

Private and Third-Party Randomization in Risk-Sensitive Equilibrium Concepts

Mickey Brautbar and Michael Kearns and Umar Syed

Department of Computer and Information Science
University of Pennsylvania

Abstract

We consider risk-sensitive generalizations of Nash and correlated equilibria in noncooperative games. We prove that, except for a class of degenerate games, unless a two-player game has a pure Nash equilibrium, it does not have a risk-sensitive Nash equilibrium. We also show that every game has a risk-sensitive correlated equilibrium. The striking contrast between these existence results is due to the different sources of randomization in Nash (private randomization) and correlated equilibria (third-party randomization).

1 Introduction

The game-theoretic approach to modeling multi-agent interaction assumes that players in a game want to maximize their expected utility. But in many settings, players instead often want to maximize some more complicated function of their utility. In this paper, we ask the following natural question: Can we extend the familiar notions of Nash and correlated equilibria to settings where players are sensitive to *risk*?

In a noncooperative game, the utility for each player depends on the actions taken by all players. In a *Nash equilibrium* of a game, each player chooses an action from a distribution, called a *strategy*, that maximizes her expected utility when she assumes the strategies of the other players are held fixed. A *correlated equilibrium* (Aumann 1987) is a well-known generalization of a Nash equilibrium. In a correlated equilibrium, a third party draws actions for each player from a joint distribution on actions, and each player then decides deterministically whether to play their recommended action or switch to another one; the joint distribution is a correlated equilibrium if no player ever has an incentive to switch. A Nash or correlated equilibrium is an inherently stable state of the game, and thus serves as both a prescriptive and descriptive characterization of the behavior of players in a multi-agent setting. It is well-known that every game has at least one Nash equilibrium (Nash 1950), and therefore also at least one correlated equilibrium.

The expected utility framework for games is obviously very general, but it does exclude the possibility that players in the game have preferences that depend on the *entire* distribution of utility, and not just on its expectation. For example, if a player is sensitive to risk, her objective might be

to choose a strategy that maximizes $E[\text{utility}] - \text{Var}[\text{utility}]$. Indeed, this is the recommendation of modern portfolio theory, and a version of this mean-variance objective is widely used by investors in financial markets. In general, we refer to the objective maximized by a player as her *preference* function.

In this paper, we define generalizations of Nash and correlated equilibria which permit players to maximize preference functions that may differ from expected utility. We will later denote these generalized equilibrium concepts as *F*-Nash and *F*-correlated equilibria, where *F* represents the players' preference functions. In this more general setting, the classical arguments for the existence of equilibria no longer apply. So under what conditions can we guarantee — or rule out — the existence of a *F*-Nash or *F*-correlated equilibrium?

To address this question, we first note that there are two major differences between the original definitions of Nash and correlated equilibria. One is that a correlated equilibrium (as the name suggests) may induce correlations among the actions chosen for the players, whereas all players' actions are independent in a Nash equilibrium. The other difference is the source of randomization: In a Nash equilibrium, each player uses private randomization to choose an action from her distribution, while in a correlated equilibrium the randomization is performed by a third party. The first difference is the only salient one, while the latter difference is largely a matter of interpretation: a correlated equilibrium in which there is no correlation among the players' actions is, by definition, a Nash equilibrium.

These differences also exist between the definitions of *F*-Nash and *F*-correlated equilibria, but in this case neither difference is superficial. Our results in this paper show that, for a large and natural class of risk-sensitive preference functions, the source of randomization has a dramatic impact on the existence of *F*-Nash and *F*-correlated equilibria. Before turning to a discussion of our main contributions, we review the existing literature.

The question of existence of *F*-Nash equilibria for alternative preference functions has long been studied. For example, it is known from the work of (Debreu 1952) that *F*-Nash equilibria always exist if the preference function for each player is continuous and concave in her strategy. However, as first observed by (Crawford 1990), many natu-

ral choices for preference functions are *convex* in a player's strategy, especially those that encode some notion of risk-sensitivity. This is unsurprising: A preference function that is convex in a player's strategy implies that, other things being equal, the player dislikes increasing her randomization, which is quite similar to saying that the player is risk-sensitive.

Many authors (e.g., (Crawford 1990), (Dekel, Safra, and Segal 1991), (Nowak 2005)) have shown that, for several well-motivated convex preference functions, F -Nash equilibria do not necessarily exist. These negative results are almost always obtained in the same way: by exhibiting, for each preference function of interest, a specific game that does not have an F -Nash equilibrium for that preference function. For example, (Nowak 2005) described a simple 2×2 game that does not have an F -Nash equilibrium for the mean-variance preference function described above. A major weakness of this 'counterexample' type of result is that it does not rule out the possibility that the counterexamples are pathological cases. In other words, could it be that these counterexamples comprise only finitely or countably many games? If so, then for any game, one could randomly perturb the game by a tiny amount and thereby obtain a nearly identical game that is guaranteed (with overwhelming probability) to have an F -Nash equilibrium. If this were the case, then current nonexistence results would have essentially no practical importance.

