

Inapproximability of NP-Complete Variants of Nash Equilibrium*

Per Austrin Mark Braverman Eden Chlamtáč

Received October 13, 2011; Revised January 6, 2013; Published January 28, 2013

Abstract: In recent work of Hazan and Krauthgamer (SICOMP 2011), it was shown that finding an ε -approximate Nash equilibrium with near-optimal value in a two-player game is as hard as finding a hidden clique of size $O(\log n)$ in the random graph $G(n, 1/2)$. This raises the question of whether a similar intractability holds for approximate Nash equilibrium without side constraints such as high value. We give evidence that asking for near-optimal value makes the problem distinctly harder: a simple algorithm finds a $1/2$ -approximate equilibrium of optimal value, but getting below $1/2$ is as hard as the Hidden Clique problem. This is in contrast to the basic problem (finding a Nash equilibrium with no optimization criteria) where more sophisticated algorithms, achieving better approximations, are known.

Unlike basic Nash equilibrium, which is in PPAD, optimal (maximum value) Nash equilibrium is NP-hard. We proceed to show that optimal Nash equilibrium is just one of several known NP-hard problems related to Nash equilibrium, all of which have approximate variants which are as hard as finding a planted clique. In particular, we show this for approximate variants of the following problems: finding a Nash equilibrium with value greater than η (for any fixed $\eta > 0$, even when the optimal Nash equilibrium has value $1 - \eta$), finding a second Nash equilibrium, and finding a Nash equilibrium with small support.

*A preliminary version of this work appeared in the Proceedings of APPROX 2011 [2]. The results in the current version are essentially the same, but several minor flaws in the proofs have been fixed. Work done while the authors were at the University of Toronto, supported by NSERC Discovery Grants. MB was also supported by a Sloan Research Fellowship.

ACM Classification: F.2

AMS Classification: 68Q17, 68W25, 91A05

Key words and phrases: Nash equilibrium, hidden clique

Finally, we consider the complexity of approximate pure Bayes-Nash equilibria in two-player games. Here we show that for general Bayesian games the problem is NP-hard. For the special case where the distribution over players' types is uniform, we give a quasi-polynomial time algorithm matched by a hardness result based on the Hidden Clique problem.

1 Introduction

The classical notion of Nash equilibrium is the most fundamental concept in the theory of non-cooperative games. In recent years, there has been much work on the computational complexity of finding a Nash equilibrium when the game is given as an input. In particular, a series of hardness results culminated in the work of Chen, Deng, and Teng [5], who showed that even for two-player (bimatrix) games, the problem of computing a Nash equilibrium is PPAD-complete, and thus unlikely to be solvable in polynomial time. We only consider bimatrix games in this paper.

Therefore, it makes sense to consider the complexity of approximate equilibria. In particular, a notion which has emerged as the focus of several works is that of an ε -approximate Nash equilibrium, or ε -equilibrium for short, where neither player can gain more than ε (additively) by defecting to a different strategy (without loss of generality, all payoffs are scaled to lie in the interval $[0, 1]$). A sampling argument of Lipton, Markakis, and Mehta [18] shows that in every game, there exist ε -equilibria with support $O(\ln n/\varepsilon^2)$, and so they can be found in quasi-polynomial time $n^{O(\ln n/\varepsilon^2)}$ by exhaustive search.

However, finding good *polynomial* time approximations has proved more challenging. While finding a $1/2$ -equilibrium turns out to be quite simple [10], more complicated algorithms have given a series of improvements [9, 3, 21], where the best currently known is the 0.3393-equilibrium shown by Tsaknakis and Spirakis [21]. A major open question in this area is whether or not there exists a PTAS for Nash equilibrium, i. e., an algorithm which for any $\varepsilon > 0$ finds an ε -equilibrium in time $n^{f(\varepsilon)}$ for some function f depending only on ε . Note that the algorithm of Lipton et al. gives a quasi-polynomial time approximation scheme (QPTAS) for the problem.

Recently, Hazan and Krauthgamer [16] have attempted to provide evidence for the optimality of the QPTAS of Lipton et al. [18], by showing a reduction from a well-studied and seemingly intractable problem (which can also be solved in quasi-polynomial time) to the related problem of finding an ε -equilibrium with near maximum value (the value of an equilibrium is the average of the payoffs of the two players).

The problem they reduce from is the *k*-Hidden Clique Problem:

Definition 1.1 (*k*-Hidden Clique Problem). Given a graph sampled from $G(n, 1/2)$ with a planted (but hidden) clique of size k , find the planted clique.

Since with high probability the maximum clique in $G(n, 1/2)$ is of size $(2 - o(1)) \log_2 n$, it is easy to see that for constant $\delta > 0$, one can distinguish between $G(n, 1/2)$ and $G(n, 1/2)$ with a planted clique of size $k > (2 + \delta) \log_2 n$ in quasi-polynomial time by exhaustive search over all subsets of $(2 + \delta) \log_2 n$ vertices. It is also not hard to extend this to an algorithm which finds the hidden clique in quasi-polynomial time.

The best polynomial time algorithm known, due to Alon, Krivelevich, and Sudakov [1], solves the $\Omega(\sqrt{n})$ -Hidden Clique Problem (that is, it solves the $c\sqrt{n}$ -Hidden Clique Problem for any constant $c > 0$,

and sufficiently large n). In fact, Feige and Krauthgamer [12] show that even extending this approach by using the Lovász-Schrijver SDP hierarchy, one still requires $\Omega(\ln n)$ levels of the hierarchy (corresponding to $n^{\Omega(\ln n)}$ running time to solve the SDP) just to find a hidden clique of size $n^{1/2-\varepsilon}$. The only approach we are aware of for breaking the $\Omega(\sqrt{n})$ barrier would still (assuming certain conjectures) only discover cliques of size $\Omega(n^c)$ for some constant $c > 0$ [13, 4].

Hazan and Krauthgamer show that finding a near-optimal ε -equilibrium is as hard as the $(C \ln n)$ -Hidden Clique Problem, for some universal constant C . By near-optimal we mean having value close to maximum possible value obtained in an actual Nash equilibrium. Subsequently, Minder and Vilenchik [19] improved this hardness result to $C = 2 + \delta$ for arbitrarily small $\delta > 0$.¹ Here, we will rely on the assumed hardness of the $(C \ln n)$ -Hidden Clique Problem for some constant C , and will not concern ourselves with optimizing the value of C .

1.1 A sharp result for near-optimal approximate Nash

It is important to note that the problem considered in [16] is not equivalent to the basic ε -equilibrium problem (in what follows, we refer to the problem of finding an exact or ε -approximate equilibrium, with no optimization objective or other restrictions, as the basic Nash equilibrium problem). In light of the results of [16, 19] it is natural to ask to what extent the hardness for near-optimal approximate equilibrium gives an indication of hardness for basic approximate equilibrium. Indeed, [16], in their concluding remarks, ask whether their methods can be used to rule out a PTAS for basic Nash equilibrium. One of the messages of the current paper is that these two problems are quite different in terms of approximability and that one should not yet be overly pessimistic about the possibility for a PTAS for basic Nash equilibrium. Indeed, while there is a polynomial time algorithm to find a 0.3393-equilibrium, we show that finding a near-optimal $(1/2 - \eta)$ -equilibrium is hard.

Theorem 1.2 (Informal statement of [Theorem 3.7](#)). *For every constant $\eta > 0$, finding a near-optimal $(1/2 - \eta)$ -approximate equilibrium is at least as hard as the $(C \ln n)$ -Hidden Clique Problem.²*

In the above theorem and in what follows, C is a universal constant unless it is explicitly stated otherwise. As mentioned above, there is a simple polynomial time algorithm to find a $1/2$ -equilibrium, and we extend this algorithm to find a $1/2$ -equilibrium with value at least that of the best exact equilibrium:

Theorem 1.3. *There exists a polynomial time algorithm to find a $1/2$ -approximate equilibrium with value at least that of the optimal (maximum value) exact equilibrium.*

Thus, [Theorem 1.2](#) is tight and unlike basic Nash equilibrium, where stronger techniques yield approximations better than $1/2$, near-optimal Nash equilibrium does not admit “non-trivial” approximations in polynomial time (assuming the Hidden Clique problem is hard).

¹There is a small caveat: the reduction of [19] only certifies the presence of a hidden clique (i. e., distinguishes the graph from the random graph $G(n, 1/2)$ with high probability) but does not identify the vertices of the clique.

²Throughout, by “at least as hard,” we mean that there is a randomized polynomial time reduction from the first problem to the second.

1.2 The bigger picture: Hardness for NP-hard variants of Nash

Just like with basic ε -equilibrium, finding a near-optimal ε -equilibrium can be done in quasi-polynomial time using the algorithm of [18]. However, the exact version—finding a maximum-value Nash equilibrium—is known to be NP-hard [15] and so seemingly harder than its basic counterpart which is in PPAD [20]. In fact, maximum value Nash is one of several optimization variants of Nash equilibrium which are NP-complete. Other variants include: determining whether a bimatrix game has more than one Nash equilibrium [15], finding a Nash equilibrium with minimum support [15], and determining whether there exists an equilibrium with value at least $1 - 1/n$ or all equilibria have value at most ε/n (even for arbitrarily small vanishing $\varepsilon = \varepsilon(n) > 0$) [6]. We show that approximate-equilibrium variants of these problems are at least as hard as Hidden Clique.

For the problem of approximating the optimal value of a Nash equilibrium, we prove the following theorem.

Theorem 1.4 (Informal statement of [Theorem 3.9](#)). *For every constant $\eta > 0$, there exist $C = C(\eta) > 0$ and $\varepsilon = \varepsilon(\eta) > 0$ such that finding an ε -equilibrium with value at least η , even in a game having an equilibrium of value $\geq 1 - \eta$, is at least as hard as the $(C \ln n)$ -Hidden Clique Problem.*

For the case of determining whether a game has more than one equilibrium, note that by continuity considerations, every two-player game has an infinite number of ε -equilibria. An appropriate approximate analogue is to consider the problem of finding two ε -equilibria with (at least) a certain total variation distance between them. We show that this problem is also at least as hard as Hidden Clique.

Theorem 1.5 (Informal statement of [Theorem 4.1](#)). *For all sufficiently small constant $\varepsilon > 0$, determining whether a game has two different ε -approximate equilibria (say, having statistical distance at least 3ε) is at least as hard as the $(C \ln n)$ -Hidden Clique Problem.*

We then move to the problem of finding an equilibrium with small support. Recall that by [18], there exist ε -equilibria with support $O(\ln n / \varepsilon^2)$. It is also known that for any $\eta > 0$, in certain two-player games all $(1/2 - \eta)$ -equilibria must have support at least $\ln n / (1 + \ln(1/\eta))$ [11] (the threshold of $1/2$ is tight, since the simple $1/2$ -equilibrium of [10] has support 3). As an approximate-equilibrium variant of the minimum support equilibrium problem, we consider the problem of finding an ε -equilibrium with support at most some threshold t , and prove the following hardness result.

Theorem 1.6. *For every $\eta > 0$ there exists $C = C(\eta) > 0$ such that finding a $(1/2 - \eta)$ -equilibrium with support at most $(\ln n)/2$ is as hard as the $(C \ln n)$ -Hidden Clique Problem.*

This can be seen as a complexity-theoretic analogue of the lower bound of [11] mentioned above. Again, this contrasts with the situation for basic Nash equilibrium, which is guaranteed to exist, and admits stronger approximations (namely, the 0.3393-equilibrium found by the algorithm of [21]).

