

Inapproximability of Nash Equilibrium

Aviad Rubinstein*

June 5, 2015

Abstract

We prove that finding an ϵ -approximate Nash equilibrium is PPAD-complete for constant ϵ and a particularly simple class of games: polymatrix, degree 3 graphical games, in which each player has only two actions.

As corollaries, we also prove similar inapproximability results for Bayesian Nash equilibrium in a two-player incomplete information game with a constant number of actions, for relative ϵ -Nash equilibrium in a two-player game, for market equilibrium in a non-monotone market, for the generalized circuit problem defined by Chen et al. [4], and for approximate competitive equilibrium from equal incomes with indivisible goods.

1 Introduction

Nash equilibrium is the central concept in Game Theory. Much of its importance and attractiveness comes from its *universality*: by Nash's Theorem [18], every finite game has at least one. The result that finding a Nash equilibrium is PPAD-complete, and therefore intractable [8, 4] casts this universality in doubt, since it suggests that there are games whose Nash equilibria, though existent, are in any practical sense inaccessible.

Can approximation repair this problem? Chen et al. [4] proved that it is also hard to find an ϵ -approximate Nash equilibrium for any ϵ that is polynomially

*UC Berkeley. I am grateful to Christos Papadimitriou for inspiring discussions, comments, and advice. I also thank Yang Cai for suggesting the connection to Bayesian Nash equilibrium in two-player games. Additional thanks go to Constantinos Daskalakis, Karthik C.S., and anonymous reviewers for helpful comments. This research was supported by NSF grants CCF0964033 and CCF1408635, and by Templeton Foundation grant 3966. This work was done in part at the Simons Institute for the Theory of Computing.

small - even for two-player games. The only remaining hope is a PTAS, i.e. an approximation scheme for any constant $\epsilon > 0$. *Whether there is a PTAS for the Nash equilibrium problem is the most important remaining open question in equilibrium computation.*

For a constant number of players, we are unlikely to prove PPAD-hardness since a quasi-polynomial time approximation algorithm exists [16]. We therefore shift our attention to games with a large number of players. For such games, there is a question of representation: the normal form representation is exponential in the number of the players. Instead, we consider two natural and well-studied concise representations:

Polymatrix games In a polymatrix game, each pair of players simultaneously plays a separate two-player game. Every player has to play the same strategy in every two-player subgame, and her utility is the sum of her subgame utilities. The game is given in the form of the payoff matrix for each two-player game.

Graphical games In a graphical game [15], the utility of each player depends only on the action chosen by a few other players. This game now naturally induces a directed graph: we say that $(i, j) \in E$ if the utility of player j depends on the strategy chosen by player i . When the maximal incoming degree is bounded, the game has a representation polynomial in the number of players and strategies.

Our results

We prove that even for games that are both polymatrix and graphical (for particularly simple graphs) finding an ϵ -approximate Nash equilibrium is intractable.

Theorem 1. *There exists a constant $\epsilon > 0$, such that given a degree 3, bipartite, polymatrix game where each player has two actions, finding an ϵ -approximate Nash equilibrium is PPAD-complete.*

In previous works [4, 8], hardness of Nash equilibrium in two-player games was obtained from hardness of a bipartite polymatrix game. Roughly speaking, the reduction lets each of the two players choose the strategies for the vertices on one side of the bipartite graphical game. This reduction incurs a polynomial blowup in the error - and indeed, as discussed earlier, we do not expect to

obtain PPAD-hardness for ϵ -approximate Nash equilibrium in two-player games. Nevertheless, in Section 8 we show that a similar technique yields an interesting corollary for ϵ -approximate Bayesian Nash equilibrium in two-player games with incomplete information.

Corollary 1 (ϵ -Bayesian Nash equilibrium). *There exists a constant $\epsilon > 0$, such that given a two-player game with incomplete information where each player has a constant number of actions, finding an ϵ -approximate Bayesian Nash equilibrium is PPAD-complete.*

In two-player complete information games, Daskalakis [7] circumvents Lipton et al.’s quasi-polynomial time algorithm by studying a notion of relative (sometimes also called *multiplicative* [13], as opposed to the more standard additive) ϵ -Nash equilibrium. Daskalakis proves that in two player games with payoffs in $[-1, 1]$, finding a relative ϵ -Nash equilibrium is PPAD-complete. One caveat of this result is that the gain from deviation is large compared to the expected payoff because the latter is small due to cancellation of positive and negative payoffs. Namely, the gain from deviation may be very small compared to the expected magnitude of the payoff. Here we answer an open question from [7] by proving that finding a relative ϵ -Nash equilibrium continues to be PPAD-complete even when all the payoffs are positive.

Corollary 2 (Relative ϵ -Nash equilibrium). *There exists a constant $\epsilon > 0$ such that finding a relative ϵ -Nash equilibrium in a bimatrix game with positive payoffs is PPAD-complete.*

The computation of Nash equilibrium is tightly related to computation of equilibrium in markets. In particular, Chen, Paparas, and Yannakakis [5] use a reduction from polymatrix games to prove PPAD-hardness for the computation of equilibrium for every family of utility functions from a very general class, which they call *non-monotone families of utility functions*. In Section 10 we prove the following hardness of approximation for market equilibrium.

