When the Players Are Not Expectation Maximizers

Amos Fiat[1] and Christos Papadimitriou[2]

[1] School of Computer Science
Tel Aviv University
fiat@tau.ac.il
[2] Computer Science Division
University of California at Berkeley
christos@cs.berkeley.edu

Abstract. Much of Game Theory, including the Nash equilibrium concept, is based on the assumption that players are expectation maximizers. It is known that if players are risk averse, games may no longer have Nash equilibria ([11,6]. We show that

1. Under risk aversion (convex risk valuations), and for almost all games, there are no mixed Nash equilibria, and thus either there is a pure equilibrium or there are no equilibria at all, and,
2. For a variety of important valuations other than expectation, it is NP-complete to determine if games between such players have a Nash equilibrium.

1 Introduction

In 1950 John Nash proved that every game has a mixed equilibrium. Myerson [17] gives a plethora of reasons as to why Nash's theorem (and his proposed framework of rationality in normal form games) underlies the foundations of modern economic thought. In recent years a computationally inspired challenge to the concept of mixed Nash equilibrium has arisen, see, e.g., [7,4], and the universality of the concept has become questionable in face of intractability results. In this paper we pursue another line of critique of the Nash equilibrium. In particular, we show that Nash's Theorem *does not hold* if the players are not expectation maximizers, in that almost all games fail to have a mixed Nash equilibrium (Theorem 3), and that it is NP-hard[1] to tell those that do from those that do not (Theorem 5). To understand our results in context, we begin by reviewing the rich literature on risk in Economics.

1.1 A Brief History of Risk

In many ways, risk is a defining characteristic of the modern world, and the analytical problems associated with it were pointed out early. Nicolas Bernoulli

[1] NP-hardness is stronger evidence of intractability than the PPAD-completeness of finding a Nash Equilibrium ([7,4]).

S. Kontogiannis, E. Koutsoupias, P.G. Spirakis (Eds.): SAGT 2010, LNCS 6386, pp. 1–14, 2010.

(1687 – 1759) posed the famous "St. Petersburg Paradox" [22], exposing the inadequacy of expectation in decision-making, and some decades later his nephew Daniel Bernoulli (1700 – 1782) [2] provided a solution by proposing to distinguish between money and the *utility of money*, and to model risk aversion by a utility function that is concave.

Two centuries later, Emil Borel [3] and John von Neumann [19] initiated the study of strategic behavior, and, two dacades later, von Neumann and Oskar Morgenstern published their "Theory of Games and Economic Behavior" [20], where they expounded their Expected Utility Theory (EUT). They postulated that the risk behavior of an agent can be modeled as a (risk) *valuation* V mapping lotteries (distributions over the reals) to the reals[2]. If this valuation satisfies some plausible axioms equivalent to *linearity*, then the agent's behavior can be captured by a utility function, and the agent behaves as a maximizer of the expectation of his utility. A few years later, John Nash extended the work of von Neumann and Morgenstern to non-zero sum non-cooperative games [18], and showed that a mixed strategy equilibrium always exists; note that Nash's Theorem is stated in the context of EUT. EUT can capture both risk-seeking and risk-averting agent behavior by having a valuation function that is convex or concave, respectively[3].

In 1948, Friedman and Savage [8] attempted to deal with criticism of expected utility and considered models where utility is either a concave or convex function of money. Portfolio theory, developed in much more empirical and less principled/axiomatic manner from 1950 onwards by Marschak [16], Markowitz [15,14] and many others [24,27], considered valuations (functions from distributions to the reals) of the form "expectation minus variance" or "expectation minus standard deviation" as a model of agent behavior in the face of financial risk. E.g., the optimal portfolio for a given expected value is the one with minimum variance. There is no way to cast such behavior within the framework of Expected Utility Theory.

Independently, in 1951 Maurice Allais [1] suggested that there are problems with "the American School", *i.e.*, he raised issues with the von Neumann-Morgenstern EUT. One of his examples was indeed the "expectation minus variance" valuation, but he also gave other empirical arguments (the *Allais paradox*) strongly suggesting that real human behavior cannot be modeled within EUT (and, consequently, certainly not by assuming that agents are expectation maximizers).

After Allais, many non-EUT valuations were proposed to address problems such as the Allais paradox, see the expository articles [26,12,13]. One such model

[2] Actually, instead of risk valuations, von Neumann and Morgenstern postulated, and axiomatized, *preferences* between any two distributions; but the valuation formalism used in this paper is essentially equivalent.

[3] There is a point of possible confusion here. By "concave" valuation we mean a function from lotteries to the reals that is concave *in the probabilities*, and *not* in the values. For example, the variance is concave in the probabilities but convex in the values.

in wide use today is *prospect theory* due to Kahneman and Tversky [10]. Prospect theory predicts that a loss of x is much more painful than a gain of x is pleasant, and, importantly, that probabilities undergo subjective modifications in the agents' calculation of expectations, not unlike our example valuations V^5 and V^6 below.

More directly related to our work, Ritzberger [23] showed that for expected utilities with rank dependent probabilities reflected risk aversion, mixed Nash Equilibria will disappear. Chen and Neilson [9] considered the flip side of this phenomena and gave conditions under which a pure strategy must exist (but this requires a compact set of pure strategies).

