

Equilibrium Play Without Mutual Knowledge of Rationality

Florian Brandl
University of Bonn

Felix Brandt
Technical University of Munich

Equilibrium play in two-player zero-sum games is usually justified via epistemic assumptions, such as mutual knowledge of rationality and beliefs, that go far beyond the rationality of the players. We propose a justification that dispenses with these assumptions. To this end, we consider solution concepts that assign to every subgame of a given game a set of plausible actions for each player, and we impose two conditions. Rationality requires that the plausible sets are supports of undominated strategies or, equivalently, that all plausible actions are best responses to a common belief about the opponent. Inheritance requires that plausible actions remain plausible when implausible actions are discarded. In two-player games, the two conditions characterize the solution concepts that consistently select supports of Nash equilibria. Zero-sum payoffs ensure that Nash equilibria—and hence the selection—are generically unique. Equilibrium play thus emerges from individual rationality and the mutual understanding that plausibility judgments persist when implausible actions are discarded.

1 Introduction

Zero-sum games occupy a special place in game theory. Von Neumann (1928) proved that every finite two-player zero-sum game admits maximin strategies that are interchangeable and yield unique payoffs—nowhere is the case for equilibrium play more compelling than in this class of games. Yet, the classic justification of maximin play is decision-theoretic rather than strategic: a maximin strategy guarantees the value against every conceivable behavior of the opponent. Why a player who faces a possibly naive opponent should nevertheless randomize according to an equilibrium strategy is much less clear, and has been debated at least since ?’s (?) “reluctant duelist”: exploiting a flawed opponent seems to require abandoning equilibrium play. Epistemic analyses answer that equilibrium behavior is warranted when the players’ rationality and their beliefs are mutually known (Aumann and Brandenburger, 1995)—assumptions that go far beyond rationality of the players themselves.

This paper offers a justification of equilibrium play in zero-sum games that dispenses with mutual knowledge of rationality. We consider solution concepts that assign to every subgame of a given game—every pair of sets of available actions—a set of plausible actions for each player, and we impose two conditions. Rationality requires the plausible sets to be supports of undominated strategies. Since a dominated strategy cannot be a best response to any belief about the opponent, rationality concerns each player in isolation and involves no assumption about what a player knows about the other player. Inheritance requires that the solution be inherited by any reduced game that is obtained by discarding actions outside the solution: plausible actions remain plausible when implausible actions are removed. In the intended interpretation, each player is rational, and it is mutually understood that plausibility judgments persist when implausible actions are discarded.

Our main result shows that for generic zero-sum games, these two conditions characterize equilibrium play: a solution concept satisfies inheritance and rationality if and only if it assigns to every subgame the supports of its unique Nash equilibrium. Optimal play thus emerges from minimal rationality and inheritance alone. The result also delimits how far the exploitation of a non-equilibrium opponent can go: a player may shade probabilities within the equilibrium support in response to anticipated mistakes, but the set of actions she can plausibly use is pinned down exactly.

Neither rationality nor inheritance refers to the zero-sum structure of the game, and we derive our main result from a more general analysis. Section 3 deals with arbitrary two-player games: if all Nash equilibria of all subgames are quasi-strict—a property that holds for generic payoff matrices—the two conditions characterize the solution concepts that consistently select supports of Nash equilibria, and the remaining multiplicity, which reflects the classical coordination problem, vanishes when Nash equilibria are unique. Section 4 concerns symmetric games, for which it is natural to restrict attention to symmetric subgames, so that the one-sided reductions underlying the general analysis cannot be used. The characterization nevertheless persists, based on the fact—derived from the index theory of Nash equilibria—that unique equilibria of symmetric games have supports of odd size. Section 5 collects the consequences for zero-sum games: uniqueness of Nash equilibria is generic in this class, which yields our main result, and tournament games, whose Nash equilibria are always unique, inherit the symmetric characterization, implying that no strict refinement of the bipartisan set satisfies inheritance.

Our result is best contrasted with rationalizability (Bernheim, 1984; Pearce, 1984). Rationalizability rests on a notion of rationality that applies to each action in isolation: an action is justifiable if it is a best response to some belief about the opponent, and different actions may be justified via different beliefs. Our rationality condition applies to sets of actions instead: a set of actions is the support of an undominated strategy if and only if all of its elements are best responses to one and the same belief (Lemma 1). All plausible actions thus have to be rationalized by a common belief. In its most transparent form, this equivalence is a theorem of the alternative: either some mixed strategy is dominated, or a single belief about the opponent equalizes—and thereby simultaneously rationalizes—all of a player’s actions. This strengthens the classical observation that an action is undominated if and only if it is a best response to some belief (Lemma 3

of Pearce, 1984, see also Lemma 60.1 of Osborne and Rubinstein, 1994) from one belief per action to one belief for all actions. The interactive assumptions can be contrasted as well. Pearce and Bernheim conclude from common knowledge of rationality that players choose rationalizable actions, which, in two-player games, are exactly the actions surviving the iterated elimination of dominated actions. Inheritance can take over the role of common knowledge of rationality: every solution concept satisfying inheritance and action-wise rationality selects rationalizable actions only, and the rationalizable sets themselves constitute the largest such solution concept. Strengthening rationality from one belief per action to one belief for all plausible actions then sharpens the prediction from rationalizability all the way down to Nash equilibrium.