While these are all negative results, we again would like to stress that there is a positive message to this story: these problems are hard because they are approximate versions of NP-complete problems. In contrast, the basic equilibrium problem in its exact form is only PPAD-complete and thus seemingly easier. Therefore, these results should *not* be viewed as evidence against a PTAS for basic Nash equilibrium.

1.3 The complexity of approximate pure Bayes-Nash equilibria

Finally, we consider the problem of approximating pure *Bayes-Nash Equilibria* (BNE) in two-player games. *Bayesian games* model the situation where the players' knowledge of the world is incomplete. In a Bayesian game, both players may be in one of a number of different states, known as *types*, representing what each player knows about the state of the world, and the payoff of each player depends on the type of both players in addition to their strategies. The types are distributed according to some joint distribution and are not necessarily independent. A pure strategy for a Bayesian game assigns to each type a strategy that the player plays when she is in that type. In a pure BNE, conditioning on any given type for a given player, the player cannot gain by changing his strategy for that type. See [Section 6](#) for precise definitions.

Conitzer and Sandholm [7] have shown that determining whether a given two-player game has a pure BNE is NP-complete. We show that this holds also for approximate pure BNE.

Theorem 1.7. *Let $\varepsilon = 0.004$. Then given a Bayesian game that admits a pure BNE, it is NP-hard to find a pure ε -BNE for the game. Moreover, it is NP-hard to distinguish games that admit a pure BNE from games that do not admit a pure ε -BNE.*

This hardness result relies heavily on the joint distribution of the players' types being non-uniform (in fact, not even product distribution). We show that when the distribution over type pairs is uniform, there is in fact a quasi-polynomial time algorithm for ε -approximate pure BNE (when a pure BNE exists).

Theorem 1.8 (Informal statement of [Theorem 6.1](#)). *For every constant $\varepsilon > 0$ there is a quasipolynomial time algorithm to find a pure ε -BNE in two-player Bayesian games with uniformly distributed types and in which a pure BNE exists.*

We remark that this algorithm extends easily to arbitrary product distributions over types but in order to keep the notation simple we restrict our attention to the uniform case.

The quasi-polynomial running time of the algorithm is possibly optimal: it follows immediately from our hardness for small support equilibrium that this problem is also as hard as Hidden Clique.

Theorem 1.9. *There exists some constant $C > 0$ such that finding a $1/30$ -approximate pure BNE in two-player Bayesian games with uniformly distributed types and in which a pure BNE exists is as hard as finding a hidden clique of size $C \ln n$.*

1.4 Organization

In [Section 3](#) we prove the results relating to approximate equilibria with large value: [Theorem 1.2](#) ([Section 3.2](#)), [Theorem 1.3](#) ([Section 3.4](#)), and [Theorem 1.4](#) ([Section 3.3](#)).

In [Section 4](#) we prove [Theorem 1.5](#) by a black-box application of [Theorem 1.4](#). In [Section 5](#) we prove [Theorem 1.6](#) using similar techniques as for the hardness results of [Section 3](#). In [Section 6](#) we prove our results for Bayesian games, [Theorems 1.7, 1.8, and 1.9](#).

2 Preliminaries

A *bimatrix game* $\mathcal{G} = (M_{\text{row}}, M_{\text{col}})$ is a game defined by two $N \times N$ matrices, M_{row} and M_{col} , and two players: the *row player* and the *column player*. We assume throughout that the game is normalized so

that both matrices have values in the interval $[0, 1]$. The row and column players choose *strategies* x and y respectively, where x, y are nonnegative N -dimensional vectors satisfying $\sum_i x_i = \sum_j y_j = 1$. A *pure strategy* is one with support 1 (i. e., a vector with 1 in one entry and 0 in the rest). The row (resp. column) player's *payoff* is given by $x^\top M_{\text{row}y}$ (resp. $x^\top M_{\text{col}y}$).

A *Nash equilibrium* is a pair of strategies (x, y) such that neither player has any incentive to deviate to a different strategy, assuming the other player does not deviate. Formally, in an equilibrium, for all i, j we have $e_i^\top M_{\text{row}y} \leq x^\top M_{\text{row}y}$ and $x^\top M_{\text{col}e_j} \leq x^\top M_{\text{col}y}$. An ε -approximate Nash equilibrium, or ε -*equilibrium* for short, is a pair of strategies x, y where each player has incentive at most ε to deviate. That is, for all i, j ,

$$e_i^\top M_{\text{row}y} \leq x^\top M_{\text{row}y} + \varepsilon \quad \text{and} \quad x^\top M_{\text{col}e_j} \leq x^\top M_{\text{col}y} + \varepsilon.$$

The *value* of a pair of strategies, denoted $v_g(x, y)$, is the average payoff of the two players, i. e.,

$$v_g(x, y) = \frac{1}{2}(x^\top M_{\text{row}y} + x^\top M_{\text{col}y}) = \sum_{i,j} x_i y_j \frac{M_{\text{row}}(i, j) + M_{\text{col}}(i, j)}{2}.$$

For a vector $x \in \mathbb{R}^n$ and $S \subseteq [n]$, we write x_S for the projection of x to the coordinates S . We write $\|x\| = \sum_{i=1}^n |x_i|$ for the ℓ_1 norm of x . Thus, for a strategy (in other words, a probability distribution) $x \in [0, 1]^n$ we write $\|x_S\|$ for the probability that the player plays an element of S .

Further, for a set $S \subseteq [n]$ of strategies, we use $v_{g|S}(x, y)$ to denote the value of (x, y) *conditioned* on both players playing in S . Formally,

$$v_{g|S}(x, y) = \mathbb{E}_{i \sim x, j \sim y} \left[\frac{M_{\text{row}}(i, j) + M_{\text{col}}(i, j)}{2} \mid i, j \in S \right] = \frac{v_g(x_S, y_S)}{\|x_S\| \cdot \|y_S\|}.$$

(If $\|x_S\| = 0$ or $\|y_S\| = 0$, $v_{g|S}(x, y)$ is undefined.)

Given an undirected graph $G = (V, E)$, and (not necessarily disjoint) vertex sets $S_1, S_2 \subseteq V$, we will denote by $E(S_1, S_2)$ the set of ordered pairs $\{(i, j) \in S_1 \times S_2 \mid \{i, j\} \in E \text{ or } i = j\}$. We will refer to $d(S_1, S_2) = |E(S_1, S_2)| / (|S_1| |S_2|)$ as the *density* of the pair (S_1, S_2) . Note that, letting A be the adjacency matrix of G (with 1's on the diagonal) and $\bar{x}, \bar{y} \in [0, 1]^n$ be the uniform probability distributions over S_1 and S_2 (respectively), we have $d(S_1, S_2) = \bar{x}^\top A \bar{y}$.

We will write $\text{HC}(n, k)$ to denote the distribution in the input to the k -Hidden Clique problem. That is, we will write $G \sim \text{HC}(n, k)$ to mean that G is sampled from $G(n, 1/2)$ with a hidden clique of size k .

Throughout, w. h. p. stands for “with high probability,” meaning $1 - o(1)$. Finally, we shall make repeated use of the standard Chernoff-Hoeffding bound [17].

Lemma 2.1 (Chernoff-Hoeffding bound). *Let X_1, \dots, X_m be independent random variables in $[0, 1]$, and let $\mu = (1/m) \mathbb{E}[\sum X_i]$. Then, for all $\varepsilon \geq 0$,*

$$\Pr \left[\frac{1}{m} \sum X_i \geq \mu + \varepsilon \right] \leq e^{-2\varepsilon^2 m}.$$

3 Approximate equilibria with near-optimal value

3.1 The reduction

In this section we describe the general reduction that we use to prove Theorems 1.2 and 1.4 and describe its properties. This reduction also forms the basis for the reductions we use to prove Theorems 1.5, 1.6 and 1.9. It is based on the reduction of [16].

As in [16] our soundness analysis proceeds by using the approximate equilibrium to find a dense bipartite subgraph of G . The following lemma shows that this is sufficient to recover the hidden clique.

Lemma 3.1 ([16, Lemma 5.3]). *There exist universal constants c_1 and c_2 such that the following holds. Let $G \sim \text{HC}(n, C \ln n)$ for some $C \geq c_1$. Then, there is a polynomial time algorithm which, given a pair of vertex sets $S_1, S_2 \subseteq [n]$ of size $c_2 \ln n$ and density $d(S_1, S_2) \geq 5/9$, reconstructs the hidden clique (with high probability over G).*

The lemma is slightly different from Lemma 5.3 of [16]: there we start with a bipartite subgraph of density $3/5$ instead of $5/9$ but this minor difference only changes the value of the constant c_2 —the lemma holds for any constant density strictly larger than $1/2$.

Let us now describe the reduction. It is controlled by three parameters $\alpha, \beta, \gamma \in (0, 1)$. Setting these parameters appropriately gives the various hardness results.

Reduction 3.2. Let $G = (V, E)$ be an n -vertex graph and A its adjacency matrix (with 1's on the diagonal). Then, for parameters $\alpha, \beta, \gamma \in (0, 1)$, we define a (random) bimatrix game $\mathcal{G} := \mathcal{G}(G, \alpha, \beta, \gamma)$ as follows.

Let $N = n^c$ where $c = (c_2 + 1) \ln(1/\beta)$ for the universal constant c_2 of Lemma 3.1. Pick a random $N \times n$ matrix B whose entries are i. i. d. $\{0, 1\}$ variables with expectation β . Then $\mathcal{G} = (M_{\text{row}}, M_{\text{col}})$, where the payoff matrices are:

$$M_{\text{row}} = \begin{pmatrix} \alpha A & 0 \\ B & \gamma J \end{pmatrix} \quad M_{\text{col}} = \begin{pmatrix} \alpha A & B^\top \\ 0 & \gamma J \end{pmatrix}, \quad (3.1)$$

and J is the all-ones $N \times N$ matrix.

We conclude this section with an additional lemma which shows how to obtain a dense bipartite subgraph given an approximate equilibrium of \mathcal{G} with certain properties. This lemma (and its proof) is analogous to [16, Lemma 5.2], but as we need it in larger generality we also give the proof.

Lemma 3.3. *Let \mathcal{G} be as in Reduction 3.2. Fix any $s \in [0, 1]$, $t \in [0, 1]$ and $\varepsilon \in [0, 1]$ such that $1 - 2t - 3\sqrt{s}/2 \geq \alpha + \varepsilon$, and let (x, y) be an ε -approximate equilibrium of \mathcal{G} with the following two properties:*

- Both $\|x_{[n]}\| \geq 1 - t$ and $\|y_{[n]}\| \geq 1 - t$.
- The conditional value $v_{\mathcal{G}|[n]}(x, y) \geq (1 - s)\alpha$.

Then, given (x, y) as above, we can find, in polynomial time, vertex sets $S_1, S_2 \subseteq [n]$ each of size $c_2 \ln n$ such that their density is $d(S_1, S_2) \geq 5/9$.

