A Direct Reduction from k-Player to 2-Player Approximate Nash Equilibrium*

Uriel Feige †
Weizmann Institute

Inbal Talgam ‡
Weizmann Institute

Abstract

We present a direct reduction from k-player games to 2-player games that preserves approximate Nash equilibrium. Previously, the computational equivalence of computing approximate Nash equilibrium in k-player and 2-player games was established via an indirect reduction. This included a sequence of works defining the complexity class PPAD, identifying complete problems for this class, showing that computing approximate Nash equilibrium for k-player games is in PPAD, and reducing a PPAD-complete problem to computing approximate Nash equilibrium for 2-player games. Our direct reduction makes no use of the concept of PPAD, thus eliminating some of the difficulties involved in following the known indirect reduction.

1 Introduction

This manuscript addresses the computation of Nash equilibrium for games represented in normal form. It is known that for 2-player games this problem is PPAD-complete [CD06b], and for k players it is in PSPACE [EY07]. Moreover, for sufficiently small ϵ, computing ϵ-well-supported Nash equilibrium for 2-player games remains PPAD-complete [CDT06], and for k players it is in PPAD [DGP09]. It follows that, for appropriate choices of ϵ, ϵ-well-supported Nash in k-player games reduces to ϵ-well-supported Nash in 2-player games. However, this chain of reductions is indirect, passing through intermediate notions other than games, and also rather complicated.

In this manuscript we present a direct, "game theoretic" polynomial-time reduction from k-player to 2-player games. In our reduction, every pure strategy of each of the k players is represented by a corresponding pure strategy of one of the 2 players. Previously, a direct reduction preserving exact Nash equilibrium was known from k-player to 3-player games [Bub79]. Such a reduction cannot exist to 2-player games due to issues of irrationality [Nas51], hence the need to consider the notion of ϵ-well-supported Nash in this context. Our reduction guarantees that for appropriate choices of ϵ_2 and ϵ_k, given any ϵ_2-well-supported Nash for the 2-player game, normalizing its probabilities according to the above correspondence immediately gives an ϵ_k-well-supported Nash for the k-player game.

The direct reduction makes no use of the concept of PPAD. This eliminates some of the difficulties involved in following the known indirect reduction. It is inevitable that unlike the indirect reduction, our reduction by itself does not establish the PPAD-completeness of computing (or approximating) Nash equilibria. Nevertheless, the new gadgets we introduce are relevant to the notion of PPAD-completeness, as they can be used in other reductions among PPAD problems. Moreover, our reduction provides an alternative proof to the proof of Daskalakis et al. [DGP09] that finding an approximate Nash equilibrium in k-player games is in PPAD.

*Work supported in part by The Israel Science Foundation (grant No. 873/08).

†Weizmann Institute of Science, PO Box 26, Rehovot 76100, Israel. Email: uriel.feige@weizmann.ac.il. The author holds the Lawrence G. Horowitz Professorial Chair at the Weizmann Institute.

‡Weizmann Institute of Science, PO Box 26, Rehovot 76100, Israel. Email: inbal.talgam@weizmann.ac.il.

In the k-player case, the payoff of each player depends on the combined behavior of the other players. We can thus view each player's set of expected payoffs as a set of multiplicative functions in the other players' strategies. In a 2-player game, however, each player interacts with a single other player, and so the expected payoffs are linear [vS07]. The described gap calls for "linearization" of k-player games, and indeed the first step of our reduction replaces the multilateral interactions among the k players with bilateral interactions among pairs of players. In the next step, two representative "super-players" replace the multiple players, resulting in a 2-player game.

In terms of techniques, the first step of the reduction uses and extends the machinery of gadget games developed by Goldberg and Papadimitriou [GP06]. We introduce a new gadget for performing approximate multiplication using linear operations, in order to bridge the gap between multiplicative and linear games. The second step of the reduction uses similar methods to [GP06] and [MT09] in order to replace multiple players by 2 players. The resulting 2-player game is a combination of a generalized Matching Pennies game [GP06] and an imitation game [MT09].

1.1 Preliminaries

Let $[n] = \{1, \dots, n\}$, and $\|v\| = \sum_i |v_i|$. For vectors $\boldsymbol{u}$ and $\boldsymbol{v}$ of length n, let $\boldsymbol{u} \otimes \boldsymbol{v}$ denote their tensor product written as a vector of length n^2, where entry $(i-1)n + j$ is $u_i v_j$. We write $x = y \pm z$ to denote $y - z \leq x \leq y + z$. For vectors, $\boldsymbol{x} = \boldsymbol{y} \pm z$ denotes $y_i - z \leq x_i \leq y_i + z$ for every i.

Normal Form Games Players of a normal form game G_k are numbered from 1 to k. Unless stated otherwise, every player has n pure strategies numbered from 1 to n. A *pure strategy profile* $\boldsymbol{s}$ is a vector of length k in $[n] \times \cdots \times [n]$, containing one pure strategy per player. $\boldsymbol{s^{-i}}$ is a pure strategy profile for all players except i, obtained from $\boldsymbol{s}$ by removing the i'th entry. A *payoff matrix* $M^i = M^i_{G_k}$ for player i is of size $n \times n^{k-1}$. Unless stated otherwise, all entries are rationals in the $[0, 1]$ range. $M^i[j, \boldsymbol{s^{-i}}]$ is the payoff player i receives for playing pure strategy j against pure strategy profile $\boldsymbol{s^{-i}}$. A *mixed strategy* $\boldsymbol{p^i}$ for player i is a probability distribution over $[n]$, denoting the probabilities with which i plays her pure strategies. Its *support* is the set of pure strategies $\{j : p^i_j > 0\}$. A *mixed strategy profile* $\boldsymbol{p} = (\boldsymbol{p^1}, \dots, \boldsymbol{p^k})$ is a set of mixed strategies for every player, and $\boldsymbol{p^{-i}}$ is a similar set for every player except i. Let $\tilde{\boldsymbol{p}}$ be the *joint mixed strategy* distribution, i.e. $\tilde{\boldsymbol{p}} = \boldsymbol{p^1} \otimes \cdots \otimes \boldsymbol{p^k}$. For every pure strategy profile $\boldsymbol{s}$, entry $\tilde{\boldsymbol{p}}[\boldsymbol{s}]$ is the probability $\prod_i p^i_{s_i}$ that every player i plays pure strategy s_i. Let $\tilde{\boldsymbol{p}}^{\boldsymbol{-i}}$ be the joint mixed strategy of all players except i. Given a mixed strategy profile $\boldsymbol{p^{-i}}$, the *expected payoff vector* $\boldsymbol{u^i_{G_k}}$ equals $M^i \tilde{\boldsymbol{p}}^{\boldsymbol{-i}}$. The j'th entry $\boldsymbol{u^i_{G_k}}[j]$ is the expected payoff player i receives for playing pure strategy j while the others play $\boldsymbol{p^{-i}}$. Thus, the expected payoffs are algebraic functions in the probabilities played by the others.

Polymatrix (Linear) Games In a polymatrix game, every player plays bilaterally against other players, and receives the sum of payoffs obtained from these bilateral interactions. Thus, polymatrix games are actually collections of 2-player games in which every player plays the same strategy in every game she participates in. Players are numbered from 1 to m; player i has $2 \leq n_i \leq n$ pure strategies and $m - 1$ rational payoff matrices $M^{i,i'}$ of size $n_i \times n_{i'}$. Entry $M^{i,i'}[j, j']$ is the payoff to player i for playing j against player i' who plays j'. If players i, i' do not interact or if their interaction is one-sided and does not influence player i's payoff, then $M^{i,i'}$ is set to be all-zeros. Given a pure strategy profile $\boldsymbol{s^{-i}}$, the total payoff to player i for playing j is $\sum_{i' \neq i} M^{i,i'}[j, \boldsymbol{s^{-i}}[i']]$. Given a mixed strategy profile $\boldsymbol{p^{-i}}$, the expected payoff vector of player i is $\boldsymbol{u^i_{G_m}} = \sum_{i' \neq i} M^{i,i'} \boldsymbol{p^{i'}}$. Equivalently, if $M^i_{G_m} = (M^{i,1} \cdots M^{i,m})$ contains all the player's payoff matrices as submatrices, then $\boldsymbol{u^i_{G_m}} = M^i_{G_m} \boldsymbol{p^{-i}}$. The expected payoffs of a player are thus linear functions in the probabilities of the others. Unlike normal form games, the size of polymatrix games is polynomial in n even when the number of players m is non-constant ($m = \text{poly}(n)$).

Nash Equilibrium, Approximations and Computational Problems A *Nash equilibrium* is a mixed strategy profile such that the players of the game cannot improve their expected payoffs by deviating from it unilaterally. The supports of a Nash equilibrium contain only pure strategies that are *best responses*, i.e. maximize the expected payoff given the mixed strategies of the other players. Formally, given a mixed strategy profile $\boldsymbol{p^{-i}}$, pure strategy j is a best response for player i if $\boldsymbol{u^i_G}[j] = \max_{j' \in [n]}\{\boldsymbol{u^i_G}[j']\}$.

Every game has a Nash equilibrium [Nas51], but finding such an equilibrium may be difficult. There are games for which every Nash equilibrium contains irrational probabilities, making it hard even to represent. This motivates the consideration of approximate instead of exact Nash equilibrium. In the context of reductions from k-player to 2-player games, there is another motivation for considering approximate Nash equilibrium. Unlike k player games, 2-player games always have a rational Nash equilibrium [Nas51, Pap07]. Thus we do not expect to find a reduction that preserves exact Nash equilibrium, direct or indirect.

There are several possible notions of approximation. We focus on the notion of ϵ-*well-supported* Nash equilibrium, a mixed strategy profile whose supports contain only ϵ-best responses, i.e. pure strategies that maximize the expected payoff up to an additive factor of ϵ. We will primarily be interested in small, non-constant values of ϵ, namely $\epsilon = 1/\text{poly}(n)$ and $\epsilon = 1/\exp(n)$. A related, computationally equivalent approximation notion is that of ϵ-Nash equilibrium - a mixed strategy profile from which deviating unilaterally cannot improve a player's expected payoff by more than ϵ [DGP09]. For other approximation notions see [EY07].

Definition 1.1 (ϵ_k-kNASH and ϵ_m-LINEAR-NASH) *Given a pair of normal form game G_k and accuracy parameter ϵ_k, the problem ϵ_k-kNASH is to find an ϵ_k-well-supported Nash equilibrium of G_k. Given a pair of polymatrix game G_m and accuracy parameter ϵ_m, the problem ϵ_m-LINEAR-NASH is to find an ϵ_m-well-supported Nash equilibrium of G_m.*

1.2 Our Results

Let $(G_{m_1}, \epsilon_{m_1}), (G_{m_2}, \epsilon_{m_2})$ be two pairs of games and accuracy parameters. The games have m_1, m_2 players respectively; the number of pure strategies of player i is n_i^1, n_i^2 respectively. The following definitions are based on the notion of *reduction scheme* defined by Bubelis [Bub79].

