A fixed point free proof of Nash's Theorem via exchangeable equilibria

Noah D. Stein, Pablo A. Parrilo, and Asuman Ozdaglar*

April 9, 2010

Abstract

We prove existence of Nash equilibria in all finite games without using fixed point theorems or path following arguments. To do so we introduce the notion of exchangeable equilibria, which are correlated equilibria with certain symmetry and factorization properties. We prove these exist by adapting Hart and Schmeidler's proof of correlated equilibrium existence. Modifying Papadimitriou's correlated equilibrium algorithm in the same way, we can compute exchangeable equilibria in polynomial time.

In an appropriate limit exchangeable equilibria converge to the convex hull of Nash equilibria, proving that these exist as well (but not in polynomial time). Exchangeable equilibria are defined in terms of symmetries of the game, so this method automatically proves the stronger statement that a symmetric game has a symmetric Nash equilibrium. The case without symmetries follows by a symmetrization argument.

1 Introduction

Nash's Theorem is one of the most fundamental results in game theory and states that any finite game has a Nash equilibrium in mixed strategies. Despite its importance, the authors of the present paper know of only two essentially different proofs. The first and most common way to prove Nash's Theorem is via a fixed point theorem, usually Brouwer's or Kakutani's. The fixed point theorem is usually proven combinatorially, say by Sperner's Lemma [12] or Gale's argument using the game of hex [5], or with (co-)homology theory, a suite of powerful but less elementary tools from algebraic topology [8].

The other known proof of Nash's Theorem is algorithmic and consists of showing that the Lemke-Howson path-following algorithm terminates at a Nash equilibrium [9]. In fact this is not so different from the fixed point proof, because Sperner's Lemma is also proven by a path-following argument. Nonetheless, both proofs have provided unique insights into

*Department of Electrical Engineering, Massachusetts Institute of Technology: Cambridge, MA 02139. nstein@mit.edu, parrilo@mit.edu, and asuman@mit.edu.

the structure of Nash equilibria and it is our hope that a different proof, which uses neither fixed point theorems nor path-following arguments, will provide further insights.

Hart and Schmeidler have proven the weaker result that correlated equilibria exist by a clever application of the Minimax Theorem [7]. For games endowed with a group action, a simple averaging argument then proves that a symmetric correlated equilibrium exists. We show that for such games Hart and Schmeidler's proof can be strengthened to produce correlated equilibria with additional symmetry and factorization properties, which we call exchangeable equilibria.

To illustrate this idea, consider the case of $k \times k$ symmetric bimatrix games ($B = A^T$). Let $X = \{xx^T \mid x \in \mathbb{R}^{k\times 1}_{\geq 0}\}$. Then we have

$$\overset{\text{Nash}}{CE \cap X} \subseteq \overset{\text{convex hull of Nash}}{\operatorname{conv}(CE \cap X)} \subseteq \overset{\text{exchangeable}}{CE \cap \operatorname{conv}(X)} \subseteq \overset{\text{correlated}}{CE},$$

where each type of (symmetric) equilibrium is defined by the set written below it. This definition shows that in some sense the exchangeable equilibria are a natural mathematical object. For examples and game theoretic interpretations of exchangeable equilibria , see the companion paper [13].

It is evident that the set of exchangeable equilibria is convex, compact, contained in the set of symmetric correlated equilibria, and contains the convex hull of the set of symmetric Nash equilibria. One can show that these containments can all be strict [13], so proving existence of exchangeable equilibria is a step in the right direction, but does not immediately prove existence of Nash equilibria.

However, we can use the same minimax techniques to prove existence of exchangeable equilibria with stronger incentive compatibility properties, which we call order m exchangeable equilibria. These converge to mixtures of symmetric Nash equilibria as m goes to infinity. In particular, this proves that symmetric Nash equilibria exist in symmetric bimatrix games.

Note that symmetry is fundamental in this argument. For example, if we had begun with a general bimatrix game and let $X = \{xy^T \mid x, y \in \mathbb{R}^{k\times 1}_{\geq 0}\}$ we would have had $\operatorname{conv}(X) = \mathbb{R}^{k\times k}_{\geq 0}$, so the exchangeable equilibria (even the order m exchangeable equilibria) would have been exactly the correlated equilibria and we would not have strengthened Hart and Schmeidler's argument at all. However, there are several ways of turning general games into symmetric games [6] and applying such a procedure proves existence of Nash equilibria in general.

Up to the step of taking m to infinity, all the steps of our proof are computationally effective. Papadimitriou has shown how to apply the ellipsoid algorithm to Hart and Schmeidler's proof to efficiently compute correlated equilibria of large games [11]. The same technique applied to our proof allows one to compute exchangeable equilibria (and order m exchangeable equilibria for fixed m) in polynomial time, even though the set of these is not polyhedral. Computing these is interesting in its own right [13] and may be useful for computing approximate Nash equilibria. However, computation is not the focus of this paper and we leave a detailed investigation of these ideas for future work.

The remainder of the paper is organized as follows. We begin with background material in Section 2. We cover the definitions of games and equilibria, give an overview of Hart and Schmeidler's existence proof so we can modify it later, and introduce group actions. In Section 3 we introduce exchangeable equilibria and prove existence of these for games under arbitrary group actions. We do the same for order m exchangeable equilibria in Section 4. We complete the argument in Section 5 by showing that the order m exchangeable equilibria converge to mixtures of Nash equilibria under a certain condition (called player transitivity) on the symmetries of the game, and then showing that we can symmetrize any game to make this condition hold. Section 6 concludes and gives directions for future work.

2 Background

This section is divided into three parts. In the first we lay out the basic definitions of finite games as well as Nash and correlated equilibria to fix notation. We assume the reader is familiar with these concepts and do not attempt to motivate them. The second part reviews Hart and Scheidler's proof of the existence of correlated equilibria [7], preparing for similar arguments later in the paper. The third part covers symmetries of games.

The concept of a symmetry of a game extends back at least to Nash's paper [10]. Symmetries are fundamental to the present paper, so we spend more time on these and give some examples. Although we use the language of group theory to discuss symmetries, it is worth noting that we do not use any but the most basic theorems from group theory (e.g., the fact that for any h in a group G, the maps $g \mapsto gh$ and $g \mapsto hg$ are bijections from G to G). Everything in this section is standard except for Definitions 2.11 and 2.18 and the remarks following the statement of Nash's Theorem.

2.1 Games and equilibria

Definition 2.1. A **(finite) game** has a finite set I of $n \geq 2$ players, each with a finite set C_i of at least two **strategies** (also called **pure strategies**) and a **utility function** $u_i : C \to \mathbb{R}$, where $C = \prod C_i$. A game is **zero-sum** if it has two players, called the **maximizer** (denoted M) and the **minimizer** (denoted m), and satisfies $u_M + u_m = 0$.

For elements of C_i we use Roman letters subscripted with the player's identity, such as s_i and t_i. We will typically use the unsubscripted letter s to denote a strategy profile (a choice of strategy for each player). For a choice of a strategy for all players except i we use the symbol s_{-i}. To denote the set of Borel probability distributions on a space X we write $\Delta(X)$. For most of the paper X will be finite so we can view $\Delta(X)$ as a convex subset of the finite-dimensional vector space $\mathbb{R}^X$ of real-valued functions on X. For $x \in X$ the probability distribution which assigns unit mass to x will be written $\delta_x \in \Delta(X)$.

Definition 2.2. A **mixed strategy** for player i is a probability distribution over his pure strategy set C_i, and the set of mixed strategies for player i is $\Delta(C_i)$. The set of **independent distributions** or **mixed strategy profiles** will be denoted $\Delta^\Pi(\Gamma) = \prod_i \Delta(C_i)$.

For independent distributions it is important that we write $\Delta^\Pi(\Gamma)$ rather than $\Delta^\Pi(C)$, because Γ specifies how C is to be thought of as a product. For example, the set $S \times S \times S$

could be viewed as a product of three copies of S, or a product of S with $S \times S$, and these lead to different notions of an independent distribution – one is a product of three terms and one is a product of two terms. This distinction will be particularly important when we define powers of games in Section 4.