Corollary 3 (Non-monotone markets). *Let \mathcal{U} be any non-monotone family of utility functions. There exists a constant $\epsilon_{\mathcal{U}} > 0$ such that given a market M where the utility of each trader is either linear or taken from \mathcal{U} , finding an $\epsilon_{\mathcal{U}}$ -tight approximate market equilibrium is PPAD-hard.*

Although our inapproximability factor is stronger than that showed by Chen et al., the results are incomparable as we require *tight* approximate equilibrium,

the more standard definition which bounds the two-sided error of the market equilibrium. Chen et al., in contrast, prove that even if we allow arbitrary excess supply, finding a $(1/n)$ -approximate equilibrium is PPAD-hard. Furthermore, for the interesting case of CES utilities with parameter $\rho < 0$, they show that there exist markets where every $(1/2)$ -tight equilibrium requires prices that are doubly-exponentially large (and thus require an exponential-size representation). Indeed, for a general non-monotone family \mathcal{U} , the problem of computing a (tight or not) approximate equilibrium may not belong to PPAD. Nevertheless, the important family of additively separable, concave piecewise-linear utilities is known to satisfy the non-monotone condition [5], and yet the computation of (exact) market equilibrium is in PPAD [28]. Therefore,

Corollary 4 (SPLC markets). *There exists a constant $\epsilon > 0$, such that finding an ϵ -tight approximate market equilibrium with additively separable, concave piecewise-linear utilities is PPAD-complete.*

En route to proving our main result, we also prove hardness of approximation for the generalized circuit problem. Generalized circuits are similar to standard algebraic circuits, the main difference being that generalized circuits contain cycles, which allow them to verify fixed points of continuous functions. A generalized circuit induces a constraint satisfaction problem, ϵ -GCIRCUIT [4]: find an assignment for the values on the lines of the circuit, that simultaneously ϵ -approximately satisfies all the constraints imposed by the gates (see Section 3.2 for a formal definition). ϵ -GCIRCUIT was implicitly proven PPAD-complete for exponentially small ϵ by Daskalakis et al. [8], and explicitly for polynomially small ϵ by Chen et al. [4]. Here we prove that it continues to be PPAD-complete for some constant ϵ .

Theorem 2 (Generalized circuit). *There exists a constant $\epsilon > 0$ such that ϵ -GCIRCUIT with fan-out 2 is PPAD-complete.*

The ϵ -GCIRCUIT problem has already proven useful in several works in recent years (e.g [4, 7, 5, 20]). We believe that Theorem 2 will lead to stronger hardness results in many applications in algorithmic game theory and economics. For example, competitive equilibrium with equal incomes (CEEI) is a well-known fair allocation mechanism [10, 27, 26]; however, for indivisible resources a CEEI may not exist. It was shown by Budish [2] that in the case of indivisible resources

there is always an allocation, called A-CEEI, that is approximately fair, approximately truthful, and approximately efficient, for some favorable approximation parameters. This approximation is used in practice to assign business school students to classes [19].

The approximation of CEEI is characterized by two parameters: α , which is a measure of the market clearing error, and β , which quantifies the inequality in initial endowments. Budish’s proof guarantees the existence of an (α^*, β) -CEEI for any $\beta > 0$ (and some parameter α^* that depends on the input). Othman et al. [20] reduced $\Theta(\beta \log(1/\beta))$ -GCIRCUIT with fan-out 2 to the problem of finding the guaranteed approximation, (α^*, β) -CEEI. Theorem 2 now gives the following corollary:

Corollary 5 (A-CEEI). *There exists a constant $\beta > 0$ such that finding an (α^*, β) -CEEI is PPAD-complete.*

1.1 Related works

This work extends our recent paper [22], where we proved PPAD-hardness of ϵ -Nash equilibrium (ϵ -NE)¹, for constant ϵ and a constant number of actions per player, over a larger class of games called *succinct games*. Succinct games can have much more complex utilities than polymatrix or graphical games of bounded degree: any utility function that can be computed in polynomial time is allowed.

We improve over the results of [22] in three ways: (1) simplicity: our new hardness result holds for a much simpler and more natural definition of games; (2) ϵ -ANE vs ϵ -NE: thanks to the transformation to bounded degree graphical games our result extends to the weaker concept (thus stronger hardness) of *ϵ -approximate* Nash equilibrium (see Section 3.1 for precise definitions); and (3) completeness: finding an approximate Nash equilibrium in a bounded-degree graphical game is not only PPAD-hard - it also belongs to PPAD [8]. For succinct games, in contrast, ϵ -NE is unlikely to belong to PPAD since it is also known to be BPP-hard [23].

On the technical side, our current construction of hard instances of Brouwer fixed points is identical to [22]. The main technical contribution in this paper

¹ ϵ -Nash equilibrium is sometimes also called ϵ -well-supported Nash equilibrium; see [7] for a short discussion on terminology.

is the adaptation of the equiangle sampling gadget of Chen et al. [4] to this particular Brouwer function.

Query complexity There are several interesting results on the query complexity of approximate Nash equilibria, where the algorithm is assumed to have black-box access to the exponential-size payoff function.

Hart and Nisan [12] prove that any deterministic algorithm needs to query at least an exponential number of queries to compute any ϵ -Nash equilibrium - and even for any ϵ -correlated equilibrium. For ϵ -correlated equilibrium, on the other hand, Hart and Nisan show a randomized algorithm that uses a number of queries polynomial in n and ϵ^{-1} .

Babichenko [1] recently showed that any randomized algorithm requires an exponential number of queries to find an ϵ -Nash equilibrium. Our proof is inspired by Babichenko's work and builds on some of his techniques.

Yet a newer paper by Goldberg and Roth [11] characterizes the query complexity of approximate coarse correlated equilibrium in games with many players. More important for our purpose is their polynomial upper bound on the query complexity of ϵ -Nash equilibria for any family of games that have *any* concise representation. This result is to be contrasted with (a) Babichenko's query complexity lower bound, which uses a larger family of games, and (b) [22] which applies exactly to this setting and gives a lower bound on the *computational complexity*.

A much older yet very interesting and closely related result is that of Hirsch, Papadimitriou, and Vavasis [14]. Hirsch et al. show that any deterministic algorithm for computing a Brouwer fixed point in the oracle model must make an exponential -in the dimension n and the approximation ϵ - number of queries for values of the function. The techniques in [14] have proven particularly useful both in [1, 22] and here.