In the proof of [Lemma 3.3](#) we shall use the following simple claim, which is analogous to [[16](#), Claim 5.3]:

Claim 3.4. *The following holds w. h. p. over \mathcal{G} : For any set $S \subset [n]$ of size $|S| < c_2 \ln n$, in any ε -equilibrium (x, y) of the above game satisfying $\|x_{[n]}\| \geq 1 - t$ and $\|y_{[n]}\| \geq 1 - t$, the probability mass the column (or row) player places on such a set S is less than $\alpha + \varepsilon + t$.*

Proof. Fix one such set S . Then with high probability, there is a row i^* of B such that $B_{i^*,j} = 1$ for all $j \in S$. Indeed, the probability that such a row does not exist is exactly

$$\left(1 - \beta^{-|S|}\right)^N < \exp(-n^{-c_2 \ln 1/\beta} N) = \exp(-n^{\ln 1/\beta}). \quad (3.2)$$

In other words, if the column player places a probability mass p on S , the row player can achieve payoff p by deviating to i^* , whereas the payoff of the row player in (x, y) is at most $\alpha(1 - t) + t \leq \alpha + t$. Thus if (x, y) is an ε -equilibrium, p is at most $\alpha + t + \varepsilon$. Finally, to show that this holds for all such sets S , it is enough to take a union bound in equation (3.2), so that (for sufficiently large n) the probability that the claim does not hold for at least one set S is at most

$$n^{c_2 \ln n} \left(1 - \beta^{-|S|}\right)^N < \exp(c_2 (\ln n)^2 - n^{\ln 1/\beta}) < \exp(-n^{\ln 1/\beta} / 2).$$

□

Proof of Lemma 3.3. Let $\tilde{x}, \tilde{y} \in [0, 1]^n$ be the strategies x and y conditioned on playing in $[n]$. That is, $\tilde{x} = x_{[n]} / \|x_{[n]}\|$, and $\tilde{y} = y_{[n]} / \|y_{[n]}\|$. The second given property of the strategy pair (x, y) can be rephrased as $\tilde{x}^\top A \tilde{y} \geq 1 - s$.

Let

$$S'_1 = \left\{ i \in [n] : e_i^\top A \tilde{y} \geq 1 - 2\sqrt{s}/3 \right\}.$$

To obtain a lower bound on the cardinality $|S'_1|$, we shall lower bound $\|x_{S'_1}\|$ by $\varepsilon + \alpha$ and then appeal to the claim. By the first given property of (x, y) , we have

$$\|x_{S'_1}\| \geq 1 - t - \|x_{[n] \setminus S'_1}\| \geq 1 - t - \|\tilde{x}_{[n] \setminus S'_1}\|.$$

We can bound the last term, $\|\tilde{x}_{[n] \setminus S'_1}\|$, from above using Markov's inequality, viz.,

$$\|\tilde{x}_{[n] \setminus S'_1}\| = \Pr_{i \sim \tilde{x}} \left[1 - e_i^\top A \tilde{y} > 2\sqrt{s}/3 \right] \leq \frac{s}{2\sqrt{s}/3} = 3\sqrt{s}/2.$$

Thus we have $\|x_{S'_1}\| \geq 1 - t - 3\sqrt{s}/2 \geq \varepsilon + \alpha + t$ and so by [Claim 3.4](#), $|S'_1| \geq c_2 \ln n$. Truncate S'_1 by taking any arbitrary subset $S_1 \subseteq S'_1$ of cardinality $|S_1| = c_2 \ln n$.

Now let $\bar{x} \in [0, 1]^n$ be the uniform distribution over S_1 . Note that $\bar{x}^\top A \tilde{y} \geq 1 - 2\sqrt{s}/3$. Let

$$S'_2 = \left\{ j \in [n] : \bar{x}^\top A e_j \geq 5/9 \right\}.$$

The argument to lower bound $|S'_2|$ is similar to the argument for $|S'_1|$. We get

$$\|\tilde{y}_{[n] \setminus S'_2}\| = \Pr_{j \sim \tilde{y}} \left[1 - \bar{x}^\top A e_j > 4/9 \right] \leq \frac{2\sqrt{s}/3}{4/9} = 3\sqrt{s}/2$$

and therefore

$$\|y_{S_2'}\| \geq 1 - t - 3\sqrt{s}/2 \geq \varepsilon + \alpha + t,$$

as desired. Again, we can truncate S_2' by taking a subset $S_2 \subseteq S_2'$ of cardinality $|S_2| = c_2 \ln n$. By construction $d(S_1, S_2) \geq 5/9$, and we are done. \square

3.2 Hardness for ε close to $1/2$

To obtain [Theorem 1.2](#) the main requirement is to set $\alpha = 1/2 + O(\eta)$. The values of β and γ are essentially irrelevant in this case—the only thing needed is that β is bounded away from both 0 and α and that $\gamma \leq 1/2$.

Lemma 3.5. *For any $t \leq 1/2$, let $\alpha = 1/2 + t$, $\gamma \leq 1/2$ and \mathcal{G} be the game of [Reduction 3.2](#). Then for any pair of strategies (x, y) with value at least $v_{\mathcal{G}}(x, y) \geq \alpha - t^2$ it holds that $\|x_{[n]}\|$ and $\|y_{[n]}\|$ are both at least $1 - t$.*

Proof. Let $p = \|x_{[n]}\|$ and $q = \|y_{[n]}\|$. As the value of any outcome outside the αA block is at most $1/2$, we have that the value of (x, y) is at most

$$pq\alpha + (1 - pq)\frac{1}{2} = tpq + \frac{1}{2},$$

so that if the value is at least $\alpha - t^2$ we have

$$tpq + \frac{1}{2} \geq \alpha - t^2 \quad \text{and} \quad pq \geq \frac{\alpha - t^2 - 1/2}{t} = 1 - t.$$

Since $p, q \in [0, 1]$, it follows that they are both at least $1 - t$. \square

Observation 3.6. Let (x, y) be any pair of strategies with value $v_{\mathcal{G}}(x, y) \geq 1/2$ and $\|x_{[n]}\| > 0$, $\|y_{[n]}\| > 0$. Then $v_{\mathcal{G}_{[n]}}(x, y) \geq v_{\mathcal{G}}(x, y)$, provided that $\gamma \leq 1/2$.

Plugging this into [Lemma 3.3](#), we can now easily complete the proof of hardness for ε close to $1/2$.

Theorem 3.7 (Detailed Statement of [Theorem 1.2](#)). *For some universal constant C and every $\eta > 0$ there exist $\delta = \Omega(\eta^2)$, $\alpha \geq 1/2$ such that the following holds. Given a graph $G = (V, E)$ we can in randomized polynomial time construct a bimatrix game \mathcal{G} with maximum average payoff α (over all strategy pairs) such that, if $G \sim \text{HC}(n, C \ln n)$, the following holds (w. h. p. over G and \mathcal{G}):*

Completeness \mathcal{G} has a Nash equilibrium (x, y) with value α .

Soundness Given any $(\frac{1}{2} - \eta)$ -equilibrium with value $\geq \alpha - \delta$, we can recover the hidden clique in polynomial time.

Proof. Given a graph $G = (V, E)$, we apply [Reduction 3.2](#) with parameters as follows. For some $t > 0$ to be determined momentarily, let $\alpha = 1/2 + t$, $\beta = \gamma = 1/3$, $\delta = t^2$.

For the completeness, we shall show that having both players play uniformly over the hidden clique is an equilibrium. For this to hold, we have to make sure that there is no row in B with average value at

least α in the positions corresponding to the clique. By [Lemma 2.1](#) and a union bound over all rows of B we can bound the probability of this happening by

$$Ne^{-2(\alpha-\beta)^2C\ln n} \leq N \cdot n^{-C/18}.$$

If C is a sufficiently large universal constant (e. g., $C = 18 \cdot (c_2 + 1) \ln 1/\beta + 1$) this probability is $o(1)$ and the completeness property follows.

For the soundness, consider any $(1/2 - \eta)$ -approximate equilibrium (x, y) with value at least $v_{\mathcal{G}}(x, y) \geq \alpha - \delta = \alpha - t^2$. By [Lemma 3.5](#) both $\|x_{[n]}\|$ and $\|y_{[n]}\|$ are at least $1 - t$. Furthermore, by [Observation 3.6](#) we have $v_{\mathcal{G}_{[n]}}(x, y) \geq \alpha - \delta \geq \alpha(1 - 2t^2)$.

Thus if $1 - 2t - 3t/\sqrt{2} \geq \alpha + \frac{1}{2} - \eta$ we can apply [Lemma 3.3](#) and extract a dense bipartite subgraph of G which can be plugged in to [Lemma 3.1](#) to obtain the hidden clique. Setting $t = \eta/6$, the desired inequality holds and we get the result. \square

3.3 Hardness of distinguishing between small and large value

For [Theorem 1.4](#) the choices of all three parameters α, β, γ of [Reduction 3.2](#) are important. We are going to set γ close to 0, and $\alpha > \beta$ both close to 1.

On a high level, the proof has the same structure as that of [Theorem 1.2](#). However, in the current setting [Lemma 3.5](#) and [Observation 3.6](#) do not apply. To arrive at similar conclusions we use a different argument, exploiting the fact that (x, y) is an ε -equilibrium. Essentially, the argument is as follows: the off-diagonal blocks (B and B^\top) are not stable, since there is too much incentive for at least one player to deviate. Therefore, most of the probability mass in an equilibrium is concentrated either in the αA block, or in the γJ block. However, in the γJ block, the value is too small. So, if the equilibrium has even slightly larger value, its mass must be concentrated in the αA block. There it has to actually have very large value, since otherwise, there is incentive for both players to deviate to B and B^\top to get reward β . The rest of the proof follows as before.

Formally, we have the following lemma, showing that (under certain conditions) any ε -equilibrium with non-negligible value must satisfy the conditions of [Lemma 3.3](#).

Lemma 3.8. *Fix a parameter $\varepsilon \in (0, 1)$, let $\alpha - \beta \leq \varepsilon$, and $\gamma = 4\sqrt{\varepsilon}$ and consider the game \mathcal{G} as in [Reduction 3.2](#).*

Then, w. h. p. over \mathcal{G} , any ε -equilibrium (x, y) with value more than $5\sqrt{\varepsilon}$ satisfies:

- *Both $\|x_{[n]}\|$ and $\|y_{[n]}\|$ are at least $1 - \sqrt{\varepsilon}$.*
- *$v_{\mathcal{G}_{[n]}}(x, y) \geq \alpha - 3\varepsilon$.*

Proof. Consider any ε -equilibrium (x, y) with value more than $5\sqrt{\varepsilon}$. Note that this trivially implies that $\varepsilon \leq 1/25$.

Let $p = \|x_{[n]}\|$ and $q = \|y_{[n]}\|$ be the probability mass that the row (resp. column) player assigns to the first n strategies in this equilibrium. We begin with the first item, i. e., the lower bound on p and q .

