# Computational aspects of Nash equilibria

Uriel Feige

May 21, 2024

## 1 Two player Nash equilibria

We sometimes use the word Nash as an abbreviation for "Nash equilibrium".

In some senses, computing Nash for two player games is an easier task than for three player games.

Consider a two player game with payoff matrices $R$ and $C$ for the row and column players. Mixed strategies for the row and column players are represented as vectors $x$ and $y$ of probabilities over choices of pure strategies. The vector associated with pure strategy $i$ for the row player is denoted by $x^i$, and it has 1 in its $i$th coordinate, and 0 elsewhere. The notation $y^j$ has a similar interpretation. Given supports $S^R$ for the row player and $S^C$ for the column player, to form a Nash, the mixed strategies $x$ and $y$ associated with them need to satisfy the following linear constraints:

- Nonnegativity: $x, y \geq 0$.
- Probability distribution: $\sum x_i = \sum y_j = 1$.
- Best response for the row player: ${x^i}^T Ry \geq {x^j}^T Ry$ for every pure strategies $i \in S^R$ and $j$. These are linear constraints for the entries of the vector $y$.
- Best response for the column player: $x^T Cy^i \geq x^T Cy^j$ for every pure strategies $i \in S^C$ and $j$. These are linear constraints for the entries of the vector $x$.

The probabilities involved in the mixed strategies are solutions to a linear program and hence (when taking a basic feasible solution) are rational and can be represented using polynomially many bits.

The Lemke-Howson algorithm (see Chapter 2 in [NRTV07]) can be used for finding a Nash equilibria in two player games. It also provides an alternative proof that in two player games, Nash equilibria must exist (and hence also proves the minimax theorem). It is not a polynomial time algorithm. There are known families of examples that require an exponential number of steps to reach a Nash equilibrium [SavaniStengel04].

It is not known whether there is a polynomial time algorithm for computing a two-player Nash. It is know that many variations of this question are NP-hard. For example, we have the following theorem.

**Theorem 1.1** *In a two-player game with non-negative payoffs, the following computational tasks are NP-hard.*

- *To determine whether there is a Nash equilibrium in which the column player has expected positive payoff.*
- *To determine whether there is a Nash equilibrium in which the row player has a "strictly" mixed strategy (rather than a pure strategy).*

**Proof:** The proof of both parts of the theorem is by the same reduction from the NP-hard max-clique problem. The input to the max-clique problem is a graph $G$ and an integer parameter $k$, and the question is whether $G$ has a clique of size (at least) $k$.

Given a graph $G$ with $n$ vertices and an integer parameter $k \leq n$, consider the following 2-player game that we refer to as the $k$-clique game. The payoff matrix for the row player has $n+1$ rows and $n$ columns. The top $n$ by $n$ submatrix is $A$, the adjacency matrix of a graph $G$. All the entries of the bottom row are $\frac{k-1}{k}$. For the column player, the top $n$ by $n$ submatrix is the identity matrix $I$, and the bottom row is all 0.

The theorem is proved by establishing the following claims.

1. If $G$ has a clique of size $k$, then there is a Nash equilibrium for the $k$-clique game in which the support for the row player has at least $k$ rows, and the expected payoff for the column player is positive.
2. If $G$ does not have a clique of size $k$, then in all Nash equilibria the support for the row player is the last row, his payoff is $\frac{k-1}{k}$, and the payoff for the column player is 0.

If $S$ is a *maximal* clique of size at least $k$, then both players playing uniformly over $S$ can be seen to be a Nash equilibrium, with payoff $1 - 1/|S|$ for the row player and $1/|S|$ for the column player. This proves claim 1.

We now prove claim 2. We shall refer to a Nash in which the row player plays row $n+1$, the payoff for the row player is $\frac{k-1}{k}$ and the payoff for the column player is 0, as a *standard* Nash equilibrium. We claim that if $G$ has no clique of size $k$, there is no non-standard Nash.

Suppose for the sake of contradiction that there is a non-standard Nash. Let $x_i$ be the probability that the $i$th column is played in this Nash, and let $S$ be the support for the mixed strategy for the column player (namely, $S = \{i | x_i > 0\}$). $S$ must also be in the support for the row player (otherwise the payoff for one of the strategies of the column player is 0, implying that this is the payoff for all the column player's strategies, implying that only row $n+1$ is played, and hence this is a standard Nash). Being in the support of a Nash, all rows of $S$ give the row player the same expected payoff, and it must be at least $\frac{k-1}{k}$ (otherwise the row player plays row $n+1$ instead). Let $y$ range over nonnegative vectors whose entries sum to 1, and recall that $A$ is the adjacency matrix of $G$. We have thus established that $\max_y y^T Ay \geq \frac{k-1}{k}$, by taking $y = (x_1, \ldots x_n)$. The well known Motzkin-Straus theorem [MotzkinStraus65] says that for every adjacency matrix $A$, the exisence of such a $y$ implies that the respective graph $G$ contains a $k$-clique.