1.2 Nash Equilibria and Risk

Whereas EUT was proposed by von Neumann and Morgenstern as a preamble to their theory of games and strategic behavior, non-EUT approaches to risk were primarily considered in non-strategic settings such as finance. The question of how non-EUT valuations impact non-cooperative game theory and Nash equilibria in particular was raised only in the 1990s by Crawford [6], who noted that, while Nash's theory holds when the agents' valuations are concave (see our Theorem 1, stated and proved here for completeness and computational emphasis), there are simple games, such as the 2×2 zero-sum game shown in (1) below, that have no Nash equilibria if the agents have convex risk valuations.

To understand the broad range of possible attitudes of strategic agents towards risk, consider the following six valuations (functions mapping distributions to the reals) modeling plausible attitudes of agents towards risk:

- V^1 If an agent is a *pure expectation maximizer*, then his valuation V^1 maps any distribution to its expected value. This is the framework used in virtually all of Game Theory.
- V^2 Most people are *risk averse.* One way to capture this would be valuation V^2, which assigns to each distribution over the reals the expectation minus the variance; this was proposed by Marschak, Markowitz, Allais, and others. We use V_2 as an exemplar of risk averse valuations; there are many variants of V^2 in which one subtracts from the expectation the standard deviation or some other increasing function of the variance, or a small multiple thereof.
- V^3 Some agents may be *risk-seeking;* for example, valuation V^3 evaluates a distribution by its expectation *plus* (an increasing function of) the variance.
- $V^4[\theta]$ An agent may be facing a costly life-saving medical procedure and his only interest in the game is to maximize the probability that the payoff is above the cost θ. This defines valuation $V^4[\theta]$.
- $V^5[p]$ Another agent may be interested in maximizing her "almost certain bottom line:" the amount of money she gets with probability at least $p = .95$, say. Let us call this valuation $V^5[p]$.
- V^6 Finally, somebody else evaluates any discrete distribution over the reals by the average between the maximum and the minimum value which occur with a nonzero probability. We call this V^6.

Note: We are not proposing these six valuations as the only possible attitudes toward risk, or even as plausible or reasonable ones; they are here only to demonstrate the range of possibilities and fix ideas. Our three results hold for very broad classes of such valuations, delimited in their statement or the discussion following the proof. Also, the last two valuations fall into a very important class proposed in Kahneman and Tversky's Prospect Theory [10], in which expected utility is maximized, albeit with the probabilities modified. Briefly, in Prospect Theory valuations are of the form $V = \sum_i u(x_i)\pi_i$, where u is an ordinary utility function, but the π_i's are *modified probabilities.* The modification is done through an increasing function $G : [0,1] \mapsto [0,1]$ with $G(0) = 0$ and $G(1) = 1$ that *modifies the cumulative probabilities.* That is, if we assume that $x_1 \leq x_2, \leq \cdots \leq x_n$, π_i is defined as $G(\sum_{j=1}^{i} p_i) - G(\sum_{j=1}^{i-1} p_i)$. It is easy to see that $V^5[p]$ corresponds to the modifier function $G(x) = 0$ if $x \leq 1-p$, and $G(x) = 1$ otherwise. And V^6 corresponds to $G(x) = \frac{1}{2}$ for $0 < x < 1$. Kahneman and Tversky speculate that "real" modifier functions, consistent with experiments, are steeply increasing at 0 and at 1, go through the $(\frac{1}{2}, \frac{1}{2})$ point, and are flat around it. The effect is that small probabilities of extreme payoffs are exaggerated. Notice that our function G defining V^6 is a stylized and exaggerated function of this form ($\pi_1 = \pi_n = \frac{1}{2}$).

Of these six risk valuations, V^1 is the one considered throughout Game Theory and, naturally, Nash's Theorem holds in it. Of the others, $V^4(\theta)$ falls squarely within the purvey of EUT: Just map the agent's payoffs to zero if they are less than θ and to one otherwise, and solve the resulting game. As it turns out, Nash's Theorem is valid under the risk-seeking valuation V^3 as well (see Theorem 1 and Proposition 2); the reason is, V^3 is *concave* in the probabilities.

The other three valuations, however, break Nash's Theorem. For example, consider the following game proposed by Crawford, which we call Γ:

$1,-1$	$0,0$
$0,0$	$r,-r$

(1)

If the agents evaluate any distribution of payoffs by a convex valuation such as V^2, then Crawford observes that there are no Nash equilibria in this game, pure or mixed. This holds for $r \neq 1$; interestingly, if $r = 1$ then there is a mixed Nash equilibrium with both players V^2 (or even with one player with valuation V^2 and the other player is an expectation maximizer).

1.3 Our Results

1. In terms of ubiquity of the Nash Equilibria, we show that
   (a) *Almost all games have no mixed Nash equilibria if the players are risk averse* (Theorem 3). By "almost all" we mean that games that do have mixed Nash equilibria form a set of measure zero in the space of all games, with utilities drawn at random; for example, any such games must have equality between certain payoff values. Moreover, even ϵ-Nash equilibria will not exist. Pure Nash equilibria may still exist; but only a $1 - \frac{1}{e}$ fraction of games have them [25]. This was known for rank dependent expected utility functions (Ritzberger, 1996, [23]).

