

Exchangeable Equilibria, Part I: Symmetric Bimatrix Games

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

September 5, 2013

Abstract

We introduce the notion of exchangeable equilibria of a symmetric bimatrix game, defined as those correlated equilibria in which players' strategy choices are conditionally independently and identically distributed given some hidden variable. The set of these equilibria is convex and lies between the symmetric Nash equilibria and the symmetric correlated equilibria. The notion of exchangeable equilibria admits several natural game-theoretic interpretations which emphasize their high degree of symmetry and suggest why they may be focal.

Keywords: Correlated equilibrium, exchangeability, equilibrium concepts, symmetric games. JEL classification codes: C72, C63, C61.

1 Introduction

The contribution of this paper is to introduce exchangeable equilibria, a solution concept for symmetric games intermediate between symmetric Nash and correlated equilibria. For simplicity we restrict to symmetric bimatrix (two-player) games throughout. The more technical multiplayer case will be treated in Part II of this paper, with computational issues covered in future work.¹

1.1 Motivation

A basic tenet of decision theory is that preferences are specific to each individual, and so interpersonal comparison of utilities is meaningless. In practice this belief is often dropped in favor of other simplifying assumptions which can be justified within the setting at hand. One important example of this phenomenon is zero-sum bimatrix games. The zero-sum condition

*The first author is at Analog Devices | Lyric Labs, reachable at noah.stein@analog.com. The remaining authors are at the Laboratory for Information and Decision Systems of the Massachusetts Institute of Technology, reachable at asuman@mit.edu and parrilo@mit.edu.

¹For more details on these topics, see the first author's doctoral thesis [26].

says not only that the players' preferences over outcomes are opposite, but also that their preferences over lotteries over outcomes are opposite. This is a very strong condition indeed, and while it may never be true exactly, it is often a reasonable approximation to model situations in which there seems to be no opportunity for cooperation.

Similarly, symmetric bimatrix games are an idealization in which we take the roles and preferences of the players to be identical. When such an approximation is natural, it seems reasonable to take this idea to its logical conclusion and treat the players as independent instances of an identical decision-making agent, or clones, as it were. Doing so, we should expect both players would make the same choice when confronted with the same situation. They interpret their observations of the environment in the same way. If the setup is such that we have reason to believe that the players will choose their strategies statistically independently, then our solution concept of choice should be symmetric Nash equilibrium. Here we explicitly use the assumption that the players are identical; in particular they have not been provided with a way to break the symmetry so as to choose an asymmetric Nash equilibrium.²

This raises a natural question: what is the appropriate solution concept to use if statistical independence of actions is not a reasonable assumption and some correlation may be expected? For a general game the answer would be correlated equilibrium³, so the obvious answer in the symmetric setting is symmetric correlated equilibrium. That is to say, we could expect any correlated equilibrium which is invariant (as a probability distribution) under interchange of the players.

This obvious answer is correct insofar as the symmetry of the situation rules out any correlated equilibria which are not symmetric. It is our goal to argue that symmetry can also rule out some of the symmetric correlated equilibria!

To see this, we begin with the tautological statement that each player bases his action on the state of the world. Perhaps he plays the same action at all states, perhaps not. Insofar as he does not have full knowledge of the state of the world his action must be based on some noisy observations of the world. Since the players are identical they both make the same measurements and interpret them in the same way, but the outcomes of these measurements may differ. In this way and only this way may the players' actions differ.

For example, the players may wish to base their actions on tomorrow's weather. This is an aspect of the underlying state of the world, albeit one which may not be known precisely to the players (it is a *hidden variable*), so they can only condition their actions on their estimates of the weather. Perfect correlation may occur if these estimates are obtained from

²General bimatrix games can be symmetrized by considering two copies of the game played simultaneously with each player taking opposite roles in the two copies. In this way results from the symmetric theory can be applied to asymmetric games. This idea of players "putting themselves in each other's shoes" incorporates a notion of fairness into the analysis of asymmetric games. Rawls explores the philosophical implications of taking this idea as the basis for a theory of justice [23], though in a way which is at odds with some fundamental game-theoretic principles [14, 6].

³See Aumann's argument in [3]. The key assumption forming the basis of that argument is that the players have common prior beliefs. While this is debatable in some contexts, it is certainly true in the symmetric setting considered here. Put another way, the philosophical or epistemic argument for correlated equilibrium is even stronger in symmetric games than in general games.

a common weather report. Weaker correlation may arise if both players make their own forecasts in terms of independent measurements of atmospheric data.

We define an *exchangeable equilibrium* of a symmetric bimatrix game to be a correlated equilibrium distribution in which the players' actions are independent and identically distributed (i.i.d.) conditioned on some hidden variable (Definition 3.1). These are exactly the correlated equilibria which could conceivably be implemented by players making independent environmental measurements, then interpreting and acting on this information in the same way.

Over the course of the paper we motivate this definition from several viewpoints: game-theoretic, probabilistic, algebraic / geometric, and (while leaving a detailed discussion for future work) computational. We begin by unpacking the definition.

Any probability distribution of two players' actions X_1, X_2 can be written as a square matrix A , with $A_{ij} := \mathbb{P}(X_1 = i \text{ and } X_2 = j)$. If the actions were i.i.d. their joint distribution could be written in matrix form as an outer product xx^T for some probability vector x . When the actions are merely conditionally i.i.d. the distribution is in fact a convex combination⁴ $\sum_i \lambda_i x_i x_i^T$ of terms of this form ($\lambda_i \geq 0$, $\sum_i \lambda_i = 1$).

We might say that any matrix of probabilities which is not symmetric explicitly breaks symmetry, whereas a symmetric matrix which is not conditionally i.i.d. implicitly breaks symmetry. We focus on the exchangeable equilibria, those correlated equilibria which do not break symmetry at all.

To illustrate this symmetry breaking, consider an $m \times m$ matrix A (symmetric or not) of probabilities specifying the joint distribution of two random variables and let e_i denote the i^{th} unit column vector. We can write $A = \sum_{i,j} A_{ij} e_i e_j^T$. The matrix $e_i e_j^T$ is an independent distribution corresponding to two constants. Therefore we can always view a sample taken according to a probability matrix A as the result of two "measurements" which are conditionally independent given some hidden variable.

Here we have taken the hidden variable to be the result of the two measurements. Symmetry is broken in the sense that players' actions are not identically distributed given the hidden variable: if the value is (i, j) the first player deterministically plays i and the second plays j . The players do not map their measurements of the hidden variable to actions in the same way. Looked at differently, they are not measuring the same aspects of the hidden variable to determine their choices of action.

In general there are many such decompositions $A = \sum_i \lambda_i x_i y_i^T$. We only wish to consider matrices A which admit a decomposition $A = \sum_i \lambda_i x_i x_i^T$ respecting the symmetry of the problem. Clearly lack of symmetry of A is an obstruction to the existence of such a decomposition, but it is not the only obstruction: see Section 1.2 for an example.

Geometrically, the set of exchangeable equilibria arises naturally when considering convex relaxations of the set of Nash equilibria. Let CE_{Sym} denote the set of symmetric correlated equilibria of a symmetric game, and X denote the set of symmetric, rank 1, elementwise

⁴The end result is the same whether we consider finite sums or infinite "convex combinations" weighted by probability distributions; see Section 2.1.

$(u_{\text{row}}, u_{\text{col}})$	Wimpy	Macho
Wimpy	(4, 4)	(1, 5)
Macho	(5, 1)	(0, 0)

Table 1: The game of Chicken. Two players ride their bikes towards each other. At the last moment, each can choose to either be Wimpy and veer off to the side, or be Macho and remain on the collision course. One player chooses a row, the other a column, and the corresponding cell of the table shows their respective utilities.

nonnegative matrices, i.e., matrices of the form xx^T for $x \in \mathbb{R}_{\geq 0}^{m \times 1}$. Then we have

$$\text{CE}_{\text{Sym}}^{\text{Nash}} \cap X \subseteq \text{conv}(\text{CE}_{\text{Sym}}^{\text{convex hull of Nash}} \cap X) \subseteq \text{CE}_{\text{Sym}}^{\text{exchangeable}} \cap \text{conv}(X) \subseteq \text{CE}_{\text{Sym}}^{\text{correlated}},$$

where each type of symmetric equilibrium is defined by the set written below it.

We will discuss three further interpretations of exchangeable equilibria in this paper, in addition to the above motivations via hidden variables or geometry. First, we show that these are the equilibria which can arise as a Bayesian observer’s predictions when players are chosen from a large pool of agents whom the observer views as interchangeable. Second, we view them as symmetric correlated equilibria of games with many players which can be broken up into identical symmetric pairwise interactions between these players. In particular the exchangeable equilibria are those symmetric correlated equilibria which do not depend on the exact number of players. Third, the exchangeable equilibria are those which can be implemented with a particular kind of correlation scheme.

