On the Complexity of Approximating a Nash Equilibrium

CONSTANTINOS DASKALAKIS, EECS, MIT

We show that computing a relatively (i.e. multiplicatively as opposed to additively) approximate Nash equilibrium in two-player games is PPAD-complete, even for constant values of the approximation. Our result is the first constant inapproximability result for Nash equilibrium, since the original results on the computational complexity of the problem [DGP06; CD06]. Moreover, it provides an apparent—assuming that PPAD is not contained in $\texttt{TIME}(n^{O(\log n)})$—dichotomy between the complexities of additive and relative approximations, as for constant values of additive approximation a quasi-polynomial-time algorithm is known [LMM03]. Such a dichotomy does not exist for values of the approximation that scale inverse-polynomially with the size of the game, where both relative and additive approximations are PPAD-complete [CDT06]. As a byproduct, our proof shows that (unconditionally) the [LMM03] sparse-support lemma cannot be extended to relative notions of constant approximation.

Categories and Subject Descriptors: F.2.0 [**Analysis of Algorithms and Problem Complexity**]: General; J.4 [**Computer Applications**]: Social and Behavioral Sciences—Economics

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Complexity, Approximation, Nash Equilibrium, PPAD-Completeness, Game Theory

ACM Reference Format:
Constantinos Daskalakis, 2013. On the Complexity of Approximating a Nash Equilibrium. *ACM Trans. Algor.* XX, XX, Article XX (January 2013), 34 pages.
DOI:http://dx.doi.org/10.1145/0000000.0000000

1. INTRODUCTION

In the wake of the intractability results for Nash equilibrium [DGP06; CD06; CDT06], researchers undertook the important—and indeed very much algorithmic—task of understanding the complexity of approximate Nash equilibria. A positive outcome to this investigation, i.e. a computationally efficient approximation algorithm, would be useful in practice. More importantly, it would alleviate the negative implications of the aforementioned hardness results to the predictive power of the Nash equilibrium, showing that strategic behavior could in principle converge to approximate equilibrium states. On the other hand, should this investigation result in inapproximability results, this would stress even more the barrier posed by computation to the universal applicability of the Nash equilibrium. Despite the timeliness of the investigation and considerable effort towards obtaining upper [LMM03; KPS06; DMP06; DMP07; KS07; FNS07; BBM07; TS07; TS10] and lower [DP09b; HK09] bounds, the approximation complexity of Nash equilibria has remained poorly understood. In this paper we make progress to

This work is supported by a Microsoft Research Faculty Fellowship, a Sloan Foundation Fellowship and NSF Awards CCF-0953960 (CAREER) and CCF-1101491. Part of this work was done while the author was in Microsoft Research New England.

DOI:http://dx.doi.org/10.1145/0000000.0000000

this question by *obtaining the first constant inapproximability results for Nash equilibrium.*

When it comes to approximation, the typical algorithmic approach is to aim for *relative*, i.e. *multiplicative*, approximations to the optimum of the objective function of interest. In a strategic game, however, it is not clear what function to approximate as each player seeks to optimize her own payoff function, which assigns to every possible selection of strategies by the players a real number representing the player's own payoff. To define approximate notions of Nash equilibrium it is more natural to instead relax the optimality conditions that the Nash equilibrium itself imposes. Recall that a Nash equilibrium is a collection of randomized, or *mixed*, strategies (these are distributions over deterministic, or *pure*, strategies,) such that each player's randomization optimizes her payoff in expectation, given the mixed strategies of the other players, and assuming that all players randomize independently from each other. Since a player's expected payoff is a convex combination of her pure strategy payoffs, in order to optimize the player must only use in her mixed strategy pure strategies with optimal expected payoff against the other players' mixed strategies.

Relaxing these conditions to allow for approximation, a *relative ϵ-Nash equilibrium* is a collection of mixed strategies, so that no player uses in her mixed strategy a pure strategy whose payoff fails to be within a relative error of ϵ from the best pure strategy payoff.[1] In an ϵ-Nash equilibrium, the expected payoff of each player is within a relative error of ϵ from her best possible payoff given the opponents' strategies. However, the latter is a strictly weaker requirement, as we can always include in the mixed strategy of a player a poorly performing pure strategy and assign to it a tiny probability so that the expected payoff from the overall mixed strategy is only trivially affected. To distinguish between the two kinds of approximation the literature has converged to the term *ϵ-approximate Nash equilibrium* for the latter, weaker kind of approximation, while the un-quantified term *approximate Nash equilibrium* is reserved for both kinds of approximation.

Despite considerable effort towards algorithms for approximate Nash equilibria, there has been a single positive result for relative approximations, providing a polynomial-time algorithm for 0.5-approximate Nash equilibrium in two-player games with positive payoffs [FNS07]. On the other hand, the investigation of the absolute-error—i.e. additive-error—counterparts of the approximate equilibrium notions defined above has been much more fruitful.[2] Indeed, while additive notions of approximation are less common in algorithms, they appear more benign in this setting. Moreover, they naturally arise in designing simplicial approximation algorithms for Nash equilibria, as an additive error guarantee is directly implied by the Lipschitz properties of Nash's function in the neighborhood of a Nash equilibrium [Sca67]. For finite values of additive approximation, the best polynomial-time algorithm known to date obtains a 0.34-approximate Nash equilibrium [TS07], and a 0.66-Nash equilibrium [KS07] in two-player games whose payoffs are normalized (by scaling) to lie in a unit-length interval.

Clearly, scaling the payoffs of a game changes the approximation guarantee of additive approximations. So the performance of algorithms for additive approximate equilibria is typically compared after an affine transformation that brings the payoffs of the input game into a unit-length interval; where this interval is located is irrelevant

[1]Given this definition, this kind of approximation also goes by the name *ϵ-approximately well-supported Nash equilibrium*. The literature has adopted the shorter name ϵ-Nash equilibrium for convenience.

[2]The additive notions of approximation require that the expected payoff from either the whole mixed strategy of a player or from every pure strategy in its support are within an additive error ϵ from the best possible payoff.

since the additive approximations are payoff-shift invariant. Unlike additive notions of approximation, relative notions are payoff-scale invariant, but not payoff-shift invariant. This distinction turns the two notions of approximation appropriate in different applications. Imagine a play of some game, say Texas hold 'em, in which a player is gaining an expected payoff of $\$1$M from her current strategy, but could improve her payoff to $\$1.1$M with some other strategy. Compare this situation to a play of the same game where the player's payoff is -$\$100$k and could become $\$0$ with a different strategy. Notice that in both cases the improvement in payoff is $\$100$k, but it is debatable whether the incentive of the player to update her strategy is the same in the two situations. If one subscribes to the theory of diminishing marginal utility of wealth [Ber38], the two situations could be very different, making the relative notion of approximation more appropriate. If instead the incentive to switch is perceived to be as strong in the two situations, then the additive notion of approximation is more fitting.

From a computational complexity standpoint, additive and relative approximations have thus far enjoyed similar fate. In two-player games, if the desired approximation scales inverse-polynomially with the size of the game, both relative and additive approximations are PPAD-complete [CDT06]. Hence, unless PPAD$\subseteq$P, there is no fully polynomial-time approximation scheme (FPTAS)[3] for either additive or relative approximations. In the other direction, for both additive and relative notions, there exist efficient algorithms for fixed finite values of ϵ. Even though progress in this frontier has stalled in the past few years, the hope for a polynomial-time approximation scheme (PTAS),[4] at least for additive approximations, ultimately stems from an elegant result due to Lipton, Markakis and Mehta [LMM03]. For any fixed ϵ, this obtains a $n^{O(\log n/\epsilon^2)}$ time algorithm for additive ϵ-Nash equilibrium in normalized bimatrix games (games with payoffs scaled to a unit-length interval,) by establishing that in all such games there exists an additive ϵ-Nash equilibrium of support-size logarithmic in the total number of strategies.[5] Given this structural result, the [LMM03] algorithm finds an approximate equilibrium by performing an exhaustive search over all pairs of mixed strategies with logarithmic support, and in fact it can also be used to optimize some linear objective over the output equilibrium. These properties of the algorithm have been exploited in recent lower bounds for the problem [HK09; DP09b], albeit these fall short from a quasi-polynomial-time lower bound for additive approximations. On the other hand, even a quasi-polynomial-time approximation algorithm is not known for constant values of relative approximation, and indeed this was posed to the author of this paper as a question by Shang-Hua Teng [Ten08].

Our Results. We show that computing a relative ϵ-Nash equilibrium in two-player games is PPAD-complete even for constant values of ϵ; namely

THEOREM 1.1 (CONSTANT RELATIVE INAPPROXIMABILITY OF NASH EQUILIBRIUM). *For any constant $\epsilon \in [0, 1)$, it is PPAD-complete to find a relative ϵ-Nash equilibrium in bimatrix games with payoffs in $[-1, 1]$. This remains true even if the expected payoffs of both players are positive in every relative ϵ-Nash equilibrium of the game.*

[3] A FPTAS is an approximation algorithm that runs in time polynomial in the size of the game and $1/\epsilon$, where ϵ is the approximation requirement.

[4] A PTAS is an approximation algorithm that runs in time polynomial in the size of the game n but not necessarily $1/\epsilon$, where ϵ is the approximation requirement. In particular, the running time may be of the form $n^{g(1/\epsilon)}$, for some positive function g.

[5] The [LMM03] argument establishes the small-support property for additive ϵ-approximate Nash equilibrium, but it can easily be extended to the stronger notion of additive ϵ-approximately well-supported Nash equilibrium [DP09b].

Our result is the first constant inapproximability result for Nash equilibrium. In particular, it precludes the existence of a quasi-polynomial-time algorithm à la [LMM03] for constant values of relative approximation, unless $\mathtt{PPAD} \subseteq \mathtt{TIME}(\mathtt{n}^{\mathtt{O}(\log \mathtt{n})})$. Under the same assumption, our result provides a dichotomy between the complexity of relative and additive constant approximations. Such a dichotomy has not been shown before, as for approximation values that scale inverse-polynomially with the size of the game the PPAD-hardness results of [CDT06] apply to both notions of approximation. Finally, an [LMM03]-style small support lemma is precluded unconditionally from our proof, which constructs games whose relative ϵ-Nash equilibria have all linear support.

THEOREM 1.2 (NON-EXISTENCE OF LOGARITHMIC SUPPORT RELATIVE ϵ-NASH EQ.). *For all $\epsilon \in [0,1)$ and $N_0 \in \mathbb{N}$ there exists a two-player game with $N \geq N_0$ pure strategies per player such that in all relative ϵ-Nash equilibria of this game the mixed strategies of both players have support of size at least $\alpha \cdot N$, where $\alpha \in (0,1)$ is some absolute constant that does not depend on ϵ, N_0, N.*

Observe that, if a game's payoffs are all positive (or all negative) and their absolute values lie in some interval $[m, M]$, where $1 \leq \frac{M}{m} \leq c$ for some absolute constant $c \geq 1$ (let us call these games *c-balanced*,) then the relative approximation problem can be polynomial-time reduced to the additive approximation problem in normalized games, i.e. games with payoffs in a unit-length interval; see the discussion following Remark 2.2. Hence, in view of the quasi-polynomial time algorithm of [LMM03] for additive approximate Nash equilibria, and unless $\mathtt{PPAD} \subseteq \mathtt{TIME}(\mathtt{n}^{\mathtt{O}(\log \mathtt{n})})$, we cannot hope to extend Theorem 1.1 to the special class of c-balanced games for any constant c. On the other hand, our result may very well extend to games with payoffs in $[0,1]$ or in $[-1,0]$, which are not c-balanced for any constant c. We believe that these classes of games are also PPAD-complete for constant values of relative approximation and that similar to ours but more tedious arguments may prove such lower bounds. We leave this as an open problem from this work.

Finally, it is not clear how tight our lower bound is with regards to the value of the approximation. For $\epsilon = 1$, a trivial algorithm yields a relative 1-Nash equilibrium for games with payoffs in $[0,1]$, while for win-lose games [AKV05] with payoffs in $\{-1,0\}$ we obtain a polynomial-time algorithm by approximating these games with zero-sum games.

THEOREM 1.3 (EFFICIENT-BUT-LOOSE APPROXIMATIONS). *There is a polynomial-time algorithm for computing a relative 1-Nash equilibrium in bimatrix games with payoffs in either $\{-1,0\}$ or $[0,1]$.*

The proof of this theorem is provided in Section E. We leave as an open question whether this upper bound can be extended to games with payoffs in $[-1,0]$ or in $[-1,1]$, or whether our lower bound can be extended to $\epsilon \geq 1$.

Results for Polymatrix Games. To obtain Theorem 1.1, we prove as an intermediate step a similar (and somewhat stronger) lower bound for the well-studied class of *(graphical) polymatrix games*, which is interesting in its own right. In a polymatrix game the players are nodes of a graph and participate in 2-player-game interactions with each of their neighbors, summing up their payoffs from all these interactions. (For a formal definition, see Section 2.3.) Polymatrix games are guaranteed to have exact Nash equilibria in rational numbers [EY07], their exact equilibrium computation problem was shown to be PPAD-complete [DGP06; EY07], an FPTAS has been precluded assuming $\mathtt{PPAD} \not\subseteq \mathtt{P}$ [CDT06], and they are poly-time solvable if they are zero-sum [DP09a; CD11]. We establish the following lower bound for these games.

THEOREM 1.4 (CONSTANT RELATIVE INAPPROXIMABILITY FOR POLYMATRIX GAMES). *For any constant $\epsilon \in [0,1)$, it is* PPAD*-complete to find a relative ϵ-Nash equilibrium of a bipartite graphical polymatrix game of bounded degree and payoffs in $[-1,1]$. This remains true even if a pure strategy guarantees positive payoff to every player, regardless of the other players' mixed strategies; i.e., it remains true even if the max-min value of every player is positive.*

Another way to describe Theorem 1.4 is this: While it is trivial for every player to guarantee positive payoff to herself using a pure strategy, it is PPAD-hard to find mixed strategies for the players so that every strategy in their support is payoff-optimal to within a factor of $(1-\epsilon)$.

Our Techniques. To show Theorem 1.4, it is natural to try to follow the approach of [DGP06] of reducing the generic PPAD-complete problem END OF THE LINE [DGP06] to a graphical polymatrix game. This was done in [DGP06] by introducing the so called *game-gadgets*: these were small polymatrix games designed to simulate in their Nash equilibria arithmetic and boolean operations and comparisons. Each game gadget consisted of a few players with two strategies each, so that the mixed strategy of each player encoded a real number in $[0,1]$. Then these players were assigned payoffs in such a way that, in any Nash equilibrium of the game, the mixed strategy of the "output player" of the gadget implemented an operation on the mixed strategies of the "input players." Many copies of these gadgets were then combined in a large circuit (with feedback) so that any stable state of this circuit provided a solution to a given END OF THE LINE instance. Unfortunately, for the construction of [DGP06] to go through, the input-output relations of the gadgets need to be accurate to within an exponentially small (in the size of the circuit) additive error; and even the more efficient construction of [CDT06] needs the approximation error to be inverse-polynomial. Alas, if we consider ϵ-Nash equilibria for constant values of ϵ, the errors in the gadgets of [DGP06] become constant, and they accumulate over long paths of the circuit in a destructive manner.

We circumvent this problem with an idea that is rather intuitive, at least in retrospect. The error accumulation is unavoidable if the gates are connected over long paths. But, can we design self-correcting gates if feedback is introduced after each operation? Indeed, our proof of Theorem 1.4 is based on a simple "gap-amplification" kernel (described in Section 3.1,) which reads both the inputs *and the outputs* of a gadget, checks if the output deviates from the prescribed behavior, and amplifies the deviation. The amplified deviation is fed back into the gadget and pushes the output value to the right direction. Using this gadget we can easily obtain an exponentially accurate (although brittle as usual [DGP06]) comparator gadget (see Section 3.3,) and exponentially accurate arithmetic gadgets (see Section 3.2.) Using our new gadgets we can easily finish the proof of Theorem 1.4 (see Section 3.5.)

The Main Challenge. The construction outlined above, while non-obvious, is in the end rather intuitive. The real challenge in establishing Theorem 1.1 lies in reducing the polymatrix games of Theorem 1.4 to two-player games. Those familiar with the hardness reductions for normal form games [GP06; DGP06; CD06; CDT06; EY07] will recognize the challenge. The "generalized matching pennies reduction" of a polymatrix game to a two-player game (more details on this construction shortly) is *not* approximation preserving, in that ϵ-Nash equilibria of the polymatrix game are reduced to $O\left(\frac{\epsilon}{n}\right)$-Nash equilibria of the 2-player game, where n is the number of nodes of the polymatrix game. Hence, even if the required accuracy for hardness in the polymatrix game is a constant, we still need inverse-polynomial accuracy in the resulting two-player game.