Definition 1.2 (Mapping between Games) *A* mapping *from G_{m_1} to G_{m_2} includes:*

- *A function $g : [m_1] \to [m_2]$ mapping players of G_{m_1} to players of G_{m_2};*
- *For every $i \in [m_1]$, an injective function $h_i : [n_i^1] \to [n^2_{g(i)}]$ mapping pure strategies of player i to distinct pure strategies of player $g(i)$.*

Definition 1.3 (Direct Reduction) *A* direct reduction *from $(G_{m_1}, \epsilon_{m_1})$ to $(G_{m_2}, \epsilon_{m_2})$ is a mapping from G_{m_1} to G_{m_2}, such that for every ϵ_{m_2}-well-supported Nash equlibrium $(\boldsymbol{q^1}, \ldots, \boldsymbol{q^{m_2}})$ of G_{m_2}, an ϵ_{m_1}-well-supported Nash equilibrium $(\boldsymbol{p^1}, \ldots, \boldsymbol{p^{m_1}})$ of G_{m_1} can be obtained by renormalizing probabilities as follows: $\boldsymbol{p^i}[j] = (1/z)\boldsymbol{q^{g(i)}}[h_i(j)]$ (where z is a normalization factor).*

Theorem 1.4 (Main) *For every $\epsilon_k < 1$, there exists a direct reduction from ϵ_k-kNASH to ϵ_2-2NASH, where $\epsilon_2 = poly(\epsilon_k/|G_k|)$. The reduction runs in polynomial time in $|G_k|$ and in $\log(1/\epsilon_k)$.*

Corollary 1.5 *There is a direct, polynomial time reduction from $(1/\exp(n))$-kNASH to $(1/\exp(n))$-2NASH, and from $(1/poly(n))$-kNASH to $(1/poly(n))$-2NASH.*

Proof of Theorem 1.4: By combining Theorem 3.1 (linearizing reduction) with Theorem 4.1 (reduction from linear to bimatrix games), and plugging in the parameters of Lemma 2.3 (logarithmic-sized linear multiplication gadget). ■

For simplicity of presentation we defer the proof of Lemma 2.3 to Section 5, and first prove in Section 2 a slightly weaker version (Lemma 2.2 - polynomial-sized linear multiplication gadget), resulting in a reduction that runs in polynomial time in $1/\epsilon_k$ instead of $\log(1/\epsilon_k)$.

1.3 Related Work

Bubelis [Bub79] shows a direct reduction from k-player to 3-player games. This reduction relies heavily on the multiplicative nature of 3-player games. Examples of direct reductions involving 2-player games include *symmetrization* [GKT50], and reduction to *imitation games* [MT09]. We use imitation games in Section 4.

PPAD-completeness results Papadimitriou introduced PPAD in 1991, motivated largely by the challenge of classifying the Nash equilibrium problem [Pap94]. Formally, PPAD is the class of total search problems polynomial-time reducible to the abstract path-following problem END OF THE LINE. Another important PPAD-complete problem is 3D-BROUWER, a discrete version of finding Brouwer fixed-points in a 3-dimensional domain (the same problem in high-dimension is known as nD-BROUWER). The known results can be summarized by the two following chains of reductions, each forming an indirect reduction (according to Definition 1.3 of directness) from k-player games to 2-player games:

- $1/\exp(n)$-kNASH $\leq$ END OF THE LINE $\leq$ 3D-BROUWER $\leq$ ADDITIVE GRAPHICAL NASH $\leq 1/\exp(n)$-2NASH
- $1/\exp(n)$-kNASH $\leq$ END OF THE LINE $\leq$ 2D-BROUWER $\leq$ nD-BROUWER $\leq 1/\text{poly}(n)$-2NASH

The reductions in the first chain are by [vdLT82, DGP09], [Pap94, DGP09], [CD06b, DGP09] and [DGP09], respectively. The reductions in the second chain are by [vdLT82, DGP09], [CD06a], [CDT06], [CDT06], respectively. For an overview of these celebrated results see [Rou10]. In comparison, our reduction can be written as:

- ϵ_k-kNASH $\leq$ ϵ_m-LINEAR-NASH $\leq$ ϵ_2-2NASH

where $\epsilon_k, \epsilon_m, \epsilon_2$ can either be all $1/\exp(n)$ or all $1/\text{poly}(n)$. Note there is gap between the second chain of reductions and our results - the second chain achieves a stronger reduction from $1/\exp(n)$-kNASH to $1/\text{poly}(n)$-2NASH. Achieving a direct version of this result by [CDT06] is an interesting open problem. Note also that our reduction from ϵ_m-LINEAR-NASH to ϵ_2-2NASH is somewhat similar to the reduction from ADDITIVE GRAPHICAL NASH to $1/\exp(n)$-2NASH, however our reduction does not require the input game to be bipartite nor does it limit the number of interactions per player.

Another open question is the complexity of ϵ_k-kNASH and ϵ_2-2NASH for *constant* values of ϵ_k, ϵ_2. As a quasi-polynomial algorithm is known [Alt94, LMM03], these problems are not believed to be PPAD-complete. The current state-of-the-art is a polynomial-time algorithm for ϵ_2-2NASH where $\epsilon_2 \approx 0.667$ [KS10]. For finding ϵ_2-Nash equilibrium rather than ϵ_2-well-supported Nash equilibrium, there is an algorithm where $\epsilon_2 \approx 0.339$ [TS07] (see also [DMP06], [BBM07], [TS09]). On the negative side, several algorithmic techniques have been ruled out [HK09], [DP09].

Reductions to 2-player games and linearization The empirical success of the Lemke-Howson algorithm [LH64] for finding Nash equilibrium in 2-player games has motivated research on extending it to a more general class of games. Daskalakis et al. show a general reduction from *succinct* games to 2-player games, which can be applied to any game in which the expected payoffs can be calculated using only $+, *, \max$ [DFP06]. Their reduction goes through the steps of the first chain of reductions above. Govindan and Wilson present a non-polynomial linearizing reduction, which reduces multiplayer games to polymatrix games while preserving approximate Nash equilibrium [GW10]. Their reduction introduces a central coordinator player, who interacts bilaterally with every player while simulating the combined behavior of the other players.

In addition, linearization is also related to Etessami and Yannakakis's formulation of PPAD as the class of fixed-point problems for piecewise-linear functions (computable by $+, \text{scale}, \max$) [EY07].

2 A Linear Multiplication Gadget

In this section we construct a linear multiplication gadget using standard gadgets as building blocks.

Theorem 2.1 (Linear Multiplication Gadget) *There exist constants $\epsilon_0 < 1, c, d$ and an increasing polynomial function f such that the following holds. For every $\epsilon < \epsilon_0$, there exists a linear multiplication gadget $G_* = G_*(\epsilon)$ of size $O(m \cdot f(\frac{1}{\epsilon}))$, such that in an ϵ-well-supported Nash equilibrium, the output of G_* equals the product of its m inputs up to an additive error of $\pm dm\epsilon^c$.*

We develop two different constructions of G_*, with two different sets of parameters ϵ_0, c, d, f.

Lemma 2.2 (Polynomial-Sized Construction) *Theorem 2.1 holds with the following parameters:*[1] $\epsilon_0 = \frac{1}{4}$, $c = 1$, $d = 19$ *and* $f(x) = x^2$.

Lemma 2.3 (Logarithmic-Sized Construction) *Theorem 2.1 holds with the following parameters:* $\epsilon_0 = \frac{1}{10^5}$, $c = \frac{1}{2}$, $d = 3$ *and* $f(x) = \log x$.

The second construction gives a smaller gadget with size $O(m \log \frac{1}{\epsilon})$ instead of $O(\frac{m}{\epsilon^2})$, but is more complicated than the first construction. The rest of this section describes the first construction and proves Lemma 2.2. Details of the second construction and the proof of Lemma 2.3 appear in Section 5.

2.1 Linear Gadgets

Goldberg and Papadimitriou developed the framework of gadgets [GP06], carefully-engineered games that simulate arithmetic calculations and are useful in many PPAD-completeness results (see, e.g., [DGP09]). The players of a gadget game are typically *binary*, representing numerical values in the range $[0, 1]$.

Definition 2.4 (Binary Player) *A* binary player *P is a player that has exactly two pure strategies 0 and 1. We say P* represents *the numerical value $p \in [0, 1]$ if her mixed strategy is $(1 - p, p)$, i.e. she plays pure strategy 1 with probability p.*

Gadget games have three kinds of binary players - one or more input players,[2] one output player, and one or more auxiliary players. The *size* of a gadget is the number of its auxiliary and output players. The values represented by the input and output players are the *input values* and *output value* of the gadget. In every ϵ-well-supported Nash equilibrium of the gadget game, the output value p is equal to the result of an arithmetic operation on the input values (up to small error). This arithmetic relation between the input and output values is the *guarantee* of the gadget. To achieve the guarantee, the output player is incentivized to play the appropriate value p, by choosing appropriate payoff values for both the output and auxiliary players. Our reductions require gadgets with *linear* guarantees, which differ slightly from the *graphical* and *additive*-graphical gadgets used in previous works.

Definition 2.5 (Linear Gadgets) *A* linear gadget *is a polymatrix gadget game with payoffs in* $[0, 1]$. *Linear gadgets simulate linear arithmetic operations, i.e. their guarantee is a linear relation between the input and output values.*

[1] The choice of $d = 19$ simplifies the proof, but it is not hard to show that for the same construction d can be replaced with a smaller value.

[2] Gadgets can also have non-binary input players, in which case the value the input players represent is considered to be the probability with which they play a certain predetermined pure strategy.

Several gadgets can be combined into a single game, much like arithmetic gates are combined into a circuit to carry out involved calculations. By setting the output player P of gadget G_1 to be an input player of gadget G_2, the value p represented by P is shared among the gadgets. We represent a combination of gadgets by a series of calculations on the input and output values. For example, if P' is the output player of G_2 representing the value p', then we write the above combination as $p' = G_2(p)$ (where in turn $p = G_1(\dots)$). The following fact explains why the same player can be an input player of multiple gadgets, but can only be the output player of a single gadget (as for auxliary players, they are considered part of the inner implementation and are thus never shared among different gadgets). It is a consequence of the gadget determining the payoff matrix of its output player, but not of its input players.

Fact 2.6 (Combining Gadgets) *For every game in which no player is the output player of more than one gadget, the guarantees of all gadgets hold simultanuously when the game is in ϵ-well-supported Nash equilibrium.*

2.1.1 Standard Gadgets

The following gadgets are constructed by Daskalakis et al. [DGP09]. To demonstrate the principle behind their construction, we include here the proof of Lemma 2.7; proofs of Lemma 2.8 and Lemma 2.9 appear in Appendix A for completeness.

