Approximate and Well-supported Approximate Nash Equilibria of Random Bimatrix Games

Panagiota N. Panagopoulou and Paul G. Spirakis

Department of Computer Engineering and Informatics, Patras University
Research Academic Computer Technology Institute, Greece
E-mails: panagopp@cti.gr, spirakis@cti.gr

Abstract

We focus on the problem of computing approximate Nash equilibria and well-supported approximate Nash equilibria in random bimatrix games, where each player's payoffs are bounded and independent random variables, not necessarily identically distributed, but with common expectations. We show that the completely mixed uniform strategy profile, i.e. the combination of mixed strategies (one per player) where each player plays with equal probability each one of her available pure strategies, is an *almost Nash equilibrium* for random bimatrix games, in the sense that it is, with high probability, an ε-well-supported Nash equilibrium where ε tends to zero as n tends to infinity.

Keywords: Approximate Nash equilibrium, random bimatrix game

1. Introduction

A non-cooperative game in strategic form consists of a set of players, and, for each player, a set of strategies available to her as well as a payoff function mapping each strategy profile (i.e. each combination of strategies, one for each player) to a real number that captures the preferences of the player over the possible outcomes of the game. The most important solution concept in non-cooperative game theory is the notion of *Nash equilibrium* [Nash (1951)]; this is a strategy profile from which no player would have an incentive to deviate, i.e. no player could increase her payoff by choosing another strategy while the rest of the players persevered their strategies.

Despite the certain existence of such equilibria [Nash (1951)], the problem of finding any Nash equilibrium even for games involving only two players has been recently proved to be complete in the PPAD class, introduced by [Papadimitriou (1991)]. This fact emerged the computation of *approximate* Nash equilibria, referred to as ε-Nash equilibria. An ε-Nash equilibrium is a strategy profile such that no deviating player could achieve a payoff higher than the one that the specific profile gives her, plus ε. A stronger notion of approximate Nash equilibria is the *ε-well-supported Nash equilibria*; these are strategy profiles such that each player plays only approximately best-response pure strategies with non-zero probability.

In this work, we focus on the problem of computing an ε-Nash equilibrium as well as an ε-well supported Nash equilibrium of a random $n \times n$ bimatrix game. In a random game, as considered in this work, the entries of each player's payoff matrix are random variables which are drawn independently from some probability distribution on the interval [0, 1]. We do not require that these random variables should be identically distributed, but we assume that, in each matrix, all entries have the same mean (μ_A for the first player and μ_B for the second player). We show that the *completely mixed, uniform strategy profile* where each player plays with equal probability each of her n available pure strategies is, with high probability, a $\sqrt{\ln n / n}$-Nash equilibrium and a $\sqrt{3 \ln n / n}$-well supported Nash equilibrium. This implies that the simple uniform randomization over the set of pure strategies of each player yields an *almost Nash equilibrium* profile, in the sense that it is, with high probability, an ε-well-supported Nash equilibrium with ε tending to zero as the number of available strategies tends to infinity.

Related Work. [Nash (1951)] introduced the concept of Nash equilibria in non-cooperative games and proved that any game possesses at least one such equilibrium. A well-known algorithm for computing a Nash equilibrium of a game with 2 players is the Lemke-Howson algorithm [Lemke and Howson (1964)], however it has exponential worst-case running time [Savani and von Stengel (2004)]. Recently, [Daskalakis and Papadimitriou (2005)] and [Daskalakis et. Al. (2006a)] showed that the problem of computing a Nash equilibrium in a game with 3 or more players is PPAD-complete. [Chen and Deng (2006)] proved that the problem is PPAD-complete for bimatrix games as well.

In [Lipton et. Al. (2003)] it was shown that, for any bimatrix game and for any *constant* ε, there exists an ε-Nash equilibrium with only logarithmic support in the number n of available pure strategies. This result directly yields a quasi-polynomial algorithm for computing such an approximate equilibrium. Moreover, as pointed out in [Althöfer (1994)], no algorithm that examines supports smaller than about $\ln n$ can achieve an approximation better than 1/4. In [Chen et. Al. (2006)] it was shown that the problem of computing a $1/n^{\Theta(1)}$-Nash equilibrium is PPAD-complete, and that bimatrix games are unlikely to have a fully polynomial time approximation scheme (unless PPAD $\subseteq$ P). However, it was conjectured that it is unlikely that finding an ε-Nash equilibrium is PPAD-complete when ε is an absolute constant.