Consider the row player's incentive to deviate by shifting the probability mass in the first n rows to the uniform distribution over the remaining rows. When the column player is playing outside the first n columns, this deviation changes the row player's payoff from 0 to γ , and when the column player is

playing in one of the first n columns the row player's payoff decreases by at most $\alpha - \beta + o(1)$ (since, by the Chernoff bound, the number of 1s in the i th column of B is $N(\beta - o(1))$ for all i w. h. p.). Let us ignore this $o(1)$ change since it is immaterial to the rest of the proof. Since this is an ε -equilibrium, we have

$$p(1-q)\gamma - pq(\alpha - \beta) \leq \varepsilon, \quad (3.3)$$

and so

$$p(1-q) \leq \frac{1}{\gamma}(\varepsilon + pq(\alpha - \beta)) \leq \frac{1}{\gamma}(\varepsilon + \alpha - \beta) \leq \sqrt{\varepsilon}/2.$$

Considering also the symmetric argument for the column player, this gives

$$\max\{p(1-p), q(1-q)\} \leq \max\{p(1-q), q(1-p)\} \leq \sqrt{\varepsilon}/2, \quad (3.4)$$

where the first inequality holds trivially for all $p, q \in [0, 1]$. The inequality $p(1-p) \leq \sqrt{\varepsilon}/2$ together with the constraint $\varepsilon \leq 1/25$ implies that either $p \leq \sqrt{\varepsilon}$ or $p \geq 1 - \sqrt{\varepsilon}$. The inequality for q gives an analogous bound. Moreover, it cannot be the case that $p < 1/2 < q$ or $q < 1/2 < p$, since then we would have $\max\{p(1-q), q(1-p)\} \geq 1/4$, contradicting (3.4). Thus, it follows that either $p, q \leq \sqrt{\varepsilon}$ or $p, q \geq 1 - \sqrt{\varepsilon}$.

Let us now exclude the first option. Suppose for contradiction that $p, q \leq \sqrt{\varepsilon}$. Then the value of the equilibrium is at most

$$pq\alpha + (p(1-q) + q(1-p)) \cdot \frac{1}{2} + (1-p)(1-q)\gamma \leq \varepsilon\alpha + \sqrt{\varepsilon}/2 + \gamma \leq \varepsilon\alpha + 4.5\sqrt{\varepsilon} < 5\sqrt{\varepsilon},$$

contradicting the assumption that (x, y) has value more than $5\sqrt{\varepsilon}$. Hence both p and q are at least $1 - \sqrt{\varepsilon}$.

It remains to give a lower bound on the conditional value $w := v_{\mathcal{G}||[n]}(x, y)$ obtained when playing inside αA . As above, consider the incentive for the row player to deviate to the uniform distribution over the rows outside $[n]$. To be more precise, we can use w instead of α in the bound (3.3). Solving for w gives

$$w \geq \beta + \frac{p(1-q)\gamma - \varepsilon}{pq} \geq \beta - \frac{\varepsilon}{pq} \geq \beta - 2\varepsilon \geq \alpha - 3\varepsilon,$$

where we used that

$$\frac{1}{pq} \leq \frac{1}{(1 - \sqrt{\varepsilon})^2} \leq \frac{25}{16} < 2$$

for $\varepsilon \leq 1/25$. □

Equipped with [Lemma 3.8](#), it is easy to finish the proof of [Theorem 1.4](#).

Theorem 3.9 (Detailed statement of [Theorem 1.4](#)). *For every constant $\eta > 0$ there exist $\varepsilon = \Omega(\eta^2)$ and $C = O(1/\eta^3)$ such that the following holds. Given a graph G , we can in randomized polynomial time construct a bimatrix game \mathcal{G} such that, if $G \sim \text{HC}(n, C \ln n)$, the following holds (w. h. p. over G and \mathcal{G}):*

Completeness *There is a Nash equilibrium (x, y) with both players earning payoff $1 - \eta$.*

Soundness *Given any ε -equilibrium with value $\geq \eta$, we can recover the hidden clique in polynomial time.*

Proof. Given a graph $G = (V, E)$, we apply [Reduction 3.2](#) with parameters as follows.

Let $\varepsilon = (\eta/6)^2$, $\alpha = 1 - \eta = 1 - 6\sqrt{\varepsilon}$, $\beta = \alpha - \varepsilon$, and $\gamma = 4\sqrt{\varepsilon}$. Assume without loss of generality that η is small enough so that $\alpha > 3/4$.

For the completeness, we proceed as in the proof of [Theorem 3.7](#). We can upper bound the probability that the uniform distribution over the hidden clique is not an equilibrium by

$$Ne^{-2(\alpha-\beta)^2C\ln n} \leq Nn^{-\Omega(C\varepsilon^2)}.$$

We have $N = n^{O(\ln 1/\beta)} = n^{O(\eta)} = n^{O(\sqrt{\varepsilon})}$. Letting C be a sufficiently large multiple of $1/\varepsilon^{1.5}$ the completeness property follows.

For the soundness analysis, take any ε -approximate equilibrium (x, y) with value at least $5\sqrt{\varepsilon}$. By [Lemma 3.8](#), w. h. p. (x, y) satisfies the conditions of [Lemma 3.3](#) with $t = \sqrt{\varepsilon}$ and $s = 3\varepsilon/\alpha \leq 4\varepsilon$. The only remaining thing to check is the condition $\alpha + \varepsilon \leq 1 - 2t - 3\sqrt{s}/2$, which is easily verified ($1 - 5\sqrt{\varepsilon}$ is an upper bound for the LHS and a lower bound for the RHS). Thus we can apply [Lemma 3.3](#) to obtain a dense subgraph of G which by [Lemma 3.1](#) allows us to recover the planted clique. \square

3.4 An algorithm for near-optimal (1/2)-approximate equilibria

In this section we prove [Theorem 1.3](#) by describing a simple algorithm to find a 1/2-approximate Nash equilibrium with at least as large a value as the best exact Nash equilibrium. This shows that the bound of $1/2 - \eta$ in [Theorem 1.2](#) is tight.

For basic 1/2-approximate equilibria (irrespective of the value), the following simple algorithm was suggested by Daskalakis, Mehta and Papadimitiou [10]. Start by choosing an arbitrary pure strategy e_i for the row player, let e_j be the column player’s best response to e_i , and let e_k be the row player’s best response to e_j . Then the following is a 1/2-equilibrium: let the column player play e_j , and let the row player play e_i with probability 1/2 and e_k with probability 1/2 (neither player can gain more than 1/2 by deviating, since each player is playing a best response strategy with probability 1/2). Thus, every bimatrix game has a 1/2-approximate equilibrium in which one of the players plays a pure strategy. We show that this is also the case for optimal value 1/2-equilibria (the difference being that in this case, the other player may play a mixed strategy with arbitrarily large support—not just two strategies).

Lemma 3.10. *For every bimatrix game which has a Nash equilibrium of value v , there exists a 1/2-approximate equilibrium with value at least v in which one of the players plays a pure strategy.*

Proof. Let M_{row} and M_{col} be the payoff matrices for the row and column players, respectively. Let (x^*, y^*) be a Nash equilibrium in this game of value $v := v_G(x^*, y^*)$, and let v_r and v_c be the payoff in this equilibrium to the row and column players, respectively (hence $v = (1/2)(v_r + v_c)$). Without loss of generality, assume $v_r \geq v_c$ (otherwise, a symmetric argument applies).

Let the pure strategy e_j be some strategy in the support of y^* for which the row player’s payoff $x^{*\top} M_{\text{row}} e_j$ is at least v_r (when the row player is playing x^* and the column player e_j). Such a strategy exists since $v_r = \mathbb{E}_{j \sim y^*} [x^{*\top} M_{\text{row}} e_j]$ is the expected payoff for the row player when the column player plays according to y^* . Furthermore, any such e_j is a best response to x^* since (x^*, y^*) is an equilibrium.

Clearly, if the pair of strategies (x^*, e_j) is a 1/2-equilibrium we are done, since both the row and column player are getting at least the same payoff as for the pair (x^*, y^*) .

If (x^*, e_j) is *not* a $1/2$ -equilibrium, this must be because the row player has incentive $\geq 1/2$ to deviate (as the column player is by definition playing a best response). Note that this implies that $v_c \leq v_r \leq 1/2$ since the row player's incentive to deviate can never be more than $1 - v_r$. Let e_k be some best response for the row player, and consider the pair of strategies $((1/2)(x^* + e_k), e_j)$. As above, this is a $1/2$ -equilibrium, since both players are playing a best response with probability $1/2$. Furthermore, since the row player gets payoff $\geq v_r$ with probability $1/2$ (when playing according to x^*) and payoff $\geq v_r + 1/2$ with probability $1/2$ (when playing according to e_k), the payoff for the row player is at least $v_r/2 + (v_r + 1/2)/2 = v_r + 1/4 \geq v_r + v_c/2$. The payoff for the column player is at least $v_c/2$ (since with probability $1/2$ the row player plays according to x^* in which case the column player gets v_c). Thus the value is at least $(1/2)(v_r + v_c/2 + v_c/2) = v$, and we are done. \square

Now our algorithm follows immediately.

Theorem 1.3 restated. There exists a polynomial time algorithm to find a $1/2$ -approximate equilibrium with value at least that of the optimal (maximum value) exact equilibrium.

Proof. Let M_{row} and M_{col} be the payoff matrices for the row and column matrices, respectively. By [Lemma 3.10](#) there exists such an equilibrium in which one player plays a pure strategy. Suppose this is the column player (otherwise, a symmetric algorithm applies). Try all possible pure strategies e_j for the column player. For each such strategy, solve the following linear program (if it is feasible):

$$\begin{array}{ll}
 \text{maximize} & \frac{1}{2}x^\top (M_c + M_{\text{row}})e_j \\
 \text{subject to} & x^\top M_{\text{col}}e_{j'} \leq x^\top M_{\text{col}}e_j + \frac{1}{2} & \forall j' \\
 & (e_i)^\top M_{\text{row}}e_j \leq x^\top M_{\text{row}}e_j + \frac{1}{2} & \forall i \\
 & x_i \geq 0 & \forall i \\
 & \sum_i x_i = 1
 \end{array}$$

For at least one strategy e_j , this LP is feasible and computes a $1/2$ -equilibrium with at least the desired value. \square

4 Finding a second equilibrium

In the following theorem, d_{TV} refers to the total variation distance between two vectors, i. e., $d_{\text{TV}}(x, y) = (1/2) \sum_i |x_i - y_i|$.

Theorem 4.1 (Detailed Statement of [Theorem 1.5](#)). *There is a $C > 0$ such that for all sufficiently small $\varepsilon > 0$ there is some $\varepsilon' = \varepsilon + O(\varepsilon^2)$ for which the following holds. Given a graph G we can in randomized polynomial time construct a bimatrix game \mathcal{G}' which admits a pure Nash equilibrium (e_i, e_j) such that, if $G \sim \text{HC}(n, C \ln n)$, the following holds (w. h. p. over G and \mathcal{G}'):*

Completeness *There is a Nash equilibrium (x, y) such that $d_{\text{TV}}(e_i, x) = d_{\text{TV}}(e_j, y) = 1$.*

Soundness *Given any ε -approximate equilibrium (x, y) of \mathcal{G}' with $d_{\text{TV}}(e_i, x) \geq \varepsilon'$ or $d_{\text{TV}}(e_j, y) \geq \varepsilon'$, we can recover the hidden clique in polynomial time.*

Remark 4.2. Note that the bound $\varepsilon + O(\varepsilon^2)$ on the statistical distance is almost tight: given any true equilibrium (x, y) there are ε -approximate equilibria (x, y') and (x', y) with $d_{\text{TV}}(x, x') \geq \varepsilon(1 - 1/n)$ and $d_{\text{TV}}(y, y') \geq \varepsilon(1 - 1/n)$ (simply take $x' = \varepsilon e_{i'} + (1 - \varepsilon)x$ and $y' = \varepsilon e_{j'} + (1 - \varepsilon)y$ for i' and j' which minimize $x_{i'}, y_{j'}$, respectively).

