The Computational Aspect of Risk in Playing Non-Cooperative Games

DEEPARNAB CHAKRABARTY SUBHASH A. KHOT RICHARD J. LIPTON
College of Computing, Georgia Tech*

NISHEETH K. VISHNOI
Computer Science Division, U.C. Berkeley†

Abstract

This paper considers the computational aspect of risk in game playing situations. The risk we consider arises due to the players being compelled to choose mixed strategies in order to ensure that they play at a strategic equilibrium, which happens often when the game has no pure strategic equilibria. More precisely, this paper studies the following question: *What is the computational complexity of finding equilibrium or near-equilibrium points where the risks of a given set of players are within specified bounds?*

1 Introduction

Risk in Playing Games. Consider a zero-sum game whose payoff matrix is presented in Figure 1. The

1	-1
-1	1
10^6	-10^6
-10^6	10^6

Figure 1: Payoff matrix of a zero-sum game with "high" and "low" risk equilibria for the row player

row player has four strategies: $(1, 2, 3, 4)$, one corresponding to each row, and the column player has two strategies: $(1, 2)$, those corresponding to the two columns. The (i, j)-th entry of the matrix denotes the dollar amount the row player pays to the column player when they pick strategies i and j respectively. It can be checked that at any strategic (Nash) equilibrium the expected payoff of both the players is 0, and that there are infinitely many such equilibria. For instance, (1) playing rows 1 and 2 with probability 1/2 each; (2) playing rows 3 and 4 with probability 1/2 each. Although for each of these strategic equilibria the payoff of the row player is 0 in expectation, a case can be made that a "risk averse" row player would prefer strategy (1) over the other. One justification being that in strategy (2), the probability that the row player ends up paying a million dollars to the column player is 1/2, while in strategy (1), all payments are bounded in absolute value by a dollar. This would not be an issue if the game could be played an infinite number of times as the average payoff would then tend to zero in the limit. But there are at least two problems with this argument involving repeating the game several times to converge to expected payoff: (1) One should then consider strategic equilibria of the repeated game which may, in general, contain many more strategies

*{deepc,khot,rjl}@cc.gatech.edu
†Correspondence address: nisheeth.vishnoi@gmail.com

than the strategic equilibria of the single shot game. Although this raises interesting issues, in this paper we will focus on the case when the game is played only once. (2) To play this game an infinite (or a large) number of times, one needs a proportional capital. In situations where a "small" row player (such as an individual) is playing against a "large" column player (such as a casino or a stock market), or the payments constitute significantly different percentages of their respective wealths, a more refined look at equilibria seems necessary to incorporate the aspect of risk averseness. It should be noted, however, that it would still be desirable to play at an equilibrium point, as otherwise the players may not have a guaranteed payoff. This is the case we will consider in this paper; the players will only consider playing at equilibrium points (or near them), and hence, it makes sense to consider risk only near these points. This raises the following question: *Are there equilibrium points where the risk of some given set of players is within a specified bound? If so, can they be found efficiently?* To formalize this question we first need a measure of risk. There have been several approaches in the game theory and economics literature to quantify risk, and we adopt the most popular one; the variance of the random variable associated to the payoff. The next step is to incorporate the notion of risk in game playing. Again, there have been numerous attempts, but the one we choose is to take risk into consideration only at equilibrium or near equilibrium points. Once, this is done, we focus on the computational aspect of this problem.

Measuring Risk. One natural (and popular) way to capture the risk associated to a random variable involving payments is its variance. Thus, in a game playing situation (such as the one discussed in the beginning), at an equilibrium point, apart from the expected payoff, each player also has an associated risk. This gives a refined way of looking at equilibria, and possibly, gives players an objective criteria for selecting an equilibria when there are many. There have been essentially two approaches to measure risk: (1) An axiomatic approach [2, 10, 7]. (2) Univariate measures such as standard deviation, variance, variance/expectation, absolute deviation and interquartile range. All of them have their advantages and disadvantages, and for a thorough discussion, the reader is referred to [2, 8]. For instance, the axiomatic approach of Pollatsek and Tversky [10] leads to the risk associated to a random variable X being a convex combination of $\mathbf{Var}[X]$ and $\mathbf{E}[X]$, while the approach of Aumann and Serrano [2] leads to the risk of a random variable X being the number ρ such that $\mathbf{E}[\exp(-X/\rho)] = 1$. The key criticism against the use of variance is that it could be insensitive to stochastic dominance: the variance of the random variables X and $X + c$, for any $c > 0$ is the same, while $X + c > X$ everywhere. In game playing situations, the notion of stochastic dominance has no meaning as the payments are a function of the strategies of all the players. Moreover, in zero-sum games, all equilibrium points have the *same* expected payoff. This being said, we note that our results are not very sensitive to the measure of risk and can be generalized to several of the aforementioned measures.