The first contribution of this paper, in Section 5, is a significant generalization of these negative results. We confine our analysis to a large class of so-called *mean-variance* preference functions, which reward higher expected utility but penalize higher variance of utility, and ask whether F -Nash equilibria exist in this case. Intuitively, we might expect that in an F -Nash equilibrium players will be disinclined to choose their actions randomly since, other things being equal, randomization increases variance. We show that this intuition is not only correct but extremely general, and limits the existence of F -Nash equilibria to a very restricted class of games — namely, those with either pure Nash equilibria, or those in which the variance experienced by a player is already "saturated" by the randomization due to other players. In fact, we prove that, in a two-player game, if each player's utility function is chosen randomly, then with probability 1 the game does not have an F -Nash equilibrium in which even a single player i randomizes their choice of action.

Our second contribution, in Section 6, is to observe that an F -correlated equilibrium always exists in a game in which each player has a convex preference function — a class which includes mean-variance preference functions. We explain that an F -correlated equilibrium in a game with convex preferences is actually a strict generalization of an *equilibrium in beliefs*, whose existence was first proved by (Crawford 1990). We show that, counterintuitively, an F -correlated equilibrium need not actually induce any correlation among the behavior of the players — the existence result is solely a consequence of the different source of randomization. The intuition is that when a third party is responsible for the randomization, the players only react deterministically to the stochastic action proposed for them by

this third party — in other words, the "variance lies elsewhere" from the perspective of each player. We further show that for some games with convex preference functions there exists an F -correlated equilibrium in which each player is better off than in any equilibrium in beliefs.

2 Preliminaries

A game has n players, indexed $i = 1, \dots, n$. Let \mathcal{A}_i be the (finite) set of actions available to player i . Let the cross-product $\mathcal{A} = \times_{i=1}^n \mathcal{A}_i$ be the set of action profiles, and let

$$\mathcal{A}_{-i} \triangleq \mathcal{A}_1 \times \dots \times \mathcal{A}_{i-1} \times \mathcal{A}_{i+1} \times \dots \times \mathcal{A}_n$$

be the set of action profiles for all players but player i . If $a \in \mathcal{A}$, then we write $a_i \in \mathcal{A}_i$ for the i th component of a , while $a_{-i} \in \mathcal{A}_{-i}$ denotes dropping a_i from a .

Let $u_i : \mathcal{A} \rightarrow \mathbb{R}$ be the utility function for player i , where $u_i(a)$ is the utility to player i under action profile a . For convenience, $u_i(a)$ can be equivalently written $u_i(a_1, \dots, a_n)$ or $u_i(a_i, a_{-i})$. We assume that utility functions are bounded, but otherwise allow them to be arbitrary.

Let $\mathcal{P}(S)$ be the set of distributions on a (finite) set S . A distribution $\mathbf{p} \in \mathcal{P}(\mathcal{A})$ is called an *action profile distribution*, where $\mathbf{p}(a)$ is the probability assigned by \mathbf{p} to action profile a . Like utility functions, for convenience $\mathbf{p}(a)$ can be equivalently written $\mathbf{p}(a_1, \dots, a_n)$ or $\mathbf{p}(a_i, a_{-i})$.

For any $\mathbf{p} \in \mathcal{P}(\mathcal{A})$, we write $\mathbf{p}_i \in \mathcal{P}(\mathcal{A}_i)$ for the marginal distribution of \mathbf{p} on \mathcal{A}_i , while $\mathbf{p}_{-i} \in \mathcal{P}(\mathcal{A}_{-i})$ denotes the marginal distribution of \mathbf{p} on \mathcal{A}_{-i} .

If $\mathbf{p} \in \mathcal{P}(\mathcal{A})$ is a product distribution (i.e. $\mathbf{p}(a) = \prod_{i=1}^n \mathbf{p}_i(a_i)$), then we call \mathbf{p}_i the *strategy* for player i , and \mathbf{p} is the *strategy profile* for all players. Also, if \mathbf{p}_i is a degenerate distribution concentrated on a single action, then we say that \mathbf{p}_i is a *pure strategy*.

For any distribution $\mathbf{p} \in \mathcal{P}(\mathcal{A})$ and action $\hat{a}_i \in \mathcal{A}_i$ we write $\mathbf{p}_{-i}|\hat{a}_i$ for the conditional distribution on \mathcal{A}_{-i} given that $a \sim \mathbf{p}$ and $a_i = \hat{a}_i$. Note that if \mathbf{p} is a product distribution, then $\mathbf{p}_{-i} = \mathbf{p}_{-i}|\hat{a}_i$ for all $\hat{a}_i \in \mathcal{A}_i$.

The *support* of $\mathbf{p}_i \in \mathcal{P}(\mathcal{A}_i)$ is defined by $\text{supp}(\mathbf{p}_i) \triangleq \{a_i \in \mathcal{A}_i : \mathbf{p}_i(a_i) > 0\}$. Also define $\Delta(\mathbf{p}_i) \triangleq \{\hat{\mathbf{p}}_i \in \mathcal{P}(\mathcal{A}_i) : \text{supp}(\hat{\mathbf{p}}_i) = \text{supp}(\mathbf{p}_i)\}$ to be set of all distributions in $\mathcal{P}(\mathcal{A}_i)$ which have the same support as \mathbf{p}_i .