For completeness, we sketch a proof of the Motzkin-Straus theorem. Since $A$ is a 0/1 matrix with 0 along the diagonal, then for every non-negative vector $y$ whose sum of entries is 1 and with support $S$, it is easy to see that $y^T Ay \leq 1 - \frac{1}{|S|}$, with equality only if $S$ is a clique and $y$ is uniform over its support. Hence $y$ needs to have support at least $k$. If $|S| > k$ and $S$ is not a clique, then consider any two vertices $u$ and $v$ that are not adjacent in $G$. For every $\delta$, replacing $y$ by changing $y_u$ to $y_u + \delta$ and changing $y_v$ to $y_v - \delta$ changes the objective $y^T Ay$ by a linear function of $\delta$. Hence $\delta$ can be chosen in a way that reduces the support of $y$ (makes either $y_u = 0$ or $y_v = 0$), without reducing the objective. Continuing in this fashion, we must end up with a clique of size at least $k$. $\square$

# 2 PPAD

The Lemke-Howson algorithm places Nash in the class PPAD [Papadimitriou94], standing for *Polynomial Parity Argument Directed*. This class captures problem that can be represented as directed graph, where vertices of the graph represent states in which an algorithm that searches for a solution can be in. (There can be exponentially many such states.) Importantly, each vertex has at most one incoming edge and at most one outgoing edge. One is given a source (a vertex with no incoming edge and one outgoing edge), and is required to output either a sink, or a different source.

Such a graph can easily be associated with the proof of the two-dimensional Sperner lemma. (The proof of Sperner's lemma presented in the previous lecture involves an undirected graph. One can make it directed by requiring that red-blue edges are crossed with red on the right-hand side.)

Two player Nash is complete for PPAD [ChenDeng06] (even for $\epsilon$-Nash for polynomially small $\epsilon$ [ChenDengTeng06]). As computing an $\epsilon$-Nash for multi-player games is in PPAD [DGP06], this implies a reduction from multi player games to two player games, preserving the concept of $\epsilon$-Nash (though $\epsilon$ needs to be very small for the two player game). For other problems in PPAD, see [Papadimitriou94].

Some problems related to PPAD computations are known to be very difficult. This is touched upon in the homework assignment. Finding the other end of a PPAD computation is PSPACE complete (since reversible computation is universal [Bennett73]). Reversible Turing machines are not in PPAD because in a PPAD problem one is allowed to output any sink (or any nonstandard source) and in reversible computation one seeks a particular sink.

# 3 Computing $\epsilon$-Nash, for $\epsilon$ not too small

In [LMM03] it is proved that for every two player game with $n$ by $n$ payoff matrices, there are $\epsilon$-Nash that are supported on $O(\frac{1}{\epsilon^2}\log n)$ strategies. Moreover, the distribution over these strategies may be taken to be uniform, if the same strategy is allowed to appear several times in the support (the support is considered as a multi-set rather than a set.) This, together with the fact that checking whether a pair of mixed strategies is an $\epsilon$-Nash can be done in polynomial time, implies that an $\epsilon$-Nash can be computed in time $n^{O(\frac{1}{\epsilon^2}\log n)}$, by exhaustive search over all possible small supports. Obstacles towards substantially improving this running time are presented in [Rubinstein16].

The proof that an $\epsilon$-Nash of small support exists is by the probabilistic method. We shall use standard bounds on large deviations for sums of independent random variables. The following is a special case of Hoeffding's inequality.

**Lemma 3.1** *Let $X_1, \ldots, X_k$ be independent random variables, bounded between 0 and 1, and let $S_k = \sum X_i$. Then for all $t > 0$,*

$$Pr(S_k - E[S_k] \geq t) \leq e^{-\frac{2t^2}{k}}$$

**Theorem 3.2** *Consider a two player game with payoff matrices $R$ and $C$ for the row and column players, where payoffs are between 0 and 1, and matrices are of order $n$. Let $c$ be a sufficiently large constant ($c = 12$ is chosen in [LMM03]). For any Nash equilibrium $x^*, y^*$ and for any $\epsilon > 0$, there exists, for every $k \geq \frac{12 \ln n}{\epsilon^2}$, a pair of $k$-uniform strategies $x', y'$, such that:*

1. *$x', y'$ is an $\epsilon$-Nash.*
2. *$|x'^T Ry' - x^{*T} Ry^*| < \epsilon$ (the row player gets almost the same payoff as in the Nash equilibrium).*
3. *$|x'^T Cy' - x^{*T} Cy^*| < \epsilon$ (the column player gets almost the same payoff as in the Nash equilibrium).*

**Proof:** For $1 \leq i, j \leq n$, let $x^i$ denote the $i$th pure strategy of the row player and $y^j$ denote the $j$th pure strategy of the column player.