Incorporating Risk. There have been several attempts to incorporate the notion of risk in game playing. The first systematic approach towards this was suggested by von Neumann and Morgenstern [12]. They show how a player's preferences, and therewith also its attitude towards risk can be captured by a single utility function. However, the strategy profile which would then optimize the utility function may not be an equilibrium point, and hence, have no guaranteed payoff. The other notable approach is due to Harsanyi and Selten [4] who introduced the concept of "risk dominance" as a way for the players to an select among several available equilibria. This notion, however, could assign a non-zero value of risk to a strategic equilibrium involving pure strategies, and is fundamentally different from the situation we consider in this paper where the risk is a statistical quantity associated to the randomization used by the players in game playing. In our case, there is no risk in playing at a pure strategic equilibrium point. Hence, our approach differs from both that of von Neumann and Morgenstern and that suggested by Harsanyi and Selten in that we employ the statistical concept of variance and assume that the players are risk averse in the sense that they prefer solutions with a lesser variance to those with a greater variance.

Our Contribution. We show that there is a procedure that, in nearly polynomial time, decides whether there are "near-equilibrium" points where the risk of each of the player is below the specified threshold.

This generalizes a result due to Lipton, Mehta and Markakis [5]. Improving on this result would imply a faster algorithm to compute near-equilibrium points for non-cooperative games, which is a well-studied and, perhaps, a computationally hard problem.

For the case of zero-sum games we also consider a different aspect of risk. Here, the two players are "unequal" and the row player is far more risk averse than the column player and would like to play at an equilibrium point where its risk is guaranteed. Thus, it makes sense for the row player to assume that the column player may play at an equilibrium point where its risk is maximized. Thus, restricted to the equilibrium strategies (all of which have the same value for zero-sum games), one can think of the players engaged in another zero-sum game associated to the risk of the row player, where one can think of the risk of the row player as its payoff to the column player. In this setting, we generalize the classic theorem of von Neumann [11] and prove a min-max theorem. More importantly, our proof is algorithmic and gives a polynomial time procedure to find such an equilibrium point, where the players are also at an equilibrium point with respect to the risk of the row player.

To the best of our knowledge, our way of incorporating risk into game playing is novel and could lead to enriching our understanding about equilibrium selection in the theory of non-cooperative games. Before proceeding to present our results and their proofs formally, in the next section we recall the basic setting in non-cooperative game theory. In Section 2 we formalize the notion of risk. Subsequently, we state and discuss our results on nonzero-sum and zero-sum games in Section 2 and Section 3 respectively.

1.1 Notational Preliminaries

Consider a game in which there are m players: $P_1, \ldots, P_m$. P_i can choose a strategy from a set Σ_i. P_i has an associated function $f_i : \Sigma_1 \times \cdots \times \Sigma_m \mapsto \mathbb{R}$, which is referred to as its "payoff" function. Its payoff depends not only on its own choice of a strategy, but also on the strategy of other players. The game is characterized by the payoff functions $(f_1, \ldots, f_m)$. In the non-cooperative setting the players want to maximize their individual payoff functions. A tuple of strategies $(\widehat{\sigma}_1, \ldots, \widehat{\sigma}_m)$ is referred to as a " strategic equilibrium" (Nash equilibrium) if no player can gain by defecting unilaterally. Formally, for every i, and every $\sigma_i \in \Sigma_i$,

$$f_i(\widehat{\sigma}_1, \ldots, \widehat{\sigma}_i, \ldots, \widehat{\sigma}_m) \geq f_i(\widehat{\sigma}_1, \ldots, \sigma_i, \ldots, \widehat{\sigma}_m).$$

In general, such an equilibrium point may not exist, even for the simple setting consisting of two players with their payoff functions satisfying $f_1 + f_2 \equiv 0$ for all points in $\Sigma_1 \times \Sigma_2$. (This special case involving two players where one's gain is other's loss is referred to as a "zero-sum" game.) However, if the players are allowed to choose their strategies probabilistically, the classic result of Nash [9] guarantees an equilibrium. In this setting, player P_i's strategy is be a probability distribution $\mathbf{p}_i$ over Σ_i. Given a choice of probability functions $(\mathbf{p}_1, \ldots, \mathbf{p}_m)$, the expected payoff of the i-th player can now be denoted by

$$\mathbf{E}_{\sigma_1 \leftarrow \mathbf{p}_1, \ldots, \sigma_m \leftarrow \mathbf{p}_m}[f_i(\sigma_1, \ldots, \sigma_m)].$$

These $\mathbf{p}_i$'s are assumed to be independent of each other. In this setting, the notion of equilibrium defined above naturally translates to one involving probability distributions. A choice of probability distributions $(\widehat{\mathbf{p}}_1, \ldots, \widehat{\mathbf{p}}_m)$ is said to be a Nash equilibrium if, for every i and every probability distribution $\mathbf{p}_i$ on Σ_i,

$$\mathbf{E}_{\sigma_1 \leftarrow \widehat{\mathbf{p}}_1, \ldots, \sigma_i \leftarrow \widehat{\mathbf{p}}_i, \ldots, \sigma_m \leftarrow \widehat{\mathbf{p}}_m}[f_i(\sigma_1, \ldots, \sigma_m)] \geq \mathbf{E}_{\sigma_1 \leftarrow \widehat{\mathbf{p}}_1, \ldots, \sigma_i \leftarrow \mathbf{p}_i, \ldots, \sigma_m \leftarrow \widehat{\mathbf{p}}_m}[f_i(\sigma_1, \ldots, \sigma_m)].$$

In general there may be many (possibly infinite number of) equilibria.