For convenience, we will sometimes write a_i to denote the degenerate distribution in $\mathcal{P}(\mathcal{A}_i)$ which is concentrated on the single action $a_i \in \mathcal{A}_i$. Context will make clear whether a_i is intended to refer to an action or a distribution.

A real-valued function $f : \mathbb{R}^k \rightarrow \mathbb{R}$ is *convex* if for all $x, y \in \mathbb{R}^k$ and $\lambda \in [0, 1]$,

$$f(\lambda x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y).$$

A function f is *concave* if $-f$ is convex.

3 Preference Functions

Players in a game are usually assumed to be interested in maximizing their expected utility. In order to generalize this to other possible objectives, we allow the *preference function* $F_i : \mathcal{P}(\mathcal{A}_i) \times \mathcal{P}(\mathcal{A}_{-i}) \rightarrow \mathbb{R}$ for player i to be an arbitrary

trary continuous and bounded¹ function which encodes the objective of player i in the game. The preference function depends on the distributions from which the players draw their actions. For example, if player i 's action is drawn from distribution $\mathbf{p}_i \in \mathcal{P}(\mathcal{A}_i)$, and the other players' action profile is drawn from distribution $\tilde{\mathbf{p}}_{-i} \in \mathcal{P}(\mathcal{A}_{-i})$, and player i wishes to maximize expected utility, then player i 's preference function is

$$F_i(\mathbf{p}_i, \tilde{\mathbf{p}}_{-i}) = E_{a_i \sim \mathbf{p}_i, a_{-i} \sim \tilde{\mathbf{p}}_{-i}}[u_i(a_i, a_{-i})].$$

Note that F_i is defined for *all* $\mathbf{p}_i \in \mathcal{P}(\mathcal{A}_i)$ and $\tilde{\mathbf{p}}_{-i} \in \mathcal{P}(\mathcal{A}_{-i})$. In other words, the distributions which are arguments to a preference function can be completely unrelated. However, as we will see in Section 4 when we define equilibria, we are usually interested in cases where they are linked in some way. For illustration, we have given examples below of preference functions that have widespread use in risk-sensitive optimization, particularly in financial markets. In these expressions, $\alpha > 0$ is a constant that controls the degree of risk-sensitivity, and for notational compactness we introduce the following definitions:

$$E_i(\mathbf{p}_i, \tilde{\mathbf{p}}_{-i}) \triangleq E_{a_i \sim \mathbf{p}_i, a_{-i} \sim \tilde{\mathbf{p}}_{-i}}[u_i(a_i, a_{-i})]$$

$$V_i(\mathbf{p}_i, \tilde{\mathbf{p}}_{-i}) \triangleq \text{Var}_{a_i \sim \mathbf{p}_i, a_{-i} \sim \tilde{\mathbf{p}}_{-i}}[u_i(a_i, a_{-i})]$$

Pref. Function	$F_i(\mathbf{p}_i, \tilde{\mathbf{p}}_{-i}) = \dots$
Markovitz (I)	$E_i(\mathbf{p}_i, \tilde{\mathbf{p}}_{-i}) - \alpha V_i(\mathbf{p}_i, \tilde{\mathbf{p}}_{-i})$
Markovitz (II)	$E_i(\mathbf{p}_i, \tilde{\mathbf{p}}_{-i}) - \alpha \sqrt{V_i(\mathbf{p}_i, \tilde{\mathbf{p}}_{-i})}$
Sharpe Ratio	$E_i(\mathbf{p}_i, \tilde{\mathbf{p}}_{-i}) / (1 + \sqrt{V_i(\mathbf{p}_i, \tilde{\mathbf{p}}_{-i})})$

The first two preference functions are based on the *Markovitz criterion* for portfolio optimization, while the *Sharpe ratio* is another widely-used criterion in portfolio theory.² Many other choices for financially-motivated preference functions are available, such as Roy's 'safety-first' criterion. We also note that our use of these functions is slightly atypical: Investors are usually interested in maximizing functions of their *rate of return*, a quantity that is related to, but technically different from, utility.

3.1 Mean-Variance Preference Functions

Rather than proving our results for specific preference functions, we will prove them for a class of risk-sensitive preference functions which subsume the examples given above.

Definition 1. F_i is a mean-variance preference function if

1. $F_i(\mathbf{p}_i, \tilde{\mathbf{p}}_{-i}) = G_i(E_i(\mathbf{p}_i, \tilde{\mathbf{p}}_{-i}), V_i(\mathbf{p}_i, \tilde{\mathbf{p}}_{-i}))$ for some function G_i that is nondecreasing in its first argument.
2. F_i is convex in its first argument.

¹The continuity and boundedness assumptions for preference functions are assumed throughout the paper, and for brevity will not be repeated.

²We have introduced the constant 1 in the denominator of the Sharpe ratio only to ensure that it is bounded.

