assumption **A.1** is a standard property of generic extensive-form games (see Kreps and Wilson, 1982 or Govindan and Wilson, 2001). For our purposes, note that the definition of excluded game relies on the component inducing a unique outcome. The new assumption is **A.2**. The statement concerning the full-dimensionality of polytope K^n in Σ^n is satisfied generically.¹³ Assumption **A.2** as a whole tells us that K^n is the closure of an admissible neighborhood in Σ^n and, therefore, admits an index. It is a technical assumption in order to well-define the index of the neighborhood.

The next example shows that **A.1** does not imply **A.2**.

Example 4.1. Consider the game in Figure 4.

The game has two equilibrium outcomes, each of which is associated to two different equilibrium components. In the first one, player 1 plays T and player 2 plays ℓa . In the second, player 1 plays B and player 2 plays ra. This game satisfies **A.1** but does not satisfy **A.2**. Recall that in the excluded game \mathbb{G}^1 of player 1 associated to the first component, the non-deviating player 2 plays, at all information sets which are on equilibrium-path the exact equilibrium distribution - therefore, player 2 must play ℓ in the leftmost information set. In this excluded game, the deviating player 1 is a dummy player (she only has one strategy $S_1^1 = \{B\}$) and player 2 has two strategies (i.e., $S_2^1 = \{\ell a, \ell b\}$).

If player 2 plays ℓa , the payoff in the excluded game is (1,1), and if player 2 plays ℓb , the payoff is (3/2,0). Therefore, $(B, \ell a)$ is the obvious equilibrium of the excluded game \mathbb{G}^1 , which gives player 1 the equilibrium payoff of the component, violating **A.2**.

The next proposition implies that restricting to games that satisfy Assumptions **A.1** and **A.2** is a mild constraint.

¹³ For generically chosen terminal payoffs of a fixed game tree, any equilibrium component in mixed strategies induces a unique outcome and contains an equilibrium profile for which any observable deviation by any player is a strictly inferior reply to that profile (cf. the first paragraph of the proof of Theorem 4.2 in Govindan and Wilson, 2002). This guarantees full-dimensionality of K^n .

Definition 4.2. A subset $\mathscr{G}' \subseteq \mathscr{G}$ is *generic* if its complement is a lower-dimensional semialgebraic set.

Proposition 4.3. The subset $\mathscr{G}' \subseteq \mathbb{R}^{N|Z|}$ for which properties **A.1** and **A.2** are satisfied is generic.

Proof. It is well-known that Assumption A.1 is generic in the space of extensive-form games (Kreps and Wilson, 1982; Govindan and Wilson, 2001). The proof that generic two-player games satisfy Assumption A.2 is relegated to Online Appendix A.

5. PROOF OF THEOREM 2.5

It is already known that a non-zero index component is hyperstable. It can be proved by observing that equivalent components have the same index and that a component with non-zero index is robust to sufficiently small payoff perturbations.

¹⁴ Henceforth, we focus on showing that a hyperstable component must have non-zero index.

The strategy of the proof is to show that if component K has zero index then K is not hyperstable. Thus, we need to find a game $\overline{\mathbb{G}}$ equivalent to \mathbb{G} and a neighborhood \overline{V} of \overline{K} (the component in $\overline{\mathbb{G}}$ equivalent to K) such that for any $\varepsilon > 0$, there is an ε -payoff-perturbation $\overline{\mathbb{G}}^{\varepsilon}$ of $\overline{\mathbb{G}}$ with no equilibrium in \overline{V} . We proceed in three steps. In Step 1, we modify the original game \mathbb{G} to an auxiliary game to show that for some n = 1, 2, the supporting polytope K^n has zero index in the excluded game \mathbb{G}^n . In Step 2, we construct a game equivalent to \mathbb{G}^n and a perturbation so that the perturbed game does not have an equilibrium in the subset of the strategy profiles equivalent to K^n . This is used in Step 3 to construct a game equivalent to the original game \mathbb{G} and a perturbation that shows that the original zero index component K is not hyperstable. Prior to each step we present the key conceptual details about the step and overall strategy of the argument.

5.1. Step 1: For some player n, the supporting polytope has index 0 in \mathbb{G}^n

From Proposition 2.3 we know that at least one player has an excluded game. We proceed under the assumption that both players have an excluded game as the proof can be easily adapted to the case in which only one has it. Observe that K^n is full-dimensional in Σ^n and admissible with respect to \mathbb{G}^n (cf. A.2). Therefore, it has a well-defined index in \mathbb{G}^n . We also assume that for each player n = 1, 2 the supporting polytope K^n contains an equilibrium of

 $[\]overline{^{14}}$ See, e.g., GW.





 \mathbb{G}^n . If for some n = 1, 2 the excluded game \mathbb{G}^n does not have an equilibrium in K^n then such a supporting polytope has index 0 and we can move to Section 5.2.

We now prove that there exists n = 1, 2 such that K^n has index 0 in \mathbb{G}^n . The objective is to use the multiplication property of the index (cf. **I.4**) to express the index of K as the product of the indexes of K^1 and K^2 in their corresponding excluded games so that if the index of Kis zero then either the index of K^1 or the index of K^2 is zero as well. With that in mind, we perturb the original game so that each player is forced to play an observable deviation (i.e. an element of the strategy set in their excluded games) with vanishing probability. We then show that the non-deviating player can effectively best-reply by separately responding to the event in which the other player is forced to play as in her excluded game, and to the event in which the other player plays according to the equilibrium path defined by the component K. This defines a selection of the best-reply that, after some manipulations invoking the properties of the index, can be expressed as a product of the best replies in \mathbb{G}^{\otimes} , \mathbb{G}^1 , and \mathbb{G}^2 .

Let us define the auxiliary extensive-form game $G^{\varepsilon_1,\varepsilon_2}$ in Figure 5. First, Nature randomizes between three states, θ_0 with probability $1 - \varepsilon_1 - \varepsilon_2$, θ_1 with probability ε_1 , and θ_2 with probability ε_2 . For each n = 1, 2, player n observes if Nature has selected θ_n but cannot distinguish between θ_0 and θ_{-n} . If Nature chooses θ_n then player n chooses a member of S_n^n . In turn, if Nature chooses either θ_0 or θ_{-n} then player n, without observing player -n's move, chooses an element from S_n . Payoffs are inherited from the original game \mathbb{G} .

Nature, by choosing θ_n , forces player *n* to play from the same strategy set as in her excluded game. Nonetheless, players have their entire strategy set available after θ_0 . Note that game $G^{0,0}$ is equivalent to the original game G. We denote the component equivalent to

K in $G^{0,0}$ by $K^{0,0}$. Since $G^{0,0}$ is equivalent to G, component $K^{0,0}$ has also index 0 (cf. Property **I.2**). Fix a neighborhood $\mathcal{O}^{0,0}$ of $K^{0,0}$, such that every equilibrium of $G^{0,0}$ in the closure of $\mathcal{O}^{0,0}$ belongs to $K^{0,0}$. Since the index is locally constant with respect to payoff perturbations, taking $\varepsilon_1, \varepsilon_2 > 0$ sufficiently small implies that the neighborhood $\mathcal{O}^{0,0}$ is admissible for the game $G^{\varepsilon_1,\varepsilon_2}$ and that it has index 0. For notational convenience, we work with behavioral strategies. A behavioral strategy for player n in game $G^{\varepsilon_1,\varepsilon_2}$ is a pair $(\tau_n, \zeta_n^n) \in \Sigma_n \times \Sigma_n^n$.

Lemma 5.1. For $\varepsilon_1, \varepsilon_2 > 0$ sufficiently small, every equilibrium $(\tau_1, \zeta_1^1, \tau_2, \zeta_2^2) \in \mathcal{O}^{0,0}$ of game $G^{\varepsilon_1, \varepsilon_2}$ satisfies $\operatorname{Supp}(\tau_n) \cap S_n^n = \emptyset$ for each n = 1, 2. In addition, any $s_n^n \in S_n^n$ is a strictly inferior reply (after θ_0) to any such equilibrium.

Proof. The proof is divided in two cases. For the first, let $(\varepsilon_1^k)_k, (\varepsilon_2^k)_k$ be positive sequences converging to 0 and $(\tau_1^k, \zeta_1^{1,k}, \tau_2^k, \zeta_2^{2,k})_{k \in \mathbb{N}} \subseteq \mathcal{O}^{0,0}$ a sequence of equilibria of $G^{\varepsilon_1^k, \varepsilon_2^k}$. We prove that for k large enough $\operatorname{Supp}(\tau_1^k) \cap S_1^1 = \emptyset$ and that, after θ_0 , every member of S_1^1 is a strictly inferior reply to τ_2^k . Passing to subsequences if necessary, we can assume that the sequence of strategy profiles converges to $(\tau_1, \zeta_1^1, \tau_2, \zeta_2^2)$. In addition, writing $\tau_1^k = (1 - \alpha_k)\tau_1^{\star,k} + \alpha_k\tau_1^{1,k}$, with $\alpha_k \ge 0$, $\operatorname{Supp}(\tau_1^{\star,k}) \subseteq (S_1 \setminus S_1^1)$ and $\operatorname{Supp}(\tau_1^{1,k}) \subseteq S_1^1$, we also assume $\psi^k \equiv \varepsilon_1^k[\varepsilon_1^k + (1 - \varepsilon_1^k)\alpha_k]^{-1} \to \psi \in [0, 1]$. Observe that the support of τ_2 is a subset of $(S_2 \setminus S_2^2)$, since $(\tau_1, \zeta_1^1, \tau_2, \zeta_2^2) \in K^{0,0}$.

Claim 5.2. Suppose $Supp(\tilde{\tau}_2) \subseteq (S_2 \setminus S_2^2)$. If $q_2^{\otimes}(\tilde{\tau}_2)$ is equivalent to $q_2^{\otimes}(\tau_2)$ in \mathbb{G}^{\otimes} , then for n = 1, 2 we have $\mathbb{G}_n(\tau_1^{\times,k}, \tilde{\tau}_2) = \mathbb{G}_n(\tau_1^{\times,k}, \tau_2)$.

Proof of Claim 5.2. Note that $(\tau_1^k, \tau_2^k) \to (\tau_1, \tau_2) \in K$ so that $\operatorname{Supp}(\tau_2) \cap S_2^2 = \emptyset$. Since $q_2^{\otimes}(\tilde{\tau}_2)$ is equivalent to $q_2^{\otimes}(\tau_2)$ in \mathbb{G}^{\otimes} both $q_2^{\otimes}(\tilde{\tau}_2)$ and $q_2^{\otimes}(\tau_2)$ give the same payoff against any player 1's strategy τ_1^{\otimes} in \mathbb{G}^{\otimes} . Given $\tau_1^{\times,k} \in \Delta(S_1 \setminus S_1^1)$, both $q_2^{\otimes}(\tilde{\tau}_2)$ and $q_2^{\otimes}(\tau_2)$ yield the same payoff against $q_1^{\otimes}(\tau_1^{\times,k})$. Therefore, $\mathbb{G}_n(\tau_1^{\times,k}, \tilde{\tau}_2) = \mathbb{G}_n^{\otimes}(q^{\otimes}(\tilde{\tau}_2), q_1^{\otimes}(\tau_1^{\times,k})) = \mathbb{G}_n^{\otimes}(q^{\otimes}(\tau_2), q_1^{\otimes}(\tau_1^{\times,k})) = \mathbb{G}_n(\tau_1^{\times,k}, \tau_2)$.

For *k* sufficiently large, $\text{Supp}(\tau_2) \subseteq \text{Supp}(\tau_2^k)$. For such a *k*, the limit τ_2 maximizes player 2's expected payoff at her information set following θ_0 and θ_1 . For $x_2 \in \Sigma_2$, player 2's expected payoff is

(4)
$$[\varepsilon_1^k + (1 - \varepsilon_1^k - \varepsilon_2^k)\alpha_k] \Big[\psi^k \mathbb{G}_2(\zeta_1^{1,k}, x_2) + (1 - \psi^k) \mathbb{G}_2(\tau_1^{1,k}, x_2) \Big] + (1 - \varepsilon_1^k - \varepsilon_2^k)(1 - \alpha_k) \Big[\mathbb{G}_2(\tau_1^{\times, k}, x_2) \Big].$$

Let $A_2 \subset \Sigma_2$ be defined as: $\tilde{\tau}_2 \in A_2$ if and only if $q_2^{\otimes}(\tilde{\tau}_2)$ is equivalent to $q_2^{\otimes}(\tau_2)$ in \mathbb{G}^{\otimes} . Since τ_2 maximizes expression (4), Claim 5.2 implies that τ_2 also solves

$$\max_{x_2 \in A_2} \psi^k \mathbb{G}_2(\zeta_1^{1,k}, x_2) + (1 - \psi^k) \mathbb{G}_2(\tau_1^{1,k}, x_2).$$

From Lemma 3.1, player 2 has a product strategy $\tau_2^{\otimes} * \tau_2^0$ that is equivalent to τ_2 in G. Since $(\tau_1, \tau_2) \in K$, then τ_2^{\otimes} is equivalent to $\hat{\tau}_2^{\otimes}$ in \mathbb{G}^{\otimes} . By the definition of payoff functions in player 1's excluded game, τ_2^0 solves

$$\max_{x_2^0 \in \Sigma_2^0} \psi^k \mathbb{G}_2^1(\zeta_1^{1,k}, x_2^0) + (1 - \psi^k) \mathbb{G}_2^1(\tau_1^{1,k}, x_2^0).$$

That is, for k large enough, τ_2^0 is a best-reply against $\psi^k \zeta_1^{1,k} + (1 - \psi^k) \tau_1^{1,k}$ in player 1's excluded game \mathbb{G}^1 . Of course, τ_2^0 is also a best-reply against the limit $\psi \zeta_1^1 + (1 - \psi) \tau_1^1$.