A simple linear-time algorithm for computing a 3/4-Nash equilibrium for any bimatrix game was presented in [Kontogiannis et. Al. (2006)]. This algorithm was then extended so as to obtain a (potentially stronger) parameterized approximation. Namely, [Kontogiannis et. Al. (2006)] suggested an algorithm that computes a

$(2+\lambda)/4$-Nash equilibrium, where λ is the minimum, among all Nash equilibria, expected payoff of either player. The suggested algorithm runs in time polynomial in the number of strategies available to the players. [Daskalakis et. Al. (2006b)] gave a simple, linear-time algorithm examining just two strategies per player and resulting in a 1/2-Nash equilibrium in any bimatrix game. For the more demanding notion of well-supported approximate equilibrium it was shown that the problem can be reduced to the case of win-lose games (where each matrix entry is either 0 or 1), and that an approximation of 5/6 is possible contingent upon a graph-theoretic conjecture.

[Bárány et. Al. (2005)] analyzed a Las Vegas algorithm for finding a Nash equilibrium in 2-player random games. In their model however the matrices entries were considered to be identically distributed (drawn either from the uniform distribution on some interval with zero mean, or from the standard Normal distribution $N(0,1)$). The randomized algorithm proposed in [Bárány et. Al. (2005)] always finds an equilibrium, and an involved analysis of its time complexity shows that it runs in polynomial time (namely $O(n^3 \ln\ln n)$) with high probability.

2. Games and Nash Equilibria

1.1 Notation

For an integer n, let $[n]=\{1,2,\dots,n\}$. For a $n\times 1$ vector $\mathbf{x}$ we denote by $x_1, x_2, \dots x_n$ the components of $\mathbf{x}$ and by $\mathbf{x}^T$ the transpose of $\mathbf{x}$. We denote by $\mathbf{e}_i$ the column vector with a 1 at the ith coordinate and 0 elsewhere; the size of $\mathbf{e}_i$ will be clear from the context. For an $n\times m$ matrix A, we denote a_{ij} the element in the i-th row and j-th column of A. Let $\mathbb{P}^n$ be the set of all probability vectors in n dimensions, i.e. $\mathbb{P}^n \equiv \left\{ \mathbf{x}\in\mathbb{R}^n : \ \sum_{i=1}^{n} x_i = 1 \text{ and } x_i \ge 0 \ \ \forall i\in[n] \right\}$.

For an event E in a sample space, denote $\Pr\{E\}$ the probability of event E occurring. For a random variable X that follows the probability distribution D, denote $E[X]$ the *expectation* of X (according to the probability distribution D).

1.2 Bimatrix games

A *noncooperative game* $\Gamma = \langle N, (S_i)_{i\in N}, (u_i)_{i\in N}\rangle$ consists of (i) a finite set of *players* N, (ii) a nonempty finite set of *pure strategies* S_i for each player $i\in N$ and (iii) a *payoff function* $u_i : \times_{i\in N} S_i \to \mathbb{R}$ for each player $i\in N$.

Bimatrix games [Lemke (1965), Lemke and Howson (1964)], denoted by $\Gamma = \langle A, B \rangle$ are a special case of 2-player games such that the payoff functions can be described by two real $n \times m$ matrices A and B. The n rows of A, B represent the pure strategies of the 1st player (*row player*) and the m columns represent the pure strategies of the 2nd player (*column player*). When the row player chooses strategy i and the column player chooses strategy j, the former gets payoff a_{ij} while the latter gets payoff b_{ij}.

