

MIT Open Access Articles

On the complexity of approximating a nash equilibrium

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Constantinos Daskalakis. 2011. On the complexity of approximating a Nash equilibrium. In Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '11). SIAM 1498-1517. SIAM ©2011

As Published: <http://dl.acm.org/citation.cfm?id=2133036.2133153&coll=DL&dl=ACM&CFID=117975919&CFTOKEN=69611958>

Publisher: Society for Industrial and Applied Mathematics

Persistent URL: <http://hdl.handle.net/1721.1/73096>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



On the Complexity of Approximating a Nash Equilibrium

Constantinos Daskalakis*
EECS and CSAIL, MIT

Abstract

We show that computing a relative—that is, multiplicative as opposed to additive—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 the problem, since the appearance of the original results on the complexity of the Nash equilibrium [8, 5, 7]. Moreover, it provides an apparent—assuming that $\text{PPAD} \not\subseteq \text{TIME}(n^{O(\log n)})$ —dichotomy between the complexities of additive and relative notions of approximation, since for constant values of additive approximation a quasipolynomial-time algorithm is known [22]. Such a dichotomy does not arise for values of the approximation that scale with the size of the game, as both relative and additive approximations are PPAD-complete [7]. As a byproduct, our proof shows that the Lipton-Markakis-Mehta sampling lemma is not applicable to relative notions of constant approximation, answering in the negative direction a question posed to us by Shang-Hua Teng [26].

1 Introduction

In the wake of the complexity results for computing a Nash equilibrium [8, 5, 7], researchers undertook the important—and indeed very much algorithmic—task of understanding the complexity of approximate Nash equilibria. A positive outcome of this investigation would be useful for applications since it would provide algorithmic tools for computing approximate equilibria; but, most importantly, it would alleviate the negative implications of the aforementioned hardness results to the predictive power of the Nash equilibrium concept. Unfortunately, since the appearance of the original hardness results, and despite considerable effort in providing upper [20, 9, 10, 21, 15, 3, 27, 28] and lower [12, 18] bounds for the approximation problem, the approximation complexity of the Nash equilibrium has remained unknown. This paper obtains the first constant inapproximability results for the problem.

*Supported by a Sloan Foundation Fellowship and NSF CAREER Award CCF-0953960. Part of this work was done while the author was in Microsoft Research New England.

When it comes to approximation, the typical algorithmic approach is to look at *relative*, that is *multiplicative*, approximations to the optimum of an objective function. In a game, there are multiple objective functions, one for each player, called her *payoff function*. In a Nash equilibrium each player plays a randomized strategy that optimizes the expected value of her objective function given the strategies of her opponents. And, since the expected payoffs are linear functions of the players' strategies, to optimize her payoff a player needs to use in the support of her own strategy only pure strategies that achieve optimal expected payoff against the opponents' strategies.

Relaxing this requirement, a relative ϵ -Nash equilibrium is a collection of mixed strategies, one for each player of the game, so that no player uses in her support any pure strategy whose payoff fails to be within a relative error of ϵ from the best response payoff.¹ Clearly, in an ϵ -Nash equilibrium, the expected payoff of every player is within a relative error ϵ from her best response payoff. However, the latter is a strictly weaker requirement; 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 the two kinds of approximation the literature has converged to ϵ -approximate Nash equilibrium as the name for the latter, weaker kind of approximation.

Despite the long line of research on algorithms for approximate equilibria cited above, there is a single positive result for relative approximations due to Feder et al. [15], which provides a polynomial-time algorithm for relative 0.5-approximate Nash equilibria in 2-player games with payoffs in $[0, M]$, for all $M > 0$. On the other hand, the investigation of the absolute-error (i.e. additive-error) counterparts of the notions of approximate equilibrium defined above has been much more fruitful.² The additive notions of approximation are

¹Given this definition, this kind of approximation also goes by the name ϵ -well supported Nash equilibrium in the literature. We adopt the shorter name ϵ -Nash equilibrium for convenience.

²In the additive notions of approximation, it is required that the expected payoff from either the whole mixed strategy of a player or from everything in its support is within an ϵ absolute,

less common in algorithms, but they appear algorithmically more benign in this setting. Moreover, they naturally arise in designing simplicial approximation algorithms for the computation of equilibria, as the additive error is directly implied by the Lipschitz properties of Nash's function in the neighborhood of a Nash equilibrium [25]. For finite values of additive approximation, the best efficient algorithm to date computes a 0.34-approximate Nash equilibrium [27], and a 0.66-well-supported equilibrium [21], when the payoffs of the two-player game 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. Hence the performance of algorithms for additive approximate equilibria is typically compared after scaling the payoffs of the input game to lie in a unit-length interval; where this interval is located is irrelevant 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 settings. Imagine a play of some game in which a player is gaining an expected payoff of \$1M from her current strategy, but could improve her payoff to \$1.1M via some other strategy. Compare this situation to a play of the same game where the player's payoff is -\$50k and could become \$50k via a different strategy. 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 [2], the two situations could be very different, making the relative notion of approximation more appropriate; if the regret is perceived to be the same in these two situations, then the additive notion of approximation becomes more fitting.

In terms of computational complexity, additive and relative approximations have thus far enjoyed a similar fate. In two-player games, if the approximation guarantee scales inverse polynomially in the size of the game then both relative and additive approximations are PPAD-complete [7]. Hence, unless $\text{PPAD} \subseteq \text{P}$, there is no FPTAS for either additive or relative approximations. In the other direction, for both additive and relative notions, we have efficient algorithms for finite fixed values of ϵ . Even though progress in this frontier has stalled in the past couple of years, the hope for a polynomial-time approximation scheme—at least for additive approximations—ultimately stems from an older elegant result due to Lipton, Markakis

and Mehta [22]. This provides a quasi-polynomial-time algorithm for normalized bimatrix games (games with payoffs scaled in a unit-length interval), by establishing that, for any fixed ϵ , there exists an additive ϵ -approximate Nash equilibrium of support-size logarithmic in the total number of strategies.³ The LMM algorithm performs an exhaustive search over pairs of strategies with logarithmic support and can therefore also optimize some objective over the output equilibrium. This property of the algorithm has been exploited in recent lower bounds for the problem [18, 12], albeit these fall short from a quasi-polynomial-time lower bound for additive approximations. On the other hand, a quasi-polynomial-time algorithm is not known for relative approximations, and indeed this was posed to us as an open problem by Shang-Hua Teng [26].

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. *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 payoffs of both players are positive in every ϵ -Nash equilibrium of the game.*

Our result is the first inapproximability result for constant values of approximation to the Nash equilibrium problem. Moreover, unless $\text{PPAD} \subseteq \text{TIME}(n^{O(\log n)})$, it precludes a quasi-polynomial-time algorithm à la [22] for constant values of relative approximation.⁴ Under the same assumption, our result provides a dichotomy between the complexity of relative and additive notions of constant approximation. Such a dichotomy has not been shown before, since for approximation values that scale inverse polynomially in the size of the game the hardness results of [7] apply to both notions.