2 Risk in General Games

Now we can formally define the notion of risk alluded to in the Section 1. We capture the notion of risk at a Nash equilibrium by the variance of the random variable associated to the payoffs. First, for brevity, denote $\mathbf{E}_{\sigma_1 \leftarrow \mathbf{p}_1, \ldots, \sigma_m \leftarrow \mathbf{p}_m}[f_i(\sigma_1, \ldots, \sigma_m)]$ by $\mathbf{E}_{\mathbf{p}_1, \ldots, \mathbf{p}_m}[f_i]$. The variance of the i-th player for these probability distributions is then defined to be

$$\mathbf{Var}_{\mathbf{p}_1, \ldots, \mathbf{p}_m}[f_i] := \mathbf{E}_{\mathbf{p}_1, \ldots, \mathbf{p}_m}[f_i^2] - \mathbf{E}^2_{\mathbf{p}_1, \ldots, \mathbf{p}_m}[f_i].$$

This is referred to as the risk of the i-th player when the game is being played at the point $(\mathbf{p}_1, \ldots, \mathbf{p}_m)$. Given a choice between two equilibria, in the setting we consider, each player prefers the equilibrium with the lesser variance. From a Computer Science perspective, the following question becomes important: Given bounds on the payoff and risks for every player, is there a Nash equilibrium which attains these bounds? Formally, is there an equilibrium point $(\mathbf{p}_1, \ldots, \mathbf{p}_m)$ such that $\mathbf{Var}_{\mathbf{p}_1, \ldots, \mathbf{p}_m}[f_i] \leq R_i$ and $\mathbf{E}_{\mathbf{p}_1, \ldots, \mathbf{p}_m}[f_i] \geq \Pi_i$ for each i? It immediately follows that this problem is at least as hard as that of computing a Nash equilibrium which has been, in a series of recent results culminating in a paper of Chen and Deng [3], established to be PPAD-complete. Hence, it is unlikely that there is an efficient algorithm for this problem.

However, if one is willing to relax the notion of Nash equilibrium to what is now known as an ε-Nash equilibrium, the problem seems to be more tractable from a computational perspective. An ε-Nash equilibrium is one in which no player has more than an ε incentive[1] in defecting. Formally, a point $(\widehat{\mathbf{p}}_1, \ldots, \widehat{\mathbf{p}}_m)$ is said to be an ε-Nash equilibrium if, for every i and every probability distribution $\mathbf{p}_i$ on Σ_i,

$$\mathbf{E}_{\widehat{\mathbf{p}}_1, \ldots, \widehat{\mathbf{p}}_i, \ldots, \widehat{\mathbf{p}}_m}[f_i] \geq \mathbf{E}_{\widehat{\mathbf{p}}_1, \ldots, \mathbf{p}_i, \ldots, \widehat{\mathbf{p}}_m}[f_i] - \varepsilon.$$

This notion has proven to be of interest in the light of the above complexity result. This was studied for zero-sum games independently by [6, 1], and for nonzero-sum games by [5]. The following result was proved in [5].

Theorem 2.1 (Computing ε-Nash Equilibria [5]) *For a game specified by payoffs $(f_1, \ldots, f_m)$ and $\Sigma_i = \{1, \ldots, n\}$ for each i, for any Nash equilibrium $(\widehat{\mathbf{p}}_1, \ldots, \widehat{\mathbf{p}}_m)$ of it, and for any small enough constant $\varepsilon > 0$, there is an $n^{O(\frac{m^2 \ln mn}{\varepsilon^2})}$ time algorithm to find a point $(\mathbf{p}'_1, \ldots, \mathbf{p}'_m)$ such that the following hold:*

1. *$(\mathbf{p}'_1, \ldots, \mathbf{p}'_m)$ is an ε-Nash equilibrium.*
2. *All the players get an expected payoff which is within an ε-additive error from the expected payoff at $(\widehat{\mathbf{p}}_1, \ldots, \widehat{\mathbf{p}}_m)$, i.e., for all $1 \leq i \leq m$, $\left|\mathbf{E}_{\widehat{\mathbf{p}}_1, \ldots, \widehat{\mathbf{p}}_m}[f_i] - \mathbf{E}_{\mathbf{p}'_1, \ldots, \mathbf{p}'_m}[f_i]\right| \leq \varepsilon$.*

This result is proved by showing that every Nash equilibrium can be well approximated by a set of strategies, such that for each player, its strategy is uniform[2] on $O(\frac{m^2 \ln mn}{\varepsilon^2})$ pure strategies. Hence, one can run over, for every player, a choice of a subset of pure strategies of size $O(\frac{m^2 \ln mn}{\varepsilon^2})$. When m is a constant, this gives a slightly super-polynomial algorithm to find all ε-Nash equilibria. We extend Theorem 2.1 to the case of the variance as follows:

Theorem 2.2 *For a game specified by payoffs $(f_1, \ldots, f_m)$ and $\Sigma_i = \{1, \ldots, n\}$ for each i, for any Nash equilibrium $(\widehat{\mathbf{p}}_1, \ldots, \widehat{\mathbf{p}}_m)$ of it, and for any small enough constant $\varepsilon > 0$, there is an $n^{O(\frac{m^2 \ln mn}{\varepsilon^2})}$ time algorithm to find a point $(\mathbf{p}'_1, \ldots, \mathbf{p}'_m)$ such that the following hold:*

1. *$(\mathbf{p}'_1, \ldots, \mathbf{p}'_m)$ is an ε-Nash equilibrium.*
2. *For all $1 \leq i \leq m$, $\left|\mathbf{E}_{\widehat{\mathbf{p}}_1, \ldots, \widehat{\mathbf{p}}_m}[f_i] - \mathbf{E}_{\mathbf{p}'_1, \ldots, \mathbf{p}'_m}[f_i]\right| \leq \varepsilon$.*
3. *The risk of all the players is within an ε-additive error from that at $(\widehat{\mathbf{p}}_1, \ldots, \widehat{\mathbf{p}}_m)$, i.e., for each $1 \leq i \leq m$, $\left|\mathbf{Var}_{\widehat{\mathbf{p}}_1, \ldots, \widehat{\mathbf{p}}_m}[f_i] - \mathbf{Var}_{\mathbf{p}'_1, \ldots, \mathbf{p}'_m}[f_i]\right| \leq \varepsilon$.*

[1] As is standard, we assume for simplicity that all f_i take values in the range $[-1, 1]$. All the results translate to when this is not the case, albeit, with appropriate scaling parameters.

[2] Possibly with repetitions.

This allows us to search through all near-equilibrium points to find one which meets the desired risk conditions imposed by the players. This is possible because the variance is also approximated to within an ε-additive error. For a constant number of players, we get an algorithm that runs in time $n^{O(\frac{\ln n}{\varepsilon^2})}$. Thus, the following theorem answers the question raised in the beginning of the section.

Corollary 2.3 *Given bounds $\Pi_1, \cdots, \Pi_m$ on the payoffs and $R_1, \cdots, R_m$ on the risk, if there is a Nash equilibrium satisfying these bounds, then in $n^{O(\frac{\ln m}{\varepsilon^2})}$ time we can find an ε-Nash equilibrium such that for every player i the payoff is at least $\Pi_i - \varepsilon$ and the risk is at most $R_i + \varepsilon$.*

The full proof of Theorem 2.2 is presented in the Appendix.

3 Risk in Zero-Sum Games

In zero-sum games, due to the fact that all equilibrium points give the same payoff to the players, we focus on the situation where the two players are unequal in their aversity to risk. First, we recall the basic setting. A zero-sum game is one in which for any choice of strategies for the players, the sum of all the payoffs is zero. Formally, for any $(\sigma_1, \ldots, \sigma_m) \in \Sigma_1 \times \cdots \times \Sigma_m$, $\sum_{i=1}^m f_i(\sigma_1, \ldots, \sigma_m) = 0$. In this case, it is convenient to assume that $\Sigma_1 = \Sigma_2 = \{1, \ldots, n\}$, and to represent the payoffs by a matrix M, whose (i,j)-th entry is $M(i,j) := f_2(i,j)$. Thus, $M(i,j)$ can be interpreted as the cost of the "row" player and the payoff of the "column" player. Hence, the expected payoff of a pair of mixed strategies, $\mathbf{p}$ for the row player and $\mathbf{q}$ for the column player, is $(-\mathbf{p}^{\mathrm{t}} M\mathbf{q}\ , \mathbf{p}^{\mathrm{t}} M\mathbf{q})$. We let $v(\mathbf{p}, \mathbf{q}) := \mathbf{p}^{\mathrm{t}} M\mathbf{q}$. It was shown by von Neumann [11] that zero-sum games always have a strategic equilibrium point (which was termed as Nash equilibrium after Nash generalized his result to nonzero-sum games). He also proved the following theorem which gives an important characterization of equilibrium points in zero-sum games.

Theorem 3.1 (von Neumann [11]) *Given any two-player zero-sum game specified by the matrix M, the following min-max theorem holds*

$$\min_{\mathbf{p}} \ \max_{\mathbf{q}} \ \mathbf{p}^{\mathrm{t}} M\mathbf{q} = \max_{\mathbf{q}} \ \min_{\mathbf{p}} \ \mathbf{p}^{\mathrm{t}} M\mathbf{q}.$$

While in general nonzero-sum games, there does not seem to be a mechanism whereby players can play at an equilibrium without a prior consensus, it is possible to play a (two player) zero-sum game without deciding an equilibrium point a priori. The row player picks any min-max strategy while the column player picks any max-min strategy and that constitutes an equilibrium point. The following corollary, which is immediate from von Neumann's Theorem, shows that all equilibrium points have the same payoff.