1.2 Preview

We illustrate the theory of exchangeable equilibria using the game Chicken (Table 1). This game has two asymmetric pure Nash equilibria and one symmetric mixed Nash equilibrium. We write the corresponding joint distributions as matrices. Each entry is the probability of choosing the corresponding strategy profile.

$$\text{NE} = \left\{ \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} 1/4 & 1/4 \\ 1/4 & 1/4 \end{bmatrix} \right\}.$$

In the same notation a probability matrix $\begin{bmatrix} p & q \\ q & r \end{bmatrix}$ with $p, q, r \geq 0$ and $p + 2q + r = 1$ is a symmetric correlated equilibrium if and only if $q \geq p, r$. These inequalities define the polytope

$$\text{CE}_{\text{Sym}} = \text{conv} \left\{ \begin{bmatrix} 0 & 1/2 \\ 1/2 & 0 \end{bmatrix}, \begin{bmatrix} 1/3 & 1/3 \\ 1/3 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 1/3 \\ 1/3 & 1/3 \end{bmatrix}, \begin{bmatrix} 1/4 & 1/4 \\ 1/4 & 1/4 \end{bmatrix} \right\}. \quad (1)$$

Consider the symmetric correlated equilibrium

$$B := \begin{bmatrix} 0 & 1/2 \\ 1/2 & 0 \end{bmatrix}.$$

Is it an exchangeable equilibrium? That is to ask, is it conditionally i.i.d.? Equivalently, is it of the form $\sum \lambda_i x_i x_i^T$?

Any matrix of this form would be positive semidefinite, but $\det B = -\frac{1}{4} < 0$, so B is not positive semidefinite. It is a symmetric correlated equilibrium which is not exchangeable. In any game positive semidefiniteness is necessary for a correlated equilibrium to be exchangeable. Known results (Theorem 2.9) show it is sufficient in games with at most four strategies.

The same test rules out the second and third extreme points listed in (1) as not exchangeable. In fact for Chicken the condition that the determinant be nonnegative ($q^2 \leq pr$) and the symmetric correlated equilibrium conditions have only one common solution:

$$\left\{ \begin{bmatrix} 1/4 & 1/4 \\ 1/4 & 1/4 \end{bmatrix} \right\} = \text{NE}_{\text{Sym}} = \text{XE}_{\text{Sym}} \subsetneq \text{CE}_{\text{Sym}}.$$

This example illustrates several general properties of the set of exchangeable equilibria that we will establish in this paper: $\text{conv}(\text{NE}_{\text{Sym}}) \subseteq \text{XE}_{\text{Sym}}$ (Proposition 3.3) with equality in games with two strategies (Theorem 3.6) and $\text{XE}_{\text{Sym}} \subsetneq \text{CE}_{\text{Sym}}$ (Proposition 3.3) with strict containment whenever asymmetric Nash equilibria exist (Theorem 3.5).

1.3 Previous work

As we will see, the idea of exchangeable equilibrium arises naturally when one combines the notion of exchangeable random variables with the correlated equilibrium solution concept. Both of these are large areas of study in their own right and we make no attempt at an exhaustive survey of the literature of one or the other. However, several papers either have led toward the idea of exchangeable equilibrium in some way or seem to be particularly related to what is accomplished by exchangeable equilibria and we review those here.

One such example is Hart and Schmeidler's minimax proof of the existence of correlated equilibria [15] or the similar proof by Nau and McCardle [21]. These arguments share the somewhat mysterious quality that they work with non-equilibrium product distributions at an intermediate stage, but in the end the product structure vanishes entirely and the result is a correlated equilibrium with no obvious extra structure. Considering what extra structure might be obtained from such an argument, in particular if we require the game to be symmetric and try to make the argument respect the symmetry as much as possible, is one route to our notion of exchangeable equilibrium. Two existence proofs along these lines will be presented in Part II of this paper.

Another example is Brandenburger's survey of epistemic game theory [7]. He mentions that the standard assumption of probabilistic independence of different players' strategies in noncooperative game theory is somewhat suspect and perhaps unnatural. In a footnote he offers the concept of exchangeability as an example of a more natural way to capture our ignorance of the distinctions between random variables. However, this is not explored further.

The notion of exchangeable equilibrium bears at least an outward resemblance to the work of Hillas, Kohlberg, and Pratt on an outside observer's assessment of the outcome

of a game [16]. They consider an observer watching a given n -player game being played repeatedly. Each time the game is played by a new set of players disjoint from the set of all previous players, so the observer views these interactions as exchangeable: he has no preconceived notions about how the order of these repeated plays relates to their outcomes. In other words, his prior distribution over actions is not affected by swapping all players in copy j of the game with the corresponding players in copy k for any j, k . The players do not have access to the history of play, though the observer does.

Hillas, Kohlberg, and Pratt characterize correlated equilibria of the original game by considering whether at each play of the game the observer can offer advice based on his past observations enabling some player to increase his expected payoff. They argue that the players should have a better understanding of the game than the observer, and prove that the limiting distribution of play is a correlated equilibrium if and only if the observer cannot offer such helpful advice.

If the observer also does not have any preconceived notions about the connections between the players in each instance of the game (i.e., his prior distribution on actions is not affected by swapping the copy of player i in instance j of the game with the copy in instance k for any i, j, k), then Hillas, Kohlberg, and Pratt show that the long-run distribution of play will be a Nash equilibrium. The paper [16] does not consider solution concepts between Nash and correlated equilibrium. However, its use of exchangeability suggests an ideological kinship with the present work. It seems likely that the concept of exchangeable equilibrium could be reinterpreted in that setup, though we do not do so here.

One interpretation of exchangeable equilibria is that they are those correlated equilibria which extend to symmetric correlated equilibria of games with an arbitrarily large number of identical interactions (Section 4.3). That is to say, they are robust to the number of players: the players could not profitably deviate even if they knew the number of players. This is in contrast to the paper of Myerson on games with many players in which the number of players is modeled probabilistically [19]. That work is in a Bayesian game setting and so not directly comparable to ours, but the distinction between which equilibria are sensitive to a given probabilistic model and which equilibria are not is worth making. In a given situation, one assumption or the other may be more natural.

Another interpretation of exchangeable equilibria is that they are those correlated equilibria in which the correlating device takes a specific form: the players choose strategies i.i.d. conditioned on some hidden variable (Section 4.1). In general it is an interesting problem to study which equilibria can arise from a particular type of correlating device. Of course the most well-studied class of correlating device makes players choose their strategies independently and corresponds to the mixed Nash equilibria. But there is also Sorin's notion of "distribution equilibria," which are correlated equilibria in which each player gets the same payoff conditional on all outcomes and which are motivated by arguments from evolutionary game theory [25]. Another good example of this type is Du's classification of the correlated equilibria which can arise when the players correlate on their hierarchy of beliefs about the play of the game [11].

1.4 Outline

We begin with Section 2, fixing notation and reviewing background material on exchangeability, conditionally i.i.d. random variables, games, and equilibria. With all this in place in Section 3 we introduce exchangeable equilibria and establish their fundamental properties. The two sections which follow are Section 4 on interpretations and Section 5 on examples. These are largely independent of each other and can be read in either order. We close with Section 6 on directions for future work. This final section also outlines the content of Parts II and III of this paper.

2 Background and notation

2.1 Probability distributions, exchangeability, and conditionally i.i.d. random variables

We write $\Delta(T)$ to denote the set of (Borel) probability distributions on T . The mass of a singleton $t \in T$ under a probability distribution $\pi \in \Delta(T)$ will be denoted $\pi(t) := \pi(\{t\})$ for simplicity. For finite T we will view $\Delta(T)$ as a set of column vectors and $\Delta(T^2) := \Delta(T \times T)$ as a set of matrices. All random variables will take values in a set T with $2 \leq |T| < \infty$ to avoid trivialities on one end and unnecessary complications on the other. We write T^n and T^∞ for the products of n and countably many copies of T , respectively.

Throughout the paper we will deal with collections of random variables which are *exchangeable*, meaning they are in some sense indistinguishable. Such is the case when the measurements they represent are unlabeled or labeled in a way which does not relate to their outcome. We will usually assume these are given to us in some arbitrary order. The defining property of exchangeability, then, is that the joint distribution does not depend on this order.

Definition 2.1. For finite n , a probability distribution in $\Delta(T^n)$ is **n -exchangeable** if it is invariant under permutations of the indices, or in other words, if it represents a sequence of random variables X_1, \dots, X_n which is equal in distribution to $X_{\sigma(1)}, \dots, X_{\sigma(n)}$ for all permutations σ on $\{1, \dots, n\}$. The set of n -exchangeable distributions is denoted $\Delta_{\text{Sym}}(T^n)$.

A typical example of an n -exchangeable sequence arises by sampling without replacement. If 4 white balls and 6 black balls are placed in a bag, shaken, and then removed one by one, the distribution of the resulting sequence of colors is 10-exchangeable.

Definition 2.2. A probability distribution in $\Delta(T^\infty)$ representing a sequence X_1, X_2, \dots is said to be **(∞) -exchangeable** if it is invariant under permutations of finitely many indices⁵, or in other words if the distribution of X_1, \dots, X_n is n -exchangeable for all $n \in \mathbb{N}$. The set of exchangeable distributions is denoted $\Delta_{\text{Sym}}(T^\infty)$. We call a finite or infinite sequence of

⁵This finiteness condition is natural in that it allows us to identify $\Delta_{\text{Sym}}(T^\infty)$ with the inverse limit $\varprojlim \Delta_{\text{Sym}}(T^n)$ using the Kolmogorov Consistency Theorem (12.1.2 in [12], for example).

random variables n -**exchangeable** or (∞) -**exchangeable**, respectively, if its distribution is.