   (b) Even if the underlying game has a mixed Nash equilibrium, arbitrarily small random errors by the players in interpreting the payoff matrix will lead to instability with high probability (Observation 4).
2. Any given game may not have a Nash equilibrium. We show that
   (a) *It is NP-complete to determine if a two-person game with non-EUT player valuations has a Nash Equilibrium*[4]. We show this for functions such as V^2, $V^5[p]$, and V^6, and we lay out broad conditions on the risk valuations under which our proof works (Corollary 6).
   (b) In contrast, for concave valuations such as V^3 — Nash equilibria are guaranteed to exist and are PPAD-complete, *i.e.*, the same complexity as Nash Equilibria under the expected utility theory (Theorem 1).

1.4 The Model

To avoid confusion, we use the terminology of payoffs and valuations rather than utility. Under expected utility theory, our "payoffs" are considered utilities and some of our valuation functions are also utilities, whereas others cannot be so expressed.

A *k-player game* G, where $k > 1$, consists of k finite sets of *strategies* $S_1, \ldots, S_k$ and k *payoff functions* $p_1, \ldots, p_k$ mapping $S = \prod_i S_i$ to $\Re$. We denote by $\Delta[S_i]$ the set of *mixed strategies* for player i.

Given a k-tuple of mixed strategies $x = (x_1, \ldots, x_k) \in \prod_i \Delta[S_i]$, for any combination of pure strategies $s = (s_1, s_2, \ldots, s_k) \in \prod S_i$ define $q_x(s) = \prod_{i=1}^k x_i(s_i)$, where $x_i(s_i)$ is the probability player i plays pure strategy s_i. Every k-tuple of mixed strategies, x, defines a *strategy distribution*, $S(x)$, over $\prod S_i$, where the probability of $s \in \prod S_i$ being played is $q_x(x)$.

Let G be a k-player game, and $V_1, \ldots, V_k$ be valuations. Given a payoff function p_i, let $R_i = p_i(s), s \in \prod_i S_i$. R_i is the range of possible payoffs for agent i over all combinations of pure strategies.

Given a k-tuple of mixed strategies $x = (x_1, \ldots, x_k)$, for every player i the strategy-distribution $S(x)$ implies a *payoff-distribution*, $P_i(x)$, the support of $P_i(x)$ is a subset of R_i, and the probability of $a \in R_i$ is

$$q_{i,x}(a) = \sum_{s \in S | p_i(s) = a} q_x(s).$$

Again, for any $1 \leq i \leq k$ we have $\sum_{a \in R_i} q_{i,x}(a) = 1$

A *(risk) valuation* is any function from payoff distributions to the reals. The functions $V^1, V^2, V^3, V^4[\theta]$, $V^5[p]$, and V^6 mentioned in the introduction are indicative examples of important valuations in the literature. Ergo, if agent i has risk valuation V_i then the value of a k tuple of mixed strategies $x = (x_1, x_2, \ldots, x_k)$ to player i is $V_i(P_i(x))$.

[4] Our NP-completeness proof uses a new gadget, based on a generalized rock-paper-scissors game, which is arguably simpler than the construction in [5], we suspect that this gadget may prove useful in other contexts as well.

Fix i and the mixed strategies of all players but that of player i, $x_{-i} \in \prod_{j \neq i} \Delta(S_j)$. For fixed $x_{-i} \in \prod_{j \neq i} \Delta(S_j)$, $x_i \in \Delta(S_i)$, we use the shorthand notation

$$V^{x_{-i}}(x_i) = V(P_i(x_i, x_{-i})).$$

Finally, a k-tuple of mixed strategies $x = (x_1, \ldots, x_k) \in \prod_i \Delta[S_i]$ is called a $(V_1, \ldots, V_k)$-*Nash equilibrium* if for every player i and every mixed strategy $x'_i \in \Delta[S_i]$,

$$V_i^{x_{-i}}(x'_i) \leq V_i^{x_{-i}}(x_i).$$

That is, changing i's mixed strategy in any way, and keeping all other mixed strategies the same, results in a distribution over payoffs whose valuation, with respect to player i, is no better than the valuation of the current distribution.

Examples: Consider the valuations $V^1, \ldots, V^6$ introduced above, and let us concentrate on the risk averse valuation V_2, expectation minus variance. Analyzing mixed Nash equilibria when players behave this way (or in any of the other five V^i's save V^1, expectation) is very tricky. Pure strategies in a mixed Nash equilibrium are not necessarily individually best responses to the other players' strategies. Also, because of the nonlinearity of the valuations considered, games are not invariant under translation or scaling by positive constants.

Suppose that Crawford's game (see Equation (1)) is played by a row player with valuation V^2, and a column player who is an expectation maximizer (*i.e.*, the column player is risk neutral). If $r = 1$, then a mixed Nash equilibrium exists in which both players randomize uniformly. As we shall see, this situation is a singular exception (see the proof of Theorem 3). But suppose that $r > 1$. From the point of view of the expectation maximizer (column player), as the probability that the row player plays down is increased from zero, playing left is the best response, up to a point in which left and right — *and any mixture in between* — are at a tie. From then on, right is the best response. A similar behavior is observed of the V^2 (row) player — with an important difference. As the column player increases from zero the probability of playing right, the row player prefers up, and at some point there is a tie between up and down.