A *mixed strategy* for player $i \in N$ is a probability distribution on the set of her pure strategies S_i. In a bimatrix game $\Gamma = \langle A, B \rangle$, a mixed strategy for the row player can be expressed as a probability vector $\mathbf{x} \in \mathbb{P}^n$ while a mixed strategy for the column player can be expressed as a probability vector $\mathbf{y} \in \mathbb{P}^m$. A *strategy profile* $(\mathbf{x}, \mathbf{y})$ is a combination of strategies, one for each player. In a given strategy profile $(\mathbf{x}, \mathbf{y})$ the players get *expected payoffs* $\mathbf{x}^T A \mathbf{y}$ (row player) and $\mathbf{x}^T B \mathbf{y}$ (column player). The *support* of a player in a strategy profile is the subset of her pure strategies that are assigned strictly positive probability.

We now define the completely mixed, uniform strategy profile:

Definition 2.1 (Completely mixed, uniform strategy profile) *Consider a* $n \times m$ *bimatrix game* $\Gamma = \langle A, B \rangle$. *The completely mixed, uniform strategy profile is the strategy profile* $(\mathbf{x}, \mathbf{y})$ *such that* $x_i = 1/n \quad \forall i \in [n]$ *and* $y_j = 1/m \quad \forall j \in [m]$.

1.3 Nash equilibria and approximate Nash equilibria

A Nash equilibrium [Nash (1951)] for a game Γ is a combination of strategies, one for each player, such that no player could increase her payoff by unilaterally changing her strategy. We formally give two equivalent definitions of a Nash equilibrium:

Definition 2.2 (Nash equilibrium) *A strategy profile* $(\mathbf{x}, \mathbf{y})$ *is a Nash equilibrium for the* $n \times m$ *bimatrix game* $\Gamma = \langle A, B \rangle$ *if*

- *For all pure strategies* $i \in [n]$ *of the row player,* $\mathbf{e}_i^T A \mathbf{y} \le \mathbf{x}^T A \mathbf{y}$ *and*
- *For all pure strategies* $j \in [m]$ *of the column player,* $\mathbf{x}^T B \mathbf{e}_j \le \mathbf{x}^T B \mathbf{y}$.

Definition 2.3 (Nash equilibrium) *A strategy profile* $(\mathbf{x}, \mathbf{y})$ *is a Nash equilibrium for the* $n \times m$ *bimatrix game* $\Gamma = \langle A, B \rangle$ *if*

- *For all pure strategies* $i \in [n]$ *of the row player,*

$$x_i > 0 \Rightarrow \mathbf{e}_i^T A \mathbf{y} \ge \mathbf{e}_k^T A \mathbf{y} \quad \forall k \in [n]$$

- *For all pure strategies* $j \in [m]$ *of the column player,*

$$y_j > 0 \Rightarrow \mathbf{x}^T B\mathbf{e}_j \geq \mathbf{x}^T B\mathbf{e}_k \quad \forall k \in [m].$$

For $\varepsilon > 0$, an ε-Nash equilibrium (or an ε-approximate Nash equilibrium) is a combination of (pure or mixed) strategies, one for each player, such that no player could increase her payoff more than ε by unilaterally changing her strategy:

Definition 2.4 (ε-Nash equilibrium) *For any* $\varepsilon > 0$ *a strategy profile* $(\mathbf{x}, \mathbf{y})$ *is an* ε*-Nash equilibrium for the* $n \times m$ *bimatrix game* $\Gamma = \langle A, B \rangle$ *if*

- *For all pure strategies* $i \in [n]$ *of the row player,* $\mathbf{e}_i^T A\mathbf{y} \leq \mathbf{x}^T A\mathbf{y} + \varepsilon$ *and*
- *For all pure strategies* $j \in [m]$ *of the column player,* $\mathbf{x}^T B\mathbf{e}_j \leq \mathbf{x}^T B\mathbf{y} + \varepsilon$.

A stronger notion of approximate Nash equilibria was introduced in [Goldberg and Papadimitriou (2006), Daskalakis et. Al. (2006a)]: For $\varepsilon > 0$, an ε-well supported Nash equilibrium (or a well-supported ε-approximate Nash equilibrium) is a combination of (pure or mixed) strategies, one for each player, such that each player plays only approximately best-response pure strategies with non-zero probability:

Definition 2.5 (ε-well-supported Nash equilibrium) *For any* $\varepsilon > 0$ *a strategy profile* $(\mathbf{x}, \mathbf{y})$ *is an* ε*-well-supported Nash equilibrium for the* $n \times m$ *bimatrix game* $\Gamma = \langle A, B \rangle$ *if*

- *For all pure strategies* $i \in [n]$, $x_i > 0 \Rightarrow \mathbf{e}_i^T A\mathbf{y} \geq \mathbf{e}_k^T A\mathbf{y} - \varepsilon \quad \forall k \in [n]$
- *For all pure strategies* $j \in [m]$, $y_j > 0 \Rightarrow \mathbf{x}^T B\mathbf{e}_j \geq \mathbf{x}^T B\mathbf{e}_k - \varepsilon \quad \forall k \in [m]$.