One example of exchangeability is a sequence of i.i.d. fair coin flips, or indeed any i.i.d. sequence. However, this does not exhaust the possibilities. Another example is the case in which $X_i = X_1$ almost surely for all i . This is exchangeable regardless of the distribution on X_1 , but is only independent in the trivial case when X_1 is deterministic.

The definition of exchangeability amounts to a collection of linear equations a given probability distribution may or may not satisfy. This means the set of exchangeable distributions is convex: an arbitrary mixture of exchangeable distributions is also exchangeable. Therefore a sequence of random variables which are i.i.d. conditioned on some other random variable is exchangeable. In the case of an infinite sequence of exchangeable random variables, the converse is also true:

Definition 2.3. A (finite or infinite) sequence of random variables is called **conditionally i.i.d.** if it is i.i.d. conditioned on some auxiliary random variable.

Theorem 2.4 (de Finetti). *An infinite sequence X_1, X_2, \dots of random variables is exchangeable if and only if it is conditionally i.i.d.*

For a proof see e.g. [24]. The upshot is that we can think of any exchangeable distribution as arising from a two-stage process⁶. First, a random element $\Theta \in \Delta(T)$ is secretly selected according to some distribution in $\Delta(\Delta(T))$. Second, the sequence of random variables X_1, X_2, \dots is constructed to be i.i.d. with distribution Θ .

The following two results summarize some facts about conditionally i.i.d. pairs X_1, X_2 and the set of all distributions thereof. Proposition 2.6 shows that such pairs take a simple form and will often be invoked implicitly.

Proposition 2.5. *Consider the subset of $\Delta(T^2)$ corresponding to distributions of pairs X_1, X_2 which are conditionally i.i.d. given a parameter taking only finitely many values, i.e., those which can be written in matrix form as*

$$\sum_{i=1}^k \lambda_i w_i w_i^T \text{ for some } k, \text{ with all } \lambda_i > 0 \text{ and } \sum_{i=1}^k \lambda_i = 1. \quad (2)$$

This set is convex, compact, and semialgebraic (described by finitely many polynomial inequalities).

Proof. Call the set in question Z . Let X be the set of i.i.d. distributions in $\Delta(T^2)$. Then X is compact and Z is the convex hull of X , so compact and convex[5]. The set X is semialgebraic. By Carathéodory's theorem any element of Z can be written as a convex combination of a bounded number of elements of X , so Z can be described by a first-order formula over the real-closed field \mathbb{R} . Quantifier elimination then shows that Z is semialgebraic [4]. \square

⁶The classic example of an exchangeable distribution which does not obviously arise in this way is the Pólya urn model [22].

Proposition 2.6. *If a pair X_1, X_2 is conditionally i.i.d. then we may assume without loss of generality that the associated parameter takes its values in a finite set. Equivalently, the joint distribution of X_1, X_2 is of the form (2). That is to say, the set described in Proposition 2.5 is the set of all conditionally i.i.d. distributions.*

Proof. Suppose the pair X_1, X_2 is i.i.d. conditioned on some parameter Θ with (not necessarily finitely-supported) distribution $\lambda \in \Delta(\Delta(T))$. The unconditional joint distribution of X_1 and X_2 is $W := \int ww^T d\lambda(w)$. By definition of the Lebesgue integral, W is the limit of conditionally i.i.d. distributions each of whose parameters takes values in some finite set. The set thereof is compact (Proposition 2.5), so it is closed and contains W . \square

How can one be sure that two random variables are not conditionally i.i.d.? In other words, given a probability distribution as a matrix, how can one certify that a decomposition of the form (2) does not exist? A weak but easily applied condition is:

Proposition 2.7. *If X_1, X_2 are conditionally i.i.d. random variables taking values in the finite set T and $\mathbb{P}(X_1 = X_2 = t) = 0$ then $\mathbb{P}(X_1 = t, X_2 = \tilde{t}) = 0$ for all $\tilde{t} \in T$.*

Proof. Let Θ be the hidden parameter conditioned on which the X_i are i.i.d. By assumption Θ satisfies $\mathbb{P}(X_1 = t \mid \Theta) = 0$ almost surely, so $\mathbb{P}(X_1 = t) = \int \mathbb{P}(X_1 = t \mid \Theta) d\mathbb{P}(\Theta) = 0$. \square

Using this condition one can immediately construct symmetric, elementwise nonnegative matrices which are not the joint distribution of any conditionally i.i.d. pair, such as $\begin{bmatrix} 0 & 1/2 \\ 1/2 & 0 \end{bmatrix}$. A more broadly applicable necessary condition is double nonnegativity.

Definition 2.8. A matrix is **doubly nonnegative** if it is symmetric, elementwise nonnegative, and positive semidefinite. The set of doubly nonnegative $m \times m$ matrices is denoted DNN_m .

The advantage of this definition is that double nonnegativity of a matrix is easy to check, and the definitions immediately yield the forward direction in:

Theorem 2.9. *If an $m \times m$ matrix is the joint distribution of a conditionally i.i.d. pair then it is doubly nonnegative and its entries sum to one. The converse holds if and only if $m \leq 4$.*

Maxfield and Minc's proof of the converse (which is not obvious from the definitions) and a counterexample in the $m = 5$ case due to Hall can be found in [18]. Diananda proved an equivalent result about the duals of these cones [10]. For our purposes, the utility of this result is that it allows us to decide quickly whether 4×4 and smaller matrices and the joint distribution of conditionally i.i.d. pairs.

2.2 Games and equilibria

Definition 2.10. A **(finite) game**, denoted by Γ throughout, consists of $n \geq 2$ players, each equipped with a finite set C_i of two or more **(pure) strategies** and a utility function $u_i : C \rightarrow \mathbb{R}$, which assigns a real number to each **strategy profile** in $C := \prod_i C_i$.

A two-player game is also called a **bimatrix game**, because the utilities can be specified by two matrices $A, B \in \mathbb{R}^C$: $A_{ij} = u_1(i, j)$ and $B_{ij} = u_2(i, j)$. A bimatrix game is **symmetric** if $B = A^T$, or in other words if $C_1 = C_2$ and $u_1(s_1, s_2) = u_2(s_2, s_1)$ for all $(s_1, s_2) \in C$. For a symmetric bimatrix game we let $m = |C_1|$. When it is notationally convenient to do so, we will assume without loss of generality that $C_1 = \{1, \dots, m\}$.

We make the standard assumptions that the number of players, strategy sets, and utility functions are common knowledge: everyone knows them, everyone knows that everyone knows them, and so on (see e.g. [13]). For symmetric bimatrix games we make the additional framing assumptions that each of a player's strategies has a distinct label, the same set of labels is used by both players, and this shared labeling is common knowledge. In other words there is a commonly known bijection between the strategy sets of the players. We will always take these sets to be equal and the distinguished bijection to be the identity. Changing this bijection has a profound effect on the game as we see in Examples 5.1 and 5.2. For further discussion of symmetry and framing see [1, 8, 27].

Definition 2.11. 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** of the game Γ will be denoted⁷ $\Delta^\Pi(\Gamma) := \prod_i \Delta(C_i)$. For a symmetric bimatrix game Γ , the set of symmetric product distributions will be denoted

$$\Delta_{\text{Sym}}^\Pi(\Gamma) := \{\pi \times \pi \mid \pi \in \Delta(C_1)\} \subset \Delta^\Pi(\Gamma).$$

For uniformity of notation we define $\Delta(\Gamma) := \Delta(C)$. We extend each player's utility from C to $\Delta(\Gamma)$ linearly, defining $u_i(\pi) = \sum_{s \in C} u_i(s)\pi(s)$ for all $\pi \in \Delta(\Gamma)$.

Definition 2.12. A mixed strategy profile $\pi = (\pi_1, \dots, \pi_n) \in \Delta^\Pi(\Gamma)$ is a **Nash equilibrium** if $u_i(\pi) \geq u_i(t_i, \pi_{-i})$ for all players i and all $t_i \in C_i$. The set of Nash equilibria is denoted $\text{NE}(\Gamma)$.

Definition 2.13. A distribution $\pi \in \Delta(\Gamma)$ is a **correlated equilibrium** if for all players i and all functions $f : C_i \rightarrow C_i$

$$\mathbb{E}_\pi [u_i(f(s_i), s_{-i}) - u_i(s)] := \sum_{s \in C} [u_i(f(s_i), s_{-i}) - u_i(s)] \pi(s) \leq 0.$$

Equivalently, π is a correlated equilibrium if for all i and all $s_i, t_i \in C_i$,

$$\sum_{s_{-i} \in C_{-i}} [u_i(t_i, s_{-i}) - u_i(s_i, s_{-i})] \pi(s) \leq 0.$$

The set of correlated equilibria is denoted $\text{CE}(\Gamma)$.

⁷We write $\Delta^\Pi(\Gamma)$ instead of $\Delta^\Pi(C)$ because Γ specifies the product structure on C .

The interpretation is that $\pi \in \Delta(\Gamma)$ is a joint distribution of recommended strategies for the n players. Each player sees only his recommended strategy and π is a correlated equilibrium when players are best off following these recommendations rather than deviating.