Neither approach requires active randomization by the players—a concern that has classically motivated the purification of mixed equilibria (Harsanyi, 1973b). For rationalizability, mixed strategies enter only through beliefs. For our result, rationality can be stated purely in terms of beliefs, and the conclusion admits the familiar interpretation of equilibrium in terms of conjectures: each player holds a single belief about the opponent and plays some pure best response to it. Since only supports are pinned down, no particular randomization is ever prescribed. The randomness instead resides in the beliefs, and here the two results differ: rationalizability justifies each action by its own, generically robust, belief, whereas our common belief must equalize all plausible actions—a knife-edge property characteristic of equilibrium conjectures. The sharper prediction thus comes at the price of knife-edge beliefs, which our theorem derives from inheritance rather than assuming them to be mutually known (cf. Aumann and Brandenburger, 1995).

Alternative approaches to characterizing equilibrium play are based on notions of consistency across games (see, e.g., Peleg and Tijs, 1996; Norde et al., 1996; Dufwenberg et al., 2001; Brandl and Brandt, 2019, 2024). In contrast to these methods, our result has the advantage of being applicable to a single game and its subgames without needing to reference additional hypothetical games. Set-valued solution concepts for zero-sum games go back to Shapley (1964); prominent examples for general games include sets that are closed under rational behavior (Basu and Weibull, 1991). The support of equilibrium play, our f^{EQ} , is known as the essential set for symmetric zero-sum games (Dutta and Laslier, 1999) and as the bipartisan set in tournament games (Laffond et al., 1993). Conditions in the spirit of inheritance have a long history in the theory of choice consistency: inheritance weakens Chernoff’s contraction condition (Chernoff, 1954), while strengthening inheritance by the reverse inclusion yields the strong superset property (Bordes, 1979), known as the outcast condition in choice theory (Aizerman and Malishevski, 1981).

2 The Model

Let $m, n \in \mathbb{N}$ be positive integers and $[m] = \{1, \dots, m\}$. For $x, y \in \mathbb{R}^k$, $x \leq y$ if $x_i \leq y_i$ for all $i \in [k]$, $x < y$ if $x \leq y$ and $x_i < y_i$ for some $i \in [k]$, and $x \ll y$ if $x_i < y_i$ for all $i \in [k]$. Moreover, $\mathbf{1}$ denotes the all-ones vector of appropriate dimension and $\text{supp}(x) = \{i \in [k] : x_i \neq 0\}$ the support of $x \in \mathbb{R}^k$. Let $\Delta(m) = \{p \in \mathbb{R}_+^m : \sum_{i \in [m]} p_i = 1\}$ be the set

of strategies over $[m]$ and, for nonempty $S \subseteq [m]$, $\Delta(S) = \{p \in \Delta(m) : \text{supp}(p) \subseteq S\}$ the set of strategies with support in S . A two-player game is a pair of matrices $A, B \in \mathbb{R}^{m \times n}$, where A_{ij} and B_{ij} are the payoffs of the row player and the column player for the action profile $(i, j) \in [m] \times [n]$; zero-sum games correspond to $B = -A$. For nonempty $S \subseteq [m]$ and $T \subseteq [n]$, the subgame $(A, B)_{S,T} = (A_{S,T}, B_{S,T})$ of (A, B) is the game in which the row player's actions are restricted to S and the column player's actions are restricted to T ; in particular, $(A, B)_{[m],[n]} = (A, B)$. We identify the strategies of the two players in $(A, B)_{S,T}$ with $\Delta(S)$ and $\Delta(T)$. A Nash equilibrium of $(A, B)_{S,T}$ is then a pair $(p, q) \in \Delta(S) \times \Delta(T)$ such that $p'^t A q \leq p^t A q$ and $p^t B q' \leq p^t B q$ for all $p' \in \Delta(S)$ and $q' \in \Delta(T)$. In products involving the matrices $A_{S,T}$ and $B_{S,T}$, strategies are restricted to the relevant coordinates, so that, for $p \in \Delta(S)$ and $q \in \Delta(T)$, $p^t A_{S,T} = ((p^t A)_j)_{j \in T}$ and $B_{S,T} q = ((B q)_i)_{i \in S}$.

A solution concept for a game (A, B) returns a subset of row player actions and a subset of column player actions for each subgame: it is a function f such that, for all nonempty $S \subseteq [m]$ and $T \subseteq [n]$, $f((A, B)_{S,T}) = (S', T')$ for some nonempty $S' \subseteq S$ and $T' \subseteq T$.¹ For such pairs of sets, we write $(S', T') \subseteq (S, T)$ if $S' \subseteq S$ and $T' \subseteq T$. If every subgame of (A, B) has a unique Nash equilibrium, one solution concept is the *Nash equilibrium support* f^{EQ} , which returns the pair of supports of the unique Nash equilibrium of every subgame.