Corollary 3.2 *Given any two-player zero-sum game specified by the matrix M, all its Nash equilibria have the same value.*

Hence, without any confusion, we can define the value of the game (denoted $v(M)$) to be the value at any Nash equilibrium. Further, let $U(M)$ denote the set of all Nash equilibria for the zero-sum game with payoff matrix M. The risk at any point $(\mathbf{p}, \mathbf{q})$ for the two players turns out to be:

$$\text{Row player}: \ \mathbf{p}^{\mathrm{t}}(-M)^{(2)}\mathbf{q} - (\mathbf{p}^{\mathrm{t}}(-M)\mathbf{q})^2, \quad \text{Column player}: \ \mathbf{p}^{\mathrm{t}} M^{(2)}\mathbf{q} - (\mathbf{p}^{\mathrm{t}} M\mathbf{q})^2.$$

Here for a matrix A, $A^{(2)}$ denotes the matrix whose (i,j)-th entry is $(A(i,j))^2$. Note that the risk for both the players is the *same*. Again, Theorem 2.2 can be used to determine quickly if there is a near equilibrium point where the risks of the players is bounded by a specified value. Now we consider the more interesting situation when the row player is far more risk-averse than the column player and would like to play at an equilibrium point where its risk is guaranteed to be bounded from above. Here, the row player would (pessimistically) assume that the column player wishes to maximize the risk of the row player. This scenario would arise in the case when the wealth (if the payoffs were in terms of money) of the column player is considerably more than that of the row player; a casino and an individual gambler, for instance. But, the column player would still restrict its play to an equilibrium point as it wishes to maximize its worst case expected payoff. Formally, the row player wishes to play the following min-max strategy $\mathbf{p}$:

$$\min_{\mathbf{p}} \max_{\mathbf{q}:(\mathbf{p},\mathbf{q})\in U(M)} \mathbf{p}^{\mathrm{t}} M^{(2)}\mathbf{q} - (\mathbf{p}^{\mathrm{t}} M\mathbf{q})^2$$

while the column player wishes to play the following max-min strategt $\mathbf{q}$:

$$\max_{\mathbf{q}} \min_{\mathbf{p}:(\mathbf{p},\mathbf{q})\in U(M)} \mathbf{p}^{\mathrm{t}} M^{(2)}\mathbf{q} - (\mathbf{p}^{\mathrm{t}} M\mathbf{q})^2.$$

In this setting, we prove the following generalization of von Neumann's Theorem showing that the above two expressions are the same. Moreover, we give a polynomial time algorithm to compute these risk-minimizing strategies for the row player.

Theorem 3.3 *Given a zero-sum game M on n strategies, there is an $n^{O(1)}$ time algorithm to compute a Nash equilibrium $(\mathbf{p},\mathbf{q})$ and a number R such that on playing at $(\mathbf{p},\mathbf{q})$, the row player (column player) is guaranteed a risk of at most (resp. at least) R when the column player (resp. row player) plays an equilibrium strategy.*

The proof of this theorem follows immediately from the proof of the following lemma.

Lemma 3.4 *Given a zero-sum game M,*

$$\min_{\mathbf{p}} \max_{\mathbf{q}:(\mathbf{p},\mathbf{q})\in U(M)} \mathbf{p}^{\mathrm{t}} M^{(2)}\mathbf{q} - (\mathbf{p}^{\mathrm{t}} M\mathbf{q})^2 = \max_{\mathbf{q}} \min_{\mathbf{p}:(\mathbf{p},\mathbf{q})\in U(M)} \mathbf{p}^{\mathrm{t}} M^{(2)}\mathbf{q} - (\mathbf{p}^{\mathrm{t}} M\mathbf{q})^2.$$

Proof: First let us present a few facts about the set $U(M)$. Note that since $U(M)$ is the set of Nash equilibria of the zero-sum game,

$$U(M) := \{(\mathbf{p},\mathbf{q}) : \min_{\mathbf{p}'} \mathbf{p}'^t M\mathbf{q} = v(M), \quad \max_{\mathbf{q}'} \mathbf{p}^t M\mathbf{q}' = v(M)\}$$

where $\mathbf{p},\mathbf{q},\mathbf{p}',\mathbf{q}'$ range over probability distributions over $\{1,\ldots,n\}$. Let $U_p := \{\mathbf{p} : \max_{\mathbf{q}'} \mathbf{p}^t M\mathbf{q}' = v(M)\}$ and $U_q := \{\mathbf{q} : \min_{\mathbf{p}'} \mathbf{p}'^t M\mathbf{q} = v(M)\}$ be the equilibrium strategies for the row player and the column player respectively. The following claim is immediate from the definition.

Claim 3.5 $U(M) = U_p \times U_q$.