3. For any nonempty convex subset $P \subseteq \mathcal{P}(\mathcal{A}_i)$ and distribution $\tilde{\mathbf{p}}_{-i} \in \mathcal{P}(\mathcal{A}_{-i})$, if $F_i(\cdot, \tilde{\mathbf{p}}_{-i})$ is constant on P , then both $E_i(\cdot, \tilde{\mathbf{p}}_{-i})$ and $V_i(\cdot, \tilde{\mathbf{p}}_{-i})$ are constant on P .

The second property is consistent with our desire that F_i encode a sensitivity to risk — if F_i is convex in its first argument, then other things being equal, this implies that player i dislikes randomization (recall our comment to this effect in Section 1).

The third property says that whenever a mean-variance preference function is constant (with respect to its first argument) on some convex set, then expected utility and variance of utility are also constant on that same set. While this property may seem harder to justify, it is the case that all the examples of risk-sensitive functions we gave above satisfy all the conditions of the definition of a mean-variance preference function. Due to lack of space, we will only provide a proof for the case of the Markovitz (I) preference function; the derivation for the other functions is similar.

Claim 1. *The Markovitz (I) preference function is a mean-variance preference function.*

Proof. Let F_i be the Markovitz (I) preference function. Fix any $\tilde{\mathbf{p}}_{-i} \in \mathcal{P}(\mathcal{A}_{-i})$ and nonempty convex subset $P \subseteq \mathcal{P}(\mathcal{A}_i)$. Also choose $\mathbf{p}_i^1, \mathbf{p}_i^2 \in P$, and let $\mathbf{p}_i^\lambda = \lambda \mathbf{p}_i^1 + (1 - \lambda) \mathbf{p}_i^2$ be a point on the line segment connecting \mathbf{p}_i^1 and \mathbf{p}_i^2 , for some $\lambda \in [0, 1]$. We will only be concerned with the behavior of $F_i(\cdot, \tilde{\mathbf{p}}_{-i})$ on this line segment, so we overload notation and define $F_i(\lambda) = F_i(\mathbf{p}_i^\lambda, \tilde{\mathbf{p}}_{-i})$ and $E_i(\lambda) = E_i(\mathbf{p}_i^\lambda, \tilde{\mathbf{p}}_{-i})$ and $V_i(\lambda) = V_i(\mathbf{p}_i^\lambda, \tilde{\mathbf{p}}_{-i})$.

We now prove that $F_i(\cdot, \tilde{\mathbf{p}}_{-i})$ is convex in its first argument. Since a function is convex if and only if it is convex on every line segment in its domain, and \mathbf{p}_i^1 and \mathbf{p}_i^2 were chosen arbitrarily, it suffices to prove that $F_i(\lambda)$ is convex on the interval $[0, 1]$. A straightforward calculation shows that

$$\frac{d^2 F_i}{d\lambda^2} = 2(E_i(0) - E_i(1))^2 \quad (1)$$

for $\lambda \in (0, 1)$. This quantity is nonnegative, implying that $F_i(\lambda)$ is a convex function on $(0, 1)$, and by continuity $F_i(\lambda)$ is convex on $[0, 1]$.

Now suppose that $F_i(\cdot, \tilde{\mathbf{p}}_{-i})$ is constant on P . We wish to prove that both $E_i(\cdot, \tilde{\mathbf{p}}_{-i})$ and $V_i(\cdot, \tilde{\mathbf{p}}_{-i})$ are constant on P . Again, since \mathbf{p}_i^1 and \mathbf{p}_i^2 were chosen arbitrarily, it suffices to show that $E_i(\lambda)$ and $V_i(\lambda)$ are constant on the interval $[0, 1]$.

Since $F_i(\lambda)$ is constant on the interval $[0, 1]$, the expression in Eq. (1) must be equal to zero, which implies that $E_i(0) = E_i(1)$. Because $E_i(\lambda)$ is a linear function, this means that $E_i(\lambda)$ is constant on $[0, 1]$. Now, by examining the definition of the Markovitz (I) preference function, we see that if both $F_i(\lambda)$ and $E_i(\lambda)$ are constant on some interval, then $V_i(\lambda)$ must be as well. \square

4 Equilibrium Concepts

Usually, the definitions of Nash and correlated equilibrium assume that each player wishes to maximize expected utility, but these definitions can be easily generalized to admit arbitrary preference functions.

The action profile distribution $\tilde{\mathbf{p}} \in \mathcal{P}(\mathcal{A})$ is an *F-Nash equilibrium* (*F-NE*) if $\tilde{\mathbf{p}}$ is a product distribution and if for all players i

$$\tilde{\mathbf{p}}_i \in \arg \max_{\mathbf{p}_i \in \mathcal{P}(\mathcal{A}_i)} F_i(\mathbf{p}_i, \tilde{\mathbf{p}}_{-i}).$$

Similarly, an action profile distribution $\tilde{\mathbf{p}} \in \mathcal{P}(\mathcal{A})$ is an *F-correlated equilibrium* (*F-CE*) if for all players i and actions $a_i \in \mathcal{A}_i$

$$a_i \in \arg \max_{\mathbf{p}_i \in \mathcal{P}(\mathcal{A}_i)} F_i(\mathbf{p}_i, \tilde{\mathbf{p}}_{-i} | a_i).$$

These definitions are generalizations in the following sense: If each F_i is the expected utility preference function, then we recover the usual definitions of Nash and correlated equilibrium.

