

A Robust Characterization of Nash Equilibrium

Florian Brandl Felix Brandt
University of Bonn Technical University of Munich

We give a robust characterization of Nash equilibrium by postulating coherent behavior across varying games: Nash equilibrium is the only solution concept that satisfies consequentialism, consistency, and rationality. As a result, every equilibrium refinement violates at least one of these properties. We moreover show that every solution concept that approximately satisfies consequentialism, consistency, and rationality returns approximate Nash equilibria. The latter approximation can be made arbitrarily good by increasing the approximation of the axioms. This result extends to various natural subclasses of games such as two-player zero-sum games, potential games, and graphical games.

1. Introduction

More than 70 years after the publication of Nash’s (1951) original work, the concept of Nash equilibrium has been engraved in economic reasoning so deeply that it is rarely questioned. But what makes Nash equilibrium stand out from the plethora of solution concepts that have been proposed? The game theory literature has come up with various answers to this question based on different approaches.

In this paper, we take a normative approach: we consider solution concepts for games in normal form with a fixed number of players and formulate conditions for solution concepts that capture coherent behavior across different games. The solution concepts we consider map every (finite) multi-player game to a non-empty set of (mixed) strategy profiles. Our first of three axioms requires that the labels of actions are irrelevant—only the payoffs matter. Call two actions of a player clones if, irrespective of the other players’ actions, they give the same payoff to all players. In other words, clones are outcome-equivalent and only discernible by their labels.

Consequentialism. A player can shift probability arbitrarily between clones, and deleting a clone neither changes the probabilities assigned to the player’s other actions nor the strategies of the other players.

If a solution concept satisfying consequentialism returns a strategy profile and we modify the game by cloning an action, then the solution concept has to return all strategy profiles in which the probabilities on the player’s uncloned actions and the other players’ strategies are unchanged.

The second axiom is motivated by situations where the players are uncertain which game will be played.

Consistency. Every strategy profile that is played in two given games with the same sets of actions for each player is also played when a coin toss decides which of the two games is played and the players choose their strategies before the coin toss.

Instead of modeling the randomization explicitly, we assume that a coin toss between two games is equivalent to the convex combination of these games.

Third, we stipulate a weak notion of rationality.

Rationality. Players never put positive probability on actions that are dominated in pure strategies.

That is, if one action of a player has a higher payoff than another action irrespective of the other players’ strategies, the latter is played with probability 0. In fact, it suffices for our characterization to assume that dominated actions receive probability at most, say, one half.

Our main result characterizes Nash equilibrium as the unique solution concept that satisfies consequentialism, consistency, and rationality. In particular, players’ behavior has to be consistent with expected utility maximization, which is not apparent from the axioms. Moreover, every refinement of Nash equilibrium violates at least one of the axioms.

Our second result shows that this characterization is robust: every solution concept that approximately satisfies the three axioms is approximately Nash equilibrium. This type of stability result is common in many areas of mathematics.¹ To make this precise, we formulate quantitatively relaxed versions of the axioms. δ -consequentialism demands that a player can shift probability arbitrarily between clones and deleting a clone does not change the probability on any player’s action (except the cloned action) by more than δ . Similarly, δ -consistency requires that if a strategy profile is played in two games, some strategy profile that differs by no more than δ is played in any convex combination of the two games. Finally, δ -rationality implies that actions that are dominated in pure strategies by at least a margin of δ are played with probability at most one half. We show that for every positive ε , there exists a positive δ such that every solution concept that satisfies the δ -versions of consequentialism, consistency, and rationality is a refinement of ε -Nash equilibrium. This result implies one direction of our main theorem. Moreover, it holds for various natural subclasses of games such as two-player zero-sum games, potential games, or graphical games.

¹A classic example is isoperimetric stability. In Euclidean space, a ball is characterized as the unique volume-maximizing shape among all well-behaved shapes with the same surface area. Isoperimetric stability strengthens this assertion by showing that any shape that is close to volume maximizing has to resemble a sphere.

2. Related Work

Which assumptions can be used to justify Nash equilibrium has been primarily studied in epistemic game theory. In this stream of research, the knowledge of individual players is modeled using Bayesian belief hierarchies, which consist of a game and a set of types for each player with each type including the action played by this type and a probability distribution over types of the other players, called the belief of this type (Harsanyi, 1967). Rather than assuming that players actively randomize, the beliefs about the types of the other players are randomized. Players are rational if they maximize expected payoff given their types and beliefs. Aumann and Brandenburger (1995) have shown that for two-player games the beliefs of every pair of types whose beliefs are mutually known and whose rationality is mutually known constitute a Nash equilibrium. This result extends to games with more than two players if the beliefs are commonly known and admit a common prior. Common knowledge assumptions in game theory have been criticized for not adequately modeling reality (see, e.g., Gintis, 2009). Barelli (2009), Hellman (2013), and Bach and Tsakas (2014) showed that the results of Aumann and Brandenburger (1995) still hold under somewhat weaker common knowledge assumptions.

Building on earlier work by Peleg and Tijs (1996), Norde et al. (1996) have characterized Nash equilibrium via one-player rationality (only utility-maximizing strategies are returned in one-player games) and a consistency condition that is orthogonal to ours because it varies the set of *players*. Their condition requires that every strategy profile s returned for an n -player game is also returned for the $(n - k)$ -player game that results when k players invariably play their strategies in s . The two axioms immediately imply that only subsets of Nash equilibria can be returned. Their results have no implications for games with a fixed number of players.

The work most closely related to ours is due to Brandl and Brandt (2019) who have characterized *maximin strategies* in two-player zero-sum games by consequentialism, consistency, and rationality. The difference between their results and ours are as follows. Solution concepts as considered by Brandl and Brandt return a set of strategies for one player, rather than a set of strategy *profiles*. Noting that Nash equilibria in zero-sum games consist of pairs of maximin strategies, their result translates to the terminology of the present paper as follows: *in zero-sum games*, every solution concept that satisfies consequentialism, consistency, and rationality returns an (exchangeable) *subset* of Nash equilibria.² Our main theorem, Theorem 1, is stronger since it *(i)* holds for any number of players, *(ii)* shows that *all* Nash equilibria have to be returned (and thus rules out equilibrium refinements), and *(iii)* is not restricted to games with rational-valued payoffs and rational-valued strategies (which are assumptions needed for the proof of Brandl and Brandt). Moreover, we show that the containment in the set of Nash equilibria *(iv)* also holds for restricted classes of games (cf. Remark 4), and *(v)* is robust with respect to small violations of the axioms (Theorem 2).

²A set of strategy profiles is exchangeable if it is a Cartesian product of a set of strategies for each player.

3. The Model

Let U be an infinite universal set of actions and denote by $\mathcal{F}(U)$ the set of finite and nonempty subsets of U . For $A \in \mathcal{F}(U)$, let Σ_A be the set of all permutations of U that fix $U \setminus A$ pointwise. If $p \in \mathbb{R}_+^U$, $\text{supp}(p) = \{a \in U : p(a) > 0\}$ denotes the support of p . Moreover, let

$$\Delta A = \{p \in \mathbb{R}_+^U : \text{supp}(p) \subseteq A \text{ and } \sum_{a \in A} p(a) = 1\}$$

be the set of probability distributions on U that are supported on A . We call ΔA the set of strategies for action set A . For two strategies $p, q \in \Delta A$, let $\|p - q\| = \sum_{a \in A} |p(a) - q(a)|$ be their ℓ_1 -distance. The ball of radius $\delta > 0$ around a set $S \subseteq \Delta A$ is $B_\delta(S) = \{p \in \Delta A : \inf\{\|p - q\| : q \in S\} < \delta\}$, the set of strategies supported on A and less than δ away from some strategy in S .³

Let $N = \{1, \dots, n\}$ be the set of players. For action sets $A_1, \dots, A_n \in \mathcal{F}(U)$, we write $A = A_1 \times \dots \times A_n$ for the corresponding set of action profiles. A game on A is a function $G : A \rightarrow \mathbb{R}^n$. For $i \in N$ and $a \in A$, $G_i(a)$ is the payoff of player i for the action profile a . We say that G is normalized if for every player i , either G_i has minimum 0 and maximum 1 or is constant at 1. Strategies and norms extend to A as follows: $\square A = \Delta A_1 \times \dots \times \Delta A_n$ and for $p, q \in \square A$, $\|p - q\| = \max_{i \in N} \|p_i - q_i\|$. We call $\square A$ the set of (strategy) profiles. The players' payoffs for a strategy profile are the corresponding expected payoffs. Thus, a strategy profile $p \in \square A$ is a Nash equilibrium of G if

$$G_i(p_i, p_{-i}) \geq G_i(q_i, p_{-i}) \text{ for all } q_i \in \Delta A_i \text{ and } i \in N.$$

For two games G and G' on $A = A_1 \times \dots \times A_n$ and $A' = A'_1 \times \dots \times A'_n$, we say that G is a blow-up of G' (G' is a blow-down of G) if every action of every player in G is payoff-equivalent to one of her actions in G' . That is, there are surjective functions $\phi_i : A_i \rightarrow A'_i$, $i \in N$, such that with $\phi = (\phi_1, \dots, \phi_n)$, $G = G' \circ \phi$. Actions in $\phi_i^{-1}(a'_i)$ for $a'_i \in A'_i$ are called clones of a'_i . Put differently, G is obtained from G' by “blowing up” each action a'_i to $|\phi_i^{-1}(a'_i)|$ clones of it.⁴ A strategy $p_i \in \Delta A_i$ induces a strategy on A'_i via the pushforward of ϕ_i : $(\phi_i)_*(p_i) = p_i \circ \phi_i^{-1}$. Then, a strategy profile $p \in \square A$ induces $\phi_*(p) = ((\phi_1)_*(p_1), \dots, (\phi_n)_*(p_n))$.

A solution concept f maps every game to a set of strategy profiles. That is, $f(G) \subseteq \square A$ for a game G on $A \in \mathcal{F}(U)^n$. If $f(G) \neq \emptyset$ for all G , f is a *total* solution concept. An example of a total solution concept is *NASH*, which returns all strategy profiles that constitute Nash equilibria. Nash (1951) has shown that every game admits at least one Nash equilibrium.

$$\text{NASH}(G) = \{p \in \square A : p \text{ is a Nash equilibrium of } G\}.$$

³Note that $B_\delta(S)$ depends on A . In our usage, A will be clear from the context.

⁴Kohlberg and Mertens (1986) have considered a more permissive notion of “blowing down” in the context of Nash equilibrium refinements for extensive-form games. Their notion of a reduced form of a normal-form game allows deleting any action that is a convex combination of other actions.

4. Characterization of Nash Equilibrium

This section defines our axioms and states the characterization of Nash equilibrium along with the more illuminating part of its proof. The remainder of the proof and all other proofs are given in the Appendix.

Consequentialism requires that if G is a blow-up of G' , a strategy profile is returned in G if and only if its pushforward is returned in G' . Equivalently, it asserts that (i) cloning an action does not change the probabilities of other actions and the strategies of the other players, and (ii) the probability on the cloned action can be distributed arbitrarily on its clones.

Definition 1 (Consequentialism). A solution concept f satisfies consequentialism if for all games G and G' such that G is a blow-up of G' with surjection $\phi = (\phi_1, \dots, \phi_n)$,

$$f(G) = \phi_*^{-1}(f(G')).$$

Consequentialism is a common desideratum in decision theory. It corresponds to the conjunction of Chernoff's (1954) *Postulate 6* (cloning of a player's actions) and *Postulate 9* (cloning of Nature's states, i.e., of opponent's actions). The latter also appears as *column duplication* (Milnor, 1954) and *deletion of repetitious states* (Arrow and Hurwicz, 1972; Maskin, 1979). In the context of social choice theory, a related condition called *independence of clones* was introduced by Tideman (1987) (see also Zavist and Tideman, 1989; Brandl et al., 2016).

If G and G' are games on the same action sets, then ϕ_i permutes the actions of player i . This special case of consequentialism, called equivariance, implies that relabeling the actions of a player results in the same relabeling of her strategies.

Definition 2 (Equivariance). A solution concept f satisfies equivariance if for all games G on A and all $\pi = (\pi_1, \dots, \pi_n)$ where π_i is a permutation of A_i ,⁵

$$f(G \circ \pi) = f(G) \circ \pi.$$

We will frequently apply equivariance to strategy profiles where each player's strategy is the uniform distribution on some subset of her actions and the permutations map each action to an action with the same probability, thus giving a new game for which the same strategy profile is returned.

Consistency requires that if a strategy profile is returned in several games with the same action sets, it should also be returned in any convex combination of these games. An inductive argument shows that this is equivalent to the restriction of the axiom to convex combinations of only two games.

Definition 3 (Consistency). A solution concept f satisfies consistency if for all games G^1, \dots, G^k on A and any $\lambda \in \mathbb{R}_+^k$ with $\sum_j \lambda_j = 1$,

$$f(G^1) \cap \dots \cap f(G^k) \subseteq f\left(\lambda_1 G^1 + \dots + \lambda_k G^k\right).$$

⁵For a strategy $p_i \in \Delta(A_i)$, $p_i \circ \pi_i$ is the strategy with $(p_i \circ \pi_i)(a_i) = p_i(\pi(a_i))$. For a strategy profile $p = (p_1, \dots, p_n)$, $p \circ \pi = (p_1 \circ \pi_1, \dots, p_n \circ \pi_n)$, and this operation extends to sets of strategy profiles pointwise.

We are not aware of game-theoretic work using this consistency axiom other than that by Brandl et al. (2016). Chernoff considers combinations of decision-theoretic situations obtained by taking unions of action sets. His *Postulate 9* states that any action that is chosen in two situations should also be chosen in such a combination. In our context, this translates to a consistency condition on the support of strategies and varying sets of actions. Closer analogs of consistency, involving convex combinations of distributions over states (i.e., strategies of Nature), have been considered as decision-theoretic axioms (see, e.g., Chernoff, 1954; Milnor, 1954; Gilboa and Schmeidler, 2003). Shapley’s (1953) characterization of the Shapley value involves an additivity axiom (which he calls *law of aggregation*) that is similar in spirit to consistency. Lastly, analogs of consistency feature prominently in several axiomatic characterizations in social choice theory, where it relates the choices for different sets of voters to each other (see, e.g., Smith, 1973; Young and Levenglick, 1978; Myerson, 1995; Brandl et al., 2016; Lackner and Skowron, 2021).