Observe that, if the absolute values of the game's payoffs lie in the set $[m, M]$, where $\frac{M}{m} < c$, for some constant c —call these games c -balanced, then the relative approximation problem can be reduced to the additive approximation problem in normalized games (that is games with payoffs in a unit-length interval) with a loss of $2c$ in the approximation guarantee (see Remark 2.1 and following discussion). Therefore, in view of [22] and unless $\text{PPAD} \subseteq \text{TIME}(n^{O(\log n)})$, we cannot hope to extend Theorem 1.1 to the special class of c -balanced games. On the other hand, our result may very well extend to the special class of games whose

³The argument also applies to the stronger notion of additive ϵ -well supported Nash equilibria [12].

⁴In fact, an LMM-style sampling lemma is precluded unconditionally from our proof, which constructs games whose relative ϵ -Nash equilibria have all linear support.

that is additive as opposed to multiplicative, error from the best response payoff.

payoff functions range in either $[0, +\infty)$ or $(-\infty, 0]$, but are not c -balanced for some constant c . We believe that this special class of games is also PPAD-complete for constant values of relative approximation and that it should be possible to remove the use of negative payoffs from our construction with similar, although more tedious arguments. As an indication, we note that in all approximate Nash equilibria of the games resulting from our construction, all players get positive payoff. Finally, we believe that our result is tight with regards to the range of values of ϵ . For $\epsilon = 1$, a trivial algorithm yields a relative 1-Nash equilibrium for games with payoffs in $\{0, 1\}$ (this is the class of win-lose games introduced in [1]), while for win-lose games with payoffs in $\{-1, 0\}$ we can obtain an interesting polynomial-time algorithm that goes through zero-sum games (we postpone the details for the full version). We believe that these upper bounds can be extended to payoffs in $[-1, 1]$.

Results for Polymatrix Games. To show Theorem 1.1, we prove as an intermediate result a similar (and somewhat stronger) lower bound for *graphical polymatrix games*, which in itself is significant. In a polymatrix game the players are nodes of a graph and participate in 2-player games with each of their neighbors, summing up the payoffs gained from each adjacent edge. These games always possess exact equilibria in rational numbers [14], their exact Nash equilibrium problem was shown to be PPAD-complete in [8, 14], an FPTAS was precluded by [7], and their zero-sum counterparts are poly-time solvable [13, 4]. We establish the following lower bound.

THEOREM 1.2. *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 deterministic strategy guarantees positive payoff to every player, regardless of the other players' choices; i.e., it remains true even if the minimax value of every player is positive.*

Another way to describe our theorem is this: While it is trivial for every player to guarantee positive payoff to himself via a deterministic 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 obtain Theorem 1.2, it is natural to try to follow the approach of [8] of reducing the generic PPAD-complete problem to a graphical polymatrix game. This was done in [8] by introducing the so called *game-gadgets*: these were small graphical 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, 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”. Unfortunately, for the construction of [8] to go through, the input-output relations of the gadgets need to be accurate to within an exponentially small additive error; and even the more efficient construction of [7] needs the approximation error to be inverse polynomial. Alas, if we consider ϵ -Nash equilibria with constant values of ϵ , the errors in the gadgets of [8] 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 in a series graph without feedback. But, can we design self-correcting gates if feedback is introduced after each operation? Indeed, our proof of Theorem 1.2 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 [8]) 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 1.2 (see Section 3.5).

The Grand Challenge. The construction outlined above, while non-obvious, is in the end rather intuitive. The real challenge to establish Theorem 1.1 lies in reducing the polymatrix games of Theorem 1.2 to two-player games. Those familiar with the hardness reductions for normal form games [17, 8, 5, 7, 14], will recognize the challenge. The “generalized matching pennies reduction” of a polymatrix game to a two-player game (more details on this shortly) is *not* approximation preserving, in that ϵ -Nash equilibria of the polymatrix game are reduced to $O(\frac{\epsilon}{n})$ -Nash equilibria of the 2-player game; as a consequence, even if the required accuracy ϵ in the polymatrix game is a constant, we still need inverse polynomial accuracy in the resulting two-player game. In fact, as we argue below, any matching pennies-style reduction is doomed to fail, if ϵ -Nash equilibria for constant values of relative approximation are

considered in the two-player game.⁵

We obtain Theorem 1.1 via a novel reduction, which in our opinion constitutes significant progress in PPAD-hardness proofs. The new reduction can obtain all results in [17, 8, 5, 7, 14], but is stronger in that it shaves a factor of n off the relative approximation guarantees. In particular, the new reduction is approximation preserving for relative approximations. Given the ubiquity of the matching pennies reduction in previous PPAD hardness proofs, we expect that our new tighter reduction will enable reductions in future research.

To explain a bit the challenge, there are two kinds of constraints that a reduction from multi-player games to two-player games needs to satisfy. The first makes sure that information about the strategies of all the nodes of the polymatrix game is reflected in the behavior of the two players of the bimatrix game. The second ensures 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 ϵ is a constant, the two requirements get coupled in ways that makes it hard to enforce both. This is why previous reductions take the approximation requirement in the bimatrix game to scale inverse polynomially in n ; in that regime the above semantics can indeed be decoupled. In our case, the use of constant approximations makes the construction and analysis extremely fragile. In a very 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 pure from these requirements in order to enforce there the equilibrium conditions of the polymatrix game. This is hard to implement and it is quite surprising that it is at all feasible. Indeed, all details in our construction are extremely finely chosen.

Overview of the Construction. We explain our approximation preserving reduction from polymatrix to bimatrix games by providing intuition about the inadequacy of existing technology. As mentioned above, all previous lower bounds for bimatrix games are based on *generalized matching-pennies* constructions. To reduce a bipartite graphical polymatrix game into a bimatrix game, such 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 the same color. Then two “lawyers” are introduced for the two color classes, whose purpose is to “represent” all the nodes in their color class. This is done by including in the strategy set of each lawyer a block of strategies cor-

responding to the strategy-set of every player in their color class, so that, if the two lawyers choose strategies corresponding to adjacent nodes of the graph, the lawyers get payoff equal to the payoff from the interaction on that edge. The hope is then that, in any Nash equilibrium of the lawyer-game, the marginal distributions of the lawyers’ strategies inside 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. To avoid this, a high-stakes matching pennies game is added to 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 inside the block they look at the payoffs of the underlying graphical game. This tie-breaking reflects the Nash equilibrium conditions of the graphical game.

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 graphical 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 graphical 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 visible. This is necessary to guarantee that the distribution of probability mass within each block is (almost) only determined by the payoffs of the graphical 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 and F.4. Threats that are similar in spirit to ours were used in an older NP-hardness proof of Gilboa and Zermel [16]. However, their construction is inadequate here as it could lead to a uniform equilibrium over the threat strategies, which cannot be mapped to an equilibrium

⁵In view of [22] additive approximations are unlikely to be PPAD-complete for constant values of ϵ . So the matching pennies—as well as any other construction—should fail in the additive case.

of the underlying polymatrix game. Indeed, it takes a lot of effort to avoid such occurrence of meaningless equilibria.