Proof. Given G , first construct the game \mathcal{G} of [Theorem 3.9](#) with parameter $\eta = 1/3$, so that given a (4ε) -approximate equilibrium of \mathcal{G} with value at least $1/3$ we can reconstruct the hidden clique (in fact, the reduction of [\[16\]](#) would also suffice, but we use this instantiation of [Theorem 3.9](#) for completeness). Let $(M_{\text{row}}, M_{\text{col}})$ be the payoff matrices of \mathcal{G} and let N' denote their dimension. Consider the new $(N' + 1)$ -strategy game $\mathcal{G}' = (M'_{\text{row}}, M'_{\text{col}})$, with the following payoff matrices.

$$M'_{\text{row}} = \begin{pmatrix} M_{\text{row}} & 0 \\ L & 1 \end{pmatrix} \quad M'_{\text{col}} = \begin{pmatrix} M_{\text{col}} & L^\top \\ 0 & 1 \end{pmatrix}, \quad (4.1)$$

where L is the $1 \times N'$ vector with each coordinate equal to λ . We set $\lambda = 1 - \eta = 2/3$. The game \mathcal{G}' is a bimatrix game of size $(N' + 1) \times (N' + 1)$.

Clearly, $(e_{N'+1}, e_{N'+1})$ is a pure Nash equilibrium of \mathcal{G}' . Note that for any mixed strategy x , we can write $d_{\text{TV}}(e_{N'+1}, x) = \|x_{[N']}\| = 1 - x_{N'+1}$. Similarly, we write $\|y_{[N']}\|$ or $1 - y_{N'+1}$ for $d_{\text{TV}}(e_{N'+1}, y)$.

The completeness case follows immediately since in that case \mathcal{G} has a Nash equilibrium with both players earning payoff $1 - \eta$. As $\lambda \leq 1 - \eta$ this is an equilibrium in \mathcal{G}' as well and since it does not use the $(N' + 1)$ 'st strategy we obtain the completeness property.

For the soundness, consider any ε -approximate equilibrium (x, y) of \mathcal{G}' , let $p = \|x_{[N']}\|$, $q = \|y_{[N']}\|$ be the probability that the row (resp. column) player plays in the original game \mathcal{G} , and assume without loss of generality that $x_{N'+1} < 1 - \varepsilon'$ (i. e., $p > \varepsilon'$), for ε' to be defined below (the case where $y_{N'+1} < 1 - \varepsilon'$ is similar). We need to show that (x, y) can be used to recover the planted clique. We first show that the assumption $p > \varepsilon'$ in fact implies that both p and q must be very large.

Claim 4.3. *Suppose $p > \varepsilon'$ for some $\varepsilon' = \varepsilon + O(\varepsilon^2)$. Then,*

$$p \geq \frac{1 - \varepsilon'}{1 + \eta} \quad \text{and} \quad q \geq \frac{1 - \varepsilon'}{1 + \eta}.$$

Proof. We argue similarly to the first part of the proof of [Lemma 3.8](#). Bounding the incentive to deviate when the row player shifts all mass in x to row $N' + 1$, we have

$$\begin{aligned} \varepsilon &\geq pq(\lambda - 1) + p(1 - q)(1 - 0) && \text{(row player's incentive to deviate)} \\ &= p(1 - (1 + \eta)q) && \text{(by choice of } \lambda), \end{aligned}$$

and similarly with the roles of p and q reversed. We now claim that

$$\max\{p(1 - (1 + \eta)p), q(1 - (1 + \eta)q)\} \leq \varepsilon. \quad (4.2)$$

If $p, q > 1/(1 + \eta)$, this is trivially true. Otherwise, it follows from the previous inequality, noting that in this case $\max\{p(1 - (1 + \eta)p), q(1 - (1 + \eta)q)\} \leq \max\{p(1 - (1 + \eta)q), q(1 - (1 + \eta)p)\}$.

Inequality [\(4.2\)](#) implies that either $p \leq \varepsilon'$ or $(1 + \eta)p \geq 1 - \varepsilon'$ for some $\varepsilon' = \varepsilon + O(\varepsilon^2)$ (for sufficiently small ε), and similarly for q . Note that the first possibility for p is ruled out by our assumption, therefore

$p \geq (1 - \varepsilon')/(1 + \eta)$. Moreover, we must have $q \geq (1 - \varepsilon')/(1 + \eta)$, since if $p \geq (1 - \varepsilon')/(1 + \eta)$ and $q \leq \varepsilon'$, then (x, y) would clearly not be an ε -equilibrium (the row player's incentive to deviate would be nearly $1/(1 + \eta)$, or more). \square

Let $(\tilde{x}, \tilde{y}) = (x_{[N']}/p, y_{[N']}/q)$ denote the strategies conditioned on playing on the first N' strategies. We need to show that (\tilde{x}, \tilde{y}) is a (4ε) -approximate equilibrium with value $> \eta$, which allows us to reconstruct the hidden clique by [Theorem 3.9](#).

Claim 4.4. (\tilde{x}, \tilde{y}) is an $\varepsilon/(pq)$ -approximate equilibrium for the original game \mathcal{G} .

Proof. To see this, suppose that one of the players, say the row player, gains $\tilde{\varepsilon}$ in \mathcal{G} by deviating to some strategy $e_{i'}$. Consider “lifting” this to a new strategy x' for \mathcal{G}' (i. e., in the strategy x' the row player plays $e_{N'+1}$ with probability $1 - p$, and $e_{i'}$ with probability p). The change in payoff the row player obtains in \mathcal{G}' by switching from x to x' can be written as

$$p(q\tilde{\varepsilon} + (1 - q) \cdot 0), \quad (4.3)$$

where the $q\tilde{\varepsilon}$ term is what the row player gains from when the column player plays on the first N' strategies, and the other term is 0 since, when the column player plays on $N' + 1$ the row player gets the same payoff (0) on all the first N' strategies. As (x, y) is an ε -approximate equilibrium in \mathcal{G}' , the expression in (4.3) must be bounded by ε and hence $\tilde{\varepsilon} \leq \varepsilon/(pq)$, and the claim follows. \square

Now, since by [Claim 4.3](#) we have $p, q \geq (1 - \varepsilon')/(1 + \eta) > 1/2$, [Claim 4.4](#) implies that (\tilde{x}, \tilde{y}) is a (4ε) -approximate equilibrium in \mathcal{G} . Thus, in order to apply [Theorem 3.9](#), it remains to lower bound the value of (\tilde{x}, \tilde{y}) .

Let r be the value of the row player in (\tilde{x}, \tilde{y}) . Note that if in (x, y) , the row player defects to row $N' + 1$, then the incentive to deviate is $\lambda - r = 1 - \eta - r$ when the column player plays in $[N']$ (and even greater when the column player plays in column $N' + 1$). Thus, since (x, y) is an ε -equilibrium we have

$$(1 - \eta - r)p \leq \varepsilon, \quad \text{or} \quad r \geq 1 - \eta - \varepsilon/p \geq 1 - \eta - \frac{\varepsilon(1 + \eta)}{1 - \varepsilon'}$$

(by [Claim 4.3](#)). Arguing similarly for the column player, we get that the value of (\tilde{x}, \tilde{y}) is at least

$$\frac{2}{3} - \frac{\varepsilon(1 + \eta)}{1 - \varepsilon'} > \frac{1}{3} = \eta$$

(for sufficiently small ε). This concludes the proof of [Theorem 4.1](#). \square

5 Small support equilibria

In this section, we prove a hardness result for finding an ε -approximate Nash equilibrium with small (logarithmic) support when one exists, even for ε close to $1/2$. Note that an ε -approximate Nash equilibrium for two-player n' -strategy games with support at most $O(\ln n'/\varepsilon)$ is guaranteed to exist by the algorithm of Lipton et al. [18]. Here we consider approximate equilibria with smaller (but still

logarithmic) support. Also, note that this is tight, since for $\varepsilon = 1/2$, we have the simple algorithm of [10], which gives a $1/2$ -equilibrium of support size 3.

Our reduction for small support equilibria involves the following construction, which is very similar to the earlier one.

Reduction 5.1. Given a graph $G = (V, E)$ with adjacency matrix A , parameters $\alpha \in [0, 1], \beta \in [0, 1]$ and $N_2 \in \mathbb{N}$, construct a game $\mathcal{G} = (M_{\text{row}}, M_{\text{col}})$ as follows.

$$M_{\text{row}} = \begin{pmatrix} \alpha A & 0 & \dots & 0 \\ B & & & \\ \vdots & & R & \\ B & & & \end{pmatrix} \quad M_{\text{col}} = \begin{pmatrix} \alpha A & B^\top & \dots & B^\top \\ 0 & & & \\ \vdots & & J - R & \\ 0 & & & \end{pmatrix}, \quad (5.1)$$

where B is an $N_1 \times n$ matrix whose entries are i. i. d. $\{0, 1\}$ variables with expectation β . As in [Reduction 3.2](#), $N_1 = n^c$ where $c = (c_2 + 1) \ln(1/\beta)$ and c_2 is the constant from [Lemma 3.1](#). Each payoff matrix of (5.1) contains N_2 identical copies of B , and R is an $N_1 N_2 \times N_1 N_2$ uniformly random $\{0, 1\}$ matrix.

Theorem 1.6 restated. For every $\eta > 0$ there exists $C = C(\eta) > 0$ such that finding a $(1/2 - \eta)$ -equilibrium with support at most $(\ln n)/2$ is as hard as the $(C \ln n)$ -Hidden Clique Problem.

Proof. Given a graph $G \sim \text{HC}(n, C \ln n)$ (for C to be defined soon), construct the game \mathcal{G} as in [Reduction 5.1](#) with the following parameters:

$$\alpha = 1/2 + \eta/8,$$

$$\beta = \alpha + (1/2 - \eta) - \eta^2/16 = 1 - \frac{7}{8}\eta - \eta^2/16.$$

We choose the dimension N_2 as $N_2 = n^{c'}$, where c' is chosen to satisfy $(\eta^2/16)^2 c' = 4c$. Since $c = (c_2 + 1) \ln 1/\beta = \Theta(\eta)$ we have that $c' = \Theta(1/\eta^3)$. We choose the density C of the hidden clique to be $C = c'/2 + 1$.

Note that the number of strategies is $n + N_1 N_2 = n + n^{c'+c}$, so that we are looking for an equilibrium with support at most $(1/2) \ln(n + n^{c'+c}) = C' \ln n$ for some $c'/2 \leq C' \leq c'/2 + 1$ (assuming η is sufficiently small).