For a game G on A and two actions $a_i, a'_i \in A_i$, we say that a_i *dominates* a'_i if $G_i(a_i, a_{-i}) > G_i(a'_i, a_{-i})$ for all $a_{-i} \in A_{-i}$; a_i is *undominated* if no action dominates it. Rationality requires that only strategy profiles in which dominated actions receive probability at most $1/2$ are returned.⁶

Definition 4 (Rationality). A solution concept f satisfies rationality if for all games G ,

$$f(G) \subseteq B_{1/2}(\Delta(\hat{A}_1)) \times \cdots \times B_{1/2}(\Delta(\hat{A}_n)),$$

where \hat{A}_i denotes the set of undominated actions of player i in G .

Note that rationality is not concerned with *mixed* strategies and thus does not rely on *expected* payoffs. Moreover, it does not need any assumptions about other players. The strengthening of rationality requiring that dominated action receive probability 0 is equivalent to Milnor’s (1954) *strong domination*, Maskin’s (1979) *Property (5)*, and weaker than Chernoff’s (1954) *Postulate 2*.

As a shorthand, we say that a solution concept is *nice* if it satisfies these three axioms. It turns out that Nash equilibrium is the only nice total solution concept.

Theorem 1. *Let f be a total solution concept that satisfies consequentialism, consistency, and rationality. Then, $f = \text{NASH}$.*

Remark 1 (Totality). Totality is only required for the inclusion $\text{NASH} \subseteq f$. The converse inclusion $f \subseteq \text{NASH}$ also holds for solution concepts that fail to be total.

Remark 2 (Equilibrium refinements). One consequence of Theorem 1 is that every refinement of Nash equilibrium violates at least one of the axioms (including totality). We discuss some examples. Since rationality is preserved under taking subsets, every

⁶For Theorem 1, it suffices to require that the probability on dominated actions is below 1. When studying the robustness of this characterization in Theorem 2, it becomes important to bound the probability on dominated actions away from 1. Other than that, any fixed bound smaller than 1 in place of $1/2$ would work as well.

refinement satisfies rationality. *Quasi-strict equilibrium* (Harsanyi, 1973) satisfies consistency, and, for two players, is total (Norde, 1999). However, it violates consequentialism (even for two players) since it does not allow for the possibility that clones of equilibrium actions are played with probability 0. A trivial example is a game where all players' utility functions are constant for all action profiles. Then, every *full support* strategy profile is a quasi-strict equilibrium, whereas consequentialism requires that every strategy profile is returned. For three or more players, quasi-strict equilibria may not exist. *Trembling-hand perfect equilibrium* (Selten, 1975) is total and satisfies consequentialism. Consequently, it is not consistent. Lastly, *strong equilibrium* (Aumann, 1959) and *coalition-proof equilibrium* (Bernheim et al., 1987) also satisfy consequentialism and violate consistency (and fail to be total even for two players).

Remark 3 (Independence of axioms). All properties in Theorem 1 are required to derive the conclusion. For each of the four axioms (including totality), there is a solution concept different from *NASH* that satisfies the three remaining axioms.

- (i) *Consequentialism*: return all strategy profiles in which every player randomizes only over actions that are best responses against uniformly randomizing opponents; satisfies totality, consistency, and rationality but violates consequentialism.
- (ii) *Consistency*: return all strategy profiles in which every player randomizes only over actions that maximize this player's highest possible payoff; satisfies totality, consequentialism, and rationality but violates consistency.
- (iii) *Rationality*: return all strategy profiles that maximize the sum of all players' payoffs; satisfies totality, consequentialism, and consistency but violates rationality.
- (iv) *Totality*: return all strategy profiles whose pushforwards are pure Nash equilibria in a blowdown of the original game; satisfies consequentialism, rationality, and consistency but violates totality.

The first three examples are neither contained in nor contain *NASH*. The last one is necessarily a refinement of *NASH* due to Remark 1. Further examples that *are* refinements or coarsenings of *NASH* are not hard to find. Quasi-strict equilibrium (for two players) violates consequentialism but satisfies consistency and rationality as discussed in Remark 2. Trembling-hand perfect equilibrium is not consistent but satisfies the other two axioms. The trivial solution concept returning all strategy profiles in all games violates rationality but satisfies the remaining two axioms. Note that rationality is so weak that even for one-player games, all three axioms are required for the characterization.

Remark 4 (Restricted classes of games). Examining the proof of the inclusion $f \subseteq \text{NASH}$, one can see that it remains valid for any class of games that is closed under blowing-up, blowing-down, and taking convex combinations. More precisely, it holds for any class of games \mathcal{G} with the following properties.

- (i) If G is a blow-up of G' , then $G \in \mathcal{G}$ if and only if $G' \in \mathcal{G}$.

- (ii) If $G_1, \dots, G_k \in \mathcal{G}$ are games on the same action profiles and $\lambda \in \mathbb{R}_+^k$ with $\sum_j \lambda_j = 1$, then $\lambda_1 G_1 + \dots + \lambda_k G_k \in \mathcal{G}$.

Various well-known classes of games satisfy these properties, for example, (strategically) zero-sum games, graphical games, and potential games. By contrast, symmetric games are not closed with respect to blow-ups.⁷ For example, cloning or permuting actions of only one player makes a symmetric game asymmetric. More generally, it is unclear how to extend the current proof approach to symmetric games.⁸

For our proof of the converse inclusion, $NASH \subseteq f$, a class of games needs to have enough games with a unique equilibrium. Suppressing technicalities, it is required that for every game $G \in \mathcal{G}$ and every equilibrium p of G , G can be written as a convex combination of games in \mathcal{G} that have p as the unique equilibrium. We have not examined which classes of games, other than the class of all games, have this property.

Remark 5 (Closures of solution concepts). One can “repair” any given solution concept by iteratively adding strategy profiles whenever there is a failure of consequentialism or consistency. Equivalently, one can define the closure of a solution concept f as the smallest solution concept containing f and satisfying consequentialism and consistency.⁹ By Theorem 1, the closure of a total refinement of Nash equilibrium is Nash equilibrium, and the closure of total non-refinements violates rationality.

The proof of Theorem 1 uses a lemma that is technical to state but illustrates how one can manipulate games using the above axioms. It shows that nice solution concepts behave as one would hope under the analog of row and column operations familiar from linear algebra. More precisely, it shows that if one adds a linear combination of some actions (with positive rational-valued coefficients) to another action, then a nice solution concept shifts probability from the former actions to the latter in proportion to the coefficients. A linear combination of actions here means a linear combination of the corresponding payoffs for all players. Figure 2 below illustrates this. A similar conclusion applies to adding new actions that are linear combinations of existing ones.

⁷A game is symmetric if all players have the same set of actions and permuting the actions in any action profile results in the same permutation of the players’ payoffs.

⁸For example, consider the following symmetric two-player game.

$$\begin{array}{c} 0 \quad 1/2 \quad 1/2 \\ 0 \quad \begin{pmatrix} 3,3 & 2,2 & 2,2 \\ 2,2 & 3,3 & 0,0 \\ 2,2 & 0,0 & 3,3 \end{pmatrix} \\ 1/2 \\ 1/2 \end{array}$$

The indicated strategy profile is not a Nash equilibrium. However, all symmetric games that are blow-ups of this game and convex combinations thereof have the same payoffs on the diagonal. Hence, clones of the second and third action of either player are not dominated in any of these games, so that no contradictions to rationality occur. This issue does not arise for symmetric *zero-sum* games (see Brandl and Brandt, 2019, Remark 3).

⁹This closure is well-defined since consequentialism and consistency are preserved under taking intersections of solution concepts.




Figure 1: Schematic depiction of the games G , \tilde{G} , and \hat{G} constructed in the proof of Lemma 1 (with $n = 2$, $|A_i| = 2$, $k_i = (2, 2)$, and $\hat{a}_i \notin A_i$ for $i = 1, 2$). \tilde{G} is obtained from G by adding $k_i(a_i)$ clones of every action a_i of player i . Then, an intermediate game \tilde{G} is constructed from \tilde{G} by permuting the actions outside of A_i and summing over the resulting games. The actions outside of A_i are now clones obtained from a convex combination (with weights k_i) of actions in A_i . Removing all but one of these clones gives \hat{G} .

Lemma 1. *Let f be a nice solution concept, G be a game on A , and $p \in f(G)$. Let $\hat{a} \in U^N$, $k \in (\mathbb{Z}_+^U \setminus \{0\})^N$, and $\kappa \in \mathbb{R}_+^N$ so that for all $i \in N$, $k_i(\hat{a}_i) > 0$ if $\hat{a}_i \in A_i$, $\text{supp}(k_i) \subseteq A_i$, $x_i := \kappa_i k_i \leq p_i$, and $x_i(\hat{a}_i) = p_i(\hat{a}_i)$. Then, there is a game \hat{G} on \hat{A} with $\hat{A}_i = A_i \cup \{\hat{a}_i\}$ so that the following holds.*

1. $\hat{p} \in f(\hat{G})$, where $\hat{p}_i = p_i - x_i + |x_i| \cdot e_{\hat{a}_i}$.¹⁰
2. For all $I \subseteq N$ and $a \in A$, $\hat{G}(\hat{a}_I, a_{-I}) = G((\frac{k_i}{|k_i|})_{i \in I}, a_{-I})$.

Condition 1 states that \hat{p}_i is obtained from p_i by shifting probability $x_i(a_i)$ from a_i to \hat{a}_i . Condition 2 ensures that playing \hat{a}_i in \hat{G} is payoff-equivalent to playing $\frac{k_i}{|k_i|}$ in G . Figure 1 illustrates the proof of Lemma 1.

Proof. For all $i \in N$ and $a_i \in A_i$, let $\hat{A}_i^{a_i} \subseteq U$ so that $|\hat{A}_i^{a_i}| = k_i(a_i)$ if $a_i \neq \hat{a}_i$, $|\hat{A}_i^{\hat{a}_i}| = \max\{k_i(\hat{a}_i) - 1, 0\}$, and all $\hat{A}_i^{a_i}$ are disjoint and disjoint from $A_i^- := A_i \setminus \{\hat{a}_i\}$. Let $\tilde{A}_i = A_i \cup \bigcup_{a_i \in A_i} \hat{A}_i^{a_i}$ and $\phi_i: \tilde{A}_i \rightarrow A_i$ so that $\phi_i^{-1}(a_i) = \{a_i\} \cup \hat{A}_i^{a_i}$. Let \tilde{G} be a game on $\tilde{A} = \tilde{A}_1 \times \cdots \times \tilde{A}_n$ so that \tilde{G} is a blow-up of G with surjection $\phi = (\phi_1, \dots, \phi_n)$. Consequentialism implies that $\tilde{p} \in \phi_*^{-1}(p) \subseteq f(\tilde{G})$, where $\tilde{p}_i = p_i - x_i + |x_i| \cdot \text{uni}(\tilde{A}_i \setminus A_i^-)$.

For all $i \in N$, let $\Sigma_i \subseteq \Sigma_{\tilde{A}_i}$ be the set of all permutations π_i of \tilde{A}_i so that π_i is the

¹⁰By $e_{\hat{a}_i}$, we denote the standard unit vector in \mathbb{R}^U with a 1 in position \hat{a}_i .

	G				\hat{G}		
	1/3	1/3	1/3		1/3	1/3	1/3
0	1, 0	1, 1	1, 0	0	1, 0	1, 1	1, 0
1/3	0, 1	2, 0	0, 1	1/6	0, 1	2, 0	0, 1
1/3	1, 1	1, 0	0, 0	1/6	1, 1	1, 0	0, 0
1/3	1, 0	0, 0	1, 1	2/3	$\frac{3}{4}, \frac{1}{2}$	$\frac{3}{4}, 0$	$\frac{1}{4}, \frac{3}{4}$

Figure 2: Example for an application of Lemma 1. Here, one half of the second and third action of the first player are added to the fourth action. That is, \hat{a}_1 is the fourth action, $k_1 = (0, 1, 1, 2)$, $\kappa_1 = 1/6$, and $x_1 = (0, 1/6, 1/6, 1/3)$; \hat{a}_2 is arbitrary, say, the first action of player 2, $k_2 = (1, 0, 0, 0)$, $\kappa_2 = 1$, and $x_2 = (1, 0, 0, 0)$.

identity on A_i^- , and let $\Sigma = \Sigma_1 \times \cdots \times \Sigma_n$. Let

$$\bar{G} = \frac{1}{|\Sigma|} \sum_{\pi \in \Sigma} \tilde{G} \circ \pi.$$

Then all actions in $\tilde{A}_i \setminus A_i^-$ are clones of each other in \bar{G} . Since \tilde{p}_i assigns the same probability to all actions in $\tilde{A}_i \setminus A_i^-$ and f is equivariant as a consequence of satisfying consequentialism, $\tilde{p} \in f(\tilde{G} \circ \pi)$ for all $\pi \in \Sigma$. Consistency then implies that $\tilde{p} \in f(\bar{G})$. Moreover, for all $I \subseteq N$, $a \in A$, and $\tilde{a} \in (\tilde{A}_1 \setminus A_1^-) \times \cdots \times (\tilde{A}_n \setminus A_n^-)$,

$$\bar{G}(\tilde{a}_I, a_{-I}) = G\left(\left(\frac{k_i}{|k_i|}\right)_{i \in I}, a_{-I}\right),$$

since any action in $\tilde{A}_i \setminus A_i^-$ is the convex combination of actions in A_i with coefficients $x_i = \kappa_i k_i$.

Lastly, one can delete all but one of the clones of any action in $\tilde{A}_i \setminus A_i^-$. For all $i \in N$, let $\hat{\phi}_i: \tilde{A}_i \rightarrow \hat{A}_i$ so that $\hat{\phi}_i$ is the identity A_i^- and $\hat{\phi}_i^{-1}(\hat{a}_i) = \tilde{A}_i \setminus A_i^-$. Let \hat{G} be a blow-up of \tilde{G} with surjection $\hat{\phi} = (\hat{\phi}_1, \dots, \hat{\phi}_n)$. Note that $\hat{p} = \hat{\phi}_*(\tilde{p})$. Consequentialism thus gives $\hat{p} \in f(\hat{G})$. Moreover, for all $I \subseteq N$ and $a \in A$,

$$\hat{G}(\hat{a}_I, a_{-I}) = \bar{G}(\hat{a}_I, a_{-I}) = G\left(\left(\frac{k_i}{|k_i|}\right)_{i \in I}, a_{-I}\right).$$

□

Starting with the assumption that f returns a non-equilibrium strategy profile in some game G , the proof of Theorem 1 has two steps, each of which uses Lemma 1. First, we construct from G a game \bar{G} where a non-equilibrium profile is returned and every player plays some distinguished action with probability close to 1 with the distinguished

action of one player, say j , not being a best response. The second step is to replace every other action of every player other than j by a convex combination of itself and the distinguished action (with sufficiently large weight on the latter). This results in a game \hat{G} in which j 's distinguished action is dominated, violating rationality. With a third application of Lemma 1 (replacing every action of j by a convex combination of itself and j 's distinguished action), one could construct a game where j plays dominated actions with probability *equal to* 1. We omit this step since it is not required for our notion of rationality.