Every ε-well-supported Nash equilibrium is also an ε-Nash equilibrium, and as pointed out in [Chen et. Al. (2006)], given an $\varepsilon/8n$-Nash equilibrium we can compute in polynomial time an ε-well-supported Nash equilibrium, for every $\varepsilon > 0$.

We define an *almost Nash equilibrium* as an ε-well-supported Nash equilibrium such that ε tends to 0 as the number of the available pure strategies tends to infinity:

Definition 2.6 (Almost Nash equilibrium) *A strategy profile* $(\mathbf{x}, \mathbf{y})$ *is an almost Nash equilibrium for the* $n \times n$ *bimatrix game* $\Gamma = \langle A, B \rangle$ *if*

- *For all pure strategies* $i \in [n]$ *of the row player,*

$$x_i > 0 \Rightarrow \mathbf{e}_i^T A\mathbf{y} \geq \mathbf{e}_k^T A\mathbf{y} - \varepsilon(n) \quad \forall k \in [n]$$

- *For all pure strategies* $j \in [n]$ *of the column player,*

$$y_j > 0 \Rightarrow \mathbf{x}^T B\mathbf{e}_j \geq \mathbf{x}^T B\mathbf{e}_k - \varepsilon(n) \quad \forall k \in [n]$$

- $\lim_{n \to \infty} \varepsilon(n) = 0$.

1.4 Random bimatrix games

As pointed out in [Chen et. Al. (2006)], the set of Nash equilibria of a bimatrix game remains precisely the same if we multiply all entries of a matrix by a positive constant or if we add the same constant to each entry. Therefore, it suffices to consider bimatrix games with normalized matrices so as to study their complexity. We adopt the normalization used in [Lipton et. Al. (2003)]: we assume that the value of each entry in the matrices is lower bounded by 0 and upper bounded by 1. Such games are referred to as *positively normalized* [Chen et. Al. (2006)]. Furthermore, in this work we focus on *random* bimatrix games:

Definition 2.7 (Random bimatrix game) *A* $n \times m$ *random bimatrix game* $\Gamma = \langle A, B \rangle$ *is* $n \times m$ *bimatrix game such that*

- *all elements of matrix* A *are independent random variables, each taking a value in the interval* $[0,1]$ *and with expectation* $\mu_A \in [0,1]$*, and*
- *all elements of matrix* B *are independent random variables, each taking a value in the interval* $[0,1]$ *and with expectation* $\mu_B \in [0,1]$.

Note that, according to the above definition, the entries of each matrix of a bimatrix game need not be identically distributed; it suffices that they all have the same expectation.

3. Approximate Nash Equilibria in Random Games

In the following, we deal with the problem of computing an ε-Nash equilibrium of a random $n \times n$ bimatrix game $\Gamma = \langle A, B \rangle$. We show that the completely mixed, uniform strategy profile is, with high probability (w.h.p.), a $\sqrt{\ln n / n}$-Nash equilibrium. For the proof we will use the following lemma:

Lemma 3.1 [Hoeffding (1963)] *If* $Y_1, \ldots, Y_m, Z_1, \ldots, Z_n$ *are independent random variables with values in the interval* $[a,b]$*, and if* $\overline{Y} = (Y_1 + \cdots + Y_m)/m$ *and* $\overline{Z} = (Z_1 + \cdots + Z_n)/n$*, then for any* $t > 0$*,*

$$\Pr\left\{\overline{Y} - \overline{Z} - \left(E[\overline{Y}] - E[\overline{Z}]\right) \geq t\right\} \leq \exp\left(-\frac{2t^2}{\left(\frac{1}{m} + \frac{1}{n}\right)(b-a)^2}\right).$$