An algorithm for $(1/2 + \delta)$ -approximate Nash in any polymatrix game and for any constant $\delta > 0$, was recently shown by Deligkas et al. [9].

2 Proof overview

We begin our proof of Theorem 1 with the ENDOFTHELINE problem over $\{0, 1\}^n$ [8]. Our first step (Section 4) is to reduce this problem to the problem of finding

an approximate fixed point of a particular continuous function, inspired by the work of Hirsch et al. [14] (this reduction also recently appeared in [22]). Our second step (Section 5) is to reduce the problem of finding an approximate fixed point of this particular function to that of finding an approximate solution to a generalized circuit; this step uses techniques introduced by Chen et al. [4]. In Section 6, we further reduce to a generalized circuit with fan-out 2^2 . Our third and final step (Section 7) is to reduce the problem of finding an approximate solution to a circuit with fan-out 2 to that of finding an approximate Nash equilibrium. This reduction follows directly from the game gadgets of Daskalakis et al. [8] which establishes the result for ϵ -NE. We complete the proof by pointing out that for bounded degree graphical games, finding an ϵ -NE reduces to finding a $\Theta(\epsilon^2)$ -ANE.

2.1 Comparison of techniques to Chen, Deng, and Teng

The second step in our reduction, going from approximate Brouwer fixed-point to approximate generalized circuit, heavily relies on the techniques used by Chen, Deng, and Teng [4]. The key technical difference in our approach compared to [4] is the particular choice of hard instances of Brouwer functions.

In order to construct the hard Brouwer function we use a technique due to Hirsch, Papadimitriou, and Vavasis [14] that embeds paths as mappings over $[0, 1]^n$, in a delicate way which we describe later.

Chen et al. [4], on the other hand, divide the unit hypercube into subcubes of constant side length, and specify a color for each subcube. The i -th color corresponds to ξ_i , i.e. the i -th unit vector, whereas a special red color corresponds to $-\sum \xi_i$. Any vector of players' mixed strategies corresponds to a distribution over neighboring subcubes, and we get a fixed point whenever the expectation of the corresponding vectors is 0^n .

Constant hardness of approximation In Chen et al.'s construction, any distribution corresponding to an exact fixed point must have support over a panchromatic neighborhood, i.e. a neighborhood with subcubes corresponding to all $n + 1$ colors. Furthermore, it is shown that it is PPAD-hard to find such panchromatic neighborhood. In fact, a panchromatic neighborhood is still nec-

²This step was missing in previous versions of this manuscript.

essary for any $\Theta(1/n)$ -approximate³ fixed point. However, a $(1/n)$ -approximate fixed point can be achieved from a distribution over only n colors: for example, taking each of the colors $i \in [n]$ with probability $1/n$ results in the expected vector $(1/n, \dots, 1/n)$.

As we show later, the Hirsch et al. construction, does not suffer from this caveat. This advantage of the Hirsch et al. construction was used before by Babichenko [1] and also in [22].

A simpler averaging gadget Both constructions of hard instances of Brouwer functions (Hirsch et al., and the one used in Chen et al.), partition the hypercube into subcubes, and define the function separately on each subcube. When we construct a circuit that approximately simulates such a Brouwer function, we have a problem near the facets of the subcubes: the approximate circuit cannot determine to which subcube the input belongs.


Chen et al. overcome this problem using *equiangle sampling*: consider many translations of the input vector by adding small multiples of the all-ones vector; compute the displacement for each translation, and average. Since each translation may be close to a facet in a different dimension, Chen et al. consider a polynomial number of translations. Thus, all translations must be polynomially close to each other - otherwise they will be too far to approximate the true input.

We avoid this problem by observing another nice property of the Hirsch et al. construction: when the input vector lies near two or more facets, the displacement is (approximately) the same, regardless of the subcube. Once we rule out such points, it suffices to sample only a constant number of points (as at most one of them may be too close to a facet). See also illustration in Figure 2.1.

Bounding the fan-out Chen et al. do not explicitly discuss the fan-out of the generalized circuit gates (it is not needed for the reduction to two-player games), and Daskalakis et al. achieve fan-out 2 by replacing gates with large fan-out with binary trees of $G_{=}$ gates. When the original fan-out is large, this technique amplifies the error by a logarithmic factor which we cannot afford.

³Approximation here is in the sense of L^∞ . In fact, it seems to give a constant hardness of approximation in L^1 , but unfortunately it is not clear how to use that for a reduction to Nash equilibrium.

Figure 1: Comparison of averaging gadgets



Daskalakis et al. [8]

Chen et al. [4]

This paper

A comparison of the averaging gadgets of [8], [4], and this paper. **x** is the point whose displacement we would like to estimate using imprecise gates and brittle comparators. Points that are too close to a facet between subcubes are denoted by **triangles**, while points that are sufficiently far are denoted by circles. Finally, in this paper we have a “safe” zone (**shaded**) around the corner where we don’t need to parse the subcube; thus we only need to avoid one facet.

In Section 6 we overcome this issue by resorting to logical gates (in particular, G_{-}) that do not amplify error. Given a real output of an arithmetic gate, we convert it to unary representation⁴ using a constant number of (fan-out 2) gates. Then we copy the unary representation using a binary tree of G_{-} gates. Finally, we convert each copy back to a real number using a constant number of gates.

2.2 Alternative proofs

We informally sketch a couple of different approaches that could lead to our main theorem or similar results.