The catch is that, because of the convexity of V^2, the points in between are *not* best responses, and there is a discontinuous jump from up to down. As a result, the trajectories of the dynamics of the two players (the two best response maps) do not intersect, and the game has no equilibrium. The same behavior is observed if the row player's valuation is expectation minus any increasing function of the variance (say, a small multiple, or square root), and if both players are risk averse.

Intuitively, the reason that $r = 1$ is a singularity (and a mixed Nash Equilibrium does in fact exist) is because the payoff-distributions don't depend on the mixed strategy chosen by the row player. So, although the strategy-distributions are in fact different, the payoff-distributions are invariant to changes in the row strategy, and the valuation to the row player is therefore also invariant to changes in the row strategy. , *i.e.*, $V^{x_{-row}}(x_{row}) = V^{x_{-row}}(x'_{row})$, for any $x_{row} = (q, 1-q)$, $x'_{row} = (q', 1-q')$, $0 \leq q, q' \leq 1$.

2 Valuation Convexity and Concavity

Valuation functions are defined on payoff distributions. We require (and use) a limited form of convexity/concavity, that of valuation V_i convex (resp. concave) with respect to mixed strategies of player i, $x_i \in \Delta(S_i)$.

A valuation function V_i is said to be convex (resp. concave) with respect to x_i if for every $k-1$ tuple of mixed strategies $x_{-i} \in \prod_{j \neq i} \Delta(S_j)$, for any two mixed strategies $x_i, x_i' \in \Delta(S_i)$, and for any $0 \leq \alpha \leq 1$.

$$V^{x_{-i}}(\alpha x_i + (1-\alpha)x_i') \leq \alpha V^{x_{-i}}(x_i) + (1-\alpha)V^{x_{-i}}(x_i');$$
$$V^{x_{-i}}(\alpha x_i + (1-\alpha)x_i') \geq \alpha V^{x_{-i}}(x_i) + (1-\alpha)V^{x_{-i}}(x_i');$$

A valuation function V_i is said to be strictly convex, or strictly concave, respectively, if for any two payoff distributions $P_i(x_i, x_{-i}) \neq P_i(x_i', x_{-i})$, and for any $0 < \alpha < 1$,

$$V^{x_{-i}}(\alpha x_i + (1-\alpha)x_i') < \alpha V^{x_{-i}}(x_i) + (1-\alpha)V^{x_{-i}}(x_i');$$
$$V^{x_{-i}}(\alpha x_i + (1-\alpha)x_i') > \alpha V^{x_{-i}}(x_i) + (1-\alpha)V^{x_{-i}}(x_i');$$

We say that V_i is *efficiently concave* in $x_i \in \Delta(S_i)$, if, for any x_{-i}, and for any strictly concave polynomial $t(x_i)$, the (unique) point $\text{argmax}_{x_i} V^{x_{-i}}(x_i) + t(x_i)$ can be computed in time polynomial in $|S_i|$, the representation of t, the total number of bits in the coefficients of $V(x_i)$, and the number of bits of precision required.

3 Computing Nash Equilibria

We now give a sufficient condition on the V_i's for $(V_1, \ldots, V_k)$-Nash equilibria to exist and be as easy to compute as ordinary Nash equilibria: it suffices for each V_i to be efficiently concave in x_i. This result is well known to economists; for example, von Neumann stated his minmax Theorem in terms of concave functions. Here we restate and prove it for computational emphasis and contrast with our main result that follows.

Theorem 1. *Let G be a k player game and let $V_1, \ldots, V_k$ be valuations, where $k > 1$. If each V_i is concave in x_i, then any k-player game has an $(V_1, \ldots, V_k)$-Nash equilibrium. If in addition the V_i's are efficiently concave, then the problem of finding an an $(V_1, \ldots, V_k)$-Nash equilibrium is in PPAD.*

Proof. Define the following function ϕ from $\prod_i \Delta[S_i]$ to itself: $\phi(x_1, \ldots, x_k) = (y_1, \ldots, y_k)$, where for each i y_i is defined as follows:

$$y_i = \text{argmax}_{z_i \in \Delta(S_i)} V_i^{x_{-i}}(z_i) - ||z_i - x_i||^2.$$

Since V_i is concave in x_i, and $-||z_i - x_i||^2$ is a strictly concave polynomial in z_i, the argmax exists and is unique, and continuous as a function of x. Therefore ϕ is a continuous function from a compact convex set to itself, and so, by

Brouwer's fixpoint theorem, ϕ has a fixpoint. It is now easy to see, arguing by contradiction, that this fixpoint is an $(V_1, \ldots, V_k)$-Nash equilibrium of G. The computational elaboration of the theorem follows easily from the fact that approximating Brouwer fixpoints is in PPAD [21]. □

We next point out that three of our original six valuation examples fall into this benign category.