To make the notation fit together we will write $\Delta(\Gamma)$ for $\Delta(C)$. We may then view $\Delta^\Pi(\Gamma)$ as the (nonconvex) subset of $\Delta(\Gamma)$ consisting of product distributions or as a convex subset of $\mathbb{R}^{\sqcup_i C_i}$. Which of these views we are using will be clear from context if not explicitly specified.

As usual we extend the domain of u_i from C to $\Delta(\Gamma)$ by linearity, defining $u_i(\pi) = \sum_{s \in C} u_i(s)\pi(s)$. Having done so we can define equilibria.

Definition 2.3. A **Nash equilibrium** is an n-tuple $(\rho_1, \ldots, \rho_n) \in \Delta^\Pi(\Gamma) = \prod_i \Delta(C_i)$ of mixed strategies, one for each player, such that $u_i(s_i, \rho_{-i}) \leq u_i(\rho_i, \rho_{-i})$ for all strategies $s_i \in C_i$ and all players i. The set of Nash equilibria of a game Γ is denoted $\mathrm{NE}(\Gamma)$.

Definition 2.4. A **correlated equilibrium** is a joint distribution $\pi \in \Delta(\Gamma)$ such that $\sum_{s_{-i} \in C_{-i}} [u_i(t_i, s_{-i}) - u_i(s)]\, \pi(s) \leq 0$ for all strategies $s_i, t_i \in C_i$ and all players i. The set of correlated equilibria of a game Γ is denoted $\mathrm{CE}(\Gamma)$.

Nash equilibria correspond exactly to the correlated equilibria which are product distributions, so viewing $\Delta^\Pi(\Gamma)$ as a subset of $\Delta(\Gamma)$ we can write $\mathrm{NE}(\Gamma) = \mathrm{CE}(\Gamma) \cap \Delta^\Pi(\Gamma)$. We introduce the existence theorems for correlated and Nash equilibria in Sections 2.2 and 2.3.

We need the Minimax Theorem at this point to define the value of a zero-sum game. It also plays an important role our proof of Nash's Theorem. The Minimax Theorem is perhaps the only result in game theory which could be said to be more fundamenal than Nash's Theorem. An elementary proof can be given using the separating hyperplane theorem [2].

Minimax Theorem. *Let U and V be finite-dimensional vector spaces with compact convex subsets $K \subset U$ and $L \subset V$. Let $\Phi : U \times V \to \mathbb{R}$ be a bilinear map. Then*

$$\sup_{x \in K} \inf_{y \in L} \Phi(x, y) = \inf_{y \in L} \sup_{x \in K} \Phi(x, y),$$

and the optima are attained.

Definition 2.5. Given a zero-sum game Γ, we can apply this theorem with $K = \Delta(C_M)$, $L = \Delta(C_m)$, and $\Phi = u_M$. The common value of these two optimization problems is called the **value** of the game and denoted $v(\Gamma)$. Maximizers on the left hand side are called **maximin strategies** and the set of such is denoted $\mathrm{Mm}(\Gamma) \subseteq \Delta(C_M)$.

The notion of strategic equivalence gives a way to turn structural information about a game into structural information about equilibria.

Definition 2.6. Two mixed strategies $\sigma_i, \tau_i \in \Delta(C_i)$ are said to be **strategically equivalent** if $u_j(\sigma_i, s_{-i}) = u_j(\tau_i, s_{-i})$ for all $s_{-i} \in C_{-i}$ and all players j.

Proposition 2.7. *If σ_i is strategically equivalent to τ_i for all i, then $(\sigma_1, \ldots, \sigma_n)$ is a Nash equilibrium if and only if $(\tau_1, \ldots, \tau_n)$ is a Nash equilibrium.*

Proof. Immediate from the definitions. ∎

2.2 Hart and Schmeidler's proof

In this section we recall the structure of Hart and Schmeidler's proof of the existence of correlated equilibria based on the Minimax Theorem [7]. The goal of this is to frame their argument in language which will allow us to extend it, redoing as little as possible of the work they have done. We will use similar arguments to prove Theorems 3.7 and 4.5.

Hart and Schmeidler's argument begins by associating with a game Γ a new zero-sum game Γ^0 and interpreting correlated equilibria of Γ in terms of this new game.

Definition 2.8. Given any game Γ, define a two-player zero-sum game Γ^0 with $C_M^0 = C$, $C_m^0 = \bigsqcup_i C_i^2$, and utilities

$$u_M^0(s,(r_i,t_i)) = -u_m^0(s,(r_i,t_i)) = \begin{cases} u_i(s) - u_i(t_i, s_{-i}) & \text{if } r_i = s_i, \\ 0 & \text{otherwise.} \end{cases}$$

Proposition 2.9. *Let Γ be any game. For any player i in Γ, $r_i \in C_i$, and $s \in C$ we have $u_M^0(s,(r_i,r_i)) = 0$, so we can bound the value of Γ^0 by $v(\Gamma^0) \leq 0$. A mixed strategy $\sigma \in \Delta(C_M^0) = \Delta(C)$ for the maximizer in Γ^0 satisfies $u_M^0(\sigma,(r_i,t_i)) \geq 0$ for all $(r_i,t_i) \in C_m^0$ if and only if $\sigma \in \mathrm{CE}(\Gamma)$. Therefore, if $v(\Gamma^0) = 0$ then $\mathrm{Mm}(\Gamma^0) = \mathrm{CE}(\Gamma)$.*

Proof. Immediate from the definitions. □

The bulk of the work of proving $v(\Gamma^0) = 0$, and hence the existence of correlated equilibria (Theorem 2.14), consists of proving Lemma 2.12. We will also use this lemma below to prove stronger statements in a similar spirit: Lemmas 3.6 and 4.4. These in turn allow us to strengthen Theorem 2.14, yielding Theorems 3.7 and 4.5.

To state Lemma 2.12, we need to define a family of auxiliary games $\gamma(y_i)$. For the purposes of the present paper, it is more important to understand the conclusion of the lemma than to remember the details of this construction.

Definition 2.10. For any player i in Γ and any nonnegative $y_i \in \mathbb{R}^{C_i \times C_i}$, define the zero-sum game $\gamma(y_i)$ with strategy sets $C_M = C_m = C_i$ and utilities

$$u_M(s_i,t_i) = -u_m(s_i,t_i) = \begin{cases} \sum_{r_i \neq t_i} y_i^{s_i,r_i} & \text{if } s_i = t_i, \\ -y_i^{s_i,t_i} & \text{otherwise.} \end{cases}$$

Definition 2.11. In a zero-sum game Γ, we say that a strategy $\sigma \in \Delta(C_M)$ for the maximizer is a **good reply to** $\theta \in \Delta(C_m)$ if $u_M(\sigma,\theta) \geq v(\Gamma)$. We say that a set $\Sigma \subseteq \Delta(C_M)$ of strategies is **good against the set** $\Theta \subseteq \Delta(C_m)$ if for all $\theta \in \Theta$ there is a $\sigma \in \Sigma$ which is a good reply to θ. If Σ is good against $\Delta(C_m)$ we say that Σ is **good**.