No averaging gadget

The simpler averaging gadget was perhaps the main breakthrough that enabled us to improve on the results in [22]. In hindsight it seems that we could completely avoid the use of averaging gadgets. As we discuss in the previous subsection, reducing from the Hirsch et al. Brouwer function allows us to treat *corners*,

⁴Unary representation of numbers with constant precision is prevalent throughout our implementation of the generalized circuit. We prefer unary representation over binary, because in the former at most one bit can be ambiguous due to the use of brittle comparators.

i.e. the intersection of two or more facets, differently. In particular, near corners we don't need to determine to which subcube our input vector belongs because the displacement is always the same. The simpler averaging gadget is used when we are close to one facet.

Alternatively, one could construct a gadget that determines that the input is close to a particular facet *-without deciding on which side of the facet-* and compute the Hirsch et al. displacement on that facet. Thus, no averaging gadget is needed, but some care is necessary when stitching together the corner gadget (which would be the same as this paper), the facet gadget, and the interior point gadget.

A richer set of gates

There is also a simpler and more intuitive reduction that gives somewhat weaker results, namely: PPAD-completeness for degree 3 graphical games with a constant number of actions per player.

The most arduous part in our proof is the second step, which takes us from the Hirsch et al. Brouwer function to a generalized circuit. In order to prove this reduction, we must design a circuit (or an algorithm) that implements the equiangle sampling and the Hirsch et al. Brouwer function using the limited set of gates allowed in Chen et al.'s definition of ϵ -GCIRCUIT. This part could be simplified using a more expressive set of gates.

A recent reduction that appeared in Eran Shmaya's blog [24] would essentially allow us to replace the ϵ -GCIRCUIT gates with any gates of bounded fan-in and fan-out that compute c -Lipschitz functions for any constant c . We briefly sketch here this reduction from an arbitrary generalized circuit to a graphical game:

Each line in the circuit corresponds to a player, and we connect all players that share a common gate. The strategies of each player correspond to an $O(\epsilon)$ -discretization of $[0, 1]$. When the player that corresponds to the output of gate G plays strategy b , and the input players play strategies a_1 and a_2 , the utility to the output player is given by

$$u(b, a_1, a_2) = -|b - G(a_1, a_2)|^2$$

It can be shown that in any $O(\epsilon^2)$ -NE of this game, only the strategies closest to $\mathbb{E}_{a_1, a_2} [G(a_1, a_2)]$ are played.

3 Preliminaries

Throughout this paper we use the max-norm as the default measure of distance. In particular, when we say that f is M -Lipschitz we mean that for every \mathbf{x} and \mathbf{y} in the domain of f , $\|f(\mathbf{x}) - f(\mathbf{y})\|_\infty \leq M \|\mathbf{x} - \mathbf{y}\|_\infty$.

A large part of our paper deals with approximate solutions to equations. We adopt the notation of writing $x = y \pm \epsilon$ to imply that $x \in (y - \epsilon, y + \epsilon)$.

The ENDOFTHELINE problem

Our reduction starts from the ENDOFTHELINE problem. This problem was implicit in [21], and explicitly defined by Daskalakis et al. [8].

Definition 1. ENDOFTHELINE: ([8]) Given two circuits S and P , with n input bits and n output bits each, such that $P(0^n) = 0^n \neq S(0^n)$, find an input $x \in \{0, 1\}^n$ such that $P(S(x)) \neq x$ or $S(P(x)) \neq x \neq 0^n$.

Theorem. (Essentially [21]) ENDOFTHELINE is PPAD-complete (for poly-size S and P).

3.1 ϵ -Nash equilibrium vs ϵ -approximate Nash equilibrium

A mixed strategy of player i is a distribution over i 's set of actions, denoted $x_i \in \Delta A_i$. We say that a vector of mixed strategies $\mathbf{x} \in \times_j \Delta A_j$ is a *Nash equilibrium* if every strategy a_i in the support of x_i is a best response to the actions of the mixed strategies of the rest of the players, x_{-i} . Formally,

$$\forall a_i \in \text{Supp}(x_i) \quad \mathbb{E}_{a_{-i} \sim x_{-i}} [u_i(a_i, a_{-i})] = \max_{a' \in A_i} \mathbb{E}_{a_{-i} \sim x_{-i}} [u_i(a', a_{-i})].$$

Equivalently, \mathbf{x} is a Nash equilibrium if each mixed strategy x_i is a best mixed response to x_{-i} :

$$\mathbb{E}_{\mathbf{a} \sim \mathbf{x}} [u_i(\mathbf{a})] = \max_{x'_i \in \Delta A_i} \mathbb{E}_{\mathbf{a} \sim (x'_i; x_{-i})} [u_i(\mathbf{a})].$$

Each of those equivalent definitions can be generalized to include approximation in a different way. (Of course, there are also other interesting generalizations of Nash equilibria to approximate settings.) We say that \mathbf{x} is an ϵ -approximate

Nash equilibrium (ϵ -ANE) if each x_i is an ϵ -best mixed response to x_{-i} :

$$\mathbb{E}_{\mathbf{a} \sim \mathbf{x}} [u_i(\mathbf{a})] \geq \max_{x'_i \in \Delta A_i} \mathbb{E}_{\mathbf{a} \sim (x'_i; x_{-i})} [u_i(\mathbf{a})] - \epsilon.$$

On the other hand, we generalize the first definition of Nash equilibrium by saying that \mathbf{x} is a ϵ -Nash equilibrium (ϵ -NE; sometimes also ϵ -well-supported Nash equilibrium) if each a_i in the support of x_i is an ϵ -best response to x_{-i} :

$$\forall a_i \in \text{Supp}(x_i) \quad \mathbb{E}_{a_{-i} \sim x_{-i}} [u_i(a_i, a_{-i})] \geq \max_{a' \in A_i} \mathbb{E}_{a_{-i} \sim x_{-i}} [u_i(a', a_{-i})] - \epsilon.$$

It is easy to see that every ϵ -NE is also an ϵ -ANE, but the converse is false. In Lemma 4 we prove that given an ϵ -ANE in a graphical game with incoming degree d_{in} , it is possible to find a $\Theta(\sqrt{\epsilon}d_{in})$ -NE. We note that this approach is unlikely to give a lower bound on the *query complexity* for ϵ -ANE: graphical games of bounded incoming degree can be learned in a polynomial number of queries; furthermore, due to [11], for any concisely represented game, it is possible to find an ϵ -NE (and hence also ϵ -ANE) in a polynomial number of queries.