The completeness follows easily. Suppose G contains a hidden clique of size $C \ln n > C' \ln n$. Then if both players play uniformly over the same subset of $C' \ln n$ clique vertices, they both achieve reward α . The probability that, say, the row player can gain $1/2 - \eta$ (i. e., get payoff $\alpha + 1/2 - \eta = \beta + \eta^2/16$) by deviating to some row in B (note that he can only deviate to rows in copies of B) is by the Chernoff bound [Lemma 2.1](#) and a union bound at most

$$N_1 e^{-(\eta^2/16)^2 C' \ln n} = n^{c - (\eta^2/16)^2 C'} \leq n^{c - (\eta^2/16)^2 c'/2} = n^{-c}.$$

Now, for the soundness, consider any $(1/2 - \eta)$ -equilibrium (x, y) . Let us first show that both players must have most of their probability concentrated in the αA block. Let $p = \|x_{[n]}\|$ and $q = \|y_{[n]}\|$ (the probabilities that each player plays in the first n rows/columns). Let us consider the two player's incentive to deviate. In the αA block, the row player achieves at most payoff α , and can achieve payoff $\beta - o(1)$ by playing uniformly over all rows in B (w. h. p. this is true for all distributions over columns in $[n]$). In particular, there exists at least one row in B in which the row player can achieve this value. Choose the row in B which maximizes this payoff. Now consider the right hand side of the payoff matrix. For any row in B , there are N_2 corresponding rows in R (one corresponding to each copy of B). Since the column player's support is at most $C' \ln n$, the probability that regardless of the column player's choice of support, there will be at least one row among these that has all 1's in the corresponding positions is at least

$$\begin{aligned} 1 - n^{C' \ln n} (1 - 2^{-C' \ln n})^{N_2} &\geq 1 - n^{C' \ln n} \exp(-2^{-C' \ln n} N_2) \\ &= 1 - n^{C' \ln n} e^{-n^{-C'+c'}} \\ &\geq 1 - n^{C' \ln n} e^{-n^{c'/2-1}}. \end{aligned}$$

Thus, w. h. p. for every row in B and every possible strategy for the column player (up to the restriction on support size), there is a row in R corresponding to the correct row in B such that the row player would achieve payoff 1 in R (by deviating to this row). In particular, this is true for the row we chose in B (where the row player can achieve value $\beta - o(1)$)—as before, we will ignore this $o(1)$.

Therefore, when the column player plays in $[n]$, the row player can gain at least $p(\beta - \alpha)$ (since we chose the row in B which maximizes the row player's payoff, there is no loss associated with deviating within B). Let $\lambda \in [0, 1]$ be the payoff that the row player receives in R before deviating (the column player receives $1 - \lambda$ here). Then when the column player plays outside $[n]$, the row player can gain at least $1 - (1 - p)\lambda$. To summarize, by deviating, the row player can gain at least

$$pq(\beta - \alpha) + (1 - q)(1 - (1 - p)\lambda).$$

Similarly, the column player's incentive to deviate is at least

$$pq(\beta - \alpha) + (1 - p)(1 - (1 - q)(1 - \lambda)).$$

On average, the two players' incentive to deviate is at least

$$pq(\beta - \alpha) + \frac{(1 - p) + (1 - q)}{2} - \frac{(1 - p)(1 - q)}{2} = pq(\beta - \alpha) + \frac{1 - pq}{2},$$

and since this incentive is at most $1/2 - \eta$, we have

$$pq(\beta - \alpha - 1/2) + 1/2 \leq 1/2 - \eta, \tag{5.2}$$

or

$$pq \geq \frac{\eta}{(1/2 - (\beta - \alpha))} = \frac{\eta}{(\eta + \eta^2/16)} > 1 - \eta/16.$$

Now it remains to bound the conditional value w that is achieved in the αA block. This we can do using the same analysis as above, but substituting w for α . Making this substitution in (5.2) and solving for w we have

$$\begin{aligned} w &\geq \frac{\eta}{pq} + \beta - \frac{1}{2} = \frac{\eta}{pq} + \frac{1}{2} - \frac{7\eta}{8} - \frac{\eta^2}{16} \\ &\geq \frac{1}{2} + \frac{\eta}{8} - \frac{\eta^2}{16} \\ &= \alpha \left(1 - \frac{\eta^2}{8 + \eta} \right) \geq \alpha \left(1 - \frac{\eta^2}{8} \right). \end{aligned}$$

We are now in a position to apply Lemma 3.3. As stated Lemma 3.3 only applies to Reduction 3.2 and not the present reduction but it can be verified that it works also in this case (what is needed is that the size of B is the same, and the presence of an approximate equilibrium with most mass in A and near-optimal value in A). We have $t = \eta/16$ and $s = \eta^2/8$ which is easily checked to satisfy the condition $1 - 2t - 3\sqrt{s}/2 \geq \alpha + (1/2 - \eta)$. Thus we conclude the existence of a dense bipartite subgraph which, as before, allows us to reconstruct the hidden clique by Lemma 3.1. \square

Note that in the completeness case of the above theorem, an exact equilibrium is not guaranteed to exist. In that sense, the above theorem has a much smaller gap between the completeness and hardness than in the other problems we have considered. However, modifying parameters in the above proof (specifically, for $\alpha = 1/4$, and $\beta = \alpha - 1/900$), such hardness can be shown for a smaller additive approximation:

Theorem 5.2. *There exists some constant $C > 0$, such that finding a $1/30$ -equilibrium with support of size at most $(\ln n)/2$ in a two-player game which admits an exact Nash equilibrium with support of size $(C \ln n)$ is as hard as the $(C \ln n)$ -Hidden Clique Problem.*

6 Computing approximate pure Bayes-Nash equilibrium

Bayesian games model the situation where the players' knowledge of the world is incomplete. In this paper we focus on Bayesian games with two players, but the results generalize to an arbitrary number of players. More details on Bayesian games can be found in most Game Theory textbooks, for example in [14].

In a *Bayesian game* the payoff of the players depends on the state of the world in addition to the players' strategies. In a situation with two players, the row player and the column player, each player is presented with a signal, called type, about the state of the world $\theta_{\text{row}} \in \Theta_{\text{row}}$ and $\theta_{\text{col}} \in \Theta_{\text{col}}$, respectively. The types are distributed according to some joint distribution \mathcal{P} and are not necessarily independent. The types determine the values in the payoff matrices $M_{\text{row}}(\theta_{\text{row}}, \theta_{\text{col}})$ and $M_{\text{col}}(\theta_{\text{row}}, \theta_{\text{col}})$, but not their dimensions. Denote the set of rows and columns in any of these matrices by S_{row} and S_{col} , respectively. Each player chooses an action $s_{\text{row}} \in S_{\text{row}}$ and $s_{\text{col}} \in S_{\text{col}}$ from their respective set of actions. The payoff function of the first player is thus $u_{\text{row}}(s_{\text{row}}, s_{\text{col}}, \theta_{\text{row}}, \theta_{\text{col}}) = M_{\text{row}}(\theta_{\text{row}}, \theta_{\text{col}})_{s_{\text{row}}, s_{\text{col}}} \in [0, 1]$. The payoff function u_{col} is defined similarly. The payoff matrices, that depend on the players' types, as well as the distribution on types is known to the players ahead of the game.

A *pure strategy* for the row player in a Bayesian game is a function (that by a slight abuse of notation) we denote by $s_{\text{row}} : \Theta_{\text{row}} \rightarrow S_{\text{row}}$ that for each type θ_{row} as observed by row player associates a strategy $s_{\text{row}}(\theta_{\text{row}})$ that the player chooses to execute. A pure strategy $s_{\text{col}} : \Theta_{\text{col}} \rightarrow S_{\text{row}}$ is defined similarly.

Denote by $\mathcal{P}_{\theta_{\text{row}}}$ the distribution on the column player's types θ_{col} conditioned on the type θ_{row} being observed. For a pair of pure strategies $(s_{\text{row}}, s_{\text{col}})$ the payoff function of the row player is given by

$$p_{\text{row}}(\theta_{\text{row}}) = \mathbb{E}_{\theta_{\text{col}} \sim \mathcal{P}_{\theta_{\text{row}}}} [u_{\text{row}}(s_{\text{row}}(\theta_{\text{row}}), s_{\text{col}}(\theta_{\text{col}}), \theta_{\text{row}}, \theta_{\text{col}})].$$

A *pure strategy Nash equilibrium* in a Bayesian game, is a pair of functions $s_{\text{row}}, s_{\text{col}}$ such that for all types observed, neither player has an incentive to deviate from his current strategy. In other words, for each θ_{row} , and for each $s'_{\text{row}} \in S_{\text{row}}$,

$$p_{\text{row}}(\theta_{\text{row}}) \geq \mathbb{E}_{\theta_{\text{col}} \sim \mathcal{P}_{\theta_{\text{row}}}} [u_{\text{row}}(s'_{\text{row}}, s_{\text{col}}(\theta_{\text{col}}), \theta_{\text{row}}, \theta_{\text{col}})],$$

and a similar condition holds for p_{col} .

Since a pure Nash equilibrium need not exist in non-Bayesian games, it need not exist in Bayesian games either. Moreover, while verifying whether a non-Bayesian two player game has a pure Nash equilibrium is trivial, verifying whether a pure Bayesian Nash equilibrium exists is NP-hard [7]. Furthermore, as the example in [7] demonstrates, this problem remains hard even when the distribution on types is uniform and the payoff does not depend on the players' types.

A *pure ε -Bayesian Nash equilibrium (ε -BNE)* is defined similarly to an ε -Nash equilibrium. For each observed type θ_{row} , the incentive to deviate should be bounded by ε :

$$p_{\text{row}}(\theta_{\text{row}}) \geq \mathbb{E}_{\theta_{\text{col}} \sim \mathcal{P}_{\theta_{\text{row}}}} [u_{\text{row}}(s'_{\text{row}}, s_{\text{col}}(\theta_{\text{col}}), \theta_{\text{row}}, \theta_{\text{col}})] - \varepsilon.$$

A similar requirement should hold for the column player.

We show that for general distributions on types and for some small constant ε , finding a pure ε -BNE in games where a pure BNE exists is still NP-hard. On the other hand, we also show that if the distribution on the players' types is uniform, whenever a pure BNE exists, a pure ε -BNE can be found in quasi-polynomial time.

6.1 General distributions on types

We show that for some constant ε , determining whether a pure strategy ε -Bayes-Nash equilibrium exists is NP-hard. Specifically, $\varepsilon = 0.004$ suffices. We prove:

Theorem 1.7 restated. Let $\varepsilon = 0.004$. Then given a Bayesian game that admits a pure BNE, it is NP-hard to find a pure ε -BNE for the game. Moreover, it is NP-hard to solve the promise problem of distinguishing games that admit a pure BNE from games that do not admit a pure ε -BNE.

Proof. We give a reduction from the problem of 3-coloring 4-regular graphs, which is known to be NP-complete [8]. Let $G = (V, E)$ be a 4-regular graph with $|V| = n$. The edges of G can be properly colored with 5 colors using Vizing's algorithm. In other words, we can compute a coloring $c^* : E \rightarrow \{1, \dots, 5\}$ such that for every two edges e_1, e_2 incident to the same vertex, $c^*(e_1) \neq c^*(e_2)$.

We now design a Bayesian game such that:

- if G is 3-colorable, then the game admits a pure BNE;
- if G is not 3-colorable, then the game admits no pure ε -BNE.

The type of each of the players corresponds to a vertex in the graph, thus $\Theta_{\text{row}} = \Theta_{\text{col}} = V$. The distribution on types is such that with probability $4/5$, (θ_1, θ_2) is a random edge in E ; with probability $1/5$, $(\theta_1, \theta_2) = (\theta, \theta)$ is the same random vertex in V . We note that in this example there is a high degree of correlation between the types of the two players.