Lemma 2.7 (Linear Threshold Gadget) *For every rational $\zeta \in [0, 1]$, there exists a linear gadget $G_{>\zeta}$ of size $O(1)$ with input p_1, such that in an ϵ-well-supported Nash equilibrium the output is 1 if $p_1 > \zeta + \epsilon$ and 0 if $p_1 < \zeta - \epsilon$, and otherwise it may be any value in $[0, 1]$.*

Proof: Let P_1, P be the input and output players of $G_{>\zeta}$ representing values p_1, p, respectively. We set the payoff matrix M^{P,P_1} to be:

$$M^{P,P_1} = \begin{pmatrix} \zeta & \zeta \\ 0 & 1 \end{pmatrix}$$

$G_{>\zeta}$ has no auxiliary players, and so this concludes the construction. We now show that when $G_{>\zeta}$ is in ϵ-well-supported Nash equilibrium, the guarantee of this gadget holds. Let $\boldsymbol{p^1} = (1 - p_1, p_1)$ be player P_1's mixed strategy in the ϵ-well-supported Nash equilibrium. The expected payoff vector $\boldsymbol{u^P}$ of player P is equal to:

$$\boldsymbol{u^P} = M^{P,P_1}\boldsymbol{p^1} = (\zeta, p_1)$$

If $p_1 > \zeta + \epsilon$, the only ϵ-best response for player P is pure strategy 1, so P's mixed strategy $(1 - p, p)$ in the ϵ-well-supported Nash equilibrium must be $(0, 1)$ and thus $p = 1$. Similarly, if $p_1 < \zeta - \epsilon$ then $(1 - p, p) = (1, 0)$ and thus $p = 0$. ■

Lemma 2.8 (Linear AND Gadget) *There exists a linear gadget $G_\wedge$ of size $O(1)$ with inputs p_1, p_2, such that in an ϵ-well-supported Nash equilibrium where $\epsilon < \frac{1}{4}$ the output is 1 if $p_1 = p_2 = 1$ and 0 if $(p_1 = 0) \vee (p_2 = 0)$, and otherwise it may be any value in $[0, 1]$.*

Lemma 2.9 (Linear Scaled-Summation Gadget) *For every rational $\zeta \in [0, 1]$, there exists a linear gadget $G_{+,*\zeta}$ of size $O(1)$ with inputs $p_1, \dots, p_m$, such that in an ϵ-well-supported Nash equilibrium the output is $\min\{\zeta(p_1 + \cdots + p_m), 1\} \pm \epsilon$.*

In addition, there exist standard gadgets for multiplication, but these are inherently nonlinear - the constructions are based on expected payoffs being multiplicative functions in players' probabilities (see, e.g., [DGP09]).

2.2 Construction

Here we show the construction of G_* that will be used to prove Lemma 2.2. We show a construction for multiplying 2 inputs, and multiplying m inputs can be achieved by connecting $m-1$ copies of this construction serially. Let P_1, P_2 be the input players representing values p_1, p_2, and let P be the output player representing value p. Let $\tau = 3\epsilon$ (for simplicity assume that $1/\tau$ is integer). We first encode every input in unary representation, with precision of up to $\pm\tau$. For this we use $2/\tau$ auxiliary players: The vectors $\boldsymbol{v^1} = (v_1^1, \ldots, v_{1/\tau}^1)$ and $\boldsymbol{v^2} = (v_1^2, \ldots, v_{1/\tau}^2)$ of values represented by auxiliary players $\{V_i^1\}$ and $\{V_i^2\}$ are the unary encodings. The i'th unary bit of p_1 is v_i^1, and it is calculated by the threshold gadget $G_{>\zeta}$ (Lemma 2.7) as follows: $v_i^1 = G_{>i\tau}(p_1)$. Similarly, $v_i^2 = G_{>i\tau}(p_2)$. Then we perform unary multiplication using the AND gadget $G_\wedge$ (Lemma 2.8). The result is a matrix U, which contains $1/\tau^2$ values $u_{i,j} = G_\wedge(v_i^1, v_j^2)$, represented by auxiliary players $\{U_{i,j}\}$. The construction is complete by summing up and scaling U's entries using the scaled-summation gadget $G_{+,*\zeta}$ (Lemma 2.9) as follows: $p = G_{+,*\tau^2}(u_{1,1}, u_{1,2}, \ldots, u_{1/\tau,1/\tau})$. This establishes the relation between the input values p_1, p_2 and the output value p of G_*. Note that the payoffs of all players are determined by the standard gadgets.

2.3 Correctness

We prove Lemma 2.2 for the case $m = 2$. Namely, we show that for every $\epsilon < 1/4$, when G_* is in ϵ-well-supported Nash equilibrium then $p = p_1p_2 \pm d\epsilon$, that G_* is linear and that the size of G_* is $O(1/\epsilon^2)$. The proof of Lemma 2.2 for general m follows, since concatenating $m-1$ copies of G_* increases the error and gadget size by a multiplicative factor of m.

Proof of Lemma 2.2 (Polynomial-Sized Construction): First note that G_* is a combination of linear gadgets and is thus itself linear. The size of G_* is $O(1/\tau^2)$, the total size of the standard gadgets ($2/\tau$ threshold gadgets $G_{>\zeta}$, $1/\tau^2$ AND gadgets $G_\wedge$, and 1 scaled-summation gadget $G_{+,*\zeta}$, all of size $O(1)$).

We assume G_* is in ϵ-well-supported Nash equilibrium where $\epsilon < 1/4$, and write the input values p_1, p_2 as integer multiples of τ plus a small error: let $p_1 = i^*\tau + \delta_1$ and $p_2 = j^*\tau + \delta_2$, where $0 \le i^*, j^* \le 1/\tau$ and $0 \le \delta_1, \delta_2 < \tau$. The following claim shows that the coefficients i^*, j^* are correctly encoded as unary vectors $\boldsymbol{v^1}, \boldsymbol{v^2}$, and is a direct consequence of the threshold gadget's guarantee (Lemma 2.7). The threshold gadget is "brittle" in the sense that for a small range of inputs it returns an arbitrary output, but the choice of $\tau = 3\epsilon$ ensures this happens for at most one unary bit.

Claim 2.10 (Unary Encoding) *$\boldsymbol{v^1}$ is of the form $(1, \ldots, 1, ?, 0, \ldots, 0)$, where $\|\boldsymbol{v^1}\| = i^* \pm 1$ and '?' denotes any value in $[0, 1]$. The same holds for $\boldsymbol{v^2}$ and j^*.*

Proof: Consider the i'th entry of $\boldsymbol{v^1}$. By construction, $v_i^1 = G_{>i\tau}(p_1)$. By Lemma 2.7, v_i^1 indicates whether $p_1 > i\tau + \epsilon$ or $p_1 < i\tau - \epsilon$, and otherwise can be any value in $[0, 1]$. Since $\tau = 3\epsilon$ we know that $p_1 > (i^* - 1)\tau + \epsilon$, therefore for every $i \le i^* - 1$ entry v_i^1 is equal to 1, and in total $\|\boldsymbol{v^1}\| \ge i^* - 1$. On the other hand we know that $p_1 < (i^* + 2) - \epsilon$, therefore for every $i \ge i^* + 2$ entry v_i^1 is equal to 0, and in total $\|\boldsymbol{v^1}\| \le i^* + 1$. Moreover, since $\tau > 2\epsilon$, there can be at most one value of i for which $p_1 = i\tau \pm \epsilon$, and so there can be at most one entry i of $\boldsymbol{v^1}$ which is an arbitrary value '?' in $[0, 1]$. ■

The rest of the proof of Lemma 2.2 is a straightforward corollary of the other gadget guarantees. Let $\|U\|$ denote the sum $\sum_{i,j} u_{i,j}$ of matrix U's entries. By construction, $p = G_{+,*\tau^2}(u_{1,1}, u_{1,2}, \ldots, u_{1/\tau,1/\tau})$, thus by the guarantee of gadget $G_{+,*\zeta}$ (Lemma 2.9), $p = \tau^2\|U\| \pm \epsilon$. We write the product p_1p_2 as an integer multiple of τ^2 up to a small error: $i^*j^*\tau^2 \le p_1p_2 < i^*j^*\tau^2 + 3\tau$. The next claim shows that $\|U\|$ gives approximately the correct coefficient of τ^2.

Claim 2.11 $(i^* - 1)(j^* - 1) \le \|U\| \le (i^* + 1)(j^* + 1)$

Proof: Consider the (i,j)'th entry of U. By construction, $u_{i,j} = G_\wedge(v_i^1, v_j^2)$. By Lemma 2.8, if $\epsilon < 1/4$ and $v_i^1 = v_j^2 = 1$ then $u_{i,j} = 1$. By Claim 2.10, there are at least $(i^*-1)(j^*-1)$ pairs i,j such that $v_i^1 = v_j^2 = 1$, and so $\|U\| \geq (i^*-1)(j^*-1)$. Similarly, by Lemma 2.8, if $\epsilon < 1/4$ and $v_i^1 = 0 \vee v_j^2 = 0$ then $u_{i,j} = 0$. By Claim 2.10, there are at most $(i^*+1)(j^*+1)$ pairs i,j such that $v_i^1 \neq 0 \wedge v_j^2 \neq 0$, and so $\|U\| \leq (i^*+1)(j^*+1)$. ■

Since $i^*, j^* \leq 1/\tau$, it follows from the above claim that $\|U\| = i^*j^* \pm (2/\tau+1)$. So $p = \tau^2 i^* j^* \pm (3\tau+\epsilon) = p_1p_2 \pm (6\tau+\epsilon) = p_1p_2 \pm d\epsilon$, where $d = 19$. This concludes the proof of Lemma 2.2, showing that G_* outputs the product of its inputs p_1p_2 up to a small error of $\pm d\epsilon$. ■

Example 2.12 *Let $p_1 = 7\tau + \epsilon/4$ and $p_2 = 2\tau + (\tau - \epsilon/8)$. First we find the unary encoding: $\boldsymbol{v^1} = (1,1,1,1,1,1,?,0,\ldots,0)$ and $\boldsymbol{v^2} = (1,1,?,0,\ldots,0)$. Then we perform unary multiplication:*

$$U' = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & ? & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & ? & 0 \\ ? & ? & ? & ? & ? & ? & ? & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}, U = \begin{pmatrix} U' & 0 \\ 0 & 0 \end{pmatrix}_{1/\tau \times 1/\tau}$$

Summing up and scaling the entries of U we get $p = 12\tau^2 + O(\epsilon)$, which is close to p_1p_2 up to $O(\epsilon)$.

3 Linearizing Multiplayer Games

In this section we show a direct reduction from k-player games to polymatrix games. Let G_k denote the input game to the reduction, and let G_m denote the corresponding output game. The reduction relies on the fact that, although G_k's expected payoffs are nonlinear in its players' probabilities, they are linear in products of its players' probabilities. A key component of our reduction is a linear multiplication gadget for computing these products, which exists according to Theorem 2.1. Let f be an increasing polynomial function as in Theorem 2.1.

Theorem 3.1 (A Linearizing Reduction) *For every $\epsilon_k < 1$, there exists a direct reduction from ϵ_k-kNASH to ϵ_m-LINEAR-NASH, where $\epsilon_m = poly(\epsilon_k/|G_k|)$. The reduction runs in polynomial time in $|G_k|$ and in $f(1/\epsilon_k)$.*

Lemma 3.2 (Recovering ϵ_k-Well-Supported Nash) *Let $(\boldsymbol{p^1}, \ldots, \boldsymbol{p^m})$ be an ϵ_m-well-supported Nash equilibrium of G_m. Then the first k mixed strategies $\boldsymbol{p^1}, \ldots, \boldsymbol{p^k}$ form an ϵ_k-well-supported Nash equilibrium of G_k.*

3.1 Preserving Expected Payoffs

The following lemma will be useful in desiging the linearizing reduction. Let G_k be a game with k players, n pure strategies each. Let G_m be a game with $m > k$ players, where the first k players have the same pure strategies as the players of G_k. Let $(\boldsymbol{p^1}, \ldots, \boldsymbol{p^k}), (\boldsymbol{p^1}, \ldots, \boldsymbol{p^m})$ be mixed strategy profiles of G_k, G_m.