Proof of Theorem 1. We prove that $f \subseteq \text{NASH}$. The proof of $\text{NASH} \subseteq f$ is given in Appendix A.

Let G be a game on A and $p \in f(G)$. Assume that G is normalized and $A_1 = \dots = A_n = B$. The former is for convenience with obvious adjustments for non-normalized G and the latter is without loss of generality since f satisfies consequentialism.

Assume for contradiction that $p \notin \text{NASH}(G)$. Then, there is a player $j \in N$ for whom p_j is not a best response. That is, $G_j(a_j^*, p_{-j}) - G_j(p) > \varepsilon > 0$ for some $a_j^* \in A_j$. Let $\bar{a} \in (U \setminus A_1) \times \dots \times (U \setminus A_n)$, $k \in (\mathbb{Z}_+^U \setminus \{0\})^N$, and $\kappa \in \mathbb{R}_+^N$ so that for all $i \in N$, $x_i := \kappa_i k_i \leq p_i$ and $|x_i| \geq 1 - \frac{\varepsilon^2}{4n}$. Note that then $|p_i - \frac{k_i}{|k_i|}| \leq \frac{\varepsilon^2}{4n}$. By Lemma 1, there is a game \bar{G} on \bar{A} with $\bar{A}_i = A_i \cup \{\bar{a}_i\}$ so that

1. $\bar{p} \in f(\bar{G})$, where $\bar{p}_i = p_i - x_i + |x_i| \cdot e_{\bar{a}_i}$, and
2. for all $I \subseteq N$ and $a \in A$, $\bar{G}(\bar{a}_I, a_{-I}) = G((\frac{k_i}{|k_i|})_{i \in I}, a_{-I})$.

In particular, $\bar{p}_j(\bar{a}_j) \geq 1 - \frac{\varepsilon^2}{4n}$ and, since G is normalized,

$$\bar{G}_j(a_j^*, \bar{a}_{-j}) - \bar{G}_j(\bar{a}_j, \bar{a}_{-j}) > \frac{\varepsilon}{2}.$$

The second step is to modify \bar{G} so that a_j^* dominates \bar{a}_j by replacing every action of every player $i \neq j$ by a convex combination of itself and \bar{a}_i with a sufficiently high weight on \bar{a}_i . For every $b \in B$, let $k^b \in (\mathbb{Z}_+^U \setminus \{0\})^N$ and $\kappa^b \in \mathbb{R}_+^N$ so that for all $i \in N \setminus \{j\}$, $k_i^b = e_b + \frac{2n}{\varepsilon} \cdot e_{\bar{a}_i}$ and $\kappa_i^b k_i^b(b) = \bar{p}_i(b)$, and $k_j^b = e_b$ and $\kappa_j^b = \bar{p}_j(b)$ (this choice for j is to make the application of Lemma 1 trivial for player j). Note that for $i \neq j$ and $x_i^b := \kappa_i^b k_i^b$,

$$\sum_{b \in A_i} x_i^b(\bar{a}_i) = \frac{2n}{\varepsilon} \sum_{b \in A_i} x_i^b(b) = \frac{2n}{\varepsilon} \sum_{b \in A_i} \bar{p}_i(b) \leq \frac{2n}{\varepsilon} \cdot \frac{\varepsilon^2}{4n} = \frac{\varepsilon}{2}.$$

Sequential application of Lemma 1 to \bar{G} , one for each $b \in B$ with $\hat{a} = (b, \dots, b)$, gives a game \hat{G} on \hat{A} so that

1. $\hat{p} \in f(\hat{G})$, where $\hat{p}_i = \bar{p}_i - \frac{2n}{\varepsilon}(1 - |x_i|) \cdot e_{\bar{a}_i} + \frac{2n}{\varepsilon} \sum_{b \in A_i} (p_i(b) - x_i(b)) \cdot e_b$ for $i \neq j$ and $\hat{p}_j = \bar{p}_j$, and
2. for all $I \subseteq N \setminus \{j\}$ and $a \in A$, $\hat{G}(a_j, a_I, \bar{a}_{-(I \cup \{j\})}) = \bar{G}(a_j, (\frac{k_i^{a_i}}{|k_i^{a_i}|})_{i \in I}, \bar{a}_{-(I \cup \{j\})})$.

The last part of the first statement implies that \hat{p}_j assigns probability close to 1 to \bar{a}_j , more precisely, $\hat{p}_j(\bar{a}_j) = \bar{p}_j(\bar{a}_j) \geq 1 - \varepsilon$. The second statement means that player $i \neq j$ playing a_i in \hat{G} is payoff equivalent to playing a_i with probability $\frac{1}{1+2n/\varepsilon}$ and \bar{a}_i with probability $\frac{2n/\varepsilon}{1+2n/\varepsilon}$ in \bar{G} . Thus, using again that G is normalized, we have for all $a \in \bar{A}$,

$$|\hat{G}_j(a_j, a_{-j}) - \bar{G}_j(a_j, \bar{a}_{-j})| \leq n \cdot \frac{1}{1+2n/\varepsilon} < \frac{\varepsilon}{2}.$$

It follows that for all $a \in \bar{A}$,

$$\hat{G}_j(a_j^*, a_{-j}) - \hat{G}_j(\bar{a}_j, a_{-j}) > 0.$$

That is, j plays the dominated action \bar{a}_j with probability at least $1 - \varepsilon$ in \hat{G} . This contradicts rationality. \square

5. Robustness of the Characterization

The goal of this section is to examine whether the characterization of the preceding section is robust with respect to small violations of the axioms. Observe that robustness is not a foregone conclusion. The fact that the three sets of solution concepts defined by our three axioms intersect only in one point, $NASH$, does not imply that the intersection of slight thickenings of these sets only contains points close to $NASH$. We start by defining approximate equilibria and approximate versions of the axioms.

The standard notion of an approximate equilibrium is ε -equilibrium. A strategy profile is an ε -equilibrium if no player can deviate to a strategy that increases her payoff by more than ε . By $NASH_\varepsilon$ we denote the solution concept returning all ε -equilibria in all games.

Definition 5 (ε -equilibrium). Let G be a game on $A = A_1 \times \cdots \times A_n$. A profile $p \in \square A$ is an ε -equilibrium of G if

$$G_i(p_i, p_{-i}) \geq G_i(q_i, p_{-i}) - \varepsilon \text{ for all } q_i \in \Delta A_i \text{ and } i \in N.$$

Likewise, there are natural approximate notions of consequentialism, consistency, and rationality. They are obtained from the exact versions defined in the preceding section by allowing for small perturbations of strategy profiles.

Recall that consequentialism requires that if G is a blow-up of G' , then a profile is returned in G if and only if its pushforward is returned in G' . Approximate consequentialism weakens this condition by requiring only that the set of returned profiles for G is close (in Hausdorff distance) to the set of profiles whose pushforward is returned in G' .

Definition 6 (δ -consequentialism). A solution concept f satisfies δ -consequentialism if for all games G and G' such that G is a blow-up of G' with surjection $\phi = (\phi_1, \dots, \phi_n)$, $f(G) = \phi_*^{-1}(\phi_*(f(G)))$, and

$$f(G) \subseteq B_\delta(\phi_*^{-1}(f(G'))) \text{ and } \phi_*^{-1}(f(G')) \subseteq B_\delta(f(G)).$$

The first assertion requires that probability can be distributed arbitrarily over clones. The first set inclusion asserts that in the game obtained by cloning actions, the solution concept can only return profiles that differ by no more than δ from some profile obtained as a blow-up of a profile that is returned in the original game. Conversely, the second inclusion requires that every blow-up of a profile returned in the original game differs by at most δ from some profile returned in the blown-up game.

Approximate consistency weakens its exact counterpart by requiring only that if a profile is returned in several games, then some profile close to it has to be returned in any convex combination of these games.

Definition 7 (δ -consistency). A solution concept f satisfies δ -consistency if for all games G^1, \dots, G^k on the same action profiles and every $\lambda \in \mathbb{R}_+^k$ with $\sum_j \lambda_j = 1$,

$$f(G^1) \cap \dots \cap f(G^k) \subseteq B_\delta \left(f \left(\lambda_1 G^1 + \dots + \lambda_k G^k \right) \right).$$

Unlike for the exact notion, δ -consistency as defined is not equivalent to its restriction to $k = 2$.¹¹

Lastly, approximate rationality asserts that actions that are dominated by a non-negligible amount are not played too frequently. If G is a game on $A = A_1 \times \dots \times A_n$ and $a_i, a'_i \in A_i$ are actions of player i , we say that a_i δ -dominates a'_i if $G(a_i, a_{-i}) \geq G(a'_i, a_{-i}) + \delta$ for all $a_{-i} \in A_{-i}$.

Definition 8 (δ -rationality). A solution concept f satisfies δ -rationality if for all games G ,

$$f(G) \subseteq B_{1/2}(\hat{A}_1^\delta) \times \dots \times B_{1/2}(\hat{A}_n^\delta),$$

where \hat{A}_i^δ denotes the set of actions of player $i \in N$ that are not δ -dominated in G .

Each of the δ -axioms coincides with the exact version defined in the previous section when $\delta = 0$. We call a solution concept δ -nice if it satisfies δ -consequentialism, δ -consistency, and δ -rationality. It is routine to check that $NASH_\varepsilon$ is δ -nice for small enough δ . Conversely, we show that for small enough δ , every δ -nice solution concept is a refinement of $NASH_\varepsilon$. The statement is restricted to normalized games and requires equivariance for reasons that we discuss in Remarks 6 and 7.

Theorem 2. *Consider solution concepts on the set of normalized games. Then, for every $\varepsilon > 0$, there is $\delta > 0$ so that if f is equivariant and satisfies δ -consequentialism, δ -consistency, and δ -rationality, then f is a refinement of $NASH_\varepsilon$.*

¹¹A proof attempt by induction fails for two reasons. First, every application of δ -consistency introduces an additive error of δ , so that $k - 1$ applications to two games only gives an error bound of $(k - 1)\delta$ on the right hand side. Second, even if $f(G_1)$, $f(G_2)$, and $f(G_3)$ have a non-empty common intersection, $f(\frac{\lambda_1}{\lambda_1 + \lambda_2} G_1 + \frac{\lambda_2}{\lambda_1 + \lambda_2} G_2)$ need not intersect with $f(G_3)$, making a further application of δ -consistency useless.

Note that δ in Theorem 2 does not depend on f . Otherwise the statement would follow from the fact that exact consequentialism, consistency, and rationality characterize Nash equilibrium. The proof is similar to that of Theorem 1. However, apart from the need to keep track of error terms arising from applications of the axioms, some steps require additional care. The proof appears in Appendix B. The comments on subclasses of games in Remark 4 remain valid for Theorem 2.

Remark 6 (Normalized games). Theorem 2 is restricted to normalized games since ε -equilibrium becomes too stringent without a bound on the payoffs. More precisely, for any game G , the set of ε -equilibria of cG shrinks to the set of exact equilibria of G as c goes to positive infinity. To see that the restriction to normalized games is indeed necessary, consider the following solution concept f for two-player games. Let $\hat{G} = (0, 0 \quad 0, -c)$ be a game where the first player has only one action, the second player has two actions, and c is a large positive number ($c \gg \frac{1}{\delta}$), and let $\hat{p} = (1, (1 - \delta, \delta))$. Observe that \hat{p} is not an ε -equilibrium for \hat{G} . Define $f(G) = \text{NASH}(G) \cup \{\hat{p}\}$ if $G = \hat{G}$ and $f(G) = \text{NASH}(G)$ otherwise. It is not hard to see that f is δ -nice.¹²

Remark 7 (Equivariance). Applying δ -consequentialism to two games that are the same up to permuting actions, one can see that it implies δ -equivariance defined analogously to δ -consequentialism. However, our proof requires exact equivariance. Whether Theorem 2 holds without this assumption is open.

Remark 8 (Converse of Theorem 2). In contrast to Theorem 1, a characterization of NASH_ε as the only δ -nice solution for some δ is not possible. If NASH_ε is δ -nice for some δ , then so is $\text{NASH}_{\varepsilon'}$ for all $\varepsilon' \leq \varepsilon$. A weaker converse to Theorem 2 would require that for every $\delta > 0$, there is $\varepsilon > 0$ so that if f is δ -nice, then $f = \text{NASH}_{\varepsilon'}$ for some $\varepsilon' \leq \varepsilon$. While this statement is obviously false in general since δ -niceness is vacuous for large δ , we do not know if it holds for small enough δ .

Remark 9 (Alternative approximate equilibrium notions). Theorem 2 fails if ε -equilibrium is replaced by some alternative notions of approximate equilibrium. For example, Theorem 2 does not hold when replacing NASH_ε by $B_\varepsilon(\text{NASH})$, the set of profiles that are ε -close to some Nash equilibria. This is because $\text{NASH}_{\varepsilon'}$ is δ -nice for small enough ε' , however, $\text{NASH}_{\varepsilon'} \not\subseteq B_\varepsilon(\text{NASH})$ for any $\varepsilon' > 0$. In words, no matter how small ε' is, there exist ε' -equilibria that are more than ε away from every Nash equilibrium.

Acknowledgments

Florian Brandl acknowledges support by the DFG under the Excellence Strategy EXC-2047. Felix Brandt acknowledges support by the DFG under grants BR 2312/11-2

¹² δ -consequentialism and δ -rationality follow from the fact that NASH is nice and the definition of the axioms. To verify δ -consistency, it suffices to consider convex combinations of games with equilibrium \hat{p} involving \hat{G} . The only games on the same action sets as \hat{G} with equilibrium \hat{p} are those where both actions of the second player give her the same payoff. Any convex combination of such games with \hat{G} has $(1, (1, 0))$ as the unique equilibrium, which conforms with δ -consistency.

and BR 2312/12-1. The authors thank Francisc Dilmé, Benny Moldovanu, and Lucas Pahl for helpful feedback. A preliminary version of this paper was presented at the Interdisciplinary CIREQ-Workshop at Université de Montréal (Montréal, March 2023) and the Microeconomic Theory Workshop at the University of Bonn (Bonn, May 2023).