In fact, as explained below, any matching pennies-style reduction is doomed to fail, if ϵ-Nash equilibria for constant values of relative[6] approximation are computed in the resulting two-player game. To obtain Theorem 1.1 we provide instead a novel reduction, which in our opinion constitutes significant progress in PPAD-hardness proofs. Our new reduction can obtain all results in [GP06; DGP06; CD06; CDT06; EY07], but is stronger in that it shaves a factor of n off of the relative approximation guarantees. In particular, our reduction is approximation preserving for relative approximations. Given the ubiquity of the matching pennies reduction in previous work, we expect that our new tighter reduction will enable PPAD-hardness proofs in future research.

To explain the challenge, there are two kinds of constraints that a reduction from polymatrix games to two-player games needs to satisfy. The first is enforcing that the mixed strategies of the two players in a Nash equilibrium encode mixed strategies for all the nodes of the polymatrix game simultaneously. The second is ensuring that the equilibrium conditions of the polymatrix game are faithfully encoded in the equilibrium conditions of the two-player game. Unfortunately, when the approximation guarantee is constant, these constraints get coupled in ways that make it hard to enforce both. This is why previous reductions take the approximation in the bimatrix game to scale inverse-polynomially in n; in that regime the above semantics can be decoupled. In our case, the use of constant approximations makes the construction and analysis extremely fragile. In a delicate and technical reduction, we use the structure of the game *outside of the equilibrium* to enforce the first set of constraints, while keeping the equilibrium states clear of these constraints to enforce there the second set of constraints. This is hard to achieve and it is quite surprising that it is at all feasible. Indeed, all details in our construction are very finely chosen.

Overview of the Construction. We explain our approximation preserving reduction from polymatrix to bimatrix games by first providing intuition about the inadequacy of existing techniques. As mentioned above, all previous lower bounds for bimatrix games are based on *generalized matching-pennies* constructions. To reduce a bipartite graphical polymatrix game to a bimatrix game these constructions work as follows. First the nodes of the polymatrix game are colored with two colors so that no two nodes sharing an edge get assigned the same color. Then two "lawyers" are introduced corresponding to the two color classes, and the purpose of each lawyer is to "represent" all the nodes in her color class. This is done by including in the strategy set of each lawyer a block of strategies corresponding to the strategy-set of every node in her color class; and the lawyer payoffs are defined so that, if the lawyers choose strategies corresponding to adjacent nodes of the polymatrix game, the lawyers get payoffs equal to the payoffs from the interaction on that edge. The hope is then that, in any Nash equilibrium of the lawyer-game, the marginal distributions of the lawyer strategies within the different blocks define a Nash equilibrium of the underlying polymatrix game.

But this naive construction may induce the lawyers to focus on the most "lucrative" nodes, playing strategies of certain nodes of the polymatrix game with zero probability. To avoid this, a high-stakes generalized matching pennies game is added onto the lawyers' payoffs, played over blocks of strategies. This game forces the lawyers to randomize (almost) uniformly among their different blocks, and only to decide how to distribute the probability mass of every block to the strategies within the block they look at the payoffs of the underlying polymatrix game. This tie-breaking reflects the Nash equilibrium conditions of the polymatrix game.

[6]In view of the quasi polynomial-time algorithm of [LMM03], constant values of additive approximation are unlikely to be PPAD-hard. So the matching pennies hardness construction—as well as any other construction—should fail in the case of constant additive approximations.

For constant values of relative approximation, this construction fails to work. Because, once the high-stakes game is added to the payoffs of the lawyers, the payoffs coming from the polymatrix game become almost invisible, since their magnitude is tiny compared to the stakes of the high-stakes game (this is discussed in detail in Section 4.1.) To avoid this problem we need a construction that forces the lawyers to randomize uniformly over their different blocks of strategies in a subtle manner that does not overwhelm the payoffs coming from the polymatrix game. We achieve this by including *threats*. These are large punishments that a lawyer can impose to the other lawyer if she does not randomize uniformly over her blocks of strategies. But unlike the high-stakes matching pennies game, these punishments essentially disappear if the other lawyer *does* randomize uniformly over her blocks of strategies; to establish this we have to argue that the additive payoff coming from the threats, which could potentially have huge contribution and overshadow the payoff of the polymatrix game, has very small magnitude at equilibrium, thus making the interesting payoff component (coming from the polymatrix game) visible. This is essential to guarantee that the distribution of probability mass within each block is (almost) only determined by the payoffs of the polymatrix game at an ϵ-Nash equilibrium, even when ϵ is constant. The details of our construction are given in Section 4.2, the analysis of threats is given in Section 4.3, and the proof is completed in Sections 4.4 through 4.7. Threats that are similar in spirit to ours were used in an older NP-hardness proof of Gilboa and Zemel [GZ89]. However, their construction is inadequate here as it could lead to a uniform equilibrium over the threat strategies, which cannot be mapped back to an equilibrium of the underlying polymatrix game. Indeed, it takes a lot of effort to avoid such occurrence of meaningless equilibria.

The Final Twist. As mentioned above, our reduction from graphical polymatrix games to bimatrix games is very fragile. As a result we actually fail to establish that the relative ϵ-Nash equilibria of the lawyer-game correspond to relative ϵ-Nash equilibria of the polymatrix game. Nevertheless, we manage to show that they can be mapped to evaluations of the gadgets used to build up the polymatrix game of Theorem 1.4 that are highly accurate; and this rescues the reduction from END OF THE LINE to bimatrix games.

2. PRELIMINARIES

2.1. Bimatrix Games and Nash Equilibrium

A two-player, or *bimatrix*, game $\mathcal{G}$ has two players, called *row* and *column*, and a finite set of m strategies, $1, \ldots, m$, available to each. If the row player chooses strategy i and the column player strategy j then the row player receives payoff R_{ij} and the column player payoff C_{ij}, where (R, C) is a pair of $m \times m$ matrices, called the *payoff matrices* of the game. The players are allowed to randomize among their strategies by choosing any probability distribution, also called a *mixed strategy*. For notational convenience let $[m] := \{1, \ldots, m\}$ and denote the set of mixed strategies of both players $\Delta^m := \{x \mid x \in \mathbb{R}_+^m, \sum_i x_i = 1\}$. If the row player randomizes according to mixed strategy $x \in \Delta^m$ and the column player according to $y \in \Delta^m$, then the row player receives an expected payoff of $x^{\mathrm{T}} R y$ and the column player an expected payoff of $x^{\mathrm{T}} C y$.

A *Nash equilibrium* of the game is then a pair of mixed strategies (x, y), $x, y \in \Delta^m$, such that $x^{\mathrm{T}} R y \geq x'^{\mathrm{T}} R y$, for all $x' \in \Delta^m$, and $x^{\mathrm{T}} C y \geq x^{\mathrm{T}} C y'$, for all $y' \in \Delta^m$. That is, if the row player randomizes according to x and the column player according to y, then none of the players has an incentive to change her mixed strategy. Equivalently, a pair

(x, y) is a Nash equilibrium iff:[7]

$$\text{for all } i \text{ with } x_i > 0: \quad e_i^{\mathrm{T}} R y \geq e_{i'}^{\mathrm{T}} R y, \text{ for all } i'; \tag{1}$$

$$\text{for all } j \text{ with } y_j > 0: \quad x^{\mathrm{T}} C e_j \geq x^{\mathrm{T}} C e_{j'}, \text{ for all } j'; \tag{2}$$

i.e. every pure strategy that the row player includes in the support of x should give him at least as large expected payoff against y as any other pure strategy would, and similarly for the column player. It was shown in the seminal paper of Nash that every game has a Nash equilibrium [Nas50].

2.2. Additively vs Relatively Approximate Nash equilibrium

It is possible to define two kinds of approximate Nash equilibria, *additive* and *relative*, by relaxing, in the additive or multiplicative sense, the defining conditions of Nash equilibrium. A pair of mixed strategies (x, y) is called an *additive ϵ-approximate Nash equilibrium* if $x^{\mathrm{T}} R y \geq x'^{\mathrm{T}} R y - \epsilon$, for all $x' \in \Delta^m$, and $x^{\mathrm{T}} C y \geq x^{\mathrm{T}} C y' - \epsilon$, for all $y' \in \Delta^m$. That is, no player has more than an additive incentive of ϵ to change her mixed strategy. A related notion of additive approximation arises by relaxing Conditions (1) and (2). A pair of mixed strategies (x, y) is called an *additive ϵ-approximately well-supported Nash equilibrium*, or simply an *additive ϵ-Nash equilibrium*, if

$$\text{for all } i \text{ with } x_i > 0: \ e_i^{\mathrm{T}} R y \geq e_{i'}^{\mathrm{T}} R y - \epsilon, \text{ for all } i'; \tag{3}$$

and similarly for the column player. That is every player allows in the support of her mixed strategy only pure strategies with expected payoff that is within an absolute error of ϵ from the payoff of the best response to the other player's strategy. Clearly, an additive ϵ-Nash equilibrium is also an additive ϵ-approximate Nash equilibrium, but the opposite implication is not always true. Nevertheless, we can show the following:

PROPOSITION 2.1 ([CDT06; DGP09A]). *Given an additive ϵ-approximate Nash equilibrium (x, y) of a game (R, C), we can compute in polynomial time an additive $\sqrt{\epsilon} \cdot (\sqrt{\epsilon} + 1 + 4u_{\max})$-Nash equilibrium of (R, C), where $u_{\max}$ is the maximum absolute value in the payoff matrices R and C.*

Clearly, if (x, y) is an additive ϵ-Nash equilibrium or an additive ϵ-approximate Nash equilibrium of a game (R, C), it remains so if any constant is added to all the entries of R or C. So additive approximate Nash equilibria are shift invariant. However, they are not scale invariant: if all the entries of R or C are scaled by a factor $\alpha > 0$, then (x, y) has approximation accuracy $\epsilon \cdot \alpha$.

The relative notions of approximation are defined similarly via multiplicative relaxations of the equilibrium conditions. We call a pair of mixed strategies (x, y) a *relative ϵ-approximate Nash equilibrium* if $x^{\mathrm{T}} R y \geq x'^{\mathrm{T}} R y - \epsilon \cdot |x'^{\mathrm{T}} R y|$, for all $x' \in \Delta^m$, and $x^{\mathrm{T}} C y \geq x^{\mathrm{T}} C y' - \epsilon \cdot |x^{\mathrm{T}} C y'|$, for all $y' \in \Delta^m$. That is, no player has a relative incentive of more than ϵ to change her mixed strategy. Similarly, a pair of mixed strategies (x, y) is called a *relative ϵ-approximately well-supported Nash equilibrium*, or simply a *relative ϵ-Nash equilibrium*, if

$$\text{for all } i \text{ s.t. } x_i > 0: e_i^{\mathrm{T}} R y \geq e_{i'}^{\mathrm{T}} R y - \epsilon \cdot |e_{i'}^{\mathrm{T}} R y|, \ \forall i'; \tag{4}$$

and similarly for the column player. (4) requires that the relative regret $\left| \frac{e_i^{\mathrm{T}} R y - e_{i'}^{\mathrm{T}} R y}{e_{i'}^{\mathrm{T}} R y} \right|$ experienced by the row player for not replacing strategy i in her support by another

[7]As always, e_i represents the unit vector along dimension i of $\mathbb{R}^m$; it represents the randomized strategy that places probability 1 to pure strategy i (and 0 to all other strategies.)

strategy i' with better payoff is at most ϵ. Notice that both definitions remain meaningful even when R, C have negative entries.

As far as the quality of the approximation goes, values of relative approximation $\epsilon \in [0,1]$ are always meaningful. Values $\epsilon > 1$ are only meaningful when the payoffs of the game are non-positive. Indeed, when the payoffs are non-negative, computing 1-Nash equilibrium is trivial (see Theorem 1.3), so it is uninteresting to look at $\epsilon > 1$; when the payoffs are unrestricted, relatively approximate Nash equilibria with $\epsilon > 1$ have unnatural properties: e.g., a pair of pure strategies that give both players -\$1 are a 2-Nash equilibrium if their best response payoffs are \$1, but not a 2-Nash equilibrium if their best response payoffs are \$0.

Clearly, a relative ϵ-Nash equilibrium is also a relative ϵ-approximate Nash equilibrium, but the opposite implication is not always true. Also, if (x, y) is a relative ϵ-Nash equilibrium or a relative ϵ-approximate Nash equilibrium of a game (R, C), it remains so if all the entries of R or C are scaled by a constant $\alpha > 0$. So relative approximate Nash equilibria are scale invariant. But they are not shift invariant.

And what is the relation between the additive and relative notions of approximation? The following is an easy observation.

Remark 2.2. Let $\mathcal{G} = (R, C)$ be a game whose payoff entries are in $[\ell, u]$, where $\ell, u > 0$. Then an additive ϵ-Nash equilibrium of $\mathcal{G}$ is a relative $\frac{\epsilon}{\ell}$-Nash equilibrium of $\mathcal{G}$, and a relative ϵ-Nash equilibrium of $\mathcal{G}$ is an additive $(\epsilon \cdot u)$-Nash equilibrium of $\mathcal{G}$. The same is true for ϵ-approximate Nash equilibria. Both statements are true if all payoff entries are in $[-u, -\ell]$.

As noted earlier, algorithms for additive approximations are usually compared after taking an affine transformation that brings the payoffs of the game to some unit-length interval. Where this interval lies is irrelevant since the additive approximations are shift invariant. In particular, we can always bring these payoffs to lie in $[-1/2, 1/2]$.[8] In this range, if we compute a relative 2ϵ-Nash equilibrium this would also be an additive ϵ-Nash equilibrium and, similarly, a relative 2ϵ-approximate Nash equilibrium would be an additive ϵ-approximate Nash equilibrium. So the computation of additive ϵ-approximations in normalized games can be polynomial-time reduced to the computation of relative 2ϵ-approximations. But the opposite need not be true. In particular, given our main result (Theorem 1.1), the quasi-polynomial time algorithm for additive approximations of [LMM03], and assuming $\texttt{PPAD} \not\subseteq \texttt{TIME}\left(n^{O(\log n)}\right)$, such a reduction is not possible. On the other hand, if all payoffs of a game lie in some interval $[-u, -\ell]$ or $[\ell, u]$, where $1 \le \frac{u}{\ell} < c$ for some absolute constant $c > 1$ and positive numbers $\ell < u$, then the computation of a relative ϵ-Nash equilibrium can be polynomial-time reduced to the computation of an additive $\frac{\epsilon}{c}$-Nash equilibrium of a normalized game.[9]

2.3. Graphical Polymatrix Games

We define a subclass of *graphical games*, called *graphical polymatrix games*, or just *polymatrix games*. In a graphical game, the players are nodes of a graph $G = (V, E)$, and each node (player) $v \in V$ has her own (finite) strategy set S_v and her own payoff function, which only depends on the strategies of the players in her neighborhood $\mathcal{N}(v)$

[8] If the affine transformation of the payoffs is chosen properly, then going back to the original payoffs results in a loss of a factor of $(u_{\max} - u_{\min})$ in the approximation guarantee, where $u_{\max}$ and $u_{\min}$ are respectively the largest and smallest payoffs of the game before the transformation.

[9] The reduction is this: First multiply all payoffs by $\frac{1}{u}$ so that they lie in $[-1, -\frac{\ell}{u}]$ or in $[\frac{\ell}{u}, 1]$. Compute an additive $\frac{\epsilon}{c}$-Nash equilibrium of the resulting (normalized) game. Remark 2.2 implies that this is also a relative ϵ-Nash equilibrium. As relative approximations are scale invariant, this is also a relative ϵ-Nash equilibrium of the original game.

in G.[10] A graphical game is called a graphical polymatrix game if, in addition, for every $v \in V$ and every pure strategy $s_v \in S_v$, the expected payoff that v gets for playing strategy s_v is a linear function of the mixed strategies of her neighbors $\mathcal{N}(v) \setminus \{v\}$ with rational coefficients; that is, there exist rational numbers $\{\alpha^{v:s_v}_{u:s_u}\}_{u \in \mathcal{N}(v) \setminus \{v\}, s_u \in S_u}$ and $\beta^{v:s_v}$ such that the expected payoff of v for playing pure strategy s_v is

$$\sum_{u \in \mathcal{N}_v \setminus \{v\}, s_u \in S_u} \alpha^{v:s_v}_{u:s_u} p(u : s_u) + \beta^{v:s_v}, \tag{5}$$

where $p(u : s_u)$ denotes the probability with which node u plays pure strategy s_u. Building on the techniques of [DGP06], it was shown in [CDT06] that computing an additive (n^{-c})-Nash equilibrium of a polymatrix game of size n is PPAD-complete, for any absolute constant $c > 0$.

(Approximate) Nash Equilibrium in Polymatrix Games.. Nash equilibrium and its approximate counterparts are defined for polymatrix games similarly to bimatrix games. E.g., a collection of mixed strategies for the nodes of a polymatrix game forms an additive ϵ-Nash equilibrium iff no player includes in the support of her mixed strategy any pure strategy whose expected payoff against the mixed strategies of the other players is not within an additive ϵ from the expected payoff of her best response to their strategies. We omit re-stating all these definitions formally.