Lemma 3.3 (Almost Equal Expected Payoffs) *If for every player $1 \leq i \leq k$, the expected payoff vectors $\boldsymbol{u^i_{G_m}}$ and $\boldsymbol{u^i_{G_k}}$ are entry-wise equal up to an additive factor of δ, and $(\boldsymbol{p^1}, \ldots, \boldsymbol{p^m})$ is an ϵ_m-well-supported Nash equilibrium of G_m, then $(\boldsymbol{p^1}, \ldots, \boldsymbol{p^k})$ is an ϵ_k-well-supported Nash equilibrium of G_k where $\epsilon_k = 2\delta + \epsilon_m$.*

Proof: Let $j \in [n]$ be a pure strategy in the support of player i ($p^i_j > 0$). We know that j is an ϵ_m-best response in G_m. Assume for contradiction that j is *not* an ϵ_k-best response in G_k, i.e. there is a pure strategy $j' \in [n], j' \neq j$ such that $\boldsymbol{u^i_{G_k}}(j') > \boldsymbol{u^i_{G_k}}(j) + \epsilon_k$. So $\boldsymbol{u^i_{G_m}}(j') + \delta > \boldsymbol{u^i_{G_m}}(j) - \delta + \epsilon_k$. Since $\epsilon_k - 2\delta = \epsilon_m$, then $\boldsymbol{u^i_{G_m}}(j') > \boldsymbol{u^i_{G_m}}(j) + \epsilon_m$, contradiction. ■

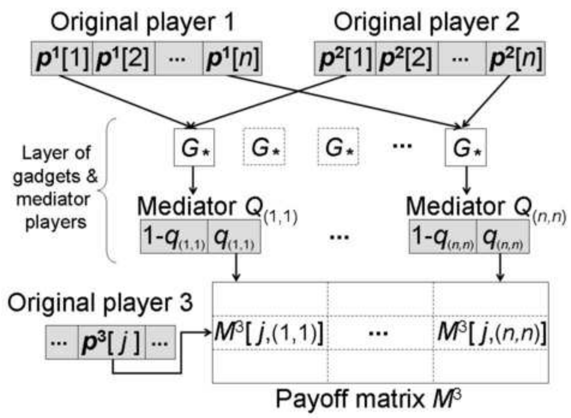

Figure 1: Linearization of a 3-Player Game - Partial View of G_m
The arrows indicate how the probabilities of original players 1 and 2 influence the expected payoff of original player 3 via a layer of gadgets and mediator players.

3.2 The Linearizing Reduction

Given an input pair (G_k, ϵ_k), we find an output pair (G_m, ϵ_m) as follows. Let $\epsilon_0 < 1, c, d$ be the constant parameters of Theorem 2.1. Then $\epsilon_m = \min\{(\epsilon_k/3n^{k-1}dk)^{1/c},\ \epsilon_0\}$. The players of G_m are:

- *Original players* - the first k players of G_m have the same pure strategies as G_k's players. $\boldsymbol{p^i}$ denotes the mixed strategy of original player i.
- *Mediator players* - for every $i \in [k]$, there is a set of n^{k-1} binary players that corresponds to the set of n^{k-1} pure strategy profiles of all original players except i. We denote by $Q_{\boldsymbol{s^{-i}}}$ the mediator player corresponding to pure strategy profile $\boldsymbol{s^{-i}}$ and by $q_{\boldsymbol{s^{-i}}}$ the represented value.
- *Gadget players* - all auxiliary players belonging to kn^{k-1} copies of the linear multiplication gadget G_*.

Every mediator player is set to be the output player of a gadget G_* as follows: $q_{\boldsymbol{s^{-i}}} = G_*(p^1_{\boldsymbol{s^{-i}}[1]}, \ldots, p^k_{\boldsymbol{s^{-i}}[k]})$. Thus, $q_{\boldsymbol{s^{-i}}}$ will be approximately equal to the probability with which the original players play the pure strategy profile $\boldsymbol{s^{-i}}$. Let $\boldsymbol{q^i}$ be the vector of values $\{q_{\boldsymbol{s^{-i}}}\}$, then it's approximately equal to $\tilde{\boldsymbol{p}}^{-\boldsymbol{i}}$, the joint mixed strategy distribution of all original players except i.

To complete the description of G_m it remains to specify the non-zero payoff matrices of the original players (all other payoffs are determined by the gadgets). In G_k, the expected payoff vector of player i is $u^i_{G_k} = M^i_{G_k}\tilde{\boldsymbol{p}}^{-\boldsymbol{i}}$. In G_m, the payoff of original player i will be influenced only by the i'th set of mediator players $\{Q_{\boldsymbol{s^{-i}}}\}$ who play $\boldsymbol{q^i}$. Instead of describing every payoff matrix $M^{i,Q_{\boldsymbol{s^{-i}}}}$ separately, (such a description appears in the proof of Lemma 3.2), we describe one large payoff matrix $M^i_{G_m}$ that contains all the others (or more precisely, all their nonzero columns) as submatrices. We want the expected payoffs in G_m to be as close as possible to those of G_k. Thus, we set $M^i_{G_m} = M^i_{G_k}$. This concludes the contruction.

3.3 Correctness

Proof of Theorem 3.1 (A Linearizing Reduction): The reduction runs in time polynomial in $|G_k| = \Theta(kn^k)$ and in $f(1/\epsilon_k)$: The running time depends on the size of the polymatrix game G_m, which is polynomial in the number of its players. There are k original players, kn^{k-1} mediator players and $kn^{k-1}O(|G_*|)$ auxiliary players. By Theorem 2.1, $|G_*| = O(k \cdot f(1/\epsilon_m))$. Since f is a polynomial function and $\epsilon_m = \text{poly}(\epsilon_k/|G_k|)$, the total number of players is indeed polynomial in $|G_k|$ and in

$f(1/\epsilon_k)$. The rest of the proof follows from Lemma 3.2. ■

Proof of Lemma 3.2 (Recovering ϵ_k-Well-supported Nash after Linearization): Let $(\boldsymbol{p^1}, \ldots, \boldsymbol{p^m})$ be an ϵ_m-well-supported Nash equilibrium of G_m. We show that the first k mixed strategies of $(\boldsymbol{p^1}, \ldots, \boldsymbol{p^m})$ form an ϵ_k-well-supported Nash equilibrium of G_k. We would like to upper bound the entry-wise distance between the payoff vectors $u^i_{G_m}, u^i_{G_k}$ so that we can apply Lemma 3.3. The proof proceeds as follows: We show that the expected payoff vector of original player i in G_m is $u^i_{G_m} = M^i_{G_m}\boldsymbol{q^i}$. Then we observe that the linear multiplication gadget G_* guarantees that vectors $\boldsymbol{q^i}$ and $\boldsymbol{\tilde{p}^{-i}}$ are close to each other, and recall that $M^i_{G_m} = M^i_{G_k}$. Since all payoffs are in $[0,1]$, we conclude that the expected payoffs $u^i_{G_k} = M^i_{G_k}\boldsymbol{\tilde{p}^{-i}}$ are preserved in G_m. The proof of Lemma 3.2 is then immediate by preservation of expected payoffs (Lemma 3.3).

We start by an alternative, more formal description of the original players' payoff matrices in G_m. Consider the payoff matrix $M^{i,Q_{\boldsymbol{s}^{-i}}}$, corresponding to the interaction between original player i and mediator $Q_{\boldsymbol{s}^{-i}}$. Since $Q_{\boldsymbol{s}^{-i}}$ is a binary player with pure strategies $\{0,1\}$, the size of $M^{i,Q_{\boldsymbol{s}^{-i}}}$ is $n \times 2$. For every $j \in [n]$ we set $M^{i,Q_{\boldsymbol{s}^{-i}}}[j,1] = M^i_{G_k}[j, \boldsymbol{s^{-i}}]$ (where $M^i_{G_k}$ is the payoff matrix of player i in game G_k), and $M^{i,Q_{\boldsymbol{s}^{-i}}}[j,0] = 0$. So the column $M^{i,Q_{\boldsymbol{s}^{-i}}}[\cdot,0]$ corresponding to the mediator's pure strategy 0 is all-zeros. The payoff matrix $M^i_{G_m}$ was defined above to contain all nonzero columns of payoff matrices $M^{i,Q_{\boldsymbol{s}^{-i}}}$, i.e., all columns $M^{i,Q_{\boldsymbol{s}^{-i}}}[\cdot,1]$. It is now not hard to verify that $M^i_{G_m} = M^i_{G_k}$, and so the alternative description is equivalent to the original one.

Claim 3.4 (Expected Payoffs Vector) *For every $i \in [k]$, the expected payoff vector of original player i in game G_m is $u^i_{G_m} = M^i_{G_m}\boldsymbol{q^i}$.*

Proof: $u^i_{G_m}$ is equal to the sum of expected payoff vectors of original player i from playing bilaterally against every mediator player in $\{Q_{\boldsymbol{s}^{-i}}\}$. Each expected payoff vector is a product of the payoff matrix $M^{i,Q_{\boldsymbol{s}^{-i}}}$ with vector $\boldsymbol{p}^{\boldsymbol{Q}_{\boldsymbol{s}^{-i}}} = (1 - q_{\boldsymbol{s}^{-i}}, q_{\boldsymbol{s}^{-i}})$ (the mixed strategy played by the binary mediator player $Q_{\boldsymbol{s}^{-i}}$). By construction of $M^{i,Q_{\boldsymbol{s}^{-i}}}$, the expected payoff vector is equal to the product of column vector $M^{i,Q_{\boldsymbol{s}^{-i}}}[\cdot,1]$ with scalar $q_{\boldsymbol{s}^{-i}}$. Therefore, the sum of expected payoff vectors over all mediators is equal to $M^i_{G_m}\boldsymbol{q^i}$. ■

We now show that $\boldsymbol{q^i}$ and $\boldsymbol{\tilde{p}^{-i}}$ are almost equal. Consider entry $q_{\boldsymbol{s}^{-i}}$ of $\boldsymbol{q^i}$. By construction, $q_{\boldsymbol{s}^{-i}} = G_*(p^1_{\boldsymbol{s}^{-i}[1]}, \ldots, p^k_{\boldsymbol{s}^{-i}[k]})$. By Theorem 2.1 and since $\epsilon_m < \epsilon_0$, the gadget G_* guarantees that $q_{\boldsymbol{s}^{-i}} = \prod_{i' \neq i} p^{i'}_{\boldsymbol{s}^{-i}[i']} \pm dk(\epsilon_m)^c$. By definition of $\boldsymbol{\tilde{p}^{-i}}$ as the joint mixed strategy distribution of all players except i we get that $q_{\boldsymbol{s}^{-i}} = \boldsymbol{\tilde{p}^{-i}}[\boldsymbol{s^{-i}}] \pm dk(\epsilon_m)^c$. Thus, $\boldsymbol{q^i} = \boldsymbol{\tilde{p}^{-i}} \pm dk(\epsilon_m)^c$.

Using the fact that the entries of $M^i_{G_m}, M^i_{G_k}$ are all in the range $[0,1]$, and that the dimensions of the matrices are $n \times n^{k-1}$, we conclude that $M^i_{G_m}\boldsymbol{q^i} = M^i_{G_k}\boldsymbol{\tilde{p}^{-i}} \pm n^{k-1}dk(\epsilon_m)^c$. We can now apply Lemma 3.3 with $\delta = n^{k-1}dk(\epsilon_m)^c$. So $(\boldsymbol{p^1}, \ldots, \boldsymbol{p^k})$ is a $(2\delta + \epsilon_m)$-well-supported Nash equilibrium of G_k, and plugging in the chosen value of ϵ_m gives ϵ_k-well-supported Nash equilibrium, as required. ■

4 Reducing Linear Games to Bimatrix Games

In this section we show how to replace the multiple players of a polymatrix game by two representative "super-players" of a bimatrix game. Let G_m denote the input game to the reduction, and let G_2 denote the corresponding output game.