The final maneuver. 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 the evaluations of the gadgets used to build up the graphical game are correct to within very high accuracy; and this rescues the reduction.

2 Preliminaries

A *bimatrix game* has two players, called *row* and *column*, and m strategies, $1, \dots, m$, 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 $[m] := \{1, \dots, m\}$ and define the set of mixed strategies $\Delta^m := \{x \in \mathbb{R}_+^m, \sum_i x_i = 1\}$. If the row player randomizes according to the mixed strategy $x \in \Delta^m$ and the column player according to the strategy $y \in \Delta^m$, then the row player receives an expected payoff of $x^T R y$ and the column player an expected payoff of $x^T C y$.

A *Nash equilibrium* of the game is a pair of mixed strategies (x, y) , $x, y \in \Delta^m$, such that $x^T R y \geq x'^T R y$, for all $x' \in \Delta^m$, and $x^T C y \geq x^T C y'$, 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: ⁶

$$(2.1) \quad \text{for all } i \text{ with } x_i > 0: e_i^T R y \geq e_{i'}^T R y, \text{ for all } i';$$

$$(2.2) \quad \text{for all } j \text{ with } y_j > 0: x^T C e_j \geq x^T C e_{j'}, \text{ for all } j'.$$

That is, every strategy that the row player includes in the support of x must give him at least as large expected payoff as any other strategy, and similarly for the column player.

It is possible to define two kinds of approximate Nash equilibria, *additive* or *relative*, by relaxing, in the additive or multiplicative sense, the defining conditions of the Nash equilibrium. A pair of mixed strategies

⁶As always, e_i represents the unit vector along dimension i , i.e. pure strategy i .

(x, y) is called an *additive ϵ -approximate Nash equilibrium* if $x^T R y \geq x'^T R y - \epsilon$, for all $x' \in \Delta^m$, and $x^T C y \geq x^T C y' - \epsilon$, 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 2.1 and 2.2. 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

$$(2.3) \quad \text{for all } i \text{ with } x_i > 0: e_i^T R y \geq e_{i'}^T R y - \epsilon, \text{ for all } i';$$

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 know the following:

PROPOSITION 2.1. [7, 8] *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 .*

The relative notions of approximation are defined by multiplicative relaxations of the equilibrium conditions. We call a pair of mixed strategies (x, y) a *relative ϵ -approximate Nash equilibrium* if $x^T R y \geq x'^T R y - \epsilon \cdot |x'^T R y|$, for all $x' \in \Delta^m$, and $x^T C y \geq x^T C y' - \epsilon \cdot |x^T C y'|$, for all $y' \in \Delta^m$. That is, no player has more than a relative incentive of 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

$$(2.4) \quad \text{for all } i \text{ s.t. } x_i > 0: e_i^T R y \geq e_{i'}^T R y - \epsilon \cdot |e_{i'}^T R y|, \forall i';$$

and similarly for the column player. Condition (2.4) implies that the relative regret $\left| \frac{e_i^T R y - e_{i'}^T R y}{e_{i'}^T R y} \right|$ experienced by the row player for not replacing a strategy i in her support by another strategy i' with better payoff is at most ϵ . Notice that this remains meaningful even if R has negative entries. Clearly, a relative ϵ -Nash equilibrium is also a relative ϵ -approximate Nash equilibrium, but the opposite implication is not always true.

And what about the relation between the additive and the relative notions of approximation? The following is an easy observation based on the above definitions.

REMARK 2.1. Let $\mathcal{G} = (R, C)$ be a game whose payoff entries have absolute values 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.

As noted earlier, algorithms for additive approximations are usually compared after scaling the payoffs of the game to some unit-length interval. Where this interval lies is irrelevant since the additive approximations are shift invariant. So by shifting we can bring the payoffs to $[-1/2, 1/2]$.⁷ 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 ϵ -approximate Nash equilibrium. So the computation of additive ϵ approximations in normalized games can be reduced to the computation of $\epsilon/2$ relative approximations. But the opposite is not clear; in fact, given our main result and [22], it is impossible assuming $\text{PPAD} \not\subseteq \text{TIME}(n^{O(\log n)})$. However, if the ratio $\frac{u}{\ell} < c$, where u, ℓ are respectively the largest, smallest in absolute value entries of the game, then the computation of a relative ϵ -Nash equilibrium can be reduced to the computation of an additive $\frac{\epsilon}{2c}$ -Nash equilibrium of a normalized game.

Graphical Polymatrix Games. As mentioned in the introduction, we use in our proof a subclass of graphical games [19], called *graphical polymatrix games*. As in graphical games, the players are nodes of a graph $G = (V, E)$ and each node $v \in V$ has her own strategy set S_v and her own payoff function, which only depends on the strategies of the players in her neighborhood $\mathcal{N}(v)$ in G . The game is called a graphical polymatrix game if, moreover, for every $v \in V$ and for 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 $\{\alpha_{u:s_u}^{v:s_v}\}_{u \in \mathcal{N}(v) \setminus \{v\}, s_u \in S_u}$ and $\beta^{v:s_v}$ such that the expected payoff of v for playing pure strategy s_v is

$$(2.5) \quad \sum_{u \in \mathcal{N}(v) \setminus \{v\}, s_u \in S_u} \alpha_{u:s_u}^{v:s_v} p_{u:s_u} + \beta^{v:s_v},$$

where $p_{u:s_u}$ denotes the probability with which node u plays pure strategy s_u .

⁷Clearly, 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 scaling.

3 Hardness of Graphical Polymatrix Games

Our hardness result for graphical games is based on developing a new set of gadgets for performing arithmetic and boolean operations, and comparisons via graphical games. Given these gadgets we can follow the construction in [8] of putting together a large graphical game that solves, via its Nash equilibria, the generic PPAD-complete problem. The main challenge is that, since we are considering constant values of relative approximation, the gadgets developed in [8] introduce a constant error per operation. And, the construction in [8]—even the more careful construction in [7]—cannot accommodate such error. We go around this problem by introducing new gadgets that are very accurate despite the fact that we are considering constant values of approximation. Our gadgets are largely dependent on the *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 the difference if a certain threshold is exceeded. Based on this gadget we construct our arithmetic and comparison gadgets, given in Sections 3.2 and 3.3. And, with a different construction, we also get boolean gadgets in Section 3.4. Due to space limitations we only state the properties of our gadgets in the following sections and defer the details of their construction to Appendix B. Moreover, we only describe the “simple” versions of our gadgets. In Appendix B, we also present the more “sophisticated” versions, in which all the players have positive minimax values. These latter gadgets are denoted with a superscript of ‘+’.

3.1 Gap Amplification.

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}_{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}_{det} , the mixed strategies of the players satisfy

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

3.2 Arithmetic Operators. We use our gap amplification gadget \mathcal{G}_{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}_{det} to compare the inputs and the output of the arithmetic operator, magnify any devia-

tion, 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}_{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}_{lin} with a bipartite graph, two input players x and y , and one output player z , such that in any relative ϵ -Nash equilibrium*

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

Using \mathcal{G}_{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}_ζ 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* [8]) comparator gadget as follows.