Each player has 6 strategies available to him: $S_{\text{row}} = S_{\text{col}} = \{1, 2, 3\} \times \{0, 1\}$. The first coordinate of $s_{\text{row}}(v)$ should be thought of as the “color” assigned by player A to vertex v , while the second coordinate is either “head” or “tails,” the use of which will be explained later. We denote the first coordinate by $s_{\text{row}}(v)_1 \in \{1, 2, 3\}$ and the second by $s_{\text{row}}(v)_2 \in \{0, 1\}$.

The payoff functions are defined as follows (s_{row} and s_{col} stand for $s_{\text{row}}(\theta_{\text{row}})$ and $s_{\text{col}}(\theta_{\text{col}})$, respectively):

$\theta_{\text{row}} \stackrel{?}{=} \theta_{\text{col}}$	$(s_{\text{row}})_1 \stackrel{?}{=} (s_{\text{col}})_1$	$(s_{\text{row}})_2 \stackrel{?}{=} (s_{\text{col}})_2$	u_{row}	u_{col}
yes	yes		1	1
yes	no	yes	0.64	0
yes	no	no	0	0.64
no	no		1	1
no	yes	yes	$0.01 \times 2^{c^*(\theta_{\text{row}}, \theta_{\text{col}})}$	0
no	yes	no	0	$0.01 \times 2^{c^*(\theta_{\text{row}}, \theta_{\text{col}})}$

In other words, the first coordinate of the strategy represents vertex color. The players are rewarded for producing a consistent 3-coloring: for the same vertex θ the colors $(s_{\text{row}}(\theta))_1$ and $(s_{\text{col}}(\theta))_1$ should match; for different vertices connected by an edge the colors $(s_{\text{row}}(\theta_{\text{row}}))_1$ and $(s_{\text{col}}(\theta_{\text{col}}))_1$ should differ. If these conditions are satisfied, both players are rewarded with a payoff of 1. If a certain edge or vertex fails to satisfy the conditions, the game becomes a zero-sum game that depends on the second coordinate of the players’ strategies. The payoff can take values of 0.02, 0.04, 0.08, 0.16, and 0.32 for edge pairs $(\theta_{\text{row}}, \theta_{\text{col}})$ and the value of 0.64 for vertex pairs (θ, θ) . What makes these values interesting every combination of these values with coefficients of $-1, 0, +1$ adds up to a non-zero integer multiple of 0.02 (except when we take all-zero coefficients). Note that for every vertex $v \in V$, when $\theta_{\text{row}} = v$, the column player’s type is distributed uniformly over v and its four neighbors. By our edge-coloring, and the way we designed our payoffs, this means that by flipping the value of its coin, the change in the row player’s payoff at this vertex is the average of five of the six values 0.02, 0.04, 0.08, 0.16, 0.32 and 0.64 with $-1, 0, +1$ coefficients, and thus is an integer multiple of $(1/5) \cdot 0.02 = 0.004$ (and is only 0 if the change is identically 0 regardless of the column player’s type), and similarly for the column player.

If the graph G admits a 3-coloring $\mathcal{C} : V \rightarrow \{1, 2, 3\}$, then the pure strategy $s_{\text{row}}(v) = s_{\text{col}}(v) = (\mathcal{C}(v), 0)$ yields the optimal possible payoff of 1 for both players, and thus is a pure BNE.

On the other hand, suppose that G is not 3-colorable. Let $s_{\text{row}}(v), s_{\text{col}}(v)$ be a set of pure strategies that is an ε -BNE. Denote the expected payoff functions under these strategies by $u_{\text{row}}(v)$ and $u_{\text{col}}(v)$. If the row player deviates in her strategy from $(s_{\text{row}}(v)_1, s_{\text{row}}(v)_2)$ to $(s_{\text{row}}(v)_1, 1 - s_{\text{row}}(v)_2)$ her payoff only changes for type pairs where the 3-coloring conditions are not satisfied (i. e., a vertex which is assigned different colors by the two players, or an edge where the two endpoints are given the same color). Let us

denote this change by $\Delta_{\text{row}}(v)$ and similarly $\Delta_{\text{col}}(v)$ for the change in payoff the column player receives when flipping the value of the coin at vertex v . The contribution of each edge $(v_{\text{row}}, v_{\text{col}})$ to $\Delta_{\text{row}}(v_{\text{row}})$ and $\Delta_{\text{col}}(v_{\text{col}})$ cancels out, as does the contribution of each vertex type pair to $\Delta_{\text{row}}(v)$ and $\Delta_{\text{col}}(v)$. Hence, summing over all edges,

$$\sum_{v \in V} (\Delta_{\text{row}}(v) + \Delta_{\text{col}}(v)) = 0. \quad (6.1)$$

Since G is not 3-colorable, for at least one vertex v , the row player's coloring (when $\theta_{\text{row}} = v$) is not always consistent with the column player's coloring. As explained earlier, in this case $\Delta_{\text{row}}(v)$ must be a non-zero integer multiple of 0.004. By (6.1) this means that there is some (possibly different) vertex v such that $\Delta_{\text{row}}(v) > 0$ or $\Delta_{\text{col}}(v) > 0$ and thus $\Delta_{\text{row}}(v) \geq 0.004$ or $\Delta_{\text{col}}(v) \geq 0.004$. This contradicts the assumption that $(s_{\text{row}}, s_{\text{col}})$ is an ε -BNE. \square

6.2 Uniform distribution on types

In this section we show that in the case where the distribution on types is uniform, a pure ε -BNE can be computed in quasi-polynomial time. This contrasts with the previously noted fact from [7] that computing a pure BNE is NP-hard even in this special case. As for other quasi-polynomial time computable approximate equilibria we have considered, whose exact variants are NP-hard, this problem is also as hard as Hidden Clique.

Theorem 1.9 restated. There exists some constant $C > 0$ such that finding a $1/30$ -approximate pure BNE in a two-player Bayesian games with uniformly distributed types and in which a pure BNE exists is as hard as finding a hidden clique of size $C \ln n$.

Proof. This follows immediately from Theorem 5.2. Let M_{row} and M_{col} be as in Reduction 5.1 with parameters set in order to achieve the $1/30$ -hardness of Theorem 5.2. Consider the Bayesian game in which both players have exactly $C' \ln n$ types, the distribution over types is uniform, and the payoff matrices for each player are always *the same* M_{row} and M_{col} as above. Note that when there is a hidden clique, the following is an pure BNE: choose a subset S of $C' \ln n$ clique vertices, and let each player play according to a one-to-one mapping from their type set to S . On the other hand, it is easy to see that every pair of pure strategies for the Bayesian game corresponds to a pair of mixed strategies (with support at most $C' \ln n$) for the original game, where the incentive to deviate for, say, the row player, is the expected incentive to deviate in the Bayesian game (over all choices of types for the row player). In particular, if a player always has incentive at most $1/30$ to deviate in the Bayesian game, their incentive will be at most $1/30$ in the original game, and so by Theorem 5.2, we can recover the hidden clique. \square

Next we give a quasi-polynomial algorithm for finding an ε -BNE in Bayesian games with uniformly distributed types. The algorithm easily generalizes to the case when the distribution on types is a product distribution, i. e., when θ_{row} is independent from θ_{col} . To simplify the presentation we further assume that the type space is of equal size for both players, i. e., $|\Theta_{\text{row}}| = |\Theta_{\text{col}}| = k$. We prove:

Theorem 6.1 (Detailed statement of Theorem 1.8). *In a two-player Bayesian game, suppose that the types are distributed uniformly on the space $\Theta_{\text{row}} \times \Theta_{\text{col}}$, and that $|\Theta_{\text{row}}| = |\Theta_{\text{col}}| = k$, and $|S_{\text{row}}| = |S_{\text{col}}| = n$. Assuming that a pure BNE exists, we can find a pure ε -BNE in time $n^{O((\ln n + \ln k)/\varepsilon^2)}$.*

Remark 6.2. The assumption in [Theorem 1.8](#) that an exact equilibrium exists can be relaxed to a pure $(\varepsilon/2)$ -BNE equilibrium existing.

Proof. The proof is similar in spirit to the quasi-polynomial ε -Nash algorithm of [\[18\]](#), but some additional work is needed. We first (approximately) guess the payoffs and the allowed strategies for both players for all possible types. We then use linear programming to produce a (not necessarily pure) $(3\varepsilon/4)$ -BNE. We then sample from this approximate non-pure BNE to obtain a pure ε -BNE.

For simplicity, denote $\Theta_{\text{row}} = \Theta_{\text{col}} = \{1, \dots, k\}$ and $S_{\text{row}} = S_{\text{col}} = \{1, \dots, n\}$. Let $s_{\text{row}} : \Theta_{\text{row}} \rightarrow S_{\text{row}}$, $s_{\text{col}} : \Theta_{\text{col}} \rightarrow S_{\text{col}}$ be a pair of pure equilibrium strategies (that we assume exist). Let p_{ij}^{row} and p_{ij}^{col} be the corresponding payoff values. In other words, p_{ij}^{row} is the expected payoff the row player gets when his type is $i \in \{1, \dots, k\}$, he plays strategy $j \in \{1, \dots, n\}$, and the column player plays according to the equilibrium. The equilibrium assumption is that for all i ,

$$p_{i, s_{\text{row}}(i)}^{\text{row}} = \max_j p_{ij}^{\text{row}},$$

and a similar condition holds for p^{col} . We first claim that we can recover all the values of p^{row} and p^{col} within an error of $\varepsilon/4$ with probability $> n^{-O((\ln(nk))/\varepsilon^2)}$.

Claim 6.3. *There is a randomized polynomial time algorithm that with probability $> n^{-O((\ln(nk))/\varepsilon^2)}$ outputs values $q_{ij}^{\text{row}}, q_{ij}^{\text{col}}$ such that for all i, j , $|q_{ij}^{\text{row}} - p_{ij}^{\text{row}}| < \varepsilon/8$ and $|q_{ij}^{\text{col}} - p_{ij}^{\text{col}}| < \varepsilon/8$.*

Proof. We show how to approximate p^{row} well with probability $> n^{-O((\ln(nk))/\varepsilon^2)}$. The claim follows by approximating p^{row} and p^{col} independently. Pick a uniformly random subset (with repetitions) of $m = \lceil (40 \ln(nk))/\varepsilon^2 \rceil$ of the column player's types $t_1, \dots, t_m \in \{1, \dots, k\}$. For each type t_r guess (uniformly) the column player's strategy s_r on this type. With probability n^{-m} we correctly guess all strategies, i. e., $s_r = s_{\text{col}}(t_r)$ for all r . Set

$$q_{ij}^{\text{row}} = \frac{1}{m} \sum_{r=1}^m u_{\text{row}}(j, s_r, i, t_r).$$

In other words, we calculate an estimate of the expected payoff p^{row} using only the value at the types we've guessed. By the Chernoff-Hoeffding bound [Lemma 2.1](#), for each (i, j) the probability

$$\mathbf{P}[|q_{ij}^{\text{row}} - p_{ij}^{\text{row}}| \geq \varepsilon/8] \leq 2 \cdot \exp(-m\varepsilon^2/32) < 1/(4nk).$$

The claim follows by union bound. □

Allowing for a blow-up of $n^{O((\ln(nk))/\varepsilon^2)}$ in running time, we may assume from now on that the correct values of q^{row} and q^{col} had been computed (up to an error of $\varepsilon/8$), since in the end we can check whether the set of pure strategies obtained is a ε -BNE.