Lemma 2.12 ([7])**.** *Fix a game Γ and consider Γ^0. If $y \in \Delta(C_m^0)$, then any strategy $\pi \in \mathrm{Mm}(\gamma(y_1)) \times \cdots \times \mathrm{Mm}(\gamma(y_n)) \subset \Delta(C_M^0)$ satisfies $u_M(\pi,y) = 0$. In particular π is good against y and $\Delta^\Pi(\Gamma)$ is good.*

Proof. Omitted. See [7] for a proof using the Minimax Theorem. □

Proposition 2.13. *If Γ is a zero-sum game and $\Sigma \subseteq \Delta(C_M)$ is good, then Γ has a maximin strategy in $\overline{\mathrm{conv}}(\Sigma)$, i.e., $\overline{\mathrm{conv}}(\Sigma) \cap \mathrm{Mm}(\Gamma) \neq \emptyset$.*

Proof. Apply the Minimax Theorem with $K = \overline{\mathrm{conv}}(\Sigma)$ and $L = \Delta(C_m)$. $\square$

Theorem 2.14 ([7])**.** *For any game Γ, the value $v(\Gamma^0) = 0$, so $\mathrm{Mm}(\Gamma^0) = \mathrm{CE}(\Gamma)$ and a correlated equilibrium of Γ exists.*

Proof. Lemma 2.12 implies $v(\Gamma^0) \geq 0$, so by Proposition 2.9 we have $v(\Gamma^0) = 0$ and $\mathrm{Mm}(\Gamma^0) = \mathrm{CE}(\Gamma)$. Apply Proposition 2.13 to Γ^0 with $\Sigma = \Delta^\Pi(\Gamma)$. $\square$

This proof merits two remarks. First of all, since $\overline{\mathrm{conv}}(\Delta^\Pi(\Gamma)) = \Delta(\Gamma)$, Proposition 2.13 does not yield any benefit in this case over directly applying the Minimax Theorem to Γ^0. Rather, we have used Proposition 2.13 to illustrate our proof strategy for Theorems 3.7 and 4.5, in which we choose Σ with $\overline{\mathrm{conv}}(\Sigma) \subsetneq \Delta(\Gamma)$.

Second, note that in this case we know that there is a maximin strategy of Γ^0 in the good set $\Delta^\Pi(\Gamma)$: this is just the statement of Nash's Theorem. However, we cannot conclude this directly because in general a good set need not include a maximin strategy. For example, in any zero-sum game the set $C_M \subsetneq \Delta(C_M)$ is a good set, but some zero-sum games such as matching pennies only have mixed maximin strategies.

2.3 Groups acting on games

In this section we recall the notion of a group acting on a game, as defined by Nash [10]. All groups will be finite throughout. In any group e will denote the identity element. The subgroup generated by group elements $g_1, \ldots, g_n$ will be denoted $\langle g_1, \ldots, g_n \rangle$. For $n \in \mathbb{N}$ we will write $\mathbb{Z}_n$ for the additive group of integers mod n and S_n for the symmetric group on n letters. We will use cycle notation to express permutations. For example $\sigma = (1\ 2\ 3)(4\ 5)(6)$ is shorthand for

$$\sigma(1) = 2,\ \sigma(2) = 3,\ \sigma(3) = 1,\ \sigma(4) = 5,\ \sigma(5) = 4,\ \text{and } \sigma(6) = 6.$$

Definition 2.15. A **left action of the group G on the set X** is a map $\cdot : G \times X \to X$ written with infix notation which satisfies the identity condition $e \cdot x = x$ and the associativity condition $g \cdot (h \cdot x) = (gh) \cdot x$. A **right action of G on X** is a map $\cdot : X \times G \to X$ such that $x \cdot e = x$ and $(x \cdot g) \cdot h = x \cdot (gh)$.

We say that an action is **linear** if it extends to an action on an ambient vector space V containing X and the map $x \mapsto x \cdot g$ on V is linear for all $g \in G$. An $x \in X$ is **G-invariant** if $x \cdot g = x$ for all $g \in G$. The set of G-invariant elements is denoted X_G.

Proposition 2.16. *If G acts linearly on the convex set X then there is a map $\mathrm{ave}_G : X \to X_G$ given by $\mathrm{ave}_G(x) = \frac{1}{|G|} \sum_{g \in G} x \cdot g$. In particular if X is nonempty then X_G is nonempty.*

Proof. For any $x \in X$, $\text{ave}_G(x)$ is a convex combination of elements $x \cdot g \in X$, hence $\text{ave}_G(x) \in X$. For any $h \in G$ we have

$$\begin{aligned}\text{ave}_G(x) \cdot h &= \left[\frac{1}{|G|}\sum_{g\in G} x \cdot g\right] \cdot h = \frac{1}{|G|}\sum_{g\in G}(x \cdot g) \cdot h = \frac{1}{|G|}\sum_{g\in G} x \cdot (gh) \\ &= \frac{1}{|G|}\sum_{g\in G} x \cdot g = \text{ave}_G(x),\end{aligned}$$

where we have used linearity, the definition of a group action, and bijectivity of $g \mapsto gh$. □

A left action of G on X induces right actions on many function spaces defined on X. For example $\mathbb{R}^X$ is the space of functions $X \to \mathbb{R}$. For $y \in \mathbb{R}^X$ we can define $y \cdot g \in \mathbb{R}^X$ by $(y \cdot g)(x) = y(g \cdot x)$. The condition that this is a right action of G on $\mathbb{R}^X$ follows immediately from the fact that we began with a left action of G on X. For finite X (the case of most interest to us), the same argument shows that G acts on $\Delta(X)$ on the right.

Definition 2.17. We say that **a group G acts on the game** Γ if the following conditions hold. The group G acts on the the left on I and $\bigsqcup_i C_i$, making $g \cdot s_i \in C_{g\cdot i}$ for $s_i \in C_i$. Such actions automatically induce a left action of G on $C = \prod_i C_i$ defined by $(g \cdot s)_{g\cdot i} = g \cdot s_i$. We require that the utilities be invariant under the induced action on the right: $u_{g\cdot i} \cdot g = u_i$, i.e., $u_{g\cdot i}(g \cdot s) = u_i(s)$ for all $i \in I$, $s \in C$, and $g \in G$. We say that G is a **symmetry group of** Γ and call elements of G **symmetries of** Γ.

Note that an action of G on a game can be fully specified by its action on $\bigsqcup_i C_i$ or on C. One way to do this is to choose G to be a subgroup of the symmetric group on $\bigsqcup_i C_i$ or C satisfying the above properties.

Definition 2.18. The **stabilizer subgroup of player** i is $G_i = \{g \in G \mid g \cdot i = i\}$, and acts on C_i on the left. We say that the action of G is **player trivial** if $G_i = G$ for all i, or in other words if $g \cdot i = i$ for all g and i. We say that the action of G is **player transitive** if for all $i, j \in I$ there exists $g \in G$ such that $g \cdot i = j$.

We illustrate the notion of group actions on a game using four examples.

Example 2.19. Let Γ be any game and G any group. Define $g \cdot s = s$ for all $g \in G$ and $s \in C$. This defines a player-trivial action of G on Γ called the trivial action.

Example 2.20. A two-player finite game is often called a **bimatrix game** because it can be described by two matrices A and B, such that if player one plays strategy i and player two plays strategy j then their payoffs are A_{ij} and B_{ij}, respectively. If these matrices are square and $B = A^T$ then we call the game a **symmetric bimatrix game**. One example is the game of chicken, which has $A = \left[\begin{smallmatrix} 4 & 1 \\ 5 & 0 \end{smallmatrix}\right] = B^T$.

To put this in the context of group actions defined above, let each player's strategy set be $C_1 = C_2 = \{1, \dots, m\}$ indexing the rows and columns of A and B. Define $g \cdot (i, j) = (j, i)$ for $(i, j) \in C$, so $g \cdot (g \cdot (i, j)) = (i, j)$. The assumption $B = A^T$ is exactly the utility

(u_1, u_2)	H_2	T_2
H_1	$(1, -1)$	$(-1, 1)$
T_1	$(-1, 1)$	$(1, -1)$

Table 1: Matching pennies. Player 1 chooses rows and player 2 chooses columns.

compatibility condition saying that this specifies an action of $G = \{e, g\} \cong \mathbb{Z}_2$ on this game. Of course, depending on the structure of A and B there may be other nontrivial symmetries as well. The element g swaps the players, so the action of G is player transitive.

Example 2.21. Consider the game of matching pennies, whose utilities are shown in Table 1. The labels H and T stand for heads and tails, respectively, and the subscripts indicate the identities of the players for notational purposes. This a bimatrix game, but it is not a symmetric bimatrix game in the sense of Example 2.20.