A solution concept f for (A, B) satisfies inheritance if for all nonempty $S \subseteq [m]$ and $T \subseteq [n]$ and all nonempty $S' \subseteq S$ and $T' \subseteq T$,

$$f((A, B)_{S,T}) \subseteq (S', T') \text{ implies } f((A, B)_{S',T'}) \subseteq f((A, B)_{S,T}). \quad (\text{inheritance})$$

That is, when actions outside the solution of a subgame are discarded, the reduced subgame inherits the solution. Observe that if $f((A, B)_{S,T}) = (S', T')$, then inheritance implies $f((A, B)_{S',T'}) = (S', T')$, since $f((A, B)_{S,T}) \subseteq f((A, B)_{S',T'}) \subseteq (S', T') = f((A, B)_{S,T})$. Dominance refers to a player's own payoffs: a strategy $p \in \Delta(S)$ of the row player is dominated in $(A, B)_{S,T}$ if there exists a strategy $\tilde{p} \in \Delta(S)$ with $p^t A_{S,T} \ll \tilde{p}^t A_{S,T}$; analogously, a strategy $q \in \Delta(T)$ of the column player is dominated in $(A, B)_{S,T}$ if there exists a strategy $\tilde{q} \in \Delta(T)$ with $B_{S,T} q \ll B_{S,T} \tilde{q}$. A solution concept f for (A, B) satisfies rationality if for all nonempty $S \subseteq [m]$ and $T \subseteq [n]$,

$$f((A, B)_{S,T}) = (S', T') \text{ implies } S' = \text{supp}(p) \text{ and } T' = \text{supp}(q) \text{ for some } \quad (\text{rationality})$$

strategies $p \in \Delta(S)$ and $q \in \Delta(T)$ that are undominated in $(A, B)_{S,T}$.

As defined, rationality requires the components of $f((A, B)_{S,T})$ to be the supports of *some* undominated strategies. This existential formulation is equivalent to the universal formulation, in which *every* strategy whose support is contained in a component of $f((A, B)_{S,T})$ is required to be undominated in $(A, B)_{S,T}$. The universal formulation clearly implies the existential one. For the converse, let $f((A, B)_{S,T}) = (S', T')$ and assume that some $p \in \Delta(S')$ is dominated in $(A, B)_{S,T}$ by $\tilde{p} \in \Delta(S)$. Then, every $p' \in \Delta(S)$ with $\text{supp}(p') = S'$ is dominated in $(A, B)_{S,T}$ by $p' + \epsilon(\tilde{p} - p)$ for sufficiently

¹Formally, f is a function of the pair (S, T) , since distinct pairs can induce identical matrices.

small $\epsilon > 0$: this is a strategy in $\Delta(S)$ because $\text{supp}(p) \subseteq S' = \text{supp}(p')$, and it dominates p' because $(p' + \epsilon(\tilde{p} - p))^t A_{S,T} = p'^t A_{S,T} + \epsilon(\tilde{p}^t A_{S,T} - p'^t A_{S,T}) \gg p'^t A_{S,T}$. Hence, S' is not the support of any undominated strategy, which contradicts the existential formulation. The same argument applies to the column player. In the following, we use both formulations of rationality interchangeably.

3 Two-Player Games

Our results rely on a correspondence between dominance and beliefs, which substantiates the interpretation of rationality given in the introduction: the components of the solution are exactly the sets of actions that are best responses to a common belief. Here, a strategy $p \in \Delta(S)$ is a best response to a belief $q \in \Delta(T)$ in $(A, B)_{S,T}$ if $p'^t Aq \leq p^t Aq$ for all $p' \in \Delta(S)$, and an action $i \in S$ is a best response to q in $(A, B)_{S,T}$ if $(Aq)_i = \max_{i' \in S} (Aq)_{i'}$. Analogously, an action $j \in T$ of the column player is a best response to a belief $p \in \Delta(S)$ in $(A, B)_{S,T}$ if $(p^t B)_j = \max_{j' \in T} (p^t B)_{j'}$.

Lemma 1. *Let $S \subseteq [m]$, $T \subseteq [n]$, and $S' \subseteq S$ be nonempty. Then, S' is the support of a strategy that is undominated in $(A, B)_{S,T}$ if and only if there exists a belief $q \in \Delta(T)$ such that every action in S' is a best response to q in $(A, B)_{S,T}$. The analogous statement holds for the column player.*

Proof. First, let $q \in \Delta(T)$ such that every action in S' is a best response to q , and let $p \in \Delta(S)$ with $\text{supp}(p) = S'$. Then, $p^t Aq = \sum_{i \in S'} p_i (Aq)_i = \max_{i' \in S} (Aq)_{i'}$, so p is a best response to q . If a strategy $\tilde{p} \in \Delta(S)$ dominated p , then $\tilde{p}^t A_{S,T} \gg p^t A_{S,T}$ and hence $\tilde{p}^t Aq > p^t Aq$, contradicting that p is a best response to q . Hence, p is undominated in $(A, B)_{S,T}$.

Second, let $p \in \Delta(S)$ with $\text{supp}(p) = S'$ be undominated in $(A, B)_{S,T}$, and let $X = \{p^t A_{S,T} : p' \in \Delta(S)\} \subseteq \mathbb{R}^T$. Because p is undominated, the convex set $X - p^t A_{S,T}$ is disjoint from the open convex set $\{z \in \mathbb{R}^T : z \gg 0\}$. The separating hyperplane theorem yields $w \in \mathbb{R}^T$ with $w \neq 0$ and $c \in \mathbb{R}$ such that $w^t(x - p^t A_{S,T}) \leq c \leq w^t z$ for all $x \in X$ and $z \gg 0$ (Boyd and Vandenberghe, 2004, Section 2.5.1). Because $w^t z$ is bounded below on $\{z \in \mathbb{R}^T : z \gg 0\}$, $w \geq 0$, and hence $c \leq \inf \{w^t z : z \gg 0\} = 0$. Thus, the belief $q \in \Delta(T)$ with $q_j = w_j / (w^t \mathbf{1})$ for all $j \in T$ satisfies $p'^t Aq \leq p^t Aq$ for all $p' \in \Delta(S)$, i.e., p is a best response to q . In particular, $p^t Aq = \max_{i' \in S} (Aq)_{i'}$ and, because $p^t Aq = \sum_{i \in S'} p_i (Aq)_i$ is an average of the $(Aq)_i$ with $i \in S'$, every action in S' is a best response to q . $\square$