In turn, $\zeta_1^{1,k}$ is a best-reply (when player 1's strategy is constrained to Σ_1^1) against τ_2^k in \mathbb{G} for every k. By continuity, ζ_1^1 is a constrained best-reply against τ_2 in \mathbb{G} and, therefore, also against the equivalent product strategy $\tau_2^{\otimes} * \tau_2^0$, which implies that ζ_1^1 is a best-reply against τ_2^0 in \mathbb{G}^1 . Using an analogous argument, if $\alpha_k > 0$ for all k, we also conclude that τ_1^1 is a best-reply against τ_2^0 in \mathbb{G}^1 . If we do not have $\alpha_k > 0$ for all k then either there exists a subsequence for which it holds or $\psi = 1$. In either case, $(\tau_2^0, \psi \zeta_1^1 + (1 - \psi)\tau_1^1)$ is an equilibrium of \mathbb{G}^1 .

From genericity property **A.2**, no equilibrium of \mathbb{G}^1 gives the same payoff as the equilibrium outcome Q to player 1. Therefore, $\tau_2^0 \notin \tilde{\partial} K_2$ and, for every $s_1^1 \in S_1^1$ we have $\mathbb{G}_1(s_1^1, \tau_2) = \mathbb{G}_1^1(s_1^1, \tau_2^0) < G_1(Q)$. We conclude that for sufficiently large k, every member of S_1^1 is a strictly inferior reply to τ_2^k . So for k large enough, $\operatorname{Supp}(\tau_1^k) \subseteq S_1 \setminus S_1^1$ which finishes the proof of the first case.

For the second case, consider positive sequences $(\varepsilon_i^k)_k$, i = 1, 2 converging to 0 and $(\sigma^k)_{k \in \mathbb{N}} \subset \mathcal{O}^{0,0}$ a sequence of equilibria of $G^{\varepsilon^{k_1}, \varepsilon_2^k}$. For sufficiently large k, using an argument symmetric to the first case, $\operatorname{Supp}(\tau_2^k) \subseteq (S_2 \setminus S_2^2)$ and S_2^2 is a set of strictly inferior replies after θ_0 to τ_1^k . From these two cases, for sufficiently small $\varepsilon_1, \varepsilon_2 > 0$, Lemma 5.1 follows.

Take $\varepsilon_1, \varepsilon_2 > 0$ as small as prescribed by Lemma 5.1. Every equilibrium of $G^{\varepsilon_1, \varepsilon_2}$ in $\mathcal{O}^{0,0}$ is such that any strategy that assigns positive probability to some member of S_n^n after θ_0 is a strictly inferior reply to the equilibrium. With abuse of notation, denote also by $G^{\varepsilon_1, \varepsilon_2}$ the game that is obtained by eliminating S_1^1 and S_2^2 from players' actions at their information set after θ_0 . Since they are always inferior replies, their elimination does not affect the index of $\mathcal{O}^{0,0}$.

For $\delta > 0$, define

$$U_n^{\delta} \equiv \Big\{ \tau_n \in \Delta(S_n \setminus S_n^n) \mid \mathbb{G}_{-n}(s_{-n}^{-n}, \tau_n) < G_{-n}(Q) - \delta \text{ for all } s_{-n}^{-n} \in S_{-n}^{-n} \Big\}.$$

Since the component K induces a unique distribution Q, the component that induces Q in \mathbb{G}^{\otimes} is a rectangle $R_1 \times R_2$ with $R_n \subseteq \Sigma_n^{\otimes}$.¹⁵ From this we obtain a δ -neighborhood $B_n^{\delta} \subseteq \Sigma_n^{\otimes}$ of R_n such that $B^{\delta} \equiv B_1^{\delta} \times B_2^{\delta}$ is a neighborhood of $R_1 \times R_2$ and does not contain any other equilibrium of \mathbb{G}^{\otimes} . Let $V_1^{\delta} \equiv (q_1^{\otimes})^{-1}(B_1^{\delta})$, $V_2^{\delta} \equiv (q_2^{\otimes})^{-1}(B_2^{\delta})$, and $V^{\delta} \equiv V_1^{\delta} \times V_2^{\delta}$. Fix in addition $\delta > 0$ sufficiently small such that for both n = 1, 2 all equilibria of player n's excluded game in K^n induce a payoff smaller than $G_n(Q) - \delta$ (such $\delta > 0$ exists due to Assumption A.2).

Claim 5.3. For $\varepsilon_1, \varepsilon_2 > 0$ sufficiently small,

- (1) $\mathcal{O}^{\delta} \equiv ([V_1^{\delta} \cap U_1^{\delta}] \times \Sigma_1^1) \times ([V_2^{\delta} \cap U_2^{\delta}] \times \Sigma_2^2)$ is admissible for $G^{\varepsilon_1, \varepsilon_2}$;
- (2) The symmetric difference $\mathcal{O}^{\delta} \Delta \mathcal{O}^{0,0}$ contains no equilibrium of $G^{\varepsilon_1,\varepsilon_2}$.

Proof. We start by proving admissibility of \mathcal{O}^{δ} for $G^{\varepsilon_1,\varepsilon_2}$. Let $(\varepsilon_n^k)_{k\in\mathbb{N}}$ converge to 0 for n = 1,2 such that (ζ^k, τ^k) is an equilibrium of $G^{\varepsilon_1^k,\varepsilon_2^k}$ in $\operatorname{cl}(\mathcal{O}^{\delta})$, where the closure is taken with respect to $(\times_{n=1,2}\Delta(S_n \setminus S_n^n) \times \Sigma_1^1 \times \Sigma_2^2)$. Consider a limit (ζ, τ) . By **A.1**, $q^{\otimes}(\tau)$ is an equilibrium of \mathbb{G}^{\otimes} that induces Q. Therefore, for k large enough, $q^{\otimes}(\tau^k) \in B^{\delta}$, which implies $\tau^k \in V^{\delta}$. Recall that $\tau_n \in \Delta(S_n \setminus S_n^n)$ and take a product strategy $\tau_n^{\otimes} * \tau_n^0$ equivalent to τ_n in \mathbb{G} . We claim that $(\zeta_{-n}^{-n}, \tau_n^0)$ is an equilibrium of the excluded game \mathbb{G}^{-n} . It is clear that ζ_{-n}^{-n} is a best-reply of player -n against τ_n^0 in \mathbb{G}^{-n} . It remains to show that τ_n^0 is a best-reply of player n against ζ_{-n}^{-n} in \mathbb{G}^{-n} . Since τ_n^k is an equilibrium strategy in $G^{\varepsilon_1^k, \varepsilon_2^k}$, it solves

(5)
$$\max_{x_n\in\Delta(S_n\setminus S_n^n)}\varepsilon_{-n}^k\mathbb{G}_n(x_n,\zeta_{-n}^{-n,k})+(1-\varepsilon_1^k-\varepsilon_2^k)\mathbb{G}_n(x_n,\tau_{-n}^k).$$

For k sufficiently large, $\operatorname{Supp}(\tau_n) \subseteq \operatorname{Supp}(\tau_n^k)$, so τ_n is also a solution to (5). Note that for every $x_n \in \Delta(S_n \setminus S_n^n)$ we have $\mathbb{G}_n(x_n, \tau_{-n}^k) = \mathbb{G}_n^{\otimes}(q_n^{\otimes}(x_n), q_{-n}^{\otimes}(\tau_{-n}^k))$. Therefore, since τ_n^{\otimes} is equivalent to $\hat{\tau}_n^{\otimes}$ in \mathbb{G}^{\otimes} , we obtain $\mathbb{G}_n(\tau_n, \zeta_{-n}^{-n,k}) = \mathbb{G}_n^{-n}(\tau_n^0, \zeta_{-n}^{-n,k})$. Taking limits in k, we obtain that τ_n^0 is a best-reply against ζ_{-n}^{-n} in \mathbb{G}^{-n} . That is, $(\tau_n^0, \zeta_{-n}^{-n})$ is an equilibrium of \mathbb{G}^{-n} as claimed. This implies $\mathbb{G}_{-n}^{-n}(\tau_n^0, \zeta_{-n}^{-n}) < \mathbb{G}_{-n}(Q) - \delta$, by Assumption A.2 and our choice of δ . Hence, $\mathbb{G}_{-n}(\tau_n, \zeta_{-n}^{-n}) = \mathbb{G}_{-n}^{-n}(\tau_n^0, \zeta_{-n}^{-n}) < \mathbb{G}_{-n}(Q) - \delta$. That is, $\tau_n^k \in U_n^{\delta}$ for sufficiently large k. This proves that for $\varepsilon_1, \varepsilon_2 > 0$ small enough \mathcal{O}^{δ} is admissible for $G^{\varepsilon_1, \varepsilon_2}$.

We now prove that for sufficiently small $\varepsilon_1, \varepsilon_2 > 0$, \mathcal{O}^{δ} contains the equilibria of $G^{\varepsilon_1, \varepsilon_2}$ that are contained in $\mathcal{O}^{0,0}$. It is sufficient to prove that for any convergent sequence $(\zeta^k, \tau^k)_{k \in \mathbb{N}} \subset \mathcal{O}^{0,0}$ of equilibria $G^{\varepsilon_1^k, \varepsilon_2^k}$ with limit (ζ, τ) and ε_n^k converging to 0 for n = 1, 2, the sequence is eventually in \mathcal{O}^{δ} . First observe that the limit must satisfy $\mathbb{G}_{-n}(\zeta_{-n}^{-n}, \tau_n) \leq G_{-n}(Q)$. Take

¹⁵ The game \mathbb{G}^{\otimes} is the normal form of the extensive-form game obtained from *G* by eliminating nodes and branches that have zero probability under *Q*. Therefore, the outcome *Q* is completely mixed in this extensiveform and is induced by a unique component of mixed strategies in \mathbb{G}^{\otimes} . Hence, as in the proof of Lemma 2.3, *Q* is induced by a unique enabling strategy profile in the enabling-form of \mathbb{G}^{\otimes} . (See Appending A for the definition of enabling strategies and enabling-form.) This implies that the component in \mathbb{G}^{\otimes} inducing *Q* in \mathbb{G}^{\otimes} is a rectangle.

a product strategy $\tau_n^{\otimes} * \tau_n^0$ equivalent to τ_n in \mathbb{G} . We claim that $(\zeta_{-n}^{-n}, \tau_n^0)$ is an equilibrium of the excluded game \mathbb{G}^{-n} . It is clear that ζ_{-n}^{-n} is a best-reply of player -n against τ_n^0 in \mathbb{G}^{-n} . It remains to show that τ_n^0 is a best-reply of player n against ζ_{-n}^{-n} in \mathbb{G}^{-n} . Since τ_n^k is an equilibrium strategy in $G^{\varepsilon_1^k, \varepsilon_2^k}$, it solves (5). Continuing the argument in the same fashion as in the first part, we obtain that $(\tau_n^0, \zeta_{-n}^{-n})$ is an equilibrium of \mathbb{G}^{-n} . This implies $\mathbb{G}_{-n}^{-n}(\tau_n^0, \zeta_{-n}^{-n}) < \mathbb{G}_{-n}(Q) - \delta$, by Assumption A.2 and our choice of δ . Hence, $\mathbb{G}_{-n}(\tau_n, \zeta_{-n}^{-n}) =$ $\mathbb{G}_{-n}^{-n}(\tau_n^0, \zeta_{-n}^{-n}) < \mathbb{G}_{-n}(Q) - \delta$. Thus (ζ^k, τ^k) belongs to \mathcal{O}^{δ} for sufficiently large k.

The proof of the result that for sufficiently small $\varepsilon_1, \varepsilon_2 > 0$, $\mathcal{O}^{0,0}$ contains the equilibria of $G^{\varepsilon_1, \varepsilon_2}$ that are contained in \mathcal{O}^{δ} can be obtained from noticing that any limit (ζ, τ) of a sequence of equilibria in \mathcal{O}^{δ} satisfies $\tau \in V^{\delta}$ and $\tau \in U^{\delta}$.