Nonetheless this game does have symmetries. The easiest to see is the map σ which interchanges the roles of heads and tails. Letting g be the permutation of $\bigsqcup_i C_i$ given in cycle notation as $g = (H_1\ T_1)(H_2\ T_2)$, we define $g \cdot s_i = g(s_i)$. Another symmetry is the permutation $h = (H_1\ H_2\ T_1\ T_2)$. These satisfy $g^2 = e$ and $h^2 = g$, so $G = \langle h \rangle \cong \mathbb{Z}_4$. Note that g acts on I as the identity whereas h swaps the players, so G acts player transitively, whereas $\langle g \rangle \cong \mathbb{Z}_2$ acts player trivially.

Example 2.22. Now we consider an example of an n-player game with symmetries. Throughout this example all arithmetic will be done mod n. For simplicity in this example we will index the players using the members of $\mathbb{Z}_n$ instead of the set $\{1, \ldots, n\}$. Each player's strategy space will be $C_i = \mathbb{Z}_n$ as well. Define

$$u_i(s_1, \ldots, s_n) = \begin{cases} 1, & \text{when } s_i = s_{i-1} + 1 \\ 0, & \text{otherwise.} \end{cases}$$

Then we can define a symmetry g by $g(s_i) = s_i + 1$, which increments each player's strategy by one mod n, but fixes the identities of the players. Clearly g is a permutation of order n.

We can define another symmetry h which maps a strategy for player i to the same numbered strategy for player $i+1$. That is to say, h acts on C by cyclically permuting its arguments. Again, h is a permutation of order n. Note that g and h commute, so together they generate a symmetry group $G \cong \mathbb{Z}_n \times \mathbb{Z}_n$. Both $\langle h \rangle \cong \mathbb{Z}_n$ and G act player transitively, whereas $\langle g \rangle \cong \mathbb{Z}_n$ acts player trivially. If n is composite and factors as $n = kl$ for $k, l > 1$ then $\langle h^k \rangle \cong \mathbb{Z}_l$ acts on Γ but neither player transitively nor player trivially.

The left actions in the definition of a group action on a game induce linear right actions on function spaces such as $\Delta(\Gamma) \subsetneq \mathbb{R}^C$ and $\Delta^\Pi(\Gamma) \subsetneq \mathbb{R}^{\sqcup_i C_i}$. The inclusion map $\mathbb{R}^{\sqcup_i C_i} \to \mathbb{R}^C$ is G-equivariant (commutes with the action of G), so with regard to this action it does not matter whether we choose to view $\Delta^\Pi(\Gamma)$ as a subset of $\mathbb{R}^{\sqcup_i C_i}$ or of $\mathbb{R}^C$.

Because of the utility compatibility conditions of a group action on a game, the actions on $\Delta(\Gamma)$ and $\Delta^\Pi(\Gamma)$ restrict to actions on the sets $\mathrm{CE}(\Gamma)$ and $\mathrm{NE}(\Gamma)$, respectively. This allows us to define the G-invariant subsets $\Delta_G(\Gamma)$, $\Delta_G^\Pi(\Gamma)$, $\mathrm{CE}_G(\Gamma)$, and $\mathrm{NE}_G(\Gamma)$. The action of the stabilizer subgroup G_i on C_i allows us to define the G-invariant subset $\Delta_{G_i}(C_i)$.

The main theorem we set out to prove is the following. This theorem is most often applied in the case where G is the trivial group, but Nash proved the general case in [10] and so shall we.

Nash's Theorem. *A game with symmetry group G has a G-invariant Nash equilibrium.*

To prove this we will use Hart and Scheidler's techniques in a new way. We will show that certain classes of symmetric games have correlated equilibria with a much higher degree of symmetry than might be expected without knowledge of Nash's Theorem. To illustrate what we mean, consider the following trivial improvement on Theorem 2.14.

Proposition 2.23. *A game with symmetry group G has a G-invariant correlated equilibrium.*

Proof. Apply Proposition 2.16 to a correlated equilibrium, which exists by Theorem 2.14. □

A priori we might not expect correlated equilibria with a greater degree of symmetry than predicted by Proposition 2.23 to exist. But viewing G-invariant Nash equilibria as correlated equilibria, we see that we can often guarantee much more. Suppose we have an n-player game which has identical strategy sets for all players and which is symmetric under cyclic permutations of the players, such as the game in Example 2.22. Then Proposition 2.23 yields a correlated equilibrium π which is invariant under cyclic permutations of the players, but need not be invariant under other permutations. On the other hand the Nash equilibrium $\rho = (\rho_1, \ldots, \rho_n)$ given by Nash's Theorem satisfies $\rho_1 = \ldots = \rho_n$ so the corresponding product distribution $\pi(s_1, \ldots, s_n) = \rho_1(s_1) \cdots \rho_1(s_n)$ is a correlated equilibrium which is invariant under arbitrary permutations of the players.

3 Exchangeable equilibria

In this section we prove the existence of correlated equilibria with this higher degree of symmetry, as well as a useful factorization property, without appealing to Nash's Theorem. The proofs of some of the propositions in this section are direct algebraic manipulations and these are omitted.

Definition 3.1. Viewing $\Delta_G^\Pi(\Gamma)$ as a nonconvex subset of the convex set $\Delta_G(\Gamma)$, we define the set of ***G*-exchangeable** probability distributions

$$\Delta_G^X(\Gamma) = \operatorname{conv} \Delta_G^\Pi(\Gamma) \subseteq \Delta_G(\Gamma).$$

We use the term "exchangeable" because of the important case where the C_i are all equal and the group G acts player transitively. Then distributions in $\Delta_G^X(\Gamma)$ are invariant under arbitrary permutations of the players. Furthermore, by De Finetti's Theorem these are exactly the distributions which can be extended to exchangeable distributions on infinitely many copies of C_1, i.e., distributions invariant under permutations of finitely many indices.

De Finetti's Theorem will not play a role in our analysis; here it merely serves to motivate Definition 3.1.

To get a feel for these sets, we will look at them in the context of some examples.

Example 2.19 (cont'd). Since G acts trivially we can ignore it entirely. Not all distributions are independent so $\Delta^{\Pi}_G(\Gamma) \subsetneq \Delta_G(\Gamma) = \Delta(\Gamma)$, but $\Delta^X_G(\Gamma) = \Delta_G(\Gamma)$. As we have seen, one inclusion is automatic. To prove the reverse note that for any $s \in C$, $\delta_s = \delta_{s_1} \cdots \delta_{s_n} \in \Delta^{\Pi}(\Gamma) = \Delta^{\Pi}_G(\Gamma)$. But for any $\pi \in \Delta(\Gamma)$ we can write $\pi = \sum_{s \in C} \pi(s)\delta_s$, and such a convex combination of the δ_s is in $\Delta^X_G(\Gamma)$ by definition.

Example 2.20 (cont'd). For a symmetric bimatrix game Γ with m strategies per player, we can view probability distributions over C as $m \times m$ nonnegative matrices with entries summing to unity. The nontrivial symmetry $g \in G$ acts by swapping the players. From the definitions we see that $\Delta_G(\Gamma)$ consists of symmetric matrices and $\Delta^{\Pi}_G(\Gamma)$ of matrices which are outer products xx^T for nonnegative column vectors $x \in \mathbb{R}^m$. The elements of $\Delta^X_G(\Gamma) = \operatorname{conv} \Delta^{\Pi}_G(\Gamma)$ are exactly the (normalized) **completely positive matrices** studied in [1]. Clearly all such matrices are symmetric, elementwise nonnegative, and positive semidefinite; it turns out the converse holds if and only if $m \leq 4$ [3].

Example 2.21 (cont'd). The map on C induced by h is the permutation

$$((H_1, H_2)\ (T_1, H_2)\ (T_1, T_2)\ (H_1, T_2)).$$

In particular, a G-invariant probability distribution must assign equal probability to all four outcomes in $C_1 \times C_2$. There is only one such distribution and it is independent, so $\Delta^{\Pi}_G(\Gamma) = \Delta^X_G(\Gamma) = \Delta_G(\Gamma)$.