Proposition 2. *The valuations V^1, V^3, and $V^4[\theta]$ are efficiently concave.*

Proof. It is easy to see that V^1 and $V^4[\theta]$ are actually linear, and thus trivially efficiently concave. For V^3 we need to establish that the variance of the distribution $P_i(x)$ is concave in x_i. Recall that for any random variable X the variance is $E(X^2) - (E(X))^2$. Since the first term is linear in x_i, we concentrate on the second term. It is easy to see that, if X is the random variable for the payoff to player i, the Hessian of $-(E(X))^2$ with respect to x_i (variables $x_i[s], s \in S_i$) is

$$\frac{\partial^2[-(E(X))^2]}{\partial x_i[s_i]\partial x_i[s_i']} = -h_{s_i} \cdot h_{s_i'},$$

where $h_{s_i} = E_x[X|s_i]$, and similarly for $h_{s_i'}$. Now it is clear that the Hessian is the tensor product of vector $(h_{s_i} : s_i \in S_i)$ with itself, negated, and thus it is trivially a negative semi-definite matrix. Hence the variance is indeed a concave function of the probabilities, and so is risk seeking valuation V^3. That it is efficiently concave is straightforward. □

Note: There is a point of confusion here. It is well known, and often useful, that the variance is a *convex* function of the *values*. However, it turns out that it is also a concave function *of the probabilities*.

Notice that adding to the expectation, instead of the variance, a positive multiple, or any concave function, of the variance (such as the standard deviation), preserves the valuation's concavity. Hence the positive result stated in the theorem applies to a broad variety of risk-seeking valuations.

4 Games with No Equilibria

But what if the valuations are not concave — for example, convex like V^2? Risk averse agents typically have strictly convex valuations. We know from Crawford that Nash equilibria may not exist, but two questions come up immediately: How prevalent are such pathologies? And even if they are prevalent, can they at least be characterized and excluded? In this and the next section we answer both questions in the negative.

For a pure strategy $s \in S_i$ let x_i^s be the strategy for agent i that plays s deterministically. Call a game *in general position* if for any player i, for any mixed strategy of the other players x_{-i}, and for any two pure strategies $s, s' \in S_i$, the payoff distributions $P_i(x_i^s, x_{-i}) \neq P_i(x_i^{s'}, x_{-i})$. Note that $P_i(x_i^s, x_{-i}) \neq P_i(x_i^{s'}, x_{-i})$ does not imply that $V_i^{x_{-i}}(x_i^s) \neq V_i^{x_{-i}}(x_i^{s'})$.

Our first negative result states that, if the players are risk-averse, then a game in general position cannot have a mixed Nash equilibrium. This result has an evocative probabilistic interpretation: In the space of all games considered as tuples of tensors, games with mixed Nash equilibria form a set of measure zero, a lower-dimensional manifold. Since it is known that the probability that a game has a pure Nash equilibrium in this space is asymptotically $1 - \frac{1}{e}$, this means that about 37% of all games have no Nash equilibrium when the players are risk averse.

Theorem 3. *Let Γ be a game in general position, and suppose that player i's valuation is strictly convex (as a function of $x_i \in \Delta[S_i]$). Then in any Nash equilibrium this player plays a pure strategy. In particular, if all agent valuations are strictly convex, there is no mixed Nash equilibrium.*

Proof. Fix the (possibly mixed) strategies of all agents but i, x^{-i}. Assume that agent i has a strictly mixed strategy, which, along with all other agent strategies, is in Nash Equilibrium.

As x_i is a strictly mixed strategy, then for all $s \in S_i$, $x_i(s) < 1$, and for at least two distinct $s, s' \in S_i$, $x_i(s) > 0$ and $x_i(s') > 0$. It follows that

$$V_i^{x_{-i}}(x_i) = V_i^{x_{-i}}\left(\sum_{s\in S} x_i^s \cdot x_i(s)\right) < \sum_{s\in S_i} x_i(s) V_i^{x_{-i}}(x_i^s) \leq \max_{s\in S_i} V_i^{x_{-i}}(x_i^s).$$

This is in contradiction to the assumption that x_i was in Nash Equilibrium (alternately that $x_i = \text{argmax}_{z_i \in \Delta(S_i)} V_i^{x_{-i}}(z_i)$). □

Theorem 3 is in no way a converse of Theorem 1; there are many valuations that are neither concave nor strictly convex — valuations V^5 and V^6, for example. However, it is not hard to see that for these two there are games in which there are no Nash equilibria. Generalizing Theorem 2 so that it comes close to being a converse of Theorem 1 is an interesting open question.

It now remains to ask: when is a game in general position? A game is not in general position if for some player i, there exists a $k-1$ tuple of mixed strategies for the other players, x_{-i}, such that for two different pure strategies for player i, $s_i, s_i' \in S_i$, we have that $P_i(x^{s_i}, x_{-i}) = P_i(x^{s_i'}, x_{-i})$. Ergo, the two distributions don't depend on the choice between s_i and s_i'. The $k-1$ tuple, x_{-i}, determines at most $N = \prod_{j\neq i} |S_j|$ probabilities for the various payoffs available if player i plays s and at most N probabilities for the various payoffs if player i plays s'.