In its most transparent form, Lemma 1 is a theorem of the alternative. For $S' = S = [m]$ and $T = [n]$, the set $[m]$ is the support of an undominated strategy if and only if no strategy in $\Delta(m)$ is dominated, as shown in the argument for the universal formulation of rationality, and all actions of the row player are best responses to a belief $q \in \Delta(n)$ if and only if $Aq = \alpha \mathbf{1}$ for some $\alpha \in \mathbb{R}$. Hence, exactly one of the following is true: some strategy in $\Delta(m)$ is dominated, or a single belief on the column player's actions rationalizes all row player actions simultaneously. Compare this to a classical observation of Pearce (1984, Lemma 3) (see also Lemma 60.1 of Osborne and Rubinstein, 1994): a

row player action is either dominated or it is a best response to some belief on the column player's actions. Therefore, if no row player action is dominated, each such action can be rationalized by some belief; Lemma 1 strengthens this conclusion from one belief per action to one belief for all actions.

In general, inheritance and rationality cannot single out one solution concept, because Nash equilibria need not be unique even in generic games: generic two-player games have a finite odd number of Nash equilibria (Harsanyi, 1973a), and this number can be larger than one on open sets of games. Consider the coordination game

$$(A, B) = \begin{pmatrix} (2, 2) & (0, 0) \\ (0, 0) & (1, 1) \end{pmatrix}$$

with Nash equilibria $(1, 1)$ and $(2, 2)$ as well as a mixed Nash equilibrium in which both players play $(\frac{1}{3}, \frac{2}{3})$. On the proper subgames, rationality uniquely determines the solution because one of the players has a single undominated action; for the full game, each of the three pairs of Nash equilibrium supports— $(\{1\}, \{1\})$, $(\{2\}, \{2\})$, and $(\{1, 2\}, \{1, 2\})$ —yields a solution concept that satisfies inheritance and rationality. What rationality and inheritance do imply—for arbitrary two-player games and without any genericity assumption—is the following.

Lemma 2. *Let $A, B \in \mathbb{R}^{m \times n}$, and let f be a solution concept for (A, B) that satisfies inheritance and rationality. Then, for every subgame $(A, B)_{S,T}$ with $f((A, B)_{S,T}) = (S', T')$, there is a Nash equilibrium (p, q) of $(A, B)_{S,T}$ such that*

$$\text{supp}(p) \subseteq S' \subseteq BR(q) \quad \text{and} \quad \text{supp}(q) \subseteq T' \subseteq BR(p),$$

where $BR(q) = \{i \in S : (Aq)_i = \max_{i' \in S} (Aq)_{i'}\}$ and $BR(p) = \{j \in T : (p^t B)_j = \max_{j' \in T} (p^t B)_{j'}\}$ denote the sets of best responses to q and p in $(A, B)_{S,T}$.

Proof. By inheritance, $(S', T') \subseteq f((A, B)_{S,T'})$ and $(S', T') \subseteq f((A, B)_{S',T})$. By rationality and Lemma 1 applied to the subgame $(A, B)_{S,T'}$, there is a belief $q \in \Delta(T')$ to which every action in the row component of $f((A, B)_{S,T'})$ —in particular, every action in S' —is a best response, i.e., $S' \subseteq BR(q)$. Analogously, there is a belief $p \in \Delta(S')$ with $T' \subseteq BR(p)$. Then, p is a best response to q in $(A, B)_{S,T}$ because $\text{supp}(p) \subseteq S' \subseteq BR(q)$, and q is a best response to p because $\text{supp}(q) \subseteq T' \subseteq BR(p)$. Hence, (p, q) is a Nash equilibrium of $(A, B)_{S,T}$ with the required inclusions. $\square$

Lemma 2 states that all plausible actions are best responses to the conjectures of a single Nash equilibrium whose supports are themselves plausible. Note that its proof uses only two instances of inheritance per subgame, namely those in which all implausible actions of exactly one player are discarded.

A Nash equilibrium (p, q) of $(A, B)_{S,T}$ is quasi-strict if $BR(q) = \text{supp}(p)$ and $BR(p) = \text{supp}(q)$. While uniqueness of Nash equilibrium is not generic, quasi-strictness of all Nash equilibria is: for almost all payoff matrices, all Nash equilibria of all subgames are regular and hence quasi-strict (Harsanyi, 1973a; van Damme, 1991).