We introduce additional terminology to distinguish interesting cases of equilibria. If each player is using a mean-variance preference function, we will refer to an *F-NE* and *F-CE* as an *MV-Nash equilibrium* (*MV-NE*) and *MV-correlated equilibrium* (*MV-CE*), respectively. Also, if $\tilde{\mathbf{p}}$ is an *F-NE*, we say it is a *non-pure equilibrium* if at least one player in $\tilde{\mathbf{p}}$ is using a non-pure strategy.

Having defined *F-NE* and *F-CE* formally, let us discuss how these concepts differ with respect to the source of randomization. We momentarily set aside the possibility that an *F-CE* may induce correlations among the behavior of the players, and consider a *product* distribution $\tilde{\mathbf{p}}$. By examining the definitions above, we see that if $\tilde{\mathbf{p}}$ is to be an *F-NE*, then each player i must prefer the strategy \mathbf{p}_i at least as much as any action. On the other hand, if $\tilde{\mathbf{p}}$ is to be an *F-CE*, then each player i must prefer an action drawn from \mathbf{p}_i at least as much as any action. As we stated in Section 1, these are equivalent statements when the preference functions are expected utility. But as we will see in the rest of the paper, this is emphatically not the case more generally.

5 Sparsity of Mean-Variance Nash Equilibria

In this section, we prove our first main result. We show that non-pure MV-Nash equilibria fail to exist in all two-player games, except in degenerate cases. We characterize a “degenerate” game in a probabilistic fashion, by showing that if all the utility values in a two-player game are chosen randomly and independently, then the probability that a non-pure MV-Nash equilibrium exists is zero. Intuitively, it is not surprising that non-pure MV-NE are so rare. We might even guess *a priori* that a player using a mean-variance preference function will generally not prefer to choose her actions randomly, since this will tend to increase her variance. In fact, if a player is randomizing her choice of action in an MV-NE, we show that the variance experienced by that player must already be “saturated” due to the behavior of the other players. Moreover, we prove that this saturation is essentially a degenerate condition. The combination of these facts is what makes non-pure MV-NE so rare.

We begin the proof of our main result with the following lemma, which shows that whenever a preference function is convex in its first argument and maximized at a non-pure strategy, the preference function must have the same value

for all strategies which share the support of the maximizing strategy.

Lemma 1. *If a preference function F_i is convex in its first argument, and $\tilde{\mathbf{p}}$ is an *F-NE*, then $F_i(\cdot, \tilde{\mathbf{p}}_{-i})$ is constant on $\Delta(\tilde{\mathbf{p}}_i)$.*

Proof. The lemma holds trivially if $\tilde{\mathbf{p}}_i$ is pure, so suppose $\tilde{\mathbf{p}}_i$ is not pure. For shorthand, let $f_i(\cdot) = F_i(\cdot, \tilde{\mathbf{p}}_{-i})$. Because $\Delta(\tilde{\mathbf{p}}_i)$ is a convex set, we know that f_i is convex over $\Delta(\tilde{\mathbf{p}}_i)$. Since $\tilde{\mathbf{p}}$ is an *F-NE*, we also know that $\tilde{\mathbf{p}}_i$ is a maximum of f_i on $\Delta(\tilde{\mathbf{p}}_i)$, and that $\tilde{\mathbf{p}}_i$ is in the interior of $\Delta(\tilde{\mathbf{p}}_i)$, by definition. Since $\tilde{\mathbf{p}}_i$ is a maximum of f_i on $\Delta(\tilde{\mathbf{p}}_i)$, the gradient of f_i must vanish at $\tilde{\mathbf{p}}_i$. And since f_i is convex, this implies that $\tilde{\mathbf{p}}_i$ is also a minimum of f_i on $\Delta(\tilde{\mathbf{p}}_i)$. This can only happen if f_i is constant on $\Delta(\tilde{\mathbf{p}}_i)$. \square

Lemma 1 has the following implication: Suppose each F_i is a mean-variance preference function. By Definition 1, if F_i is constant on some convex set (with respect to its first argument), then the variance of utility is constant on that same set. So Lemma 1 says that a player in an MV-NE uses a non-pure strategy only if randomization doesn’t add to variance, i.e., the variance is already “saturated” for her by the other players.

We are now ready to prove the sparsity of non-pure MV-NE in two-player games.