Viewing $\text{NE}(\Gamma)$ and $\text{CE}(\Gamma)$ as subsets of $\Delta(\Gamma)$, we have $\text{NE}(\Gamma) = \text{CE}(\Gamma) \cap \Delta^\Pi(\Gamma)$. **Nash's Theorem** states that Nash equilibria exist, hence so do correlated equilibria [20]. In fact in the same paper Nash also proved that symmetric bimatrix games admit symmetric Nash equilibria as we define below.

Definition 2.14. Let Γ be a symmetric bimatrix game. A **symmetric Nash equilibrium**⁸ is a $(\pi_1, \pi_2) \in \text{NE}(\Gamma)$ with $\pi_1 = \pi_2$. The set of such equilibria is denoted $\text{NE}_{\text{Sym}}(\Gamma) \subseteq \Delta_{\text{Sym}}^\Pi(\Gamma)$. We will at times refer to π_1 itself as a symmetric Nash equilibrium. A **symmetric correlated equilibrium** is a $\pi \in \text{CE}(\Gamma)$ such that $\pi(s_1, s_2) = \pi(s_2, s_1)$ for all $s \in C$, and the set of these equilibria is denoted $\text{CE}_{\text{Sym}}(\Gamma)$.

In symbols, $\text{NE}_{\text{Sym}}(\Gamma) = \text{CE}(\Gamma) \cap \Delta_{\text{Sym}}^\Pi(\Gamma)$ and $\text{CE}_{\text{Sym}}(\Gamma) = \text{CE}(\Gamma) \cap \Delta_{\text{Sym}}(\Gamma)$. Throughout the remainder of the paper we will be discussing symmetric bimatrix games and symmetric Nash and correlated equilibria; we will occasionally drop the word “symmetric” for brevity.

From the definitions we obtain the standard facts:

Proposition 2.15. *For any game Γ the set $\text{NE}(\Gamma)$ is compact and the set $\text{CE}(\Gamma)$ is a compact convex polytope. For a symmetric game the same is true of $\text{NE}_{\text{Sym}}(\Gamma)$ and $\text{CE}_{\text{Sym}}(\Gamma)$.*

3 Basics of exchangeable equilibria

We use Γ to denote a symmetric bimatrix game throughout.

Definition 3.1. A **(symmetric) exchangeable equilibrium** of Γ is a correlated equilibrium such that conditioned on some hidden variable the two players choose i.i.d. actions. The set of such equilibria is denoted

$$\text{XE}_{\text{Sym}}(\Gamma) := \text{CE}(\Gamma) \cap \text{conv} \Delta_{\text{Sym}}^\Pi(\Gamma).$$

The text definition agrees with the equation by Proposition 2.6. In this paper we leave a notion of “asymmetric exchangeable equilibrium” undefined, so we are free to drop the word “symmetric.”

Here we record some of the geometric properties of exchangeable equilibria and their relationship to symmetric Nash and correlated equilibria. For examples, see Section 5, which is largely independent of the intervening Section 4 on interpretations.

Proposition 3.2. *The set $\text{XE}_{\text{Sym}}(\Gamma)$ is compact, convex, and semialgebraic.*

⁸Note that symmetric equilibrium concepts are defined as usual in terms of robustness to unilateral (and so in a sense, asymmetric) deviations, *not* simultaneous identical deviations.

Proof. The set $\text{CE}(\Gamma)$ is a compact, convex polytope (Proposition 2.15) and $\text{conv } \Delta_{\text{Sym}}^{\Pi}(\Gamma)$ is closed, convex, and semialgebraic (Proposition 2.5). \square

Proposition 3.3. *The equilibrium sets are nested:*

$$\text{conv}(\text{NE}_{\text{Sym}}(\Gamma)) \subseteq \text{XE}_{\text{Sym}}(\Gamma) \subseteq \text{CE}_{\text{Sym}}(\Gamma).$$

Proof. A symmetric Nash equilibrium is an i.i.d. correlated equilibrium, so it is an exchangeable equilibrium. The set of exchangeable equilibria is convex, so it contains all mixtures of these. Exchangeable equilibria are symmetric correlated equilibria by definition. \square

The set of symmetric correlated equilibria is always a compact convex polytope and the convex hull of the symmetric Nash equilibria is generically so, since there are generically finitely many symmetric Nash equilibria. The set of exchangeable equilibria need not be polyhedral: see Example 5.3 below. It can be shown that this set is always semialgebraic, i.e., described by finitely many polynomial inequalities (proof omitted).

Theorem 3.4. *There exists a rational⁹ exchangeable equilibrium: $\text{XE}_{\text{Sym}}(\Gamma) \cap \mathbb{Q}^{m \times m} \neq \emptyset$.*

Proof. Nash's Theorem [20] gives a symmetric Nash equilibrium (w, w) . It is known that w can be taken to be in \mathbb{Q}^m . For example we can take w to be an extreme point of a certain polyhedron defined by linear inequalities with rational coefficients, by a symmetric version of the argument in [17]. Then $ww^T \in \text{XE}_{\text{Sym}}(\Gamma) \cap \mathbb{Q}^{m \times m}$. \square

The existence of exchangeable equilibria can be proven, without invoking Nash's Theorem or fixed point results, by extending the methods of [15]. We defer details to Part II of this paper for two reasons. First, the proofs go through in the more general context treated there. Second, it is not clear whether such methods can prove the existence of a *rational* exchangeable equilibrium, a result specific to bimatrix games.

While symmetric games always admit symmetric Nash equilibria, asymmetric Nash equilibria can also appear (e.g. in Chicken and Examples 5.2 and 5.3 below). Games with asymmetric equilibria always have strict separation between symmetric correlated and exchangeable equilibria:

Theorem 3.5. *If $\text{NE}_{\text{Sym}}(\Gamma) \subsetneq \text{NE}(\Gamma)$ then $\text{XE}_{\text{Sym}}(\Gamma) \subsetneq \text{CE}_{\text{Sym}}(\Gamma)$.*

Proof. Let $(x, y) \in \text{NE}(\Gamma)$ with $x \neq y$. Both xy^T and yx^T are correlated equilibria, so $W := \frac{1}{2}(xy^T + yx^T) \in \text{CE}_{\text{Sym}}(\Gamma)$. Since $x, y \in \Delta(C_1)$ are distinct, there is a z with $z^T x < 0$ and $z^T y > 0$. This means W is not positive semidefinite so not conditionally i.i.d.:

$$z^T W z = \frac{1}{2}(z^T x y^T z + z^T y x^T z) = (z^T x)(z^T y) < 0. \quad \square$$

In general exchangeable equilibria need not be mixtures of symmetric Nash equilibria (Example 5.3 below), but for two-strategy symmetric bimatrix games this is always the case:

⁹Here we mean that the coordinates are rational in the algebraic sense, i.e., ratios of integers. By definition exchangeable equilibria are rational in the game-theoretic incentive-compatibility sense.

Theorem 3.6. *If Γ has $m = 2$ strategies per player then $\text{conv}(\text{NE}_{\text{Sym}}(\Gamma)) = \text{XE}_{\text{Sym}}(\Gamma)$.*

Proof. Let $A = \begin{bmatrix} x & y \\ z & w \end{bmatrix}$ be the payoff matrix for the row player, $b := x - z$ and $c := w - y$. An exchangeable equilibrium $W = \begin{bmatrix} p & q \\ q & r \end{bmatrix}$ must satisfy the correlated equilibrium constraints and be conditionally i.i.d., which is the same as double nonnegativity per Theorem 2.9. Altogether the conditions are:

$$\begin{aligned} pr &\geq q^2 && \text{(semidefiniteness),} \\ p, q, r &\geq 0 && \text{(nonnegativity),} \\ bp &\geq cq && \text{(incentive constraint \#1),} \\ cr &\geq bq && \text{(incentive constraint \#2), and} \\ p + 2q + r &= 1 && \text{(normalization).} \end{aligned}$$

Let us show that any extreme point W of $\text{XE}_{\text{Sym}}(\Gamma)$ is in $\text{conv}(\text{NE}_{\text{Sym}}(\Gamma))$. If $pr = q^2$ then $\text{rank}(W) = 1$, so it is a Nash equilibrium. If $q = 0$ then we can write $W = p \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} + r \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$ as a convex combination of symmetric pure Nash equilibria.

The remaining case is that semidefiniteness and nonnegativity are not tight. Extremality requires that the remaining three linear conditions on p , q , and r be tight and linearly independent. Multiplying the tight incentive constraints yields $(bc)pr = (bc)q^2$. But $pr > q^2$, so $b = 0$ or $c = 0$. Then $bp = cq$ and $p, q > 0$ give $b = c = 0$. The incentive constraints are trivial, hence linearly dependent. \square

4 Interpretations of exchangeable equilibria

We next study various game-theoretic setups in which exchangeable equilibria arise naturally as outcomes of games. Of course these are all related (some more obviously than others), but each has something different to add to the overall picture. Taken together, the fact that these interpretations each give rise to the same solution concept leads us to conclude that exchangeable equilibria are natural objects to study.

4.1 Hidden variable interpretation

In this section we formalize the discussion of hidden variables given in the introduction. We do so by first fixing terminology for a more explicit notion of correlated equilibrium (which we call external), then studying these in a symmetric setting.