2.4. The Approximate Circuit Evaluation Problem

We define the APPROXIMATE CIRCUIT EVALUATION problem, which was shown to be PPAD-complete in [DGP06]. Our definition is based on the notions of a *generalized circuit* and *approximate circuit evaluation*, given in Definitions 2.3 and 2.4 below. These definitions were implicit in [DGP06] and were made more explicit in [CD06].

Definition 2.3 (*Generalized Circuit*). A *generalized circuit* is a collection of nodes and gates, which are interconnected so that nodes only connect to gates and gates only connect to nodes. It is made up of the following types of gates:

— arithmetic gates: the *addition* and *subtraction* gates, denoted by $\mathcal{C}_+$ and $\mathcal{C}_-$ respectively, have two input nodes and one output node (for the gate $\mathcal{C}_-$ one of the input nodes is designated to be the "positive" input); for $\zeta \geq 0$ the *scale by* ζ gate, $\mathcal{C}_{\times\zeta}$, has one input and one output node, and the *set equal to* ζ, $\mathcal{C}_\zeta$, gate has one output node;
— comparison gates: the *comparison* gate, $\mathcal{C}_>$, has two input nodes (one of which is designated to be the "positive" input) and one output node;
— boolean gates: the *OR* gate, $\mathcal{C}_\vee$, has two input nodes and one output node, and the *NOT* gate, $\mathcal{C}_\neg$, has one input and one output node.

Additionally a node can be the output node of at most one gate.

Definition 2.4 (*Approximate Circuit Evaluation*). Given a generalized circuit $\mathcal{C}$ and some constant c, an approximate evaluation of the circuit with accuracy 2^{-cn} is an assignment of $[0, 1]$ values to the nodes of the circuit such that the inputs and outputs of the various gates of the circuit satisfy the following

— $\mathcal{C}_+$: if the input nodes have values x, y and the output node has value z then

$$z = \min\{1\ , x + y\} \pm 2^{-cn};^{11}$$

[10] We include v in $\mathcal{N}(v)$.
[11] Throughout this paper, whenever we write $x = y \pm \delta$, for two reals x, y and a positive real δ, we mean that $x \in [y - \delta, y + \delta]$.

— $\mathcal{C}_-$: if the input nodes have values x, y, where x is the value of the positive input, and the output node has value z then

$$z = \max\{0\ , x - y\} \pm 2^{-cn};$$

— $\mathcal{C}_{\times\zeta}$: if the input node has value x and the output node value z then $z = \min\{1\ , \zeta \cdot x\} \pm 2^{-cn}$;
— $\mathcal{C}_\zeta$: if the output node has value z then $z = \min\{1\ , \zeta\} \pm 2^{-cn}$;
— $\mathcal{C}_>$: if the input nodes have values x, z, where z is the value of the positive input, and the output node has value t then

$$z \geq x + 2^{-cn} \quad \Rightarrow \quad t = 1;$$
$$z \leq x - 2^{-cn} \quad \Rightarrow \quad t = 0;$$

— $\mathcal{C}_\vee$: the values x, y of the input nodes and the value z of the output node satisfy:

$$\text{if } x, y \in \{0, 1\},\ \text{then } z = x \vee y;$$

— $\mathcal{G}_\neg$: the value x of the input player and the value z of the output player satisfy

$$\text{if } x \in \{0, 1\},\ \text{then } z = 1 - x;$$

Definition 2.5. (APPROXIMATE CIRCUIT EVALUATION Problem) Given a generalized circuit $\mathcal{C}$ and a constant c, find an approximate evaluation of the circuit $\mathcal{C}$ with accuracy 2^{-cn}.

THEOREM 2.6 ([DGP06]). APPROXIMATE CIRCUIT EVALUATION *is* PPAD-*complete*

3. HARDNESS OF GRAPHICAL POLYMATRIX GAMES

Our PPAD-hardness proof for graphical polymatrix games is a reduction from the APPROXIMATE CIRCUIT EVALUATION problem, and is based on developing a collection of *game-gadgets*, graphical polymatrix games that simulate arithmetic, boolean, and comparison gates. More precisely, each gadget is a polymatrix game comprising a set of players designated as “input players,” a single player designated as “the output player,” and maybe other players designated as “intermediate players.” These players are involved in pairwise interactions that satisfy the following properties: (i) the payoff of every input player does not depend on the mixed strategy of any other player of the gadget (this is useful for building circuits from game-gadgets where the output player of one gadget is the input player of another gadget); and (ii) in any (approximate) Nash equilibrium of the gadget the mixed strategy of the output player encodes the (approximate) outcome of an arithmetic or boolean operation, or a comparison on the mixed strategies of the input players.

Given such gadgets we can follow the approach of [DGP06], constructing a large graphical polymatrix game that solves, at any Nash equilibrium, an instance of the APPROXIMATE CIRCUIT EVALUATION problem. The main challenge is that, since we are considering constant values of relative approximation, the gadgetry developed in [DGP06] introduces a constant error per operation. And, the construction of [DGP06]—even the more careful one of [CDT06]—cannot accommodate such large error. We go around this problem by introducing new gadgets that are highly accurate even when constant values of relative approximation are considered. Our gadgets build largely on a *gap amplification* gadget given in the next section, which compares the mixed strategy of a player with a linear function of the mixed strategies of two other players, and magnifies their difference if it exceeds a certain threshold. Based on this gadget we construct highly accurate arithmetic and comparison gadgets (Sections 3.2 and 3.3.) And, with a different construction, we also get highly accurate boolean gadgets in Section 3.4. To help the flow of the argument, we only state the properties of our gadgets in the following sections and defer the details of their construction to Appendix A. Moreover, for simplicity, we only describe the "simple versions" of our gadgets here. In Appendix A, we also present "sophisticated versions," in which additionally all participating players have positive max-min values. These latter gadgets are denoted with a superscript of '$+$'.

3.1. Gap Amplification

A key construction that enables our hardness result is the *detector gadget* $\mathcal{G}_{\rm det}$ described in the following lemma. $\mathcal{G}_{\rm det}$ performs a form of gap amplification that enables building highly accurate gadgets, in the additive sense, from very weak multiplicative guarantees. The gadget compares the mixed strategy of a player with a linear function of the mixed strategies of two other players. If the absolute difference exceeds an exponentially small threshold, the gadget outputs a deterministic 1 or 0 reflecting whether the difference is positive or negative. The gadget is presented in the following lemma and its application to construct highly accurate gadgets is given in the next sections.

LEMMA 3.1 (DETECTOR GADGET). *Fix $\epsilon \in [0,1)$, $\alpha,\beta,\gamma \in [-1,1]$, and $c \in \mathbb{N}$. There exists $n_0 \in \mathbb{N}$, such that for all $n > n_0$, there exists a graphical polymatrix game $\mathcal{G}_{\rm det}$ with three input players x, y and z, one intermediate player w, and one output player t, and two strategies per player, 0 and 1, such that in any relative ϵ-Nash equilibrium of $\mathcal{G}_{\rm det}$, the mixed strategies of the players satisfy*

$$p(z:1) - [\alpha p(x:1) + \beta p(y:1) + \gamma] \geq 2^{-cn} \qquad \Rightarrow \quad p(t:1) = 1;$$
$$p(z:1) - [\alpha p(x:1) + \beta p(y:1) + \gamma] \leq -2^{-cn} \qquad \Rightarrow \quad p(t:1) = 0.$$

3.2. Arithmetic Operators

We use our detector gadget $\mathcal{G}_{\rm det}$ to construct highly accurate—in the additive sense—arithmetic operators, such as plus, minus, multiplication by a constant, and setting a value. We use the gadget $\mathcal{G}_{\rm det}$ to compare the inputs and the output of the arithmetic operator, magnify any deviation, and correct—with the appropriate feedback—the output, if it fails to comply with the right value. In this way, we use our gap amplification gadget to construct highly accurate arithmetic operators, despite the weak guarantees that a relative ϵ-Nash equilibrium provides, for constant ϵ's. We start with a generic affine operator gadget $\mathcal{G}_{\rm lin}$.

LEMMA 3.2 (AFFINE OPERATOR). *Fix $\epsilon \in [0,1)$, $\alpha,\beta,\gamma \in [-1,1]$, and $c \in \mathbb{N}$. There exists $n_0 \in \mathbb{N}$, such that for all $n > n_0$, there is a graphical polymatrix game $\mathcal{G}_{\rm lin}$ with a bipartite graph, two input players x and y, and one output player z, such that in any relative ϵ-Nash equilibrium*

$$p(z:1) \geq \quad \max\{0,\ \min\{1, \alpha p(x:1) + \beta p(y:1) + \gamma\}\} - 2^{-cn};$$
$$p(z:1) \leq \quad \min\{1,\ \max\{0, \alpha p(x:1) + \beta p(y:1) + \gamma\}\} + 2^{-cn}.$$

Using $\mathcal{G}_{\rm lin}$ we obtain highly accurate arithmetic operators.

LEMMA 3.3 (ARITHMETIC GADGETS). *Fix $\epsilon \geq 0$, $\zeta \geq 0$, and $c \in \mathbb{N}$. There exists $n_0 \in \mathbb{N}$, such that for all $n > n_0$, there exist graphical polymatrix games $\mathcal{G}_+, \mathcal{G}_-, \mathcal{G}_{\times\zeta}, \mathcal{G}_\zeta$ with bipartite graphs, two input players x and y, and one output player z, such that in any relative ϵ-Nash equilibrium*

— *the game $\mathcal{G}_+$ satisfies:* $p(z:1) = \min\{1\ , p(x:1) + p(y:1)\} \pm 2^{-cn}$;
— *the game $\mathcal{G}_-$ satisfies:* $p(z:1) = \max\{0\ , p(x:1) - p(y:1)\} \pm 2^{-cn}$;

— *the game* $\mathcal{G}_{\times\zeta}$ *satisfies:* $p(z:1) = \min\{1\ ,\zeta \cdot p(x:1)\} \pm 2^{-cn}$;
— *the game* $\mathcal{G}_{\zeta}$ *satisfies:* $p(z:1) = \min\{1\ ,\zeta\} \pm 2^{-cn}$.

3.3. Brittle Comparator

Also from $\mathcal{G}_{\det}$ it is quite straightforward to construct a (*brittle* [DGP06]) comparator gadget as follows. The term "brittle" refers to the fact that the output of the comparator is unrestricted if its inputs are closer than 2^{-cn}.

LEMMA 3.4 (COMPARATOR GADGET). *Fix* $\epsilon \in [0, 1)$*, and* $c \in \mathbb{N}$*. There exist* $n_0 \in \mathbb{N}$*, such that for all* $n > n_0$*, there exists a graphical polymatrix game* $\mathcal{G}_>$ *with bipartite interaction graph, two input players* x *and* z*, and one output player* t*, such that in any relative* ϵ*-Nash equilibrium of* $\mathcal{G}_>$

$$p(z:1) - p(x:1) \geq 2^{-cn} \quad \Rightarrow \quad p(t:1) = 1;$$
$$p(z:1) - p(x:1) \leq -2^{-cn} \quad \Rightarrow \quad p(t:1) = 0.$$

3.4. Boolean Operators

In a relatively straightforward manner that does not require our gap amplification gadget we also construct boolean operators. We only need to describe a game for `or` and `not`. Using these we can always obtain `and`.

LEMMA 3.5 (BOOLEAN OPERATORS). *Fix* $\epsilon \in [0, 1)$*. There exist graphical polymatrix games* $\mathcal{G}_\vee, \mathcal{G}_\neg$ *with bipartite graphs, two input players* x *and* y*, and one output player* z*, such that in any relative* ϵ*-Nash equilibrium*

— *if* $p(x:1), p(y:1) \in \{0,1\}$*, the game* $\mathcal{G}_\vee$ *satisfies* $p(z:1) = p(x:1) \vee p(y:1)$;
— *if* $p(x:1) \in \{0,1\}$*, the game* $\mathcal{G}_\neg$ *satisfies* $p(z:1) = 1 - p(x:1)$.

3.5. Proof of Theorem 1.4

We reduce an instance of the PPAD-complete problem APPROXIMATE CIRCUIT EVALUATION (see Section 2.4) to a graphical polymatrix game, by replacing every gate of the given circuit with the corresponding gadget. The nature of our gadgets guarantees that a relative ϵ-Nash equilibrium of the resulting polymatrix game corresponds to a highly accurate evaluation of the circuit, completing the hardness proof. The inclusion in PPAD follows from the fact that the exact ($\epsilon = 0$) Nash equilibrium problem lies in PPAD [EY07].

PROOF OF THEOREM *1.4:* From [EY07], it follows that computing an exact Nash equilibrium of a graphical polymatrix game is in PPAD. Since exact Nash equilibria are also relative ϵ-Nash equilibria, inclusion in PPAD follows immediately.

So we only need to justify the PPAD-hardness of our problem. To do this, we reduce from the APPROXIMATE CIRCUIT EVALUATION problem, i.e. the following: Given a circuit consisting of the gates *plus*, *minus*, *scale by a constant*, *set equal to a constant*, *compare*, *or*, and *not* (see Definition 2.3,) find values for the nodes of the circuit satisfying the input-output relations of the gates to within an additive error of 2^{-cn}. (See Definition 2.4 for the precise input-output relations that need to be satisfied.) Notice that these relations are in direct analogy to the input-output relations of our gadgets from Lemmas 3.3, 3.4 and 3.5. So, given any circuit, it is easy to set up, using the gadgets $\mathcal{G}_+$, $\mathcal{G}_-$, $\mathcal{G}_{\times\zeta}$, $\mathcal{G}_\zeta$, $\mathcal{G}_>$, $\mathcal{G}_\vee$, $\mathcal{G}_\neg$ of Lemmas 3.3, 3.4 and 3.5, a bipartite graphical polymatrix game $\mathcal{GG}$ with the same functionality as the circuit: every node of the circuit corresponds to a player, the players participate in arithmetic, comparison and logical gadgets depending on the types of gates with which the corresponding nodes of the circuit are connected, and given any relative ϵ-Nash equilibrium of the graphical game we can obtain an approximate circuit evaluation by interpreting the probabilities with which every player plays strategy 1 as the value of the corresponding node of the circuit. This concludes the PPAD-hardness proof. If we also want to enforce that every node of our graphical game has a positive max-min value we can use in our construction the "sophisticated versions" $\mathcal{G}_{+}^{+}$, $\mathcal{G}_{-}^{+}$, $\mathcal{G}_{\times\zeta}^{+}$, $\mathcal{G}_{\zeta}^{+}$, $\mathcal{G}_{>}^{+}$, $\mathcal{G}_{\vee}^{+}$, $\mathcal{G}_{\neg}^{+}$ of our gadgets given in Appendix A. □

For future reference, we denote by $\mathcal{GG}$ the graphical polymatrix game constructed in the proof of Theorem 1.4 using the "simple gadgets" and by $\mathcal{GG}^+$ the graphical polymatrix game constructed using the "sophisticated gadgets."

4. HARDNESS OF TWO-PLAYER GAMES

4.1. The Main Technical Challenge

To show Theorem 1.1, we need to encode the bipartite graphical polymatrix game $\mathcal{GG}$, built using the gadgets $\mathcal{G}_{>}$, $\mathcal{G}_{+}$, $\mathcal{G}_{-}$, $\mathcal{G}_{\times\zeta}$, $\mathcal{G}_{\zeta}$, $\mathcal{G}_{\vee}$, $\mathcal{G}_{\neg}$ in the proof of Theorem 1.4, into a bimatrix game, whose relative ϵ-Nash equilibria correspond to approximate evaluations of the circuit encoded by $\mathcal{GG}$. A construction similar to the one we are after, but for *additive* ϵ-Nash equilibria, was described in [DGP06; CD06]. But, that construction is not helpful in our setting, since it cannot accommodate constant values of ϵ as we will discuss shortly. Before that, let us get our notation straight.

Suppose that the bipartite graphical polymatrix game $\mathcal{GG}$ has graph $G = (V_L \cup V_R, E)$, where V_L, V_R are respectively the "left" and "right" sides of the graph, and payoffs of the linear form described in (5). Without loss of generality, let us also assume that both sides of the graph have n players, $|V_L| = |V_R| = n$; if not, we can add isolated players to make up any shortfall. To reduce $\mathcal{GG}$ into a bimatrix game, it is natural to "assign" the players on the two sides of the graph to the two players of the bimatrix game. To avoid confusion, in the remaining of this paper we are going to refer to the players of the graphical game as "vertices" or "nodes" and reserve the word "player" for the players of the bimatrix game. Also, for notational convenience, let us label the row and column players of the bimatrix game by 0 and 1 respectively, and define $\rho : V_L \cup V_R \to \{0, 1\}$ to be the function mapping vertices to players as follows: $\rho(v) = 0$, if $v \in V_L$, and $\rho(v) = 1$, if $v \in V_R$.