Theorem 1. *Let $A, B \in \mathbb{R}^{m \times n}$ such that every Nash equilibrium of every subgame of (A, B) is quasi-strict, and let f be a solution concept for (A, B) that satisfies inheritance. Then, f satisfies rationality if and only if, for every subgame, $f((A, B)_{S,T})$ is the pair of supports of some Nash equilibrium of $(A, B)_{S,T}$.*

Proof. For the direction from left to right, Lemma 2 yields a Nash equilibrium (p, q) of $(A, B)_{S,T}$ with $\text{supp}(p) \subseteq S' \subseteq BR(q)$ and $\text{supp}(q) \subseteq T' \subseteq BR(p)$, and quasi-strictness turns both chains of inclusions into equalities. For the direction from right to left, let $f((A, B)_{S,T}) = (\text{supp}(p), \text{supp}(q))$ for some Nash equilibrium (p, q) of $(A, B)_{S,T}$. Then, every action in $\text{supp}(p)$ is a best response to the belief q and every action in $\text{supp}(q)$ is a best response to the belief p , so rationality follows from Lemma 1. $\square$

Under the assumptions of Theorem 1, inheritance and rationality thus characterize the inheritance-consistent selections of Nash equilibrium supports rather than a single solution concept, and the coordination game above shows that this multiplicity cannot be avoided. The remaining freedom is the classical coordination problem: the axioms discipline the plausible actions to be supported by the conjectures of some Nash equilibrium of every subgame, but which equilibrium the players settle on is a matter of convention. This conclusion parallels the consistency-based characterization of the Nash equilibrium correspondence by Norde et al. (1996). If every subgame has a unique Nash equilibrium, however, the selection is unique and the characterization takes the following form.

Corollary 1. *Let $A, B \in \mathbb{R}^{m \times n}$ such that every subgame of (A, B) has a unique Nash equilibrium, and let f be a solution concept for (A, B) . Then, f satisfies inheritance and rationality if and only if $f = f^{\text{EQ}}$.*

Proof. Since every two-player game has a quasi-strict Nash equilibrium (Norde, 1999), the unique Nash equilibrium of every subgame of (A, B) is quasi-strict. If f satisfies inheritance and rationality, then, by Theorem 1, $f((A, B)_{S,T})$ is the pair of supports of the unique Nash equilibrium of $(A, B)_{S,T}$ for every subgame, i.e., $f = f^{\text{EQ}}$. Conversely, f^{EQ} satisfies inheritance because a Nash equilibrium of $(A, B)_{S,T}$ whose supports are contained in (S', T') is also a Nash equilibrium of $(A, B)_{S',T'}$ and hence the unique one; rationality then follows from Lemma 1. $\square$

Corollary 1 covers, for example, all games in which every subgame is solvable by the iterated elimination of strictly dominated actions, such as the prisoner's dilemma.

The analysis remains intrinsically two-player, however. In a two-player game, a belief about the opponent is itself a mixed strategy of the opponent, so a pair of mutually justifying beliefs has exactly the form of a Nash equilibrium. With three or more players, a player's belief is instead a distribution over profiles of opponents' actions. Such beliefs may correlate the opponents' actions, and different players may induce different marginal beliefs about the same third player. Thus, a direct many-player extension would require additional assumptions ensuring both agreement and independence of conjectures. This is the same two-player/many-player divide identified by Aumann and Brandenburger (1995): in two-player games, mutual knowledge of rationality, payoffs, and conjectures

suffices, whereas their many-player result uses a common prior together with common knowledge of conjectures to recover a Nash equilibrium.

4 Symmetric Games

For symmetric games, it is natural to restrict attention to symmetric subgames, in which both players have the same set of actions; the one-sided reductions used in the proof of Lemma 2 are then unavailable. A two-player game is symmetric if $n = m$ and $B = A^t$, so that the game is invariant under exchanging the roles of the players. Its symmetric subgames are the subgames $(A, A^t)_{S,S}$ for nonempty $S \subseteq [m]$, and a symmetric solution concept for the symmetric subgames of (A, A^t) is a function f with $f((A, A^t)_{S,S}) = (S', S')$ for some nonempty $S' \subseteq S$, where inheritance and rationality are imposed for symmetric subgames only. In symmetric games, both players have the same best responses to any given belief z , since $(z^t B)_j = (z^t A^t)_j = (Az)_j$; we write $BR_E(z) = \{i \in E : (Az)_i = \max_{i' \in E} (Az)_{i'}\}$ for this common set of best responses in $(A, A^t)_{E,E}$. Moreover, if (p, q) is a Nash equilibrium of a symmetric subgame, then so is (q, p) ; hence, a unique Nash equilibrium of $(A, A^t)_{S,S}$ is of the form (σ_S, σ_S) , and we refer to $\text{supp}(\sigma_S)$ as its support. The following lemma, whose proof relies on the index theory of Nash equilibria, is the key structural fact for the symmetric case.

Lemma 3. *Let (A, A^t) be a symmetric two-player game such that every symmetric subgame has a unique Nash equilibrium. Then, the support of the unique Nash equilibrium of every symmetric subgame is of odd size.*

Proof. Consider a symmetric subgame $(A, A^t)_{S,S}$ with unique Nash equilibrium (σ_S, σ_S) , and let $E = \text{supp}(\sigma_S)$. The restriction of σ_S to E is a completely mixed Nash equilibrium of $(A, A^t)_{E,E}$, which by assumption is the unique Nash equilibrium of $(A, A^t)_{E,E}$. It therefore suffices to show that a symmetric $k \times k$ game whose unique Nash equilibrium (x, x) is completely mixed has k odd; we write A for the $k \times k$ payoff matrix.