4.1.1 External correlated equilibria

To study hidden variables we consider an equivalent definition of correlated equilibria in which the given game is augmented by adding a correlating device which sends the players pre-play signals. The framework is illustrated in Figure 1. There is some underlying (random) state of the world which is not assumed to be known to the players. The information available to each player is an arbitrary function of this state and some private random noise;

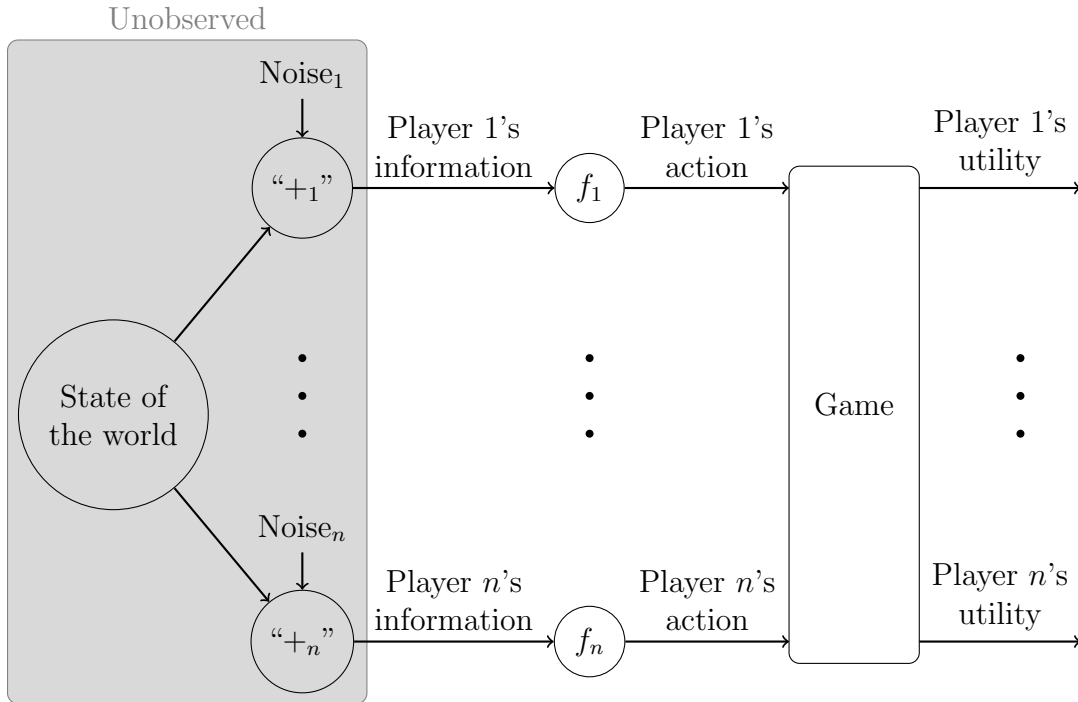


Figure 1: In a correlation scheme, each player receives noisy information about the state of the world (the “+_{*i*}” indicates the state is being combined with the noise somehow, not necessarily additively) and chooses his action in the game as a function f_i of this information. If no player can improve his utility by playing a different function of his information, we call all this data (the functions and the information structure together) an external correlated equilibrium.

the state and all the noise signals are assumed independent. This random noise could for example be measurement noise, or the outcome of some private coin tosses on which a player will base his action – anything which we wish to explicitly assume the other players cannot access.

For mathematical simplicity we typically think of the state of the world, the noises, and the players’ information as all taking values in some finite sets. If this is not the case we must make some measurability restrictions. We speak informally and avoid assigning symbols to everything involved to prevent an explosion of notation.

Each player chooses a function f_i mapping his information to actions. We refer to all the data together – the state distribution, the noise distributions, the maps from these to information, and the f_i – as a **correlation scheme**. We assume these data are known to the players: only the realizations of the random quantities are hidden. If no player can improve his expected utility by unilaterally deviating to a different function f'_i , we refer to the correlation scheme as an **external correlated equilibrium**. When we wish to emphasize the distinction we will refer the content of Definition 2.13 as an **internal correlated**

equilibrium. These two notions were studied by Aumann in [2] and [3], respectively¹⁰, and are closely related:

Proposition 4.1. *The distribution of actions in an external correlated equilibrium is an internal correlated equilibrium and every internal correlated equilibrium arises in this way.*

Proof. Given an external correlated equilibrium, no player can gain by deviating from f_i . In particular no player can gain by deviating to $\zeta_i \circ f_i$ for any $\zeta_i : C_i \rightarrow C_i$. Therefore if we push the f_i back into the unobserved part of the model in Figure 1, merging it into the noise-adding stage, the result is an external correlated equilibrium in which all players choose the identity function from their information to their action. This coincides with the definition of what it means for the distribution over information / actions to be an internal correlated equilibrium.

Conversely, given a $\pi \in \Delta(\Gamma)$ we can design a correlation scheme in which the state of the world is a random strategy profile distributed according to π , each player’s information is equal to his personal component of this strategy profile (so the “noise adding” step just strips away the other players’ choices of strategy – no additional randomness is needed), and each f_i is the identity on C_i . By the definitions this is an external correlated equilibrium if and only if π is an internal correlated equilibrium. \square

For a given correlation scheme we say that a player **knows the state of the world** if the state of the world is a function of (measurable with respect to) his information. This means that player knows all relevant information which he can know in theory, but not the outcomes of the other players’ private environment measurements or coin tosses. We have built some redundancy into the model in the sense that we could have chosen to push all the random noise into the state of the world, making each player’s information a deterministic function of the state. However, the resulting model would not allow for private coin tosses. Compare the following statements:

Proposition 4.2. *In an external correlated equilibrium, if each player knows the state of the world then the outcome conditioned on the state is a Nash equilibrium almost surely, and every Nash equilibrium arises in this way.*

Proposition 4.3. *In an external correlated equilibrium, if each player knows the state of the world and all the noise signals are constant (or equivalently are considered part of the state) then the outcome conditioned on the state is a pure Nash equilibrium almost surely, and every pure Nash equilibrium arises in this way.*

The proofs of these propositions amount to little more than repeating the definitions. The main idea is that if the players commonly know the state, then conditioned on this knowledge their play is independent and each is best replying to his opponent. That is to say, the outcome is conditionally a Nash equilibrium, which must be pure if the players cannot privately randomize.

¹⁰Aumann refers to both notions simply as *correlated equilibrium*. We know of no standard terminology for differentiating the two, and so introduce the terms *internal* and *external* here.

4.1.2 Symmetric external correlated equilibria

In the introduction we assumed that the players are copies of an identical decision-making agent. Thus players perform identical experiments and flip identical coins, interpreting and acting on the results in the same way. Randomness in these outcomes is the one way in which differences in the players' actions arise. Stated precisely:

Definition 4.4. We say that a correlation scheme is **symmetric** if all players have the same noise distribution, mapping from state and noise to information, and mapping f_i from information to action. We say that an external correlated equilibrium is **symmetric** if it arises from a symmetric correlation scheme.

The notion of correlation schemes is well-studied [2], but to the authors' knowledge this symmetric version has not been considered. Under a symmetric correlation scheme, the players' noise distributions are i.i.d., so their information, and hence actions, are i.i.d. conditioned on the state. We immediately obtain the expected generalization of Proposition 4.2:

Proposition 4.5. *In a symmetric external correlated equilibrium, if each player knows the state of the world then the outcome conditioned on the state is a symmetric Nash equilibrium almost surely, and every symmetric Nash equilibrium arises in this way.*

On the other hand the generalization of Proposition 4.1 looks a bit different:

Proposition 4.6. *The distribution of actions in a symmetric external correlated equilibrium is an exchangeable equilibrium and every exchangeable equilibrium arises in this way.*

Proof. By Proposition 4.1 this distribution is a correlated equilibrium. It is also i.i.d. conditioned on the state, so it is exchangeable.

Conversely, let π be any exchangeable equilibrium. Then we can write $\pi = \sum_{i=1}^k \lambda_i x_i x_i^T$ for some $x_i \in \Delta(C_1)$ and some probability vector λ . Construct a symmetric correlation scheme with k states of the world chosen according to the distribution λ . In state i , choose each player's information i.i.d. according to x_i . Let the mappings from information to action all be the identity map. Then the distribution over actions is π , and since this is in particular an internal correlated equilibrium, the resulting correlation scheme is a symmetric external correlated equilibrium. \square

Under symmetry assumptions, the widely used equivalence between external and internal correlated equilibria breaks down. Instead, symmetric external correlated equilibria give rise to the more natural notion of exchangeable equilibria.

Thus we conclude that there are games (such as Examples 5.2 and 5.3 below, or any symmetric game admitting asymmetric Nash equilibria as shown in Theorem 3.5) admitting symmetric internal correlated equilibria which cannot be realized by any symmetric external correlated equilibria because they fail to be exchangeable. Implementing such equilibria requires implicit symmetry-breaking.

$(u_{\text{Row}}, u_{\text{Column}})$	My name is Alice	My name is Bob
My name is Alice	(0, 0)	(1, 1)
My name is Bob	(1, 1)	(0, 0)

Table 2: Anti-coordinating based on player identity.

4.2 Unknown opponent interpretation

As in the previous section on the hidden variable interpretation, here we again focus on games in which the symmetry includes not only the action spaces and payoffs, but also the roles of the players in a wider sense which we now discuss in more detail. Consider the anti-coordination game of Example 5.2 with its strategies named as in Table 2. We will first give two examples of choices of players without this symmetry and then an example with this symmetry. Note that because of the symmetry of the payoffs, we do not have to select who will be the row player and who will be the column player. In particular this information is not available to help the players coordinate on an equilibrium.