Next, we formulate a linear program that obtains a $(3\varepsilon/4)$ -BNE with payoff values close to $q^{\text{row}}, q^{\text{col}}$. The variables of the program are X_{ij}^{row} and X_{ij}^{col} , where X_{ij}^{row} corresponds to the probability that the row player plays strategy j on type i and similarly for X_{ij}^{col} . For each type i let

$$M_i^{\text{row}} := \max_j q_{ij}^{\text{row}}$$

be the maximum possible payoff the row player can attain on type i under payoffs q_{ij}^{row} . We only allow X_{ij}^{row} to be non-zero when the corresponding payoff $q_{ij}^{\text{row}} > M_i^{\text{row}} - \varepsilon/8$. This guarantees that the solution we obtain is a $(3\varepsilon/4)$ -BNE. In addition, we enforce the payoffs to be close to q_{ij}^{row} . For each i, j we have the constraint

$$\left| q_{ij}^{\text{row}} - \frac{1}{k} \cdot \sum_{y,z} X_{yz}^{\text{col}} \cdot u_{\text{row}}(j, z, i, y) \right| < \frac{\varepsilon}{4},$$

and a similar condition on the X^{row} 's. This linear program is feasible since the pure BNE we assumed exists is a solution to it. Denote the resulting expected payoffs by v_{ij}^{row} and v_{ij}^{col} .

The solution thus obtained is a $(3\varepsilon/4)$ -BNE. To see this, we actually observe that the equilibrium is well supported: for each type i , each strategy j with $X_{ij}^{\text{row}} \neq 0$ the payoff

$$v_{ij}^{\text{row}} \geq M_i^{\text{row}} - \frac{3\varepsilon}{8} \geq \max_{j'} v_{ij'}^{\text{row}} - \frac{5\varepsilon}{8}.$$

To complete the proof we now obtain a set of pure strategies $s'_{\text{row}} : \Theta_{\text{row}} \rightarrow S_{\text{row}}$, $s'_{\text{col}} : \Theta_{\text{col}} \rightarrow S_{\text{col}}$ by sampling $j = s'_{\text{row}}(i)$ according to the probability distribution X_{ij}^{row} .

If $k = \Omega((\ln(nk))/\varepsilon^2)$, then once again by the Chernoff-Hoeffding bound the payoffs will be $\varepsilon/8$ -close to the payoffs v_{ij}^{row} , and thus the resulting game will be in an ε -BNE. On the other hand if $k = O((\ln(nk))/\varepsilon^2)$, i. e., the number of types is small, we can find the exact BNE by brute force, completing the proof of the theorem. \square

Acknowledgments

We are grateful to Kristoffer Arnsfelt Hansen for bringing errors in a previous version of this paper to our attention, and to several anonymous reviewers for their helpful comments.

References

- [1] NOGA ALON, MICHAEL KRIVELEVICH, AND BENNY SUDAKOV: Finding a large hidden clique in a random graph. *Random Structures & Algorithms*, 13(3-4):457–466, 1998. Preliminary version in SODA'98. [doi:10.1002/(SICI)1098-2418(199810/12)13:3/4<457::AID-RSA14>3.0.CO;2-W] 118
- [2] PER AUSTRIN, MARK BRAVERMAN, AND EDEN CHLAMTÁČ: Inapproximability of NP-complete variants of Nash equilibrium. In *Proc. 14th Internat. Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'11)*, pp. 13–25. Springer, 2011. [doi:10.1007/978-3-642-22935-0_2] 117
- [3] HARTWIG BOSSE, JAROSLAW BYRKA, AND EVANGELOS MARKAKIS: New algorithms for approximate Nash equilibria in bimatrix games. *Theoret. Comput. Sci.*, 411(1):164–173, 2010. Preliminary version in WINE'07. [doi:10.1016/j.tcs.2009.09.023] 118

- [4] S. CHARLES BRUBAKER AND SANTOSH VEMPALA: Random tensors and planted cliques. In *Proc. 13th Internat. Workshop on Randomization and Computation (RANDOM'09)*, pp. 406–419, 2009. [[doi:10.1007/978-3-642-03685-9_31](https://doi.org/10.1007/978-3-642-03685-9_31)] 119
- [5] XI CHEN, XIAOTIE DENG, AND SHANG-HUA TENG: Settling the complexity of computing two-player Nash equilibria. *J. ACM*, 56(3):14:1–14:57, 2009. Preliminary versions in FOCS'06 by Chen and Deng and by Chen, Deng, and Teng. [[doi:10.1145/1516512.1516516](https://doi.org/10.1145/1516512.1516516)] 118
- [6] VINCENT CONITZER AND TUOMAS SANDHOLM: Complexity results about Nash equilibria. In *Proc. 18th Internat. Joint Conf. on Artificial Intelligence (IJCAI'03)*, pp. 765–771, 2003. Article accessible at IJCAI. [[ACM:1630659.1630770](https://doi.org/10.1145/1630659.1630770)] 120
- [7] VINCENT CONITZER AND TUOMAS SANDHOLM: New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2):621–641, 2008. [[doi:10.1016/j.geb.2008.02.015](https://doi.org/10.1016/j.geb.2008.02.015)] 121, 135, 137
- [8] DAVID P. DAILEY: Uniqueness of colorability and colorability of planar 4-regular graphs are NP-complete. *Discrete Mathematics*, 30(3):289–293, 1980. [[doi:10.1016/0012-365X\(80\)90236-8](https://doi.org/10.1016/0012-365X(80)90236-8)] 135
- [9] CONSTANTINOS DASKALAKIS, ARANYAK MEHTA, AND CHRISTOS H. PAPADIMITRIOU: Progress in approximate Nash equilibria. In *Proc. 8th ACM Conf. on Electronic Commerce (EC'07)*, pp. 355–358. ACM Press, 2007. [[doi:10.1145/1250910.1250962](https://doi.org/10.1145/1250910.1250962)] 118
- [10] CONSTANTINOS DASKALAKIS, ARANYAK MEHTA, AND CHRISTOS H. PAPADIMITRIOU: A note on approximate Nash equilibria. *Theoret. Comput. Sci.*, 410(17):1581–1588, 2009. Preliminary version in WINE'06. [[doi:10.1016/j.tcs.2008.12.031](https://doi.org/10.1016/j.tcs.2008.12.031)] 118, 120, 128, 132
- [11] TOMÁS FEDER, HAMID NAZERZADEH, AND AMIN SABERI: Approximating Nash equilibria using small-support strategies. In *Proc. 8th ACM Conf. on Electronic Commerce (EC'07)*, pp. 352–354, 2007. [[doi:10.1145/1250910.1250961](https://doi.org/10.1145/1250910.1250961)] 120
- [12] URIEL FEIGE AND ROBERT KRAUTHGAMER: The probable value of the Lovász–Schrijver relaxations for Maximum Independent Set. *SIAM J. Comput.*, 32(2):345–370, 2003. [[doi:10.1137/S009753970240118X](https://doi.org/10.1137/S009753970240118X)] 119
- [13] ALAN M. FRIEZE AND RAVI KANNAN: A new approach to the planted clique problem. In *IARCS Ann. Conf. on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'08)*, pp. 187–198. Schloss Dagstuhl, 2008. [[doi:10.4230/LIPIcs.FSTTCS.2008.1752](https://doi.org/10.4230/LIPIcs.FSTTCS.2008.1752)] 119
- [14] DREW FUDENBERG AND JEAN TIROLE: *Game Theory*. MIT Press, 1991. 134
- [15] ITZHAK GILBOA AND EITAN ZEMEL: Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989. [[doi:10.1016/0899-8256\(89\)90006-7](https://doi.org/10.1016/0899-8256(89)90006-7)] 120

- [16] ELAD HAZAN AND ROBERT KRAUTHGAMER: How hard is it to approximate the best Nash equilibrium? *SIAM J. Comput.*, 40(1):79–91, 2011. Preliminary version in [SODA'09](#). [[doi:10.1137/090766991](#)] [118](#), [119](#), [123](#), [124](#), [130](#)
- [17] WASSILY Hoeffding: Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.*, 58(301):13–30, 1963. [[doi:10.1080/01621459.1963.10500830](#)] [122](#)
- [18] RICHARD J. LIPTON, EVANGELOS MARKAKIS, AND ARANYAK MEHTA: Playing large games using simple strategies. In *Proc. 4th ACM Conf. on Electronic Commerce (EC'03)*, pp. 36–41. ACM Press, 2003. [[doi:10.1145/779928.779933](#)] [118](#), [120](#), [131](#), [138](#)
- [19] LORENZ MINDER AND DAN VILENCHIK: Small clique detection and approximate Nash equilibria. In *Proc. 13th Internat. Workshop on Randomization and Computation (RANDOM'09)*, pp. 673–685. Springer, 2009. [[doi:10.1007/978-3-642-03685-9_50](#)] [119](#)
- [20] CHRISTOS H. PAPADIMITRIOU: On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. System Sci.*, 48(3):498–532, 1994. [[doi:10.1016/S0022-0000\(05\)80063-7](#)] [120](#)
- [21] HARALAMPOS TSAKNAKIS AND PAUL G. SPIRAKIS: An optimization approach for approximate Nash equilibria. *Internet Mathematics*, 5(4):365–382, 2008. Preliminary version in [WINE'07](#). [[doi:10.1080/15427951.2008.10129172](#)] [118](#), [120](#)

AUTHORS

Per Austrin
 associate professor
 KTH Royal Institute of Technology, Stockholm, Sweden
austrin@kth.se
<http://www.csc.kth.se/~austrin>

Mark Braverman
 assistant professor
 Department of Computer Science, Princeton University, USA
mbraverm@cs.princeton.edu
<http://www.cs.princeton.edu/~mbraverm/>

Eden Chlamtáč
 assistant professor
 Department of Computer Science, Ben Gurion University of the Negev, Be'er Sheva, Israel
chlamtac@cs.bgu.ac.il
<http://www.cs.bgu.ac.il/~chlamtac>

ABOUT THE AUTHORS

PER AUSTRIN graduated from [KTH Royal Institute of Technology](#) in Sweden, 2008; his advisor was [Johan Håstad](#). Following this he did post-docs at [Institut Mittag-Leffler](#), [New York University](#), and the [University of Toronto](#). In his spare time he enjoys writing tepid biographic blurbs.

MARK BRAVERMAN graduated from the [University of Toronto](#) in Canada in 2008 under the supervision of [Stephen Cook](#). He spent two years as a postdoc with [Microsoft Research New England](#) and a year as an assistant professor at the University of Toronto before coming to [Princeton](#).

EDEN CHLAMTÁČ graduated from [Princeton University](#) in Princeton, NJ, in 2008; his advisor was [Sanjeev Arora](#). He then spent two years as a postdoc at the [Weizmann Institute of Science](#), and one and three semesters as a postdoc at the [University of Toronto](#) and [Tel-Aviv University](#), respectively, before coming to [Ben Gurion University](#). If memory serves, when he used to have spare time he would enjoy playing piano and studying languages which rarely did him much good.