Theorem 3.2 *Let* $\Gamma = \langle A, B \rangle$ *be a* $n \times m$ *random bimatrix game and let* $(\mathbf{x}, \mathbf{y})$ *be the completely mixed, uniform strategy profile for* Γ*. Then, for any* $\varepsilon > 0$*,*

$$\Pr\left\{(\mathbf{x},\mathbf{y}) \text{ is an } \varepsilon\text{-Nash equilibrium}\right\} \geq 1-\left(n\exp(-2m\varepsilon^2)+m\exp(-2n\varepsilon^2)\right).$$

Proof. Fix $r \in [n]$. Then it holds that

$$\mathbf{e}_r^T A\mathbf{y} = \sum_{j\in[m]} a_{rj}y_j = \frac{1}{m}\sum_{j\in[m]} a_{rj} \quad \text{and} \quad \mathbf{x}^T A\mathbf{y} = \sum_{i\in[n]}\sum_{j\in[m]} a_{ij}x_iy_j = \frac{1}{nm}\sum_{i\in[n]}\sum_{j\in[m]} a_{ij}.$$

Now,

$$\mathbf{e}_r^T A\mathbf{y} - \mathbf{x}^T A\mathbf{y} = \frac{1}{m}\sum_{j\in[m]} a_{rj} - \frac{1}{nm}\sum_{i\in[n]}\sum_{j\in[m]} a_{ij} = \frac{n-1}{nm}\sum_{j\in[m]} a_{rj} - \frac{1}{nm}\sum_{i\in[n],i\neq r}\sum_{j\in[m]} a_{ij},$$

and therefore, for any $\varepsilon > 0$,

$$\Pr\left\{\mathbf{e}_r^T A\mathbf{y} - \mathbf{x}^T A\mathbf{y} \geq \varepsilon\right\} = \Pr\left\{\frac{\sum_{j\in[m]} a_{rj}}{m} - \frac{\sum_{i\in[n],i\neq r}\sum_{j\in[m]} a_{ij}}{(n-1)m} \geq \frac{n}{n-1}\varepsilon\right\}.$$

Now we can apply Lemma 3.1: all random variables are mutually independent, a=0 and b=1, $t = n\varepsilon/(n-1)$. Moreover,

$$E\left[\frac{1}{m}\sum_{j\in[m]} a_{rj}\right] = E\left[\frac{1}{(n-1)m}\sum_{i\in[n],i\neq r}\sum_{j\in[m]} a_{ij}\right] = \mu_A.$$

So

$$\Pr\left\{\mathbf{e}_r^T A\mathbf{y} - \mathbf{x}^T A\mathbf{y} \geq \varepsilon\right\} \leq \exp\left(-\frac{2\frac{n^2}{(n-1)^2}\varepsilon^2}{\left(\frac{1}{m}+\frac{1}{(n-1)m}\right)(1-0)^2}\right) \leq \exp(-2m\varepsilon^2)$$

$$\Pr\left\{\exists r \in [n] : \mathbf{e}_r^T A\mathbf{y} - \mathbf{x}^T A\mathbf{y} \geq \varepsilon\right\} \leq \sum_{r\in[n]} \Pr\left\{\mathbf{e}_r^T A\mathbf{y} - \mathbf{x}^T A\mathbf{y} \geq \varepsilon\right\} \leq n\exp(-2m\varepsilon^2).$$

Following the same reasoning as above we can show that, for any $\varepsilon > 0$,

$$\Pr\left\{\exists c \in [m] : \mathbf{x}^T B\mathbf{e}_c - \mathbf{x}^T B\mathbf{y} \geq \varepsilon\right\} \leq m\exp(-2n\varepsilon^2).$$

Therefore, for any $\varepsilon > 0$,

$$\Pr\left\{(\mathbf{x},\mathbf{y}) \text{ is an } \varepsilon\text{-Nash equilibrium}\right\} \geq 1-\left(n\exp(-2m\varepsilon^2)+m\exp(-2n\varepsilon^2)\right).$$

Corollary 3.3 *Consider a* $n\times n$ *random bimatrix game* Γ. *Then the completely mixed, uniform strategy profile is, with high probability, a* $\sqrt{\ln n/n}$ *-Nash equilibrium for* Γ.