Suppose first that this game were played by two people named Alice and Bob whose names were common knowledge. It is reasonable to predict that the players would have no trouble coordinating on the pure strategy Nash equilibrium in which each player correctly identifies his own name. Suppose secondly that this game were played by two friends named Alice. It may still be possible for them to coordinate based on, say, the common knowledge that one of them has the middle name Roberta and the other does not.

Suppose finally that two arbitrary Alices are selected from the record-breaking crowd at this year’s Conference of Game Theorists Named Alice. They are sequestered in separate rooms and asked to choose the action they would play in this game against an opponent chosen from the same conference who has been given the same information. It is common knowledge that both players are attending the conference, but no further information is given. As Bayesian observers, how should we expect them to play? The standard Bayesian way to capture our ignorance of any distinctions between the possible players is to say that the distribution over outcomes should be symmetric. Based on the conference title we can assume common knowledge of rationality, so by Aumann’s result [3] we should expect the Alices to play a correlated equilibrium. In particular we should not expect them to do something foolish like both playing “My name is Alice” with probability one.

What else can we say about the outcome? The best symmetric correlated equilibrium in terms of payoff is $W = \begin{bmatrix} 0 & 1/2 \\ 1/2 & 0 \end{bmatrix}$. Is this a reasonable solution?

We claim that W would not be played. Suppose for a contradiction that W were the expected distribution of outcomes and consider asking a third Alice – recall the record attendance – what action (“Alice” or “Bob”) she would play in the game. The chosen strategies A_1 , A_2 , and A_3 of the three Alices would all be random variables, and their joint distribution would be symmetric under arbitrary permutations of the Alices. In particular all the pairwise distributions of two of these random variables would be given by W . Consulting W , we see that the events $E_1 = \{A_2 = A_3\}$, $E_2 = \{A_1 = A_3\}$, and $E_3 = \{A_1 = A_2\}$ would

each occur with probability zero, therefore so would their union $E = E_1 \cup E_2 \cup E_3$. But there are only two strategies in the game, so in any realization at least two of the A_i must be equal by the pigeonhole principle. That is to say, E must occur with probability one regardless of the distribution of the A_i , a contradiction.

On the other hand, the mixed strategy Nash equilibrium in which each player randomizes over her actions with equal probability does not suffer from this difficulty: one can create a sequence of i.i.d. Bernoulli(1/2) random variables for any number of Alices. Such a sequence has a distribution which is exchangeable and has the desired mixed equilibrium as its marginal distribution corresponding to any choice of two players.

Strictly speaking we have only argued that the observed correlated equilibrium should extend to an N -exchangeable sequence A_1, \dots, A_N , where N is the number of Alices at the conference. However, with N assumed to be large, it is more natural to consider equilibria which could be extended to an arbitrary number of attendees. In particular this ensures that we select only the equilibria which are robust to the exact size of the pool and the players' knowledge thereof.

Being a bit more formal, we consider situations in which the players are drawn from a pool of N interchangeable agents ($2 \leq N \leq \infty$). We interpret interchangeability as meaning that all N agents choose a hypothetical action for the game, regardless of whether they are selected to play, and the joint distribution of these actions is N -exchangeable. Whichever two agents are chosen to play the game, we assume they choose their actions in their best interests, in the sense that their joint distribution is a correlated equilibrium. That is to say, the N agents choose their actions according to an N -exchangeable equilibrium:

Definition 4.7. For $2 \leq N \leq \infty$ an **N -exchangeable equilibrium** of Γ is an N -exchangeable distribution of random variables X_i such that the marginal distribution of X_1 and X_2 is a correlated equilibrium of Γ . The set of these equilibria is denoted $\text{XE}_{\text{Sym}}(\Gamma, N)$.

The sets of N -exchangeable equilibria shrink as N increases, in the sense that if $N_2 > N_1$ and X_1, \dots, X_{N_2} is an N_2 -exchangeable equilibrium, then X_1, \dots, X_{N_1} is an N_1 -exchangeable equilibrium. Letting N go to infinity, De Finetti's Theorem gives:

Corollary 4.8. *The distribution of two random variables X_1, X_2 is an exchangeable equilibrium if and only if it can be extended to a distribution of an infinite sequence X_1, X_2, X_3, \dots which is an ∞ -exchangeable equilibrium.*

Therefore we may view exchangeable equilibria as exactly those correlated equilibria arising when players are selected from a large pool of interchangeable agents and their identities are unknown to each other. It makes sense to say that the notion of exchangeable equilibria eliminates those correlated equilibria which are “unreasonable” on the grounds of symmetry arguments such as the one made above. Of course, in settings such as the Alice and Bob example where these kinds of symmetry arguments do not apply, we may expect other equilibria.

4.3 Many player interpretation

Another way to view exchangeable equilibria is as limits of symmetric correlated equilibria of N -player games consisting of many identical pairwise interactions. Given a symmetric bimatrix game Γ we define a symmetric N -player game $\Gamma^{(N)}$ in which player i has strategy set $C_i^{(N)} = C_1$ and utility function

$$u_i^{(N)}(s_1, \dots, s_N) = \sum_{j \neq i} u_1(s_i, s_j).$$

That is to say, each player plays Γ against each other player, choosing the same strategy in each instance. If $N = 2$ we recover the original game: $\Gamma^{(2)} = \Gamma$.

Example 4.9. We give another interpretation of the anti-coordination game Γ whose utilities are shown in Table 2. The story is different from the one in the previous section, so we simply refer to the strategies as ‘A’ and ‘B’. The utilities of the N -player game $\Gamma^{(N)}$ are given by

$$u_i^{(N)}(s_1, \dots, s_N) = \#\{j : s_j \neq s_i\},$$

so each player wants to choose whichever strategy is chosen by fewer of his opponents. This is a version of “The Minority Game” introduced by Challet and Zhang [9]. One interpretation is that the N players are choosing between two equally good restaurants, A and B , so each player wants to eat at the less crowded restaurant.

The game $\Gamma^{(N)}$ is symmetric under arbitrary permutations of the N players, so it is natural to focus on correlated equilibria of this game which are also symmetric, i.e., N -exchangeable. We now show that these are exactly the N -exchangeable equilibria of Definition 4.7. The reason is that the utility for a player in $\Gamma^{(N)}$ is summed over all his $N - 1$ interactions, so by linearity only the bivariate marginal of an N -exchangeable distribution matters to determine whether it is a correlated equilibrium of $\Gamma^{(N)}$. A simple computation then shows that the correct condition on this bivariate marginal is that it be a correlated equilibrium of Γ .

Proposition 4.10. *For finite¹¹ N symmetric correlated equilibria of $\Gamma^{(N)}$ are the same as N -exchangeable equilibria of Γ .*

Proof. For any $f : C_1 \rightarrow C_1$ we have

$$\begin{aligned} \mathbb{E} u_1^{(N)}(f(X_1), X_2, \dots, X_N) &= \mathbb{E} \sum_{i=2}^N u_1(f(X_1), X_i) = \sum_{i=2}^N \mathbb{E} u_1(f(X_1), X_i) \\ &= \sum_{i=2}^N \mathbb{E} u_1(f(X_1), X_2) = (N - 1) \mathbb{E} u_1(f(X_1), X_2), \end{aligned}$$

with the same analysis holding for all the other players. Therefore there exists an f which can improve the payoff of some player in $\Gamma^{(N)}$ under recommendations X_1, \dots, X_N if and only if there exists such an f in Γ under recommendations X_1, X_2 . \square

¹¹We do not interpret ∞ -exchangeable equilibria directly as symmetric correlated equilibria of some game to avoid technical difficulties related to games with infinitely many players.

Thus exchangeable equilibria correspond to symmetric correlated equilibria of N -player extensions of the game for arbitrary N . Given that the original game was symmetric, it is reasonable to expect that equilibria with a high degree of symmetry would be focal (apt to be chosen over other equilibria because they attract attention by being “better” in some way). This is another reason we might expect players to play an exchangeable equilibrium.

Example 4.9. (cont’d) Which equilibria should we expect to be played in The Minority Game? There are an abundance of pure equilibria of $\Gamma^{(N)}$; these are exactly the strategy profiles in which $\lfloor \frac{N}{2} \rfloor$ players choose one restaurant and $\lceil \frac{N}{2} \rceil$ choose the other. These all have the disadvantage of not being symmetric in the players. While such equilibria can be justified with an evolutionary model [9], they do not make as much sense in a symmetric one-shot game.

We can symmetrize by constructing a correlated equilibrium which picks one of these pure Nash equilibria uniformly at random. This yields an N -exchangeable equilibrium π^N . One can show that if N is odd then π^{N+1} is the unique extension of π^N to an $(N + 1)$ -exchangeable distribution (the extra player in $\Gamma^{(N+1)}$ goes to the less populated restaurant). On the other hand, if N is even then there is no $(N + 1)$ -exchangeable $\pi \in \Delta_{\text{Sym}}(C_1^{N+1})$ extending π^N . This means π^N cannot be extended to an $(N + 2)$ -exchangeable sequence for any N .