We now construct an expression for the index of \mathcal{O}^{δ} that allows us to conclude that for some n = 1, 2, the supporting polytope K^n has zero index with respect to \mathbb{G}^n . When player nbest-replies against $(\tau_{-n}, \zeta_{-n}^{-n})$ in game $G^{\varepsilon_1, \varepsilon_2}$ at her information set after θ_0 , player n chooses $\tau_n \in \Delta(S_n \setminus S_n^n)$ that maximizes

(6)
$$\varepsilon_{-n}\mathbb{G}_n(\tau_n,\zeta_{-n}^{-n}) + (1-\varepsilon_1-\varepsilon_2)\mathbb{G}_n(\tau_n,\tau_{-n}).$$

Since actions in S_n^n , n = 1, 2 are no longer available after θ_0 , the function $\mathbb{G}_n(\tau_n, \tau_{-n})$ only depends on $q_n^{\otimes}(\tau_n)$ and $q_{-n}^{\otimes}(\tau_{-n})$. In order to compute the best-reply τ_n , player n only requires ζ_{-n}^{-n} and τ_{-n}^{\otimes} . The correspondence that assigns to each $(\zeta_{-n}^{-n}, \tau_{-n}^{\otimes})$ the optimal $\tau_n \in \Delta(S_n \setminus S_n^n)$ is

$$\mathrm{BR}_n^{\otimes,\varepsilon_1,\varepsilon_2}:\Sigma_{-n}^{-n}\times\Sigma_{-n}^{\otimes}\Longrightarrow\Delta(S_n\setminus S_n^n).$$

In addition, define the correspondence $BR_n^{\times} : \Delta(S_{-n} \setminus S_{-n}^{-n}) \Rightarrow \Sigma_n^n$ by:

$$\mathrm{BR}_n^{\times}(\tau_{-n}) \equiv \arg\max_{\zeta_n^n \in \Delta(S_n^n)} \mathbb{G}_n(\tau_{-n},\zeta_n^n).$$

Define $BR^{\times} \equiv BR_1^{\times} \times BR_2^{\times}$. Let id_n be the identity in Σ_n^n . Player *n*'s best-reply correspondence $BR_n^{\varepsilon_1,\varepsilon_2}$ in game $G^{\varepsilon_1,\varepsilon_2}$ is

$$\mathrm{BR}_n^{\times} \times \left[\mathrm{BR}_n^{\otimes,\varepsilon_1,\varepsilon_2} \circ (\mathrm{id}_{-n} \times q_{-n}^{\otimes})\right] : \Sigma_{-n}^{-n} \times \Delta(S_{-n} \setminus S_{-n}^{-n}) \Longrightarrow \Sigma_n^n \times \Delta(S_n \setminus S_n^n).$$

And the best-reply correspondence in $G^{\varepsilon_1,\varepsilon_2}$ is $BR^{\varepsilon_1,\varepsilon_2} \equiv BR_1^{\varepsilon_1,\varepsilon_2} \times BR_2^{\varepsilon_1,\varepsilon_2}$.

Instead of using $BR^{\varepsilon_1,\varepsilon_2}$ directly to compute the index of \mathcal{O}^{δ} we use a selection of $BR^{\varepsilon_1,\varepsilon_2}$. That is, a new correspondence whose graph is a subset of the graph of $BR^{\varepsilon_1,\varepsilon_2}$. To define it, consider first the correspondence $\phi_n^{\varepsilon_1,\varepsilon_2}: \Sigma_{-n}^{-n} \times \Sigma_{-n}^{\otimes} \Rightarrow \Sigma_n^{\otimes} \times \Sigma_n^0$ that assigns to each $(\zeta_{-n}^{-n}, \tau_{-n}^{\otimes}) \in \mathbb{C}$ $\Sigma_{-n}^{-n} \times \Sigma_{-n}^{\otimes}$ the set of pairs $(\tau_n^{\otimes}, \tau_n^0) \in \Sigma_n^{\otimes} \times \Sigma_n^0$ that satisfy

(7)
$$\tau_n^{\otimes} \in q_n^{\otimes} \circ \mathrm{BR}_n^{\otimes, \varepsilon_1, \varepsilon_2}(\zeta_{-n}^{-n}, \tau_{-n}^{\otimes}), \text{ and }$$

(8)
$$\tau_n^0 \in \arg\max_{\tilde{\tau}_n^0 \in \Sigma_n^0} \mathbb{G}_n(\hat{\tau}_n^{\otimes} * \tilde{\tau}_n^0, \zeta_{-n}^{-n}).$$

Let $e_n : \Sigma_n^{\otimes} \times \Sigma_n^0 \to \Delta(S_n \setminus S_n^n)$ be defined by $e_n(\tau_n^{\otimes}, \tau_n^0) \equiv \tau_n^{\otimes} * \tau_n^0$. Define also $w_n : \Sigma_n^n \times \Delta(S_n \setminus S_n^n) \to \Sigma_n^n \times \Sigma_n^{\otimes} \times \Sigma_n^0$ by $w_n(\zeta_n^n, \tau_n) \equiv (\zeta_n^n, q_n^{\otimes}(\tau_n), q_n^0(\tau_n))$. As usual, $e \equiv e_1 \times e_2$, $w \equiv w_1 \times w_2$, and $\phi^{\varepsilon_1, \varepsilon_2} \equiv \phi_1^{\varepsilon_1, \varepsilon_2} \times \phi_2^{\varepsilon_1, \varepsilon_2}$. Consider the correspondence

$$\Phi^{\varepsilon_1,\varepsilon_2} \equiv (e \circ \phi^{\varepsilon_1,\varepsilon_2} \circ w) : \Sigma_1^1 \times \Sigma_2^2 \times \Delta(S_1 \setminus S_1^1) \times \Delta(S_2 \setminus S_2^2) \Rightarrow \Delta(S_1 \setminus S_1^1) \times \Delta(S_2 \setminus S_2^2).$$

Lemma 5.4. $\Phi^{\varepsilon_1, \varepsilon_2}$ is a nonempty, contractible, and compact valued, upper hemicontinuous correspondence.

Proof. The fact that $\Phi^{\varepsilon_1,\varepsilon_2}$ is nonempty, compact-valued, and upper hemicontinuous follows from standard arguments. We show that values are contractible. Fix (τ,ζ) and some $\check{\tau} \in \Phi^{\varepsilon_1,\varepsilon_2}(\tau,\zeta)$. Then $\check{\tau} = \check{\tau}^{\otimes} * \check{\tau}^0$. Consider the homotopy $J = J_1 \times J_2 : \Phi^{\varepsilon_1,\varepsilon_2}(\tau,\zeta) \times [0,1] \rightarrow \Phi^{\varepsilon_1,\varepsilon_2}(\tau,\zeta)$ that for each $\tilde{\tau} = (\tilde{\tau}_1,\tilde{\tau}_2) = (\tilde{\tau}_1^{\otimes} * \tilde{\tau}_1^0, \tilde{\tau}_2^{\otimes} * \tilde{\tau}_2^0) \in \Phi^{\varepsilon_1,\varepsilon_2}(\tau,\zeta)$ is defined by $J_n(t,\tilde{\tau}_n) = [(1-t)\tilde{\tau}_n^{\otimes} + t\check{\tau}_n^{\otimes}] * [(1-t)\tilde{\tau}_n^0 + t\check{\tau}_n^0]$ for each n = 1,2. The homotopy is well-defined. Moreover, $J_n(0, \tilde{\tau}_n^{\otimes} * \tilde{\tau}_n^0) = \tilde{\tau}_n^{\otimes} * \tilde{\tau}_n^0) = \check{\tau}_n^{\otimes} * \check{\tau}_n^0$, which shows that $\Phi^{\varepsilon_1,\varepsilon_2}(\tau,\zeta)$ is contractible valued.

Lemma 5.5. $BR^{\times} \times \Phi^{\varepsilon_1, \varepsilon_2}$ is a selection of $BR^{\varepsilon_1, \varepsilon_2}$.

Proof. Let $(\check{\zeta},\check{\tau}) \in [BR^{\times} \times \Phi^{\varepsilon_1,\varepsilon_2}](\zeta,\tau)$. Since $\check{\zeta} \in BR^{\times}(\zeta,\tau)$ we only need to show $\check{\tau}_n \in BR_n^{\otimes,\varepsilon_1,\varepsilon_2}(\zeta_{-n}^{-n},q_{-n}^{\otimes}(\tau))$ for n = 1, 2. By definition, $\check{\tau}_n = \check{\tau}_n^{\otimes} * \check{\tau}_n^0$, where

$$\check{\tau}_n^{\otimes} \in q_n^{\otimes} \left[\mathrm{BR}_n^{\otimes,\varepsilon_1,\varepsilon_2}(\zeta_{-n}^{-n}, q_{-n}^{\otimes}(\tau_{-n})) \right].$$

Therefore, there exists

$$\tau'_{n} \in \arg\max_{\tilde{\tau}_{n}} \left\{ \varepsilon_{-n} \mathbb{G}_{n}(\tilde{\tau}_{n}, \zeta_{-n}^{-n}) + (1 - \varepsilon_{1} - \varepsilon_{2}) \mathbb{G}_{n}(q_{n}^{\otimes}(\tilde{\tau}_{n}), q_{-n}^{\otimes}(\tau_{-n})) \right\}$$

such that $q_n^{\otimes}(\tau'_n) = \check{\tau}_n^{\otimes}$. Take $\tau_n^* = \tau_n^{*,\otimes} * \tau_n^{*,0}$ equivalent to τ'_n in G. Then $\check{\tau}_n^{\otimes}$ must be equivalent to $\tau_n^{*,\otimes}$ in \mathbb{G}^{\otimes} . Recall $\operatorname{Supp}(\tau_n^{*,\otimes}) \subseteq \operatorname{Supp}(\hat{\tau}_n^{\otimes})$ because $\hat{\tau}_n^{\otimes}$ is completely mixed. Hence,

$$\check{\tau}_{n}^{0} \in \arg\max_{\tau_{n}^{0}} \sum_{s_{n}^{\otimes}, s_{n}^{0}} \hat{\tau}_{n}^{\otimes}(s_{n}^{\otimes}) \tau_{n}^{0}(s_{n}^{0}) \mathbb{G}_{n}(s_{n}^{\otimes}, s_{n}^{0}, \zeta_{-n}^{-n}) \subseteq \arg\max_{\tau_{n}^{0}} \sum_{s_{n}^{\otimes}, s_{n}^{0}} \tau_{n}^{*, \otimes}(s_{n}^{\otimes}) \tau_{n}^{0}(s_{n}^{0}) \mathbb{G}_{n}(s_{n}^{\otimes}, s_{n}^{0}, \zeta_{-n}^{-n}).$$

That is,

$$\check{\tau}_n^{\otimes} * \check{\tau}_n^0 \in \arg\max_{\tilde{\tau}_n} \Big\{ \varepsilon_{-n} \mathbb{G}_n(\tilde{\tau}_n, \zeta_{-n}^{-n}) + (1 - \varepsilon_1 - \varepsilon_2) \mathbb{G}_n(q_n^{\otimes}(\tilde{\tau}_n), q_{-n}^{\otimes}(\tau_{-n})) \Big\},\$$

which shows $\check{\tau}_n \in \mathrm{BR}_n^{\otimes,\varepsilon_1,\varepsilon_2}(\zeta_{-n}^{-n},q_{-n}^{\otimes}(\tau_{-n}))$ for n = 1,2 and, therefore, $(\check{\zeta},\check{\tau}) \in \mathrm{BR}^{\varepsilon_1,\varepsilon_2}(\zeta,\tau)$. \Box

Letting BR^{\otimes} be the best-reply correspondence in the game \mathbb{G}^{\otimes} , let us define $\Phi \equiv \Phi^{0,0} = e \circ (BR_2^1 \times BR^{\otimes} \times BR_1^2) \circ w$. Similarly to $\Phi^{\varepsilon_1, \varepsilon_2}$, standard arguments show that Φ is upper hemicontinuous as well as nonempty, compact, and contractible valued.

Lemma 5.6. The following statements hold:

- (1) The neighborhood \mathcal{O}^{δ} is admissible for BR[×] × Φ .
- (2) For ε_1 , $\varepsilon_2 > 0$ sufficiently small, \mathcal{O}^{δ} is admissible for BR[×] × $\Phi^{\varepsilon_1, \varepsilon_2}$.
- (3) For ε_1 , $\varepsilon_2 > 0$ sufficiently small, $BR^{\times} \times \Phi^{\varepsilon_1, \varepsilon_2}$ assigns the same index to \mathcal{O}^{δ} as $BR^{\times} \times \Phi$ does.

Proof. To prove (1), let (ζ, τ) be a fixed point of $BR^{\times} \times \Phi$, where $(\zeta, \tau) \in cl(\mathcal{O}^{\delta})$. Then $\tau_n = \tau_n^{\otimes} * \tau_n^0$ and $\tau_n^{\otimes} \in B^{\delta}$. This implies that τ_n^{\otimes} is equivalent to $\hat{\tau}_n^{\otimes}$ in \mathbb{G}^{\otimes} . Therefore, $\zeta_{-n}^{-n} \in BR_{-n}^{\times}(\tau_n) = BR_{-n}^{-n}(\tau_n^0)$. Moreover, by construction, $\tau_n^0 \in BR_n^{-n}(\zeta_{-n}^{-n})$. Hence, $(\tau_n^0, \zeta_{-n}^{-n})$ is an equilibrium of \mathbb{G}^{-n} . Since $\mathbb{G}_{-n}(\zeta_{-n}^{-n}, \tau_n) \leq G_{-n}(Q) - \delta$ we obtain $\mathbb{G}_{-n}^{-n}(\zeta_{-n}^{-n}, \tau_n^0) \leq \mathbb{G}_{-n}(Q) - \delta$. By **A.2** and our choice of δ , we have $\mathbb{G}_{-n}^{-n}(\zeta_{-n}^{-n}, \tau_n^0) < G_{-n}(Q) - \delta$ and finally $\mathbb{G}_{-n}(s_{-n}^{-n}, \tau_n) < G_{-n}(Q) - \delta$ for every $s_{-n}^{-n} \in S_{-n}^{-n}$. That is, $(\zeta, \tau) \in \mathcal{O}^{\delta}$, which proves admissibility of \mathcal{O}^{δ} .