Now, here is a straightforward way to define the reduction: For every vertex v, we can include in the strategy set of player $\rho(v)$ two strategies denoted by $(v : 0)$ and $(v : 1)$, where strategy $(v : s)$ has the intended meaning "vertex v plays strategy s," for $s = 0, 1$. (Recall that the players of $\mathcal{G}G$ have two strategies, denoted 0 and 1.) We call the pair of strategies $(v : 0)$ and $(v : 1)$ the *block of strategies of player $\rho(v)$ corresponding to vertex v*. We can then define the payoffs of the bimatrix game in terms of the payoffs of the polymatrix game as follows (using the notation of Eq. (5)):

$$U_{\rho(v)}((v:s)\ ,\ (v':s')) := \begin{cases} \alpha^{v:s}_{v':s'} + \frac{1}{n}\beta^{v:s}, & \text{if } (v,v') \in E; \\ \frac{1}{n}\beta^{v:s}, & \text{if } (v,v') \notin E. \end{cases} \tag{6}$$

In other words, if the players $\rho(v)$ and $1 - \rho(v)$ of the bimatrix game play strategies $(v : s)$ and $(v' : s')$ respectively, then they are given payoffs equal to the payoffs $\alpha^{v:s}_{v':s'}$ and $\alpha^{v':s'}_{v:s}$ that the nodes v and v' would have got on edge (v, v') of the polymatrix game (if such edge exists), if they chose strategies s and s' respectively. The additive payoff term $\beta^{v:s}$ is scaled by a factor of $\frac{1}{n}$ for technical reasons, and the amount of $\frac{1}{n}\beta^{v:s}$ is always given to player $\rho(v)$, if she plays strategy $(v : s)$.

Observe that, if we could (somehow) guarantee that—in any Nash equilibrium of the bimatrix game thus defined—the players assign the same probability mass on their different blocks of strategies, then the marginal distributions within each block would jointly define a Nash equilibrium of the graphical game. Indeed, given our definition of the payoff function (6), in order to distribute the probability mass of the block corresponding to v to the strategies $(v : 0)$ and $(v : 1)$ inside that block, player $\rho(v)$ would have to respect the Nash equilibrium conditions of node v. Such rationale goes through as long as the players randomize uniformly, or even close to uniformly, among their different blocks of strategies. If they don't, then all bets are off ...

To make sure that the players randomize uniformly over blocks of strategies, the construction of [GP06; DGP06], which was used in all subsequent papers [CD05; DP05; CD06; CDT06; EY07], has the players of the bimatrix game play, on the side, a high-stakes matching pennies game over blocks of strategies. This forces them to randomize almost uniformly among their blocks and makes the above argument go through. To be more precise, let us define two arbitrary permutations $\pi_L : V_L \to [n]$ and $\pi_R : V_R \to [n]$, and define $\pi : V_L \cup V_R \to [n]$ as $\pi(v) = \pi_L(v)$, if $v \in V_L$, and $\pi(v) = \pi_R(v)$, if $v \in V_R$. Given π, the matching pennies game is incorporated in the construction by assigning the following payoffs to the players

$$\widetilde{U}_{\rho(v)}((v : s)\ ,\ (v' : s')) := U_{\rho(v)}((v : s), (v' : s')) + (-1)^{\rho(v)} \cdot M \cdot \mathbb{1}_{\pi(v)=\pi(v')}, \tag{7}$$

where M is chosen to be *much larger* (a polynomial in n factor larger is sufficient) than the payoffs of the graphical game. Observe that, if we ignored the payoffs coming from the graphical game in Eq. (7), the resulting game would be a generalized *matching pennies* game over blocks of strategies; and it is not hard to see that, in any Nash equilibrium of this game, both players would assign probability $1/n$ to each block. Given that M is chosen much larger than the payoffs of the graphical game, even if we do not ignore these payoffs, every block still receives roughly $1/n$ probability mass at a Nash equilibrium. And, if ϵ is sufficiently small (inverse-polynomial in n,) the same is true of ϵ-Nash equilibria. Moreover, it can be argued [GP06; DGP06] that, in every ϵ-Nash equilibrium of the bimatrix game, the marginal distributions within each block of strategies comprise jointly an ϵ'-Nash equilibrium of the graphical game, where ϵ and ϵ' are polynomially related.

The above construction works well as long as ϵ is inverse-polynomial in n. But, it seems that an inverse-polynomial in n value of ϵ is truly needed. If ϵ is constant, then additive ϵ-Nash equilibria do not guarantee that the players will randomize uniformly over their different blocks of strategies, or even that they will assign non-zero probability mass to each block [LMM03; DP09b]. Hence, we cannot argue anymore that the marginal distributions over blocks comprise an approximate equilibrium of the graphical game (as these distributions may not even be well-defined). On the other hand, if we consider relative ϵ-Nash equilibria for constant values of ϵ, then the different strategies within a block *always* give payoffs that are within a relative error ϵ from each other, for trivial reasons, since their payoff is overwhelmed by the high-stakes game. So the marginal distributions over blocks cannot be informative about the Nash equilibria of the underlying graphical game. And, if we try to decrease the value of M to make the payoffs of the graphical game visible, we cannot guarantee anymore that the players of the bimatrix game randomize uniformly over their different blocks, and the construction still fails.

To accommodate constant values of ϵ, we need a different approach. Our high-level idea is the following. We include in the definition of the game *threats*. These are large punishments that one player can impose to the other player if she does not randomize uniformly over her blocks of strategies. But, unlike the high-stakes matching pennies game of [DGP06] and subsequent works, these punishments essentially disappear if the player *does* randomize uniformly over her blocks of strategies; and this is necessary to guarantee that in a relative ϵ-Nash equilibrium of the bimatrix game the allocation of probability mass within each block is (almost) only determined by the payoffs of the graphical game, even when ϵ is a constant.

The details of our construction are given in Section 4.2, and in Section 4.3 we analyze the effect of the threats on the equilibria of the game. In particular, in Lemmas 4.1 and 4.3 we show that the threats force the players of the bimatrix game to randomize (exponentially close to) uniformly over their blocks of strategies. Unfortunately, to guarantee this we need to choose the magnitude of the punishment-payoffs to be exponentially larger than the magnitude of the payoffs of the underlying graphical game. Hence, the punishment-payoffs could in principle overshadow the graphical-game payoffs, turn the payoffs of the two players negative at equilibrium, and prevent any correspondence between the equilibria of the bimatrix and the polymatrix game. Yet, we show in Lemma 4.4 that in a relative ϵ-Nash equilibrium of the bimatrix game, the threat strategies are played with small-enough probability that the punishment-payoffs are of the same order as the payoffs from the underlying graphical game. This opens up the road to establishing the correspondence between the approximate equilibria of the bimatrix and polymatrix games. However, we are *unable* to establish this correspondence, i.e. we fail to show that the marginal distributions used by the players of the bimatrix game in their different blocks of strategies constitute an approximate Nash equilibrium of the underlying graphical game.[12] But, we *can* show (see Lemma 4.5) that these marginal distributions satisfy a weaker condition, namely they jointly define a highly accurate (in the additive sense) evaluation of the circuit encoded by the graphical game. This is enough to establish our PPAD-completeness result (completed in Section 4.7.)

4.2. Our Construction

We do the following modifications to the game defined in (6):

— For every vertex v, we introduce a third strategy to the block of strategies of player $\rho(v)$ corresponding to v; we call that strategy v^* and we are going to use it to make sure that both players of the bimatrix game have positive payoff in every relative ϵ-Nash equilibrium.
— For every vertex v, we also introduce a new strategy bad_v in the set of strategies of player $1-\rho(v)$. The strategies $\{\mathrm{bad}_v\}_{v\in V_L}$ are going to be used as threats to make sure that player 0 randomizes uniformly among her blocks of strategies. Similarly, we will use the strategies $\{\mathrm{bad}_v\}_{v\in V_R}$ in order to force player 1 to randomize uniformly among her blocks of strategies.
— The payoff functions $\widehat{U}_0(\cdot;\cdot)$ and $\widehat{U}_1(\cdot;\cdot)$ of respectively players 0 and 1 are defined in detail in Figure 4 of the appendix, for some H, U and d to be determined shortly. The reader can study the definition of the functions in detail, however it is easier to think of our game in terms of the expected payoffs that the players receive for playing

[12]Even though the un-normalized marginal distributions actually do satisfy the approximate equilibrium conditions of the polymatrix game, the fragility of relative approximations prevents us from showing that the normalized marginal distributions also do (despite the normalization factors being inverse-exponentially close.)

different pure strategies as follows

$$\mathcal{E}\left(\widehat{U}_{p:v^*}\right) = -U \cdot p_{\mathrm{bad}_v} + 2^{-dn}; \tag{8}$$

$$\mathcal{E}\left(\widehat{U}_{p:(v:s)}\right) = -U \cdot p_{\mathrm{bad}_v} + \sum_{(v,v')\in E} \sum_{s'=0,1} \alpha^{v:s}_{v':s'} \cdot p_{v':s'} + \frac{1}{n}\beta^{v:s}; \tag{9}$$

$$\mathcal{E}\left(\widehat{U}_{p:\mathrm{bad}_v}\right) = H \cdot (p_v - 1/n). \tag{10}$$

In the above, we denote by $\mathcal{E}\left(\widehat{U}_{p:v^*}\right)$, $\mathcal{E}\left(\widehat{U}_{p:(v:s)}\right)$ and $\mathcal{E}\left(\widehat{U}_{p:\mathrm{bad}_v}\right)$ the expected payoff that player p receives for playing strategies v^*, $(v : s)$ and bad_v respectively (where it is assumed that p is allowed to play these strategies, i.e. $p = \rho(v)$ for the first two to be meaningful, and $p = 1 - \rho(v)$ for the third.) We also use $p_{v:0}$, $p_{v:1}$ and p_{v^*} to denote the probability with which player $\rho(v)$ plays strategies $(v : 0)$, $(v : 1)$ and (v^*), and p_{bad_v} to denote the probability with which player $1-\rho(v)$ plays strategy bad_v. Finally, we let $p_v = p_{v:0} + p_{v:1} + p_{v^*}$.

Choice of Constants: Since we are considering relative approximate Nash equilibria we can assume without loss of generality that the payoffs of all players in the graphical game $\mathcal{GG}$ are at most 1 (otherwise we can scale down all the utilities of $\mathcal{GG}$ to make this happen.) Let us then choose $H := 2^{hn}$, $U := 2^{un}$, d, and $\delta := 2^{-dn}$, where $h, u, d \in \mathbb{N}$, $h > u > d > c' > c$, and c, c' are the constants chosen in the definition of the gadgets used in the construction of $\mathcal{GG}$ (as specified in the proofs of Lemmas 3.3, 3.4 and 3.5.) Let us also choose a sufficiently large n_0, such that for all $n > n_0$ the inequalities of Figure 5 of the appendix are satisfied. These inequalities are needed for technical purposes in the analysis of the bimatrix game.

4.3. The Effect of the Threats

We show that the threats force the players to randomize uniformly over the blocks of strategies corresponding to the different nodes of $\mathcal{GG}$, in every relative ϵ-Nash equilibrium. One direction is intuitive: if player $\rho(v)$ assigns more than $1/n$ probability to block v, then player $1-\rho(v)$ receives a lot of incentive to play strategy bad_v; this incurs a negative loss in expected payoff for all strategies of block v, making $\rho(v)$ loose her interest in this block. The opposite direction is less intuitive and more fragile, since there is no explicit threat (in the definition of the payoff functions) for under-using a block of strategies. The argument has to look at the global implications that under-using a block of strategies has and requires arguing that in every relative ϵ-Nash equilibrium the payoffs of both players are positive (Lemma 4.2); this will also become handy later. Observe that Lemma 4.1 is not sufficient to imply Lemma 4.3 directly, since besides their blocks of strategies corresponding to the nodes of $\mathcal{GG}$ the players of the bimatrix game also have strategies of the type bad_v, which are not contained in these blocks. The proofs of the following lemmas can be found in the appendix.

LEMMA 4.1. *In any relative ϵ-Nash equilibrium with $\epsilon \in [0, 1)$, for all $v \in V_L \cup V_R$, $p_v \leq \frac{1}{n} + \delta$.*

LEMMA 4.2. *In any relative ϵ-Nash equilibrium with $\epsilon \in [0, 1)$, both players of the game get expected payoff at least $(1 - \epsilon)2^{-dn}$ from every strategy in their support.*

LEMMA 4.3. *In any relative ϵ-Nash equilibrium with $\epsilon \in [0, 1)$, for all $v \in V_L \cup V_R$, $p_v \geq \frac{1}{n} - 2n\delta$.*

4.4. Mapping Equilibria to Approximate Gadget Evaluations

Almost There. Let us consider a relative ϵ-Nash equilibrium of our bimatrix game $\mathcal{G}$, where $\{p_{v:0}, p_{v:1}, p_{v*}\}_{v \in V_L \cup V_R}$ are the probabilities that this equilibrium assigns to the blocks corresponding to the different nodes of $\mathcal{G}G$. For every v, we define

$$U'_{v*} := 2^{-dn}; \quad \text{and}$$
$$U'_{(v:s)} := \sum_{(v,v') \in E} \sum_{s'=0,1} \alpha^{v:s}_{v':s'} \cdot p_{v':s'} + \frac{1}{n}\beta^{v:s}, \text{ for } s = 0, 1.$$
$$\text{so that } \mathcal{E}\left(\widehat{U}_{p:v^*}\right) = -U \cdot p_{\mathrm{bad}_v} + U'_{v*}; and$$
$$\mathcal{E}\left(\widehat{U}_{p:(v:s)}\right) = -U \cdot p_{\mathrm{bad}_v} + U'_{(v:s)}, \text{ for } s = 0, 1.$$

In Appendix D.2 we show the following.

LEMMA 4.4 (UN-NORMALIZED NASH EQUILIBRIUM CONDITIONS FOR $\mathcal{G}G$). *Fix an arbitrary v. Let*

$$\sigma_{\max} \in \arg \max_{\sigma \in \{v*, (v:0), (v:1)\}} \{U'_\sigma\}.$$

(In particular, observe that $U'_{\sigma_{\max}} > 0$.) Then, in any relative ϵ-Nash equilibrium with $\epsilon \in [0, 1)$, for all $\sigma \in \{v*, (v:0), (v:1)\}$,

$$U'_\sigma < (1-\epsilon)U'_{\sigma_{\max}} \quad \Rightarrow \quad p_\sigma = 0. \tag{11}$$

Notice the subtlety in Condition (11). If we replace U'_σ and $U'_{\sigma_{\max}}$ with $\mathcal{E}\left(\widehat{U}_{p:\sigma}\right)$ and $\mathcal{E}\left(\widehat{U}_{p:\sigma_{\max}}\right)$, then it is automatically true, since it corresponds to the relative ϵ-Nash equilibrium conditions of the game $\mathcal{G}$. But, to remove the term $-U \cdot p_{\mathrm{bad}_v}$ from $\mathcal{E}\left(\widehat{U}_{p:\sigma}\right)$ and $\mathcal{E}\left(\widehat{U}_{p:\sigma_{\max}}\right)$ and maintain Condition (11), we need to make sure that this term is not too large so that it overshadows the true relative magnitude of the underlying values of the U''s. And Lemma 4.2, comes to our rescue: since the payoff of every player is positive at equilibrium, at least one of the U''s has absolute value larger than $U \cdot p_{\mathrm{bad}_v}$; and this is enough to save the argument. Indeed, the property of our construction that the threats approximately disappear at equilibrium is really important here.

4.5. The Trouble

Given Lemma 4.4, the *un-normalized* probabilities $\{p_{v:0}, p_{v:1}, p_{v*}\}_{v \in V_L \cup V_R}$ satisfy the relative ϵ-Nash equilibrium conditions of the graphical game $\mathcal{G}\mathcal{G}$ (in fact, of the game $\mathcal{G}\mathcal{G}^+$ with three strategies $0, 1, *$ per player—see the proof of Theorem 1.4.) Hence, it is natural to try to normalize these probabilities, and argue that their normalized counterparts also satisfy the relative ϵ-Nash equilibrium conditions of $\mathcal{G}\mathcal{G}^+$. After all, given Lemmas 4.1 and 4.3, the normalization would essentially result in multiplying all the U''s by n. It turns out that the (exponentially small) variation of $\pm\delta$ in the different p_v's and the overall fragility of relative approximations makes this approach problematic. Indeed, we fail to establish that after normalization the equilibrium conditions of $\mathcal{G}\mathcal{G}^+$ are satisfied.

4.6. The Final Maneuver

Rather than worrying about the ϵ-Nash equilibrium conditions of $\mathcal{G}\mathcal{G}$ or $\mathcal{G}\mathcal{G}^+$, we argue instead that we can obtain a highly accurate evaluation of the circuit encoded by these games (a weaker condition). We consider first the following transformation, which merges the strategies $(v:0)$ and $v*$:

$$\hat{p}(v:1) := \frac{p_{v:1}}{p_v}; \quad \hat{p}(v:0) := \frac{p_{v:0}+p_{v*}}{p_v}. \tag{12}$$

We argue next that the normalized values $\{\hat{p}(v:1)\}_v$ correspond to a highly accurate evaluation of the circuit encoded by the game $\mathcal{GG}$. We do this by studying the input-output conditions of each of the gadgets used in our construction of $\mathcal{GG}$. For example, for all appearances of the gadget $\mathcal{G}_{\mathrm{det}}$ inside $\mathcal{GG}$ we show the following.