In particular, there is no symmetric way to extend any π^N to a correlated equilibrium of all the games $\Gamma^{(k)}$ simultaneously. Therefore expecting the players to actually play one of the π^N means assuming that the value of $N \pm 1$ (depending on the parity of N) is common knowledge. This assumption seems dubious in this model.

We claim it is more natural to look for a solution which is robust to the value of N . This could be accomplished in a number of ways, such as by fixing some value N which is much greater than the actual expected number of players and using π^N (which only solves the problem in an approximate sense), or by assuming the players have a probability distribution over possible values of N [19]. Exchangeable equilibria are those which are the most robust to the value of N ; namely, they are correlated equilibria of $\Gamma^{(N)}$ for all N .

The unique exchangeable equilibrium of the anti-coordination game in Table 2 is the mixed Nash equilibrium (proven in Example 5.2 below). That is to say, the only Bayesian rational strategy of the players in The Minority Game which is symmetric and makes sense regardless of the number of players is for everyone to choose a restaurant uniformly at random.

Finally, it is worth noting that the fact that the exchangeable equilibria correspond to symmetric correlated equilibria of some N -player extension for all N seems to be robust to how exactly we formulate $\Gamma^{(N)}$. For example, the above formulation corresponds to a complete graph of interactions, in the sense that each of the N players plays Γ against the remaining $N - 1$ players. We could alter this to construct a game based on any other graph of interactions. An exchangeable equilibrium is simultaneously a symmetric correlated equilibrium of all these games. Furthermore, we could allow the players in these extensions to consider different deviations for different opponents in this interaction graph, in which case the exchangeable equilibria would be exactly the correlated equilibria in which the players

choose the same action against all opponents. The proofs of both of these facts are similar to the proof of Proposition 4.10.

4.4 Sealed envelope implementation

For $1 \leq N \leq \infty$ we can assign another interpretation to $\text{XE}_{\text{Sym}}(\Gamma, N)$ in terms of which correlated equilibria can be implemented by a certain type of correlation scheme. Suppose Γ is to be played by two players who have access to a common pile of N indistinguishable sealed envelopes, each of which contains a slip of paper on which an element of C_1 is written. The contents of these envelopes play the role of the recommendations in the formulation of (internal) correlated equilibrium.

Each player is allowed to choose one envelope and base his choice of action in Γ on its contents, which he examines privately. We assume that there is some mechanism in place preventing both players from choosing the same envelope (e.g. a random one of them is required to select a different envelope).

The fact that the envelopes are indistinguishable means that the distribution over their contents should be exchangeable. This distribution is an n -exchangeable equilibrium if and only if it is a Nash equilibrium of this larger game (including the envelope selection process) for both players to choose distinct envelopes in an arbitrary fashion and then play the strategy written inside the envelope they choose. In particular, they cannot exploit the knowledge of which envelope the other player will choose in any way.

5 Examples

In this section we compare the sets of Nash, exchangeable, and correlated equilibria of some small example games. The main tool for doing so is Theorem 2.9.

We warm up with two 2×2 examples, including the anti-coordination game discussed in the introduction, illustrating the general fact that $\text{XE}_{\text{Sym}} = \text{conv}(\text{NE}_{\text{Sym}})$ for such games (Theorem 3.6). We then move on to our main example, in which we show that $\text{conv}(\text{NE}_{\text{Sym}})$, XE_{Sym} , and CE_{Sym} can all be distinct for 3×3 games. This example also shows that XE_{Sym} need not be polyhedral.

All these examples are symmetric bimatrix games of identical interest, i.e., they satisfy $A = A^T = B$. It is interesting that such behavior arises in these games, because identical interest games ($A = B$) are often thought of as somewhat trivial: they always have pure Nash equilibria, for example. However, the pure Nash equilibria need not be symmetric, so when restricting attention to symmetric equilibria identical interest games seem to be a less trivial class.

Example 5.1. Consider the coordination game with utility matrices $A_1 = B_1 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. Parametrize symmetric probability matrices as $\begin{bmatrix} p & q \\ q & r \end{bmatrix}$ with $p, q, r \geq 0$ and $p + 2q + r = 1$. Then the correlated equilibrium conditions are exactly $p \geq q$ and $r \geq q$. Note that together these imply that $pr \geq q^2$, so any symmetric correlated equilibrium is automatically positive semidefinite (the principal minors are nonnegative), hence conditionally i.i.d. by Theorem 2.9

$(u_1 = u_2)$	a	b	c
a	2	2	0
b	2	1	2
c	0	2	2

Table 3: A game for which the convex hull of the Nash equilibria is strictly contained in the set of exchangeable equilibria which is itself strictly contained in the set of correlated equilibria.

and so exchangeable. Therefore the set of symmetric exchangeable equilibria is the same as the set of symmetric correlated equilibria and a simple computation shows that these sets are equal to

$$\text{CE}_{\text{Sym}} = \text{XE}_{\text{Sym}} = \text{conv} \left\{ \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 1/4 & 1/4 \\ 1/4 & 1/4 \end{bmatrix} \right\} = \text{conv}(\text{NE}_{\text{Sym}}).$$

Example 5.2. Consider the anticoordination game with utility matrices $A_2 = B_2 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ discussed in Section 4.2 and again in the context of The Minority Game in Section 4.3. Up to a positive affine transformation of the utilities this game is the same as the game Chicken analyzed in Section 1.2, and so has the same equilibria:

$$\begin{aligned} \left\{ \begin{bmatrix} 1/4 & 1/4 \\ 1/4 & 1/4 \end{bmatrix} \right\} &= \text{NE}_{\text{Sym}} = \text{XE}_{\text{Sym}} \\ &\subsetneq \text{CE}_{\text{Sym}} = \text{conv} \left\{ \begin{bmatrix} 0 & 1/2 \\ 1/2 & 0 \end{bmatrix}, \begin{bmatrix} 1/3 & 1/3 \\ 1/3 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 1/3 \\ 1/3 & 1/3 \end{bmatrix}, \begin{bmatrix} 1/4 & 1/4 \\ 1/4 & 1/4 \end{bmatrix} \right\}. \end{aligned}$$

The games in Examples 5.1 and 5.2 could be said to be isomorphic as bimatrix games but not as symmetric bimatrix games, in the sense that one can be transformed into the other by relabeling the strategies, but not by using the same relabeling for both players. This accounts for the fact that their sets of symmetric exchangeable equilibria are not isomorphic. The difference corresponds to the fact that when these games are played by many players simultaneously, there is a symmetric (with respect to the players) way to break symmetry between the two pure strategies in the case of the coordination game but not in the case of the anticoordination game.

Also note that in both cases we had $\text{XE}_{\text{Sym}} = \text{conv}(\text{NE}_{\text{Sym}})$ as predicted by Theorem 3.6. We now construct a 3×3 game with $\text{conv}(\text{NE}_{\text{Sym}}) \subsetneq \text{XE}_{\text{Sym}} \subsetneq \text{CE}_{\text{Sym}}$.

Example 5.3. Consider the game in Table 3. First, we show strict containment in $\text{XE}_{\text{Sym}} \subsetneq \text{CE}_{\text{Sym}}$. Second, we show strict containment in $\text{conv}(\text{NE}_{\text{Sym}}) \subsetneq \text{XE}_{\text{Sym}}$. Third, we observe that XE_{Sym} is not polyhedral. These results are summarized in Figure 2.

The matrix

$$W^1 = \frac{1}{4} \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

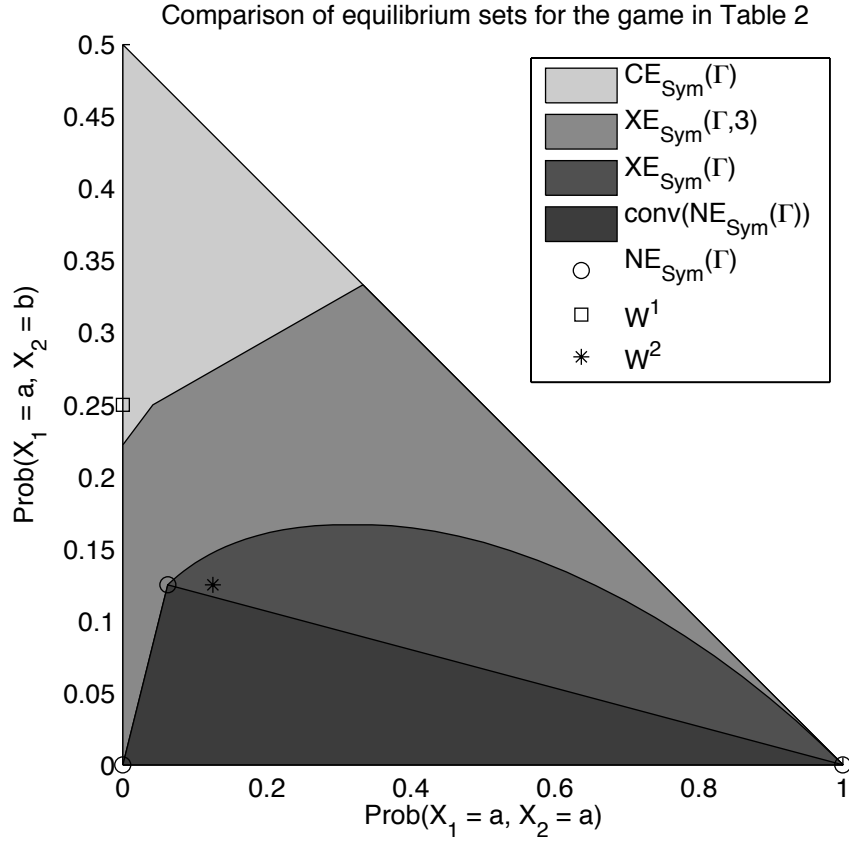


Figure 2: Comparison of equilibrium sets for the game in Table 3. These sets are naturally sets of symmetric 3×3 matrices (or $3 \times 3 \times 3$ tensors in the case of $\text{XE}_{\text{Sym}}(\Gamma, 3)$), but we have chosen a projection into two dimensions which highlights the separation between the sets. The set of exchangeable equilibria is not polyhedral despite the fact that the other sets are.