Since best responses are pure strategies, the sets U_p and U_q can also be written as

$$\begin{aligned} U_p &= \{\mathbf{p} : \mathbf{p}^t M\mathbf{e}_j \le v(M) \ \ \forall j \in [n]\} && (1)\\ U_q &= \{\mathbf{q} : \mathbf{e}_i^t M\mathbf{q} \ge v(M) \ \ \forall i \in [n]\}. && (2)\end{aligned}$$

Here $\mathbf{e}_i$ denotes the n-dimensional vector with its i-th coordinate equal to 1 and all the other coordinates equal to 0. Thus, both U_p and U_q are polytopes. Moreover, given any probability distribution, it is efficiently checkable if it is in U_p or U_q, or a certificate can be found efficiently in the case it doesn't. That is, U_p, U_q possess a separation oracle which runs in time $n^{O(1)}$. Let the sets S and T denote the vertices of U_p and U_q respectively. Let N denote the matrix $M^{(2)}$. To prove the lemma, it is sufficient to prove

$$\min_{\mathbf{p}\in U_p} \max_{\mathbf{q}\in U_q} \mathbf{p}^t N\mathbf{q} = \max_{\mathbf{q}\in U_q} \min_{\mathbf{p}\in U_p} \mathbf{p}^t N\mathbf{q}. \tag{3}$$

Let us consider the l.h.s. of the above equation. Fix a $\mathbf{p}\in U_p$. Then, the remaining expression is maximizing a linear function over the polytope U_q. Since this maximum will be attained at a vertex, the l.h.s. can be re-written as follows:

$$\min_{\mathbf{p}\in U_p} \max_{\mathbf{q}\in U_q} \mathbf{p}^{\mathrm{t}} M\mathbf{q} = \{\min R : \ R \ge \mathbf{p}^{\mathrm{t}} M\mathbf{t}_i, \ \forall \ \mathbf{t}_i \in T, \ \mathbf{p}\in U_p\}. \tag{4}$$

Thus, this min-max expression can be written as a linear program (LP). Also, given $(R, \mathbf{p})$, one can check feasibility by checking if $\mathbf{p} \in U_p$ (done above) and if so, maximizing a linear function over the set U_q and checking if it is less than R. If the point is not feasible, then the above also returns the inequality which is violated. That is, there is a polynomial time separation oracle implying that the LP can be solved in polynomial time via the ellipsoid method. To prove (3), we note that the r.h.s. of (3) can also be written as the following LP:

$$\max_{\mathbf{q}\in U_q} \min_{\mathbf{p}\in U_p} \mathbf{p}^{\mathrm{t}} N\mathbf{q} = \{\max R : \; R \leq \mathbf{s}_j^{\mathrm{t}} N\mathbf{q}, \; \forall \, \mathbf{s}_j \in S, \; \mathbf{q} \in U_q\}. \tag{5}$$

As we prove now, LP (5) is the dual of LP (4). The Strong Duality Theorem then implies (3). Note that the constraint $\mathbf{p} \in U_p$ in LP (4) can be written as a set of linear constraints by introducing variables $\{\lambda_1, \lambda_2, \cdots, \lambda_a\}$ and replacing $\mathbf{p} \in U_p$ by the constraints $\mathbf{p} = \sum_{\mathbf{s}_j\in S} \lambda_j \mathbf{s}_j$ and $\sum_{\mathbf{s}_j\in S} \lambda_j = 1$, and the non-negativity constraints for λ_j's. Thus, we get,

$$\begin{aligned} \min_{\mathbf{p}\in U_p} \max_{\mathbf{q}\in U_q} \mathbf{p}^{\mathrm{t}} N\mathbf{q} \quad = \quad & \min R \\ & R \geq (\sum_{\mathbf{s}_j\in S} \lambda_j \mathbf{s}_j^{\mathrm{t}}) N\mathbf{t}_i \;, \quad \forall \; \mathbf{t}_i \in T \\ & \sum_{\mathbf{s}_j\in S} \lambda_j = 1 \\ & \lambda_j \geq 0 \;, \quad 1 \leq j \leq a. \end{aligned}$$

The dual of the above is:

$$\begin{aligned} & \max R \\ & R \leq \mathbf{s}_j^{\mathrm{t}} N (\sum_{\mathbf{t}_i\in T} \mu_i \mathbf{t}_i) \;, \quad \forall \; \mathbf{s}_j \in S \\ & \textstyle\sum_{\mathbf{t}_i\in T} \mu_i = 1 \\ & \mu_i \geq 0 \;, \quad 1 \leq i \leq b. \end{aligned}$$

which is precisely the LP (5). ∎

Acknowledgments. We thank Adam Kalai and Ehud Kalai for discussions regarding the notion of risk.

References

[1] I. Althofer. On sparse approximations to randomized strategies and convex combinations, 1994.

[2] R. Aumann and R. Serrano. An economic index of riskiness. working paper, (http://www.brown.edu/Departments/Economics/Papers/2006/2006-20-paper.pdf), 2006.

[3] X. Chen and X. Deng. Settling the complexity of 2-player nash-equilibrium. In *Proceedings of FOCS*, 2006.

[4] J. Harsanyi and R. Selten. *A General Theory of Equilibrium Selection in Games.* MIT Press, 1988.

[5] R. J. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *Proceedings of ACM conference in Electronic Commerce*, 2003.

[6] R. J. Lipton and Neal E. Young. Simple strategies for large zero-sum games with applications to complexity theory. In *Proceedings of ACM Symposium on Theory of Computing*, 1994.