LEMMA 4.5. *Suppose that $x,y,z,w,t \in V_L \cup V_R$, so that x, y and z are inputs to some game $\mathcal{G}_{\mathrm{det}}$ with $\alpha,\beta,\gamma \in [-1,1]$, w is the intermediate node of $\mathcal{G}_{\mathrm{det}}$, and t is the output node. Then the values $\hat{p}$ obtained from a relative ϵ-Nash equilibrium of the bimatrix game using Eq.* (12) *satisfy*

$$\hat{p}(z:1) - [\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma] \geq 2^{-cn} \quad \Rightarrow \quad \hat{p}(t:1) = 1; \tag{13}$$

$$\hat{p}(z:1) - [\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma] \leq -2^{-cn} \quad \Rightarrow \quad \hat{p}(t:1) = 0. \tag{14}$$

Given Lemma 4.5, we immediately obtain that the output of all comparator gadgets is highly accurate.

COROLLARY 4.6. *Suppose $x,z,w,t \in V_L \cup V_R$, so that x, z are inputs to a comparator game $\mathcal{G}_{>}$, w is the intermediate node, and t the output node. Then*

$$\hat{p}(z:1) \geq \hat{p}(x:1) + 2^{-cn} \quad \Rightarrow \quad \hat{p}(t:1) = 1;$$

$$\hat{p}(z:1) \leq \hat{p}(x:1) - 2^{-cn} \quad \Rightarrow \quad \hat{p}(t:1) = 0.$$

Next, we study the gadget $\mathcal{G}_{\mathrm{lin}}$.

LEMMA 4.7. *Suppose $x,y,z,w,w',t \in V_L \cup V_R$, so that x, y are inputs to the game $\mathcal{G}_{\mathrm{lin}}$ with parameters α, β and γ, z is the output player, and w, w', t' are the intermediate nodes (as in Figure 2). Then*

$$\hat{p}(z:1) \geq \max\{0, \min\{1\,, \alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma\}\} - 2^{-cn}; \tag{15}$$

$$\hat{p}(z:1) \leq \min\{1, \max\{0,\ \alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma\}\} + 2^{-cn}. \tag{16}$$

Given Lemma 4.7, we obtain that all arithmetic gadgets are also highly accurate.

COROLLARY 4.8. *Suppose $x,y,z \in V_L \cup V_R$, where x, y are the inputs and z is the output of an arithmetic game. Then*

— *if the game is $\mathcal{G}_{+}$, then $\hat{p}(z:1) = \min\{1\,, \hat{p}(x:1) + \hat{p}(y:1)\} \pm 2^{-cn}$;*
— *if the game is $\mathcal{G}_{-}$, then $\hat{p}(z:1) = \max\{0\,, \hat{p}(x:1) - \hat{p}(y:1)\} \pm 2^{-cn}$;*
— *if the game is $\mathcal{G}_{\times\zeta}$, then $\hat{p}(z:1) = \min\{1\,, \zeta \cdot \hat{p}(x:1)\} \pm 2^{-cn}$;*
— *if the game is $\mathcal{G}_{\zeta}$, then $\hat{p}(z:1) = \min\{1\,, \zeta\} \pm 2^{-cn}$.*

Finally, we analyze the boolean operators.

LEMMA 4.9. *Suppose $x,y,w,z \in V_L \cup V_R$, where x, y are the inputs, w is the intermediate node, and z is the output of a boolean game $\mathcal{G}_{\vee}$ or $\mathcal{G}_{\neg}$ (as in Figure 3). Then*

— *if $\hat{p}(x:1), \hat{p}(y:1) \in \{0,1\}$, the game $\mathcal{G}_{\vee}$ satisfies $\hat{p}(z:1) = \hat{p}(x:1) \vee \hat{p}(y:1)$;*
— *if $\hat{p}(x:1) \in \{0,1\}$, the game $\mathcal{G}_{\neg}$ satisfies $\hat{p}(z:1) = 1 - \hat{p}(x:1)$;*

It follows from the above that the values $\{\hat{p}(v:1)\}_{v \in V_L \cup V_R}$ correspond to an approximate evaluation of the circuit encoded by the graphical game $\mathcal{GG}$. This is sufficient to conclude the proof of the PPAD-hardness part of Theorem 1.1, since finding such an evaluation is PPAD-hard (see Section 2.4.) On the other hand, finding an exact Nash equilibrium of a bimatrix game is in PPAD [Pap94], hence finding a relative ϵ-Nash equilibrium is also in PPAD. We provide a detailed proof of Theorem 1.1 in the next section.

4.7. Completing the Proof of Theorems 1.1 and 1.2.

PROOF OF THEOREM *1.1:* It follows from [Pap94] that computing an exact Nash equilibrium of a bimatrix game is in PPAD. Since exact Nash equilibria are also relative ϵ-Nash equilibria, inclusion in PPAD follows immediately.

Hence, we only need to argue the PPAD-hardness of the problem. Given a pair $(\mathcal{C}, c)$, where $\mathcal{C}$ is a generalized circuit (see Definition 2.3) and c a positive constant (such pair constitutes an instance of the APPROXIMATE CIRCUIT EVALUATION problem defined in Section 2.4,) we construct a bipartite graphical polymatrix game $\mathcal{GG}$ using the reduction in the proof of Theorem 1.4. The game $\mathcal{GG}$ has graph $G = (V_L \cup V_R, E)$, where V_L and V_R are the left and right sides of the bipartition, and consists of the gadgets $\mathcal{G}_+$, $\mathcal{G}_-$, $\mathcal{G}_{\times\zeta}$, $\mathcal{G}_\zeta$, $\mathcal{G}_>$, $\mathcal{G}_\vee$, $\mathcal{G}_\neg$. Now, using the reduction outlined in Section 4.2, we can construct a bimatrix game $\mathcal{G}$ with the following property: Given any relative ϵ-Nash equilibrium of the game $\mathcal{G}$, we can compute (using Equation (12)) values $\{\hat{p}(v:1)\}_{v \in V_L \cup V_R}$ for the nodes of the graphical game, corresponding to approximate evaluations of the gadgets $\mathcal{G}_+$, $\mathcal{G}_-$, $\mathcal{G}_{\times\zeta}$, $\mathcal{G}_\zeta$, $\mathcal{G}_>$, $\mathcal{G}_\vee$, $\mathcal{G}_\neg$ (as specified by Corollaries 4.6 and 4.8 and Lemma 4.9.) These values comprise then an approximate evaluation of the circuit $\mathcal{C}$. Since the APPROXIMATE CIRCUIT EVALUATION problem is PPAD-hard it follows that finding a relative ϵ-Nash equilibrium of bimatrix game $\mathcal{G}$ is also PPAD-hard.

To complete the proof of Theorem 1.1 we note that, by virtue of Lemma 4.2, all players have positive payoffs in every relative ϵ-Nash equilibrium of $\mathcal{G}$. □

PROOF OF THEOREM *1.2:* The proof is an immediate consequence of Lemma 4.3 and our choice of the parameter δ. □

5. CONCLUSIONS AND OPEN PROBLEMS

In this paper we establish the first constant inapproximability results for Nash equilibrium following a long line of research on lower [DGP06; CD06; CDT06; HK09; DP09b] and upper bounds [LMM03; KPS06; DMP06; DMP07; KS07; FNS07; BBM07; TS07; TS10] for the problem. Several questions are raised by our work. First, we have shown lower bounds for relative ϵ-approximately well-supported Nash equilibrium. Can these lower bounds be extended to the weaker notion of relative ϵ-approximate Nash equilibrium? Second, in the bimatrix/polymatrix games constructed in our lower bounds (Theorems 1.1 and 1.4) the payoff functions of the players range in $[-1, 1]$. Can our lower bounds be strengthened to games with payoff functions that range in $[0, 1]$? Support for this possibility stems from the fact that all relative approximate Nash equilibria of our hard instances of Theorems 1.1 and 1.4 give positive payoff to all players. In fact, our hard instances of polymatrix games guarantee positive max-min values to all players. Third, we show constant inapproximability results for relative ϵ-Nash equilibrium for all $\epsilon \in [0, 1)$. How tight is the range of values of ϵ for which our lower bounds apply? Is there a polynomial-time algorithm for $\epsilon = 1$? Finally, our lower bounds only apply to relative approximations, while constant additive approximations are computable in quasi-polynomial-time [LMM03]. Is there a quasi-polynomial-time lower bound?

APPENDIX

A. COMPLETE GADGETS OF SECTION 3

We provide both the "simple" and "sophisticated" versions of our gadgets introduced in Section 3. In particular, Lemma A.1 of this section implies Lemma 3.1 of Section 3, Lemma A.2 implies Lemma 3.2, Lemma A.3 implies 3.3, Lemma A.4 implies 3.4, and

Lemma A.5 implies 3.5. We remind the reader that the "sophisticated" versions of our gadgets are constructed so that every non-input player has a positive max-min value.

LEMMA A.1 (DETECTOR GADGET). *Fix* $\epsilon \in [0, 1)$, $\alpha, \beta, \gamma \in [-1, 1]$, *and* $c \in \mathbb{N}$. *There exist* $c', n_0 \in \mathbb{N}$, *such that for all* $n > n_0$:

— *there exists a graphical polymatrix game* $\mathcal{G}_{\rm det}$ *with three input players* x, y *and* z, *one intermediate player* w, *and one output player* t, *and two strategies per player,* 0 *and* 1, *such that in any relative* ϵ*-Nash equilibrium of* $\mathcal{G}_{\rm det}$, *the mixed strategies of the players satisfy the following*

$$p(z:1) - [\alpha p(x:1) + \beta p(y:1) + \gamma] \geq 2^{-cn} \Rightarrow \quad p(t:1) = 1; \tag{17}$$

$$p(z:1) - [\alpha p(x:1) + \beta p(y:1) + \gamma] \leq -2^{-cn} \Rightarrow \quad p(t:1) = 0; \tag{18}$$

— *there exists a graphical polymatrix game* $\mathcal{G}^+_{\rm det}$ *with the same characteristics as* $\mathcal{G}_{\rm det}$, *except that every player has three strategies* 0, 1, *and* $*$, *and such that Properties* (17) *and* (18) *are satisfied, in any relative* ϵ*-Nash equilibrium, and moreover every (non-input) player receives a positive payoff of* $2^{-c'n}$ *if she plays strategy* $*$, *regardless of the strategies of the other players of the game.*

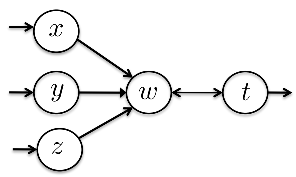

Fig. 1. The detector gadgets $\mathcal{G}_{\rm det}$ and $\mathcal{G}^+_{\rm det}$.

PROOF OF LEMMA *A.1:* The graphical structure of the games $\mathcal{G}_{\rm det}$ and $\mathcal{G}^+_{\rm det}$ is shown in Figure 1, where the direction of the edges denotes direct payoff dependence. The construction of the games $\mathcal{G}_{\rm det}$ and $\mathcal{G}^+_{\rm det}$ is similar, so we are only going to describe the construction of $\mathcal{G}^+_{\rm det}$. A trivial adaptation of this construction—by just removing all the $*$ strategies—gives the construction of $\mathcal{G}_{\rm det}$. Let us choose $c' > c, n_0$ such that $(1-\epsilon)2^{-cn} > 2^{-c'n}$, for all $n > n_0$.

Since the players x, y and z are input players, to specify the game we only need to define the payoffs of the players w and t. The payoff of player w is defined as follows:

— $u(w:*) = 2^{-c'n}$;
— $u(w:0) = \mathbb{1}_{z:1} - \alpha \cdot \mathbb{1}_{x:1} - \beta \cdot \mathbb{1}_{y:1} - \gamma$;
— $u(w:1) = 2^{-c'n} \cdot \mathbb{1}_{t:1}$;

where $\mathbb{1}_A$ denotes the indicator function of the event A. The payoff of player t is defined so that she always prefers to disagree with w:

— $u(t:*) = 2^{-c'n}$;
— $u(t:0) = \mathbb{1}_{w:1}$;
— $u(t:1) = \mathbb{1}_{w:0}$;

Clearly, both w and t receive a payoff of $2^{-c'n}$ if they play strategy $*$ regardless of the strategies of the other players of the game. So, we only need to argue that (17) and (18) are satisfied. Observe that the expected payoff of player w is $p(z:1) - [\alpha p(x:1) + \beta p(y:$

$1)+\gamma]$ for playing 0 and $2^{-c'n}\cdot p(t:1)$ for playing 1, while the expected payoff of player t is $p(w:1)$ for playing 0 and $p(w:0)$ for playing 1.

To argue that (17) is satisfied, suppose that in some relative ϵ-Nash equilibrium we have

$$p(z:1)-[\alpha p(x:1)+\beta p(y:1)+\gamma]\geq 2^{-cn}.$$

Then the expected payoff of player w is at least 2^{-cn} for playing 0, while it is at most $2^{-c'n}$ from strategies 1 and $*$. But, $(1-\epsilon)2^{-cn}>2^{-c'n}$, for all $n>n_0$. Hence, in any relative ϵ-Nash equilibrium, it must be that $p(w:0)=1$. Given this, the expected payoff of player t is 1 for playing strategy 1, while her expected payoff from strategy 0 is 0 and from strategy $*$ is $2^{-c'n}$. Hence, in a relative ϵ-Nash equilibrium, it must be that $p(t:1)=1$. So (17) is satisfied.

To show (18), suppose that in some relative ϵ-Nash equilibrium

$$p(z:1)-[\alpha p(x:1)+\beta p(y:1)+\gamma]\leq -2^{-cn}.$$

Then the expected payoff of player w is at most -2^{-cn} for playing 0, while she gets $2^{-c'n}$ for playing $*$ and ≥ 0 for playing 1. So, in any relative ϵ-Nash equilibrium $p(w:0)=0$ (recall that $\epsilon<1$.) Hence, the expected payoff to player t for playing strategy 1 is 0, while she gets at least $2^{-c'n}$ for playing $*$ and $p(w:1)$ for playing 0. So, in any relative ϵ-Nash equilibrium $p(t:1)=0$ (where we used again that $\epsilon<1$.) □

LEMMA A.2 (AFFINE OPERATOR). *Fix $\epsilon\in[0,1)$, $\alpha,\beta,\gamma\in[-1,1]$, and $c\in\mathbb{N}$. There exists $n_0,c'\in\mathbb{N}$, such that for all $n>n_0$*

— *there is a graphical polymatrix game $\mathcal{G}_{\rm lin}$ with a bipartite graph, two input players x and y, and one output player z, such that in any relative ϵ-Nash equilibrium*

$$p(z:1)\geq \max\{0,\min\{1,\alpha p(x:1)+\beta p(y:1)+\gamma\}\}-2^{-cn}; \qquad (19)$$
$$p(z:1)\leq \min\{1,\max\{0,\alpha p(x:1)+\beta p(y:1)+\gamma\}\}+2^{-cn}; \qquad (20)$$

— *there also exists a graphical polymatrix game $\mathcal{G}^{+}_{\rm lin}$ with the same characteristics as $\mathcal{G}_{\rm lin}$, except that every player has three strategies 0, 1, and $*$, and such that properties* (19) *and* (20) *are satisfied in any relative ϵ-Nash equilibrium, and moreover every (non-input) player receives a positive payoff of $2^{-c'n}$ if she plays strategy $*$, regardless of the strategies of the other players of the game.*

PROOF OF LEMMA *A.2:* $\mathcal{G}_{\rm lin}$ and $\mathcal{G}^{+}_{\rm lin}$ have the graphical structure shown in Figure 2. They are obtained by adding feedback to the gadgets $\mathcal{G}_{\rm det}$ and $\mathcal{G}^{+}_{\rm det}$ respectively through a new player w' who is introduced to relay feedback while maintaining the graph bipartite. We describe the nature of this feedback by specifying the payoffs of players w' and z. Again we are only going to describe the gadget $\mathcal{G}^{+}_{\rm lin}$, and the description of $\mathcal{G}_{\rm lin}$ is the same, except that the strategies $*$ are removed. Let us choose c',n_0 such that $(1-\epsilon)2^{-cn}>2^{-c'n}$, for all $n>n_0$. We assign to player w' the following payoff:

— $u(w':*)=2^{-c'n}$;
— $u(w':0)=\mathbb{1}_{t:1}$;
— $u(w':1)=1-\mathbb{1}_{t:1}$;

and we assign to player z the following payoff:

— $u(z:*)=2^{-c'n}$;
— $u(z:0)=\mathbb{1}_{w':0}$;
— $u(z:1)=\mathbb{1}_{w':1}$.