References

- K. J. Arrow and L. Hurwicz. An optimality criterion of decision-making under ignorance. In C. F. Carter and J. L. Ford, editors, *Uncertainty and expectations in economics: essays in honour of G.L.S. Shackle*, pages 1–11. Basil Blackwell, 1972.
- R. J. Aumann. Acceptable points in general cooperative n-person games. In A. W. Tucker and R. D. Luce, editors, *Contributions to the Theory of Games IV*, volume 40 of *Annals of Mathematics Studies*, pages 287–324. Princeton University Press, 1959.
- R. J. Aumann and A. Brandenburger. Epistemic conditions for Nash equilibrium. *Econometrica*, 63(5):1161–1180, 1995.
- C. W. Bach and E. Tsakas. Pairwise epistemic conditions for Nash equilibrium. *Games and Economic Behavior*, 85:48–59, 2014.
- P. Barelli. Consistency of beliefs and epistemic conditions for nash and correlated equilibria. *Games and Economic Behavior*, 67(2):363–375, 2009.
- B. D. Bernheim, B. Peleg, and M. D. Whinston. Coalition-proof nash equilibria I. Concepts. *Journal of Economic Theory*, 42(1):1–12, 1987.
- F. Brandl and F. Brandt. Justifying optimal play via consistency. *Theoretical Economics*, 14(4):1185–1201, 2019.
- F. Brandl, F. Brandt, and H. G. Seedig. Consistent probabilistic social choice. *Econometrica*, 84(5):1839–1880, 2016.
- H. Chernoff. Rational selection of decision functions. *Econometrica*, 22(4):422–443, 1954.
- L.-B. Cui, W. Li, and M. K. Ng. Birkhoff–von Neumann theorem for multistochastic tensors. *SIAM Journal of Matrix Analysis and Applications*, 35(3), 2014.
- I. Gilboa and D. Schmeidler. A derivation of expected utility maximization in the context of a game. *Games and Economic Behavior*, 44(1):172–182, 2003.
- H. Gintis. *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences*. Princeton University Press, 2009.
- J. C. Harsanyi. Games with incomplete information played by “Bayesian” players, part I. *Management Science*, 50(12):1804–1817, 1967.

- J. C. Harsanyi. Oddness of the number of equilibrium points: A new proof. *International Journal of Game Theory*, 2(1):235–250, 1973.
- Z. Hellman. Weakly rational expectations. *Journal of Mathematical Economics*, 49(6):496–500, 2013.
- E. Kohlberg and J.-F. Mertens. On the strategic stability of equilibria. *Econometrica*, 54:1003–1037, 1986.
- M. Lackner and P. Skowron. Consistent approval-based multi-winner rules. *Journal of Economic Theory*, 192:105173, 2021.
- E. Maskin. Decision-making under ignorance with implications for social choice. *Theory and Decision*, 11(3):319–337, 1979.
- J. Milnor. Games against nature. In *Decision Processes*, chapter 4, pages 49–59. Wiley, 1954.
- R. B. Myerson. Axiomatic derivation of scoring rules without the ordering assumption. *Social Choice and Welfare*, 12(1):59–74, 1995.
- J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.
- 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.
- B. Peleg and S. Tijs. The consistency principle for games in strategic form. *International Journal of Game Theory*, 25(1):13–34, 1996.
- A. Schrijver. *Theory of Linear and Integer Programming*. Wiley, 1998.
- R. Selten. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1):25–55, 1975.
- L. S. Shapley. A value for n-person games. *Annals of Math Studies*, 28:307–317, 1953.
- J. H. Smith. Aggregation of preferences with variable electorate. *Econometrica*, 41(6):1027–1041, 1973.
- T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.
- H. P. Young and A. Levenglick. A consistent extension of Condorcet’s election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978.
- T. M. Zavist and T. N. Tideman. Complete independence of clones in the ranked pairs rule. *Social Choice and Welfare*, 6(2):167–173, 1989.

APPENDIX

A. Omitted Proof From Section 4

We prove the missing direction of Theorem 1, that is, $NASH \subseteq f$. The main idea of the proof is simple: for every game G and every equilibrium p of G , show that G can be written as a convex combination of games for which p is the unique equilibrium. Since $f \subseteq NASH$ and f is total, we know that f has to return unique equilibria. Consistency thus gives $p \in f(G)$.

The work lies in finding a suitable representation as a convex combination. A first observation is that it suffices to prove that $NASH \subseteq f$ holds for games where the payoff functions of all players but one are 0. More formally, we say that G is a player i payoff game if for all $j \neq i$, $G_j \equiv 0$. Then, the following holds.

Lemma 2 (Reduction to player i payoff games). *Let G be a game and $p \in NASH(G)$. For $i \in N$, let G^i be the game with $G_i^i = G_i$ and $G_j^i \equiv 0$ for all $j \neq i$. Then, $p \in NASH(G^i)$.*

Proof. First, p_i is a best response to p_{-i} in G^i since it is a best response in G and $G_i^i = G_i$. Second, for all $j \neq i$, p_j (and any other strategy for that matter) is a best response to p_{-j} in G^i since $G_j^i \equiv 0$. Hence, $p \in NASH(G)$. \square

So if we can show that for every player i payoff game G , $NASH(G) \subseteq f(G)$, we can use Lemma 2 and consistency to conclude that the same conclusion holds for all games. While this reduction is convenient, it is not as powerful as it may seem since player i payoff games have a unique equilibrium only if all players other than i have only a single action. Thus, even when decomposing player i payoff games into games with a unique equilibrium, one needs to consider games with non-zero payoffs for all players.

A.1. Reduction to Deterministic Slice-Stochastic Tensors

The next step is a further reduction showing that it is sufficient to consider the case when G_i is a *slice-stochastic tensor*. To motivate this notion, recall that the well-known Birkhoff-von Neumann theorem states that every bistochastic matrix can be written as a convex combination of permutation matrices.¹³ There are different ways one might try to generalize this statement to higher-order tensors. For example, one might say that a tensor $T: A_1 \times \cdots \times A_n \rightarrow \mathbb{R}_+$ is n -stochastic if for all $i \in N$ and $a_{-i} \in A_{-i}$, $\sum_{a_i \in A_i} T(a_i, a_{-i}) = 1$ (which is to say that every “tube” of T sums to 1). However, with this definition, for $n \geq 3$, it is not true that every n -stochastic tensor can be written as a convex combination of n -stochastic tensors taking values in $\{0, 1\}$ (see Cui et al., 2014). We thus opt for a different generalization of bistochastic matrices.

Definition 9 (Slice-stochastic tensors). Let $A \in \mathcal{F}(U)^n$ with $|A_1| = \cdots = |A_n|$. A tensor $T: A \rightarrow \mathbb{R}$ is *slice-stochastic* for $i \in N$ if

¹³A matrix $M \in \mathbb{R}_+^{m \times m}$ is bistochastic if the rows sums and column sums are 1.

- (i) for all $a_{-i} \in A_{-i}$, $\sum_{a_i \in A_i} T(a_i, a_{-i}) = 1$,
- (ii) for all $a_i \in A_i$, $\sum_{a_{-i} \in A_{-i}} T(a_i, a_{-i}) = m^{n-2}$, and
- (iii) for all $a \in A$, $0 \leq T(a) \leq 1$.

We say that T is a *deterministic* slice-stochastic tensor if it is slice-stochastic and takes values in $\{0, 1\}$.

For $n = 2$, T is a bistochastic matrix if and only if it is slice-stochastic for some $i = 1, 2$. Note that if T is slice-stochastic for i , then

$$\sum_{a \in A} T(a) = \sum_{a_{-i} \in A_{-i}} 1 = \sum_{a_i \in A_i} m^{n-2} = m^{n-1}.$$

We omit writing “for i ” when i is clear from the context.

It turns out that the Birkhoff-von Neumann theorem does extend to slice-stochastic tensors of any order. That is, every slice-stochastic tensor is a convex combination of deterministic slice-stochastic tensors. This will allow us to reduce the problem $NASH \subseteq f$ to payoff functions of the latter type.

Lemma 3 (Birkhoff-von Neumann theorem for slice-stochastic tensors). *Let $A \in \mathcal{F}(U)^n$ with $|A_1| = \dots = |A_n|$. Let $T: A \rightarrow \mathbb{R}$ be a slice-stochastic tensor for $i \in N$. Then, there are tensors $T^1, \dots, T^K: A \rightarrow \{0, 1\}$ that are slice-stochastic for i and $(\lambda^1, \dots, \lambda^K) \in \Delta([K])$ so that*

$$T = \sum_{k \in [K]} \lambda^k T^k.$$

Proof. Viewing T as an element of \mathbb{R}^A , T is slice-stochastic for i if it is a solution to the linear feasibility program

$$Mx \leq v,$$

where the matrix M and the vector v are given by the constraints of type (i), (ii), and (iii). Thus, M has $2|A_{-i}| + 2|A_i| + 2|A| = 2(m^{n-1} + m + m^n)$ rows (2 for each constraint of each of the three types) and $|A| = m^n$ columns; the number of rows of v is the same as for M . Note that M is of the form $M = (\tilde{M}^\top, -\tilde{M}^\top)^\top$ for some matrix \tilde{M} , since each constraint gives two rows in M where one is the negative of the other. (For $n = 3$ and $m = 2$, the matrix \tilde{M} is depicted in Figure 3.)

We want to show that the polytope defined by $Mx \leq v$ has integral vertices. Since v is integral (in fact, $\{-1, 0, 1\}$), it suffices to show that M is totally unimodular.¹⁴ But M is totally unimodular if and only if \tilde{M} is totally unimodular. Now \tilde{M} is totally unimodular if for every subset R of rows of \tilde{M} , there is an assignment $\sigma: R \rightarrow \{-1, 1\}$ of signs to the rows in R so that for all $a \in A$,

$$\sum_{r \in R} \sigma(r) \tilde{M}_{r,a} \in \{-1, 0, 1\}. \tag{1}$$

¹⁴ M is totally unimodular if every square submatrix of M has determinant -1 , 0 , or 1 .

$$\begin{array}{l} \text{type (i)} \\ \text{type (ii)} \\ \text{type (iii)} \end{array} \left\{ \begin{array}{cccccc} 1 & 1 & & & & \\ & & 1 & 1 & & \\ & & & & 1 & 1 \\ & & & & & 1 & 1 \\ \hline 1 & & 1 & & 1 & & 1 \\ & 1 & & 1 & & 1 & \\ \hline & & & I_8 & & & \end{array} \right)$$

Figure 3: The matrix \tilde{M} in the proof of Lemma 3 for $n = 3$ and $m = 2$. I_8 denotes the identity matrix with $8 = m^n$ rows and columns. Each of the four pairs of columns separated by dashed lines corresponds to some $a_{-i} \in A_{-i}$. One can check that \tilde{M} is totally unimodular by verifying the equivalent condition (1).

A proof of this result appears for example in the book by Schrijver (1998).

This condition is easy to check in the present case. Let R be a subset of rows of \tilde{M} . It is easy to see that we may assume that R does not contain rows corresponding to constraints of type (iii) since those rows only contain of single 1 or -1 and can thus always be signed so as to not introduce violations of (1). We then define σ as follows.

- For each $r \in R$ corresponding to a constraint of type (i) for $a_{-i} \in A_{-i}$, let $\sigma(r) = 1$.
- For each $r \in R$ corresponding to a constraint of type (ii) for $a_i \in A_i$, let $\sigma(r) = -1$.

Then, for each column index $a \in A$, there is a most one 1 and at most one -1 in the sum in (1), which concludes the proof. \square

To show that considering games where every player's payoff function is slice-stochastic games is sufficient, we examine how certain changes to the payoff functions influence the set of equilibria. For $\alpha \in \mathbb{R}_{++}^N$ and $\beta \in \mathbb{R}^N$, we write $\alpha G + \beta$ for the game with payoff function $(\alpha G + \beta)_i = \alpha_i G_i + \beta_i$ for all $i \in N$. Moreover, we say that $T: A \rightarrow \mathbb{R}$ is *constant* for $i \in N$ if for all $a_{-i} \in A_{-i}$, $T(\cdot, a_{-i})$ is constant. Then, a game G is *constant* if for all $i \in N$, G_i is constant for i .

Lemma 4 (Invariance of equilibria). *Let G be a game on A . Then,*

- (i) *for all $\alpha \in \mathbb{R}_{++}^N$ and $\beta \in \mathbb{R}^N$, $\text{NASH}(G) = \text{NASH}(\alpha G + \beta)$, and*
- (ii) *for all constant games \tilde{G} on A , $\text{NASH}(G) = \text{NASH}(G + \tilde{G})$.*

Proof. The statement follows from a straightforward calculation. \square

For $n = 2$, G is constant if every column for G_1 and every row of G_2 is constant. Hence, Lemma 4(ii) asserts that adding a constant to some column of G_1 or row of G_2 does not change the set of equilibria.

The second type of modification of payoff functions concerns multiplication by tensors. The following notation will be convenient.

Definition 10 (Hadamard product). Let $A \in \mathcal{F}(U)^n$ and $T, T': A \rightarrow \mathbb{R}$. Then, the Hadamard product $T * T': A \rightarrow \mathbb{R}$ of T and T' is defined by

$$(T * T')(a) = T(a)T'(a)$$

for all $a \in A$.

Lemma 5 (Contortion of equilibria). Let G be a game on A . Let $i \in N$, $q \in \mathbb{R}_{++}^{A_i}$, and $T_q: A \rightarrow \mathbb{R}$ so that for all $a \in A$, $T_q(a) = q(a_i)$. Then,

$$p \in \text{NASH}(G) \quad \text{if and only if} \quad (\tilde{p}_i, p_{-i}) \in \text{NASH}(G^{i, T_q}),$$

where $G_j^{i, T_q} = G_j$, for all $j \neq i$, $G_j^{i, T_q} = G_j * T_q$, and for all $a_i \in A_i$, $\tilde{p}_i(a_i) = (p_i(a_i)/q(a_i))/(\sum_{a'_i \in A_i} p_i(a'_i)/q(a'_i))$.

Proof. The statement follows from a straightforward calculation. \square

Note that T_q in Lemma 5 is constant for all $j \neq i$. For $n = 2$, the game G^{1, T_q} is obtained from G by multiplying the a_1 row of the matrix G_2 by $q(a_1)$ for all $a_1 \in A_1$.

The next lemma show that in some sense, it is enough to consider games where the payoff function of every player is slice-stochastic. This is explained in more detail after the lemma.