Since x is completely mixed, all pure strategies are best responses to x , so $Ax = v\mathbf{1}$ for some $v \in \mathbb{R}$. We first show that (x, x) is regular, in the sense that the matrix $\begin{pmatrix} A & -\mathbf{1} \\ \mathbf{1}^t & 0 \end{pmatrix}$ is nonsingular. Otherwise, there is $(d, \alpha) \neq 0$ with $Ad = \alpha\mathbf{1}$ and $\mathbf{1}^t d = 0$, where $d \neq 0$ because $d = 0$ would force $\alpha = 0$. Then, $x + \epsilon d$ is a completely mixed strategy for small enough $\epsilon > 0$, and $A(x + \epsilon d) = (v + \epsilon\alpha)\mathbf{1}$, so $(x + \epsilon d, x + \epsilon d)$ is a second completely mixed Nash equilibrium, contradicting uniqueness.

The equilibrium index of a Nash equilibrium was introduced by ? and can be defined via a determinant (?, Section 6): it assigns to each regular Nash equilibrium of a finite game a value in $\{+1, -1\}$ such that the indices of all Nash equilibria sum to $+1$. The determinant formula immediately yields index $(-1)^{k+1}$ for a regular completely mixed equilibrium of a symmetric $k \times k$ game. Since (x, x) is the unique Nash equilibrium and is regular, its index is $+1$; hence $(-1)^{k+1} = +1$, so k is odd.² $\square$

²For symmetric zero-sum games, where A is skew-symmetric, the index theory can be avoided: from $Ax = v\mathbf{1}$ and $v = x^t Ax = 0$ we get $Ax = 0$, so A is singular; were k even, the kernel of A would have dimension at least two, because skew-symmetric matrices have even rank, and would contain a vector

Theorem 2. *Let (A, A^t) be a symmetric two-player game such that every symmetric subgame of (A, A^t) has a unique Nash equilibrium, and let f be a symmetric solution concept for the symmetric subgames of (A, A^t) , where f^{EQ} returns the pair of supports of the unique Nash equilibrium of every symmetric subgame. Then, f satisfies inheritance and rationality if and only if $f = f^{\text{EQ}}$.*

Proof. That f^{EQ} satisfies inheritance and rationality follows as in the proof of Corollary 1, with all subgames replaced by symmetric subgames.

For the converse, let f satisfy inheritance and rationality, and consider a symmetric subgame $(A, A^t)_{S,S}$ with $f((A, A^t)_{S,S}) = (S', S')$. By inheritance, $f((A, A^t)_{S',S'}) = (S', S')$, and rationality together with Lemma 1 yields beliefs $p, q \in \Delta(S')$ such that $S' \subseteq BR_{S'}(q)$ and $S' \subseteq BR_{S'}(p)$. Then, p is a best response to q and q is a best response to p in $(A, A^t)_{S',S'}$, so (p, q) is a Nash equilibrium of $(A, A^t)_{S',S'}$ and hence coincides with the unique Nash equilibrium $(\sigma_{S'}, \sigma_{S'})$. Since unique Nash equilibria are quasi-strict (Norde, 1999), $S' \subseteq BR_{S'}(\sigma_{S'}) = \text{supp}(\sigma_{S'}) \subseteq S'$. Thus, $\text{supp}(\sigma_{S'}) = S'$, the size of S' is odd by Lemma 3, and $(A\sigma_{S'})_s = v$ for all $s \in S'$, where $v = \max_{s \in S'} (A\sigma_{S'})_s$.

Next, let $x \in S \setminus S'$ and $E = S' \cup \{x\}$. By inheritance, $(S', S') \subseteq f((A, A^t)_{E,E})$, and rationality together with Lemma 1 yields a belief $q' \in \Delta(E)$ with $S' \subseteq BR_E(q')$. We claim that $(A\sigma_{S'})_x < v$. Suppose otherwise and consider $z_\lambda = (1 - \lambda)\sigma_{S'} + \lambda q'$ for $\lambda \in [0, 1]$. At both endpoints, $(Az_\lambda)_s$ is the same for all $s \in S'$ —for $z_0 = \sigma_{S'}$ because $(A\sigma_{S'})_s = v$ for all $s \in S'$, and for $z_1 = q'$ because $S' \subseteq BR_E(q')$ —and hence for every λ ; denote this common value by u_λ . Now, $(Az_0)_x \geq u_0$ by supposition and $(Az_1)_x \leq u_1$ because $S' \subseteq BR_E(q')$, so, by continuity, $(Az_{\lambda^*})_x = u_{\lambda^*}$ for some $\lambda^* \in [0, 1]$. Then, every action in E is a best response to z_{λ^*} , so $(z_{\lambda^*}, z_{\lambda^*})$ is a Nash equilibrium of $(A, A^t)_{E,E}$ and hence coincides with (σ_E, σ_E) ; quasi-strictness implies $\text{supp}(\sigma_E) = BR_E(\sigma_E) = E$. This contradicts Lemma 3, because $|E| = |S'| + 1$ is even.

Hence, $(A\sigma_{S'})_x < v$ for all $x \in S \setminus S'$, so $BR_S(\sigma_{S'}) = S'$ and $(\sigma_{S'}, \sigma_{S'})$ is a Nash equilibrium of $(A, A^t)_{S,S}$. It therefore coincides with (σ_S, σ_S) , and $f((A, A^t)_{S,S}) = (S', S') = (\text{supp}(\sigma_S), \text{supp}(\sigma_S)) = f^{\text{EQ}}((A, A^t)_{S,S})$. $\square$

In the next section, we apply this result to symmetric zero-sum games and, in particular, to tournament games, whose Nash equilibria are always unique.