Example 2.22 (cont'd). Recall that in this game there are n players and the C_i are the same for all i. The group G permutes the players cyclically. Therefore the elements of $\Delta^{\Pi}_G(\Gamma)$ are invariant under arbitrary permutations of the players, hence so are the elements of $\Delta^X_G(\Gamma)$. (The converse statement is false; that is to say, there are probability distributions over C which are invariant under arbitrary permutations of the players but are not in $\Delta^X_G(\Gamma)$. This is analogous to the presence in Example 2.20 of symmetric elementwise nonnegative matrices which are not positive semidefinite, hence not completely positive.) On the other hand, an element of $\Delta_G(\Gamma)$ need only be invariant under cyclic permutations of the players.

Definition 3.2. The set of G**-exchangeable equilibria** of a game Γ is

$$\mathrm{XE}_G(\Gamma) = \mathrm{CE}(\Gamma) \cap \Delta^X_G(\Gamma).$$

When G can be inferred from context we simply refer to **exchangeable equilibria**.

It is immediate from the definitions that $\operatorname{conv}(\mathrm{NE}_G(\Gamma)) \subseteq \mathrm{XE}_G(\Gamma) \subseteq \mathrm{CE}_G(\Gamma)$. There are examples in which all of these inclusions are strict [13], so proving non-emptiness of $\mathrm{XE}_G(\Gamma)$ does not immediately prove non-emptiness of $\mathrm{NE}_G(\Gamma)$. Nonetheless, this is an important step and the main result of this section.

The proof that a G-exchangeable equilibrium exists proceeds along the same lines as the correlated equilibrium existence proof in Section 2.2. We again consider the zero-sum game

Γ^0 and prove that a certain set is good in this game (Lemma 3.6). The difference is that the action of G yields a smaller good set, $\Delta^{\Pi}_G(\Gamma)$. To prove this lemma we need the following three symmetry results, which have straightforward proofs.

Proposition 3.3. *If G acts on Γ then G acts player trivially on Γ^0 by $g \cdot (s, (r_i, t_i)) = (g \cdot s, (g \cdot r_i, g \cdot t_i))$.*

Proposition 3.4. *If G acts player trivially on a zero-sum game, then a set $\Sigma \subseteq \Delta_G(C_M)$ is good if and only if it is good against $\Delta_G(C_m)$.*

Proof. For all $g \in G$, $\sigma \in \Delta_G(C_M)$, and $\theta \in \Delta(C_m)$ we have $u_M(\sigma, \theta \cdot g) = u_M(\sigma \cdot g, \theta \cdot g) = u_M(\sigma, \theta)$, so $u_M(\sigma, \theta) = u_M(\sigma, \mathrm{ave}_G(\theta))$. □

Proposition 3.5. *The map $y \mapsto \mathrm{Mm}(\gamma(y))$ is natural in the sense that if $\sigma : C_i \to C_j$ is a bijection and $y_i = y_j \circ (\sigma, \sigma)$, then composition with σ maps $\mathrm{Mm}(\gamma(y_j))$ to $\mathrm{Mm}(\gamma(y_i))$.*

Lemma 3.6. *If G acts on the game Γ then in the game Γ^0 the set $\Delta^{\Pi}_G(\Gamma)$ is good.*

Proof. By Proposition 3.3 and Proposition 3.4, it suffices to consider only $y \in \Delta_G(C^0_m)$, and show that there is a $\pi \in \Delta^{\Pi}_G(\Gamma)$ which is good against y. Lemma 2.12 states that any $\pi \in S(y) = \mathrm{Mm}(\gamma(y_1)) \times \cdots \times \mathrm{Mm}(\gamma(y_n)) \subset \Delta^{\Pi}(\Gamma)$ is good against y.

By Proposition 3.5 the action of G on $\Delta^{\Pi}(\Gamma)$ restricts to an action of G on $S(y)$ since $y \in \Delta_G(C^0_m)$. Viewing $S(y)$ as a convex subset of $\mathbb{R}^{\sqcup_i C_i}$, Proposition 2.16 shows the invariant subspace $S_G(y) \subseteq \Delta^{\Pi}_G(\Gamma)$ is nonempty, so $\Delta^{\Pi}_G(\Gamma)$ is good. □

Theorem 3.7. *A game with symmetry group G has a G-exchangeable equilibrium.*

Proof. By Theorem 2.14, $\mathrm{Mm}(\Gamma^0) = \mathrm{CE}(\Gamma)$. Lemma 3.6 shows we can apply Proposition 2.13 to Γ^0 with $\Sigma = \Delta^{\Pi}_G(\Gamma)$, proving that $\mathrm{Mm}(\Gamma^0) \cap \Delta^X_G(\Gamma) = \mathrm{XE}_G(\Gamma)$ is nonempty. □

4 Higher order exchangeable equilibria

In this section we begin with a game Γ and artificially add symmetries to produce games $\Pi^m\Gamma$ and $\Xi^m\Gamma$ with larger symmetry groups for each $m \in \mathbb{N}$. Having constructed these games, we can exploit our knowledge of their structure to improve Theorem 3.7 and show that there are distributions which are simultaneously exchangeable equilibria of both $\Pi^m\Gamma$ and $\Xi^m\Gamma$. We call such distributions order m exchangeable equilibria.

We then use a compactness argument to exhibit a distribution which is simultaneously an order m exchangeable equilibrium for all $m \in \mathbb{N}$, called an order ∞ exchangeable equilibrium. We will see in the next section that for player-transitive symmetry groups, an order ∞ exchangeable equilibrium is just a mixture of symmetric Nash equilibria.

Most of the work in this section consists of making the proper definitions. Once that is done, the proofs are rather short.

4.1 Order m G-exchangeable equilibria

To define order m G-exchangeable equilibria we will need two notions of a power of a game.

Definition 4.1. For $m \in \mathbb{N}$, the m^{th} **power of** Γ, denoted $\Pi^m\Gamma$, is a game in which m independent copies of Γ are played simultaneously. More specifically, $\Pi^m\Gamma$ has mn players labeled by pairs i, j, $1 \leq i \leq n$, $1 \leq j \leq m$, strategy spaces $\Pi^m C_{ij} = C_i$ for all i, j and utilities $\Pi^m u_{ij}(s_{11}, \ldots, s_{nm}) = u_i(s_{1j}, s_{2j}, \ldots, s_{nj})$.

The **contracted** m^{th} **power of** Γ, denoted $\Xi^m\Gamma$, is a game in which m copies of Γ are played simultaneously, but all by the same set of players. Specifically, $\Xi^m\Gamma$ has n players, strategy spaces $\Xi^m C_i = C_i^m$ with generic element $(s_{i1}, \ldots, s_{im})$ for all i, and utilities $\Xi^m u_i(s_{11}, \ldots, s_{nm}) = \sum_j u_i(s_{1j}, s_{2j}, \ldots, s_{nj})$.

To motivate the definition of an order m G-exchangeable equilibrium, we first need to establish a few properties of these powers. One can construct examples showing that in general none of the inclusions in this proposition can be reversed. In particular, no containment holds between $\mathrm{XE}_{G\times S_m}(\Pi^m\Gamma)$ and $\mathrm{XE}_{G\times S_m}(\Xi^m\Gamma)$ in either direction. This is connected to the fact that the inclusion between the sets of correlated equilibria of $\Pi^m\Gamma$ and $\Xi^m\Gamma$ goes in the opposite direction from the inclusion between the sets of Nash equilibria.