Observation 4. *A sufficient condition for the game to be in general position is if all payoffs for player i are distinct, i.e., $p_i(s_i, s_{-i}) \neq p_i(s_i', s_{-i})$ for all $s_i, s_i' \in S_i$ and $s_{-i} \in \prod_{j\neq i} S_j$. The two distributions $P_i(x_i^s, x_{-i})$ and $P_i(x_i^{s'}, x_{-i})$ cannot be the same because the support sets are different. This is not a necessary condition, e.g., Crawford's game with $r \neq 1$ is in general position.*

This also implies that if a small random error is added to every payoff value then the game will be in general position with high probability. □

5 The Complexity of Risk Aversion

Theorem 5. *Given a 2-player game and non-concave valuations V, V', it is NP-complete to tell if the game has an (V, V')-Nash equilibrium.*

Proof. We give below the proof for the case in which both players' valuations are V^2; Subsequently, we give extensions and generalizations.

NP-completeness is via reduction from 3SAT. Given a 3SAT instance ϕ with n variables $x_1, \ldots, x_n$ and a set of clauses $C = \{c_1, \ldots, c_m\}$, where each c_j is a subset of size three of the set of literals $L = \{+x_1, -x_1, +x_2, \ldots, -x_n\}$, we construct a 2-player game G_ϕ as follows. Both players have the same strategy set $S_1 = S_2 = L \cup C \cup \{f_1, f_2\}$, and their utilities are as follows ($M = 4n^2$):

- For every variable $x_i \in V$,

  $$\begin{aligned} p_1(+x_i, +x_i) &= p_2(+x_i, +x_i) = p_1(-x_i, -x_i) \\ &= p_2(-x_i, -x_i) = M. \end{aligned}$$

  However,

  $$\begin{aligned} p_1(+x_i, -x_i) &= p_2(+x_i, -x_i) = p_1(-x_i, +x_i) \\ &= p_2(-x_i, +x_i) = M - 2n. \end{aligned}$$

  Also, for every two variables $x_i, x_j \in V$ with $i \neq j$, $p_1(\pm x_i, \pm x_j) = M + g(i, j)$, and $p_2(\pm x_i, \pm x_j) = M - g(i, j)$, where $g(i, j) = 1$ if $j = 1 + 1 \bmod n$, -1 if $j = i - 1 \bmod n$, and 0 otherwise.

  The game restricted to L is a generalized rock-paper-scissors zero-sum game (with payoffs translated by M) in which the signs of the literals do no matter, except that both players are incentivized not to play opposite literals of the same variable. It is easy to see that there are 2^n (V^2, V^2)-Nash equilibria of this game. In each, the two players choose the same truth assignment (n literals, one for each variable) and play every literal with probability $\frac{1}{n}$. The V_2 valuation for this payoff distribution is $M - \frac{2}{n}$.
- The purpose of the strategies in C is to force the truth assignment chosen by the equilibrium in L to satisfy ϕ; if it does not, there is a strategy in C (the violated clause) that breaks the equilibrium by presenting a better alternative. The utilities are as follows: For any $c, c' \in C, p_1(c, c') = p_2(c', c) = M - 2n$ (it is disadvantageous for both players to both play in C). Also, for any literal λ and any clause c, $p_1(\lambda, c) = p_2(c, \lambda) = M - 2n$ (it is also disadvantageous to play a literal if the opponent is playing a clause.) Now, to encode the 3SAT instance, for any literal λ and clause c such that $\lambda \in c$, $p_1(c, \lambda) = p_2(\lambda, c) = M - \frac{2}{n}$, whereas if $\lambda \notin c$ we have $p_1(c, \lambda) = p_2(\lambda, c) = M + n - \frac{2}{n}$.
- Finally, the last two strategies f_1, f_2 provide an alternative game, to be played if an equilibrium does not exist in $L \cup C$ — but this game is essentially Crawford's game Γ (see Equation (1)), known to have no Nash equilibria for risk averse players. In particular,

- $p_1(x, f_j) = p_2(f_j, x) = 0$ for all $x \in L \cup C, j = 1, 2$;
- $p_1(f_j, x) = p_2(x, f_j) = M - \frac{2}{n}$ for all $x \in L \cup C, j = 1, 2$;
- $p_1(f_1, f_1) = \frac{M}{2} + 1$;
- $p_1(f_1, f_2) = \frac{M}{2}$;
- $p_1(f_2, f_1) = \frac{M}{2}$;
- $p_1(f_2, f_2) = \frac{M}{2} + 2$.
- $p_2(f_1, f_1) = \frac{M}{2} - 1$;
- $p_2(f_1, f_2) = \frac{M}{2}$;
- $p_2(f_2, f_1) = \frac{M}{2}$;
- $p_2(f_2, f_2) = \frac{M}{2} - 2$.