3.2 The ϵ -GCIRCUIT problem

Generalized circuits are similar to the standard algebraic circuits, the main difference being that generalized circuits contain cycles, which allow them to verify fixed points of continuous functions. We restrict the class of generalized circuits to include only a particular list of gates described below. Formally,

Definition 2. [Generalized circuits, [4]] A *generalized circuit* \mathcal{S} is a pair (V, \mathcal{T}) , where V is a set of nodes and \mathcal{T} is a collection of gates. Every gate $T \in \mathcal{T}$ is a 5-tuple $T = G(\zeta \mid v_1, v_2 \mid v)$, in which $G \in \{G_\zeta, G_{\times\zeta}, G_-, G_+, G_<, G_\vee, G_\wedge, G_-\}$ is the type of the gate; $\zeta \in \mathbb{R} \cup \{nil\}$ is a real parameter; $v_1, v_2 \in V \cup \{nil\}$ are the first and second input nodes of the gate; and $v \in V$ is the output node.

The collection \mathcal{T} of gates must satisfy the following important property: For every two gates $T = G(\zeta \mid v_1, v_2 \mid v)$ and $T' = G'(\zeta' \mid v'_1, v'_2 \mid v')$ in \mathcal{T} , $v \neq v'$.

Alternatively, we can think of each gate as a constraint on the values on the incoming and outgoing wires. We are interested in the following constraint satisfaction problem: given a generalized circuit, find an assignment to all the

wires that simultaneously satisfies all the gates. When every gate computes a continuous function of the incoming wires, a solution must exist by Brouwer's fixed point theorem.

In particular, we are interested in the approximate version of this CSP, where we must approximately satisfy every constraint.

Definition 3. Given a generalized circuit $\mathcal{S} = (V, \mathcal{T})$, we say that an assignment $\mathbf{x}: V \rightarrow [0, 1]$ ϵ -approximately satisfies \mathcal{S} if for each of the following gates, \mathbf{x} satisfies the corresponding constraints:

Gate	Constraint
$G_\zeta (\alpha \mid a)$	$\mathbf{x}[a] = \alpha \pm \epsilon$
$G_{\times\zeta} (\alpha \mid a \mid b)$	$\mathbf{x}[b] = \alpha \cdot \mathbf{x}[a] \pm \epsilon$
$G_= (\mid a \mid b)$	$\mathbf{x}[b] = \mathbf{x}[a] \pm \epsilon$
$G_+ (\mid a, b \mid c)$	$\mathbf{x}[c] = \min(\mathbf{x}[a] + \mathbf{x}[b], 1) \pm \epsilon$
$G_- (\mid a, b \mid c)$	$\mathbf{x}[c] = \max(\mathbf{x}[a] - \mathbf{x}[b], 0) \pm \epsilon$
$G_{<} (\mid a, b \mid c)$	$\mathbf{x}[c] = \begin{cases} 1 \pm \epsilon & \mathbf{x}[a] < \mathbf{x}[b] - \epsilon \\ 0 \pm \epsilon & \mathbf{x}[a] > \mathbf{x}[b] + \epsilon \end{cases}$
$G_{\vee} (\mid a, b \mid c)$	$\mathbf{x}[c] = \begin{cases} 1 \pm \epsilon & \mathbf{x}[a] = 1 \pm \epsilon \text{ or } \mathbf{x}[b] = 1 \pm \epsilon \\ 0 \pm \epsilon & \mathbf{x}[a] = 0 \pm \epsilon \text{ and } \mathbf{x}[b] = 0 \pm \epsilon \end{cases}$
$G_{\wedge} (\mid a, b \mid c)$	$\mathbf{x}[c] = \begin{cases} 1 \pm \epsilon & \mathbf{x}[a] = 1 \pm \epsilon \text{ and } \mathbf{x}[b] = 1 \pm \epsilon \\ 0 \pm \epsilon & \mathbf{x}[a] = 0 \pm \epsilon \text{ or } \mathbf{x}[b] = 0 \pm \epsilon \end{cases}$
$G_{\neg} (\mid a \mid b)$	$\mathbf{x}[b] = \begin{cases} 1 \pm \epsilon & \mathbf{x}[a] = 0 \pm \epsilon \\ 0 \pm \epsilon & \mathbf{x}[a] = 1 \pm \epsilon \end{cases}$

Given a generalized circuit $\mathcal{S} = (V, \mathcal{T})$, ϵ -GCIRCUIT is the problem of finding an assignment that ϵ -approximately satisfies it.

Brittle comparators Intuitively, in order for (approximate) solutions to the circuit problem to correspond to (approximate) equilibria, all our gates should implement continuous (Lipschitz) functions. The gate $G_{<} (\mid a, b \mid c)$, for example, approximates that the function $c(a, b) = \begin{cases} 1 & a < b \\ 0 & a \geq b \end{cases}$, which is not continuous. To overcome this problem, Daskalakis et al. [8] defined the *brittle comparator*: when a is (ϵ -) larger than b , it outputs 0; when b is (ϵ -) larger

than a , it outputs 1. However, when a and b are (ϵ -approximately) equal, its behaviour is undefined.