[7] R. D. Luce. Several possible measures of risk. *Theory and Decision*, 12(3):217–228, 1980.

[8] M. Machina and M. Rothschild. *The New Palgrave Dictionary of Economics*, chapter Risk. forthcoming.

[9] J. F. Nash. Non-cooperative games. *Annals of Mathematics*, pages 286–295, 1951.

[10] A. Pollatsek and A. Tversky. A theory of risk. *Journal of Mathematical Psychology*, 7(3):540–553, 1970.

[11] J. von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1928.

[12] J. von Nuemann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

A Computing ε-Nash equilibrium with Variances and Proof of Theorem 2.2

We just prove Theorem 2.2 for the case of two players. The proof easily generalizes to m players.

Theorem A.1 *For a game specified by payoffs (f_1, f_2) and $\Sigma_1 = \Sigma_2 = \{1, \ldots, n\}$, for any Nash equilibrium $(\widehat{\mathbf{p}}_1, \widehat{\mathbf{p}}_2)$ of it, and for any small enough constant $\varepsilon > 0$, there exists, for every $k = \Omega(\frac{\ln n}{\varepsilon^2})$, a pair of k-uniform strategies $(\mathbf{p}'_1, \mathbf{p}'_2)$ such that the following hold:*

1. *$(\mathbf{p}'_1, \mathbf{p}'_2)$ is an ε-Nash equilibrium.*
2. *For $i = 1, 2$, $\left|\mathbf{E}_{\widehat{\mathbf{p}}_1,\widehat{\mathbf{p}}_2}[f_i] - \mathbf{E}_{\mathbf{p}'_1,\mathbf{p}'_2}[f_i]\right| \leq \varepsilon$.*
3. *For $i = 1, 2$, $\left|\mathbf{Var}_{\widehat{\mathbf{p}}_1,\widehat{\mathbf{p}}_2}[f_i] - \mathbf{Var}_{\mathbf{p}'_1,\mathbf{p}'_2}[f_i]\right| \leq 2\varepsilon$.*

Proof: Given a Nash equilibrium $(\widehat{\mathbf{p}}_1, \widehat{\mathbf{p}}_2)$, form a multi-set A by drawing k independent samples (possibly with repetitions) from the distribution $\widehat{\mathbf{p}}_1$. Similarly, form B by drawing k independent samples from $\widehat{\mathbf{p}}_2$. The value of k will be specified later. Note that for any $\sigma_1 \leftarrow \widehat{\mathbf{p}}_1$ and $\sigma_2 \leftarrow \widehat{\mathbf{p}}_2$ and *any* function $f : \Sigma_1 \times \Sigma_2 \mapsto [-1, 1]$, $f(\sigma_1, \sigma_2)$ is a random variable with expectation $\mathbf{E}_{\widehat{\mathbf{p}}_1,\widehat{\mathbf{p}}_2}[f]$. Moreover, the payoff of player 1 on playing any *pure* strategy τ_1 when player 2 plays σ_2 is also a random variable whose expectation is precisely $\mathbf{E}_{\sigma_2 \leftarrow \widehat{\mathbf{p}}_2}[f_1(\tau_1, \sigma_2)]$. Define $\mathbf{p}'_1$ ($\mathbf{p}'_2$) to be the mixed strategy of playing each pure strategy in A (resp. B) with probability $1/k$. We will now show that for a suitably large k, $\mathbf{p}'_1, \mathbf{p}'_2$ satisfy the conditions of the theorem.

We first show that the payoff for any player obtained by playing $\mathbf{p}'_1, \mathbf{p}'_2$ is within $\varepsilon/2$ of the payoff obtained by playing $\widehat{\mathbf{p}}_1, \widehat{\mathbf{p}}_2$ with high probability (w.h.p.). This will prove part (2) above. To see this, note that the payoff $\mathbf{E}_{\mathbf{p}'_1,\mathbf{p}'_2}[f_i]$ of player i ($i = 1, 2$) obtained by playing $\mathbf{p}'_1, \mathbf{p}'_2$ is a random variable given by

$$\mathbf{E}_{\mathbf{p}'_1,\mathbf{p}'_2}[f_i] = \frac{1}{k^2} \sum_{\sigma_1 \in A, \sigma_2 \in B} f_i(\sigma_1, \sigma_2).$$

Note that $\mathbf{E}_{\mathbf{p}'_1,\mathbf{p}'_2}[f_i]$ itself is a sum of k^2 *independent* random variables each having expectation $\mathbf{E}_{\widehat{\mathbf{p}}_1,\widehat{\mathbf{p}}_2}[f_i]$. Thus, the standard Chernoff-Hoeffding bounds give that with probability at least $1 - 2e^{-k^2\varepsilon^2/8}$, we have for $i = 1, 2$

$$|\mathbf{E}_{\mathbf{p}'_1,\mathbf{p}'_2}[f_i] - \mathbf{E}_{\widehat{\mathbf{p}}_1,\widehat{\mathbf{p}}_2}[f_i]| \leq \varepsilon/2 \quad \text{for } i = 1, 2. \tag{6}$$