– Each player can choose either to play the strategies in $L \cup C$ (we call this "playing ϕ"), or choose one of the strategies f_1, f_2 (called "playing f"). If both players play f, notice that they end up playing a version of the game Γ described above, translated upwards by $\frac{M}{2}$.
– Suppose that the two players play f with probability x and y, respectively. Then it is easy to see that the V^2 value of the first player is approximately

$$\frac{\tilde{M}}{2}xy + \tilde{M}(1-y) - \frac{\tilde{M}^2}{4}xy - (1-y)\tilde{M}^2 + \frac{\tilde{M}^2}{4}x^2y^2 + \tilde{M}^2(1-y)^2 - \tilde{M}^2xy(1-y),$$

where by $\tilde{M}$ we denote $M(1 \pm O(\frac{1}{n}))$. By looking at the polynomial in x that results if we ignore $O(\frac{1}{n})$ terms and normalize by M^2

$$(\frac{y^2}{4}) \cdot x^2 + (y^2 - \frac{3y}{4}) \cdot x - y(1-y),$$

we notice (just by looking at the quadratic term) that, for all values of $y \in (0, 1]$, it attains its maximum (in x) at either $x = 0$ or $x = 1$. Hence, for large enough n, and if $y > 0$, player 1 either chooses purely ϕ, or purely f. If $y = 0$, then it is easy to see that $x = 0$ as well, because otherwise player 2 would be better off playing f ($y = 1$). Hence, player 1 plays either purely ϕ or purely f. By symmetry, the same holds for player 2. We conclude that, for large enough n, the only possible (V^2, V^2)-Nash equilibria have either both players playing ϕ, or both playing f.
– But of course, since both players choosing f entails playing a translated version of game Γ, there can be no (V^2, V^2)-Nash equilibria in which both players play f. Also, one player playing ϕ and the other f cannot be an equilibrium (it is a disaster for the player playing ϕ). We conclude that in any (V^2, V^2)-Nash equilibrium both players must play ϕ.
– So, let us assume that both players play ϕ. If one of them plays a strategy in C, then the other player cannot be playing strategies in L (because the f strategies would fare better for the other player than those in L, which now have V^2 payoff strictly less than $M - \frac{2}{n}$). Hence, the other player must be playing only C, and so the first player's payoff with the strategy in c is again lower than that of the f strategies, and this cannot be an equilibrium.

- We conclude that if an equilibrium exists its support is a subset of L, and hence it is a truth assignment. We claim that it is satisfying. Because, if not, there is a clause c that is not satisfied — that is, no literal in it is played by both players. In this case, c has V^2 payoff equal to $M + n - \frac{2}{n}$, and so the players should rather play c. We conclude that if the game has an (V^2, V^2)-equilibrium, then ϕ is satisfiable.
- Conversely, if ϕ is satisfiable, we claim that both players playing the literals in the satisfying truth assignment with probability $\frac{1}{n}$ is an (V^2, V^2)-Nash equilibrium. The (V^2, V^2) payoff of each player is $M - \frac{2}{n}$. Playing the f strategies instead would have the same payoff. And playing any strategy in C would have (V^2, V^2) payoff at most $M - \frac{2}{n}$, because at least one literal in the clause is played with probability $\frac{1}{n}$, and this brings the payoff down to $M - \frac{2}{n}$. □

Other Valuations. Looking at the proof of Theorem 5 we note that the following holds:

Corollary 6. *Suppose that V and V' are risk valuations with the following properties:*

- *There are games with arbitrarily large payoffs that have no (V, V')-Nash equilibria;*
- *The only (V, V')-Nash equilibrium of the generalized rock-paper-scissors game, even if shifted by M, is the uniform play;*
- *For large enough M and for any game whose payoff matrices are within L_2 distance one from*

M, M	$0, M$
$M, 0$	M, M

the only (V, V')-Nash equilibria are pure.

Then it is NP-complete to tell if a game has a (V, V')-Nash equilibrium. □

6 Discussion and Open Problems

No equilibria means that agents may be in a state of constant flux, and NP-completeness implies that they can't even realize their predicament. For a game matrix Γ and valuation functions $\{V_i\}$, if Nash Equilibria do not exist then there is some $f(\Gamma, \{V_i\})$ such that for all $\varepsilon < f(\Gamma, \{V_i\})$ no ε-Nash equilibrium exists. Thus, $f(\Gamma, \{V_i\})$ is a measure of the instability of the player dynamics and large values may help explain dramatic instabilities in certain games. It is easy to choose r for Crawford's game or jittered versions of matching pennies, with risk valuations V^2, for which $f(\Gamma, \{V^2\})$ equal to some 10% of the total value — very strong motivation for chaotic behavior. Measuring the instability of a game played by risk-averse players, and comparing this to observed behavior in strategic games, even markets, seems interesting.

Also, as pointed out in Observation 4, incomplete or partial information about the payoffs, if interpreted as making some random error about the value of the payoffs, results in games in general position and no mixed equilibria. It seem promising to characterize the instability of the resulting game as a function of the magnitude of error. Intuitively, instability should grow with error size.

The possibility that mixed Nash equilibria may not exist makes weaker solution concepts, such as the *correlated equilibrium,* more attractive. Unfortunately, it is easy to see that even those may not exist. For example, Crawford's uneven matching pennies game played by players whose risk valuation is V_6 has no correlated equilibria. We conjecture that the circumstances under which correlated equilibria fail to exist are much less common, but that it is still NP-complete to tell if one exists or not.

One technical open problem suggested by this work is to determine the extent of the valuations for which NP-completeness holds. We have not strived to state the most general theorem possible here, but we believe that our proof can be generalized in many directions. Note that, by the corollary, NP-completeness holds even even when all players except for one are expectation maximizers.