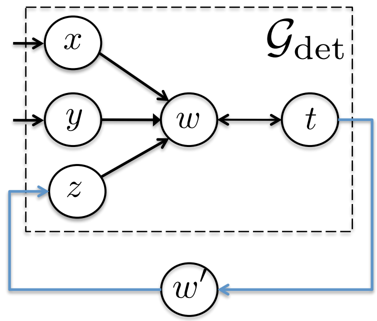

Fig. 2. The affine operator gadgets $\mathcal{G}_{\rm lin}$ and $\mathcal{G}^{+}_{\rm lin}$.

Now, we proceed to argue that (19) and (20) are satisfied. We distinguish three cases:

— $[\alpha p(x:1)+\beta p(y:1)+\gamma]\leq 0$: In this case we have

$$\max\{0\,,\min\{1\,,\alpha p(x:1)+\beta p(y:1)+\gamma\}\}=0,$$
$$\min\{1\,,\max\{0,\ \alpha p(x:1)+\beta p(y:1)+\gamma\}\}=0.$$

So, clearly, (19) is satisfied. To show (20), suppose for a contradiction that

$$p(z:1)>\ \min\{1\,,\max\{0,\ \alpha p(x:1)+\beta p(y:1)+\gamma\}\}+2^{-cn}. \tag{21}$$

The above implies, $p(z:1)>[\alpha p(x:1)+\beta p(y:1)+\gamma]+2^{-cn}$; hence, as in the proof of Lemma A.1, $p(t:1)=1$. Given this, the expected payoff of w' is 1 for playing 0, while at most $2^{-c'n}$ for playing $*$ or 1. But, $(1-\epsilon)1>2^{-c'n}$, for all $n>n_0$. Hence, in a relative ϵ-Nash equilibrium it must be that $p(w':0)=1$. Now, the expected payoff of player z is 1 for playing 0, and at most $2^{-c'n}$ for playing $*$ or 1. Hence, in a relative ϵ-Nash equilibrium it must be that $p(z:0)=1$. Hence, $p(z:1)=0$, which contradicts (21).

— $0\leq[\alpha p(x:1)+\beta p(y:1)+\gamma]\leq 1$: In this case we have

$$\max\{0\,,\min\{1\,,\alpha p(x:1)+\beta p(y:1)+\gamma\}\}=\alpha p(x:1)+\beta p(y:1)+\gamma,$$
$$\min\{1\,,\max\{0,\ \alpha p(x:1)+\beta p(y:1)+\gamma\}\}=\alpha p(x:1)+\beta p(y:1)+\gamma.$$

Suppose for a contradiction that

$$p(z:1)>[\alpha p(x:1)+\beta p(y:1)+\gamma]+2^{-cn}. \tag{22}$$

As in the proof of Lemma A.1, this implies $p(t:1)=1$. Given this, the expected payoff of w' is 1 for playing 0, while at most $2^{-c'n}$ for playing $*$ or 1. Hence, in a relative ϵ-Nash equilibrium it must be that $p(w':0)=1$. Now, the expected payoff of player z is 1 for playing 0, and at most $2^{-c'n}$ for playing $*$ or 1. Hence, in a relative ϵ-Nash equilibrium it must be that $p(z:0)=1$. Hence, $p(z:1)=0$, which contradicts (22), and therefore (20) is satisfied. To argue that (19) is satisfied, suppose for a contradiction that

$$p(z:1)<[\alpha p(x:1)+\beta p(y:1)+\gamma]-2^{-cn}. \tag{23}$$

As in the proof of Lemma A.1, this implies $p(t:1)=0$. Given this, the expected payoff of w' is 1 for playing 1, while at most $2^{-c'n}$ for playing $*$ or 0. Hence, in a relative ϵ-Nash equilibrium it must be that $p(w':1)=1$. Now, the expected payoff of player z is 1 for playing 1, and at most $2^{-c'n}$ for playing $*$ or 0. Hence, in a relative

ϵ-Nash equilibrium it must be that $p(z:1)=1$, which contradicts (23). Hence, (19) is satisfied.

— $[\alpha p(x:1)+\beta p(y:1)+\gamma]>1$: In this case,

$$\max\{0\,,\min\{1\,,\alpha p(x:1)+\beta p(y:1)+\gamma\}\}=1,$$
$$\min\{1\,,\max\{0,\ \alpha p(x:1)+\beta p(y:1)+\gamma\}\}=1.$$

So, automatically (20) is satisfied. To show (19), suppose for a contradiction that

$$p(z:1)<\ \max\{0,\min\{1,\alpha p(x:1)+\beta p(y:1)+\gamma\}\}-2^{-cn}. \tag{24}$$

The above implies, $p(z:1)<[\alpha p(x:1)+\beta p(y:1)+\gamma]-2^{-cn}$. As in the proof of Lemma A.1, this implies $p(t:1)=0$. Given this, the expected payoff of w' is 1 for playing 1, while at most $2^{-c'n}$ for playing $*$ or 0. Hence, in a relative ϵ-Nash equilibrium it must be that $p(w':1)=1$. Now, the expected payoff of player z is 1 for playing 1, and at most $2^{-c'n}$ for playing $*$ or 0. Hence, in a relative ϵ-Nash equilibrium it must be that $p(z:1)=1$, which contradicts (24). Hence, (19) is satisfied.

□

LEMMA A.3 (ARITHMETIC GADGETS). *Fix $\epsilon\geq 0$, $\zeta\geq 0$, and $c\in\mathbb{N}$. There exists $c',n_0\in\mathbb{N}$, such that for all $n>n_0$:*

— *there exist graphical polymatrix games $\mathcal{G}_+,\mathcal{G}_-,\mathcal{G}_{\times\zeta},\mathcal{G}_\zeta$ with bipartite interaction graphs, two input players x and y, and one output player z, such that in any relative ϵ-Nash equilibrium*
  - *the game $\mathcal{G}_+$ satisfies* $p(z:1)=\min\{1\,,p(x:1)+p(y:1)\}\pm 2^{-cn}$;
  - *the game $\mathcal{G}_-$ satisfies* $p(z:1)=\max\{0\,,p(x:1)-p(y:1)\}\pm 2^{-cn}$;
  - *the game $\mathcal{G}_{\times\zeta}$ satisfies* $p(z:1)=\min\{1\,,\zeta\cdot p(x:1)\}\pm 2^{-cn}$;
  - *the game $\mathcal{G}_\zeta$ satisfies* $p(z:1)=\min\{1\,,\zeta\}\pm 2^{-cn}$;

— *there also exist graphical polymatrix games $\mathcal{G}_+^+$, $\mathcal{G}_-^+$, $\mathcal{G}_{\times\zeta}^+$, $\mathcal{G}_\zeta^+$ with the same characteristics as the graphical games $\mathcal{G}_+$, $\mathcal{G}_-$, $\mathcal{G}_{\times\zeta}$, $\mathcal{G}_\zeta$, except that every player has three strategies 0, 1, and $*$, and such that the above properties are satisfied in any relative ϵ-Nash equilibrium, and moreover every (non-input) player receives a positive payoff of $2^{-c'n}$ if she plays strategy $*$, regardless of the strategies of the other players of the game.*

PROOF OF LEMMA *A.3:* All the gadgets are obtained from $\mathcal{G}_{\rm lin}$ and $G^+_{\rm lin}$ with the appropriate setting of the parameters α, β and γ. For $\mathcal{G}_+$ and $\mathcal{G}_+^+$, set $\alpha=\beta=1$ and $\gamma=0$. For $\mathcal{G}_-$ and $\mathcal{G}_-^+$ set $\alpha=1$, $\beta=-1$ and $\gamma=0$. For $\mathcal{G}_{\times\zeta}$ and $\mathcal{G}_{\times\zeta}^+$ set $\alpha=\zeta$ and $\beta=\gamma=0$. Finally, for $\mathcal{G}_\zeta$ and $\mathcal{G}_\zeta^+$ set $\alpha=\beta=0$ and $\gamma=\zeta$. □

LEMMA A.4 (COMPARATOR GADGET). *Fix $\epsilon\in[0,1)$, and $c\in\mathbb{N}$. There exist $c',n_0\in\mathbb{N}$, such that for all $n>n_0$:*

— *there exists a graphical polymatrix game $\mathcal{G}_>$ with bipartite interaction graph, two input players x and z, and one output player t, such that in any relative ϵ-Nash equilibrium of $\mathcal{G}_>$*

$$p(z:1)-p(x:1)\geq 2^{-cn}\quad\Rightarrow\quad p(t:1)=1; \tag{25}$$
$$p(z:1)-p(x:1)\leq -2^{-cn}\quad\Rightarrow\quad p(t:1)=0; \tag{26}$$

— *there also exists a graphical polymatrix game $\mathcal{G}_>^+$ with the same characteristics as $\mathcal{G}_>$, except that every player has three strategies 0, 1, and $*$, and such that the above* *properties are satisfied in any relative ϵ-Nash equilibrium, and moreover every (non-input) player receives a positive payoff of $2^{-c'n}$ if she plays strategy $*$, regardless of the strategies of the other players of the game.*

PROOF OF LEMMA *A.4:* $\mathcal{G}_>$ and $\mathcal{G}_>^+$ are obtained from $\mathcal{G}_{\rm det}$ and $\mathcal{G}_{\rm det}^+$ respectively, by setting $\alpha = 1$, $\beta = \gamma = 0$. □

LEMMA A.5 (BOOLEAN OPERATORS). *Fix $\epsilon \in [0, 1)$ and $c' \in \mathbb{N}$. There exists $n_0 \in \mathbb{N}$, such that for all $n > n_0$:*

— *there exist graphical polymatrix games $\mathcal{G}_\vee, \mathcal{G}_\neg$ with bipartite interaction graphs, two input players x and y, and one output player z, such that in any relative ϵ-Nash equilibrium*
  — *if $p(x : 1), p(y : 1) \in \{0, 1\}$, the game $\mathcal{G}_\vee$ satisfies $p(z : 1) = p(x : 1) \vee p(y : 1)$;*
  — *if $p(x : 1) \in \{0, 1\}$, the game $\mathcal{G}_\neg$ satisfies $p(z : 1) = 1 - p(x : 1)$;*

— *there also exist graphical polymatrix games $\mathcal{G}_\vee^+, \mathcal{G}_\neg^+$ with the same characteristics as the games $\mathcal{G}_\vee, \mathcal{G}_\neg$, except that every player has three strategies 0, 1, and $*$, and such that the above properties are satisfied in any relative ϵ-Nash equilibrium, and moreover every (non-input) player receives a positive payoff of $2^{-c'n}$ if she plays strategy $*$, regardless of the strategies of the other players of the game.*

PROOF OF LEMMA *A.5:* The structure of the graphical games $\mathcal{G}_\vee, \mathcal{G}_\neg$, $\mathcal{G}_\vee^+, \mathcal{G}_\neg^+$ is shown in Figure 3. We are going to describe $\mathcal{G}_\vee^+, \mathcal{G}_\neg^+$; the other games are obtained by dropping strategy $*$. We choose c', n_0 such that $1 - \epsilon > 2^{-c'n}$. To define the game $\mathcal{G}_\vee^+$, we give

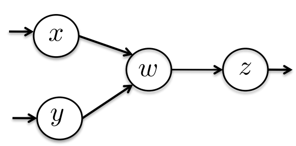

Fig. 3. The gadgets $\mathcal{G}_\vee, \mathcal{G}_\neg$, $\mathcal{G}_\vee^+, \mathcal{G}_\neg^+$.

player w the following payoff function:

$$u(w : 0) = 2^{-c'n}; \quad u(w : *) = 2^{-c'n}; \quad u(w : 1) = \mathbb{1}_{x:1} + \mathbb{1}_{y:1};$$

we also give player z an incentive to agree with player w as follows

$$u(z : 0) = \mathbb{1}_{w:0}; \quad u(z : *) = 2^{-c'n}; \quad u(z : 1) = \mathbb{1}_{w:1}.$$

Now suppose that, in some relative ϵ-Nash equilibrium, $p(x : 1), p(y : 1) \in \{0, 1\}$ and $p(x : 1) \vee p(y : 1) = 1$. Then the expected payoff to player w is at least 1 for choosing strategy 1, and $2^{-c'n}$ for choosing strategy 0 or $*$. Since, $1 - \epsilon > 2^{-c'n}$, it follows that $p(w : 1) = 1$. Given this, the expected payoff to player z is 1 for playing 1 and at most $2^{-c'n}$ for choosing strategy $*$ or 0. Hence, $p(z : 1) = 1$. On the other hand, if $p(x : 1) \vee p(y : 1) = 0$, the expected payoff to player w is $2^{-c'n}$ for choosing strategies 0 or $*$, and 0 for choosing strategy 1. Hence, $p(w : 1) = 0$. Given this, the expected payoff to player z is 0 for choosing strategy 1, $2^{-c'n}$ for choosing strategy $*$, and $p(w : 0)$ for choosing strategy 0. Hence, $p(z : 1) = 0$. So, $p(z : 1) = p(x : 1) \vee p(y : 1)$.

In the game $\mathcal{G}_\neg^+$ player w has the following payoff function:

$$u(w : 0) = \mathbb{1}_{x:1}; \quad u(w : *) = 2^{-c'n}; \quad u(w : 1) = 1 - \mathbb{1}_{x:1};$$

and we give player z an incentive to agree with player w as follows

$$u(z:0) = \mathbb{1}_{w:0}; \quad u(z:*) = 2^{-c'n}; \quad u(z:1) = \mathbb{1}_{w:1}.$$

Now suppose that, in some relative ϵ-Nash equilibrium, $p(x:1) = 1$. Then the expected payoff to player w is 1 for choosing strategy 0, and at most $2^{-c'n}$ for choosing strategy $*$ or 1. Since, $1-\epsilon > 2^{-c'n}$, it follows that $p(w:0) = 1$. Given this, the expected payoff to player z is 1 for playing 0 and at most $2^{-c'n}$ for choosing strategy $*$ or 1. Hence, $p(z:1) = 0 = 1 - p(x:1)$. On the other hand, if $p(x:1) = 0$, the expected payoff to player w is 1 for choosing strategy 1, and at most $2^{-c'n}$ for choosing strategies $*$ or 0. Hence, $p(w:1) = 1$. Given this, the expected payoff to player z is 1 for choosing strategy 1, $2^{-c'n}$ for choosing strategy $*$, and 0 for choosing strategy 0. Hence, $p(z:1) = 1$. So, $p(z:1) = 1 - p(x:1)$. □

B. THE BIMATRIX GAME IN OUR CONSTRUCTION

See Figure 4.

$$\widehat{U}_p(\sigma\ ;\ \sigma') := \begin{cases} -U + 2^{-dn}, & \text{if } \sigma = v^*,\ \sigma' = \text{bad}_v,\ \rho(v) = p; \\ 2^{-dn}, & \text{if } \sigma = v^*,\ \sigma' \neq \text{bad}_v,\ \rho(v) = p; \\ U_p(\sigma\ ;\ \sigma'), & \text{if } \sigma = (v:s),\ \sigma' = (v':s'),\ \rho(v) = 1 - \rho(v') = p; \\ -U + \frac{1}{n}\beta^{v:s}, & \text{if } \sigma = (v:s),\ \sigma' = \text{bad}_v,\ \rho(v) = p; \\ \frac{1}{n}\beta^{v:s}, & \text{if } \sigma = (v:s),\ \sigma' = \text{bad}_{v'},\ \rho(v) = \rho(v') = p; \\ H \cdot (1 - \frac{1}{n}), & \text{if } \sigma = \text{bad}_v,\ \sigma' \in \{(v:0), (v:1), v^*\},\ \rho(v) = 1 - p; \\ H \cdot (-\frac{1}{n}), & \text{if } \sigma = \text{bad}_v,\ \sigma' \notin \{(v:0), (v:1), v^*\},\ \rho(v) = 1 - p; \end{cases} \tag{27}$$

Fig. 4. The payoffs of the bimatrix game in our construction, where $\widehat{U}_p(\sigma\ ;\ \sigma')$ denotes the payoff of player p, when she plays strategy σ and her opponent player σ'.

C. CHOOSING THE RIGHT CONSTANTS

See Figure 5

$$H \cdot \delta \cdot (1-\epsilon)^2 > 1$$
$$U\frac{1}{n} > 1$$
$$H > U \cdot n(1+\epsilon)$$
$$U\delta > 1$$
$$2^{-cn} \geq \frac{2^{-cn}}{2(1-2n^2\delta)} + 5 \cdot \frac{2n^2\delta}{(1-2n^2\delta)}$$
$$(1-\epsilon)2^{-cn} > 2^{-c'n}2(1+n\delta)$$
$$(1-\epsilon)2^{-cn} > 2n2^{-dn}$$
$$(1-\epsilon)2^{-c'n} > n2^{-dn}$$

Fig. 5. We choose a sufficiently large n_0, so that the above inequalities are satisfied for all $n > n_0$, for the choices of H, U, d, δ, c and c' made in Section 4.2.