Lemma 6 (Universality of slice-stochastic tensors). Let $A \in \mathcal{F}(U)^n$ with $|A_1| = \dots = |A_n|$. Let $p \in \square A$ so that all p_i have full support. Then, there are $T_1, \dots, T_n: A \rightarrow \mathbb{R}_{++}$ so that T_i is constant for all $j \neq i$ and for all games G on A and $p \in \text{NASH}(G)$, there is a game \bar{G} on A so that the following hold.

(i) For all $i \in N$, $\bar{G}_i = G_i * T^{-i} + S_i$, where T^{-i} is the Hadamard product of all T_j with $j \neq i$, and $S_i: A \rightarrow \mathbb{R}$ is constant for i .

(ii) For all $i \in N$, \bar{G}_i is slice-stochastic for i .

(iii) $\bar{p} \in \text{NASH}(\bar{G})$, where \bar{p} and \bar{p}_i is the uniform distribution on A_i .

Proof. For all $i \in N$, let $T_i: A \rightarrow \mathbb{R}_{++}$ so that for all $a \in A$, $T_i(a) = \varepsilon p_i(a_i)$, where $\varepsilon > 0$ is small compared to (the reciprocal of) the largest payoff in G and $|A| = m^n$. Define a game \hat{G} on A so that for all $i \in N$, $\hat{G}_i = G_i * T^{-i}$. It follows from Lemma 5 that $\bar{p} \in \text{NASH}(\hat{G})$.

For all $i \in N$, let $S_i: A \rightarrow \mathbb{R}$ so that for all $a \in A$,

$$S_i(a) = \frac{1}{m} \left(1 - \sum_{a'_i \in A_i} \hat{G}_i(a'_i, a_{-i}) \right).$$

Note that S_i is constant for i .

Let \bar{G} be the game so that for all $i \in N$, $\bar{G}_i = \hat{G}_i + S_i$. By Lemma 4(ii), $\bar{p} \in \text{NASH}(\bar{G})$ and so for all $i \in N$, $\sum_{a_{-i} \in A_{-i}} \bar{G}_i(\cdot, a_{-i})$ is constant on A_i . Moreover, for all $a_{-i} \in A_{-i}$, $\sum_{a_i \in A_i} \bar{G}_i(a_i, a_{-i}) = 1$ by the choice of S_i . Lastly, since ε is small, $\hat{G}_i(a) \approx 0$ and $S_i(a) \approx \frac{1}{m}$ for all $a \in A$. Hence, \bar{G}_i is slice-stochastic. We have thus shown that \bar{G} satisfies all three conditions in the statement of the lemma. \square

The condition (iii) in Lemma 6 is redundant since the strategy profile where every player distributes uniformly automatically an equilibrium if the payoff function of every player is slice-stochastic (by Item (ii) in the definition of slice-stochasticity). We chose to state it to make it explicit.

The way in which we will use Lemma 6 is as follows. Given a game G with the same number of actions for every player and a full support equilibrium $p \in \text{NASH}(G)$, we want to show that $p \in f(G)$ if f is a total solution concept satisfying consequentialism, consistency, and rationality. To do this, we show that the game \bar{G} obtained from the lemma (by virtue of having slice-stochastic tensors as payoff functions) can essentially be written as a convex combination of games for which \bar{p} is the unique equilibrium. Applying the inverse transformation of that in (i) to each of the summands and using Lemma 5 shows that G can essentially be written as a convex combination of games for which p is the unique equilibrium. Using consistency and the fact that $f \subseteq \text{NASH}$, this shows that $p \in f(G)$. The caveat “essentially” refers to the fact that one may need to multiply the G_i by positive scalars, add constant games to G , and clone actions to get the desired representation as a convex combination. Since neither of the first two operations changes the set of equilibria and the effect of introducing clones is controlled by consequentialism, this does not cause problems. Similarly, the restriction that every player has the same number of actions in G is without loss of generality by consequentialism.

Together, Lemma 2, Lemma 3, and Lemma 6 show the following. If we want to show that $p \in f(G)$ whenever $p \in \text{NASH}(G)$ and p has full support, then it is enough to do so in the case when G is a player i payoff game with G_i deterministic slice-stochastic. The full support assumption will be successively eliminated via Lemma 10, Lemma 11, and Lemma 12.

A.2. Decomposition of Deterministic Slice-Stochastic Tensors

The first step is to construct a sufficiently rich class of games that have the strategy profile where every player randomizes uniformly as their unique equilibrium. This will be the class of cyclic games and almost cyclic games. In cyclic games, the payoff of every player only depends on the action of one other player and the dependencies form a cycle. Roughly, a player gets payoff 1 if she matches the action of the preceding player and 0 otherwise. The fact that every such game has uniform randomization as an equilibrium is not hard to see. Uniqueness is achieved by imposing a restriction on the notion of “matching” (see Definition 13(i)). Almost cyclic games differ only insofar that there is one exceptional player whose payoff not only depends on the action of the preceding player in the cycle but (for few action profiles) also on the actions of all other players. Making the concepts above precise requires two definitions.

Definition 11 (Fixed subsets). Let A be a set and $\pi \in \Sigma_A$. Then, $B \subseteq A$ is a *fixed subset* of π if $\pi(B) = B$. We say that π has *no non-trivial fixed subset* if its only fixed subsets are \emptyset and A .

For any two permutations π, π' , $\pi' \circ \pi$ has a non-trivial fixed subset if and only if $\pi \circ \pi'$ does. At least when A is finite, permutations without a non-trivial fixed subset always

exist. Any cyclic permutation is an example. Also, for is any permutation π , there is a permutation π' so that $\pi' \circ \pi$ has no non-trivial fixed subset.

Definition 12 (Permutation sets). Let $A_1, \dots, A_n \in \mathcal{F}(U)$ with $|A_1| = \dots = |A_n|$ and let $A = A_1 \times \dots \times A_n$. A set $A^* \subseteq A$ is a *permutation set* if for all $i \in N$ and $a'_i \in A_i$, there is exactly one $a \in A^*$ with $a_i = a'_i$.

Another way of saying that A^* is a permutation set is that the projection from A^* onto each A_i is bijective. If $A_1 = \dots = A_n = [m]$, yet another way is requiring that there are permutations π_1, \dots, π_{n-1} of $[m]$ so that $A^* = \{(k, \pi_1(k), \dots, \pi_{n-1} \circ \dots \circ \pi_1(k)) : k \in [m]\}$. This characterization motivates the terminology.

Definition 13 ((Almost) cyclic games). Let $A = [m]^n$ and G be a game on A . We say that G is *cyclic* if there are $\pi_1, \dots, \pi_n \in \Sigma_{[m]}$ and $\alpha \in \mathbb{R}_{++}^N$ so that the following holds.

- (i) $\pi_n \circ \dots \circ \pi_1$ has no non-trivial fixed subset,
- (ii) for all $j \in N$ and $a \in A$,

$$G_j(a) = \begin{cases} \alpha_j & \text{if } a_j = \pi_{j-1}(a_{j-1}), \\ 0 & \text{otherwise.} \end{cases}$$

For $n \geq 3$, we say that G is *almost cyclic* if there are $i \in N$, $\pi_1, \dots, \pi_n \in \Sigma_{[m]}$, $A^* \subseteq A$, and $\alpha \in \mathbb{R}_{++}^N$ so that the following holds.

- (i) $\pi_n \circ \dots \circ \pi_1$ has no non-trivial fixed subset,
- (ii) for all $j \neq i$ and $a \in A$,

$$G_j(a) = \begin{cases} \alpha_j & \text{if } a_j = \pi_{j-1}(a_{j-1}), \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad G_i(a) = \begin{cases} \alpha_i & \text{if } a_i = \pi_{i-1}(a_{i-1}) \text{ and } a \notin A^* \\ 0 & \text{otherwise,} \end{cases}$$

- (iii) A^* is a permutation set and for all $a \in A^*$, $a_i = \pi_{i-1}(a_{i-1})$ and there is $j \neq i, i+1$ with $a_j \neq \pi_{j-1}(a_{j-1})$.

Unless otherwise noted, we assume that $\alpha = (1, \dots, 1)$.

Example 1 ((Almost) cyclic games for $n = 2, 3$). Let $n = 2$ and consider the following game G where both players have three actions. The first (second) entry in each cell denotes the payoff of player 1 (player 2).

$$\begin{pmatrix} 1, 0 & 0, 1 & 0, 0 \\ 0, 0 & 1, 0 & 0, 1 \\ 0, 1 & 0, 0 & 1, 0 \end{pmatrix}$$

Then, G is a cyclic game with $\pi_1 = (123)$ and $\pi_2 = (1)(2)(3)$. Almost cyclic games for two players are degenerate in the sense that the payoff function of the exceptional player is 0.

Now let $m = n = 3$, $i = 1$, $\pi_1 = (123)$, and $\pi_2 = \pi_3 = (1)(2)(3)$. Clearly, $\pi_3 \circ \pi_2 \circ \pi_1$ has no non-trivial fixed subset. One can check that the permutation set $A^* = \{(1, 2, 1), (2, 3, 2), (3, 1, 3)\}$ satisfies (iii). The payoff function of player 1 is shown below (player 1 chooses the matrix, player two the row, and player 3 the column; the entries corresponding to A^* are marked in boldface).

$$\begin{pmatrix} 1 & 0 & 0 \\ \mathbf{0} & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} \quad \begin{pmatrix} 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & \mathbf{0} & 0 \end{pmatrix} \quad \begin{pmatrix} 0 & 0 & \mathbf{0} \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{pmatrix}$$

Lemma 7 (Equilibria of (almost) cyclic games). *Let G be a cyclic game or an almost cyclic game. Then, G has a unique Nash equilibrium where every player randomizes uniformly over all of her actions.*

Proof. We prove the statement for almost cyclic games. The proof for cyclic games is easier. More specifically, for cyclic games, Case 1 and Case 2 below can be combined into one that is proved in the same way as Case 1.

Let $n \geq 3$. Let $i \in N$, $\pi_1, \dots, \pi_n \in \Sigma_{[m]}$, A^* , and α be as in the definition of almost cyclic games. By Lemma 4, it is without loss of generality to assume that $\alpha = (1, \dots, 1)$. Let p be the strategy profile where p_j is the uniform distribution on A_j for all $j \in N$.

It is easy to see that p is an equilibrium since for all $a_i \in A_i$,

$$\sum_{a_{-i} \in A_{-i}} G_i(a_i, a_{-i}) = m^{n-2} - 1,$$

where the -1 comes from the fact that A^* is a permutation set. Moreover, for all $j \neq i$ and $a_j \in A_j$,

$$\sum_{a_{-j} \in A_{-j}} G_j(a_j, a_{-j}) = m^{n-2}.$$

Now we show that p is the unique equilibrium. Assume that p' is an equilibrium. For all $j \in N$, let $B_j = \arg \max_{a_j \in A_j} p'_j(a_j)$. Note that for all $j \neq i$, the payoff of j only depends on her own strategy and the strategy of $j-1$. Hence, $a_j \in A_j$ is a best response to p'_{-j} if $a_j \in \pi_{j-1}(B_{j-1})$. So $B_j \subseteq \text{supp}(p'_j) \subseteq \pi_{j-1}(B_{j-1})$. In particular,

$$|\text{supp}(p'_i)| \geq |\text{supp}(p'_{i+1})| \geq \dots \geq |\text{supp}(p'_{i-1})|. \quad (2)$$

We distinguish two cases.

Case 1. Suppose that $|\text{supp}(p'_j)| \geq 2$ for some $j \neq i$. By (2), $|\text{supp}(p'_{i+1})| \geq 2$. Since $n \geq 3$ and A^* is a permutation set, it follows that $G_i(a_i, p'_{-i}) > 0$ for all $a_i \in$

$\pi_{i-1}(\text{supp}(p'_{i-1}))$. Hence, $\text{supp}(p'_i) \subseteq \pi_{i-1}(\text{supp}(p'_{i-1}))$, and so $|\text{supp}(p'_i)| \leq |\text{supp}(p'_{i-1})|$. It follows that all inequalities in (2) hold with equality. But then for all $j \in N$, $\text{supp}(p'_j) = \pi_{j-1}(\text{supp}(p'_{j-1}))$, and so $\text{supp}(p'_1)$ is a fixed subset of $\pi_n \circ \dots \circ \pi_1$. By assumption, this is only possible if $A_1 = \text{supp}(p'_1)$, which in turn implies that $\text{supp}(p'_j) = A_j$ for all $j \in N$. Hence, $p' = p$.

Case 2. Suppose that $|\text{supp}(p'_j)| = 1$ for all $j \neq i$ and let $a'_j \in A_j$ so that $p'_j(a'_j) = 1$. Note that $a'_j = \pi_{j-1}(a'_{j-1})$ for all $j \neq i, i+1$. Moreover, we have that $G_i(\pi_{i-1}(a'_{i-1}), a'_{-i}) = 1$ unless $(\pi_{i-1}(a'_{i-1}), a'_{-i}) \in A^*$. But by the second part of (iii) in the definition of almost cyclic games, $(\pi_{i-1}(a'_{i-1}), a'_{-i}) \notin A^*$. Hence, $G_i(\pi_{i-1}(a'_{i-1}), a'_{-i}) = 1$. By (ii) and the fact that $|\text{supp}(p'_{i-1})| = 1$, $G_i(a_i, p'_{-i}) = 0$ unless $a_i = a'_i = \pi_{i-1}(a'_{i-1})$. Thus, a'_i is the unique best response of player i to p'_{-i} . Since p' is an equilibrium, it follows that $p'_i(a'_i) = 1$ and $|\text{supp}(p'_i)| = 1$. Finally, since p'_{i+1} is a best response to p'_i , we have that $a'_{i+1} = \pi_i(a'_i)$. So $a'_j = \pi_{j-1}(a'_{j-1})$ for all $j \in N$, which means that $\{a'_1\}$ is a fixed subset of $\pi_n \circ \dots \circ \pi_1$ and contradicts (i). \square

Definition 14 (Permutation tensors and permutation games). Let $A_1, \dots, A_n \in \mathcal{F}(U)$ with $|A_1| = \dots = |A_n|$ and let $A = A_1 \times \dots \times A_n$. A function $T: A \rightarrow \mathbb{R}$ is a *permutation tensor* if there is a permutation set $A^* \subseteq A$ so that $T(a) = 1$ for all $a \in A^*$ and $T(a) = 0$ for all $a \in A \setminus A^*$. A game G is a permutation game if there is $i \in N$ so that G is a player i payoff game and G_i is a permutation tensor.

For $n = 2$, a permutation tensor is a permutation matrix.