To prove (2), fix a neighborhood W of the graph of $BR^{\times} \times \Phi$ such that $W \cap \{(\tau, \zeta, \tau, \zeta) \mid (\tau, \zeta) \in \operatorname{cl}(\mathcal{O}^{\delta}) \setminus \mathcal{O}^{\delta}\} = \emptyset$. For $\eta > 0$, define the set of η -replies against $\tau^{\otimes} \in \Sigma^{\otimes}$ by $BR^{\otimes,\eta}(\tau^{\otimes}) \equiv \{\tilde{\tau}^{\otimes} \in \Sigma^{\otimes} \mid \mathbb{G}_{n}^{\otimes}(\tilde{\tau}_{n}^{\otimes}, \tau_{-n}^{\otimes}) \geq \max_{s_{n}^{\otimes}}\mathbb{G}_{n}^{\otimes}(s_{n}^{\otimes}, \tau_{-n}^{\otimes}) - \eta$ for $n = 1, 2\}$. Fix $\eta > 0$ sufficiently small such that $\operatorname{Graph}(BR^{\times} \times [e \circ (BR_{2}^{1} \times BR^{\otimes,\eta} \times BR_{1}^{2}) \circ w]) \subseteq W$. Take now $\varepsilon_{1}, \varepsilon_{2} > 0$ sufficiently small such that for each $n = 1, 2, s_{n} = (s_{n}^{\otimes}, s_{n}^{0})$ and s_{-n}^{-n} ,

$$\varepsilon_{-n}\mathbb{G}_n(s_n, s_{-n}^{-n}) + (1 - \varepsilon_1 - \varepsilon_2)\mathbb{G}_n^{\otimes}(s^{\otimes}) \ge \mathbb{G}_n^{\otimes}(s^{\otimes}) - \eta.$$

Then, $(\tilde{\zeta}, \tilde{\tau}) \in [BR^{\times} \times \Phi^{\varepsilon_1, \varepsilon_2}](\zeta, \tau)$ implies $(\zeta, \tau, \tilde{\zeta}, \tilde{\tau}) \in \operatorname{Graph}(BR^{\times} \times e \circ [BR_2^1 \times BR^{\otimes, \eta} \times BR_1^2] \circ w)$, which implies $\operatorname{Graph}(BR^{\times} \times \Phi^{\varepsilon_1, \varepsilon_2}) \subseteq W$. This shows that, for sufficiently small $\varepsilon_1, \varepsilon_2 > 0$, the neighbourhood \mathcal{O}^{δ} is admissible for $BR^{\times} \times \Phi^{\varepsilon_1, \varepsilon_2}$.

To prove (3), consider the neighborhood W defined in the previous paragraph. There exists then a neighborhood $V \subseteq W$ of Graph $(BR^{\times} \times \Phi)$ such that any two continuous maps whose graphs are contained in V are homotopic by a homotopy contained in W. By our choice of W, any two such maps assign the same index to \mathcal{O}^{δ} . Using a reasoning analogous to the one of the previous item, take then ε_1 and ε_2 sufficiently small so that Graph $(BR^{\times} \times \Phi^{\varepsilon_1, \varepsilon_2}) \subseteq V$. Approximate $BR^{\times} \times \Phi^{\varepsilon_1, \varepsilon_2}$ by a continuous function f with Graph $(f) \subset V$ such that the index it assigns to \mathcal{O}^{δ} determines the index of \mathcal{O}^{δ} with respect to $BR^{\times} \times \Phi^{\varepsilon_1, \varepsilon_2}$. This index is then equal to the index of \mathcal{O}^{δ} with respect to $BR^{\times} \times \Phi$. This concludes the proof.

Let
$$h_n(\tau_n, t) = (1 - t)\tau_n + t[\hat{\tau}_n^{\otimes} * q_n^0(\tau_n)]$$
 and define

$$H: \left(\prod_{n=1,2} \left[\Delta(S_n \setminus S_n^n) \times \Sigma_n^n\right]\right) \times [0, 1] \Rightarrow \prod_{n=1,2} \left[\Delta(S_n \setminus S_n^n) \times \Sigma_n^n\right] \text{ by}$$

$$H(\tau, \zeta, t) = BR^{\times}(h(\tau, t)) \times \Phi(\tau, \zeta) \text{ for every } (\tau, \zeta, t).$$

The map *H* is a homotopy of correspondences. It is a nonempty, contractible, compact valued, upper hemicontinuous correspondence. Letting $\bar{e}(\zeta, \tau^{\otimes}, \tau^{0}) = (\zeta, e(\tau^{\otimes}, \tau^{0}))$, observe that for t = 1, the homotopy equals $\bar{e} \circ (BR^{1} \times BR^{\otimes} \times BR^{2}) \circ w$ and for t = 0 it equals $BR^{\times} \times \Phi$.

Lemma 5.7. \mathcal{O}^{δ} is admissible for H and therefore $\bar{e} \circ (BR^1 \times BR^* \times BR^2) \circ w$ assigns the same index to \mathcal{O}^{δ} as $BR^* \times \Phi$.

Proof. Let $(\tau,\zeta) \in \operatorname{cl}(\mathcal{O}^{\delta})$ satisfy $(\tau,\zeta) \in \operatorname{BR}^{\times}(h(\tau,t)) \times \Phi(\tau,\zeta)$ for some $t \in [0,1]$. Then $\tau_n = \tau_n^{\otimes} * \tau_n^0$. First note that $\tau \in V^{\delta}$, since $\tau_n \in \Phi_n(\tau,\zeta)$ and the strategy τ_n^{\otimes} is equivalent to $\hat{\tau}_n^{\otimes}$. It follows $\zeta_n^n \in \operatorname{BR}_n^{\times}(h_{-n}(\tau_{-n},t)) = \operatorname{BR}_n^{\times}(\hat{\tau}_{-n}^{\otimes} * \tau_{-n}^0) = \operatorname{BR}_{-n}^n(\tau_{-n}^0)$. On the other hand, $\tau_{-n}^0 \in \operatorname{BR}_{-n}^n(\zeta_n^n)$ so that (ζ_n^n, τ_{-n}^0) is an equilibrium of \mathbb{G}^n . By **A.2** and our choice of δ we have $\mathbb{G}_n^n(\zeta_n^n, \tau_{-n}^0) < G_n(Q) - \delta$. Therefore, $\mathbb{G}_n(s_n^n, \tau_{-n}^0) < G_n(Q) - \delta$ for every s_n^n . This proves $(\zeta, \tau) \in \mathcal{O}^{\delta}$.

To finish Step 1, let

$$\hat{U}_{-n}^{\delta} \equiv \left\{ \tau_{-n}^{0} \in \Sigma_{-n}^{0} \mid \mathbb{G}_{n}^{n}(s_{n}^{n}, \tau_{-n}^{0}) < G_{n}(Q) - \delta \text{ for all } s_{n}^{n} \in S_{n}^{n} \right\} \text{ and } \hat{\mathcal{O}}^{\delta} \equiv w^{-1} \Big[\Sigma_{1}^{1} \times \Sigma_{2}^{2} \times B^{\delta} \times \hat{U}^{\delta} \Big].$$

Lemma 5.8. $\hat{\mathcal{O}}^{\delta}$ is admissible for $\bar{e} \circ (BR^1 \times BR^{\otimes} \times BR^2) \circ w$ and contains the same fixed points of $\bar{e} \circ (BR^1 \times BR^{\otimes} \times BR^2) \circ w$ as \mathcal{O}^{δ} .

Proof. Let $(\tau,\zeta) \in \operatorname{cl}(\hat{\mathcal{O}}^{\delta})$ be a fixed point of $\bar{e} \circ (\operatorname{BR}^1 \times \operatorname{BR}^{\otimes} \times \operatorname{BR}^2) \circ w$. Then $\tau_n = \tau_n^{\otimes} * \tau_n^0$ and $\tau_n^{\otimes} \in B_n^{\delta}$, which implies that τ_n^{\otimes} is equivalent to $\hat{\tau}_n^{\otimes}$ in \mathbb{G}^{\otimes} . The profile (ζ_n^n, τ_{-n}^0) is an equilibrium of \mathbb{G}^n so that $\mathbb{G}_n^n(\zeta_n^n, \tau_{-n}^0) < G_n(Q) - \delta$ because of **A.2** and our choice of δ . Therefore, for every $s_n^n \in S_n^n$ we have $\mathbb{G}_n^n(s_n^n, \tau_n^0) < G_n(Q) - \delta$. Hence $(\tau, \zeta) \in \hat{\mathcal{O}}^{\delta}$ which proves that $\hat{\mathcal{O}}^{\delta}$ is admissible.

If $(\tau,\zeta) \in \hat{\mathcal{O}}^{\delta}$ is a fixed point of $\bar{e} \circ (\mathbb{BR}^1 \times \mathbb{BR}^{\otimes} \times \mathbb{BR}^2) \circ w$ then $\tau_n = \tau_n^{\otimes} * \tau_n^0$, so that $\tau_n^{\otimes} \in B_n^{\delta}$ and τ_n^{\otimes} is equivalent to $\hat{\tau}_n^{\otimes}$. Moreover, (ζ_n^n, τ_{-n}^0) is an equilibrium of \mathbb{G}^n and $\mathbb{G}_n(s_n^n, \tau_{-n}) = \mathbb{G}_n^n(s_n^n, \tau_{-n}^0) < G_n(Q) - \delta$, for every s_n^n . Therefore, $(\zeta, \tau) \in \mathcal{O}^{\delta}$. On the other hand, if $(\zeta, \tau) \in \mathcal{O}^{\delta}$ is a fixed point then again $\tau_n = \tau_n^{\otimes} * \tau_n^0$ and $\tau_n^{\otimes} \in B_n^{\delta}$ so that τ_n^{\otimes} is equivalent to $\hat{\tau}_n^{\otimes}$ in \mathbb{G}^{\otimes} . Therefore, $\mathbb{G}_n^n(s_n^n, \tau_n^0) = \mathbb{G}_n(s_n^n, \tau_n^0) < G_n(Q) - \delta$ for every s_n^n . That is, $\tau_n^0 \in \hat{U}_n^{\delta}$ and $(\zeta, \tau) \in \hat{\mathcal{O}}^{\delta}$.

Note now that Property **I.6** implies that $\bar{e} \circ (BR^1 \times BR^{\otimes} \times BR^2) \circ w|_{\hat{\mathcal{O}}^{\delta}}$ has index zero. In turn, Property **I.5** implies that the index of $BR^1|_{\Sigma_1^1 \times \hat{U}_2^{\delta}} \times BR^2|_{\Sigma_2^2 \times \hat{U}_1^{\delta}} \times BR^{\otimes}|_{B^{\delta}}$ is zero. Given Proposition 2.3, Property **I.4** implies that for some n = 1, 2, the index of $BR^n|_{\Sigma_n^n \times \hat{U}_{-n}^{\delta}}$ is zero. Since BR^n is the best-reply correspondence of the excluded game \mathbb{G}^n , the index of the supporting polytope in player *n*'s excluded game is zero.

Example 5.9. We continue with Example 3.4. The "*Out*"-component induces supporting polytope $K^1 = \Delta(\{L, R\}) \times \{\alpha \ell + (1 - \alpha)r \mid 0 \le \alpha \le 2/3\}$ in player 1's excluded game. Note that such a supporting polytope contains two equilibria, (R, r) and the mixed $(\frac{3}{4}L + \frac{1}{4}R, \frac{1}{4}\ell + \frac{3}{4}r)$. As anticipated in the Introduction, the former has index +1 while the latter has index -1. Therefore, the supporting polytope has index 0.

5.2. Step 2: Perturbing the excluded game so that it has no equilibria in the supporting polytope

From Step 1, assume without loss of generality that K^1 has index 0. Lemma 5.10 below guarantees that, when that is the case, there is a game $\bar{\mathbb{G}}^1$ equivalent to \mathbb{G}^1 and a payoff perturbation of $\bar{\mathbb{G}}^1$ such that, letting \bar{K}^1 be the set of strategies equivalent to K^1 in $\bar{\mathbb{G}}^1$, the resulting perturbed game does not have an equilibrium in \bar{K}^1 . Such a perturbed game also satisfies some additional properties that are used in Section 5.3.

To state the Lemma, define $\Sigma_2^1 \equiv \Delta(S_2^0)$ so that we can write the set of mixed strategy profiles in \mathbb{G}^1 as $\Sigma^1 = \Sigma_1^1 \times \Sigma_2^1$. Define also $K_2^{1,\eta} \equiv \{\sigma_2^1 \in \Sigma_2^1 \mid \forall s_1^1 \in S_1^1, \mathbb{G}_1^1(s_1^1, \sigma_2^1) \leq G_1(Q) - \eta\}$.