Theorem 4.1 (Linear to Bimatrix) *For every $\epsilon_m < 1$, there exists a direct reduction from ϵ_m-LINEAR-NASH to ϵ_2-2NASH, where $\epsilon_2 = poly(\epsilon_m/|G_m|)$. The reduction runs in polynomial time in $|G_m|$ and in $\log(1/\epsilon_m)$.*

Lemma 4.2 (Recovering ϵ_m-Well-Supported Nash) *For every ϵ_2-well-supported Nash equilibrium $(\boldsymbol{x}, \boldsymbol{y})$ of G_2, partitioning $\boldsymbol{y}$ into subvectors $\boldsymbol{y^1}, \ldots, \boldsymbol{y^m}$ of lengths $n_1, \ldots, n_m$ and normalizing gives an ϵ_m-well-supported Nash equilibrium $(\boldsymbol{y^1}/\|\boldsymbol{y^1}\|, \ldots, \boldsymbol{y^m}/\|\boldsymbol{y^m}\|)$ of G_m.*

4.1 Imitation Games and Block ϵ-Uniform Games

The following definitions and lemmas will be useful in proving Theorem 4.1. An *imitation game* is a bimatrix game in which both players have N pure strategies, and the payoff matrix of player 2 is equal to the $N \times N$ identity matrix I_N. We call player 1 the *leader* and player 2 the *imitator*. A similar lemma to the following was proved in [MT09] for the case of exact Nash equilibrium.

Lemma 4.3 (Imitation) *Let $(\boldsymbol{x}, \boldsymbol{y})$ be an ϵ_2-well-supported Nash equilibrium of an imitation game G_2 where $\epsilon_2 \leq 1/N$. Then support$(\boldsymbol{y}) \subseteq$ support$(\boldsymbol{x})$.*

Proof: Assume pure strategy j is not in support$(\boldsymbol{x})$, i.e. $x_j = 0$. The expected payoff vector of the imitator is $\boldsymbol{u^2_{G_2}} = I_N \boldsymbol{x} = \boldsymbol{x}$, and so for pure strategy j the expected payoff is $x_j = 0$. Since $\boldsymbol{x}$ is a probability distribution vector with N entries of which one is assumed to be zero, there exists a pure strategy $j' \neq j$ for which the imitator's expected payoff is $\boldsymbol{u^2_{G_2}}[j'] = x_{j'} \geq 1/(N-1) > \epsilon$. The difference between the expected payoffs is more than ϵ, so j cannot be an ϵ-best response for the imitator and so does not belong to support$(\boldsymbol{y})$. We conclude that support$(\boldsymbol{y}) \subseteq$ support$(\boldsymbol{x})$, as required. ■

We call a bimatrix game *block ϵ-uniform* if player 1's payoff matrix A is of the following form:

- *Block matrix:* A is composed of m^2 blocks, where block (i, i'), denoted $A^{i,i'}$, is of size $n_i \times n_{i'}$;
- *Very negative diagonal:* The i'th diagonal block $A^{i,i}$ is equal to $-\alpha E_{n_i}$, where $\alpha = 8m^2/\epsilon$ and E_{n_i} is the all-ones matrix of size $n_i \times n_i$;
- $[0, 1]$ *entries:* All other entries of A are arbitrary values in the range $[0, 1]$.

For a similar construction see the generalized Matching Pennies game of [GP06]. If $\boldsymbol{x}, \boldsymbol{y}$ is a mixed strategy profile of an ϵ-block-uniform game, we denote by $\boldsymbol{x^1}, \ldots, \boldsymbol{x^m}$ and $\boldsymbol{y^1}, \ldots, \boldsymbol{y^m}$ the mixed strategy blocks of size $n_1, \ldots, n_m$. We say that block i belongs to the support of mixed strategy $\boldsymbol{x}$ if there is some pure strategy in block i that belongs to this support. The following lemma shows that in a block ϵ-uniform game, the weight of player 2 is ϵ-uniformly divided among all blocks i in support$(\boldsymbol{x})$.

Lemma 4.4 (ϵ-Uniform Weight Distribution) *Let $\boldsymbol{x}, \boldsymbol{y}$ be an ϵ_2-well-supported Nash equilibrium of a block ϵ_2-uniform game G_2. If block $i \in [m]$ belongs to the support of $\boldsymbol{x}$, then for every $i' \in [m]$, $\|\boldsymbol{y^i}\| \leq \|\boldsymbol{y^{i'}}\| + (1 + \epsilon_2)/\alpha$.*

Proof: The expected payoff vector $u^1_{G_2}$ of player 1 is $A\boldsymbol{y}$. By construction of matrix A, the expected payoff vector for playing pure strategies in block i is $\sum_{i' \in [m]} A^{i,i'} \boldsymbol{y^{i'}}$. The domininant vector in this sum is $A^{i,i}\boldsymbol{y^i}$, whose entries are all $-\alpha\|\boldsymbol{y^i}\|$. The entries of every other vector $A^{i,i'}\boldsymbol{y^{i'}}$ in the sum are in the range $[0, \|\boldsymbol{y^{i'}}\|]$, and since $\boldsymbol{y}$ is a distribution vector, the total contribution to the sum is at most $\sum_{i' \in [m]} \|\boldsymbol{y^{i'}}\| = 1$. Thus, the expected payoff for playing any pure strategy in block i is in the range $[-\alpha\|\boldsymbol{y^i}\|, -\alpha\|\boldsymbol{y^i}\| + 1]$. Assume for contradiction that $\|\boldsymbol{y^i}\| > \|\boldsymbol{y^{i'}}\| + (1 + \epsilon_2)/\alpha$. Then the expected payoff for playing a pure strategy in block i is at most $-\alpha(\|\boldsymbol{y^{i'}}\| + (1 + \epsilon_2)/\alpha) + 1$, while the expected payoff for playing in block i' is at least $-\alpha(\|\boldsymbol{y^{i'}}\|)$. The difference is more than ϵ_2, contradicting the assumption that i belongs to support$(\boldsymbol{x})$. ■

If a game is both imitation and block ϵ-uniform, then the weight of player 2 is divided ϵ-uniformly among all blocks in $[m]$.

Corollary 4.5 (Imitation and Block ϵ-Uniform) *Let $\boldsymbol{x}, \boldsymbol{y}$ be an ϵ_2-well-supported Nash equilibrium of a block ϵ_2-uniform imitation game G_2, where $\epsilon_2 \leq 1/N$. Then for every two blocks $i, i' \in [m]$, $\|\boldsymbol{y^i}\| = \|\boldsymbol{y^{i'}}\| \pm (1 + \epsilon_2)/\alpha$.*

Proof: Since $\boldsymbol{y}$ is a distribution vector, there exists a block $i \in [m]$ such that $\|\boldsymbol{y^i}\| \geq 1/m$. So i belongs to the support of $\boldsymbol{y}$, and by Lemma 4.3, i also belongs to the support of $\boldsymbol{x}$. By Lemma 4.4, $1/m \leq \|\boldsymbol{y^i}\| \leq \|\boldsymbol{y^{i'}}\| + (1+\epsilon_2)/\alpha$ for every $i' \in [m]$. Since $(1+\epsilon_2)/\alpha < 1/m$ we conclude that $0 < \|\boldsymbol{y^{i'}}\|$ for every i'. Thus by Lemma 4.3 all blocks are in support($\boldsymbol{x}$) and get almost uniform weight. ■

4.2 The Reduction

Given an input pair (G_m, ϵ_m), we show how to find an output pair (G_2, ϵ_2), where G_2 has payoffs in the range $[-\alpha, 1]$. To complete the reduction, G_2 can then be normalized by adding α to all payoffs and scaling by $1/(\alpha+1)$ (ϵ_2 is also scaled). Let $N = \sum_{i=1}^{m} n_i$ be the total number of pure strategies in G_m. Let $\epsilon_2 = \epsilon_m/N$. The pure strategies of every player in G_2 are the set $[N]$. The payoffs are chosen such that G_2 is both an imitation game and a block ϵ_2-uniform game:

$$A = \begin{pmatrix} -\alpha E_{n_1} & M^{1,2} & \cdots & M^{1,m} \\ M^{2,1} & -\alpha E_{n_2} & & M^{2,m} \\ \vdots & & \ddots & \vdots \\ M^{m,1} & M^{m,2} & \cdots & -\alpha E_{n_m} \end{pmatrix}_{N \times N}, B = \begin{pmatrix} I_{n_1} & 0 & \cdots & 0 \\ 0 & I_{n_2} & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & I_{n_m} \end{pmatrix}_{N \times N}$$

where $M^{i,i'}$ is the payoff matrix of player i for interacting with player i' in G_m.

4.3 Correctness

Proof of Theorem 4.1 (Linear to Bimatrix): First note that the reduction runs in time polynomial in $|G_m| = \Theta(N^2)$ and in $\log(1/\epsilon_m)$: The running time depends on the size of the bimatrix game G_2, whose payoff matrices are of size N^2 with entries of size $O(\log \alpha)$. It's enough to prove Lemma 4.2 for the unnormalized game G_2 and $\epsilon_2 = \epsilon_m/N$; this immediately gives a proof for $\epsilon_2 = \epsilon_m/N(\alpha+1)$ after normalizing the payoffs from $[-\alpha, 1]$ to $[0, 1]$.[3] Since $\epsilon_m/N(\alpha+1) = \text{poly}(\epsilon_m/N)$, Theorem 4.1 follows.

■

Proof of Lemma 4.2 (Recovering ϵ_m-Well-supported Nash from Bimatrix Game): Let $(\boldsymbol{x}, \boldsymbol{y})$ be an ϵ_2-well-supported Nash equilibrium played in G_2, and let $(\boldsymbol{y^1}/\|\boldsymbol{y^1}\|, \ldots, \boldsymbol{y^m}/\|\boldsymbol{y^m}\|)$ be a mixed strategy profile played in G_m. We show that this mixed strategy profile is actually an ϵ_m-well-supported Nash equilibrium of G_m.

For every player i of G_m, we define an injective function $h_i : [n_i] \to [N]$ to be $h_i(j) = \sum_{i'<i} n_{i'} + j$. So h_i maps the j'th pure strategy of player i in G_m to the j'th pure strategy in block i of player 1 in G_2. We now show that player i's expected payoff for playing j in G_m is closely related to player 1's expected payoff for playing $h_i(j)$ in G_2, assuming strategy profiles $(\boldsymbol{x}, \boldsymbol{y})$ and $(\boldsymbol{y^1}/\|\boldsymbol{y^1}\|, \ldots, \boldsymbol{y^m}/\|\boldsymbol{y^m}\|)$ are being played in G_2 and G_m, respectively. In fact, the expected payoffs are the same up to shifting by $\alpha\|\boldsymbol{y^i}\|$ (the contribution from the diagonal of player 1's payoff matrix A), scaling by m (the number of blocks on which y is uniformly distributed), and small additive errors. As in Section 3, the fact that the expected payoffs are preserved, even up to shift and scale, is enough for one game's ϵ-well-supported Nash equilibrium to imply the other's.

Claim 4.6 (Expected Payoffs are Preserved up to Shift and Scale) $\boldsymbol{u^i_{G_m}}[j] = m \cdot (\boldsymbol{u^1_{G_2}}[h_i(j)] + \alpha\|\boldsymbol{y^i}\|) \pm m^2(1+\epsilon_2)/\alpha$.