Proof. Let $(\mathbf{x},\mathbf{y})$ denote the completely mixed, uniform strategy profile for Γ. By

Theorem 3.2, $\Pr\{(\mathbf{x},\mathbf{y}) \text{ is an } \varepsilon\text{-Nash equilibrium}\} \geq 1-2n\exp(-2n\varepsilon^2)$. Setting $\varepsilon = \sqrt{\ln n/n}$, it follows that the uniform strategy profile is a $\sqrt{\ln n/n}$-Nash equilibrium for $\langle A, B\rangle$ with probability

$$\Pr\left\{(\mathbf{x},\mathbf{y}) \text{ is a } \sqrt{\frac{\ln n}{n}}\text{-Nash equilibrium}\right\} \geq 1-2n\exp\left(-2n\frac{\ln n}{n}\right) = 1-\frac{2}{n}.$$

4. *Well-supported Nash Equilibria in Random Games*

We will now show that the completely mixed, uniform strategy profile is, w.h.p., a $\sqrt{3\ln n/n}$-well supported Nash equilibrium for a $n\times n$ random bimatrix game.

Theorem 4.1 *Let $\Gamma = \langle A, B\rangle$ be a $n\times n$ random bimatrix game. Then the completely mixed, uniform strategy profile is a $\sqrt{3\ln n/n}$-well supported Nash equilibrium for Γ, with high probability.*

Proof. By definition, the completely mixed, uniform strategy profile $(\mathbf{x},\mathbf{y})$ is an ε-well supported Nash equilibrium if and only if, for all $i, j \in [n]$,

$$\begin{array}{rclcrcl}
\mathbf{e}_i^T A\mathbf{y} & \geq & \mathbf{e}_j^T A\mathbf{y} - \varepsilon & & \mathbf{x}^T B\mathbf{e}_i & \geq & \mathbf{x}^T B\mathbf{e}_j - \varepsilon \\
\mathbf{e}_j^T A\mathbf{y} - \mathbf{e}_i^T A\mathbf{y} & \leq & \varepsilon & \text{and} & \mathbf{x}^T B\mathbf{e}_j - \mathbf{x}^T B\mathbf{e}_i & \leq & \varepsilon \\
\frac{1}{n}\sum_{k=1}^n a_{jk} - \frac{1}{n}\sum_{k=1}^n a_{ik} & \leq & \varepsilon & & \frac{1}{n}\sum_{k=1}^n b_{kj} - \frac{1}{n}\sum_{k=1}^n b_{ki} & \leq & \varepsilon .
\end{array}$$

Now fix $i, j \in [n]$. Then

$$E\left[\frac{1}{n}\sum_{k=1}^n a_{jk}\right] = \mu_A \quad \text{and} \quad E\left[\frac{1}{n}\sum_{k=1}^n b_{kj}\right] = \mu_B$$

and therefore, by Lemma 3.1,

$$\Pr\left\{\frac{1}{n}\sum_{k=1}^n a_{jk} - \frac{1}{n}\sum_{k=1}^n a_{ik} \geq \varepsilon\right\} \leq \exp\left(-\frac{2\varepsilon^2}{\left(\frac{1}{n}+\frac{1}{n}\right)(1-0)^2}\right) = \exp\left(-n\varepsilon^2\right),$$

and similarly

$$\Pr\left\{\frac{1}{n}\sum_{k=1}^n b_{kj} - \frac{1}{n}\sum_{k=1}^n b_{ki} \geq \varepsilon\right\} \leq \exp\left(-n\varepsilon^2\right).$$

Thus

$$\begin{aligned}&\Pr\{(\mathbf{x},\mathbf{y}) \text{ is a } \varepsilon\text{-well-supported Nash equilibrium}\}\\&\quad = 1-\Pr\left\{\exists i,j : \mathbf{e}_i^T A\mathbf{y} \geq \mathbf{e}_j^T A\mathbf{y}-\varepsilon \text{ or } \mathbf{x}^T B\mathbf{e}_i \geq \mathbf{x}^T B\mathbf{e}_j-\varepsilon\right\}\\&\quad \geq 1-\binom{n}{2}\cdot 2\cdot \exp\left(-n\varepsilon^2\right) = 1-n(n-1)\exp\left(-n\varepsilon^2\right).\end{aligned}$$