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 $\mathcal{G}_>$*

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

3.4 Boolean Operators. In a different manner that does not require our gap amplification gadget we construct boolean operators. We only need to describe a game for **or** and **not**. Using these we can 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 1.2. We reduce an instance of the PPAD-complete APPROXIMATE CIRCUIT EVALUATION problem (see Appendix A) to a graphical polymatrix game, by replacing every gate of the given circuit with the corresponding gadget. The construction of our gadgets guarantees that a relative ϵ -Nash equilibrium of the 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 is in PPAD [14]. Details are given in Appendix C.

4 Hardness of Two-Player Games

4.1 Challenges. To show Theorem 1.1, we need to encode the bipartite graphical polymatrix game \mathcal{GG} , built up 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 1.2, 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 [8, 5]. But, that construction is not helpful in our setting, since it cannot accommodate constant values of ϵ as we discuss shortly. Before that, let us get our notation right.

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 as in (2.5). 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 referring to 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 $\rho : V_L \cup V_R \rightarrow \{0, 1\}$ 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 $(v : s)$ is intended to mean “vertex v plays strategy s ”, for $s = 0, 1$. 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 as follows (using the notation for the payoffs from (2.5)):

$$(4.6) \quad U_{\rho(v)}((v : s), (v' : s')) := \begin{cases} \alpha_{v':s'}^{v:s} + \beta^{v:s}, & \text{if } (v, v') \in E; \\ \beta^{v:s}, & \text{if } (v, v') \notin E. \end{cases}$$

In other words, if the players of the bimatrix game play $(v : s)$ and $(v' : s')$, they are given payoff equal to the payoff that the nodes v and v' get along the edge (v, v') , if they choose strategies s and s' and the edge (v, v') exists; and they always get the additive term $\beta^{v:s}$.

Observe then that, if we could magically guarantee that in any Nash equilibrium the players 0 and 1 randomize uniformly over 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 functions (4.6), the way player $\rho(v)$ distributes the probability mass of the block corresponding to v to the strategies $(v : 0)$ and $(v : 1)$ needs to respect the Nash equilibrium conditions at v . This goes through as long as the players randomize uniformly, or even close to uniformly, among their blocks of strategies. If they don't, then all bets are off . . .

To make sure that the players randomize uniformly over their blocks of strategies, the construction of [17, 8, 6, 11, 5, 7, 14] makes the players 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 approximately work. To be more precise, let us define two arbitrary permutations $\pi_L : V_L \rightarrow [n]$ and $\pi_R : V_R \rightarrow [n]$, and define $\pi : V_L \cup V_R \rightarrow [n]$ as $\pi(v) = \pi_L(v)$, if $v \in V_L$, and $\pi(v) = \pi_R(v)$, if $v \in V_R$. Given this, the matching pennies game is included in the construction by giving the following payoffs to the players

$$(4.7) \quad \tilde{U}_{\rho(v)}((v : s), (v' : s')) := U_{\rho(v)}((v : s), (v' : s')) + (-1)^{\rho(v)} \cdot M \cdot \mathbb{1}_{\pi(v)=\pi(v')},$$

where M is chosen to be *much larger* than the payoffs

of the graphical game. Notice that, if we ignored the graphical game payoffs from (4.7), 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 assign probability $1/n$ to each block. The same is approximately true if we do not ignore the payoffs coming from the graphical game, as long as M is chosen large enough to overwhelm these payoffs. Still, in this case every block receives roughly $1/n$ probability mass; and if ϵ is small enough (inverse polynomial in n) there may still be regret for not distributing that mass to the best strategies within the block. In particular, we can argue that, for every ϵ -Nash equilibrium of the bimatrix game, the marginal distributions of the blocks 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 the game description. But, it seems that an inverse polynomial value of ϵ is really needed. If ϵ is constant, then the additive notion of approximation cannot 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. Hence, we cannot argue anymore that the marginal distributions comprise an approximate equilibrium of the graphical game. If we consider relative ϵ -Nash equilibria, the different strategies inside a block *always* give payoffs that are within a relative ϵ from each other, for trivial reasons, since their payoff is overwhelmed by the high-stakes game. So again, the marginal distributions are not informative about the Nash equilibria of the graphical game. 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 fails.

To accommodate constant values of ϵ , we need a different approach. Our idea is roughly 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 [8, 5], these punishments (almost) disappear if the player does randomize uniformly over her blocks of strategies; and this is necessary to guarantee that at an ϵ -equilibrium of the bimatrix game the distribution of probability mass within each block is (almost) only determined by the payoffs of the graphical game, even when ϵ is 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 do 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 an 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 the polymatrix games. However, we *are not able* to establish this correspondence, i.e. we fail to show that the marginal distributions used by the players of the bimatrix game within their different blocks constitute an approximate Nash equilibrium of the underlying graphical game.⁸ But, we *can* show (see Lemma 4.5) that the marginal distributions define jointly a highly accurate (in the additive sense) evaluation of the circuit encoded by the graphical game; and this is enough to establish our PPAD-completeness result (completed in Section F.4 of the appendix).

4.2 Our Construction. We do the following modifications to the game defined by (4.6):

- 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 $\{\text{bad}_v\}_{v \in V_L}$ 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 $\{\text{bad}_v\}_{v \in V_R}$ in order to force player 1 to randomize uniformly among her blocks of strategies.

⁸While we can show that the un-normalized marginals satisfy the approximate equilibrium conditions of the polymatrix game, the fragility of relative approximations prevents us from showing that the normalized marginals do, despite the fact that the normalization factors are equal to within an exponential error.

- The payoff functions $\widehat{U}_0(\cdot; \cdot)$ and $\widehat{U}_1(\cdot; \cdot)$ of players 0 and 1 respectively are defined in Figure 4 of the appendix, for some H, U and d to be decided shortly. The reader can study the definition of the functions in detail, however it is much easier to think of our game in terms of the expected payoffs that the players receive for playing different pure strategies as follows

$$(4.8) \quad \mathcal{E} \left(\widehat{U}_{p:v^*} \right) = -U \cdot p_{\text{bad}_v} + 2^{-dn};$$

$$(4.9) \quad \mathcal{E} \left(\widehat{U}_{p:(v:s)} \right) = -U \cdot p_{\text{bad}_v} + \sum_{(v,v') \in E} \sum_{s'=0,1} \alpha_{v':s'}^{v:s} \cdot p_{v':s'} + \frac{1}{n} \beta^{v:s};$$

$$(4.10) \quad \mathcal{E} \left(\widehat{U}_{p:\text{bad}_v} \right) = H \cdot (p_v - 1/n).$$

In the above, we denote by $\mathcal{E} \left(\widehat{U}_{p:v^*} \right)$, $\mathcal{E} \left(\widehat{U}_{p:(v:s)} \right)$ and $\mathcal{E} \left(\widehat{U}_{p:\text{bad}_v} \right)$ 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 by which player $\rho(v)$ plays strategies $(v : 0)$, $(v : 1)$ and (v^*) , and by p_{bad_v} the probability by which player $1 - \rho(v)$ chooses strategy bad_v . Finally, we let $p_v = p_{v:0} + p_{v:1} + p_{v^*}$.