Proof: By construction of matrix A, player 1's expected payoff vector for playing pure strategies in block i is $A^{i,i}\boldsymbol{y^i} + \sum_{i' \neq i} M^{i,i'}\boldsymbol{y^{i'}}$. The entries of vector $A^{i,i}\boldsymbol{y^i}$ are $-\alpha\|\boldsymbol{y^i}\|$, and the sum $\sum_{i' \neq i} M^{i,i'}\boldsymbol{y^{i'}}$ equals

[3] Note that every $\epsilon_m/N(\alpha+1)$-well-supported Nash equilibrium of the normalized game is an ϵ_m/N-well-supported Nash equilibrium of the unnormalized game.

$\|\boldsymbol{y^{i'}}\| \cdot \sum_{i' \neq i} M^{i,i'} \boldsymbol{y^{i'}} / \|\boldsymbol{y^{i'}}\| = \|\boldsymbol{y^{i'}}\| \cdot \boldsymbol{u^i_{G_m}}$. By Corollary 4.5, and since $(\boldsymbol{x}, \boldsymbol{y})$ in an ϵ_2-well-supported Nash equilibrium of G_2, player 2's weight is distributed evenly over the m blocks up to $(1+\epsilon_2)/\alpha$. It is not hard to see that Corollary 4.5 implies, for every $i' \in [m]$, that $\|\boldsymbol{y^{i'}}\| = 1/m \pm (1+\epsilon_2)/\alpha$. Plugging in, we get that entry $h_i(j)$ in player 1's expected payoff vector is $\boldsymbol{u^1_{G_2}}[h_i(j)] = -\alpha\|\boldsymbol{y^i}\| + (1/m \pm (1+\epsilon_2)/\alpha) \cdot \boldsymbol{u^i_{G_m}}[j]$. The proof is complete by noting that player i's expected payoff $\boldsymbol{u^i_{G_m}}[j]$ in G_m is bounded by m, and by rearranging. ■

It's left to show that preservation of expected payoffs for playing $h_i(j)$ and j up to shift and scale is enough to ensure that $(\boldsymbol{y^1}/\|\boldsymbol{y^1}\|, \ldots, \boldsymbol{y^m}/\|\boldsymbol{y^m}\|)$ is an ϵ_m-well-supported Nash equilibrium of G_m. More precisely, we show that if pure strategy $h_i(j)$ is an ϵ_2-best response for player 1 in G_2, then pure strategy j is an ϵ_m-best response for player i in G_m. We can then invoke Claim 4.3 by which player 2 only plays pure strategies that are ϵ_2-best responses for player 1, and conclude that mixed strategy $\boldsymbol{y^i}/\|\boldsymbol{y^i}\|$ contains only ϵ_m-best responses for player i in G_m.

Assume for contradiction that j is not an ϵ_m-best response for player i in G_m. Then there exists another pure strategy $j' \in [m]$ such that $\boldsymbol{u^i_{G_m}}[j] < \boldsymbol{u^i_{G_m}}[j'] - \epsilon_m$. But by Claim 4.6 this implies $\boldsymbol{u^1_{G_2}}[h_i(j)] < \boldsymbol{u^1_{G_2}}[h_i(j')] + 2m(1+\epsilon_2)/\alpha - \epsilon_m/m$. By choice of ϵ_2 and α, $\epsilon_2 \leq \epsilon_m/m - 2m(1+\epsilon_2)/\alpha$. Thus, $h_i(j)$ cannot be an ϵ_2-best response for player 1 in G_2, contradiction. This completes the proof of Lemma 4.2. ■

5 A Logarithmic-Sized Linear Multiplication Gadget

In this section we prove Lemma 2.3 by showing an alternative contruction of a linear multiplication gadget. The main difference from the construction shown in Section 2 is that the unary encoding is replaced by binary encoding. However, this introduces a new difficulty, since every gadget that performs binary bit extraction is inherently *brittle*, i.e., its output is arbitrary for certain inputs. We use the bit extraction gadget of [DGP09], and overcome the brittleness using standard methods of averaging (somewhat simplified by introducing a new median gadget).

5.1 Linear Gadgets

We introduce several linear gadgets that will be useful for the construction. Additional gadgets that are known from previous works can be found in Appendix A. Throughout, we denote the input, output and auxiliary players of a gadget by $P_1, \ldots, P_m, P, W, W_1, \ldots, W_l$, and the values they represent by $p_1, \ldots, p_m, p, w, w_1, \ldots, w_l$.

The following gadget G_{mask} treats its first input as a binary mask for its second input (i.e., performs multiplication between a binary input and an arbitrary input while maintaining linearity). Furthermore, it guarantees that if the second input is close to zero, the output will be close to zero as well.

Lemma 5.1 (Linear Mask Gadget) *There exists a linear gadget G_{mask} such that in every ϵ-well-supported Nash equilibrium:*

$$p = \begin{cases} p_2 \pm \epsilon & \text{if } p_1 = 1 \\ 0 & \text{if } p_1 = 0 \\ 0 \pm 3\epsilon & \text{if } p_2 = 0 \pm 2\epsilon \end{cases}$$

Proof: Nonzero payoff matrices:

$$M^{P,W} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, M^{W,P} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, M^{W,P_1} = \begin{pmatrix} 2 & 0 \\ 0 & 0 \end{pmatrix}, M^{W,P_2} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$$

Expected payoff vectors:

- $\boldsymbol{u^W} = M^{W,P}\boldsymbol{p^P} + M^{W,P_1}\boldsymbol{p^{P_1}} + M^{W,P_2}\boldsymbol{p^{P_2}} = (p + 2(1-p_1), p_2)$;
- $\boldsymbol{u^P} = M^{P,W}\boldsymbol{p^W} = (1-w, w)$.

Assume G_{mask} is in ϵ-well-supported Nash equilibrium. First we show that p_1 is a binary mask for p_2. If $p_1 = 0$, the only ϵ-best response for player W is 0, and so $w = 0$ and also $p = 0$ (actually this holds whenever $p_1 < 1/2$). If $p_1 = 1$, player W's expected payoffs for playing strategies 0 and 1 are p and p_2 respectively. We claim that p must be equal to p_2 up to $\pm\epsilon$. Indeed, if $p > p_2 + \epsilon$, $w = 0$ and so $p = 0$, contradiction. Similarly, if $p < p_2 - \epsilon$, $w = 1$ and so $p = 1$, contradiction.

It is left to show that if p_2 is close to zero then p is also close to zero. Assume for contradiction that $p_2 \leq 2\epsilon$ and $p > 3\epsilon$. Pure strategy 1 must be an ϵ-best response for P, therefore $w \geq (1-w) - \epsilon$. So pure strategy 1 must be an ϵ-best response for W, therefore $p_2 \geq p + 2(1-p_1) - \epsilon$. Plugging in we get $2\epsilon \geq p_2 \geq p + 2(1-p_1) - \epsilon > 3\epsilon + 2(1-p_1) - \epsilon = 2\epsilon + 2 - 2p_1$. This implies that $p_1 > 1$, contradiction. ■

Lemma 5.2 (Linear Max Gadget) *There exists a linear gadget $G_{\max}$ such that in every ϵ-well-supported Nash equilibrium, $p = \max\{p_1, p_2\} \pm 4\epsilon$.*

Proof: The construction is by combining gadgets:

$$\begin{aligned} w_1 &= G_{<}(p_1, p_2) \\ w_2 &= G_{-}(p_2, p_1) \\ w_3 &= G_{\text{mask}}(w_1, w_2) \\ p &= G_{+,*1}(w_3, p_1) \end{aligned}$$

The correctness follows almost immediately from the guarantees of the combined gadgets (see Lemma 2.9, Lemma 5.1, Lemma A.1, Lemma A.2). The idea is to set the output p to be approximately equal to $w_1(p_2 - p_1) + p_1$, where w_1 is an indicator whether $p_1 < p_2$. Assume $G_{\max}$ is in ϵ-well-supported Nash equilibrium. If $p_1 < p_2 - \epsilon$ then $w_1 = 1$, and so p is approximately equal to p_2. Similarly, if $p_1 > p_2 + \epsilon$ then $w_1 = 0$, and so p is approximately equal to p_1. In the case that $|p_2 - p_1| \leq \epsilon$, w_1 may receive any arbitrary value, but $w_2 \leq 2\epsilon$ and so by the guarantee of G_{mask}, $w_3 \leq 3\epsilon$. The product $w_1(p_2 - p_1)$ is calculated by G_{mask} in order to maintain the linearity of the construction. ■

Lemma 5.3 (Linear Min Gadget) *There exists a linear gadget $G_{\min}$ such that in every ϵ-well-supported Nash equilibrium, $p = \min\{p_1, p_2\} \pm 8\epsilon$.*

Proof: The construction is by combining gadgets:

$$\begin{aligned} w_1 &= G_{1-x}(p_1) \\ w_2 &= G_{1-x}(p_2) \\ w_3 &= G_{\max}(w_1, w_2) \\ p &= G_{1-x}(w_3) \end{aligned}$$

The correctness follows immediately by the guarantees of the combined gadgets (see Lemma 5.2, Lemma A.3). ■

Lemma 5.4 (Linear Median Gadget) *There exists a linear gadget G_{median} such that in every ϵ-well-supported Nash equilibrium, $p = median\{p_1, p_2, p_3\} \pm 20\epsilon$.*

Proof: The construction is by combining gadgets:

$$\begin{aligned} w_1 &= G_{\max}(p_1, p_2) \\ w_2 &= G_{\min}(p_1, p_2) \\ w_3 &= G_{\min}(p_3, w_1) \\ p &= G_{\max}(w_2, w_3) \end{aligned}$$

Assume G_{median} is in ϵ-well-supported Nash equilibrium. We use the following notation to prove correctness: $s = \min\{p_1, p_2, p_3\}$, $m = \text{median}\{p_1, p_2, p_3\}$ and $l = \max\{p_1, p_2, p_3\}$, such that $s \le m \le l$. The values w_2 and w_3 are equal to two of the three values $\{p_1, p_2, p_3\} = \{s, m, l\}$, up to an error of $\pm 12\epsilon$ introduced by the maximum and minimum gadgets. To see this, notice that if $w_1 = p_{i\in\{1,2\}} \pm 4\epsilon$, then $w_2 = p_{j\in\{1,2\}, j\ne i} \pm 4\epsilon$ and $w_3 = \min\{p_3, p_i \pm 4\epsilon\} \pm 8\epsilon$ (by Lemma 5.2 and Lemma 5.3). Now, if $m < l - 4\epsilon$, both w_2 and w_3 must be strictly smaller than l, so the two values they are approximately equal to must be s and m. Taking their maximum therefore results in the median value: $\max\{w_2, w_3\} = m \pm 12\epsilon$, and so $p = m \pm (12\epsilon + 4\epsilon)$ (Lemma 5.2). If, however, $m \ge l - 4\epsilon$, then w_2 and w_3 are approximately equal to either $\{s, m\}$, $\{s, l\}$ or $\{m, l\}$. Taking the maximum of $\{w_2, w_3\}$ therefore results in either $m \pm (12\epsilon + 4\epsilon)$ or $l \pm (12\epsilon + 4\epsilon) = m \pm 20\epsilon$. ■

5.2 Brittle Construction

We first construct a brittle multiplication gadget denoted by $\tilde{G}_*$. Let $\beta = \frac{1}{2}\log(1/\epsilon)$ be a parameter of the construction (assume for simplicity that β is integral). The output p of $\tilde{G}_*$ is guaranteed to be $p_1 p_2 \pm O(\sqrt{\epsilon})$, but only as long as the input p_1 is far enough from any integer multiple of $2^{-\beta}$. For every $1 \le i \le \beta$, let B_i, S_i, W_i be auxiliary players representing values b_i, s_i, w_i respectively. $\tilde{G}_*$ sets $b_1, \ldots, b_\beta$ to be the β most significant bits of input p_1 using the bit-extraction gadget G_{bit} (Lemma A.6). Then, $\tilde{G}_*$ calculates $p_2 \sum_{i=1}^{\beta} \left(b_i 2^{-i}\right)$, which equals $p_2\left(p_1 \pm 2^{-\beta}\right)$. The calculations are carried out in the following order:

- The values $\left\{p_2 2^{-i}\right\}_{i\in[\beta]}$ are calculated using the scaling gadget $G_{*\zeta}$ (Lemma A.5), i.e., $s_i = G_{*2^{-i}}(p_2)$.
- For every $i \in [\beta]$, $p_2 2^{-i}$ is multiplied by the extracted bit b_i using the mask gadget G_{mask} (Lemma 5.1), i.e., $w_i = G_{\text{mask}}(b_i, s_i)$.
- The values $\{p_2 2^{-i} b_i\}_{i\in[\beta]}$ are summed up using the summation gadget $G_+ = G_{+,*1}$ (Lemma 2.9), i.e., $p = G_+(w_1, \ldots, w_\beta)$.