Proposition 4.2. *Let Γ be a game with symmetry group G and fix $m \in \mathbb{N}$. Then both powers $\Pi^m\Gamma$ and $\Xi^m\Gamma$ are games with symmetry group $G \times S_m$ and they satisfy:*

- $\Delta^X_{G\times S_m}(\Pi^m\Gamma) \subseteq \Delta^X_{G\times S_m}(\Xi^m\Gamma)$ *and*
- $\mathrm{NE}_{G\times S_m}(\Pi^m\Gamma) \subseteq \mathrm{NE}_{G\times S_m}(\Xi^m\Gamma) \subseteq \mathrm{CE}_{G\times S_m}(\Xi^m\Gamma) \subseteq \mathrm{CE}_{G\times S_m}(\Pi^m\Gamma)$.

Proof. Both powers are invariant under arbitrary permutations of the copies and under symmetries in G applied to all of the copies simultaneously. In fact in the case of $\Pi^m\Gamma$ we can apply a different symmetry in G to each copy independently so that $\Pi^m\Gamma$ is invariant under the larger group $G \wr S_m$ (the wreath product of G and S_m), but we will not need this fact.

Since $G \times S_m$ acts on $\Pi^m C$ and $\Xi^m C$ in the same way, $\Delta_{G\times S_m}(\Pi^m\Gamma) = \Delta_{G\times S_m}(\Xi^m\Gamma)$. The game $\Pi^m\Gamma$ has more players than $\Xi^m\Gamma$, so $\Delta^\Pi_{G\times S_m}(\Pi^m\Gamma)$ has stronger independence conditions than $\Delta^\Pi_{G\times S_m}(\Xi^m\Gamma)$. Therefore $\Delta^\Pi_{G\times S_m}(\Pi^m\Gamma) \subsetneq \Delta^\Pi_{G\times S_m}(\Xi^m\Gamma)$. Taking convex hulls yields the relation between the exchangeable distributions.

To prove the containments between the equilibrium sets, it is easiest to consider the third containment first. Any strategy deviation available to a copy of player i in $\Pi^m\Gamma$ can be applied by player i in $\Xi^m\Gamma$ to the corresponding copy of the game, proving (the contrapositive of) the containment of the correlated equilibrium sets.

Call strategy deviations of this type "limited". If players of $\Xi^m\Gamma$ choose mixed strategies independently, then the correlated equilibrium constraints for limited deviations suffice to imply these constraints for all deviations, due to the additive separability of the utility functions in $\Xi^m\Gamma$. That is to say,

$$\mathrm{CE}_{G\times S_m}(\Pi^m\Gamma) \cap \Delta^\Pi_{G\times S_m}(\Xi^m\Gamma) = \mathrm{CE}_{G\times S_m}(\Xi^m\Gamma) \cap \Delta^\Pi_{G\times S_m}(\Xi^m\Gamma).$$

Combined with the definitions and the containments already proven, this proves the rest of the containments. $\square$

Suppose $(\rho_1, \ldots, \rho_n)$ were a G-invariant Nash equilibrium of Γ. Taking the product of the independent distribution $\rho = \rho_1 \cdots \rho_n$ with itself m times, we obtain a distribution ρ^m which is clearly a $(G \times S_m)$-invariant Nash equilibrium of both $\Pi^m\Gamma$ and $\Xi^m\Gamma$. This serves to motivate the following definition in the sense that we should expect it not to be vacuous.

Definition 4.3. The set of **order** m G**-exchangeable equilibria of** Γ is

$$\mathrm{XE}_G^m(\Gamma) = \mathrm{XE}_{G\times S_m}\left(\Pi^m\Gamma\right) \cap \mathrm{XE}_{G\times S_m}\left(\Xi^m\Gamma\right) = \Delta^X_{G\times S_m}\left(\Pi^m\Gamma\right) \cap \mathrm{CE}\left(\Xi^m\Gamma\right).$$

We discuss some properties of these before proving they exist. To unpack the definition, let X_i^j, $1 \le i \le n$, $1 \le j \le m$ be random variables distributed according to $\pi \in \Delta_G^X(\Pi^m\Gamma)$, so X_i^j takes values in C_i. Then π is an order m G-exchangeable equilibrium if and only if X_i^1 is a best response for player i to X_{-i}^1 given that player i knows $X_i^1, \ldots, X_i^m$. Equivalently by the symmetry of π, X_i^j is a best response to X_{-i}^k given $X_i^1, \ldots, X_i^m$ for all j and k.

Averaging over possible values of X_i^m, we see that these conditions imply X_i^1 is also a best response to X_{-i}^1 given only $X_i^1, \ldots, X_i^{m-1}$. So the marginalization map from $\Delta^X_{G\times S_m}(\Pi^m\Gamma)$ to $\Delta^X_{G\times S_{m-1}}(\Pi^{m-1}\Gamma)$ sends $\mathrm{XE}_G^m(\Gamma)$ into $\mathrm{XE}_G^{m-1}(\Gamma)$. In this way we can view the sets of higher order exchangeable equilibria as being nested.

We will show that the set of order m G-exchangeable equilibria approaches the closure of the convex hull of the set of Nash equilibria in some sense as m increases (Theorem 5.1 below). The corresponding statement with the set of $(G \times S_m)$-exchangeable equilibria of either the m^{th} power or of the m^{th} contracted power in place of the order m G-exchangeable equilibria is false.

To see this, consider the natural marginalization map which sends an element of $\Delta_{G\times S_n}\left(\Pi^m\Gamma\right) = \Delta_{G\times S_n}\left(\Xi^m\Gamma\right)$ to $\Delta_G(\Gamma)$. One can show that for all Γ, G, and m, both the image of $\mathrm{XE}_{G\times S_m}\left(\Pi^m\Gamma\right)$ and the image of $\mathrm{XE}_{G\times S_m}\left(\Xi^m\Gamma\right)$ under this map are equal to $\mathrm{XE}_G(\Gamma)$. One can also give an explicit example showing that $\mathrm{XE}_G(\Gamma)$ may be strictly larger than $\mathrm{conv}(\mathrm{NE}_G(\Gamma))$ [13]. The lack of convergence of $\mathrm{XE}_{G\times S_m}(\Pi^m\Gamma)$ or $\mathrm{XE}_{G\times S_m}(\Xi^m\Gamma)$ to the convex hull of the Nash equilibria is the motivation for the definition of order m G-exchangeable equilibria.

Lemma 4.4. *If G acts on the game Γ then in the game $(\Xi^m\Gamma)^0$ the set $\Delta^\Pi_{G\times S_m}(\Pi^m\Gamma)$ is good.*

Proof. By Lemma 3.6, $\Delta^\Pi_{G\times S_m}(\Xi^m\Gamma)$ is good. The utilities in $\Xi^m\Gamma$ are additively separable, so any mixed strategy in $\Delta^\Pi_{G\times S_m}(\Xi^m\Gamma)$ is strategically equivalent for the maximizer in $(\Xi^m\Gamma)^0$ to a mixed strategy in $\Delta^\Pi_{G\times S_m}(\Pi^m\Gamma)$. $\square$

Theorem 4.5. *A game with symmetry group G has an order m G-exchangeable equilibrium for all $m \in \mathbb{N}$.*

Proof. By Theorem 2.14, $\mathrm{Mm}((\Xi^m\Gamma)^0) = \mathrm{CE}(\Xi^m\Gamma)$. Lemma 4.4 shows we can apply Proposition 2.13 to $(\Xi^m\Gamma)^0$ with $\Sigma = \Delta^\Pi_{G\times S_m}(\Pi^m\Gamma)$, so $\mathrm{Mm}((\Xi^m\Gamma)^0) \cap \Delta^X_{G\times S_m}(\Pi^m\Gamma) = \mathrm{XE}_G^m(\Gamma)$ is nonempty. $\square$

4.2 Order ∞ G-exchangeable equilibria

Next we use a compactness argument to prove existence of an order ∞ G-exchangeable equilibrium, a distribution which is in some sense an order m G-exchangeable equilibrium for all finite m. As we have defined them the $\mathrm{XE}_G^m(\Gamma)$ are distributions over different numbers of copies of C, so they are not directly comparable and we can't just take their intersection.