Lemma 5.10. There exists a game $\bar{\mathbb{G}}^1 = (\bar{S}_1^1, \bar{S}_2^0, \bar{\mathbb{G}}_1^1, \bar{\mathbb{G}}_2^1)$ equivalent to \mathbb{G}^1 such that:

- (1) Player 1's pure strategy set \bar{S}_1^1 is the collection of vertices of a polyhedral refinement \mathscr{P}_1 of a simplicial subdivision of $\bar{\Sigma}_1^1$.
- (2) Player 2's pure strategy set \bar{S}_2^0 is the collection of vertices of a polyhedral refinement \mathscr{P}_2 of Σ_2^1 for which no polyhedron of \mathscr{S}_2 that intersects $\tilde{\partial}K_2^1$ has a point in common with $K_2^{1,\eta}$.
- (3) There exists $\delta > 0$, a δ -perturbation $\overline{\mathbb{G}}^{1,\delta}$ of $\overline{\mathbb{G}}^1$ and, for each $\varepsilon > 0$, an ε -perturbation $\overline{\mathbb{G}}^{1,\delta,\varepsilon}$ of $\overline{\mathbb{G}}^{1,\delta}$ such that for n = 1,2:

$$\lim_{\varepsilon \to 0} \bar{\mathbb{G}}_n^{1,\delta,\varepsilon}(s_1,s_2) = \bar{\mathbb{G}}_n^{1,\delta}(s_1,s_2) = \bar{\mathbb{G}}_n^1(s_1,s_2) \text{ if } s_2 \notin K_2^{1,\eta}, \text{ and}$$
$$\lim_{\varepsilon \to 0} \bar{\mathbb{G}}_n^{1,\delta,\varepsilon}(s_1,s_2) = \bar{\mathbb{G}}_n^{1,\delta}(s_1,s_2) \ge \bar{\mathbb{G}}_n^1(s_1,s_2) \text{ if } s_2 \in K_2^{1,2\eta}.$$

- (4) For each n = 1,2 and $\varepsilon > 0$ small enough, if mixed strategies $\sigma_n \neq \tilde{\sigma}_n$ are equivalent in $\bar{\mathbb{G}}^{1,\delta}$ and σ_n 's support is within the vertices of a polyhedron of \mathscr{P}_n whereas $\tilde{\sigma}_n$'s is not, then, for every mixed strategy σ_{-n} , we have $\bar{\mathbb{G}}_n^{1,\delta,\varepsilon}(\sigma_n,\sigma_{-n}) > \bar{\mathbb{G}}_n^{1,\delta,\varepsilon}(\tilde{\sigma}_n,\sigma_{-n})$.
- (5) For $\varepsilon > 0$ small enough, the game $\bar{\mathbb{G}}^{1,\delta,\varepsilon} = (\bar{S}_1^1, \bar{S}_2^0, \bar{\mathbb{G}}_1^{1,\delta,\varepsilon}, \bar{\mathbb{G}}_2^{1,\delta,\varepsilon})$ does not have an equilibrium in \bar{K}^1 .

Proof. See Online Appendix B.

The construction of $\bar{\mathbb{G}}^1$ in Lemma 5.10 follows closely the construction presented in Theorem 3 in GW. However, there are a few differences stemming from the fact that the supporting polytope may contain more than one equilibrium component of the excluded game. Parts (1) and (2) simply record that the strategies obtained in $\bar{\mathbb{G}}^1$ are duplicates originating from vertices of a polyhedral subdivision of the strategy sets. This is needed so that (3) and (4) can be meaningfully stated and used in Section 5.3. Chiefly, in Part (3), s_2 can be viewed as a vertex of the polyhedral subdivision \mathscr{P}_2 of Σ_2^1 ; and (4) refers explicitly to the vertices of the polyhedral subdivision. The main difference with the construction in GW is (3). Within $\bar{K}^{1,2\eta}$ the payoff perturbations are not necessarily arbitrarily small, i.e., the inequality may be strict. This contrasts with GW where all payoff perturbations can be made as small as desired.¹⁶

Example 5.11. In Example 3.4 we established that the supporting polytope $K^1 = \Delta(\{L, R\}) \times \{\alpha \ell + (1 - \alpha)r \mid 0 \le \alpha \le 2/3\}$ has index 0. Therefore, we can find a game equivalent to the subgame in Figure 1 and a perturbation so that the perturbed game has no equilibrium in K^1 . Since such a subgame is relatively simple, we can accomplish this by directly perturbing it to the following game:

	$ar{\ell}$	\bar{r}		
Ī	3,1	2,0		
\bar{R}	0,0	1,3		

5.3. Step 3: The equivalent game and its perturbation

We construct a game \bar{G} equivalent to G and a perturbation of \bar{G} that has no equilibrium in a neighborhood of the equivalent component of equilibria \bar{K} . The main idea is to perturb G so that, with vanishing probability, player 1 is forced to play the game $\bar{\mathbb{G}}^{1,\delta,\varepsilon}$ that has been defined in Step 2 and that has no equilibrium in \bar{K}^1 . We show that, in this perturbed game, players' behavior induce an equilibrium of $\bar{\mathbb{G}}^{1,\delta,\varepsilon}$ so that player 2's strategy cannot belong to the supporting polytope. This, in turn, implies that the perturbed game of \bar{G} cannot have an equilibrium in the fixed original neighborhood of \bar{K} .

¹⁶ This is a critical difference. As in GW, we obtain the payoff-perturbations from a map that does not have fixed points on K^1 . In GW such a map is chosen to be "close" to the best-reply correspondence at every point of the domain. However, in our setting, it is possible for K^1 to contain many equilibrium components with nonzero index. In those cases, a map without fixed point on K^1 can be "close" to the best-reply correspondence in the complement of K^1 but "far" from it on K^1 . (See Figure 2 for an illustration. The dotted line is far from the solid line in the interval (x_1, x_2) so that it does not cross the 45° line.) This translates to "large" payoff perturbations in K^1 .



Henceforth, extend the original game \mathbb{G} to the equivalent game $\overline{\mathbb{G}}$ in which player 1's strategy set is $S_1 \cup \overline{S}_1^1$ and player 2's strategy set is $S_2 \cup \overline{S}_2$, where

$$\bar{S}_2 \equiv \left\{ \hat{\tau}_2^{\otimes} * \bar{s}_2^0 \, | \, \bar{s}_2^0 \in \bar{S}_2^0 \right\}.$$

The extension is by equivalence, i.e., by viewing each strategy in \bar{S}_1^1 as a duplicate of a strategy in Σ_1^1 , and similarly, viewing each strategy in \bar{S}_2 as a duplicate of a strategy in Σ_2 . Let $\bar{\mathbb{G}}_1$ and $\bar{\mathbb{G}}_2$ be the corresponding bilinear extensions of \mathbb{G}_1 and \mathbb{G}_2 to $\Delta(S_1 \cup \bar{S}_1^1) \times \Delta(S_2 \cup \bar{S}_2)$.

For a given $\varepsilon > 0$, we define the extensive-form game \bar{G}^{ε} whose normal-form is $\bar{\mathbb{G}}^{\varepsilon}$. The extensive-form is represented in Figure 6. Nature moves first and chooses between κ_1 with probability $(1 - \varepsilon)$ and κ_2 with probability ε . Player 1 observes Nature's choices but player 2 does not. After κ_1 , player 1 chooses from $(S_1 \setminus S_1^1) \cup \bar{S}_1^1$. In turn, if Nature chooses κ_2 , player 1 has to choose an action from the set \bar{S}_1^1 . Player 2 observes neither Nature's move nor player 1's moves, and selects from $S_2 \cup \bar{S}_2$. In the payoff assignment of Figure 6, the symbol $\mathbb{I}_{[X]}$ represents the indicator function that equals 1 if X is true and 0 otherwise. Furthermore, the value of B > 0 is chosen sufficiently large such that $-B < \min_{\bar{s}_2 \in \bar{S}_2} \bar{\mathbb{G}}_2^{1,\delta,\varepsilon}(\bar{s}_1^1,\bar{s}_2)$ for every $\bar{s}_1^1 \in \bar{S}_1^1$ and $\varepsilon > 0$.

When $\varepsilon = 0$, we write the game \bar{G}^0 simply as \bar{G} and note that it is equivalent to the original game G. Let \bar{K} be the equilibrium component equivalent to K in \bar{G} and let $\bar{\mathcal{O}}$ be a neighborhood of \bar{K} in the mixed strategy set of \bar{G}^{ε} whose closure contains no equilibrium of \bar{G}^{ε} in its boundary. Since $\bar{\mathcal{O}}$ contains no equilibria of \bar{G} in its boundary, this is also the case for the game \bar{G}^{ε} for sufficiently small $\varepsilon > 0$. We conclude the proof of Theorem 2.5 by proving that for $\varepsilon > 0$ small enough there is no equilibrium of \bar{G}^{ε} in $\bar{\mathcal{O}}$.

Similarly to Step 1, we work with behavioral strategies of game \bar{G}^{ε} . A behavioral strategy profile is a member of $\Delta((S_1 \setminus S_1^1) \cup \bar{S}_1^1) \times \bar{\Sigma}_1^1 \times \Delta(S_2 \cup \bar{S}_2)$. Suppose to the contrary that there exists a sequence $(\varepsilon_k)_{k \in \mathbb{N}}$ of positive numbers converging to 0 and a sequence $(\sigma^k)_{k \in \mathbb{N}} \subset \bar{\mathcal{O}}$ converging to $\sigma \in \bar{K}$ such that for every k the profile σ^k is an equilibrium of $\bar{\mathbb{G}}^{\varepsilon_k}$. Passing to a subsequence if necessary, assume $(\sigma_1^k, \sigma_2^k) = (\tau_1^k, \bar{\zeta}_1^{1,k}, \sigma_2^k) \to (\tau_1, \bar{\zeta}_1^1, \sigma_2)$. Write $\tau_1^k \equiv (1 - \alpha_k)\tau_1^{\times,k} + \alpha_k\bar{\tau}_1^{1,k}$ where $\tau_1^{\times,k} \in \Delta(S_1 \setminus \bar{S}_1^1)$, $\bar{\tau}_1^{1,k} \in \Delta(\bar{S}_1^1)$ and $\alpha_k \in [0,1]$. Similarly, write $\sigma_2^k \equiv (1 - \beta_k)\bar{\tau}_2^k + \beta_k\tau_2^k$, where $\bar{\tau}_2^k \in \Delta(\bar{S}_2)$, $\tau_2^k \in \Sigma_2$, and $\beta_k \in [0,1]$. Let $\varphi^k \equiv \varepsilon_k[\varepsilon_k + (1 - \varepsilon_k)\alpha_k]^{-1}$. Again, passing to the corresponding subsequence if necessary, assume $(\tau_1^{\times,k}, \bar{\tau}_1^{1,k}) \to (\tau_1^{\times}, \bar{\tau}_1^1)$, $(\bar{\tau}_2^k, \tau_2^k) \to (\bar{\tau}_2, \tau_2)$, $\beta^k \to \beta$ and $\varphi^k \to \varphi$.

Claim 5.12. For k large enough, $\alpha_k = 0$.

Proof. To the contrary, suppose there is a subsequence of $(\sigma^k)_{k \in \mathbb{N}}$ for which the corresponding α_k 's are all strictly positive. From Lemma 3.1 there is a product strategy $\sigma_2^{\otimes} * \sigma_2^0$ in \mathbb{G} that is equivalent to σ_2 in $\overline{\mathbb{G}}$. First observe $\sigma_2^0 \in \tilde{\partial}K_2^1$. Otherwise, $\sigma_2^0 \in K_2^1 \setminus \tilde{\partial}K_2^1$ which, for k sufficiently large, implies $\alpha_k = 0$ since τ_1^k is an equilibrium strategy. Player 2's expected payoff against σ_1^k in $\overline{G}^{\varepsilon_k}$ for a typical strategy $y_2 = (1-b)\overline{x}_2 + bx_2$, with $\overline{x}_2 \in \Delta(\overline{S}_2)$, $x_2 \in \Sigma_2$, and $b \in [0,1]$ is

$$(9) \quad [\varepsilon_{k} + (1 - \varepsilon_{k})\alpha_{k}] \left(\varphi^{k} \left[(1 - b)\bar{\mathbb{G}}_{2}^{1,\delta,\varepsilon_{k}}(\bar{\zeta}_{1}^{1,k}, q_{2}^{0}(\bar{x}_{2})) - bB \right] + (1 - \varphi^{k})\bar{\mathbb{G}}_{2}(\bar{\tau}_{1}^{1,k}, y_{2}) \right) + (1 - \varepsilon_{k})(1 - \alpha_{k})\bar{\mathbb{G}}_{2}(\tau_{1}^{\times,k}, y_{2}).$$

Player 2's equilibrium strategy σ_2^k maximizes (9). The limit strategy is $\sigma_2 = (1-\beta)\bar{\tau}_2 + \beta\tau_2$. For k sufficiently large, σ_2 also maximizes (9) because $\operatorname{Supp}(\sigma_2) \subseteq \operatorname{Supp}(\sigma_2^k)$. Denote by A_2^{\otimes} the subset of strategies $\sigma_2' \in \Delta((S_2 \setminus S_2^2) \cup \bar{S}_2)$ for which $q_2^{\otimes}(\sigma_2')$ is equivalent to $\hat{\tau}_2^{\otimes}$ in game \mathbb{G}^{\otimes} . Since the payoff to player 2 from all strategies in A_2^{\otimes} against $\tau_1^{\times,k}$ in \mathbb{G} is the same,

(10)
$$\sigma_{2} \in \arg\max_{y_{2} \in A_{2}^{\infty}} \varphi^{k} \left[(1-b) \bar{\mathbb{G}}_{2}^{1,\varepsilon_{k}}(\bar{\zeta}_{1}^{1,k},q_{2}^{0}(\bar{x}_{2})) - bB \right] + (1-\varphi^{k}) \bar{\mathbb{G}}_{2}(\bar{\tau}_{1}^{1,k},y_{2}).$$

Now, $q_2^{\otimes}(\sigma_2)$ is equivalent to $\hat{\tau}_2^{\otimes}$ in \mathbb{G}^{\otimes} which implies that $q_2^{\otimes}(\tau_2)$ is also equivalent to $\hat{\tau}_2^{\otimes}$ in \mathbb{G}^{\otimes} , so that $\tau_2 \in A_2^{\otimes}$. Therefore, there exists a product strategy $\hat{\tau}_2^{\otimes} * \bar{\tau}_2^0$ which is equivalent to τ_2 in \mathbb{G} . Hence, $\hat{\tau}_2^{\otimes} * \bar{\tau}_2^0 \in \Delta(\bar{S}_2)$ and