5 Zero-Sum Games

A zero-sum game is a two-player game with $B = -A$; we identify it with the row player's payoff matrix $A \in \mathbb{R}^{m \times n}$, so that A_{ij} is the row player's payoff and $-A_{ij}$ the column player's payoff for the action profile (i, j) . The column player's best responses to a belief $p \in \Delta(S)$ in $A_{S,T}$ are then the actions $j \in T$ that minimize $(p^t A)_j$, and a strategy $q \in \Delta(T)$ of the column player is dominated in $A_{S,T}$ if there exists a strategy $\tilde{q} \in \Delta(T)$ with $A_{S,T} \tilde{q} \ll A_{S,T} q$. Call a zero-sum game $A \in \mathbb{R}^{m \times n}$ *generic* if every subgame $A_{S,T}$

d with $\mathbf{1}^t d = 0$ that is not a multiple of x , making $x + \varepsilon d$ a second completely mixed Nash equilibrium for small $\varepsilon > 0$ and contradicting uniqueness.

has a unique Nash equilibrium.³ In contrast to arbitrary two-player games, uniqueness of Nash equilibrium is generic for zero-sum games: the set of non-generic zero-sum games is contained in a finite union of lower-dimensional algebraic subsets of $\mathbb{R}^{m \times n}$ (zero sets of nonzero polynomials in the entries of A) and therefore has measure 0 and is nowhere dense. Our main result for zero-sum games is now an immediate consequence of Corollary 1.

Corollary 2. *Let $A \in \mathbb{R}^{m \times n}$ be a generic zero-sum game, and let f be a solution concept for A . Then, f satisfies inheritance and rationality if and only if $f = f^{\text{EQ}}$.*

Proof. By definition, every subgame of a generic zero-sum game has a unique Nash equilibrium, so the statement follows from Corollary 1.⁴ $\square$

Remark 1 (Independence of the axioms). Rationality and inheritance are independent. The solution concept that returns, in every subgame, the actions that maximize a player’s total payoff over the opponent’s available actions—equivalently, the best responses to the uniform belief about the opponent—satisfies rationality by Lemma 1, but violates inheritance. For example, for the generic game

$$A = \begin{pmatrix} 1 & 2 \\ 4 & -2 \end{pmatrix},$$

it returns $(\{1\}, \{2\})$, since the first row maximizes $\sum_j A_{ij}$ and the second column minimizes $\sum_i A_{ij}$; but for the subgame $A_{\{1\}, \{1,2\}}$, it returns $(\{1\}, \{1\})$, because the column player’s best response to the first row is the first column. Hence, $(\{1\}, \{2\}) \subseteq (\{1\}, \{1,2\})$, but $(\{1\}, \{2\}) \not\subseteq (\{1\}, \{1\})$. Conversely, the trivial solution concept, which returns (S, T) for every subgame $A_{S,T}$, satisfies inheritance, because $f(A_{S,T}) \subseteq (S', T')$ implies $(S', T') = (S, T)$, but it violates rationality for every generic game in which one of the players has more than one action: if the column player has two available actions, then any one-row, two-column subgame has distinct payoffs by genericity, and every full-support column strategy is dominated by the pure action with the lower payoff for the row player; the case of two row actions in a one-column subgame is analogous.

A zero-sum game is symmetric as a two-player game if and only if $A = -A^t$, i.e., A is skew-symmetric: then $B = -A = A^t$, so the game is both zero-sum and symmetric. Theorem 2 therefore applies to symmetric zero-sum games without any further assumption on the supports.

³For zero-sum games, this is the subgame-wise version of regularity: a finite zero-sum game is regular if and only if it has a unique Nash equilibrium. Indeed, if there are two equilibria, the convexity of the sets of optimal strategies yields a continuum of equilibria, so equilibria are not isolated. Conversely, a unique zero-sum equilibrium is quasi-strict: any unused pure best response could be added to an equilibrium strategy to obtain another equilibrium. If the indifference equations on the equilibrium supports were singular, a nonzero support-preserving perturbation would keep all used actions indifferent and, by quasi-strictness, would yield another equilibrium for sufficiently small perturbations.

⁴For zero-sum games, the existence of quasi-strict Nash equilibria used in the proof of Corollary 1 goes back to Bohnenblust et al. (1950).

Corollary 3. *Let $A \in \mathbb{R}^{m \times m}$ be a symmetric zero-sum game such that every symmetric subgame of A has a unique Nash equilibrium, and let f be a symmetric solution concept for the symmetric subgames of A . Then, f satisfies inheritance and rationality if and only if $f = f^{\text{EQ}}$.*

Proof. Since A is skew-symmetric, the game is symmetric as a two-player game, so the statement is a special case of Theorem 2. $\square$

By a theorem of Laffond et al. (1997), every symmetric zero-sum game whose off-diagonal entries are odd integers has a unique Nash equilibrium. Since this property is inherited by symmetric subgames, all such games satisfy the assumption of Corollary 3. The most prominent examples are *tournament games*, whose off-diagonal entries are 1 or -1 , and Corollary 3 then has an interesting consequence for tournament solutions. A tournament is a complete and asymmetric relation $\succ$ on a finite set X of alternatives; its tournament game is the symmetric zero-sum game $A \in \mathbb{R}^{X \times X}$ with $A_{ab} = 1$ if $a \succ b$, $A_{ab} = -1$ if $b \succ a$, and $A_{aa} = 0$. The support of the unique Nash equilibrium (p^*, p^*) of a tournament game is known as the bipartisan set $BP(X)$ (Laffond et al., 1993). A *tournament solution* F maps every tournament to a nonempty subset of its alternatives; for $S \subseteq X$, we write $F(S)$ for the choice of F from the restriction of $\succ$ to S . F refines BP if $F(S) \subseteq BP(S)$ for every tournament on every set S , and F satisfies $\hat{\alpha}_{\supseteq}$, the inheritance part of the strong superset property $\hat{\alpha}$, if $F(X) \subseteq S \subseteq X$ implies $F(X) \subseteq F(S)$. Brandt et al. (2018, Corollary 2) have shown that no strict refinement of BP satisfies stability, i.e., $\hat{\alpha}$ and $\hat{\gamma}$. This impossibility already holds when only $\hat{\alpha}_{\supseteq}$ is required.