Brittleness introduces difficulties in the transition from continuous to discrete solutions. This challenge is overcome by an averaging gadget, which is described in detail in Section 5.

3.3 The Hirsch et al. construction

In order to construct a hard instance of Brouwer function, we use techniques introduced by Hirsch, Papadimitriou, and Vavasis [14]. Hirsch et al. showed that the number of deterministic value-oracle queries required to find a Brouwer fixed point is exponential in the number of precision digits, the Lipschitz constant of the function, and -most important to us- the dimension of the domain. We reproduce here an overview of their construction, focusing on the details that will be used in our proof.

The continuous Brouwer function is denoted by f , while the associated displacement function is denoted by $g(\mathbf{x}) = f(\mathbf{x}) - \mathbf{x}$. The domain of f is the n -dimensional (solid) hypercube. The hypercube is divided into subcubes, of side length h .

The first dimension is special. All the action takes place in a designated $(n - 1)$ -dimensional layer of subcubes that has the same values in the first coordinate; this layer is called the *slice*. Within the slice, we also ignore the subcubes that are near the boundary of the hypercube (those compose the *frame*); we are left with the subcubes in the center of the slice, which we call the *picture*.

In [14], a single path is embedded inside the picture; this embedded path is called the *tube*. The *home subcube*, the subcube that corresponds to the beginning of the path, is special: all the flow from all subcubes that do not belong to the tube leads to this subcube. We mention that constructing this flow outside the picture is relatively simple, and mostly ignore it for the rest of the paper (with the exception of a brief overview in Section 3.3.5).

3.3.1 Our modification

We briefly describe our modifications here, and present them in more detail in Section 4.2. Instead of a single path, we embed many paths and also cycles in the slice; instead of one tube, we now have many tubes. The home subcube

continues to be special in the sense described above, and it is the starting point of one of our paths.

We also make a couple of changes in notation: It will be convenient to take the hypercube to be of dimension $2n + 2$; and we let the special dimension be the last one, i.e. the $(2n + 2)$ -th.

Finally we choose the subcube side length, h , as long as possible (we will soon fix $h = 1/6$, following [1]). In particular, our picture contains 2^{2n+1} subcubes.

3.3.2 Key properties

For the purpose of adapting this construction, the most important property is that most of the slice has the same default displacement: directly upwards, i.e. $g(\mathbf{x}) = \delta \boldsymbol{\xi}_{2n+2}$, where $\boldsymbol{\xi}_{2n+2}$ is the $(2n + 2)$ -unit vector, and δ is some small parameter (constant in this paper). Formally,

Fact 1. $g(\mathbf{x}) = \delta \boldsymbol{\xi}_{2n+2}$, for every \mathbf{x} such that at least one of the following holds:

1. \mathbf{x} lies on a corner, i.e. the intersection of two or more facets of a subcube;
2. \mathbf{x} lies on an outer facet of a tube subcube, i.e. a facet other than the two facets that continue the path; or
3. \mathbf{x} lies in a subcube that does not belong to any tube.

Intuitively, Property 2 implies that *all subcubes -whether they belong to the tube or not- look the same from the outside* (except for the two facets that continue the path). This essentially means that we are free to embed many paths and cycles in the slice.

Property 1, stating that *all corners look the same*, is key to our modified construction, because it liberates us from having to disambiguate a position of a point near the corners (that is, from deciding exactly to which subcube it belongs). This is useful because those points that are close to corners, are precisely the ones that are hard to determine using “noisy” gates.

Finally, the following theorem summarizes the important parameters of f .

Theorem 3. (Parameters of f [14]⁵)

1. g is 79-Lipschitz (thus, f is 80-Lipschitz)

⁵See also [1] for the choice of constants. Please note the change in variable names: $h, M, 2^{-p}$ in [14] correspond to $\delta, \lambda, \epsilon$, respectively, in [1].

2. $\|g(\mathbf{x})\|_\infty \geq 1/88$ for every \mathbf{x} that does not correspond to the endpoint of the path
3. The value of g at each point \mathbf{x} depends only on whether the path passes through the subcube corresponding to \mathbf{x} , and in which direction.

3.3.3 Max-norm Geometry

As we mentioned earlier, throughout this paper we work with the max-norm. This has some implications that may contradict our geometric intuition. For example: *in a max-norm world, a circle is a square.*

Max-norm interpolation Given coordinates $x, y \geq 0$, we define the *max-norm angle*⁶ that point (x, y) forms with the X -axis (in the XY -plane) as

$$\theta_{\max}(x, y) = \frac{y}{x + y}$$

The max-norm angle is useful for interpolation. Given the values of f on two neighboring facets of a subcube, we can extend f to all points of the subcube by *angular interpolation*: interpolate according to the max-norm angle $\theta_{\max}(x_i, x_j)$ where x_i and x_j are the respective distances from the two facets. When f is defined on two opposite facets, we can simply use *Cartesian interpolation*, which again means to interpolate according to the distance from each facet.

Max-norm local polar coordinates Given a point $\mathbf{z} \in \mathbb{R}^n$ we define a new *local max-norm polar coordinate* system around \mathbf{z} . Every $\mathbf{x} \in \mathbb{R}^n$ is transformed into $\langle r, \mathbf{p} \rangle_{\mathbf{z}} \in \mathbb{R} \times \mathbb{R}^n$ where $r = \|\mathbf{x} - \mathbf{z}\|$ is the max-norm radius, and $\mathbf{p} = (\mathbf{x} - \mathbf{z})/r$ is the max-norm unit vector that points from \mathbf{z} in the direction of \mathbf{x} .

3.3.4 Definition of g

Finally, we are ready to explicitly define the displacement function, g . The mapping is defined so that in the center of the tube, the flow goes along the direction of the path; slightly outside the center, the flow points towards the center of the tube; further away from the center, the flow goes *against* the direction of the path; at the outer boundary of the tube, as we previously described, the flow goes upwards.