(11)
$$\bar{\mathbb{G}}_{2}(\bar{\tau}_{1}^{1,k},(1-\beta)\bar{\tau}_{2}+\beta(\hat{\tau}_{2}^{\otimes}*\bar{\tau}_{2}^{0}))=\bar{\mathbb{G}}_{2}(\bar{\tau}_{1}^{1,k},\sigma_{2}),$$

(12)
$$\bar{\mathbb{G}}_{2}^{1,\delta,\varepsilon_{k}}(\bar{\zeta}_{1}^{1,k},(1-\beta)\bar{\tau}_{2}+\beta(\hat{\tau}_{2}^{\otimes}*\bar{\tau}_{2}^{0}))>(1-\beta)\bar{\mathbb{G}}_{2}^{1,\delta,\varepsilon_{k}}(\bar{\zeta}_{1}^{1,k},q_{2}^{0}(\bar{\tau}_{2}))-\beta B$$

Therefore, it must be that $\beta = 0$, so $\sigma_2 = \hat{\tau}_2^{\otimes} * \bar{\tau}_2^0 \in \Delta(\bar{S}_2)$. Furthermore, $\bar{\tau}_2^0$'s support must be within the set of vertices of a polyhedron in \mathscr{P}_2 , otherwise player 2 can do strictly better by Lemma 5.10 (4). Hence, every $\bar{s}_2^0 \in \bar{S}_2^0$ in $\bar{\tau}_2^0$'s support satisfies $\bar{s}_2^0 \notin \bar{K}_2^{1,\eta}$ (because of (2) in Lemma 5.10). On the other hand, both $\bar{\tau}_1^{1,k}$ and $\bar{\zeta}_1^{1,k}$ must be optimal against σ_2^k for every k. The last two facts together with Lemma 5.10 (3) imply that at the limit $\varepsilon_k \to 0$, for $k \to +\infty$, we have

$$\bar{\zeta}_1^1 \in \arg \max_{x_1 \in \Delta(\bar{S}_1^1)} \bar{\mathbb{G}}_1^1(x_1, \bar{\tau}_2^0) \quad \text{and} \quad \bar{\tau}_1^1 \in \arg \max_{x_1 \in \Delta(\bar{S}_1^1)} \bar{\mathbb{G}}_1(x_1, \bar{\tau}_2).$$

Since $\bar{\tau}_2 = \hat{\tau}_2^{\otimes} * \bar{\tau}_2^0$ strategy $\bar{\tau}_1^1$ is, in fact, a best-reply against $\bar{\tau}_2^0$ in the excluded game $\bar{\mathbb{G}}^1$. Thus, $\varphi \bar{\zeta}_1^1 + (1-\varphi) \bar{\tau}_1^1$ is also a best-reply against $\bar{\tau}_2^0$ in $\bar{\mathbb{G}}^1$. Recall $\bar{\tau}_2^0 \in \tilde{\partial} K_2^1$, hence by Assumption **A.2**, player 2 must have a profitable deviation $\tilde{\tau}_2^0$ in $\bar{\mathbb{G}}^1$, i.e.,

$$\begin{split} \bar{\mathbb{G}}_{2}^{1}(\varphi\bar{\zeta}_{1}^{1}+(1-\varphi)\bar{\tau}_{1}^{1},\tilde{\tau}_{2}^{0}) > \bar{\mathbb{G}}_{2}^{1}(\varphi\bar{\zeta}_{1}^{1}+(1-\varphi)\bar{\tau}_{1}^{1},\bar{\tau}_{2}^{0}) &= \varphi\bar{\mathbb{G}}_{2}^{1}(\bar{\zeta}_{1}^{1},\bar{\tau}_{2}^{0}) + (1-\varphi)\bar{\mathbb{G}}_{2}^{1}(\bar{\tau}_{1}^{1},\bar{\tau}_{2}^{0}) &= \\ \lim_{k\to\infty} \left[\varphi^{k}\bar{\mathbb{G}}_{2}^{1,\delta,\varepsilon_{k}}(\bar{\zeta}_{1}^{1,k},\bar{\tau}_{2}^{0}) + (1-\varphi^{k})\bar{\mathbb{G}}_{2}(\bar{\tau}_{1}^{1,k},\hat{\tau}_{2}^{0}*\bar{\tau}_{2}^{0})\right], \end{split}$$

where the last equality uses Lemma 5.10 (3) since $\bar{\tau}_2^0 \notin \bar{K}_2^{1,\eta}$. Similarly,

$$\bar{\mathbb{G}}_{2}^{1}(\varphi\bar{\zeta}_{1}^{1}+(1-\varphi)\bar{\tau}_{1}^{1},\tilde{\tau}_{2}^{0}) = \varphi\bar{\mathbb{G}}_{2}^{1}(\bar{\zeta}_{1}^{1},\tilde{\tau}_{2}^{0}) + (1-\varphi)\bar{\mathbb{G}}_{2}^{1}(\bar{\tau}_{1}^{1},\tilde{\tau}_{2}^{0}) \leq \lim_{k \to \infty} \left[\varphi^{k}\bar{\mathbb{G}}_{2}^{1,\delta,\varepsilon_{k}}(\bar{\zeta}_{1}^{1,k},\tilde{\tau}_{2}^{0}) + (1-\varphi^{k})\bar{\mathbb{G}}_{2}(\bar{\tau}_{1}^{1,k},\hat{\tau}_{2}^{\otimes}*\tilde{\tau}_{2}^{0})\right],$$

where we have used again Lemma 5.10 (3) in the last inequality. From this we have a contradiction given that $\hat{\tau}_2^{\otimes} * \tilde{\tau}_2^0$ does strictly better than $\bar{\tau}_2$ for large enough k.

Claim 5.13. For k large enough, $(\bar{\zeta}_1^{1,k}, q_2^0(\bar{\tau}_2^k))$ is an equilibrium of $\bar{\mathbb{G}}^{1,\delta,\varepsilon_k}$.

Proof. From Claim 5.12, for k sufficiently large, $\alpha_k = 0$. For such a k, player 2's expected payoff from playing $\sigma_2^k = (1 - \beta_k)\overline{\tau}_2^k + \beta_k \tau_2^k$ against σ_1^k is equal to

$$(13) \quad \varepsilon_{k} \left[(1-\beta_{k})\bar{\mathbb{G}}_{2}^{1,\delta,\varepsilon_{k}}(\bar{\zeta}_{1}^{1,k},q_{2}^{0}(\bar{\tau}_{2}^{k})) - \beta_{k}B \right] + (1-\varepsilon_{k}) \left[\bar{\mathbb{G}}_{2}(\tau_{1}^{\times,k},(1-\beta_{k})\bar{\tau}_{2}^{k} + \beta_{k}\tau_{2}^{k}) \right] = \\ \varepsilon_{k} \left[(1-\beta_{k})\bar{\mathbb{G}}_{2}^{1,\delta,\varepsilon_{k}}(\bar{\zeta}_{1}^{1,k},q_{2}^{0}(\bar{\tau}_{2}^{k})) - \beta_{k}B \right] + (1-\varepsilon_{k}) \left[(1-\beta_{k})\mathbb{G}_{2}^{\otimes}(q_{1}^{\otimes}(\tau_{1}^{\times,k}),\hat{\tau}_{2}^{\otimes}) + \beta_{k}\bar{\mathbb{G}}_{2}(\tau_{1}^{\times,k},\tau_{2}^{k}) \right].$$

Repeating similar arguments to the proof of Claim 5.12 we obtain $\beta_k \to 0$. From (13) we therefore see that, for k large enough, $q_2^0(\bar{\tau}_2^k)$ must be a best-reply to $\bar{\zeta}_1^{1,k}$ in $\bar{\mathbb{G}}^{1,\delta,\varepsilon_k}$. In turn, $\bar{\zeta}_1^{1,k}$ must also be best-reply to $q_2^0(\bar{\tau}_2^k)$. That is, $(\bar{\zeta}_1^{1,k}, q_2^0(\bar{\tau}_2^k))$ is an equilibrium of $\bar{\mathbb{G}}^{1,\delta,\varepsilon_k}$. \Box

Let $(\bar{\zeta}_1^1, q_2^0(\bar{\tau}_2))$ be a limit of $(\bar{\zeta}_1^{1,k}, q_2^0(\bar{\tau}_2^k))_{k \in \mathbb{N}}$, that is, an equilibrium of $\bar{\mathbb{G}}^{1,\delta,0}$. By construction, $(\bar{\zeta}_1^1, q_2^0(\bar{\tau}_2))$ belongs to the complement of \bar{K}^1 . Therefore, by definition of \bar{K}^1 , there exists a pure strategy \bar{s}_1^1 such that $\bar{\mathbb{G}}_1^{1,\delta,0}(\bar{s}_1^1, \hat{\tau}_2^{\otimes} * q_2^0(\bar{\tau}_2))) = \bar{\mathbb{G}}_1(\bar{s}_1^1, \hat{\tau}_2^{\otimes} * q_2^0(\bar{\tau}_2)) > G_1(Q)$, where the equality comes from Lemma 5.10 (3). Hence, $\sigma \notin \bar{K}$, which implies $\sigma^k \notin \bar{\mathcal{O}}$ for sufficiently





large k. This contradicts our initial assumption. We conclude that there is no sequence of equilibria of \bar{G}^{ε_k} that converges to some point in $\bar{\mathcal{O}}$.

Example 5.14. We finish the example started in Section 1.1 using the same construction as in Figure 6. Since the game equivalent to player 1's excluded game has the same strategy space as the excluded game itself, it is enough to "plug" the perturbed version found in Example 5.11 after κ_2 . See Figure 7 for the resulting game. Following the same line of reasoning as in the general case above, note that player 1 cannot be made indifferent between In and Out at the limit of a sequence of equilibria as ε goes to zero. If that was the case, player 2 would be playing ℓ with probability close to 2/3 along the sequence which, in turn, implies that player 1 would be playing T and \overline{T} with probability 1 and, therefore, player 2 would have a profitable deviation to playing ℓ with probability 1. Hence, if there is an equilibrium close to the "Out"-component then, for any $\varepsilon > 0$, player 1 must strictly prefer playing Out and player 2 would be choosing her strategy as if she was best-replying to player 1's move after κ_2 . That is, player 2 would be choosing ℓ with probability 1. But then, player 1 would have an incentive to deviate to In and then choose T. We conclude that the perturbed game has no equilibrium close to the "Out"-component and, correspondingly, that such a component is not hyperstable in the original game.

6. Applications

In this section we analyze some examples including a few prominent models in the economics literature and show how index computation can be used to eliminate non-hyperstable equilibria. We begin by revisiting the game in Figure 1 and showing that the recent refinement concept of *sequentially stable outcomes* of Dilmé (2024) diverges from hyperstability. We then move to study two finitely repeated games. In particular, Example 6.3 illustrates





how relying only on pure-strategy subgame-perfect equilibria to analyze these games may lead to selecting equilibrium outcomes that have index zero. We conclude by analyzing some classical signaling games. This serves both to compare hyperstabily to other equilibrium concepts and to highlight how hyperstable components are easy to compute through their index. In preparation for the examples, we note a novel way to compute the index of a component using the indexes of the supporting polytopes in their corresponding excluded games. The following result is a corollary from the arguments in Section 5.1.

Proposition 6.1. Let *G* be a two-player extensive-form game with normal form \mathbb{G} . Let *K* be a component for which assumptions **A.1** and **A.2** are satisfied. Letting $K^{\otimes} = q^{\otimes}(K)$ be the projected component in game \mathbb{G}^{\otimes} , the index of component *K* in game \mathbb{G} , written as $\operatorname{Ind}_{\mathbb{G}}(K)$, can be decomposed as follows

(14)
$$\operatorname{Ind}_{\mathbb{G}}(K) = \operatorname{Ind}_{\mathbb{G}^1}(K^1) \times \operatorname{Ind}_{\mathbb{G}^{\otimes}}(K^{\otimes}) \times \operatorname{Ind}_{\mathbb{G}^2}(K^2),$$

where, by convention, $\operatorname{Ind}_{\mathbb{G}^n}(K^n) \equiv 1$ if player *n* does not have an excluded game.

Example 6.2. Sequential equilibrium is one of the most prominent solution concepts in extensive-form games as it extends the principle of backward induction to imperfect information games. However, it has shortcomings which have been extensively discussed in the literature (Kohlberg and Mertens, 1986; Cho and Kreps, 1987). Sequentially stable outcomes, recently defined by (Dilmé, 2024), refine the set of sequential equilibrium outcomes.

Roughly, an outcome ω is sequentially stable if for *any* vanishing sequence of behavioral trembles, there exists a sequence of ε -sequential equilibria of the (trembled)-game (with

vanishing $\varepsilon > 0$) inducing outcomes which converge to ω (see Dilmé, 2024, p.9 for the definition). Note that while the definition of hyperstability considers payoff perturbations explicitly, by invoking ε -sequential equilibria, sequentially stable outcomes considers payoff perturbations implicitly.