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 just scale all the utilities down by an appropriate factor 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 $\rho(v)$ lose 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 also becomes handy later. Observe that Lemma 4.1 is not sufficient to imply Lemma 4.3 directly, since apart from 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 G . For every v , we define

$$U'_{v*} := 2^{-dn}; \quad \text{and}$$

$$U'_{(v:s)} := \sum_{(v,v') \in E} \sum_{s'=0,1} \alpha_{v':s'}^{v:s} \cdot p_{v':s'} + \frac{1}{n} \beta^{v:s}, \quad \text{for } s = 0, 1.$$

$$\text{so that } \mathcal{E}(\widehat{U}_{p:v*}) = -U \cdot p_{\text{bad}_v} + U'_{v*};$$

$$\mathcal{E}(\widehat{U}_{p:(v:s)}) = -U \cdot p_{\text{bad}_v} + U'_{(v:s)}, \quad \text{for } s = 0, 1.$$

In the appendix we show the following.

LEMMA 4.4. *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 $\sigma \in \{v, (v:0), (v:1)\}$,*

$$(4.11) \quad U'_\sigma < (1 - \epsilon)U'_{\sigma_{\max}} \quad \Rightarrow \quad p_\sigma = 0.$$

Notice the subtlety in Condition (4.11). If we replace U'_σ and $U'_{\sigma_{\max}}$ with $\mathcal{E}(\widehat{U}_{p:\sigma})$ and $\mathcal{E}(\widehat{U}_{p:\sigma_{\max}})$, 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 $-U \cdot p_{\text{bad}_v}$ from $\mathcal{E}(\widehat{U}_{p:\sigma})$ and $\mathcal{E}(\widehat{U}_{p:\sigma_{\max}})$ and maintain Condition (4.11), 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 $U \cdot p_{\text{bad}_v}$; 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.

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{GG} (in fact, of the game \mathcal{GG}^+ with three strategies $0, 1, *$ per player—see the proof of Theorem 1.2). 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{GG}^+ . 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 the relative approximations makes this approach problematic. Indeed, we fail to establish that after the normalization the equilibrium conditions of \mathcal{GG}^+ are satisfied.

The Final Maneuver. Rather than worrying about the ϵ -Nash equilibrium conditions of \mathcal{GG}^+ , we argue that we can obtain a highly accurate evaluation of the circuit encoded by \mathcal{GG}^+ . We consider the following transformation, which merges the strategies $(v:0)$ and $v*$:

$$(4.12) \quad \hat{p}(v:1) := \frac{p_{v:1}}{p_v}; \quad \hat{p}(v:0) := \frac{p_{v:0} + p_{v*}}{p_v}.$$

We argue that the normalized values \hat{p} 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}_{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}_{det} with $\alpha, \beta, \gamma \in [-1, 1]$, w is the intermediate node of \mathcal{G}_{det} , and t is the output node. Then the values \hat{p} obtained from a relative ϵ -Nash equilibrium of the bimatrix game*

as shown above satisfy

$$(4.13) \quad \hat{p}(z : 1) - [\alpha\hat{p}(x : 1) + \beta\hat{p}(y : 1) + \gamma] \geq 2^{-cn} \\ \Rightarrow \hat{p}(t : 1) = 1;$$

$$(4.14) \quad \hat{p}(z : 1) - [\alpha\hat{p}(x : 1) + \beta\hat{p}(y : 1) + \gamma] \leq -2^{-cn} \\ \Rightarrow \hat{p}(t : 1) = 0.$$

Given Lemma 4.5, we immediately obtain that the output of all comparator gadgets is highly accurate (Corollary F.1). Via similar arguments we can show that the gadget \mathcal{G}_{lin} (Lemma F.1)—and hence all the arithmetic gadgets (Corollary F.2)—are highly accurate, and also that the boolean gadgets are accurate (Lemma F.2). It follows 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 [8] (see Appendix A). On the other hand, finding an exact Nash equilibrium of a bimatrix game is in PPAD [24], hence finding a relative ϵ -Nash equilibrium is also in PPAD. To complete the proof of Theorem 1.1, note that, by virtue of Lemma 4.2, all players have positive payoffs in every relative ϵ -Nash equilibrium of the game. See Section F.4 of the appendix for the detailed proof of Theorem 1.1.

References

- [1] T. G. Abbott, D. Kane and P. Valiant. On the Complexity of Two-Player Win-Lose Games. *FOCS*, 2005.
- [2] D. Bernoulli. Exposition of a new theory on the measurement of risk. *Comentarii Academiae Scientiarum Imperialis Petropolitanae*, 1738.
- [3] H. Bosse, J. Byrka and E. Markakis. New Algorithms for Approximate Nash Equilibria in Bimatrix Games. *WINE*, 2007.
- [4] Y. Cai and C. Daskalakis. A Minmax Theorem for Multiplayer Zero-Sum Games. *Manuscript*, 2010.
- [5] X. Chen and X. Deng. Settling the Complexity of Two-Player Nash Equilibrium. *FOCS*, 2006.
- [6] X. Chen and X. Deng. 3-NASH is PPAD-Complete. *ECCC Report*, 2005.
- [7] X. Chen, X. Deng and S. Teng. Computing Nash Equilibria: Approximation and Smoothed Complexity. *FOCS*, 2006.
- [8] C. Daskalakis, P. W. Goldberg and C. H. Papadimitriou. The Complexity of Computing a Nash Equilibrium. *STOC*, 2006.
- [9] C. Daskalakis, A. Mehta, and C. H. Papadimitriou. A Note on Approximate Nash Equilibria. *WINE*, 2006.
- [10] C. Daskalakis, A. Mehta, and C. H. Papadimitriou. Progress in Approximate Nash Equilibria. *EC*, 2007.
- [11] C. Daskalakis and C. H. Papadimitriou. Three-Player Games Are Hard. *ECCC Report*, 2005.
- [12] C. Daskalakis and C. H. Papadimitriou. On Oblivious PTAS's for Nash Equilibrium. *STOC*, 2009.
- [13] C. Daskalakis and C. H. Papadimitriou. On a Network Generalization of the Minmax Theorem. *ICALP*, 2009.
- [14] K. Etessami and M. Yannakakis. On the Complexity of Nash Equilibria and Other Fixed Points (Extended Abstract). *FOCS*, 2007.
- [15] T. Feder, H. Nazerzadeh and A. Saberi. Approximating Nash Equilibria Using Small-Support Strategies. *EC*, 2007.
- [16] I. Gilboa and E. Zemel. Nash and Correlated Equilibria: Some Complexity Considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.
- [17] P. W. Goldberg, C. H. Papadimitriou. Reducibility among equilibrium problems. *STOC*, 2006.
- [18] E. Hazan and R. Krauthgamer. How hard is it to approximate the best Nash equilibrium? *SODA*, 2009.
- [19] M. J. Kearns, M. L. Littman and S. P. Singh. Graphical Models for Game Theory. *UAI*, 2001.
- [20] S. C. Kontogiannis, P. N. Panagopoulou and P. G. Spirakis. Polynomial Algorithms for Approximating Nash Equilibria of Bimatrix Games. *WINE*, 2006.
- [21] S. C. Kontogiannis and P. G. Spirakis. Efficient Algorithms for Constant Well Supported Approximate Equilibria in Bimatrix Games. *ICALP*, 2007.
- [22] R. J. Lipton, E. Markakis and A. Mehta. Playing Large Games Using Simple Strategies. *EC*, 2003.
- [23] J. Nash. Non-Cooperative Games. *Annals of Mathematics*, 54(2):286–295, 1951.
- [24] C. 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.
- [25] H. E. Scarf. The Approximation of Fixed Points of a Continuous Mapping. *SIAM Journal on Applied Mathematics*, 15(5):1328–1343, 1967.
- [26] S. Teng. Private Communication. September 2008.
- [27] H. Tsaknakis and P. G. Spirakis. An Optimization Approach for Approximate Nash Equilibria. *WINE*, 2007.
- [28] H. Tsaknakis and P. G. Spirakis. A Graph Spectral Approach for Computing Approximate Nash Equilibria. *CoRR* abs/0909.4686, 2009.