5.3 Correctness of Brittle Construction

Lemma 5.5 *Let p_1, p_2 be the input values of $\tilde{G}_*$. If p_1 is $3\beta\epsilon$-far from every integer multiple of $2^{-\beta}$, then in every ϵ-well-supported Nash equilibrium where $\epsilon \le \frac{1}{10^3}$, the output value p equals $p_1 p_2 \pm 2\sqrt{\epsilon}$. The size of $\tilde{G}_*$ is $O(\beta)$.*

Proof: First observe that since $\beta = \frac{1}{2}\log(1/\epsilon)$ (i.e., $2^{-\beta} = \sqrt{\epsilon}$) and $\epsilon \le \frac{1}{10^3}$, it holds that $2 \cdot 3\beta\epsilon < 2^{-\beta}$, and so p_1 can indeed be $3\beta\epsilon$-far from any integer multiple of $2^{-\beta}$. We write p_1 as $p_1 = \sum_{i=1}^{\beta} b_i^* 2^{-i} + \delta$, where $3\beta\epsilon < \delta < 2^{-\beta} - 3\beta\epsilon$. Assume $\tilde{G}_*$ is in ϵ-well-supported Nash equilibrium where $\epsilon \le \frac{1}{10^3}$. By Lemma 5.1, Lemma A.6 and Lemma A.5, we know that for every $i \in [\beta]$:

$$\begin{aligned} b_i &= b_i^* \\ s_i &= p_2 2^{-i} \pm \epsilon \\ w_i &= s_i b_i \pm \epsilon = p_2 2^{-i} b_i^* \pm 2\epsilon \end{aligned}$$

By Lemma 2.9 and the value of δ:

$$\begin{aligned} p &= \min\left\{1, \sum_{i=1}^{\beta} w_i\right\} \pm \epsilon \\ &= \min\left\{1, \sum_{i=1}^{\beta} \left(p_2 2^{-i} b_i^* \pm 2\epsilon\right)\right\} \pm \epsilon \\ &= \min\left\{1, p_2 \sum_{i=1}^{\beta} \left(2^{-i} b_i^*\right) \pm 2\beta\epsilon\right\} \pm \epsilon \\ &= \min\left\{1, p_2\left(p_1 - \delta\right) \pm 2\beta\epsilon\right\} \pm \epsilon \\ &= p_1 p_2 \pm \left(2^{-\beta} + 3\beta\epsilon\right) \end{aligned}$$

Plugging in $\beta = \frac{1}{2}\log(1/\epsilon)$, we get that $p = p_1 p_2 \pm 2\sqrt{\epsilon}$, as required. Size of $\tilde{G}_*$: The bit-extraction gadget G_{bit} requires $O(\beta)$ vertices (Lemma A.6), and the number of auxiliary vertices $\{S_i\}$ and $\{W_i\}$ is also $O(\beta)$. The other gadgets are of constant size. ■

5.4 Robust Construction

When p_1 is close to a multiple of $2^{-\beta}$, $\tilde{G}_*$'s output may be arbitrary. To circumvent this issue, the ultimate multiplication gadget G_* applies $\tilde{G}_*$ three times, each time with a slightly perturbed copy of the input p_1. The perturbation guarantees that at most one of the three copies of p_1 is close to an integer multiple of $2^{-\beta}$, so that at least two of $\tilde{G}_*$'s three outputs are approximately correct. The difficulty is that we don't know which of the three outputs is approximately correct and which is arbitrary. We overcome this difficulty by taking the median of the three outputs as the final result, which is now guaranteed to be approximately equal to the required output $p_1 p_2$.

The inputs to $\tilde{G}_*$ are set to be (up to $\pm O(\epsilon)$): $(\tilde{p_1}, p_2)$, $(\tilde{p_1} - \Delta, p_2)$ and $(\tilde{p_1} - 2\Delta, p_2)$, where $\tilde{p_1} = \max\{p_1, 2\Delta\}$ and $\Delta = 7\beta\epsilon$. This is achieved by combining the following gadgets:

- First, the value of $\tilde{p_1}$ is set: $c_1 = G_{:=}(2\Delta + 7\epsilon), \tilde{p_1} = G_{\max}(p_1, c_1)$.
- Then, two additional inputs are prepared: $c_2 = G_{:=}(\Delta), c_3 = G_{:=}(2\Delta), d_1 = G_{-}(\tilde{p_1}, c_2), d_2 = G_{-}(\tilde{p_1}, c_3)$.
- $\tilde{G}_*$ is applied: $w_1 = \tilde{G}_*(\tilde{p_1}, p_2), w_2 = \tilde{G}_*(d_1, p_2), w_3 = \tilde{G}_*(d_2, p_2)$.
- The median is found: $p = G_{\text{median}}(w_1, w_2, w_3)$.

5.5 Correctness of Robust Construction

Proof of Lemma 2.3 (Logarithmic-Sized Construction): First note that G_* is a combination of linear gadgets and is thus itself linear. The size of G_* is $O(\beta) = O(\log \frac{1}{\epsilon})$, since the brittle multiplication gadget $\tilde{G}_*$ requires $O(\beta)$ vertices (Lemma 5.5), and the number of other auxiliary vertices is constant.

Assume G_* is in ϵ-well-supported Nash equilibrium where $\epsilon \leq \frac{1}{10^5}$. By the gadget guarantees we know that $\tilde{p_1} \geq 2\Delta + 2\epsilon$, and that $d_1 = \tilde{p_1} - \Delta \pm 2\epsilon$ and $d_2 = \tilde{p_1} - 2\Delta \pm 2\epsilon$ (Lemma 5.2, Lemma A.2 and Lemma A.4). Since $\epsilon < \frac{1}{10^5}, \Delta = 7\beta\epsilon$ and $\beta = \frac{1}{2}\log\frac{1}{\epsilon}$, it can easily be verified that $\tilde{p_1} > d_1 > d_2 \geq 0$ and that the distance between each consecutive pair is $\Delta \pm 4\epsilon$.

Claim 5.6 *At most one of $\tilde{p_1}, d_1, d_2$ can be $3\beta\epsilon$-close to a multiple of $2^{-\beta}$.*

Proof: Since the distance $\Delta \pm 4\epsilon$ is larger than $2 \cdot 3\beta\epsilon$, if one of $\tilde{p_1}, d_1, d_2$ is $3\beta\epsilon$-close to a certain multiple $k2^{-\beta}$, then the other two must be $3\beta\epsilon$-far from $k2^{-\beta}$. Furthermore, since the distance is smaller than

$(2^{-\beta} - 2 \cdot 3\beta\epsilon)/2$, the other two must be $3\beta\epsilon$-far from the nearby multiples $(k-1)\,2^{-\beta}$ and $(k+1)\,2^{-\beta}$ as well. ■

By Lemma 5.5 and Claim 5.6, at most one of w_1, w_2, w_3 can be arbitrary. There are two cases:

- The median is not the arbitrary value. Assume without loss of generality that the median is w_3 (since it is the furthest from $\tilde{p}_1 p_2$). By Lemma 5.4 and Lemma 5.5:

$$\begin{aligned} p &= w_3 \pm 20\epsilon \\ &= d_2 p_2 \pm \left(2\sqrt{\epsilon} + 20\epsilon\right) \\ &= \tilde{p}_1 p_2 \pm \left(2\sqrt{\epsilon} + 2\Delta + 22\epsilon\right) \end{aligned}$$

- The median is the arbitrary value. Assume without loss of generality that the non-arbitrary values are w_2 and w_3 (the furthest from $\tilde{p}_1 p_2$). The median is between these values, so we may assume without loss of generality that it is equal to w_3, and proceed as in the previous case.

We have seen that in both cases, p is close to $\tilde{p}_1 p_2$. It is now left to verify that $\tilde{p}_1$ is close to p_1.

Claim 5.7 $\tilde{p}_1 = p_1 \pm (2\Delta + 11\epsilon)$.

Proof: If $\max\{p_1, c_1\} = p_1$ then $\tilde{p}_1 = p_1 \pm 4\epsilon$ (Lemma 5.2). Otherwise, $0 \le p_1 \le c_1 = 2\Delta + 6\epsilon \pm \epsilon$ (Lemma A.4) and $\tilde{p}_1 = c_1 \pm 4\epsilon$. ■

We conclude that $p = p_1 p_2 \pm (2\sqrt{\epsilon} + 4\Delta + 37\epsilon) = p_1 p_2 \pm 3\sqrt{\epsilon}$, as required. ■

References

[Alt94] I. Althofer. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and its Applications*, 240:9–19, 1994.

[BBM07] H. Bosse, J. Byrka, and E. Markakis. New algorithms for approximate Nash equilibria in bimatrix games. In *WINE*, 2007.

[Bub79] V. Bubelis. On equilibria in finite games. *International Journal of Game Theory*, 8(2):65–79, 1979.

[CD06a] X. Chen and X. Deng. On the complexity of 2d discrete fixed point problem. In *ICALP*, 2006.

[CD06b] X. Chen and X. Deng. Settling the complexity of 2-player Nash-equilibrium. In *FOCS*, 2006.

[CDT06] X. Chen, X. Deng, and S. Teng. Computing Nash equilibria: Approximation and smoothed complexity. In *FOCS*, 2006.

[DFP06] C. Daskalakis, A. Fabrikant, and C. H. Papadimitriou. The game world is flat: The complexity of Nash equilibria in succinct games. In *ICALP*, 2006.

[DGP09] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

[DMP06] C. Daskalakis, A. Mehta, and C. H. Papadimitriou. Progress in approximate Nash equilibria. In *EC*, 2006.

[DP09] C. Daskalakis and C. H. Papadimitriou. On oblivious ptas's for Nash equilibrium. In *STOC*, 2009.

[EY07] K. Etessami and M. Yannakakis. On the complexity of Nash equilibria and other fixed points. In *FOCS*, 2007.

[GKT50] D. Gale, H. W. Kuhn, and A. W. Tucker. On symmetric games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, pages 81–87. Princeton, 1950.

[GP06] P. W. Goldberg and C. H. Papadimitriou. Reducibility among equilibrium problems. In *STOC*, 2006.

[GW10] S. Govindan and R. Wilson. A decomposition algorithm for n-player games. *Economic Theory*, 42(1):97–117, 2010.

[HK09] E. Hazan and R. Krauthgamer. How hard is it to approximate the best Nash equilibrium? In *SODA*, 2009.