⁶Our max-norm angle was called *unit* in [14].

We first define g on facets. Let $\langle r, \mathbf{p} \rangle_{\mathbf{z}}$ be a point on the facet centered at \mathbf{z} , and suppose that the tube enters the subcube through \mathbf{z} , advancing in the positive i -th coordinate. We define

$$g(\langle r, \mathbf{p} \rangle_{\mathbf{z}}) = \begin{cases} \delta \boldsymbol{\xi}_i & r = 0 \\ -\delta \mathbf{p} & r = h/8 \\ -\delta \boldsymbol{\xi}_i & r = h/4 \\ \delta \boldsymbol{\xi}_{2n+2} & r = h/2 \end{cases} \quad (1)$$

(Recall that h is the subcube side length, and δ is some small constant.) For $r \in (0, h/8)$, interpolate between $\delta \boldsymbol{\xi}_i$ and $-\delta \mathbf{p}$ ([14] call this *radial interpolation*), and similarly for $r \in (h/8, h/4)$ and $r \in (h/4, h/2)$. See also illustration in Figure 2.

The home subcube is defined using (1) as if the tube enters from above, i.e. coming down the $(2n + 2)$ -dimension, and exits through another facet (in one of the first $(2n + 1)$ dimensions) in the direction of the path. For all other starting and end points, we define $g(\mathbf{x}) = \delta \boldsymbol{\xi}_{2n+2}$ on the facet opposite the one that continues the tube.

In the last couple of paragraphs we defined g on two facets for each subcube that belongs to the tubes; for all other points in the tubes we interpolate (either Cartesian or angular interpolation) between those two facets. Outside the tubes, $g(\mathbf{x}) = \delta \boldsymbol{\xi}_{2n+2}$.


3.3.5 Outside the picture

For completeness, we briefly describe Hirsch et al.'s construction outside the picture. As we mentioned earlier, this part is much simpler than the picture itself, and we will ignore it after this section.

For all points in the frame and below the slice, the displacement points directly upward, i.e. $g(\mathbf{x}) = \delta \boldsymbol{\xi}_{2n+2}$. Moving above the slice, let $\mathbf{z}[\text{top}]$ be the point on the top facet of the hypercube which is directly above the center of the home subcube. For all points $\langle r, \mathbf{p} \rangle_{\mathbf{z}[\text{top}]}$ on the top facet of the hypercube, define the displacement as follows:


$$g(\langle r, \mathbf{p} \rangle_{\mathbf{z}[\text{top}]}) = \begin{cases} -\delta \boldsymbol{\xi}_{2n+2} & r = 0 \\ -\delta \mathbf{p} & r \geq h/8 \end{cases}$$

Figure 2: A facet of the Hirsch et al. construction



An illustration of the displacement on a facet between two subcubes in a tube; the direction of the path is into the paper. In the **center**, the displacement points into the paper; **slightly further outside**, the displacement points towards the center; **further outside**, the displacement points out of the paper; finally in the **outer layer**, the displacements points in the special $2n + 2$ dimension.

Figure 3: Outside the picture



An illustration of the displacement outside the picture.

and interpolate for $r \in (0, h/8)$.

Notice that the definition of g on the slice from the previous subsection, implies that all the points in the top facet of the slice, except for the top of the home subcube, point directly upwards. Let $\mathbf{z}[\text{home}]$ denote the center of the top facet of the home subcube. We therefore have that for any $\langle r, \mathbf{p} \rangle_{\mathbf{z}[\text{home}]}$ in the top facet of the slice,

$$g\left(\langle r, \mathbf{p} \rangle_{\mathbf{z}[\text{home}]}\right) = \begin{cases} -\delta \boldsymbol{\xi}_{2n+2} & r = 0 \\ -\delta \mathbf{p} & r = h/8 \\ \delta \boldsymbol{\xi}_{2n+2} & r \geq h/4 \end{cases}$$

where we again interpolate radially for r in $(0, h/8)$ and $(h/8, h/4)$.

Finally, to complete the construction above the slice, simply interpolate (using Cartesian interpolation) between the top facets of the slice and the hypercube. See also illustration in Figure 3.

4 Finding an approximate Brouwer fixed point is PPAD-hard

In the first step of the proof, we show that finding a fixed point of the modified Hirsch et al. instance is PPAD-hard. This reduction recently appeared in [22].

Definition 4. Given an arithmetic circuit that computes an M -Lipschitz function $f: [0, 1]^n \rightarrow [0, 1]^n$, $\text{SUCCINCTBROUWER}(n, M, \epsilon)$ is the problem of computing an ϵ -approximate fixed point of f (an \mathbf{x} such that $\|f(\mathbf{x}) - \mathbf{x}\| \leq \epsilon$).

Theorem 4. ([22]) $\text{SUCCINCTBROUWER}(n, 80, 1/88)$ is PPAD-hard.

Proof. In the first step, we embed the ENDOFTHELINE as a collection H of vertex-disjoint paths on the $(2n + 1)$ -dimensional hypercube graph. Given H , our second step is to construct a continuous mapping $f: [0, 1]^{2n+2} \rightarrow [0, 1]^{2n+2}$ whose fixed points correspond to ends of paths in H . This step is done using a technique introduced by Hirsch et al. [14].

4.1 Embedding the ENDOFTHELINE graph as paths in $\{0, 1\}^{2n+1}$

Our first step in the reduction is to embed an ENDOFTHELINE graph as vertex-disjoint paths on the $(2n + 1)$ -dimensional hypercube graph. We first recall that the input to the ENDOFTHELINE problem is given as two circuits S and P , which define a directed graph over G over $\{0, 1\}^n$. Given S and P , we construct a collection H of vertex-disjoint paths over the $(2n + 1)$ -dimensional hypercube graph, such that each starting or end point of a path in H corresponds to a unique starting or end point of a line in G .