We form a multi-set $A$ by sampling with repetitions $k$ strategies for the row player, from the distribution $x^*$. We form a multi-set $B$ by sampling with repetitions $k$ strategies for the column player, from the distribution $y^*$. The strategy $x'$ is uniformly mixed over $A$, and $y'$ is uniformly mixed over $B$. We claim that with positive probability $x', y'$ satisfy the theorem.

Consider any pure strategy $x^i$ for the row player. Then Lemma 3.1 easily shows that for every constant $b > 0$ it holds that $Pr(|{x^i}^T Ry' - {x^i}^T Ry^*| \geq b\epsilon) \leq 2e^{-\frac{2k^2b^2\epsilon^2}{k}} = 2n^{-2cb^2}$. Hence, taking $b = \frac{1}{2}$ and $c$ sufficiently large, the union bound over $n$ choices of $i$ and $n$ choices of $j$ shows that with positive probability, for all $i$ and $j$ we have that:

1. $|{x^i}^T Ry' - {x^i}^T Ry^*| \leq \frac{\epsilon}{2}$.
2. $|{x'}^T Cy^j - {x^*}^T Cy^j| \leq \frac{\epsilon}{2}$.

All inequalities in the statement of the theorem are straightforward consequences of the above two sets of inequalities. $\square$

# 4 Brief survey of additional related work

The following theorem is due to Bubelis [Bubelis79].

**Theorem 4.1** *Every game $G$ with $d > 3$ players and rational payoffs can be reduced in polynomial time to a game $G'$ with three players and rational payoffs in which player 1 in $G'$ simulates all players in $G$ in the following sense. In every Nash for $G'$, the mixed strategy $x$ for player 1 corresponds (after scaling by $d$) to a Nash of $G$. Every Nash profile $x$ of $G$ corresponds (after scaling by $1/d$) to a mixed strategy of player 1 in $G'$ that is part of a Nash for $G'$.*

We introduced a notion of $\epsilon$-Nash. A related notion of an $\epsilon$-close Nash is perhaps too strong, as there is evidence that checking whether a strategy of profiles in an $\epsilon$-close Nash might not be doable in polynomial time. See details in [EtessamiYannakakis07].

# References


[Bennett73] C. H. Bennett, Logical reversibility of computation, IBM Journal of Research and Development, vol. 17, no. 6, pp. 525-532, 1973.

[Bubelis79] V. Bubelis. On Equilibria in Finite Games. *Int. Journal of Game Theory, Vol 8, Issue 2, page 65- 79*, 1979.

[ChenDeng06] Xi Chen, Xiaotie Deng: Settling the Complexity of Two-Player Nash Equilibrium. FOCS 2006: 261-272

[ChenDengTeng06] Xi Chen, Xiaotie Deng, Shang-Hua Teng: Computing Nash Equilibria: Approximation and Smoothed Complexity. FOCS 2006: 603-612

[DGP06] Constantinos Daskalakis, Paul W. Goldberg, Christos H. Papadimitriou: The complexity of computing a Nash equilibrium. STOC 2006: 71-78

[EtessamiYannakakis07] Kousha Etessami, Mihalis Yannakakis: On the Complexity of Nash Equilibria and Other Fixed Points. FOCS 2007: 113-123

[LMM03] Richard J. Lipton, Evangelos Markakis, Aranyak Mehta: Playing large games using simple strategies. ACM Conference on Electronic Commerce 2003: 36–41.

[MotzkinStraus65] Motzkin T.S., Straus E.G.: Maxima for graphs and a new proof of a theorem of Turan. Can. J. Math. 17, 533–540 (1965).

[NRTV07] Noam Nisan, Tim Roughgarden, Eva Tardos and Vijay V. Vazirani (Editors), Algorithmic Game Theory, Cambridge University Press, 2007.

[Papadimitriou] Christos Papadimitriou. Computational Complexity. Addison-Wesley. 1994.

[Papadimitriou94] Christos H. Papadimitriou: On the Complexity of the Parity Argument and Other Inefficient Proofs of Existence. J. Comput. Syst. Sci. 48(3): 498-532 (1994).

[Rubinstein16] Aviad Rubinstein: Settling the Complexity of Computing Approximate Two-Player Nash Equilibria. FOCS 2016: 258–265.

[SavaniStengel04] Rahul Savani, Bernhard von Stengel: Exponentially Many Steps for Finding a Nash Equilibrium in a Bimatrix Game. FOCS 2004: 258-267