Define a map $\mu_m : \Delta(\Delta_G^\Pi(\Gamma)) \to \Delta_{G\times S_m}^X(\Pi^m\Gamma)$ as follows. For $\pi \in \Delta(\Delta_G^\Pi(\Gamma))$, let R be a random variable taking values in $\Delta_G^\Pi(\Gamma)$ distributed according to π. Let X_i^j, $1 \le i \le n$ and $1 \le j < \infty$ be random variables taking values in C_i which are independent given R and distributed according to R_i. Let $\mu_m(\pi)$ be the joint distribution of X_i^j for $1 \le i \le n$ and $1 \le j \le m$. In terms of equations, the set of strategy profiles of $\Pi^m\Gamma$ is C^m and for $(s^1,\ldots,s^m) \in C^m$ we have

$$[\mu_m(\pi)]\,(s^1,\ldots,s^m) = \int R(s^1)\cdots R(s^m)d\pi(R).$$

Define $A_m = \mu_m^{-1}(\mathrm{XE}_G^m(\Gamma))$. Elements of A_m are representations of order m G-exchangeable equilibria as mixtures of independent G-invariant distributions.

Definition 4.6. The set of order ∞ G-exchangeable equilibria is $\mathrm{XE}_G^\infty(\Gamma) = \bigcap_{m=1}^\infty A_m$.

Theorem 4.7. *A game with symmetry group G has an order ∞ G-exchangeable equilibrium.*

Proof. Endow $\Delta(\Delta_G^\Pi(\Gamma))$ with the topology of weak convergence, which makes it into a compact metric space since $\Delta_G^\Pi(\Gamma)$ is (11.5.4 and 11.5.5 in [4]). For any $(s^1,\ldots,s^m) \in C^m$ the map $\Delta_G^\Pi(\Gamma) \to \mathbb{R}$ given by $R \mapsto R(s^1)\cdots R(s^m)$ is a polynomial, hence continuous, so the map μ_m is continuous by definition of weak convergence.

Each A_m is convex and nonempty by definition of order m exchangeable equilibria and Theorem 4.5. Each set $\mathrm{XE}_G^m(\Gamma)$ is closed so the A_m are compact. They are also nested $A_1 \supseteq A_2 \supseteq A_3 \supseteq \ldots$ as per the discussion after Definition 4.3, so they have nonempty, compact, convex intersection. $\square$

5 Nash's Theorem

5.1 The player-transitive case

Theorem 5.1. *If G acts player transitively on Γ, then $\mathrm{XE}_G^\infty = \Delta(\mathrm{NE}_G(\Gamma))$.*

Proof. If $\sigma \in \mathrm{NE}_G(\Gamma)$ then $\sigma^m \in \mathrm{XE}_G^m(\Gamma)$, so $\delta_\sigma \in A_m$ for all m and $\delta_\sigma \in \mathrm{XE}_G^\infty(\Gamma)$. But $\mathrm{XE}_G^\infty(\Gamma)$ is convex and weakly closed, so $\Delta(\mathrm{NE}_G(\Gamma)) = \overline{\mathrm{conv}}\{\delta_\sigma \mid \sigma \in \mathrm{NE}_G(\Gamma)\} \subseteq \mathrm{XE}_G^\infty(\Gamma)$.

For the converse let $\pi \in \mathrm{XE}_G^\infty(\Gamma)$ and define R and X_i^j as above. By definition of an order m G-exchangeable equilibrium, for any $1 \le j \le m$ the strategy X_i^j is a best response to the random conditional distribution $P_i^m = \mathbb{P}(X_{-i}^1 \mid X_i^1,\ldots,X_i^m)$ almost surely. Let Y_i^m be the random variable taking values in $\Delta(C_i)$ which is the empirical distribution of $X_i^1,\ldots,X_i^m$. Then Y_i^j is a best response to P_i^k whenever $j \le k$.

Let Σ_i be the completion of the σ-algebra generated by $X_i^1, X_i^2, \ldots$ and define $P_i^\infty = \mathbb{P}(X_{-i}^1 \mid \Sigma_i)$. Then $P_i^k \to P_i^\infty$ almost surely as k goes to infinity (Theorem 10.5.1 in [4]). Therefore Y_i^j is a best response to P_i^∞ for all j. By the strong law of large numbers, Y_i^j converges almost surely to R_i as j goes to infinity, so R_i is a best response to P_i^∞ and is measurable with respect to Σ_i.

Since G acts player transitively, for any player j we have $R_j = R_i \cdot g$ for some $g \in G$, hence R_j is measurable with respect to Σ_i and so is R. But the X_i^j are independent conditioned on R, so $P_i^\infty = \mathbb{E}(\mathbb{P}(X_{-i}^1 \mid R) \mid \Sigma_i) = \mathbb{P}(X_{-i}^1 \mid R) = R_{-i}$. This shows that R_i is a best response to R_{-i} almost surely for all i, so $R \in \mathrm{NE}_G(\Gamma)$ almost surely. In particular, $\pi \in \Delta(\mathrm{NE}_G(\Gamma))$. $\square$

If G is the trivial group one can show that $\mu_1(\mathrm{XE}_G^\infty(\Gamma)) = \mathrm{CE}(\Gamma)$ and $\mu_1(\Delta(\mathrm{NE}_G(\Gamma))) = \mathrm{conv}(\mathrm{NE}(\Gamma))$. These sets are different for some games (e.g., chicken), so the above theorem can fail without the player-transitivity assumption.

Nash's Theorem (player-transitive case)**.** *A game with player-transitive symmetry group has a G-invariant Nash equilibrium.*

Proof. Combine Theorems 4.7 and 5.1, noting that $\Delta(\emptyset) = \emptyset$. $\square$

5.2 Arbitrary symmetry groups

In this section we show how to embed an arbitrary game Γ with symmetry group G in a game Γ^{Sym} with a player-transitive symmetry group, preserving the existence of G-invariant Nash equilibria. This allows us to drop the player-transitivity assumption from the previous section, proving Nash's Theorem in full generality.

There are a variety of ways to symmetrize games. The one we have chosen is a natural n-player generalization of von Neumann's tensor sum symmetrization discussed in [6]. The idea is that each of the n players in Γ^{Sym} plays all the roles of the players in Γ simultaneously. The players in Γ^{Sym} play $n!$ copies of Γ, one for each assignment of players in Γ^{Sym} to roles in Γ. A player's utility in Γ^{Sym} is the sum of his utilities over the copies.

Definition 5.2. Given an n-player game Γ with strategy sets C_i and utilities u_i we define its **symmetrization** Γ^{Sym} to be the n-player game with strategy sets $C_i^{\mathrm{Sym}} = C$ and utilities

$$u_i^{\mathrm{Sym}}(s) = \sum_{\tau \in S_n} u_{\tau(i)}\left(\delta(\tau \star s)\right),$$

where $s = (s^1, \ldots, s^n) \in C^{\mathrm{Sym}} = C^n$, $\star : S_n \times C^{\mathrm{Sym}} \to C^{\mathrm{Sym}}$ is defined by $(\tau \star s)^k = s^{\tau^{-1}(k)}$, and $\delta : C^{\mathrm{Sym}} \to C$ is defined by $[\delta(s)]_k = s_k^k$.

We now show that Γ^{Sym} is a game with player-transitive symmetry group. We will use $\star$ to denote the action on Γ^{Sym} to distinguish it from the action $\cdot$ on Γ.

Proposition 5.3. *If Γ is a game with symmetry group G then Γ^{Sym} is a game with player-transitive symmetry group $G \times S_n$, where $\sigma \in S_n$ acts by $\star$ as defined above and $g \in G$ acts by*

$$g \star (s^1, \ldots, s^n) \mapsto (g \cdot s^1, \ldots, g \cdot s^n).$$

Proof. First note that $\star$ defines an action of G on C^{Sym}. For $\sigma, \tau \in S_n$ we have

$$(\tau \star (\sigma \star s))^k = (\sigma \star s)^{\tau^{-1}(k)} = s^{\sigma^{-1}(\tau^{-1}(k))} = s^{(\tau\sigma)^{-1}(k)} = ((\tau\sigma) \star s)^k,$$

so $\star$ is also an action of S_n on C^{Sym}. These actions commute, so together they define an action $\star$ of $G \times S_n$ on C^{Sym}. Note that the induced actions on players are $\sigma \star i = \sigma(i)$ and $g \star i = i$.