We show that every permutation game can be written as a convex combination of cyclic games and almost cyclic games up to an additive constant. Note that even though permutation games are player i payoff games, the games in the decomposition are not.

Lemma 8 (Decomposition of permutation games). *Let $A_1, \dots, A_n \in \mathcal{F}(U)$ so that $|A_1| = \dots = |A_n| = m$, and let $A = A_1 \times \dots \times A_n$. Let $i \in N$ and $A^* \subseteq A$ be a permutation set. Let G be the permutation game for i and A^* on A . Then, G can be written as a convex combination of cyclic games, almost cyclic games, and a constant $\beta \in \mathbb{R}^N$.*

Proof. Let G be a game as in the statement of the lemma. For simplicity, assume that $i = n$ and for all $j \in N$, $A_j = [m]$.

Let $\pi_1, \dots, \pi_n \in \Sigma_{[m]}$ so that $A^* = \{(k, \pi_1(k), \dots, \pi_{n-1} \circ \dots \circ \pi_1(k)) \in [m]^N : k \in [m]\}$ and $\pi_n \circ \dots \circ \pi_1$ has no non-trivial fixed subset. This is possible since we may first choose π_1, \dots, π_{n-1} so that the first condition holds (using that A^* is a permutation set) and then choose π_n so that $\pi_n \circ \dots \circ \pi_1$ has no non-trivial fixed subset.

Let $\hat{A} = \{a \in A : a_n = \pi_{n-1}(a_{n-1})\}$. Note that $|\hat{A}| = m^{n-1}$ and $A^* \subseteq \hat{A}$. Let $B^1, \dots, B^M \subseteq A$ be a partition of $\hat{A} \setminus A^*$ into permutation sets, where $M = m^{n-2} - 1$. (Note that $M = 0$ if $n = 2$.) For example, one may take

$$B^{s-\{1,n\}} = \{(a_1, a_2 + s_2, \dots, a_{n-1} + s_{n-1}, \pi_{n-1}(a_{n-1} + s_{n-1})) \in [m]^N : a \in A^*\},$$

where $s_{-\{1,n\}} \in [m]^{N \setminus \{1,n\}} \setminus \{0\}$. Since $s_{-\{1,n\}} \neq 0$, the last condition in (iii) holds for B^l . For all $l \in [M]$, let G^l be the almost cyclic game for $i = n$, π_1, \dots, π_n , B^l , and $\alpha = (1, \dots, 1)$. In particular, $G_n^l(a) = 0$ for all $a \in B^l$ and $G_n^l(a) = 1$ for all $a \in \hat{A} \setminus B^l$.

Now let

$$\hat{G} = \sum_{l \in [M]} G^l.$$

Then, the following hold.

- For all $a \in A^*$, $G_n(a) = M$; for all $a \in \hat{A} \setminus A^*$, $G_n(a) = M - 1$.
- For all $j \neq n$ and $a \in A$ with $a_j = \pi_{j-1}(a_{j-1})$, $G_j(a) = M$.
- For all $j \in N$ and $a \in A$ with $a_j \neq \pi_{j-1}(a_{j-1})$, $G_j(a) = 0$.

Recall that for all $j \in N$ and $(a_{j-1}, a_j) \in [m]^2$, there is a cyclic game G' for some $\pi'_1, \dots, \pi'_n \in \Sigma_{[m]}$ with $a_j = \pi'_{j-1}(a_{j-1})$ and arbitrary $\alpha' \in \mathbb{R}_{++}^N$ (see the remarks after Definition 11). For all $j \in N$ and $(a_{j-1}, a_j) \in [m]^2$ with $a_j \neq \pi_{j-1}(a_{j-1})$, let G^{j,a_{j-1},a_j} be a cyclic game for some $\pi'_1, \dots, \pi'_n \in \Pi_{[m]}$ with $a_j = \pi'_{j-1}(a_{j-1})$, $\alpha'_{j'} = 1$ for all $j' \neq j$, and $\alpha'_j = M + 1$ if $j \neq n$ and $\alpha'_j = M$ if $j = n$. Then, let

$$\tilde{G} = \hat{G} + \sum G^{j,a_{j-1},a_j} + \sum G', \quad (3)$$

where the first sum ranges over all $j \in N$ and $(a_{j-1}, a_j) \in [m]^2$ with $a_j \neq \pi_{j-1}(a_{j-1})$, and the second sum ranges over all cyclic games with $\alpha = (1, \dots, 1)$ whose tuple of permutations does not already appear in one of the games in the first sum. Since for all $j \in N$ and $(a_{j-1}, a_j) \in [m]^2$, the number of cyclic games with $a_j = \pi'_{j-1}(a_{j-1})$ is the same (assuming α is fixed), the following hold.

- For all $a^* \in A^*$ and $a \in A \setminus A^*$, $\tilde{G}_n(a^*) = \tilde{G}_n(a) + 1$.
- For all $j \neq n$ and $a, a' \in A$, $G_j(a) = G_j(a')$.

Hence, $G = \tilde{G} + \beta$ for some $\beta \in \mathbb{R}^N$. More explicitly, if L is the total number of games in the first and second sum in (3), then

- For all $a \in A^*$, $\tilde{G}_n(a) = M + \frac{L}{m}$; for all $a \in A \setminus \bar{A}^*$, $G_n(a) = M - 1 + \frac{L}{m}$.
- For all $j \neq i$ and $a \in A$, $\tilde{G}_j(a) = M + \frac{L}{m}$.

(The denominator m comes from the fact that for a permutation game G' , the fraction of actions for which $G'_j(a) = \alpha_j$ is $\frac{1}{m}$.) This proves G can be written as a sum of games of the claimed types. Multiplying each game in that sum by the same appropriately chosen positive scalar gives the representation of G as a convex combination. \square

The last lemma in this section roughly shows that every game with deterministic slice-stochastic tensors as payoff functions can be written as a convex combination of permutation games. This conclusion is not literally true. More precisely, we show that

every game with deterministic slice-stochastic tensors as payoff functions is the blow-down of a convex combination of permutation games (with the same number of clones of every action for every player). By Lemma 7 and Lemma 8, each of the games in this sum can in turn be written as a convex combination of games for which uniform randomization is the unique equilibrium. Consistency and consequentialism thus imply that uniform randomization has to be returned in the original game, possibly alongside other equilibria.

Lemma 9 (Decomposition of slice-stochastic games). *Let $A_1, \dots, A_n \in \mathcal{F}(U)$ with $|A_1| = \dots = |A_n| = m$, and let $A = A_1 \times \dots \times A_n$. Let G be a game on A so that G_i is a deterministic slice-stochastic tensor for all $i \in N$. Let $p \in \square A$, where p_i is the uniform distribution on A_i . Then, up to cloning actions, multiplying by positive scalars, and adding constants, G can be written as a convex combination of games for which p is the unique equilibrium. More precisely, there are $\alpha \in \mathbb{R}_{++}^N$, $\beta \in \mathbb{R}^N$, and a game \bar{G} so that*

- (i) G is a blow-down of $\alpha \bar{G} + \beta$ with surjection ϕ ,
- (ii) $\phi_*(\hat{p}) = p$, where \hat{p}_j is the uniform distribution on the actions of player j in \bar{G} , and
- (iii) \bar{G} is a convex combination of games for which \hat{p} is the unique equilibrium.

Proof. First we observe that it suffices to consider the case that G is a player i payoff game (meaning that $G_j = 0$ for all $j \neq i$). The idea now is to “blow up” G by introducing m^{n-2} actions for each action of every player in G . Then, one can define a permutation game \hat{G} on the larger action sets so that for every action profile a in G for which player i has payoff 1, there is exactly one blow up \hat{a} of a for which i has payoff 1 in \hat{G} . Since \hat{G} is a permutation game, we know from Lemma 8 that, up to adding constants, it can be written as a convex combination of games for which the strategy profile where every player plays the uniform distribution is the unique equilibrium. Now permuting all actions in \hat{G} that come from blowing up the same action in G , summing up over all the resulting games, and multiplying by a positive scalar gives a game \bar{G} that is a blow-up of G . This achieves the desired decomposition. (Remark 10 explains why the blowing up is necessary.)

As noted above, we may assume that there is $i \in N$ so that G is a player i payoff game. Let $A^* = \{a \in A : G_i(a) = 1\}$ be the actions for which player i has payoff 1 in G . For all $j \in N$ and $a_j \in [m]$, let $B_j^{a_j} = \{a_{-j} \in A_{-j} : (a_j, a_{-j}) \in A^*\}$ be the set of opponents action profiles for which i gets payoff 1 when j plays a_j . Since G_i is slice-stochastic, $|B_j^{a_j}| = m^{n-2}$. (For $j = i$, this is (ii) in the definition of slice-stochastic games; for $j \neq i$, for all a_j and $a_{-\{i,j\}}$, there is exactly one a_i so that $G_i(a_i, a_j, a_{-\{i,j\}}) = 1$ by (i).) Let $\hat{A}_j = \bigcup_{a_j \in A_j} \{a_j\} \times B_j^{a_j}$ and $\hat{A} = \hat{A}_1 \times \dots \times \hat{A}_n$. Note that \hat{A}_j has size m^{n-1} . Now let and

$$\hat{A}^* = \{((a_1, a_{-1}), \dots, (a_n, a_{-n})) \in \hat{A} : a \in A^*\}.$$

It is easy to see that \hat{A}^* is a permutation set in \hat{A} . Let \hat{G} be the permutation game for i on \hat{A} for the permutation set \hat{A}^* . By Lemma 8, we have that \hat{G} can be written as a sum of cyclic games, almost cyclic games, and a constant. Moreover, by Lemma 7, we know that \hat{p} is the unique equilibrium of cyclic and almost cyclic games, where \hat{p}_j is the uniform distribution on \hat{A}_j .

For all $j \in N$, let $\phi_j: \hat{A}_j \rightarrow A_j$ be projection onto the first coordinate, and let $\phi = (\phi_1, \dots, \phi_n)$. That is, $\phi_j((a_j, a_{-j})) = a_j$. Let $\Sigma_j \subseteq \Sigma_{\hat{A}_j}$ be the set of all permutations π_j so that $\phi_j(\hat{a}_j) = \phi_j(\pi_j(\hat{a}_j))$ for all $\hat{a}_j \in \hat{A}_j$ (that is, all permutations that keep the first coordinate fixed). Let $\Sigma = \Sigma_1 \times \dots \times \Sigma_n$. So we have that for all $\hat{a} \in \hat{A}$ and $\pi \in \Sigma$,

$$G_i(\phi(\hat{a})) = G_i((\phi \circ \pi)(\hat{a})).$$

Consider the game

$$\bar{G} = \frac{M}{|\Sigma|} \sum_{\pi \in \Sigma} \hat{G} \circ \pi,$$

where $M = m^{(n-1)n}$. (The factor M is necessary since for all $a \in A$, there are M profiles $\hat{a} \in \hat{A}$ with $\phi(\hat{a}) = a$.)

Since each $\hat{a} \in \hat{A}^*$ is completely determined by its first coordinates, for all $a \in A^*$, there is exactly one $\hat{a} \in \hat{A}^*$ with $\phi(\hat{a}) = a$. Thus, the following hold.

- For all $\hat{a} \in \hat{A}$ with $\phi(\hat{a}) \in A^*$, $\bar{G}_i(\hat{a}) = 1$.
- For all $\hat{a} \in \hat{A}$ with $\phi(\hat{a}) \in A \setminus A^*$, $\bar{G}_i(\hat{a}) = 0$.
- For all $j \neq i$ and $\hat{a} \in \hat{A}$, $\bar{G}_j(\hat{a}) = 0$.

So for all $j \in N$ and $a_j \in A_j$, all actions in $\{a_j\} \times B_j^{a_j}$ are clones of each other in \bar{G} . This gives that \bar{G} is a blow-up of G . Moreover, $\phi_*(\hat{p}) = p$ since for all $j \in N$ and $a_j \in A_j$, there is the same number of clones (namely m^{n-2}) in \hat{A}_j . This gives the desired decomposition of G . \square

Remark 10 (Decomposition into permutation games). The “blowing up” of G to \bar{G} in the proof of Lemma 9 by introducing m^{n-2} clones of every action is necessary since not every deterministic slice-stochastic tensor can be written as a sum of permutation tensors. For example, let $n = 3$, $m = 2$, $i = 1$, and $G_1(a) = 1$ for $a \in \{(1, 1, 1), (1, 2, 2), (2, 1, 2), (2, 2, 1)\} \subseteq \{1, 2\}^3$ and $G_1(a) = 0$ otherwise. Then, there is no permutation set $B \subseteq \{1, 2\}^3$ so that $G_1(a) = 1$ for all $a \in B$.

So far, we have shown the following. To prove that $p \in f(G)$ whenever $p \in \text{NASH}(G)$ and p has full support, it suffices to show this for the case when each G_i is slice-stochastic by Lemma 6. By the generalization of the Birkhoff-von Neumann theorem, Lemma 3, the fact that in every such game uniform randomization is an equilibrium, and consistency, we can further restrict to G_i ’s that are deterministic slice-stochastic. Now Lemma 9 shows that all those games can in essence be written as convex combinations of games for which uniform randomization is the unique equilibrium and, thus, has to be returned by f . Consistency then gives the desired conclusion. The last step is to extend the statement beyond full support equilibria.

A.3. Reduction to Full Support Equilibria

We show that it suffices to prove that all full support equilibria have to be returned, which we have done in the previous section. There are three steps to the argument. First, we reduce to equilibria with support equal to the set of all rationalizable actions, then to quasi-strict equilibria, and lastly to arbitrary equilibria.¹⁵ The strategy is always to write a game with some type of equilibrium as a convex combination of games where the same equilibrium is of the type in the preceding step. For example, Lemma 10 shows that any nice total solution concept has to return all equilibria whose support consists of all rationalizable actions.

Lemma 10 (Equilibria with full support on rationalizable actions). *Let f be a nice total solution concept. Let G be a game on A . Let $\bar{A}_i \subseteq A_i$ be the sets of rationalizable actions of $i \in N$ in G and $\bar{A} = \bar{A}_1 \times \cdots \times \bar{A}_n$. Then, if $p \in \text{NASH}(G)$ so that for all $i \in N$, $\text{supp}(p_i) = \bar{A}_i$, then $p \in f(G)$.*

Proof. Consider a decreasing sequence of action profiles obtained by successively removing dominated actions until no more deletions are possible. That is, let $(A_1^0, \dots, A_n^0), \dots, (A_1^K, \dots, A_n^K) \in 2^{A_1} \times \cdots \times 2^{A_n}$ so that for all $k \in [K]$, there is $i \in N$ for which the following holds.