D. OMITTED DETAILS FROM SECTION 4

D.1. Analysis of Threats.

PROOF OF LEMMA *4.1:* Let $v_a \in \arg\max_v\{p_v\}$ and, for a contradiction, suppose that $p_{v_a} > \frac{1}{n} + \delta$. Now define the set

$$S = \left\{ v \,\middle|\, \rho(v) = \rho(v_a),\ p_v - \frac{1}{n} \geq (1-\epsilon) \cdot \left(p_{v_a} - \frac{1}{n} \right) \right\}.$$

Since $p_{v_a} > 1/n$, there must be some v_b, with $\rho(v_b) = \rho(v_a)$, such that $p_{v_b} < 1/n$.

Now the expected payoff of player $1 - \rho(v_a)$ for playing any strategy $\mathrm{bad}_v, v \in S$ is at least $H \cdot \delta \cdot (1-\epsilon)$ and, by assumption, $H \cdot \delta \cdot (1-\epsilon)^2 > 1$. So, in any relative ϵ-Nash equilibrium of the game, player $1 - \rho(v_a)$ will not play any strategy of the form $(v : s)$, since her expected payoff from these strategies is at most 1 (recall that all payoffs of $\mathcal{G}G$ were scaled in Section 4.2 to be smaller than 1). She will also not play any strategy of the form v^* as her expected payoff from these strategies is at most 2^{-dn}. So player $1 - \rho(v_a)$ only uses strategies of the form bad_v in her mixed strategy in any relative ϵ-Nash equilibrium of the game. Moreover, she will not use any strategy of the form bad_v where $v \notin S$, because by the definition of the set S she is better off playing strategy bad_{v_a} by more than a relative ϵ. Finally, $|S| < n$, since $v_b \notin S$. Hence, there must be some $v_c \in S$, such that $p_{\mathrm{bad}_{v_c}} > 1/n$.

Let's go back now to player $p(v_a)$. Her expected payoff from strategy v_b^* is at least 2^{-dn} (since we argued that $p_{\mathrm{bad}_{v_b}} = 0$), while her expected payoff from strategies $v_c : 0$, $v_c : 1$ and v_c* is at most $-U\frac{1}{n} + 1 < 0$, since $p_{\mathrm{bad}_{v_c}} > 1/n$ and we assumed that all payoffs in the graphical polymatrix game are at most 1. Hence, in any relative ϵ-Nash equilibrium, it must be that $p_{v_c} = 0$, which is a contradiction since we assumed that $v_c \in S$. $\square$

PROOF OF LEMMA *4.2:* Let us fix some player p of the bimatrix game. We distinguish the following cases:

— There exist v_a, v_b, with $\rho(v_a) = \rho(v_b) = p$, such that $p_{v_a} \geq 1/n$ and $p_{v_b} < 1/n$: The payoff of player $1-p$ from strategy bad_{v_a} is ≥ 0, while her payoff from strategy bad_{v_b} is < 0. Hence, in any relative ϵ-Nash equilibrium, player $1-p$ plays strategy bad_{v_b} with probability 0. So, the payoff of player p for playing strategy v_b^* is at least 2^{-dn}. Hence, her payoff must be at least $(1-\epsilon)2^{-dn}$ from every strategy in her support.

— $p_v < 1/n$, for all v with $\rho(v) = p$: Let $v_a \in \arg\min_{v:\rho(v)=p}\{p_v\}$. Let then $\phi_a := 1/n - p_{v_a}$. Observe that the expected payoff of player $1-p$ is $-H\phi_a$ for playing strategy bad_{v_a}, while her expected payoff from every strategy v^*, $\rho(v) = 1-p$ is at least $-U \cdot p_{\mathrm{bad}_v} \geq -U \cdot n \cdot \phi_a$. Since $U \cdot n(1+\epsilon) < H$ it follows that $-U \cdot n \cdot \phi_a(1+\epsilon) > -H\phi_a$. So player $1-p$ is going to play strategy bad_{v_a} with probability 0 in any relative ϵ-Nash equilibrium. Hence, the expected payoff of player p for playing strategy v_a^* will be 2^{-dn}. Hence, her payoff must be at least $(1-\epsilon)2^{-dn}$ from every strategy in her support.

— $p_v = 1/n$, for all v with $\rho(v) = p$: It must be that $p_{\mathrm{bad}_v} = 0$, for all v with $\rho(v) = 1-p$. Hence, the expected payoff of player $1-p$ is at least 2^{-dn} from every v^*, while her expected payoff is 0 from every strategy bad_v. So, player $1-p$ is going to play all strategies bad_v with probability 0. So, the expected payoff of player p is at least 2^{-dn} from every strategy v^*, $\rho(v) = p$. Hence, her payoff must be at least $(1-\epsilon)2^{-dn}$ from every strategy in her support.

$\square$

PROOF OF LEMMA *4.3:* Let $v_a \in \arg\min_v\{p_v\}$ and, for a contradiction, suppose that $p_{v_a} < \frac{1}{n} - 2n\delta$. Using Lemma 4.1, it follows that there must exist some v_b with $\rho(v_b) =$

$1 - \rho(v_a)$ such that

$$p_{\mathrm{bad}_{v_b}} \geq \frac{1}{n}(2n - (n-1))\delta > \delta. \tag{28}$$

Then the payoff that player $\rho(v_b)$ gets from all her strategies in the block corresponding to v_b is at most $-U\delta + 1 < 0$ (since $p_{\mathrm{bad}_{v_b}} > \delta$ and the payoffs from the graphical game are at most 1). Hence, by Lemma 4.2 it follows that in any relative ϵ-Nash equilibrium, it must be that $p_{v_b} = 0$. But then the payoff of player $\rho(v_a)$ from strategy bad_{v_b} is $-H1/n < 0$. And by Lemma 4.2 again, it must be that $p_{\mathrm{bad}_{v_b}} = 0$. This contradicts (28). □

D.2. Un-normalized Graphical-Game Equilibrium Conditions from Relative Equilibria of the Bimatrix Game.

PROOF OF LEMMA *4.4:* Notice first that $\sigma_{\max} \in \arg\max_{\sigma \in \{v*,(v:0),(v:1)\}} \left\{\mathcal{E}\left(\widehat{U}_{p:\sigma}\right)\right\}$. Next, from Lemmas 4.2 and 4.3, it follows that

$$-U \cdot p_{\mathrm{bad}_v} + U'_{\sigma_{\max}} > 0. \tag{29}$$

Now, for a given $\sigma \in \{v*, (v:0), (v:1)\} \setminus \{\sigma_{\max}\}$, we distinguish the following cases:

— $-U \cdot p_{\mathrm{bad}_v} + U'_\sigma < 0$: This implies that the expected payoff to player $\rho(v)$ for playing strategy σ is negative, while the expected payoff from strategy $\sigma_{\max}$ is positive (see Equation (29)), so the implication is true.

— $-U \cdot p_{\mathrm{bad}_v} + U'_\sigma \geq 0$: We have

$$\begin{aligned}\frac{-U \cdot p_{\mathrm{bad}_v} + U'_\sigma}{-U \cdot p_{\mathrm{bad}_v} + U'_{\sigma_{\max}}} &< \frac{-U \cdot p_{\mathrm{bad}_v} + (1-\epsilon)U'_{\sigma_{\max}}}{-U \cdot p_{\mathrm{bad}_v} + U'_{\sigma_{\max}}} \\ &= \frac{-U \cdot p_{\mathrm{bad}_v} + U'_{\sigma_{\max}} - \epsilon U'_{\sigma_{\max}}}{-U \cdot p_{\mathrm{bad}_v} + U'_{\sigma_{\max}}} \\ &= 1 - \epsilon \cdot \frac{U'_{\sigma_{\max}}}{-U \cdot p_{\mathrm{bad}_v} + U'_{\sigma_{\max}}} \\ &\leq 1 - \epsilon.\end{aligned}$$

Hence, player $\rho(v)$ will assign probability 0 to strategy σ.

□

D.3. Approximate Gate-Evaluations from Relative Equilibria of the Bimatrix Game.

PROOF OF LEMMA *4.5:* We show (13) first. Suppose $\hat{p}(z:1) - [\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma] \geq 2^{-cn}$. This implies the following

$$\frac{p_{z:1}}{p_z} - \left[\alpha \cdot \frac{p_{x:1}}{p_x} + \beta \cdot \frac{p_{y:1}}{p_y} + \gamma\right] \geq 2^{-cn}; \tag{30}$$

Now we show

CLAIM D.1. *The above imply:*

$$p_{z:1} - \left[\alpha \cdot p_{x:1} + \beta \cdot p_{y:1} + \frac{\gamma}{n}\right] \geq \frac{2^{-cn}}{2n}. \tag{31}$$

PROOF. Indeed, suppose that

$$p_{z:1} - \left[\alpha \cdot p_{x:1} + \beta \cdot p_{y:1} + \frac{\gamma}{n}\right] < \frac{2^{-cn}}{2n}.$$

Then

$$\begin{aligned}
&\frac{p_{z:1}}{p_z} - \left[\alpha \cdot \frac{p_{x:1}}{p_x} + \beta \cdot \frac{p_{y:1}}{p_y} + \gamma\right] \\
&< \frac{2^{-cn}}{2np_z} - \alpha \cdot \left[\frac{p_{x:1}}{p_x} - \frac{p_{x:1}}{p_z}\right] - \beta \cdot \left[\frac{p_{y:1}}{p_y} - \frac{p_{y:1}}{p_z}\right] - \left[\gamma - \frac{\gamma}{np_z}\right] \\
&\le \frac{2^{-cn}}{2np_z} - \alpha \cdot p_{x:1}\left[\frac{p_z - p_x}{p_x p_z}\right] - \beta p_{y:1} \cdot \left[\frac{p_z - p_y}{p_y p_z}\right] - \gamma\left[\frac{p_z - 1/n}{p_z}\right] \\
&\le \frac{2^{-cn}}{2np_z} + |\alpha| p_{x:1} \cdot \frac{|p_z - p_x|}{p_x p_z} + |\beta| p_{y:1} \cdot \frac{|p_z - p_y|}{p_y p_z} + |\gamma| \cdot \frac{|p_z - 1/n|}{p_z} \\
&\le \frac{2^{-cn}}{2(1-2n^2\delta)} + |\alpha| p_{x:1} \cdot \frac{4n\delta}{p_x p_z} + |\beta| p_{y:1} \cdot \frac{4n\delta}{p_y p_z} + |\gamma| \cdot \frac{2n\delta}{p_z} \\
&\le \frac{2^{-cn}}{2(1-2n^2\delta)} + |\alpha| \cdot \frac{4n\delta}{p_z} + |\beta| \cdot \frac{4n\delta}{p_z} + |\gamma| \cdot \frac{2n\delta}{p_z} \\
&\le \frac{2^{-cn}}{2(1-2n^2\delta)} + (2|\alpha| + 2|\beta| + |\gamma|) \cdot \frac{2n\delta}{p_z} \\
&\le \frac{2^{-cn}}{2(1-2n^2\delta)} + (2|\alpha| + 2|\beta| + |\gamma|) \cdot \frac{2n^2\delta}{(1-2n^2\delta)} \\
&\le 2^{-cn}. \qquad \text{(using Figure 5)}
\end{aligned}$$

This is a contradiction to (30). □

Given (31) we have

$$\begin{aligned}
&U'_{w:0} \ge \frac{2^{-cn}}{2n}; \\
&U'_{w:1} = 2^{-c'n} p_{t:1} \le 2^{-c'n}\frac{1}{n}(1+n\delta); \;\; \text{(using Lemma 4.1)} \\
&U'_{w*} = 2^{-dn}.
\end{aligned}$$

From Lemma 4.4, it follows then that $p_{w:1} = p_{w*} = 0$. Hence, $p_{w:0} = p_w$. Given this, we have

$$\begin{aligned}
&U'_{t:0} = 0; \\
&U'_{t:1} = p_{w:0} = p_w \ge 1/n(1-2n^2\delta); \;\; \text{(using Lemma 4.3)} \\
&U'_{t*} = 2^{-dn}.
\end{aligned}$$

Hence, Lemma 4.4 implies $p_{t:1} = p_t$. So that $\hat{p}(t:1) = 1$.

We show (14) similarly. Suppose $\hat{p}(z:1) - [\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma] \le -2^{-cn}$. This implies the following

$$\frac{p_{z:1}}{p_z} - \left[\alpha \cdot \frac{p_{x:1}}{p_x} + \beta \cdot \frac{p_{y:1}}{p_y} + \gamma\right] \le -2^{-cn}. \tag{32}$$

Next we show

CLAIM D.2. *The above imply:*

$$p_{z:1} - \left[\alpha \cdot p_{x:1} + \beta \cdot p_{y:1} + \frac{\gamma}{n}\right] \le -\frac{2^{-cn}}{2n}. \tag{33}$$

PROOF. Indeed, suppose that

$$p_{z:1} - \left[\alpha \cdot p_{x:1} + \beta \cdot p_{y:1} + \frac{\gamma}{n}\right] > -\frac{2^{-cn}}{2n}.$$

Then

$$\begin{aligned}
&\frac{p_{z:1}}{p_z} - \left[\alpha \cdot \frac{p_{x:1}}{p_x} + \beta \cdot \frac{p_{y:1}}{p_y} + \gamma\right] \\
&\quad > -\frac{2^{-cn}}{2np_z} - \alpha \cdot \left[\frac{p_{x:1}}{p_x} - \frac{p_{x:1}}{p_z}\right] - \beta \cdot \left[\frac{p_{y:1}}{p_y} - \frac{p_{y:1}}{p_z}\right] - \left[\gamma - \frac{\gamma}{np_z}\right] \\
&\quad = -\frac{2^{-cn}}{2np_z} - \alpha \cdot p_{x:1}\left[\frac{p_z - p_x}{p_x p_z}\right] - \beta p_{y:1} \cdot \left[\frac{p_z - p_y}{p_y p_z}\right] - \gamma\left[\frac{p_z - 1/n}{p_z}\right] \\
&\quad \geq -\frac{2^{-cn}}{2np_z} - |\alpha| p_{x:1} \cdot \frac{|p_z - p_x|}{p_x p_z} - |\beta| p_{y:1} \cdot \frac{|p_z - p_y|}{p_y p_z} - |\gamma| \cdot \frac{|p_z - 1/n|}{p_z} \\
&\quad \geq -\frac{2^{-cn}}{2(1-2n^2\delta)} - |\alpha| p_{x:1} \cdot \frac{4n\delta}{p_x p_z} - |\beta| p_{y:1} \cdot \frac{4n\delta}{p_y p_z} - |\gamma| \cdot \frac{2n\delta}{p_z} \\
&\quad \geq -\frac{2^{-cn}}{2(1-2n^2\delta)} - |\alpha| \cdot \frac{4n\delta}{p_z} - |\beta| \cdot \frac{4n\delta}{p_z} - |\gamma| \cdot \frac{2n\delta}{p_z} \\
&\quad \geq -\frac{2^{-cn}}{2(1-2n^2\delta)} - (2|\alpha| + 2|\beta| + |\gamma|) \cdot \frac{2n\delta}{p_z} \\
&\quad \geq -\frac{2^{-cn}}{2(1-2n^2\delta)} - (2|\alpha| + 2|\beta| + |\gamma|) \cdot \frac{2n^2\delta}{(1-2n^2\delta)} \\
&\quad \geq -2^{-cn} \qquad \text{(using Figure 5)}
\end{aligned}$$

This is a contradiction to (32). □

Given (33) we have $U'_{w:0} \leq -\frac{2^{-cn}}{2n}$. But, $U'_{w*} = 2^{-dn}$. Hence, by Lemma 4.4 we have that $p_{w:0} = 0$. Given this, we have $U'_{t:1} = p_{w:0} = 0$. But, $U'_{t*} = 2^{-dn}$. Hence, $p_{t:1} = 0$. So that $\hat{p}(t:1) = 0$. □

PROOF OF LEMMA *4.7:* The proof proceeds by considering the following cases as in the proof of Lemma A.2:

— $[\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma] \leq 0$: In this case we have

$$\max\{0\,, \min\{1\,, \alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma\}\} = 0,$$
$$\min\{1\,, \max\{0,\ \alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma\}\} = 0.$$

So, clearly, (15) is satisfied. To show (16), suppose for a contradiction that

$$\hat{p}(z:1) > \min\{1, \max\{0, \alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma\}\} + 2^{-cn}. \tag{34}$$

The above implies, $\hat{p}(z:1) > [\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma] + 2^{-cn}$. By Lemma 4.5, this implies $\hat{p}(t:1) = 1$, so $p_{t:1} = p_t$. Given this, $U'_{w':0} = p_{t:1} = p_t \geq \frac{1}{n}(1 - 2n^2\delta)$, while $U'_{w'*} = 2^{-dn}$ and $U'_{w':1} = p_{t:0} = 0$. Lemma 4.4 implies then $p_{w':1} = p_{w'*} = 0$, so that $p_{w':0} = p_{w'}$. Now, $U'_{z:0} = p_{w':0} = p_{w'} > \frac{1}{n}(1-2n^2\delta)$ (using Lemma 4.3), while $U'_{z*} = 2^{-dn}$ and $U'_{z:1} = p_{w':1} = 0$. Invoking Lemma 4.4 we get $p_{z:1} = 0$, hence $\hat{p}(z:1) = 0$ which contradicts (34).