Setting $\varepsilon = \sqrt{3\ln n/n}$ we get

$$\begin{aligned}&\Pr\{(\mathbf{x},\mathbf{y}) \text{ is a } \sqrt{\frac{3\ln n}{n}}\text{-well-supported Nash equilibrium}\}\\&\quad \geq 1-n(n-1)\exp\left(-n\frac{3\ln n}{n}\right) \geq 1-n^2\cdot\frac{1}{n^3} = 1-\frac{1}{n}.\end{aligned}$$

Corollary 4.2 *The completely mixed, uniform strategy profile is an almost Nash equilibrium for random bimatrix games, with high probability.*

5. Conclusions

We have shown that the simple strategy profile where each player plays with equal probability each of her available pure strategies is possibly an almost Nash equilibrium for the random games we consider. Our model of a random bimatrix game is less restrictive than the one considered in [Bárány et. Al. (2005)], since it allows for non-identically distributed payoffs. The results presented here give a clear and straightforward method for computing almost Nash equilibrium strategies in random games, without searching among possible supports or computing the probabilities that the pure strategies of each player should be played with.

Acknowledgments. This work was partially supported by the EU within the 6th Framework Programme under contract 015964 “Algorithmic Principles for Building Efficient Overlay Computers” (*AEOLUS*), and by the General Secretariat for Research and Technology of the Greek Ministry of Development within the Programme *PENED 2003*.

References

Althöfer, I., (1994), *On sparse approximations to randomized strategies and convex combinations*, Linear Algebra and Applications 199, 339–355.

Bárány, I., Vempala, S., Vetta, A., (2005), *Nash equilibria in random games*, Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’05), pp. 123–131.

Chen, X., Deng, X., (2006), *Settling the complexity of 2-player Nash equilibrium*, Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06).

Chen, X., Deng, X., Teng, S. H., (2006), *Computing Nash equilibria: Approximation and smoothed complexity*, Electronic Colloquium on Computational Complexity (ECCC).

Daskalakis, C., Goldberg, P., Papadimitriou, C., (2006a), *The complexity of computing a Nash equilibrium*, Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC'06), pp. 71–78.

Daskalakis, C., Mehta, A., Papadimitriou, C., (2006b), *A note on approximate Nash equilibria*, Proceedings of the 2nd Workshop on Internet and Network Economics (WINE'06), pp. 297–306.

Daskalakis, C., Papadimitriou, C., (2005), *Three-player games are hard*, Electronic Colloquium on Computational Complexity (ECCC).

Goldberg, P., Papadimitriou, C., (2006), *Reducibility among equilibrium problems*, Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC'06), pp. 61–70.

Hoeffding, W., (1963), *Probability inequalities for sums of bounded random variables*, Journal of the American Statistical Association 58, 13–30.

Kontogiannis, S., Panagopoulou, P. N., Spirakis, P. G., (2006), *Polynomial algorithms for approximating Nash equilibria of bimatrix games*, Proceedings of the 2nd Workshop on Internet and Network Economics (WINE'06), pp. 286–296.

Lemke, C. E., (1965), *Bimatrix equilibrium points and mathematical programming.* Management Science 11, 681–689.

Lemke, C. E., Howson, J. T., (1964), *Equilibrium points of bimatrix games*, J. Soc. Indust. Appl. Math. 12, 413–423.

Lipton, R. J., Markakis, E., Mehta, A., (2003), *Playing large games using simple startegies*, Proceedings of the 4th ACM Conference on Electronic Commerce (EC'03), pp. 36–41.

Nash, J., (1951), *Noncooperative games*, Annals of Mathematics 54, 289–295.

Papadimitriou, C. H., (1991), *On inefficient proofs of existence and complexity classes*, Proceedings of the 4th Czechoslovakian Symposium on Combinatorics.

Savani, R., von Stengel, B., (2004), *Exponentially many steps for finding a Nash equilibrium in a bimatrix game*, Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS'04), pp. 258–267.