[KS10] S. C. Kontogiannis and P. G. Spirakis. Well supported approximate equilibria in bimatrix games. *Algorithmica*, 57(4):653–667, 2010.

[LH64] C. E. Lemke and J. T. Howson. Equilibrium points of bimatrix games. *SIAM Journal of Applied Mathematics*, 12:413–423, 1964.

[LMM03] R. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *EC*, 2003.

[MT09] A. McLennan and R. Tourky. Imitation games and computation. *Games and Economic Behavior*, 2009.

[Nas51] J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:289–295, 1951.

[Pap94] 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.

[Pap07] C. H. Papadimitriou. The complexity of finding Nash equilibria. In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 2. Cambridge University Press, 2007.

[Rou10] T. Roughgarden. Computing equilibria: A computational complexity perspective. *Economic Theory*, 42(1):193–236, 2010.

[TS07] H. Tsaknakis and P. G. Spirakis. An optimization approach for approximate Nash equilibria. In *WINE*, 2007.

[TS09] H. Tsaknakis and P. G. Spirakis. A graph spectral approach for computing approximate Nash equilibria. In *ECCC*, 2009.

[vdLT82] G. van der Laan and A. J. J. Talman. On the computation of fixed points in the product space of unit simplices and an application to noncooperative n-person games. *Mathematics of Operations Research*, 7(1):1–13, 1982.

[vS07] B. von Stengel. Equilibrium computation for two-player games in strategic and extensive form. In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 3. Cambridge University Press, 2007.

A Standard Gadgets

The following gadgets are constructed by Daskalakis et al. [DGP09]. We denote the input and output players by P_1, P_2, P, and the values they represent by p_1, p_2, p.

Proof of Lemma 2.8 (Linear AND Gadget):

Nonzero payoff matrices:

$$M^{P,P_1} = M^{P,P_2} = \begin{pmatrix} \frac{3}{16} & \frac{3}{16} \\ 0 & \frac{1}{2} \end{pmatrix}$$

Expected payoff vectors:

- $\boldsymbol{u^P} = M^{P,P_1}\boldsymbol{p^1} + M^{P,P_2}\boldsymbol{p^2} = (3/4, (p_1 + p_2)/2)$.

Assume $G_{>\zeta}$ is in ϵ-well-supported Nash equilibrium where $\epsilon < 1/4$. If $p_1 = p_2 = 1$, the only ϵ-best response for player P is pure strategy 1, so $p = 1$. Similarly, if $(p_1 = 0) \vee (p_2 = 0)$ then $p = 0$. ■

Proof of Lemma 2.9 (Linear Scaled-Summation Gadget):

Let $P_1, \ldots, P_m, P, W$ be the input players, output player and auxiliary player of $G_{+,*\zeta}$ respectively, representing values $p_1, \ldots, p_m, p, w$. Nonzero payoff matrices:

$$M^{W,P} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, M^{W,P_i} = \begin{pmatrix} 0 & 0 \\ 0 & \zeta \end{pmatrix}, M^{P,W} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

Expected payoff vectors:

- $\boldsymbol{u^W} = M^{W,P}\boldsymbol{p^P} + \sum_{i\in[m]} M^{W,P_i}\boldsymbol{p^{P_i}} = (p, \zeta \sum_{i\in[m]} p_i)$;
- $\boldsymbol{u^P} = M^{P,W}\boldsymbol{p^W} = (1 - w, w)$.

Assume $G_{+,*\zeta}$ is in ϵ-well-supported Nash equilibrium. If player W plays full support $(0 < w < 1)$, then both of W's pure strategies 0 and 1 must be ϵ-best responses, and so $p = \boldsymbol{u^W}[0] = \boldsymbol{u^W}[1] \pm \epsilon = \zeta \sum_{i\in[m]} p_i \pm \epsilon$, as required. If $w = 0$, the only ϵ-best response for player P is pure strategy 0, so $p = 0$. Similarly, if $w = 1$ then $p = 1$. Case analysis:

- $\epsilon < \zeta \sum_{i\in[m]} p_i < 1 - \epsilon$: Assume for contradiction that W does not play full support. Without loss of generality, assume $w = 1$. But then $p = 1$ and $\boldsymbol{u^W}[0] = 1 > \zeta \sum_{i\in[m]} p_i + \epsilon = \boldsymbol{u^W}[1] + \epsilon$, contradiction. Similarly, $w = 0$ leads to contradiction.
- $\zeta \sum_{i\in[m]} p_i \leq \epsilon$: Player W can either play full support or pure strategy 0 (if $w = 1$ then $p = 1$ and the only ϵ-best response for W is pure strategy 0, contradiction). If $w = 0$ then $p = 0 = \zeta \sum_{i\in[m]} p_i \pm \epsilon$, as required.
- $\zeta \sum_{i\in[m]} p_i \geq 1 - \epsilon$: Similarly to the previous case, W can either play full support or pure strategy 1, and if $w = 1$ then $p = 1 = \min\{1, \zeta \sum_{i\in[m]} p_i\} \pm \epsilon$, as required.

■

Lemma A.1 (Linear Comparison Gadget) *There exists a linear comparison gadget $G_<$ of size $O(1)$, such that in every ϵ-well-supported Nash equilibrium, $p = 1$ if $p_1 < p_2 - \epsilon$ and $p = 0$ if $p_1 > p_2 + \epsilon$.*

Proof: Nonzero payoff matrices:

$$M^{P,P_1} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, M^{P,P_2} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$$

Expected payoff vectors:

- $\boldsymbol{u^P} = M^{P,P_1}\boldsymbol{p^1} + M^{P,P_2}\boldsymbol{p^2} = (p_1, p_2)$.

Assume $G_<$ is in ϵ-well-supported Nash equilibrium. If $p_1 < p_2 - \epsilon$, the only ϵ-best response for player P is pure strategy 1, so $p = 1$. Similarly, if $p_1 > p_2 + \epsilon$ then $p = 0$. ■

Lemma A.2 (Linear Minus Gadget) *There exists a linear subtraction gadget G_- of size $O(1)$, such that in every ϵ-well-supported Nash equilibrium, $p = \max\{0, p_2 - p_1\} \pm \epsilon$.*

Proof: Nonzero payoff matrices:

$$M^{W,P} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, M^{W,P_1} = \begin{pmatrix} 0 & 0 \\ 0 & -1 \end{pmatrix}, M^{W,P_2} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, M^{P,W} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

Expected payoff vectors:

- $\boldsymbol{u^W} = M^{W,P}\boldsymbol{p^P} + M^{W,P_1}\boldsymbol{p^1} + M^{W,P_2}\boldsymbol{p^2} = (p, p_2 - p_1)$;
- $\boldsymbol{u^P} = M^{P,W}\boldsymbol{p^W} = (1 - w, w)$.

Assume G_- is in ϵ-well-supported Nash equilibrium. As in the proof of Lemma 2.9, it is not hard to show that either player W plays full support (so both of W's pure strategies must be ϵ-best responses and $p = p_2 - p_1 \pm \epsilon$), or one of the following happens:

- $p_2 - p_1 > 1 - \epsilon$: Player W can play pure strategy 1, and then $p = 1 = p_2 - p_1 \pm \epsilon$, as required.
- $p_2 - p_1 < \epsilon$: Player W can play pure strategy 0, and then $p = 0 = \max\{0, p_2 - p_1\} \pm \epsilon$, as required.

■

Lemma A.3 (Linear Complementary Gadget) *There exists a linear complementary gadget G_{1-x} of size $O(1)$, such that in every ϵ-well-supported Nash equilibrium, $p = 1 - p_1 \pm \epsilon$.*

Proof: Nonzero payoff matrices:

$$M^{W,P} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, M^{W,P_1} = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, M^{P,W} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

Expected payoff vectors:

- $\boldsymbol{u^W} = M^{W,P}\boldsymbol{p^P} + M^{W,P_1}\boldsymbol{p^1} = (p, 1 - p_1)$;
- $\boldsymbol{u^P} = M^{P,W}\boldsymbol{p^W} = (1 - w, w)$.

Assume G_{1-x} is in ϵ-well-supported Nash equilibrium. As in the proof of Lemma 2.9, it is not hard to show that either player W plays full support (so both of W's pure strategies must be ϵ-best responses and $p = 1 - p_1 \pm \epsilon$), or one of the following happens:

- $1 - p_1 > 1 - \epsilon$: Player W can play pure strategy 1, and then $p = 1 = 1 - p_1 \pm \epsilon$, as required.
- $1 - p_1 < \epsilon$: Player W can play pure strategy 0, and then $p = 0 = 1 - p_1 \pm \epsilon$, as required.

■

Lemma A.4 (Linear Assignment Gadget) *For every rational $\zeta \in [0, 1]$, there exists a linear assignment gadget $G_{:=\zeta}$ of size $O(1)$, such that in every ϵ-well-supported Nash equilibrium, $p = \zeta \pm \epsilon$.*

Proof: Nonzero payoff matrices:

$$M^{W,P} = \begin{pmatrix} 0 & 1 \\ \zeta & \zeta \end{pmatrix}, M^{P,W} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

Expected payoff vectors:

- $\boldsymbol{u^W} = M^{W,P}\boldsymbol{p^P} = (p, \zeta)$;
- $\boldsymbol{u^P} = M^{P,W}\boldsymbol{p^W} = (1-w, w)$.

Proof of correctness as in Lemma A.3. ■

Lemma A.5 (Linear Scaling Gadget) *For every rational $\zeta \in [0,1]$, there exists a linear scaling gadget $G_{*\zeta}$ of size $O(1)$, such that in every ϵ-well-supported Nash equilibrium, $p = \zeta p_1 \pm \epsilon$.*

Proof: Nonzero payoff matrices:

$$M^{W,P} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, M^{W,P_1} = \begin{pmatrix} 0 & 0 \\ 0 & \zeta \end{pmatrix}, M^{P,W} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

Expected payoff vectors:

- $\boldsymbol{u^W} = M^{W,P}\boldsymbol{p^P} + M^{W,P_1}\boldsymbol{p^1} = (p, \zeta p_1)$;
- $\boldsymbol{u^P} = M^{P,W}\boldsymbol{p^W} = (1-w, w)$.

Proof of correctness as in Lemma A.3. ■

The following gadget has multiple output players, denoted by $B_1, \ldots, B_\beta$ and representing values $b_1, \ldots, b_\beta$. It extracts the first β bits of its input, provided the distance of p_1 from any multiple of $2^{-\beta}$ is at least $3\beta\epsilon$.

Lemma A.6 (Linear Bit Extraction Gadget) *For every integer $\beta > 0$, there exists a linear bit extraction gadget G_{bit} of size $O(\beta)$, such that given input $p_1 = \sum_{i\in[\beta]} b_i^* 2^{-i} + \delta$ where $3\beta\epsilon < \delta < 2^{-\beta} - 3\beta\epsilon$, in every ϵ-well-supported Nash equilibrium where $\epsilon = O\left(2^{-(\beta+\log\beta)}\right)$, $b_i = b_i^*$ for every $i \in [\beta]$.*

Proof: The construction is by combining linear gadgets:

$$\begin{aligned} x_1 &= G_{:=}(p_1) \\ \forall i: b_i &= G_{>2^{-i}}(x_i) \\ \forall i: w_i &= G_{*2^{-i}}(b_i) \\ \forall i: x_{i+1} &= G_{-}(x_i, w) \end{aligned}$$

The correctness follows from the guarantees of the combined gadgets, and by induction on i. See [DGP09, Lemma 19] for details. ■