References

1. Allais, M.: Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole americaine. Econometrica 21(4), 503–546 (1953), `http://www.jstor.org/stable/1907921`
2. Bernoulli, D.: Exposition of a new theory on the measurement of risk. Commentaries of the Imperial Academy of Science of Saint Petersburg (1738); Republished Econometrica 22(1), 23–36 (1954), `http://www.jstor.org/stable/1909829`
3. Borel, E.: The theory of play and integral equations with skew symmetric kernels. In: Translated by Leonard J. Savage, University of Chicago, from Comptes Rendus de l'Academie des Sciences, December 19, vol. 173, pp. 1304–1308 (1921); Reprinted in Econometrica 21(1), 97–100 (1953), `http://www.jstor.org/stable/1906946`
4. Chen, X., Deng, X.: Settling the complexity of two-player nash equilibrium. In: FOCS 2006: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science, pp. 261–272. IEEE Computer Society, Washington (2006)
5. Conitzer, V., Sandholm, T.: New complexity results about nash equilibria. Games and Economic Behavior 63(2), 621–641 (2008), `http://econpapers.repec.org/RePEc:eee:gamebe:v:63:y:2008:i:2:p:621-641`
6. Crawford, V.P.: Equilibrium without independence. Journal of Economic Theory 50(1), 127–154 (1990), `http://www.sciencedirect.com/science/article/B6WJ3-4CYH5N4-CF/2/0071841ad1283474807beca912feb550`
7. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a nash equilibrium. ACM Commun. 52(2), 89–97 (2009)
8. Friedman, M., Savage, L.J.: The utility analysis of choices involving risk. Journal of Political Economy 56 (1948), `http://econpapers.repec.org/RePEc:ucp:jpolec:v:56:y:1948:p:279`
9. Chen, H.-C.: Pure-strategy equilibria with non-expected utility players. Theory and Decision 46, 201–212 (1999), `http://www.ingentaconnect.com/content/klu/theo/1999/00000046/00000002/00152924`

10. Kahneman, D., Tversky, A.: Prospect theory: An analysis of decision under risk. Econometrica 47(2), 263–291 (1979), http://www.jstor.org/stable/1914185
11. Karni, E., Safra, Z.: Dynamic consistency, revelations in auctions and the structure of preferences. Review of Economic Studies 56(3), 421–433 (1989), http://econpapers.repec.org/RePEc:bla:restud:v:56:y:1989:i:3:p:421-33
12. Machina, M.J.: Choice under uncertainty: Problems solved and unsolved. The Journal of Economic Perspectives 1(1), 121–154 (1987), http://www.jstor.org/stable/1942952
13. Machina, M.J.: Dynamic consistency and non-expected utility models of choice under uncertainty. Journal of Economic Literature 27(4), 1622–1668 (1989), http://ideas.repec.org/a/aea/jeclit/v27y1989i4p1622-68.html
14. Markowitz, H.: Portfolio Selection: Efficient Diversification in Investments. John Wiley, Chichester (1959)
15. Markowitz, H.: Portfolio selection. The Journal of Finance 7(1), 77–91 (1952), http://www.jstor.org/stable/2975974
16. Marschak, J.: Rational behavior, uncertain prospects, and measurable utility. Econometrica 18(2), 111–141 (1950), http://www.jstor.org/stable/1907264
17. Myerson, R.B., Myerson, R.B.: Nash equilibrium and the history of economic theory. Journal of Economic Literature 37, 1067–1082 (1999)
18. Nash, J.: Non-cooperative games. The Annals of Mathematics 54(2), 286–295 (1951), http://www.jstor.org/stable/1969529
19. Neumann, J.v.: Zur Theorie der Gesellschaftsspiele. Math. Ann. 100(1), 295–320 (1928), http://dx.doi.org/10.1007/BF01448847
20. Neumann, J.V., Morgenstern, O.: Theory of Games and Economic Behavior. Princeton University Press, Princeton (1944), http://jmvidal.cse.sc.edu/library/neumann44a.pdf
21. Papadimitriou, C.H.: On the complexity of the parity argument and other inefficient proofs of existence. J. Comput. Syst. Sci. 48(3), 498–532 (1994)
22. Translated by Richard J. Pulskamp, N.B.: Correspondence of nicolas bernoulli concerning the st. petersburg game (1713-1732) (1999), http://cerebro.xu.edu/math/sources/montmort/stpetersburg.pdf
23. Ritzberger, K.: On games under expected utility with rank dependent probabilities. Theory and Decision 40(1), 1–27 (1996)
24. Sharpe, W.F.: A simplified model for portfolio analysis. Management Science 9(2), 277–293 (1963), http://www.jstor.org/stable/2627407
25. Stanford, W.: The limit distribution of pure strategy nash equilibria in symmetric bimatrix games. Mathematics of Operations Research 21(3), 726–733 (1996)
26. Starmer, C.: Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. Journal of Economic Literature 38(2), 332–382 (2000), http://ideas.repec.org/a/aea/jeclit/v38y2000i2p332-382.html
27. Varian, H.: A portfolio of nobel laureates: Markowitz, miller and sharpe. The Journal of Economic Perspectives 7(1), 159–169 (1993), http://www.jstor.org/stable/2138327