- $A_i^k \subsetneq A_i^{k-1}$ and for all $j \neq i$, $A_j^k = A_j^{k-1}$.
- For all $a_i \in A_i^{k-1} \setminus A_i^k$, there is an action $\psi(a_i) \in A_i^k$ that dominates a_i when restricting G to $A_1^{k-1} \times \cdots \times A_n^{k-1}$.
- For all $j \in N$, $A_j^0 = A_j$ and $A_j^K = \bar{A}_j$.

For all $i \in N$ and $a_i \in A_i \setminus \bar{A}_i$, let $\bar{\psi}(a_i) = \psi^s(a_i) \in \bar{A}_i$, where $s \in \mathbb{N}$ is the unique power so that $\psi^s(a_i) \in \bar{A}_i$ (here we mean ψ applied s times).

Denote by \bar{G} the game G restricted to action profiles in \bar{A} . We make several reductions. By consequentialism, we may assume that $|\bar{A}_1| = \cdots = |\bar{A}_n|$. By Lemma 5 and Lemma 6, we may assume that for all $i \in N$, \bar{G}_i is slice-stochastic and p_i is the uniform distribution on \bar{A}_i . By consequentialism, Lemma 9, and Lemma 4, we may further assume that \bar{G} can be written as a convex combination of games for which p is the unique equilibrium. That is, there are games $\bar{G}^1, \dots, \bar{G}^M$ on \bar{A} so that for all $m \in [M]$, p is the unique equilibrium of \bar{G}^m , and

$$\bar{G} = \frac{1}{M} \sum_{m \in [M]} \bar{G}^m. \quad (4)$$

For all $m \in [M]$, we define a game G^m on A so that for all $i \in N$ and $a \in A$,

$$G_i^m(a) = \begin{cases} \bar{G}_i^m(a) & \text{if } a \in \bar{A}, \\ G_i(a) + \bar{G}_i^m(\bar{\psi}(a_i), a_{-i}) - G_i(\bar{\psi}(a_i), a_{-i}) & \text{if } a \in (A_i \setminus \bar{A}_i) \times \bar{A}_{-i}, \text{ and} \\ G_i(a) & \text{if } a_{-i} \in A_{-i} \setminus \bar{A}_{-i}. \end{cases}$$

¹⁵We say that an action (profile) is rationalizable if it survives iterated elimination of strictly dominated actions.

Since $\bar{\psi}(a_i) = \bar{\psi}(a'_i)$ for $a_i, a'_i \in A_i \setminus \bar{A}_i$ with $\psi(a_i) = a'_i$, we have that \bar{A} is the set of rationalizable action profiles of G^m . It follows that p is the unique equilibrium in G^m and so $p \in f(G^m)$. Observe that for $a \in (A_i \setminus \bar{A}_i) \times \bar{A}_{-i}$, we have by (4) that

$$\sum_{m \in [M]} \bar{G}_i^m(\bar{\psi}(a_i), a_{-i}) - G_i(\bar{\psi}(a_i), a_{-i}) = 0.$$

Hence, $G = \frac{1}{m} \sum_{m \in [M]} G^m$. Consistency then implies that $p \in f(G)$. \square

A profile $p \in \square A$ is a quasi-strict equilibrium of G if $p \in \text{NASH}(G)$ and

$$G_i(a_i, p_{-i}) > G_i(a'_i, p_{-i}) \text{ for all } a_i \in \text{supp}(p_i), a'_i \in A_i \setminus \text{supp}(p_i), \text{ and } i \in N.$$

We show that nice total solution concepts have to return quasi-strict equilibria.

Lemma 11 (Quasi-strict equilibria). *Let f be a nice total solution concept. Let G be a game on A . Then, if $p \in \text{NASH}(G)$ is quasi-strict, then $p \in f(G)$.*

Proof. For all $i \in N$, let $\hat{A}_i = \text{supp}(p_i)$ and $\hat{A} = \hat{A}_1 \times \cdots \times \hat{A}_n$. By consequentialism, we may assume that $|\hat{A}_i| = 2|A_i \setminus \hat{A}_i|$ and the number of clones of each action in \hat{A}_i is even. Moreover, by Lemma 6 and the remarks thereafter, we may assume that p_i is the uniform distribution on \hat{A}_i . Write $\hat{A}_i = \{a_i^1, \dots, a_i^K, b_i^1, \dots, b_i^K\}$ so that a_i^k and b_i^k are clones for all $k \in [K]$ and $A_i \setminus \hat{A}_i = \{c_i^1, \dots, c_i^K\}$. The idea is to write G as a convex combination of two games G^1, G^2 for which all actions in $A_i \setminus \hat{A}_i$ are dominated and p is an equilibrium of G^1, G^2 . Lemma 10 and consistency of f will then give that $p \in f(G)$.

Let $i \in N$. Since p is quasi-strict and p_{-i} is uniform on \hat{A}_{-i} , we have for all $a_i \in \hat{A}_i$ and $a'_i \in A_i \setminus \hat{A}_i$,

$$\sum_{a_{-i} \in \hat{A}_{-i}} G_i(a_i, a_{-i}) > \sum_{a_{-i} \in \hat{A}_{-i}} G_i(a'_i, a_{-i}). \quad (5)$$

For all $k \in [K]$, let $v_i^k \in \mathbb{R}^{A_{-i}}$ so that

$$\sum_{a_{-i} \in \hat{A}_{-i}} v_i^k(a_{-i}) = 0 \quad (6)$$

and for all $a_{-i} \in A_{-i}$,

$$G_i(a_i^k, a_{-i}) + v_i^k(a_{-i}) = G_i(b_i^k, a_{-i}) + v_i^k(a_{-i}) > G_i(c_i^k, a_{-i}). \quad (7)$$

By (5), such v_i^k exist. (Note that the sum in (6) is taken over \hat{A}_{-i} and (7) is required to hold for all action profiles in A_{-i} .)

Now define games G^1, G^2 on A as follows. For all $i \in N$ and $a \in A$,

$$G_i^1(a) = \begin{cases} G_i(a_i^k, a_{-i}) + v_i^k(a_{-i}) & \text{if } a_i = a_i^k \text{ for some } k \in [K], \\ G_i(b_i^k, a_{-i}) - v_i^k(a_{-i}) & \text{if } a_i = b_i^k \text{ for some } k \in [K], \text{ and} \\ G_i(c_i^k, a_{-i}) & \text{if } a_i = c_i^k \text{ for some } k \in [K]. \end{cases}$$

Define G^2 similarly with the roles of a_i^k and b_i^k exchanged. By (6), p is an equilibrium of G^1, G^2 , and by (7), all actions in $A_i \setminus \hat{A}_i$ are dominated. More specifically, in G^1 , each c_i^k is dominated by a_i^k and in G^2 , each c_i^k is dominated by b_i^k . Hence, the set of rationalizable action profiles in G^1, G^2 is \hat{A} . By Lemma 10, $p \in f(G^1) \cap f(G^2)$. Since $G = 1/2G^1 + 1/2G^2$, consistency implies that $p \in f(G)$. \square

Lemma 11 together with consequentialism allows us to push slightly beyond quasi-strict equilibria. If p is an equilibrium of a game G so that for every player i , every action of i that is a best response against p_{-i} is either in the support of p_i or a clone of such an action, then we get from consequentialism that $p \in f(G)$. In that case, we say that p is essentially quasi-strict.

Definition 15 (Essentially quasi-strict equilibrium). Let G be a game on A . An equilibrium p of G is *essentially quasi-strict* if there is a blow-down G' of G with surjection ϕ so that $\phi_*(p)$ is a quasi-strict equilibrium of G' .

Similarly, one could define essentially unique and essentially full support equilibria, but we will not need these notions. Note that if a solution concept satisfies consequentialism and returns all quasi-strict equilibria, then it also has to return all essentially quasi-strict equilibria. We use this fact in the proof of the last step: if a solution concept satisfies consequentialism and consistency and returns all quasi-strict equilibria, then it in fact has to return all equilibria.

Lemma 12 (Reduction from quasi-strict to all equilibria). *Let f be a solution concept that satisfies consequentialism and consistency so that $p \in f(G)$ whenever p is a quasi-strict equilibrium of G . Then, $NASH \subseteq f$.*

Proof. Let G be a game on A and $p \in NASH(G)$. For $i \in N$, let $\hat{A}_i = \text{supp}(p_i)$ and for any game G' on A with $p \in NASH(G')$, let

$$\bar{A}_i(G') = \{a_i \in A_i \setminus \hat{A}_i : G'_i(a_i, p_{-i}) = G'_i(p_i, p_{-i}) \text{ and } a_i \text{ is not a clone of an action in } \hat{A}_i\}.$$

That is, $\bar{A}_i(G')$ is the set of actions that are best responses against p_{-i} and not in the support of p_i or clones of actions in the support of p_i . We write $\bar{A}_i = \bar{A}_i(G)$. Note that p is an essentially quasi-strict equilibrium of G' if $\bar{A}_i(G') = \emptyset$ for all $i \in N$.

We prove that $p \in f(G)$ by induction on the number of players for which $\bar{A}_i \neq \emptyset$. If $\bar{A}_i = \emptyset$ for all $i \in N$, the statement follows from the assumption that f satisfies consequentialism and returns quasi-strict equilibria. Otherwise, let $i \in N$ with $\bar{A}_i \neq \emptyset$. We write G as a convex combination of two games G^1, G^2 so that $p \in NASH(G^1) \cap NASH(G^2)$, and

$$\{j \in N : \bar{A}_i(G^l) \neq \emptyset\} \subseteq \{j \in N : \bar{A}_i \neq \emptyset\} \setminus \{i\} \quad (8)$$

for $l = 1, 2$.

By Lemma 6 and the remarks thereafter, we may assume that p_i is the uniform distribution on \hat{A}_i . Moreover, by consequentialism, we may assume that $\hat{A}_i =$

$\{a_i^1, \dots, a_i^K, b_i^1, \dots, b_i^K\}$, $\bar{A}_i = \{c_i^1, \dots, c_i^K\}$, and a_i^k, b_i^k are clones in G for all $k \in [K]$. Let G^1, G^2 be games on A so that for all $j \in N$ and $a \in A$,

$$G_j^1(a) = \begin{cases} G_j(c_i^k, a_{-i}) & \text{if } a_i = a_i^k \text{ or } a_i = c_i^k \text{ for some } k \in [K], \\ G_j(a_i^k, a_{-i}) + G_j(b_i^k, a_{-i}) - G_j(c_i^k, a_{-i}) & \text{if } a_i = b_i^k \text{ for some } k \in [K], \text{ and} \\ G_j(a_i, a_{-i}) & \text{if } a_i \in A_i \setminus (\hat{A}_i \cup \bar{A}_i), \end{cases}$$

and

$$G_j^2(a) = \begin{cases} G_j(c_i^k, a_{-i}) & \text{if } a_i = b_i^k \text{ or } a_i = c_i^k \text{ for some } k \in [K], \\ G_j(a_i^k, a_{-i}) + G_j(b_i^k, a_{-i}) - G_j(c_i^k, a_{-i}) & \text{if } a_i = a_i^k \text{ for some } k \in [K], \text{ and} \\ G_j(a_i, a_{-i}) & \text{if } a_i \in A_i \setminus (\hat{A}_i \cup \bar{A}_i). \end{cases}$$

Then, $G = 1/2G^1 + 1/2G^2$ and for all $k \in [K]$, a_i^k, c_i^k are clones in G^1 and b_i^k, c_i^k are clones in G^2 . Moreover, $p \in \text{NASH}(G^1) \cap \text{NASH}(G^2)$ by the definition of \bar{A}_i . To see that p_i is a best response to p_{-i} , recall that for all $k \in [K]$,

$$G_i(a_i^k, p_{-i}) = G_i(b_i^k, p_{-i}) = G_i(p_i, p_{-i}) = G_i(c_i^k, p_{-i}),$$

since c_i^k is assumed to be in \bar{A}_i . So for all $a_i \in A_i$, $G_i^1(a_i, p_{-i}) = G_i^2(a_i, p_{-i}) = G_i(a_i, p_{-i})$. Also, for $j \neq i$ and $a_j \in A_j$, we have

$$\begin{aligned} G_j(a_j, p_{-j}) &= \frac{1}{2K} \sum_{k \in [K]} G_j(a_j, a_i^k, p_{-\{i,j\}}) + G_j(a_j, b_i^k, p_{-\{i,j\}}) \\ &= \frac{1}{2K} \sum_{k \in [K]} G_j^1(a_j, a_i^k, p_{-\{i,j\}}) + G_j^1(a_j, b_i^k, p_{-\{i,j\}}) \\ &= G_j^1(a_j, p_{-j}), \end{aligned}$$

where we use that p_i is uniform on \hat{A}_i for the first equality, and the definition of G^1 for the second equality. In particular, p_j is a best response to p_{-j} in G^1 . The same holds for G^2 with a similar argument. Thus, (8) holds. Now, by induction, $p \in f(G^1) \cap f(G^2)$, and so by consistency, we get $p \in f(G)$. \square

The fact that any nice total solution concept is a coarsening of *NASH* now follows from Lemma 11 and Lemma 12. This finishes the proof of Theorem 1.

B. Omitted Proofs From Section 5

Throughout this section, we consider equivariant solution concepts that are δ -nice for some small enough δ . Let G be a game with action profiles A and fix a strategy profile $p \in f(G)$. Lemma 13 below then states that a new game \hat{G} can be obtained by adding an extra action \hat{a}_i to the action set of every player $i \in N$ such that the following holds.

1. There is a profile $q \in \square A$ close to p so that for all $i \in N$, if i plays \hat{a}_i in \hat{G} , all players get the same payoff as if i played q_i in G .

2. There is a profile $\hat{p} \in f(\hat{G})$ so that for all $i \in N$, \hat{p}_i has all but a small fraction of probability on \hat{a}_i .