To show that this is an action of $G \times S_n$ on Γ^{Sym} it suffices to show that the utilities of Γ^{Sym} are invariant under the action of any $\sigma \in S_n$ and any $g \in G$. To see the former, let $\sigma \in S_n$. Then we have

$$\begin{aligned} u^{\mathrm{Sym}}_{\sigma\star i}(\sigma \star s) &= \sum_{\tau \in S_n} u_{\tau(\sigma(i))}\left(\delta(\tau \star (\sigma \star s))\right) = \sum_{\tau \in S_n} u_{(\tau\sigma)(i)}\left(\delta((\tau\sigma) \star s\right) = \sum_{\tau \in S_n} u_{\tau(i)}\left(\delta(\tau \star s)\right) \\ &= u^{\mathrm{Sym}}_i(s), \end{aligned}$$

where we have used in the penultimate equation the fact that S_n is a group, so the map $\tau \mapsto \tau\sigma$ is a bijection. To see the latter, let $g \in G$ and let $\gamma \in S_n$ be the permutation induced by g on the set of players in Γ. Then we have $\delta(g \star s) = g \cdot \delta(\gamma^{-1} \star s)$, so

$$\begin{aligned} u^{\mathrm{Sym}}_{g\star i}(s) &= \sum_{\tau \in S_n} u_{\tau(i)}\left(\delta(\tau \star (g \star s))\right) = \sum_{\tau \in S_n} u_{\tau(i)}\left(\delta(g \star (\tau \star s))\right) = \sum_{\tau \in S_n} u_{\tau(i)}\left(g \cdot \delta(\gamma^{-1} \star (\tau \star s))\right) \\ &= \sum_{\tau \in S_n} u_{(\gamma^{-1}\tau)(i)}\left(\delta((\gamma^{-1}\tau) \star s)\right) = \sum_{\tau \in S_n} u_{\tau(i)}\left(\delta(\tau \star s)\right) = u^{\mathrm{Sym}}_i(s), \end{aligned}$$

where the third-to-last equation follows because g is a symmetry of Γ. Clearly S_n acts transitively on the set of players. $\square$

Nash's Theorem. *A game with symmetry group G has a G-invariant Nash equilibrium.*

Proof. Let Γ be a game with symmetry group G. Then Γ^{Sym} is a game with player-transitive symmetry group $G \times S_n$ by Proposition 5.3, so it has a $(G \times S_n)$-symmetric Nash equilibrium by the player-transitive version of Nash's Theorem. By definition of the action of $G \times S_n$ on Γ^{Sym}, this Nash equilibrium is of the form $(\rho, \dots, \rho)$, with $\rho \in \Delta_G(\Gamma)$. Notice that for each player i, each utility $u^{\mathrm{Sym}}_k(s^1, \dots, s^n)$ is a sum of functions which only depend on s^i_j for a single value of j. Thus ρ is strategically equivalent to the product of its marginals $\rho_1 \times \cdots \times \rho_n \in \Delta^{\Pi}_G(\Gamma)$. Therefore we can take the Nash equilibrium $(\rho, \dots, \rho)$ to be such that $\rho \in \Delta^{\Pi}_G(\Gamma)$ by Proposition 2.7.

It remains to verify that $\rho \in \mathrm{NE}_G(\Gamma)$. For any $s^i \in C$ we can compute

$$\begin{aligned} u^{\mathrm{Sym}}_i(\rho, \dots, \rho, s^i, \rho, \dots, \rho) &= \sum_{\tau \in S_n} u_{\tau(i)}\left(\rho_1, \dots, \rho_{\tau(i)-1}, s^i_{\tau(i)}, \rho_{\tau(i)+1}, \dots, \rho_n\right) \\ &= (n-1)! \sum_{j=1}^{n} u_j(\rho_1, \dots, \rho_{j-1}, s^i_j, \rho_{j+1}, \dots, \rho_n). \end{aligned}$$

For each value of j we can vary the s^i_j component of s^i independently and it is a best response for player i to play ρ in Γ^{Sym} if the rest of the players play ρ, so we must have

$$u_j(\rho_1, \ldots, \rho_{j-1}, s_j, \rho_{j+1}, \ldots, \rho_n) \leq u_j(\rho)$$

for all players j and all $s_j \in C_j$, i.e., $\rho \in \mathrm{NE}_G(\Gamma)$. □

6 Conclusion

We have shown that by studying group actions on games and introducing the notion of exchangeable equilibrium, we can extend Hart and Schmeidler's methods and prove Nash's Theorem. To the authors' knowledge, this is the first proof of this theorem which uses convexity-based methods (i.e., the minimax theorem). Previous proofs use path-following arguments or fixed-point theorems, which are essentially equivalent to path-following arguments by Sperner's Lemma.

This new proof invites new approaches for computing or approximating Nash equilibria. One can rewrite the existence proof above for (order m) exchangeable equilibria in terms of linear programs and separation arguments instead of the Minimax Theorem and apply the ellipsoid algorithm, just as Papadimitriou has done for Hart and Schmeidler's proof of the existence of correlated equilibria [11]. This shows that exchangeable equilibria can be computed in polynomial time, at least under some assumptions on the parameters. For example, order m exchangeable equilibria of symmetric bimatrix games can be computed in polynomial time for fixed m.

We have seen that in the player-transitive case order m exchangeable equilibria converge to convex combinations of Nash equilibria as m goes to infinity. There are a variety of ways one could imagine "rounding" exchangeable equilibria to try to produce approximate Nash equilibria. We leave the analysis of such procedures, along with the question of which assumptions on G allow computation of exchangeable equilibria in polynomial time, for future work.

The power of these methods suggests that exchangeable equilibria should not merely be viewed as a step on the way to Nash equilibria. Rather, they deserve further study in their own right. We consider several interpretations of exchangeable equilibria and the applications they suggest in [13].

References

[1] A. Berman and N. Shaked-Monderer. *Completely positive matrices.* World Scientific Publishing Co. Pte. Ltd., River Edge, NJ, 2003.

[2] D. P. Bertsekas, A. Nedić, and A. E. Ozdaglar. *Convex Analysis and Optimization.* Athena Scientific, Belmont, MA, 2003.

[3] P. H. Diananda. On non-negative forms in real variables some or all of which are non-negative. *Mathematical Proceedings of the Cambridge Philosophical Society*, 58(1):17 – 25, January 1962.

[4] R. M. Dudley. *Real Analysis and Probability.* Cambridge University Press, New York, 2002.

[5] D. Gale. The game of Hex and the Brouwer fixed-point theorem. *The American Mathematical Monthly*, 86(10), December 1979.

[6] D. Gale, H. W. Kuhn, and A. W. Tucker. On symmetric games. In *Contributions to the Theory of Games*, volume 1, pages 81 – 87. Princeton University Press, Princeton, NJ, 1950.

[7] S. Hart and D. Schmeidler. Existence of correlated equilibria. *Mathematics of Operations Research*, 14(1), February 1989.

[8] A. Hatcher. *Algebraic Topology.* Cambridge University Press, New York, NY, 2002.

[9] C. E. Lemke and J. T. Howson, Jr. Equilibrium points in bimatrix games. *SIAM Journal on Applied Math*, 12:413 – 423, 1964.

[10] J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286 – 295, September 1951.

[11] C. H. Papadimitriou. Computing correlated equilibria in multi-player games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, New York, NY, 2005. ACM Press.

[12] E. Sperner. Neuer Beweis für die Invarianz der Dimensionszahl und des Gebietes. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 6(1):265 – 272, December 1928.

[13] N. D. Stein, A. Ozdaglar, and P. A. Parrilo. Exchangeable equilibria, part I: the symmetric bimatrix case. In preparation.