— $0 \leq [\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma] \leq 1$: In this case we have

$$\begin{aligned}&\max\{0\,,\min\{1\,,\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma\}\}\\&\quad = \alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma,\\&\min\{1\,,\max\{0,\ \alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma\}\}\\&\quad = \alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma.\end{aligned}$$

Suppose now that

$$\hat{p}(z:1) > [\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma] + 2^{-cn}. \tag{35}$$

By Lemma 4.5, this implies $\hat{p}(t:1) = 1$, so $p_{t:1} = p_t$. Given this, $U'_{w':0} = p_{t:1} = p_t \geq \frac{1}{n}(1 - 2n^2\delta)$, while $U'_{w'*} = 2^{-dn}$ and $U'_{w':1} = p_{t:0} = 0$. Lemma 4.4 implies then $p_{w':1} = p_{w'*} = 0$, so that $p_{w':0} = p_{w'}$. Now, $U'_{z:0} = p_{w':0} = p_{w'} > \frac{1}{n}(1 - 2n^2\delta)$ (using Lemma 4.3), while $U'_{z*} = 2^{-dn}$ and $U'_{z:1} = p_{w':1} = 0$. Invoking Lemma 4.4 we get $p_{z:1} = 0$, hence $\hat{p}(z:1) = 0$ which contradicts (35). Hence, (16) is satisfied.
To show (15), suppose for a contradiction that

$$\hat{p}(z:1) < [\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma] - 2^{-cn}. \tag{36}$$

From Lemma 4.5 it follows that $p_{t:1} = 0$. Given this, $U'_{w':0} = 0$, $U'_{w':1} = 1/n$ and $U'_{w'*} = 2^{-dn}$. So it follows from Lemma 4.4 that $p_{w':1} = p_{w'}$. Now, $U'_{z:1} = p_{w':1} = p_{w'} > \frac{1}{n}(1 - 2n^2\delta)$ (using Lemma 4.3), while $U'_{z:0} = 0$, $U'_{z*} = 2^{-dn}$. So from Lemma 4.4 we have that $p_{z:1} = p_z$, and therefore $\hat{p}(z:1) = 1$, which contradicts (36). Hence, (15) is satisfied.

— $[\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma] > 1$: In this case,

$$\begin{aligned}&\max\{0\,,\min\{1\,,\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma\}\} = 1,\\&\min\{1\,,\max\{0,\ \alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma\}\} = 1.\end{aligned}$$

So, automatically (16) is satisfied. To show (15), suppose for a contradiction that

$$\begin{aligned}&\hat{p}(z:1) <\\&\max\{0, \min\{1, \alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma\}\} - 2^{-cn}.\end{aligned} \tag{37}$$

The above implies, $\hat{p}(z:1) < [\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma] - 2^{-cn}$. From Lemma 4.5 it follows that $p_{t:1} = 0$. Given this, $U'_{w':0} = 0$, $U'_{w':1} = 1/n$ and $U'_{w'*} = 2^{-dn}$. So it follows from Lemma 4.4 that $p_{w':1} = p_{w'}$. Now, $U'_{z:1} = p_{w':1} = p_{w'} > \frac{1}{n}(1 - 2n^2\delta)$ (using Lemma 4.3), while $U'_{z:0} = 0$, $U'_{z*} = 2^{-dn}$. So from Lemma 4.4 we have that $p_{z:1} = p_z$, and therefore $\hat{p}(z:1) = 1$, which contradicts (37). Hence, (15) is satisfied.

□

PROOF OF LEMMA *4.9:* We analyze $\mathcal{G}_\vee$ first. Suppose that $\hat{p}(x:1), \hat{p}(y:1) \in \{0,1\}$ and $\hat{p}(x:1) \vee \hat{p}(y:1) = 1$. Then $U'_{w:1} = p_{x:1} + p_{y:1} \geq \frac{1}{n}(1 - 2n^2\delta)$ (using also Lemma 4.3). On the other hand, $U'_{w*} = 2^{-dn}$ and $U'_{w:0} = \frac{1}{n}2^{-c'n}$. Hence, from Lemma 4.4 we get $p_{w:0} = p_{w*} = 0$ and $p_{w:1} = p_w$. Given this, $U'_{z:0} = 0$, $U'_{z*} = 2^{-dn}$ and $U'_{z:1} = p_{w:1} = p_w \geq \frac{1}{n}(1 - 2n^2\delta)$ (using Lemma 4.3.) Hence, from Lemma 4.4 we get $p_{z:1} = p_z$, i.e $\hat{p}(z:1) = 1 = \hat{p}(x:1) \vee \hat{p}(y:1)$.

Now, suppose that $\hat{p}(x:1) \vee \hat{p}(y:1) = 0$. This implies $p_{x:1} = p_{y:1} = 0$. Hence, $U'_{w:1} = p_{x:1} + p_{y:1} = 0$, while $U'_{w*} = 2^{-dn}$ and $U'_{w:0} = \frac{1}{n}2^{-c'n}$. From Lemma 4.4 we get $p_{w:1} = 0$. Given this, $U'_{z:1} = 0$, while $U'_{z*} = 2^{-dn}$. Hence from Lemma 4.4 we get $p_{z:1} = 0$, i.e $\hat{p}(z:1) = 0 = \hat{p}(x:1) \vee \hat{p}(y:1)$.

We proceed to analyze $\mathcal{G}_\neg$. Suppose that $\hat{p}(x:1) = 1$, i.e. $p_{x:1} = p_x$. Then $U'_{w:0} = p_{x:1} = p_x \geq \frac{1}{n}(1 - 2n^2\delta)$ (using Lemma 4.3.) On the other hand, $U'_{w*} = 2^{-dn}$ and $U'_{w:1} =$

$\frac{1}{n} - p_{x:1} \le 2n\delta$ (using Lemma 4.3 again.) Hence, from Lemma 4.4 we get $p_{w:1} = p_{w*} = 0$ and $p_{w:0} = p_w$. Given this, $U'_{z:1} = 0$, $U'_{z*} = 2^{-dn}$ and $U'_{z:0} = p_{w:0} = p_w \ge \frac{1}{n}(1 - 2n^2\delta)$ (using Lemma 4.3.) Hence, from Lemma 4.4 we get $p_{z:1} = 0$, i.e $\hat{p}(z:1) = 0 = 1 - \hat{p}(x:1)$.

Suppose now $\hat{p}(x:1) = 0$, i.e. $p_{x:1} = 0$. Then $U'_{w:0} = p_{x:1} = 0$, $U'_{w*} = 2^{-dn}$ and $U'_{w:1} = \frac{1}{n} - p_{x:1} = 1/n$. Hence, from Lemma 4.4 we get $p_{w:0} = p_{w*} = 0$ and $p_{w:1} = p_w$. Given this, $U'_{z:0} = 0$, $U'_{z*} = 2^{-dn}$ and $U'_{z:1} = p_{w:1} = p_w \ge \frac{1}{n}(1 - 2n^2\delta)$ (using Lemma 4.3.) Hence, from Lemma 4.4 we have $p_{z:1} = 1$, i.e $\hat{p}(z:1) = 1 = 1 - \hat{p}(x:1)$. □

E. UPPER BOUNDS FOR WIN-LOSE GAMES

We already mentioned in the introduction that a polynomial-time algorithm for computing relative $\frac{1}{2}$-approximate Nash equilibria in bimatrix games with payoffs in $[0,1]$ was given in [FNS07]. Computing relative ϵ-Nash equilibria is more challenging, since we are not allowed to use in the support a strategy that is not approximately optimal. Here we present two easy upper bounds for relative 1-Nash equilibria in two-player games with payoffs in $[0,1]$ or in $\{0,-1\}$. The case $[0,1]$ is trivial as any pair of strategies (x,y) forms a relative 1-Nash equilibrium. For games with payoffs in $\{0,-1\}$ this is not the case. Indeed, we cannot see an immediate argument providing a relative 1-Nash equilibrium for these games. We suggest instead the following algorithm, which computes a relative 1-Nash equilibrium by approximating the given game with a constant-sum game. In spirit this is similar to the approach of [KS07].

(1) if there is a pair of strategies i, j such that $R_{ij} = C_{ij} = 0$, output strategies i and j; /* *this is clearly a pure Nash equilibrium* */
(2) otherwise, define the following (-1)-sum game (R', C') as follows:
— for all i, j such that $(R_{ij}, C_{ij}) \neq (-1,-1)$, set $R'_{ij} = R_{ij}$ and $C'_{ij} = C_{ij}$;
— for all i, j such that $(R_{ij}, C_{ij}) = (-1,-1)$, set $R'_{ij} = -1/2$ and $C'_{ij} = -1/2$.
(3) output any exact Nash equilibrium (x,y) of the game (R', C'); /**since* (R', C') *is constant sum, we can compute in polynomial-time an exact Nash equilibrium of this game** */

Let us argue now that, if a pair of strategies (x,y) is output in the last step of the algorithm, then this pair is indeed a relative 1-Nash equilibrium of the original game (R,C). We are only going to show the equilibrium conditions for the row player; a similar argument applies for the column player. Since (x,y) is an exact Nash equilibrium of the game (R', C') it must be that

$$\text{for all } i \text{ such that } x_i > 0: \quad e_i^{\mathrm{T}} R' y \ge e_{i'}^{\mathrm{T}} R' y, \text{ for all } i'. \tag{38}$$

We shall use (38) to argue that

$$\text{for all } i \text{ such that } x_i > 0: \quad e_i^{\mathrm{T}} R y \ge e_{i'}^{\mathrm{T}} R y - 1 \cdot |e_{i'}^{\mathrm{T}} R y|, \text{ for all } i'. \tag{39}$$

This is enough to justify that (x,y) is a relative 1-Nash equilibrium of the game (R,C). Notice that since $e_{i'}^{\mathrm{T}} R y \le 0$, (39) is equivalent to

$$\text{for all } i \text{ such that } x_i > 0: \quad e_i^{\mathrm{T}} R y \ge 2 e_{i'}^{\mathrm{T}} R y, \text{ for all } i'. \tag{40}$$

To justify (40), let us define the matrices R^α and R^β as follows. For all i, j:

$$R^\alpha_{ij} = \begin{cases} -\frac{1}{2}, & \text{if } R_{ij} = C_{ij} = -1 \\ 0, & \text{otherwise} \end{cases}$$

and

$$R^\beta = R - 2R^\alpha.$$

It is easy to see then that $R' = R^\alpha + R^\beta$, while $R = 2R^\alpha + R^\beta$. Now fix some i such that $x_i > 0$ and an arbitrary i'. (38) implies

$$2 \cdot e_i^{\mathrm{T}}(R^\alpha + R^\beta)y \geq 2 \cdot e_{i'}^{\mathrm{T}}(R^\alpha + R^\beta)y. \tag{41}$$

Since $e_i^{\mathrm{T}} R^\beta y \leq 0$, we have

$$e_i^{\mathrm{T}}(2R^\alpha + R^\beta)y \geq 2 \cdot e_i^{\mathrm{T}}(R^\alpha + R^\beta)y.$$

Similarly, because $e_{i'}^{\mathrm{T}} R^\alpha y \leq 0$,

$$2e_{i'}^{\mathrm{T}}(R^\alpha + R^\beta)y \geq 2 \cdot e_{i'}^{\mathrm{T}}(2R^\alpha + R^\beta)y.$$

Hence, (41), gives

$$e_i^{\mathrm{T}}(2R^\alpha + R^\beta)y \geq 2 \cdot e_{i'}^{\mathrm{T}}(2R^\alpha + R^\beta)y,$$

i.e.

$$e_i^{\mathrm{T}} Ry \geq 2 \cdot e_{i'}^{\mathrm{T}} Ry,$$

so (40) is satisfied. Hence, (x, y) is a 1-Nash equilibrium of the game (R, C).

REFERENCES

Timothy G. Abbott, Daniel Kane, and Paul Valiant. On the Complexity of Two-Player Win-Lose Games. In *the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2005.

Hartwig Bosse, Jaroslaw Byrka, and Evangelos Markakis. New Algorithms for Approximate Nash Equilibria in Bimatrix Games. In *the 3rd International Workshop on Internet and Network Economics (WINE)*, 2007.

Daniel Bernoulli. Exposition of a New Theory on the Measurement of Risk. *Comentarii Academiae Scientiarum Imperialis Petropolitanae*, 1738.

Xi Chen and Xiaotie Deng. 3-Nash is PPAD-Complete. *Electronic Colloquium on Computational Complexity (ECCC)*, (134), 2005.

Xi Chen and Xiaotie Deng. Settling the Complexity of Two-Player Nash Equilibrium. In *the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006.

Yang Cai and Constantinos Daskalakis. A Minmax Theorem for Multiplayer Zero-Sum Games. In *the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2011.

Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Computing Nash Equilibria: Approximation and Smoothed Complexity. In *the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006.

Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The Complexity of Computing a Nash Equilibrium. In *the 38th Annual ACM Symposium on Theory of Computing (STOC)*, 2006.

Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The Complexity of Computing a Nash Equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The Complexity of Computing a Nash Equilibrium. *Communications of the ACM*, 52(2):89–97, 2009.

Constantinos Daskalakis, Aranyak Mehta, and Christos H. Papadimitriou. A Note on Approximate Nash Equilibria. In *the 2nd International Workshop on Internet and Network Economics (WINE)*, 2006.

Constantinos Daskalakis, Aranyak Mehta, and Christos H. Papadimitriou. Progress in Approximate Nash Equilibria. In *8th ACM Conference on Electronic Commerce (EC)*, 2007.

Constantinos Daskalakis and Christos H. Papadimitriou. Three-Player Games Are Hard. *Electronic Colloquium on Computational Complexity (ECCC)*, (139), 2005.

Constantinos Daskalakis and Christos H. Papadimitriou. On a network generalization of the minmax theorem. In *the 36th International Colloquium on Automata, Languages and Programming (ICALP)*, 2009.

Constantinos Daskalakis and Christos H. Papadimitriou. On Oblivious PTAS's for Nash Equilibrium. In *the 41st Annual ACM Symposium on Theory of Computing (STOC)*, 2009.

Kousha Etessami and Mihalis Yannakakis. On the Complexity of Nash Equilibria and Other Fixed Points (Extended Abstract). In *the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2007.

Tomás Feder, Hamid Nazerzadeh, and Amin Saberi. Approximating Nash Equilibria Using Small-Support Strategies. In *the 8th ACM Conference on Electronic Commerce (EC)*, 2007.

Paul W. Goldberg and Christos H. Papadimitriou. Reducibility Among Equilibrium Problems. In *the 38th Annual ACM Symposium on Theory of Computing (STOC)*, 2006.

Itzhak Gilboa and Eitan Zemel. Nash and Correlated Equilibria: Some Complexity Considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.

Elad Hazan and Robert Krauthgamer. How Hard is it to Approximate the Best Nash Equilibrium? In *the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2009.

Spyros C. Kontogiannis, Panagiota N. Panagopoulou, and Paul G. Spirakis. Polynomial Algorithms for Approximating Nash Equilibria of Bimatrix Games. In *the 2nd International Workshop on Internet and Network Economics (WINE)*, 2006.

Spyros C. Kontogiannis and Paul G. Spirakis. Efficient Algorithms for Constant Well Supported Approximate Equilibria in Bimatrix Games. In *the 34th International Colloquium on Automata, Languages and Programming (ICALP)*, 2007.

Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. Playing Large Games Using Simple Strategies. In *the 4th ACM Conference on Electronic Commerce (EC)*, 2003.

John F. Nash. Equilibrium Points in n-Person Games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.

Christos H. Papadimitriou. On the Complexity of the Parity Argument and Other Inefficient Proofs of Existence. *Journal of Computer and System Sciences*, 48(3):498–532, 1994.

Herbert E. Scarf. The Approximation of Fixed Points of a Continuous Mapping. *SIAM Journal on Applied Mathematics*, 15(5):1328–1343, 1967.

Shang-Hua Teng. Private Communication. 2008.

Haralampos Tsaknakis and Paul G. Spirakis. An Optimization Approach for Approximate Nash Equilibria. In *the 3rd International Workshop on Internet and Network Economics (WINE)*, 2007.

Haralampos Tsaknakis and Paul G. Spirakis. Practical and Efficient Approximations of Nash Equilibria for Win-Lose Games Based on Graph Spectra. In *6th International Workshop on Internet and Network Economics (WINE)*, 2010.