A The Approximate Circuit Evaluation Problem

We define the APPROXIMATE CIRCUIT EVALUATION problem, which was shown to be PPAD-complete in [8]. Our definition is based on the notions of the *generalized circuit* and the *approximate circuit evaluation*, given in Definitions A.1 and A.2 below. All these definitions were implicit in [8] and were made explicit in [5].

DEFINITION A.1. (GENERALIZED CIRCUIT) *A circuit \mathcal{C} is called a generalized circuit if it is built up using:*

- *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 C_- one of the input nodes is designated to be the “positive” input); the scale by ζ gate, $C_{\times\zeta}$, has one input and one output node, and the set equal to ζ , C_ζ , gate has one output node;

- comparison gates: the comparison gate, $C_{>}$, has two input nodes (one of which is designated to be the “positive” input) and one output node;
- boolean gates: the OR gate, C_\vee , has two input nodes and one output node, and the NOT gate, C_- , has one input and one output node.

DEFINITION A.2. (APPROXIMATE CIRCUIT EVALUATION)

Given a generalized circuit 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

- C_+ : if the input nodes have values x, y and the output node has value z then

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

- 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

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

- $C_{\times\zeta}$: if the input node has value x and the output node value z then $z = \min\{1, \zeta \cdot x\} \pm 2^{-cn}$;

- C_ζ : if the output node has value z then $z = \min\{1, \zeta\} \pm 2^{-cn}$;

- $C_{>}$: 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

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

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

- C_\vee : the values x, y of the input nodes and the value z of the output node satisfy:

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

- C_- : 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 A.3. (APPROXIMATE CIRCUIT EVALUATION Problem) Given a generalized circuit C and a constant c , find an approximate evaluation of the circuit C with accuracy 2^{-cn} .

B Gadgets of Section 3

LEMMA B.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}_{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}_{det} , the mixed strategies of the players satisfy the following

$$(2.15) \quad \begin{aligned} p(z : 1) - [\alpha p(x : 1) + \beta p(y : 1) + \gamma] &\geq 2^{-cn} \\ \Rightarrow p(t : 1) &= 1; \end{aligned}$$

$$(2.16) \quad \begin{aligned} p(z : 1) - [\alpha p(x : 1) + \beta p(y : 1) + \gamma] &\leq -2^{-cn} \\ \Rightarrow p(t : 1) &= 0; \end{aligned}$$

- there exists a graphical polymatrix game $\mathcal{G}_{\text{det}}^+$ with the same characteristics as \mathcal{G}_{det} , except that every player has three strategies 0, 1, and *, and such that Properties (2.15) and (2.16) 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.




Figure 1: The detector gadgets \mathcal{G}_{det} and $\mathcal{G}_{\text{det}}^+$.

Proof of Lemma B.1: The graphical structure of the games \mathcal{G}_{det} and $\mathcal{G}_{\text{det}}^+$ is shown in Figure 1, where the direction of the edges denotes direct payoff dependence. The construction of the games \mathcal{G}_{det} and $\mathcal{G}_{\text{det}}^+$ is similar, so we are only going to describe the construction of $\mathcal{G}_{\text{det}}^+$. A trivial adaptation of this construction—by just removing all the * strategies—gives the construction of \mathcal{G}_{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 (2.15) and (2.16) 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 (2.15) 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 (2.15) is satisfied.

To show (2.16), 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 $\epsilon < 1$). \square

LEMMA B.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}_{lin} with a bipartite graph, two input players x and y , and one*

output player z , such that in any relative ϵ -Nash equilibrium

$$(2.17)$$

$$p(z : 1) \geq$$

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

$$(2.18)$$

$$p(z : 1) \leq$$

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

- *there also exists a graphical polymatrix game $\mathcal{G}_{\text{lin}}^+$ with the same characteristics as \mathcal{G}_{lin} , except that every player has three strategies 0, 1, and $*$, and such that properties (2.17) and (2.18) 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.*

Proof of Lemma B.2: \mathcal{G}_{lin} and $\mathcal{G}_{\text{lin}}^+$ have the graphical structure shown in Figure 2. They are obtained by adding feedback to the gadgets \mathcal{G}_{det} and $\mathcal{G}_{\text{det}}^+$ respectively through a new player w' who is introduced to keep 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 $\mathcal{G}_{\text{lin}}^+$, and the description of \mathcal{G}_{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




Figure 2: The affine operator gadgets \mathcal{G}_{lin} and $\mathcal{G}_{\text{lin}}^+$.

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:

- $u(z : *) = 2^{-c'n}$;

- $u(z : 0) = \mathbb{1}_{w':0}$;
- $u(z : 1) = \mathbb{1}_{w':1}$.

Now, we proceed to argue that (2.17) and (2.18) 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, (2.17) is satisfied. To show (2.18), suppose for a contradiction that

$$(2.19) \quad p(z : 1) > \min\{1, \max\{0, \alpha p(x : 1) + \beta p(y : 1) + \gamma\}\} + 2^{-cn}.$$

The above implies, $p(z : 1) > [\alpha p(x : 1) + \beta p(y : 1) + \gamma] + 2^{-cn}$; hence, by Lemma B.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 (2.19).