Theorem 1. *Consider a two-player game where, for each player $i \in \{1, 2\}$ and action profile $a \in \mathcal{A}$, the utility $u_i(a) \in \mathbb{R}$ is drawn i.i.d. from an absolutely continuous distribution (with respect to Lebesgue). Then with probability 1 the game does not have a non-pure MV-NE.*

Proof. Let $\tilde{\mathbf{p}} = (\tilde{\mathbf{p}}_1, \tilde{\mathbf{p}}_2)$ be a strategy profile that is a non-pure MV-NE. Without loss of generality, assume that $|\text{supp}(\tilde{\mathbf{p}}_1)| \geq |\text{supp}(\tilde{\mathbf{p}}_2)|$. Let $k = |\text{supp}(\tilde{\mathbf{p}}_1)|$, and note that we must have $k > 1$, or else $\tilde{\mathbf{p}}$ would be a pure strategy profile. By Lemma 1, there is a constant C such that $F_1(\mathbf{p}_1, \tilde{\mathbf{p}}_2) = C$ for all $\mathbf{p}_1 \in \Delta(\tilde{\mathbf{p}}_1)$. Therefore, by Definition 1, we have that for all actions $a_1, a'_1 \in \text{supp}(\tilde{\mathbf{p}}_1)$

$$E_{a_2 \sim \tilde{\mathbf{p}}_2}[u(a_1, a_2)] - E_{a_2 \sim \tilde{\mathbf{p}}_2}[u(a'_1, a_2)] = 0 \quad (2)$$

$$\text{Var}_{a_2 \sim \tilde{\mathbf{p}}_2}[u(a_1, a_2)] - \text{Var}_{a_2 \sim \tilde{\mathbf{p}}_2}[u(a'_1, a_2)] = 0 \quad (3)$$

By the definition of variance

$$\begin{aligned} \text{Var}_{a_2 \sim \tilde{\mathbf{p}}_2}[u(a_1, a_2)] &= \\ E_{a_2 \sim \tilde{\mathbf{p}}_2}[u(a_1, a_2)^2] - E_{a_2 \sim \tilde{\mathbf{p}}_2}[u(a_1, a_2)]^2 \end{aligned} \quad (4)$$

and therefore Eq. (2)-(4) together imply

$$E_{a_2 \sim \tilde{\mathbf{p}}_2}[u(a_1, a_2)^2] - E_{a_2 \sim \tilde{\mathbf{p}}_2}[u(a'_1, a_2)^2] = 0. \quad (5)$$

for all actions $a_1, a'_1 \in \text{supp}(\tilde{\mathbf{p}}_1)$.

The rest of the proof will be an application of the following well-known mathematical facts:

1. If the entries of a matrix $M \in \mathbb{R}^{(k-1) \times k}$ are drawn i.i.d. from an absolutely continuous distribution, then for any fixed vector $c \in \mathbb{R}^k$, with probability 1, the rows of M are linearly independent, and furthermore, the vector c doesn’t belong to the linear span of the rows of M .

2. For any multivariate polynomial $P(x_1, \dots, x_k)$ that is not identically zero, if each x_i is drawn i.i.d. from an absolutely continuous distribution, then $P(x_1, \dots, x_k) \neq 0$ with probability 1 (the set of roots of P is an algebraic variety and therefore has measure zero under a product distribution of absolutely continuous random variables).

We now show that these two facts imply that, with probability 1, Eq. (2) and Eq. (5) are not true simultaneously. Note that Eq. (2) and Eq. (5) each specify $k - 1$ equations. Via a suitable renaming of variables, the i th equation specified by Eq. (2) has the form $\sum_{j=1}^k \lambda_j x_{ij} = 0$, and the i th equation specified by Eq. (5) has the form $\sum_{j=1}^k \lambda_j y_{ij} = 0$, where each $x_{ij}, y_{ij} \in \mathbb{R}$ is drawn i.i.d. from an absolutely continuous distribution, and $\sum_{j=1}^k \lambda_j = 1$.

In other words, we have $X\lambda = b$ and $Y\lambda = b$, where $X, Y \in \mathbb{R}^{k \times k}$, the last row of both X and Y is the all-ones vector, and $b = (0, 0, 0, \dots, 0, 1)$. Based on the first fact above, both X and Y are invertible with probability 1, so we have $\lambda = X^{-1}b = Y^{-1}b$. Now, based on Cramer's rule, we get that $\frac{\det(X_b)}{\det(X)} = \frac{\det(Y_b)}{\det(Y)}$, where X_b (resp. Y_b) is the matrix X (resp. Y) when we replace its first column by b . By simple algebra this is equivalent to demanding that $\det(X) \det(Y_b) - \det(Y) \det(X_b) = 0$. Notice that the left-hand side of this equation is a multivariate polynomial in the x_{ij} 's and y_{ij} 's. It is easy to show that this polynomial is not zero for at least one realization of the variables, and thus the polynomial is not identically zero. Therefore, by the second fact above, this equation is not satisfied with probability 1. \square

The previous theorem essentially rules out non-pure MV-NE, but says nothing about pure MV-NE. So how common are pure MV-NE? Not any more common than pure Nash equilibria, by the following theorem, whose proof is omitted as it is entirely straightforward.