Lemma 13. *Let $\delta > 0$ and f an equivariant solution concept that satisfies δ -consequentialism and δ -consistency. Let G be a game on A and $p \in f(G)$. Then, there is a game \hat{G} with action set $\hat{A}_i = A_i \cup \{\hat{a}_i\}$ for all $i \in N$ so that the following holds.*

1. *There is $q \in B_{3\delta}(p)$ such that $\hat{G}(\hat{a}_I, a_{-I}) = G(q_I, a_{-I})$ for all $I \subseteq N$ and $a_{-I} \in A_{-I}$.*
2. *There is $\hat{p} \in f(\hat{G})$ such that $\hat{p}_i(\hat{a}_i) \geq 1 - 3\delta$ for all $i \in N$.*

Proof. For all $i \in N$, let $k_i = |A_i| \lceil \frac{1}{\delta} \rceil$.¹⁶ For each $a_i \in A_i$, let $A_i^{a_i} \in \mathcal{F}(U)$, all disjoint and disjoint from A_i , with $|A_i^{a_i}| = k_i$. Let G' be a game with action set $A'_i = A_i \cup \bigcup_{a_i \in A_i} A_i^{a_i}$ for all $i \in N$, and G a blow-down of G' with surjections $\phi_i: A'_i \rightarrow A_i$ so that $\phi_i^{-1}(a_i) = \{a_i\} \cup A_i^{a_i}$ for all $a_i \in A_i$ and $i \in N$. That is, G' results from G by adding k_i clones of a_i for each action a_i .

Since f satisfies δ -consequentialism, $\phi_*^{-1}(p) \subseteq B_\delta(f(G'))$. We construct $p' \in f(G')$ so that p'_i assigns probability at least $1 - \delta$ uniformly to some subset \tilde{A}_i of $\bigcup_{a_i \in A_i} A_i^{a_i}$. Let $\tilde{p} \in B_\delta(p) \cap \phi_*(f(G'))$, which exists since f satisfies δ -consequentialism. Let $l_i = \lfloor k_i \tilde{p}_i \rfloor \in \{0, \dots, k_i\}^{A_i}$ and $r_i = k_i \tilde{p}_i - l_i \in [0, 1)^{A_i}$. For each $a_i \in A_i$, choose a subset $\tilde{A}_i^{a_i}$ of $A_i^{a_i}$ with $|\tilde{A}_i^{a_i}| = l_i(a_i)$ and let $\tilde{A}_i = \bigcup_{a_i \in A_i} \tilde{A}_i^{a_i}$. Let $p' \in \square A'$ such that for all $i \in N$ and $a_i \in A_i$,

$$p'_i(a'_i) = \begin{cases} \frac{1}{k_i} & \text{for } a'_i \in \tilde{A}_i^{a_i}, \text{ and} \\ \frac{r_i(a_i)}{k_i} \frac{1}{|\{a_i\} \cup A_i^{a_i} - \tilde{A}_i^{a_i}|} & \text{for } a'_i \in \{a_i\} \cup A_i^{a_i} - \tilde{A}_i^{a_i}. \end{cases}$$

Observe that for all $a_i \in A_i$,

$$p'(\{a_i\} \cup A_i^{a_i}) = \frac{|\tilde{A}_i^{a_i}|}{k_i} + \frac{r_i(a_i)}{k_i} = \frac{l_i(a_i)}{k_i} + \frac{r_i(a_i)}{k_i} = \tilde{p}_i(a_i).$$

Hence, p' is well-defined and $\phi_* p' = \tilde{p}$. By the choice of \tilde{p} , $p' \in f(G')$.

Recall that $r_i \in [0, 1)^{A_i}$, and so

$$\frac{|r_i|}{k_i} < \frac{|A_i|}{k_i} \leq \delta.$$

Let $q \in \square A$ with $q_i = \frac{\tilde{p}_i - r_i}{|\tilde{p}_i - r_i|}$. Since $\frac{|r_i|}{k_i} < \delta$, it follows that $q \in B_{2\delta}(\tilde{p}) \subseteq B_{3\delta}(p)$.

For all $i \in N$, let $\tilde{\Sigma}_{A'_i} \subseteq \Sigma_{A'_i}$ be the permutations that map the set $\{a_i\} \cup A_i^{a_i} - \tilde{A}_i$ to itself for all $a_i \in A_i$, and let $\tilde{\Sigma}_{A'} = \tilde{\Sigma}_{A'_1} \times \dots \times \tilde{\Sigma}_{A'_n}$. Then, $p' = p' \circ \pi$ for all $\pi \in \tilde{\Sigma}_{A'}$ since p'_i is a uniform distribution on \tilde{A}_i and the uniform distribution on $\{a_i\} \cup A_i^{a_i} - \tilde{A}_i^{a_i}$

¹⁶For $x \in \mathbb{R}$, $\lceil x \rceil$ is the smallest integer that is at least as large as x . Similarly, $\lfloor x \rfloor$ is the largest integer that is at most as large as x .

for every $a_i \in A_i$ and $i \in N$. Thus, equivariance of f implies that $p' = p' \circ \pi \in f(G' \circ \pi)$. Let

$$\bar{G} = \frac{1}{|\tilde{\Sigma}_{A'}|} \sum_{\pi \in \Pi} G' \circ \pi$$

It follows from δ -consistency that $p' \in B_\delta(f(\bar{G}))$. Let $\bar{p} \in f(\bar{G}) \cap B_\delta(p')$. By construction, for all $i \in N$, all actions in \tilde{A}_i are clones of each other in \bar{G} , and for each $a_i \in A_i$, all actions in $\{a_i\} \cup A_i^{a_i} - \tilde{A}_i^{a_i}$ are clones of each other. Let \hat{G} be the blow-down of \bar{G} with action set $\hat{A}_i = A_i \cup \{\hat{a}_i\}$ for all $i \in N$ and surjections $\hat{\phi}_i: A'_i \rightarrow \hat{A}_i$ so that $\hat{\phi}_i^{-1}(\hat{a}_i) = \tilde{A}_i$ and $\hat{\phi}_i^{-1}(a_i) = \{a_i\} \cup A_i^{a_i} - \tilde{A}_i^{a_i}$ for all $a_i \in A_i$. Since f satisfies δ -consequentialism, there is $\hat{p} \in f(\hat{G})$ so that $\hat{\phi}_*(\bar{p}) \in B_\delta(\hat{p})$. Hence, for all $i \in N$,

$$\sum_{a_i \in A_i} \hat{p}_i(a_i) \leq \sum_{a_i \in A_i} \bar{p}_i(a_i \cup A_i^{a_i} - \tilde{A}_i^{a_i}) + \delta \leq \sum_{a_i \in A_i} p'_i(a_i \cup A_i^{a_i} - \tilde{A}_i^{a_i}) + 2\delta \leq 3\delta,$$

where the last inequality uses the definition of p'_i and $\frac{|r_i|}{k_i} \leq \delta$. Equivalently, $\hat{p}_i(\hat{a}_i) \geq 1 - 3\delta$ for all $i \in N$. Moreover, by construction of \hat{G} , $\hat{G}(\hat{a}_I, a_{-I}) = G(q_I, a_{-I})$ for all $I \subseteq N$, and $a_{-I} \in A_{-I}$. \square

Now consider a game \hat{G} with action profiles \hat{A} and let $\hat{a} \in \hat{A}$, $a_i \in \hat{A}_i$, and $i \in N$ such that action a_i yields strictly more payoff against \hat{a}_{-i} than action \hat{a}_i . Lemma 14 below shows that if there is $\hat{p} \in f(\hat{G})$ so that \hat{p}_j assigns probability close to 1 to \hat{a}_j for each $j \in N$, then there is a game G and $p \in f(G)$ such that p_i assigns probability close to 1 to a dominated action.

Lemma 14. *Let $\varepsilon, \delta > 0$ with $4(\lceil \frac{1}{\varepsilon} \rceil + 1)\delta \leq (1 - 2\delta)$ and $2\delta \leq \varepsilon$. Let f be an equivariant solution concept that satisfies δ -consequentialism and δ -consistency. Let \hat{G} be a normalized game with action sets $\hat{A}_j = A_j \cup \{\hat{a}_j\}$ for all $j \in N$ and let $i \in N$ and $a_i \in A_i$ so that $\hat{G}_i(a_i, \hat{a}_{-i}) > \hat{G}_i(\hat{a}_i, \hat{a}_{-i}) + \varepsilon$. Then, if $\hat{p} \in f(\hat{G})$ with $\hat{p}_j(\hat{a}_j) \geq 1 - \delta$ for all $j \in N$, there is a game \bar{G} and $\bar{p} \in f(\bar{G})$ so that \bar{p}_i assigns probability at least $1 - 3\delta$ to an action that is δ -dominated in \bar{G} .*

Proof. Let $M = 2(\lceil \frac{1}{\varepsilon} \rceil + 1)$. For each $j \in N$, let $A_j = \{a_j^1, \dots, a_j^{|A_j|}\}$. Let G' be a game with action set $A'_i = \hat{A}_i$ for i and $A'_j = \hat{A}_j \cup \{a_j^{k,l}: k \in [|A_j|], l \in [M]\}$ so that each $a_j^{k,l}$ is a clone of \hat{a}_j for all $j \in N \setminus \{i\}$. That is, let $\phi_i: A'_i \rightarrow \hat{A}_i$ be the identity, and for all $j \in N \setminus \{i\}$, let $\phi_j: A'_j \rightarrow \hat{A}_j$ be the identity on A_j and $\phi_j^{-1}(\hat{a}_j) = \{\hat{a}_j\} \cup \{a_j^{k,l}: k \in [|A_j|], l \in [M]\}$. Then, \hat{G} is a blow-down of G' with surjection $\phi = (\phi_1, \dots, \phi_n)$. Since f satisfies δ -consequentialism, there is $\hat{p}' \in B_\delta(\hat{p})$ so that $p' \in f(G')$ whenever $\phi_*(p') = \hat{p}'$. In particular, there is $p' \in f(G')$ with

$$\begin{aligned} \sum_{k=1}^{|A_j|} p'_j(a_j^k) &\leq 2\delta && \text{for all } j \in N, \text{ and} \\ p'_j(a_j^{k,l}) &= \hat{p}'_j(a_j^k) && \text{for all } k \in [|A_j|], l \in [M], \text{ and } j \in N \setminus \{i\}. \end{aligned}$$

The latter condition can be satisfied since

$$M \sum_{k=1}^{|A_j|} \hat{p}'(a_j^k) \leq 2M\delta \leq (1 - 2\delta) \leq \hat{p}'(\hat{a}_j) = \sum_{a_j \in \phi_j^{-1}(\hat{a}_j)} p'(a_j),$$

for all $j \in N \setminus \{i\}$. Let $\Sigma'_i \subseteq \Sigma_{A'_i}$ be the set consisting of the identity permutation on U , and for all $j \in N \setminus \{i\}$, let $\Sigma'_j \subseteq \Sigma_{A'_j}$ be the set of permutations that map $\{a_j^k\} \cup \{a_j^{k,l} : l \in [M]\}$ to itself for all $k \in [|A_j|]$. Hence, a large set of clones of \hat{a}_j is permuted with each action a_j^k in all possible ways. Let $\Sigma' = \Sigma'_1 \times \cdots \times \Sigma'_n$. Note that, by construction, $p' \circ \pi = p'$ for $\pi \in \Sigma'$. Let

$$\bar{G} = \frac{1}{|\Sigma'|} \sum_{\pi \in \Sigma'} G' \circ \pi.$$

For all $j \in N \setminus \{i\}$ and $k \in [|A_j|]$, $\{a_j^k\} \cup \{a_j^{k,l} : l \in [M]\}$ is a set of clones in \bar{G} . Since $\hat{G}_i(a_i, \hat{a}_{-i}) > \hat{G}_i(\hat{a}_i, \hat{a}_{-i}) + \varepsilon$, $M = 2(\lceil \frac{1}{\varepsilon} \rceil + 1)$, and \hat{G} is normalized, $\bar{G}_i(a_i, a_{-i}) > \bar{G}_i(\hat{a}_i, a_{-i}) + \frac{\varepsilon}{2}$ for all $a_{-i} \in A'_{-i}$. Hence, a_i δ -dominates \hat{a}_i . Lastly, since f satisfies δ -consistency, there is $\bar{p} \in f(\bar{G})$ so that for all $j \in N$,

$$\bar{p}_j(\hat{a}_j) \geq p'_j(\hat{a}_j) - \delta \geq 1 - 3\delta.$$

□

Proof of Theorem 2. Given $\varepsilon > 0$, let $\delta > 0$ so that

$$3(n-1)\delta \leq \frac{\varepsilon}{4}, \tag{9}$$

$$4 \left(\left\lceil \frac{3}{\varepsilon} \right\rceil + 1 \right) 3\delta \leq (1 - 6\delta), \text{ and} \tag{10}$$

$$1 - 9\delta > 1/2. \tag{11}$$

Assume that f satisfies δ -consequentialism, δ -consistency, and δ -rationality, but is *not* a refinement of $NASH_\varepsilon$ on the set of normalized games. Then, there is a normalized game G on A , $i \in N$, $a_i \in A_i$, and $p \in f(G)$ so that

$$G_i(a_i, p_{-i}) > G_i(p_i, p_{-i}) + \varepsilon. \tag{12}$$

Observe that for all $q \in B_{3\delta}(p)$ and $p'_i \in \Delta A_i$,

$$|G_i(p'_i, p_{-i}) - G_i(p'_i, q_{-i})| \leq 3(n-1)\delta. \tag{13}$$

By Lemma 13, there is a game \hat{G} with action set $\hat{A}_j = A_j \cup \{\hat{a}_j\}$ for all $j \in N$ so that the following holds.

1. There is $q \in B_{3\delta}(p)$ such that $\hat{G}(\hat{a}_I, a_{-I}) = G(q_I, a_{-I})$ for all $I \subseteq N$ and $a_{-I} \in A_{-I}$.

2. There is $\hat{p} \in f(\hat{G})$ such that $\hat{p}_j(\hat{a}_j) \geq 1 - 3\delta$ for all $j \in N$.

Applying 1 with $I = N \setminus \{i\}$, it follows from (9), (12), and (13) that

$$\begin{aligned} \hat{G}_i(a_i, \hat{a}_{-i}) &= G_i(a_i, q_{-i}) \geq G_i(a_i, p_{-i}) - 3(n-1)\delta \\ &> G_i(p_i, p_{-i}) + 3/4\varepsilon \geq G_i(p_i, q_{-i}) + \frac{2}{4}\varepsilon \\ &\geq G_i(q_i, q_{-i}) + 1/4\varepsilon = \hat{G}_i(\hat{a}_i, \hat{a}_{-i}) + 1/4\varepsilon. \end{aligned}$$

By (9) and (10), Lemma 14 applied to $\frac{\varepsilon}{4}$ and 3δ gives a game G' and $\bar{p} \in f(G')$ so that \bar{p}_i assigns probability at least $1 - 9\delta$ to a 3δ -dominated action. Since $1 - 9\delta > 1/2$ by (11), this contradicts δ -rationality. \square