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

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

Suppose for a contradiction that

$$(2.20) \quad p(z : 1) > [\alpha p(x : 1) + \beta p(y : 1) + \gamma] + 2^{-cn}.$$

From Lemma B.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 (2.20), and therefore (2.18) is satisfied. To argue that (2.17) is satisfied, suppose for a contradiction that

$$(2.21) \quad p(z : 1) < [\alpha p(x : 1) + \beta p(y : 1) + \gamma] - 2^{-cn}.$$

From Lemma B.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 (2.21). Hence, (2.17) is satisfied.

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

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

So, automatically (2.18) is satisfied. To show (2.17), suppose for a contradiction that

$$(2.22) \quad p(z : 1) < \max\{0, \min\{1, \alpha p(x : 1) + \beta p(y : 1) + \gamma\}\} - 2^{-cn}.$$

The above implies, $p(z : 1) < [\alpha p(x : 1) + \beta p(y : 1) + \gamma] - 2^{-cn}$. From Lemma B.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 (2.22). Hence, (2.17) is satisfied.

□

LEMMA B.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}_ζ 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}_ζ^+ with the same characteristics as the graphical games \mathcal{G}_+ , \mathcal{G}_- , $\mathcal{G}_{\times\zeta}$, \mathcal{G}_ζ , 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 B.3: All the gadgets are obtained from \mathcal{G}_{lin} and $\mathcal{G}_{\text{lin}}^+$ 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}_ζ and \mathcal{G}_ζ^+ set $\alpha = \beta = 0$ and $\gamma = \zeta$. \square

LEMMA B.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 $\mathcal{G}_>$

$$(2.23) \quad p(z : 1) - p(x : 1) \geq 2^{-cn} \quad \Rightarrow \quad p(t : 1) = 1;$$

$$(2.24) \quad p(z : 1) - p(x : 1) \leq -2^{-cn} \quad \Rightarrow \quad p(t : 1) = 0;$$

- 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 B.4: $\mathcal{G}_>$ and $\mathcal{G}_>^+$ are obtained from \mathcal{G}_{det} and $\mathcal{G}_{\text{det}}^+$ respectively, by setting $\alpha = 1$, $\beta = \gamma = 0$. \square

LEMMA B.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

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

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

- 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 B.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 player w




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

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, $p(z : 1) = 1 - p(x : 1)$. \square

C Proof of Theorem 1.2

Proof of Theorem 1.2: From [14], 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 the problem. To do this, we reduce from the APPROXIMATE CIRCUIT EVALUATION problem (implicit in [8], explicit in [5]) which is roughly the following (see Appendix A for a detailed definition): Given a circuit consisting of the gates *plus*, *minus*, *scale by a constant*, *set equal to a constant*, *compare*, *or*, and *not*, find values for the nodes of the circuit satisfying the input-output relations of the gates to within an additive error of 2^{-cn} —analogously to the input-output relations of the gadgets specified in Lemmas 3.3, 3.4, and 3.5. It was shown in [8] that there is a constant c such that the APPROXIMATE CIRCUIT EVALUATION problem is PPAD-complete. But, given any circuit, it is easy to set up, using the gadgets \mathcal{G}_+ , \mathcal{G}_- , $\mathcal{G}_{\times\zeta}$, \mathcal{G}_ζ , $\mathcal{G}_>$, \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. To make sure that every node in our graphical game has positive minimax value we can use in our construction the sophisticated versions \mathcal{G}_+^+ , \mathcal{G}_-^+ , $\mathcal{G}_{\times\zeta}^+$, \mathcal{G}_ζ^+ , $\mathcal{G}_>^+$, \mathcal{G}_\vee^+ , \mathcal{G}_\neg^+ of our gadgets given in Appendix B. Call the resulting game \mathcal{GG}^+ for future reference. \square

D The Bimatrix Game in Our Construction

See Figure 4.

E Choosing the Right Constants

See Figure 5

F Omitted Details from Section 4

F.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

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

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 $\text{bad}_{v,v} \in S$ 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. Also, player $1 - \rho(v_a)$ will not play any strategy of the form $\text{bad}_{v,v}$, $v \notin S$, because by the definition of the set S she is better off playing strategy bad_{v_a} by more than a relative ϵ . Moreover, $|S| < n$, since $v_b \notin S$. Hence, there must be some $v_c \in S$, such that $p_{\text{bad}_{v_c}} > 1/n$.

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 $p_{\text{bad}_{v_b}} = 0$), 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 $p_{\text{bad}_{v_c}} > 1/n$ 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 $-U \cdot p_{\text{bad}_v} \geq -U \cdot n \cdot \phi_a$. Since $U \cdot n(1 + \epsilon) < H$ it follows that $-U \cdot n \cdot \phi_a(1 + \epsilon) > -H\phi_a$. So

$$(4.25) \quad \widehat{U}_p(\sigma; \sigma') := \begin{cases} -U + 2^{-dn}, & \text{if } \sigma = v^*, \sigma' = \text{bad}_v, \rho(v) = p; \\ 2^{-dn}, & \text{if } \sigma = v^*, \sigma' \neq \text{bad}_v, \rho(v) = p; \\ U_p(\sigma; \sigma')|_{\beta^{v:s} \leftarrow \frac{1}{n}\beta^{v:s}}, & \text{if } \sigma = (v:s), \sigma' = (v':s'), \rho(v) = 1 - \rho(v') = p; \\ -U + \frac{1}{n}\beta^{v:s}, & \text{if } \sigma = (v:s), \sigma' = \text{bad}_v, \rho(v) = p; \\ \frac{1}{n}\beta^{v:s}, & \text{if } \sigma = (v:s), \sigma' = \text{bad}_{v'}, \rho(v) = \rho(v') = p; \\ H \cdot (1 - \frac{1}{n}), & \text{if } \sigma = \text{bad}_v, \sigma' \in \{(v:0), (v:1), v^*\}, \rho(v) = 1 - p; \\ H \cdot (-\frac{1}{n}), & \text{if } \sigma = \text{bad}_v, \sigma' \notin \{(v:0), (v:1), v^*\}, \rho(v) = 1 - p; \end{cases}$$

with $U_p(\sigma; \sigma')|_{\beta^{v:s} \leftarrow \frac{1}{n}\beta^{v:s}}$, we denote that the term $\beta^{v:s}$ in $U_p((v:s); (v':s'))$ is scaled down by a factor of n ; this is done for technical reasons that we discuss later.

Figure 4: The payoffs of the bimatrix game in our construction.

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

Figure 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.

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 $p_{\text{bad}_v} = 0$, 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.

□

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

$$(6.26) \quad p_{\text{bad}_{v_b}} \geq \frac{1}{n}(2n - (n - 1))\delta > \delta.$$

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 $p_{\text{bad}_{v_b}} > \delta$ 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 $p_{\text{bad}_{v_b}} = 0$. This contradicts (6.26). □

F.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

$$(6.27) \quad -U \cdot p_{\text{bad}_v} + U'_{\sigma_{\max}} > 0.$$

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

- $-U \cdot p_{\text{bad}_v} + U'_\sigma < 0$: This implies that the expected payoff to player $\rho(v)$ for playing strategy σ is negative, while the expected payoff from strategy σ_{\max} is positive (see Equation (6.27)), so the implication is true.
- $-U \cdot p_{\text{bad}_v} + U'_\sigma \geq 0$: We have

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

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

□

F.3 Approximate Circuit-Evaluations from Relative Equilibria of the Bimatrix Game. We argue first that the values \hat{p} satisfy the input-output relations of the gadget \mathcal{G}_{det} (Lemma 4.5).