Theorem 2. *If $\tilde{\mathbf{p}}$ is a pure MV-NE, then $\tilde{\mathbf{p}}$ is a pure Nash equilibrium.*

Summing up, we see that there are essentially two kinds of MV-NE in two-player games: Those that correspond to pure Nash equilibria, and degenerate cases. We note that this conclusion is particularly unexpected in the case of two-player zero-sum games. In a zero-sum game, one player's utility is the negative of the other player's utility. Therefore, the variance of utility is always the *same* for both players. It is counterintuitive, but nonetheless true, that adding the same term to both players' preference functions destroys nearly all the equilibria.

The preceding analysis suggests that non-pure MV-NE might not *ever* exist. Below we give an example proving otherwise: the well-known zero-sum game 'Matching Pennies'.

	H	T
H	+1, -1	-1, +1
T	-1, +1	+1, -1

The unique Nash equilibrium of this game is a strategy profile in which each player plays each action with equal

probability. It is easy to check that this strategy profile is also an MV-NE when each F_i is the Markovitz (I) preference function given in Section 3. In fact, under this strategy profile, each player experiences a variance of 1, which is the largest possible value for variance in this game — so here we have an example of the “saturation” property discussed earlier.

6 Existence of Mean-Variance CE

In the previous section, we showed that non-pure MV-Nash equilibria are extremely uncommon in two-player games. The proof hinged on the combination of two facts: (1) a mean-variance preference function is convex in its first argument, which means, roughly, that it penalizes players who play an action randomly; (2) in a non-pure MV-NE, at least one player prefers randomizing over her actions at least as much as playing any single action deterministically.

In this section, we prove that, in striking contrast to MV-Nash equilibria, MV-correlated equilibria always exist. Interestingly, the reason is *not* that an MV-CE allows correlations among the players' action choices (although this does have other advantages, as we will explain at the end of this section). Intuitively, the reason is that an MV-CE does not require each player to perform her own randomization. Instead, a third party is responsible for choosing an action randomly, and the players only need to prefer the action chosen for them at least as much as any other action. This condition is substantially easier to meet — essentially, the players themselves do not pay a penalty for introducing variance, because the “variance lies elsewhere”. Indeed, we prove that an F -CE exists whenever each preference function F_i is convex in its first argument. The F -CE concept is a strict generalization of a closely related concept in the economics literature known as an *equilibrium in beliefs*, and the existence proof follows immediately from this relationship. This proof was first discovered by (Crawford 1990), but it is simple and useful for our exposition, so we include it for completeness.

We begin by stating a well-known extension of Nash's original result (Nash 1950) on the existence of Nash equilibria.

Theorem 3 (Nash (1950); Debreu (1952)). *If each preference function F_i is linear in its first argument, then an F -NE exists.*

Note that Theorem 3 requires that the preference functions be linear in their first argument, while the mean-variance preference functions we described in Section 3 are convex in their first argument. Nonetheless, we are able to apply Theorem 3 by ‘linearizing’ each convex preference function, and then observing that the two versions of each preference function agree on any pure strategy.

Theorem 4. *If each preference function F_i is convex in its first argument, then an F -CE equilibrium exists.*

Proof. Define the linearization \bar{F}_i of F_i to be the following:

$$\bar{F}_i(\mathbf{p}_i, \mathbf{p}_{-i}) = \sum_{a_i \in \mathcal{A}_i} \mathbf{p}_i(a_i) F_i(a_i, \mathbf{p}_{-i}) \quad (6)$$

Clearly \bar{F}_i is linear in its first argument. By Theorem 3, there exists a product distribution $\tilde{\mathbf{p}} \in \mathcal{P}(\mathcal{A})$ such that for all players i

$$\tilde{\mathbf{p}}_i \in \arg \max_{\mathbf{q}_i \in \mathcal{P}(\mathcal{A}_i)} \bar{F}_i(\mathbf{q}_i, \tilde{\mathbf{p}}_{-i}) \quad (7)$$

Moreover, by Eq. (7) and the linearity of \bar{F}_i in its first argument, for any player i and actions $a_i, a'_i \in \text{supp}(\tilde{\mathbf{p}}_i)$, we must have

$$F_i(a_i, \tilde{\mathbf{p}}_{-i}) = F_i(a'_i, \tilde{\mathbf{p}}_{-i}) \quad (8)$$

Now fix any player i , action $a_i \in \mathcal{A}_i$ such that $\tilde{\mathbf{p}}_i(a_i) > 0$, and $\mathbf{q}_i \in \mathcal{P}(\mathcal{A}_i)$, and consider

$$\begin{aligned} F_i(a_i, \tilde{\mathbf{p}}_{-i}|a_i) &= F_i(a_i, \tilde{\mathbf{p}}_{-i}) \\ &= \bar{F}_i(a_i, \tilde{\mathbf{p}}_{-i}) \\ &= \bar{F}_i(\mathbf{q}_i, \tilde{\mathbf{p}}_{-i}) \\ &\geq \bar{F}_i(\mathbf{q}_i, \tilde{\mathbf{p}}_{-i}) \\ &\geq F_i(\mathbf{q}_i, \tilde{\mathbf{p}}_{-i}) \\ &= F_i(\mathbf{q}_i, \tilde{\mathbf{p}}_{-i}|a_i) \end{aligned}$$

where we used, in order: $\tilde{\mathbf{p}}$ is a product distribution; Eq. (6); Eq. (8); Eq. (7); convexity of F_i in its first argument; $\tilde{\mathbf{p}}$ is a product distribution.