Next, we show that $\mathbf{p}'_1, \mathbf{p}'_2$ form an ε-Nash w.h.p. That is, we show that w.h.p. the payoff of player i doesn't deviate by more than ε when it deviates. Its enough to show this when the player deviates to pure strategies. Let us look at the expected payoff to player 1 when it plays any pure strategy τ_1 while player 2 plays $\mathbf{p}'_2$.

$$\mathbf{E}_{\sigma_2 \leftarrow \mathbf{p}'_2}[f_1(\tau_1, \sigma_2)] \quad = \quad \frac{1}{k} \sum_{\sigma_2 \in B} f_1(\sigma_1, \sigma_2).$$

Note that, as above, each term in the summation of the r.h.s. is a random variable with expectation $\mathbf{E}_{\sigma_2 \leftarrow \widehat{\mathbf{p}}_2}[f_1(\tau_1, \sigma_2)]$. Thus, again via Chernoff-Hoeffding bounds we get that, for any pure strategy τ_1, with probability at least $1 - e^{-k\varepsilon^2/8}$ the following holds

$$\begin{aligned} \mathbf{E}_{\sigma_2 \leftarrow \mathbf{p}'_2}[f_1(\tau_1, \sigma_2)] &\leq \mathbf{E}_{\sigma_2 \leftarrow \widehat{\mathbf{p}}_2}[f_1(\tau_1, \sigma_2)] + \varepsilon/2 \\ &\leq \mathbf{E}_{\widehat{\mathbf{p}}_1, \widehat{\mathbf{p}}_2}[f_1] + \varepsilon/2, \end{aligned}$$

where the second inequality follows from definition of the Nash equilibrium. This implies that for any distribution $\mathbf{p}_1$ on n pure strategies, with probability at least $1 - ne^{-k\varepsilon^2/8}$, we have

$$\mathbf{E}_{\mathbf{p}_1, \mathbf{p}'_2}[f_1] \leq \mathbf{E}_{\widehat{\mathbf{p}}_1, \widehat{\mathbf{p}}_2}[f_1] + \varepsilon/2.$$

Using Equation (6), we conclude that with probability at least $1 - [ne^{-k\varepsilon^2/8} + e^{-k^2\varepsilon^2/8}]$ we have $\mathbf{E}_{\mathbf{p}_1, \mathbf{p}'_2}[f_1] \leq \mathbf{E}_{\mathbf{p}'_1, \mathbf{p}'_2}[f_1] + \varepsilon$. Similarly, arguing for player 2 gives us that with probability at least $1 - [2ne^{-k\varepsilon^2/8} + 2e^{-k^2\varepsilon^2/8}]$, $\mathbf{p}'_1, \mathbf{p}'_2$ is an ε-Nash equilibrium. To argue about the variance, note that Equation (6) holds for *any* function f_i, and in particular for f_i^2. Thus, with probability at least $1 - 4e^{-k^2\varepsilon^2/8}$, we get that for $i = 1, 2$

$$\begin{aligned} \mathbf{E}_{\mathbf{p}'_1, \mathbf{p}'_2}[f_i^2] &\leq \mathbf{E}_{\widehat{\mathbf{p}}_1, \widehat{\mathbf{p}}_2}[f_i^2] + \varepsilon/2, \text{ and} \\ \mathbf{E}_{\mathbf{p}'_1, \mathbf{p}'_2}[f_i] &\geq \mathbf{E}_{\widehat{\mathbf{p}}_1, \widehat{\mathbf{p}}_2}[f_i] - \varepsilon/2. \end{aligned}$$

The second equation above implies $(\mathbf{E}_{\mathbf{p}'_1, \mathbf{p}'_2}[f_i])^2 \geq (\mathbf{E}_{\widehat{\mathbf{p}}_1, \widehat{\mathbf{p}}_2}[f_i])^2 - \varepsilon$ since f_i is bounded in absolute value by 1. By the definition of variance, we get that with probability at least $1 - 4e^{-k^2\varepsilon^2/8}$, for $i = 1, 2$

$$\mathbf{Var}_{\mathbf{p}'_1, \mathbf{p}'_2}[f_i] \leq \mathbf{Var}_{\widehat{\mathbf{p}}_1, \widehat{\mathbf{p}}_2}[f_i] + 2\varepsilon.$$

Summarizing, we get $\mathbf{p}'_1, \mathbf{p}'_2$ satisfies the conditions of the theorem with probability at least

$$1 - 8e^{-k^2\varepsilon^2/8} - 2ne^{-k\varepsilon^2/8}.$$

Choosing $k = \Omega(\ln n/\varepsilon^2)$ makes this probability less than $1/n^{100}$. This implies that there exists sets A, B of size at most k which define, $\mathbf{p}'_1, \mathbf{p}'_2$, satisfying the properties in the theorem. Further, these sets can be found in time $n^{O(k)}$. This completes the proof of the theorem. ∎

Remark A.2 *A straightforward generalization of Theorem 2.2 is to higher moments: not only does the point* $(\mathbf{p}'_1, \mathbf{p}'_2)$ *approximate the second moments, but also higher moments. The bounds of course worsen with the number of moments.*