Proof of Lemma 4.5: We show (4.13) first. Suppose $\hat{p}(z : 1) - [\alpha\hat{p}(x : 1) + \beta\hat{p}(y : 1) + \gamma] \geq 2^{-cn}$. This implies the following

$$(6.28) \quad \frac{p_{z:1}}{p_z} - \left[\alpha \cdot \frac{p_{x:1}}{p_x} + \beta \cdot \frac{p_{y:1}}{p_y} + \gamma \right] \geq 2^{-cn};$$

Now we show

CLAIM F.1. *The above imply:*

$$(6.29) \quad p_{z:1} - \left[\alpha \cdot p_{x:1} + \beta \cdot p_{y:1} + \frac{\gamma}{n} \right] \geq \frac{2^{-cn}}{2n}.$$

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

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

This is a contradiction to (6.28).

Given (6.29) we have

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

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

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

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

We show (4.14) similarly. Suppose $\hat{p}(z : 1) - [\alpha\hat{p}(x : 1) + \beta\hat{p}(y : 1) + \gamma] \leq -2^{-cn}$. This implies the following

$$(6.30) \quad \frac{p_{z:1}}{p_z} - \left[\alpha \cdot \frac{p_{x:1}}{p_x} + \beta \cdot \frac{p_{y:1}}{p_y} + \gamma \right] \leq -2^{-cn};$$

Now we show

CLAIM F.2. *The above imply:*

$$(6.31) \quad p_{z:1} - \left[\alpha \cdot p_{x:1} + \beta \cdot p_{y:1} + \frac{\gamma}{n} \right] \leq -\frac{2^{-cn}}{2n}.$$

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

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

This is a contradiction to (6.30).

Given (6.31) 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 $\hat{p}(t:1) = 0$. \square

We immediately obtain from Lemma 4.5 that the input-output relations of the comparator gadget are fine.

COROLLARY F.1. *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*

$$\begin{aligned}
\hat{p}(z:1) \geq \hat{p}(x:1) + 2^{-cn} & \Rightarrow \hat{p}(t:1) = 1; \\
\hat{p}(z:1) \leq \hat{p}(x:1) - 2^{-cn} & \Rightarrow \hat{p}(t:1) = 0.
\end{aligned}$$

Now, we study the gadget \mathcal{G}_{lin} .

LEMMA F.1. *Suppose $x, y, z, w, w', t \in V_L \cup V_R$, so that x, y are inputs to the game \mathcal{G}_{lin} with parameters α, β and γ , z is the output player, and w, w', t' are the intermediate nodes (as in Lemma B.2, Figure 2). Then*

$$(6.32) \quad \hat{p}(z:1) \geq \max\{0, \min\{1, \alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma\}\} - 2^{-cn};$$

$$(6.33) \quad \hat{p}(z:1) \leq \min\{1, \max\{0, \alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma\}\} + 2^{-cn}.$$

Proof of Lemma F.1: The proof proceeds by considering the following cases as in the proof of Lemma B.2:

- $[\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma] \leq 0$: In this case we have
$$\begin{aligned} \max\{0, \min\{1, \alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma\}\} &= 0, \\ \min\{1, \max\{0, \alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma\}\} &= 0. \end{aligned}$$

So, clearly, (6.32) is satisfied. To show (6.33), suppose for a contradiction that

$$(6.34) \quad \hat{p}(z:1) > \min\{1, \max\{0, \alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma\}\} + 2^{-cn}.$$

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 (6.34).

- $0 \leq [\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma] \leq 1$: 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

$$(6.35) \quad \hat{p}(z:1) > [\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma] + 2^{-cn}.$$

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 (6.35). Hence, (6.33) is satisfied.

To show (6.32), suppose for a contradiction that

$$(6.36) \quad \hat{p}(z:1) < [\alpha\hat{p}(x:1) + \beta\hat{p}(y:1) + \gamma] - 2^{-cn}.$$

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 (6.36). Hence, (6.32) 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 (6.33) is satisfied. To show (6.32), suppose for a contradiction that

$$(6.37) \quad \hat{p}(z : 1) < \max\{0, \min\{1, \alpha\hat{p}(x : 1) + \beta\hat{p}(y : 1) + \gamma\}\} - 2^{-cn}.$$

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 (6.37). Hence, (6.32) is satisfied.

□

Given Lemma F.1, we obtain that all arithmetic gadgets are highly accurate.

COROLLARY F.2. *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}_ζ , then $\hat{p}(z : 1) = \min\{1, \zeta\} \pm 2^{-cn}$.

Finally, we analyze the boolean operators.

LEMMA F.2. *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}_\wedge (as in Figure 3, Lemma B.5). 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}_\wedge satisfies $\hat{p}(z : 1) = 1 - \hat{p}(x : 1)$;

Proof of Lemma F.2: 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}_\wedge . 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 $U'_{w:1} = \frac{1}{n} - p_{x:1} \leq 2n\delta$ (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 $U'_{z:0} = p_{w:0} = p_w \geq \frac{1}{n}(1 - 2n^2\delta)$ (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 $U'_{z:1} = p_{w:1} = p_w \geq \frac{1}{n}(1 - 2n^2\delta)$ (using Lemma 4.3). Hence, from Lemma 4.4 we have $p_{z:1} = 1$, i.e. $\hat{p}(z : 1) = 1 = 1 - \hat{p}(x : 1)$. □

F.4 Completing the Proof of Theorem 1.1.

From [24], it follows 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 justify the PPAD-hardness of the problem. Given a pair (\mathcal{C}, c) , where \mathcal{C} is a generalized circuit (see Definition A.1) and c a positive constant (this pair is an instance of the APPROXIMATE CIRCUIT EVALUATION problem defined in Appendix A), we construct a bipartite graphical polymatrix game \mathcal{GG} using the reduction in the proof of Theorem 1.2. 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 $\mathcal{G}_+, \mathcal{G}_-, \mathcal{G}_{\times\zeta}, \mathcal{G}_\zeta, \mathcal{G}_>, \mathcal{G}_\vee, \mathcal{G}_\wedge$. 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 (4.12)) 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 $\mathcal{G}_+, \mathcal{G}_-, \mathcal{G}_{\times\zeta}, \mathcal{G}_\zeta, \mathcal{G}_>, \mathcal{G}_\vee, \mathcal{G}_\wedge$ (as specified by Corollaries F.1 and F.2, and Lemma F.2). These values comprise then an approximate evaluation of the circuit \mathcal{C} . Since the APPROXIMATE CIRCUIT EVALUATION problem is PPAD-complete it follows that finding a relative ϵ -Nash equilibrium is also PPAD-complete.