Comparing the first and last line in the chain above proves that $\tilde{\mathbf{p}}$ is an F -CE. \square

Interestingly, although the definition of an F -CE permits correlations among the players' actions, note that the proof of Theorem 4 does *not* imply that such an equilibrium can exist. It only establishes the existence of F -CE that are product distributions (note that, in general, an F -CE can be a product distribution without being an F -NE, unlike the situation for CE and NE).

We now discuss an example which illustrates that a F -CE need not be a product distribution. Moreover, our example will show that it is possible for all players in a game to benefit from correlating their actions. Taken together, these facts demonstrate that an F -CE is a strictly more general — and potentially more useful — equilibrium concept than an equilibrium in beliefs.

Consider the well-known 'Chicken' game:

	C	D
C	6, 6	2, 7
D	7, 2	0, 0

This game has the following interpretation: The players are driving two cars which are headed towards each other. If a player swerves, she is a 'chicken'; otherwise she 'dares'. The best outcome for a player is to dare while the other player chickens. If both players dare, they collide head-on.

This game has three Nash equilibria, including two pure Nash equilibria. The pure Nash equilibria occur when one player 'dares' and the other 'chickens'. In the non-pure Nash equilibrium, each player 'dares' with probability $1/3$. For ease of comparison, let us write out the distribution on action profiles induced by this non-pure equilibrium:

$$\begin{aligned} \tilde{\mathbf{p}}^{NE}(C, C) &= 4/9 & \tilde{\mathbf{p}}^{NE}(C, D) &= 2/9 \\ \tilde{\mathbf{p}}^{NE}(D, C) &= 2/9 & \tilde{\mathbf{p}}^{NE}(D, D) &= 1/9 \end{aligned}$$

Now suppose each player is using the Markovitz (I) preference function. It can be shown that this game has no non-pure MV-NE (indeed, even if it did, we know from our results in Section 5 that a small perturbation of the utilities would cause all non-pure MV-NE to disappear).

For a suitable choice of the risk parameter α , a product distribution that is an MV-CE for this game is one in which each player 'dares' with probability $1/4$:

$$\begin{aligned} \tilde{\mathbf{p}}^{MV-CE-1}(C, C) &= 9/16 & \tilde{\mathbf{p}}^{MV-CE-1}(C, D) &= 3/16 \\ \tilde{\mathbf{p}}^{MV-CE-1}(D, C) &= 3/16 & \tilde{\mathbf{p}}^{MV-CE-1}(D, D) &= 1/16 \end{aligned}$$

Note that each player places less weight on 'dare' in this equilibrium, because doing so is helpful for reducing the variance they experience. However, the Markovitz (I) objective can be improved even more by correlating the players' actions to ensure that both players never dare simultaneously:

$$\begin{aligned} \tilde{\mathbf{p}}^{MV-CE-2}(C, C) &= 3/5 & \tilde{\mathbf{p}}^{MV-CE-2}(C, D) &= 1/5 \\ \tilde{\mathbf{p}}^{MV-CE-2}(D, C) &= 1/5 & \tilde{\mathbf{p}}^{MV-CE-2}(D, D) &= 0/5 \end{aligned}$$

7 Conclusion and Future Work

We have studied the existence of risk-sensitive generalizations of Nash and correlated equilibria. In marked contrast to classical results, we have shown risk-sensitive Nash equilibria seldom exist, while risk-sensitive correlated equilibria always do. We argued that this dichotomy is due to the differing sources of randomization for each type of equilibrium.

While we have shown that risk-sensitive Nash equilibria seldom exist, we have not yet ruled out the possibility that an *approximate* equilibrium always exists; we leave this for future work. Also, we speculate that it may be possible to compute a risk-sensitive correlated equilibrium in polynomial time.

8 Acknowledgments

We thank Shreyas Sundaram for his ideas and advice, and the anonymous reviewers for their helpful comments.

References

- Aumann, R. J. 1987. Correlated equilibrium as an expression of bayesian rationality. *Econometrica* 55(1):1–18.
- Crawford, V. P. 1990. Equilibrium without independence. *Journal of Economic Theory* 50(1):127–154.
- Debreu, G. 1952. A social equilibrium existence theorem. *Proc. of the National Academy of Science* 38:886–893.
- Dekel, E.; Safra, Z.; and Segal, U. 1991. Existence and dynamic consistency of Nash equilibrium with non-expected utility preferences. *Journal of Economic Theory* 55(2):229–246.
- Nash, J. F. 1950. Equilibrium points in n-person games. In *Proc. of the National Academy of Sciences*.
- Nowak, A. 2005. Notes on risk-sensitive Nash equilibria. *Annals of International Society of Dynamic Games* 7:95–109.