Sequentially stable outcomes are a natural refinement of sequential equilibrium. They appropriately refine the set of equilibria in some examples while perhaps being equally tractable. Unfortunately, similarly to sequential equilibrium, sequentially stable outcomes rely on the specifics of the extensive-form. Hence, contrasting with hyperstability, they do not satisfy invariance.¹⁷ We can illustrate this difference returning to Figure 1. It can be easily checked that "Out" is a sequential equilibrium outcome. It is also a sequentially stable outcome (which can be seen by applying Dilmé, 2024, Proposition 4.3, Part 2). However, "Out" is not selected by either concept in the well-known equivalent representation of the game provided in Figure 8 (cf. Hillas, 1994). After adding the redundant strategy X, the unique equilibrium of the subgame that follows In is (L, ℓ) , making $(In - L, \ell)$ the unique Sequential Equilibrium (or Sequentially Stable) outcome. Not surprisingly, the forward induction outcome $(In - L, \ell)$ is the unique equilibrium selected by hyperstability both in in the representation of Figure 1 and in the representation of Figure 8. (Hyperstability can be computed directly in the latter game by noticing that player 1's excluded game associated with "Out" does not have any equilibrium in the supporting polytope, implying that the supporting polytope has index 0 and, by Proposition 6.1, component "Out" also has index 0, i.e., it is not hyperstable.)

We move to study a finitely repeated game. As we show below, finitely repeated games have a recursive structure that aids the computation of the index of equilibrium outcomes.¹⁸

Example 6.3. The repeated game with stage-game in Figure 9 is taken from Mailath and Samuelson (2006, Example 4.4.1). The stage game has 3 equilibria in pure strategies (*A*, *A*), (*B*, *B*) and (*C*, *C*) and 4 equilibria in mixed strategies $(\frac{1}{5}A + \frac{4}{5}B, \frac{3}{7}A + \frac{4}{7}B)$, $(\frac{1}{5}A + \frac{4}{5}C, \frac{1}{3}A + \frac{2}{3}C)$, $(\frac{3}{4}B + \frac{1}{4}C, \frac{1}{4}B + \frac{3}{4}C)$ and $(\frac{1}{17}A + \frac{12}{17}B + \frac{4}{17}C, \frac{3}{13}A + \frac{4}{13}B + \frac{6}{13}C)$. Pure strategy equilibria are strict and therefore have index +1. Mixed strategy equilibria that are not completely mixed have

¹⁷ In addition, as already discussed by Dilmé, sequentially stable outcomes do not satisfy admissibility either. ¹⁸ An interesting result in the literature on finitely repeated games is Osborne (1990). Proposition 1 in this paper provides a sufficient condition for a pure Nash equilibrium outcome (i.e., an outcome where no randomizations are made at any period by both players) of a finitely repeated, two-player game *not* to be Kohlberg-Mertens stable. In particular, this condition implies that this outcome has an index of zero: any non-zero index equilibrium component in mixed strategies contains a Kohlberg-Mertens stable set. If an outcome is not Kohlberg-Mertens stable, then the associated index is zero.

FIGURE 9. Stage-game for a repeated game with player specific punishments

	A	B	C
\boldsymbol{A}	4,4	0,0	0,0
B	0,0	3,1	0,0
C	2,2	0,0	1,3

index -1.¹⁹ Since indexes must add up to one, the completely mixed equilibrium has index +1.

Examples such as the finitely-repeated Prisoners Dilemma do not allow for equilibria where in some stage players play a profile different from the unique equilibrium of the stage game. This contrasts with this example in which the stage-game has equilibria that can be used differently to punish/reward players in the associated finitely repeated game, and allows for subgame perfect equilibrium outcomes in which a profile that is not necessarily an equilibrium of the stage-game is played in early stages. Indeed, Mailath and Samuelson (2006, p. 113) note that, if the game in Figure 9 is played twice, then it is possible to support (C, A) in the first period in subgame perfect equilibrium by playing (A, A) in the second stage if no player deviated in the first stage, (B, B) if the second player deviated, and (C, C) if the first player deviated. We show in this particular case that supporting the non-equilibrium profile in the first stage is not plausible (in the sense of hyperstability). To be precise, we show that the component K which induces the outcome "6,6" in which players play on-path (C, A) and then (A, A) is not hyperstable.

First, we note that such an outcome is not isolated in the set of equilibrium outcomes. Indeed, if "2,6" is the outcome in which players play on-path (C,C) twice, the same equilibrium component induces the continuum of outcomes α "6,6"+ $(1 - \alpha)$ "2,6" for $0 \le \alpha \le 1$. Therefore, Assumption **A.1** is not satisfied and we cannot directly apply the results in this paper to analyze the equilibrium component K. Nonetheless, we can use Property **I.1** and perturb player 2's payoffs slightly so that, *in the first stage*, her payoff is $3 - \varepsilon$ under profile (C,C) with $\varepsilon > 0$. Call the resulting perturbed game $G(\varepsilon)$. The index of any small enough admissible neighbourhood \mathcal{O} of K in the original game coincides with the index of \mathcal{O} in $G(\varepsilon)$ provided $\varepsilon > 0$ is small enough. Furthermore, in $G(\varepsilon)$, the unique component in \mathcal{O} does induce a unique outcome and, in this unique outcome, players play on path (C,A) and then (A,A).

¹⁹ The latter can be seen by deleting the unused pure strategy for both players for each equilibrium as it is an inferior response. The resulting truncated game has 3 equilibria, 2 of them strict and one in completely mixed strategies which, therefore, has index -1.

	A	В	C		$C\!A$	CB	CC
AA	8,8	4,4	4,4	A	$5,7-\varepsilon$	$1,3-\varepsilon$	$1,3-\varepsilon$
AB	4,4	7,5	4,4	В	$1,3-\varepsilon$	$4, 4-\varepsilon$	$1,3-\varepsilon$
AC	6,6	4,4	5,7	C	$3,5-\varepsilon$	$1,3-\varepsilon$	$2,6-\varepsilon$

FIGURE 10. Excluded games \mathbb{G}^1 (left) and $\mathbb{G}^2(\varepsilon)$ (right).

Let us compute player 1's excluded game under such an outcome. Since it does not depend on ε , we simply denote it \mathbb{G}^1 . Any deviation that involves player 1 playing *B* in the first stage gives player 1 a payoff of at most 3, which is strictly smaller that her payoff of 6 under the equilibrium component. Therefore, we may only consider player 1's deviations in which she plays *A* in the first stage. After eliminating duplicates, player 1's (simplified) excluded game is represented in Figure 10 (left). Note that payoffs are obtained from Figure 9 after adding to each payoff vector the vector (4,4) and, consequently, the set of Nash equilibria of \mathbb{G}^1 and the stage game in Figure 9 coincide. Player 1's supporting polytope consists of those strategy profiles in \mathbb{G}^1 that yield player 1 a payoff smaller or equal than 6. This polytope includes equilibrium (*AC*,*C*) as well as every mixed strategy equilibrium. Hence, the index of the supporting polytope in player 1's excluded game is -1.

In turn, consider player 2's excluded game $\mathbb{G}^2(\varepsilon)$ in Figure 10 (right). Analogously to \mathbb{G}^1 , this (simplified) excluded game is constructed by noticing that every deviation in which player 2 plays *B* is an inferior response to the equilibrium outcome and by eliminating duplicates. Payoffs are obtained from Figure 9 after adding to each payoff vector the vector $(1, 3 - \varepsilon)$. That is, both games also have the same Nash equilibria. Player 2's supporting polytope consists of those strategy profiles in $\mathbb{G}^2(\varepsilon)$ that yield player 2 a payoff smaller than 6. This polytope includes every equilibrium of $\mathbb{G}^2(\varepsilon)$ but (A, CA) and, therefore, the index of the supporting polytope in player 2's excluded game is zero. Using Proposition 6.1, the index of the component *K* that induces the outcome in which players play (A, C) in the first stage is zero for $\varepsilon \ge 0$ small enough. That is, (A, C) cannot be sustained in the first stage in a hyperstable equilibrium component in the two-fold repetition of the game in Figure 9.

Example 6.4. van Damme (1989) shows that in the two-fold repetition of the stage game in Figure 11 the outcomes in which players either play (T,L) or (B,R) twice is not stable in the sense of Kohlberg and Mertens (1986) whereas the outcome in which players alternate is, in turn, stable. In addition, van Damme (1989, p. 429) also points out that "[c]omputing the set of all stable equilibrium paths seems to be laborious, however, even in this most simple conceivable case." Here, we illustrate how the properties of the index and the results in this

paper can be used to compute the index of every equilibrium component in this repeated game. We also note that every nonzero index component contains a stable set (both in the sense of Kohlberg and Mertens, 1986, and in the sense of Mertens, 1989.)

Call G the two-fold repetition of the game in Figure 11. The stage game of G has 3 equilibria. The two pure equilibria (T,L) and (B,R) are strict and therefore have index +1. The completely mixed equilibrium $(\frac{4}{5}T + \frac{1}{5}B, \frac{1}{5}L + \frac{4}{5}R)$ has index -1.

First, every equilibrium component in which players play a completely mixed action profile in the first period has nonzero index. Indeed, if in such a component players always play the completely mixed equilibrium of the stage game after any realization in the first stage, then this equilibrium component induces a completely mixed outcome which, therefore has nonzero index (cf. Proposition 2.3). On the other hand, if for some realization in the first stage, players play a strict equilibrium in the second stage then every observable deviation consists of deviating in the second stage from the prescribed strict equilibrium. Those deviations are inferior responses to the equilibrium outcome and after eliminating them the outcome is completely mixed and, therefore, has nonzero index.

Let "8,2" be the equilibrium outcome in which players play (T,L) twice. Using a similar argument as above, deviating only in the second stage is an inferior response to the equilibrium outcome. After eliminating those deviations from G, we can see that player 1 has (after eliminating duplicates) only 2 observable deviations, BT and BB (that is, playing Bin the first stage and then playing either T or B in the subgame that follows (B,L)). Player 2's strategy set in player 1's excluded game is $\{L,R\}$. Note that if player 1 deviates to either BT or BB in G then the payoff vector accrued in the first stage is (0,0) given that player 2 is playing L in such a stage. Therefore, the excluded game \mathbb{G}^1 coincides with the game in Figure 11.

Player 1's payoff in \mathbb{G}^1 is always strictly smaller than player 1's payoff of 8 if (T,L) is played twice. Hence, the supporting polytope K^1 coincides with the entire strategy space in \mathbb{G}^1 which implies that K^1 has index +1 in \mathbb{G}^1 . In turn, consider player 2's observable deviations RL and RR. The excluded game \mathbb{G}^2 again coincides with the game in Figure 11.

The supporting polytope K^2 requires player 1 to play B with probability smaller or equal than 0.5. The only equilibrium that lies outside that region is (B,R) which is strict and, therefore, has index +1. Thus, the supporting polytope K^2 has zero index in game \mathbb{G}^2 . Proposition 6.1 implies that the component "8,2" has zero index and, therefore, it is not hyperstable. The analogous argument of course holds for the equilibrium component "2,8". It also holds for the equilibrium outcome in which players play the completely mixed equilibrium of the stage game after playing one of the strict equilibria in the first stage. FIGURE 11. Coordination stage-game

	L	R		
Т	4,1	0,0		
В	0,0	1,4		

Finally, consider now the two equilibrium outcomes "5,5", one in which players play (T,L) first and (B,R) second, and the other in which this order is reversed. In both cases player 1 and player 2's excluded games coincide with the game in Figure 11. Thus, every payoff in the two excluded games associated to the corresponding equilibrium component is smaller for both players than the payoff 5 induced by the equilibrium outcome. This implies that both supporting polytopes have index +1 and, therefore, the two components that induce the equilibrium payoff vector (5,5) also have index +1 and are hyperstable.

We now turn our attention to signaling games. Signaling games have driven the creation of a wealth of refinement concepts to eliminate implausible equilibria. Some of the most well-known are the Intuitive Criterion, D1, D2, Never-weak-best-reply (NWBR), and Universal Divinity. Hyperstability is strictly stronger than any of these criteria and, as we show, it is easy to compute by using properties of the index. The next proposition is known but we recall it for completeness.

Proposition 6.5. Let G be a finite signaling game. Suppose the terminal payoffs of G are chosen according to **A.1** and **A.2** and Q is an equilibrium outcome which is hyperstable. Then Q satisfies the Intuitive Criterion, D1, D2, NWBR and Universal Divinity.

The result follows from observing that a stable outcome in the sense of Kohlberg and Mertens (1986) satisfies all the criteria listed above and that every hyperstable outcome is also Kohlberg-Mertens stable.

Example 6.6. Consider the Beer-Quiche Game from Cho and Kreps (1987) (see Figure 12). The game has two equilibrium outcomes. In the first one, Sender chooses Beer (B) for both of his types and Receiver, seeing this, does Not Fight (NF). Off the equilibrium path, Receiver Fights (F) with probability at least 0.5. This equilibrium outcome survives all the classical refinements listed in Cho and Kreps and there are good reasons for it to be selected (see Cho and Kreps, 1987, pp. 184-185). In the other equilibrium outcome, Sender chooses Quiche (Q) for both of his types and, upon seeing this, Receiver does not fight. To prevent the strong type (S) from deviating to Beer (B), Receiver fights with probability at least 0.5 off the equilibrium path. Cho and Kreps provide intuitive reasons under which Receiver's

LUCAS PAHL AND CARLOS PIMIENTA

FIGURE 12. Beer-and-Quiche Game



off-equilibrium beliefs are implausible (see p.185) and this equilibrium outcome does not survive the Intuitive Criterion. Therefore, it is not hyperstable (see Proposition 6.5). But we can verify this directly by computing the index of the second equilibrium outcome.