is a correlated equilibrium because both players get their maximum payoff with probability one. By Proposition 2.7 it is not conditionally i.i.d., so it is a correlated equilibrium which is not exchangeable. In fact we can say more: W^1 cannot be extended to a 3-exchangeable distribution. Suppose X_1, X_2, X_3 were random variables with a 3-exchangeable distribution taking values in $\{a, b, c\}$ such that the distribution of X_1 and X_2 were W^1 . Then

$$\begin{aligned} \frac{1}{4} = \mathbb{P}(X_1 = a) &\leq \mathbb{P}((X_1 = a, X_2 \neq b) \vee (X_1 = a, X_3 \neq b) \vee (X_2 = b, X_3 = b)) \\ &\leq \mathbb{P}(X_1 = a, X_2 \neq b) + \mathbb{P}(X_1 = a, X_3 \neq b) + \mathbb{P}(X_2 = b, X_3 = b) \\ &= 2W_{11}^1 + 2W_{13}^1 + W_{22}^1 = 0, \end{aligned}$$

which is a contradiction, so no such 3-exchangeable distribution exists.

One can also verify that

$$W^2 = \frac{1}{8} \begin{bmatrix} 1 & 1 & 0 \\ 1 & 2 & 1 \\ 0 & 1 & 1 \end{bmatrix} = \frac{1}{8} \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix}$$

is a correlated equilibrium, and the exhibited factorization shows that W^2 is conditionally i.i.d., hence exchangeable. Suppose for a contradiction that W^2 were a convex combination of Nash equilibria. Then at least one of the Nash equilibria in the convex combination would have to assign positive probability to the strategy profile (b, b) . Suppose player 2 did not play c with positive probability in such a Nash equilibrium. Given this information, player 1 prefers a to b , so player 1 cannot choose b with positive probability in such an equilibrium, a contradiction. Symmetric arguments show that each player must play all his strategies with positive probability. Therefore this Nash equilibrium has full support. But W^2 has entries which are zero, hence this Nash equilibrium cannot be included in an expression of W^2 as a convex combination of Nash equilibria. Thus W^2 is not a convex combination of Nash equilibria.

This argument also shows that the only symmetric Nash equilibrium which assigns positive probability to b is the one with full support, which a simple computation shows to be $[\frac{1}{4} \ \frac{1}{2} \ \frac{1}{4}]$. The only other symmetric Nash equilibria are $[1 \ 0 \ 0]$ and $[0 \ 0 \ 1]$. There can be no symmetric Nash equilibrium which assigns zero probability to b but positive probability to both a and c . Such an equilibrium would have to assign equal probability to a and c , but b is the unique best response to such a mixture.

As is well known, the set of correlated equilibria of a game is always polyhedral, hence so is the set of k -exchangeable equilibria. The set of Nash equilibria is generically finite, so its convex hull is polyhedral for generic games (and for this game in particular). It is visually evident, and can be proven algebraically, that the projection of the set of exchangeable equilibria pictured in Figure 2 is not polyhedral. In fact it is an algebraic curve of degree 11 which factors into three linear components, a quadratic component, and a degree six component over \mathbb{Q} . Two of the linear components are easily visible (the bottom and left of the convex hull of the symmetric Nash equilibria), and the third corresponds to the maximum y value, attained along the line segment joining $(\frac{11}{36}, \frac{1}{6})$ to $(\frac{1}{3}, \frac{1}{6})$. The quadratic component

corresponds to the curved portion of the boundary to the right of this maximum and is defined by the vanishing of $x^2 + 2xy + 4y^2 - x$. The degree six component is the curved portion of the boundary to the left of the maximum; we omit its equation for brevity.

6 Conclusions and Future Work

In the case of symmetric bimatrix games, we have argued that the exchangeable equilibria are a natural mathematical object with a variety of game-theoretic interpretations. The set of exchangeable equilibria lies between the sets of Nash and correlated equilibria, and its structural properties are a mix of the two. It is convex like the correlated equilibria, but its semialgebraicity and nonpolyhedrality remind one more of Nash equilibria.

The exchangeable equilibria are not the only object living in the gap between the Nash and correlated equilibria. There are also, for example, Sorin’s distribution equilibria [25], which neither contain nor are contained in the exchangeable equilibria. We believe there is room for much more insight to be gained interpolating between Nash and correlated equilibria.

In this paper we have largely ignored the question of what “exchangeable equilibrium” should mean in symmetric games with more than two players. We extend the theory to such games in Part II, studying what is required of the symmetry structure to obtain theoretical results like those of Section 3 and interpretations along the lines of Section 4 of this paper. This more abstract setting is a natural one for proving existence of exchangeable equilibria by adapting Hart and Schmeidler’s argument [15]. We also modify the sealed envelope interpretation of Section 4.4 to design a hierarchy of nested convex sets called higher order exchangeable equilibria interpolating between the exchangeable equilibria and the convex hull of the symmetric Nash equilibria, to which they converge in the limit.

These ideas and relevant computational questions are all explored further in the first author’s doctoral thesis [26].

Acknowledgements

The authors would like to acknowledge Muhamet Yildiz and Dirk Bergemann for their helpful insights and suggestions on this topic. This research was funded by the National Science Foundation under Award 1027922 and the Air Force Office of Scientific Research under grant FA9550-11-1-0305.

References

- [1] C. Alós-Ferrer and C. Kuzmics. Hidden symmetries and focal points. Working paper, June 2010.

- [2] R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67 – 96, 1974.
- [3] R. J. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55(1):1 – 18, January 1987.
- [4] S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in Real Algebraic Geometry*. Springer, 2nd edition, 2006.
- [5] D. P. Bertsekas, A. Nedić, and A. E. Ozdaglar. *Convex Analysis and Optimization*. Athena Scientific, Belmont, MA, 2003.
- [6] K. Binmore. Social contract I: Harsanyi and Rawls. *The Economic Journal*, 99:84–102, 1989.
- [7] A. Brandenburger. Knowledge and equilibrium in games. *The Journal of Economic Perspectives*, 6(4):83–101, Autumn 1992.
- [8] A. Casajus. Focal points in framed strategic forms. *Games and Economic Behavior*, 32:263–291, 2000.
- [9] D. Challet and Y.-C. Zhang. Emergence of cooperation and organization in an evolutionary game. *Physica A*, 246:407 – 418, 1997.
- [10] 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.
- [11] S. Du. Correlated equilibrium and higher order beliefs about play. Working paper, January 2011.
- [12] R. M. Dudley. *Real Analysis and Probability*. Cambridge University Press, New York, 2002.
- [13] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1991.
- [14] J. Harsanyi. Can the maximum principle serve as a basis for morality? a critique of John Rawls’s theory. *American Political Science Review*, 69(2):594–606, 1975.
- [15] S. Hart and D. Schmeidler. Existence of correlated equilibria. *Mathematics of Operations Research*, 14(1):18 – 25, February 1989.
- [16] J. Hillas, E. Kohlberg, and J. Pratt. Correlated equilibrium and Nash equilibrium as an observer’s assessment of the game. Harvard Business School working paper #08-005, July 2007.
- [17] O. L. Mangasarian. Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics*, 12(4):778 – 780, December 1964.

- [18] J. E. Maxfield and H. Minc. On the matrix equation $X'X = A$. *Proceedings of the Edinburgh Mathematical Society (Series 2)*, 13:125 – 129, 1962.
- [19] R. B. Myerson. Population uncertainty and Poisson games. *International Journal of Game Theory*, 27:375 – 392, 1998.
- [20] J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286 – 295, September 1951.
- [21] R. F. Nau and K. F. McCardle. Coherent behavior in noncooperative games. *Journal of Economic Theory*, 50:424 – 444, 1990.
- [22] G. Pólya. Sur quelques points de la théorie des probabilités. *Annales de l'I.H.P.*, 1(2):117 – 161, 1930.
- [23] J. Rawls. *A Theory of Justice*. The Belknap Press of Harvard University Press, Cambridge, MA, revised edition, 1999.
- [24] M. J. Schervish. *Theory of Statistics*. Springer-Verlag, New York, 1995.
- [25] S. Sorin. Distribution equilibrium I: Definition and examples. Unpublished, December 1998.
- [26] N. D. Stein. *Exchangeable Equilibria*. PhD thesis, Massachusetts Institute of Technology, 2011.
- [27] R. Sugden. A theory of focal points. *The Economic Journal*, 105(430):533–550, May 1995.