Corollary 4. *No strict refinement of BP satisfies $\hat{\alpha}_{\supseteq}$.*

Proof. Let F be a tournament solution that refines BP and satisfies $\hat{\alpha}_{\supseteq}$, and consider the tournament game A of an arbitrary tournament on a finite set X . Then, f with $f(A_{S,S}) = (F(S), F(S))$ for all nonempty $S \subseteq X$ is a symmetric solution concept for the symmetric subgames of A , and $\hat{\alpha}_{\supseteq}$ is precisely inheritance for f . Moreover, f satisfies rationality: because the unique Nash equilibrium (p_S, p_S) of $A_{S,S}$ is quasi-strict, every action in $BP(S)$ —and hence every action in $F(S) \subseteq BP(S)$ —is a best response to the belief p_S in $A_{S,S}$, so $F(S)$ is the support of an undominated strategy by Lemma 1. Hence, Corollary 3 implies $f = f^{\text{EQ}}$, and, in particular, $F(X) = BP(X)$. Since the tournament was arbitrary, $F = BP$. $\square$

Acknowledgments

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/14-1 and EXC-2047. This project was inspired by unpublished work on impartial ordinalism and the essential set by Klaus Nehring.

References

M. Aizerman and A. V. Malishevski. General theory of best variants choice: Some aspects. *IEEE Transactions on Automatic Control*, 26(5):1030–1040, 1981.

- R. J. Aumann and A. Brandenburger. Epistemic conditions for Nash equilibrium. *Econometrica*, 63(5):1161–1180, 1995.
- K. Basu and J. W. Weibull. Strategy subsets closed under rational behavior. *Economics Letters*, 36:141–146, 1991.
- B. D. Bernheim. Rationalizable strategic behavior. *Econometrica*, 52(4):1007–1028, 1984.
- H. F. Bohnenblust, S. Karlin, and L. S. Shapley. Solutions of discrete, two-person games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, volume 1, pages 51–72. Princeton University Press, 1950.
- G. Bordes. Some more results on consistency, rationality and collective choice. In J.-J. Laffont, editor, *Aggregation and Revelation of Preferences*, chapter 10, pages 175–197. North-Holland, 1979.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- F. Brandl and F. Brandt. Justifying optimal play via consistency. *Theoretical Economics*, 14(4):1185–1201, 2019.
- F. Brandl and F. Brandt. An axiomatic characterization of Nash equilibrium. *Theoretical Economics*, 19(4):1473–1504, 2024.
- F. Brandt, M. Brill, H. G. Seedig, and W. Suksompong. On the structure of stable tournament solutions. *Economic Theory*, 65(2):483–507, 2018.
- H. Chernoff. Rational selection of decision functions. *Econometrica*, 22(4):422–443, 1954.
- M. Dufwenberg, H. Norde, H. Reijnierse, and S. H. Tijs. The consistency principle for set-valued solutions and a new direction for normative game theory. *Mathematical Methods of Operations Research*, 54(1):119–131, 2001.
- B. Dutta and J.-F. Laslier. Comparison functions and choice correspondences. *Social Choice and Welfare*, 16(4):513–532, 1999.
- J. C. Harsanyi. Oddness of the number of equilibrium points: A new proof. *International Journal of Game Theory*, 2(1):235–250, 1973a.
- J. C. Harsanyi. Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory*, 2(1):1–23, 1973b.
- G. Laffond, J.-F. Laslier, and M. Le Breton. More on the tournament equilibrium set. *Mathématiques et sciences humaines*, 31(123):37–44, 1993.
- G. Laffond, J.-F. Laslier, and M. Le Breton. A theorem on symmetric two-player zero-sum games. *Journal of Economic Theory*, 72(2):426–431, 1997.
- H. Norde. Bimatrix games have quasi-strict equilibria. *Mathematical Programming*, 85: 35–49, 1999.

- H. Norde, J. Potters, H. Reijnierse, and D. Vermeulen. Equilibrium selection and consistency. *Games and Economic Behavior*, 12(2):219–225, 1996.
- M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- D. G. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050, 1984.
- B. Peleg and S. H. Tijs. The consistency principle for games in strategic form. *International Journal of Game Theory*, 25(1):13–34, 1996.
- L. S. Shapley. Some topics in two-person games. In M. Dresher, L. S. Shapley, and A. W. Tucker, editors, *Advances in Game Theory*, volume 52 of *Annals of Mathematics Studies*, pages 1–29. Princeton University Press, 1964.
- E. van Damme. *Stability and Perfection of Nash Equilibria*. Springer-Verlag, 2nd edition, 1991.
- J. von Neumann. Zur Theorie der Gesellschaftspiele. *Mathematische Annalen*, 100(1):295–320, 1928.