Sender's excluded game associated to this equilibrium outcome is in the right-hand side of Figure 12 where QB, for instance, is the deviation in which sender chooses Q after S and B after W. First note that the supporting polytope consists of those strategy profiles in the excluded game whose payoff to Sender is smaller than or equal to $\frac{21}{10}$ (i.e. her payoff under the QQ - NF component). In this excluded game, BB is strictly dominated. Furthermore, the excluded game has three equilibria, (BQ,NF), (QB,F), and $(\frac{1}{10}BQ + \frac{9}{10}QB, \frac{3}{8}NF + \frac{5}{8}F)$. Sender's payoff under these equilibria is, respectively, 3, $\frac{9}{5}$, and $\frac{15}{8}$. That is, only the last two belong to the supporting polytope and, since strict equilibria have index +1 and the sum of all the indexes must be equal to +1, the index of the supporting polytope is zero. Proposition 6.1 implies that QQ - NF has index zero and, consequently, it is not hyperstable.

Example 6.7. Consider the game in Cho and Kreps (1987, Figure IV, p. 207) and represented in Figure 13. The game has two equilibrium outcomes. In the first, both types send m' and, off the equilibrium path, Receiver plays in the convex hull of (2/3,0,1/3), (1/6,1/2,1/3), (0,1/2,1/2) and (0,0,1) where probabilities in each vector correspond, in order, to strategies r_1 , r_2 and r_3 . This first component satisfies the Intuitive Criterion, D1, D2, and Universal Divinity. In the second equilibrium outcome both types send m and Receiver plays r_2 . Since the second equilibrium outcome is induced by a strict equilibrium (hence with index+1) and the indexes of all equilibria always add up to +1 (Property **I.3**) the first component has zero index. While simple, this reasoning requires knowing all the equilibrium outcomes but, as before, we can also directly compute its index.





Sender's excluded game in the component is on the right-hand side of Figure 13. Its unique equilibrium is (mm, r_2) and it lies outside the supporting polytope which, consequently, has index zero. Therefore, by Proposition 6.1, the m'm' component has index zero and, therefore, it is not hyperstable.

Example 6.8. We move to the example in Cho and Kreps (1987, p.216) represented in Figure 14. This example illustrates how Kohlberg-Mertens stability is strictly stronger than NWBR, D1, D2, Universal Divinity, and the Intuitive Criterion. Similarly to the previous example, this game has two equilibrium outcomes. The first one is the strict equilibrium $(mm'm', r_1)$, where mm'm' means player 1 plays m in his leftmost information set, and m' in the middle and rightmost information sets. Hence, the other equilibrium outcome in which all three types send m' must have index 0 and not be hyperstable. The formal argument sketched by Cho and Kreps that verifies that the second equilibrium outcome is not Kohlberg-Mertens stable (although satisfying all other refinement criteria) is more involved than what we showed here and, because of the unintuitive nature of their characterization of stability, less convincing.²⁰ In contrast, using the properties of the index, identifying hyperstability is easy. Its intuitiveness comes from the fact that it captures payoff robustness in every equivalent game.

The supporting polytope in Sender's excluded game is a polyhedron obtained by permuting the vector (2/3, 1/3, 0) in Receiver's mixed strategy set. This supporting polytope contains two equilibria in its interior. Both indices can be explicitly computed using Shapley's formula (Shapley, 1974). One has index +1 and the other index -1, which implies that the supporting polytope in Sender's excluded game has index zero and, by Proposition 6.1, the

²⁰ This is a position Cho and Kreps express themselves, see last paragraph of p. 220.

LUCAS PAHL AND CARLOS PIMIENTA



FIGURE 14. A signaling game with three types

outcome in which every type chooses m' has zero index and, therefore, fails to be hyperstable.

Example 6.9. We conclude with a finite game version of Spence's job-market model in Figure 15. There are two equilibrium outcomes, one pooling and one separating and hyperstability selects the separating equilibrium outcome. We can analyze the game in the same manner as the Beer-and-Quiche game to show that the pooling equilibrium (where the Sender sends e for both types, and the Receiver replies with w on equilibrium path, and off equilibrium path, with a probability of w at least 1/2) has index 0.

Interestingly, when we consider the general Spence model (Cho and Kreps, 1987, Section V) with any number of types and a continuum of education levels and wages, the same reasoning can be used to eliminate any pooling or semi-pooling equilibrium outcome, leaving only the fully separating one (the so-called Riley outcome) left. To show that, consider a (semi-)pooling equilibrium outcome where a certain number of types pool at education level e_p . To prevent the highest of the pooling types, denoted θ^* , from deviating from e_p to a higher education level $e > e_p$ the associated equilibrium wage $w(e) < (\theta^* e)$ is offered (off-path) after e. This renders a deviation of all pooling types but θ^* to e a strictly inferior reply (this follows from single-crossing). Therefore, we can eliminate all corresponding pure strategies that assign e to all such pooling types. After this elimination and under the (semi-)pooling equilibrium outcome, firms know that they are facing a worker of type θ^* upon observing a deviation to education level e. Backward induction now implies that firms best-reply with wage level ($\theta^* e$) > w(e) which, in turn, upsets the (semi-)pooling equilibrium.

We emphasize that the characterization in our main theorem is not known to hold in this infinite dimensional setting. The purpose here is only to remark that the same two principles invoked previously to compute the index in the finite game setting (elimination





of strictly inferior replies and backward induction) can also be invoked in the Spence model to eliminate pooling and semi-pooling equilibria.

APPENDIX A. PROOF OF PROPOSITION 2.3

We introduce enabling strategies. Fix a two-player finite game tree Γ . For each player n = 1, 2 and $z \in Z$, let $A_n(z)$ be the set of Player n's actions that precede z. The set of Player n's actions that are \prec -maximal in $A_n(z)$ for some $z \in Z$ are denoted L_n and called *Player n*'s *last actions in* Γ . If $L_n = \emptyset$ then n is a dummy player. Player n's last action that precedes z is $\ell_n(z)$. Given last action $i \in L_n$, let $u \in U_n$ satisfy $i \in A_n(u)$. The subset $S_n(i) \subset S_n$ consists of those strategies that at every information set u' of Player n that precede u prescribe the unique action leading to u. Define $q_n^e : \Sigma_n \to [0,1]^{L_n}$ so that for every $\sigma_n \in \Sigma_n$ and $i \in L_n$ we have $q_n^e(\sigma_n) \equiv (\sum_{s_n \in S_n(i)} \sigma_n(s_n))_{i \in L_n}$. The map q_n^e is affine and therefore $C_n \equiv q_n^e(\Sigma_n)$ is a polytope in \mathbb{R}^{L_n} . The polytope C_n is the enabling strategy set of Player n. If Nature moves in Γ we analogously denote the set of Nature's last moves by L_0 . Denote $L \equiv L_0 \times L_1 \times L_2$.

Given extensive-form payoffs $G \in \mathscr{G}$ and a profile of enabling strategies $(c_1, c_2) \in C \equiv C_1 \times C_2$, we construct payoffs are as follows. First, for every Nature's last move $i_0 \in L_0$, define $c_0(i_0) \equiv \prod_{i'_0 \in A_0, i' \leq i_0} \rho(i'_0)$. Each of Player *n*'s last action in L_n is associated with a unique sequence of Player *n*'s past actions that lead to the information set in which such a last action is available. Therefore, a vector of last actions $i = (i_0, i_1, i_2)$ either defines a path from the root to a terminal node or it does not. In the latter case, define $g_n(i) \equiv 0$. In the former case, letting *z* be the terminal nodes that is reached by such path, define $g_n(i) \equiv G_n(z)$. Player *n*'s enabling payoff function $V_n^e : C \to \mathbb{R}$ is defined by

(15)
$$V_n^e(c) \equiv \sum_{i \in L} g_n(i) c_0(i_0) c_1(i_1) c_2(i_2).$$

Note that the payoff function V_n^e is affine in each coordinate and, therefore, defines a polytopeform game $\mathfrak{V}^e = (C_1, C_2, V_1^e, V_2^e)$ which is called the *enabling-form of G*.

To prove Proposition 2.3, let d_n be the dimension of C_n . Since K induces a unique outcome with full support, then $q^e(K) = p$ is an isolated equilibrium in enabling strategies, located in the interior of C. For each n = 1, 2, take a d_n -simplex F_n contained in the relative interior of the enabling strategy set C_n and containing p_n in its relative interior. If we restrict the enabling payoff function V_n^e to D_n , then this defines a normal-form game with an isolated completely mixed equilibrium p, with the same index as the polytope-form index of the enabling-form (cf. Pahl, 2023). It is known that completely mixed isolated equilibria of normal-form games have indexes +1 or -1, since the payoff matrices are non-singular (and Shapley's formula gives their index, see Shapley, 1974). In particular, p has non-zero index. The equivalence between the index in normal and polytope forms (cf. Pahl, 2023) then implies that K has non-zero index.

REFERENCES

- A. Balthasar. *Geometry and equilibria in bimatrix games*. PhD thesis, London School of Economics and Political Science, 2009.
- I.-K. Cho and D. M. Kreps. Signaling games and stable equilibria. The Quarterly Journal of Economics, 102(2):179–221, 1987.
- S. DeMichelis and F. Germano. On the indices of zeros of nash fields. *Journal of Economic Theory*, 94(2):192–217, 2000.
- S. Demichelis and K. Ritzberger. From evolutionary to strategic stability. Journal of Economic Theory, 113(1):51-75, 2003.
- F. Dilmé. Sequentially stable outcomes. Econometrica, 92(4):1097-1134, 2024.
- A. Dold. Lectures on algebraic topology. Berlin: Springer-Verlag, 1972.
- S. Govindan and R. Wilson. Equivalence and invariance of the index and degree of nash equilibria. Games and Economic Behavior, 21(1):56-61, 1997a. ISSN 0899-8256. doi: https: //doi.org/10.1006/game.1997.0516. URL https://www.sciencedirect.com/science/article/ pii/S0899825697905169.
- S. Govindan and R. Wilson. Uniqueness of the index for nash equilibria of two-player games. *Economic Theory*, 10:541-549, 1997b. ISSN 0938-2259. doi: 10.1007/s001990050173.
- S. Govindan and R. Wilson. Direct proofs of generic finiteness of Nash equilibrium outcomes. *Econo*metrica, 69(3):765–69, May 2001.
- S. Govindan and R. Wilson. Structure theorems for game trees. *Proceedings of the National Academy* of Sciences, 99(13):9077–9080, 2002.

- S. Govindan and R. Wilson. Essential equilibria. Proceedings of the National Academy of Sciences, 102(43):15706-15711, 2005. doi: 10.1073/pnas.0506796102. URL https://www.pnas.org/doi/ abs/10.1073/pnas.0506796102.
- S. Govindan, R. Laraki, and L. Pahl. On sustainable equilibria. *Journal of Economic Theory*, 213: 105736, 2023. ISSN 0022-0531. doi: https://doi.org/10.1016/j.jet.2023.105736. URL https://www.sciencedirect.com/science/article/pii/S0022053123001321.
- E. Hauk and S. Hurkens. On forward induction and evolutionary and strategic stability. *Journal of Economic Theory*, 106(1):66-90, 2002. ISSN 0022-0531. doi: https://doi.org/10.1006/jeth.2001. 2856. URL https://www.sciencedirect.com/science/article/pii/S002205310192856X.
- J. Hillas. How much of forward induction is implied by backward induction and ordinality? University of Auckland, New Zealand, 1994.
- E. Kohlberg and J.-F. Mertens. On the strategic stability of equilibria. *Econometrica: Journal of the Econometric Society*, 54(5):1003–1037, 1986.
- D. M. Kreps and R. Wilson. Sequential equilibria. Econometrica, 50(4):863-894, 1982. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1912767.
- G. J. Mailath and L. Samuelson. *Repeated games and reputations: long-run relationships*. Oxford university press, 2006.
- A. McLennan. Fixed points of contractible valued correspondences. International Journal of Game Theory, 18(2):175–184, 1989.
- A. McLennan. The index+ 1 principle. Working Paper, 2016.
- A. McLennan. Advanced fixed point theory for economics, volume 25. Springer, 2018.
- J.-F. Mertens. Stable equilibria: A reformulation part i. definition and basic properties. Mathematics of Operations Research, 14(4):575-625, 1989. ISSN 0364765X, 15265471. URL http: //www.jstor.org/stable/3689732.
- J. Nash. Non-cooperative games. Annals of Mathematics, 54(2):286-295, 1951. ISSN 0003486X. URL http://www.jstor.org/stable/1969529.
- B. O'Neill. Essential sets and fixed points. American Journal of Mathematics, 75:497, 1953. ISSN 0002-9327. doi: 10.2307/2372499.
- M. Osborne and A. Rubinstein. A course in game theory. The MIT press, 1994.
- M. J. Osborne. Signaling, forward induction, and stability in finitely repeated games. Journal of Economic Theory, 50(1):22–36, 1990.
- L. Pahl. Polytope-form games and index/degree theories for extensive-form games. *Games and Economic Behavior*, 141:444–471, 2023.
- K. Ritzberger. The theory of normal form games from the differentiable viewpoint. *International Journal of Game Theory*, 23:207–236, 1994.
- L. S. Shapley. A note on the lemke-howson algorithm. *Pivoting and Extension: In honor of A.W. Tucker*, pages 175–189, 1974.

LUCAS PAHL AND CARLOS PIMIENTA

- E. van Damme. Stable equilibria and forward induction. *Journal of Economic Theory*, 48(2):476–496, 1989.
- 44