In order to construct our embedding we divide the $2n + 1$ coordinates as follows: the first n coordinates store the current vertex \mathbf{u} , the next n coordinates for the next vertex in the line, \mathbf{v} , and finally, the last coordinate b stores a compute-next vs copy bit. When $b = 0$, the path proceeds to update $\mathbf{v} \leftarrow S(\mathbf{u})$, bit-by-bit. When this update is complete, the value of b is changed to 1. Whenever $b = 1$, the path proceeds by copying $\mathbf{u} \leftarrow \mathbf{v}$ bit-by-bit, and then changes that value of b again. Finally, when $\mathbf{u} = \mathbf{v} = S(\mathbf{u})$ and $b = 0$, the path reaches an end point.

Notice that the paths in H do not intersect. Furthermore, given a vector in $\mathbf{p} \in \{0, 1\}^{2n+1}$, we can output in polynomial time whether \mathbf{p} belongs to a path in H , and if so which are the previous and consecutive vectors in the path. It

is therefore PPA-hard to find a starting or end point of any path in H other than 0^n .

4.2 Continuous mapping on $[0, 1]^{2n+2}$

Given a collection H of non-intersecting paths in $\{0, 1\}^{2n+1}$, we construct a mapping $f: [0, 1]^{2n+2} \rightarrow [0, 1]^{2n+2}$ using a modification of the Hirsch et al. mapping described in Section 3.3. Essentially, we construct many tubes (some of which may form closed cycles) in the same slice.

If any path in H passes through some $\mathbf{p} \in \{0, 1\}^{2n+1}$, then \mathbf{p} corresponds to a subcube on which f is defined exactly the same as in [14]. Likewise, every end point of a path in H corresponds to a subcube on which f is defined exactly the same as the unique end point in [14].

Our construction differs in the starting points of the paths in H . In [14], the starting point corresponds to the home subcube, which is *globally unique*: the flow from any subcube outside the path leads towards that point. Indeed we cannot imitate this behaviour for the starting point of every path in H . However, this does not complicate our reduction, because all other starting points are also legitimate solutions of ENDOFTHELINE. In particular we can construct f in a similar fashion to the end points of the paths, thereby creating additional fixed points.

For each starting point, consider the corresponding subcube of $[0, 1]^{2n+2}$. The values of f on the facet F_1 in the direction of the path, are already determined by the next subcube in the path. Now, we let $g(\mathbf{x}) = \delta\xi_{2n+2}$ uniformly for every point \mathbf{x} on the opposite facet, F_0 . For any \mathbf{x}' between the facets we interpolate by taking the weighted average of the values $f(\mathbf{x}'_0)$ and $f(\mathbf{x}'_1)$ on the projections of \mathbf{x}' on F_0 and F_1 , respectively. (This corresponds to [14]'s Cartesian interpolation). Notice that this subcube also satisfies $g(\mathbf{x}) = \delta\xi_{2n+2}$ for all points on facets that do not face the rest of the path.

Outside the subcubes corresponding to the additional paths, all the properties of the mapping are trivially preserved. On the interface between any subcube in an additional path, and any other subcube which is not consecutive in the same path, all the properties are again preserved since adding the paths does not change the value of f near those facets. Finally within each path, f is constructed exactly the same way as the single path in [14], and therefore it is easy to see that all the properties continue to hold - except of course the new

fixed points at the starting points of the additional paths.

This completes our reduction. It follows from Theorem 3 and the simple modifications described above, that every ϵ -approximate fixed point of f corresponds to a solution to ENDOFTHELINE.

□

5 From Hirsch et al. to ϵ -GCIRCUIT

Proposition 1. *There exists a constant $\epsilon > 0$ such that ϵ -GCIRCUIT is PPAD-complete.*

Proof. We continue to denote the Hirsch et al. mapping by $f: [0, 1]^{2n+2} \rightarrow [0, 1]^{2n+2}$, and its displacement, $f(\mathbf{y}) - \mathbf{y}$ by $g(\mathbf{y})$. We design a generalized circuit \mathcal{S} that computes f , and verifies that the output is equal to the input. We show that every ϵ -approximate solution to \mathcal{S} corresponds to an $O(\epsilon^{1/4})$ -approximate fixed point of f .

Recall that the Hirsch et al. construction divides the hypercube into equal-sized subcubes. In particular, in our construction we divide the hypercube into subcubes of length $1/6$. Furthermore, all the paths in H are embedded in the 2^{2n+1} subcubes that belong to the picture. For ease of exposition, we present a construction that only works for points in the picture, i.e. $\mathbf{y} \in [2/6, 4/6]^{2n+1} \times [2/6, 1/2]$. The circuit can be extended in a similar fashion to deal with all $\mathbf{y} \in [0, 1]^{2n+2}$.

The most challenging part of the construction is the extraction of the information about the local subcube: is it part of the tube? if so, which are the entrance and exit facets? This is done by extracting the binary representation of the current subcube, and feeding it to the circuit that computes H . Notice that whenever we have valid logic bits, i.e. $\mathbf{x}[b] < \epsilon$ or $\mathbf{x}[b] > 1 - \epsilon$, we can perform logic operations on them without increasing the error.

Once we know the behaviour of the path on the current subcube, we simply have to implement the local Hirsch et al. mapping, for which we have a closed form description, using the available gates in the definition of generalized circuits. Since this definition does not include multiplication and division, we implement them in Algorithms 1 and 2, respectively. Notice that for Algorithm 2, in any ϵ -approximate solution, $\mathbf{x}[c] = \mathbf{x}[a] / \mathbf{x}[b] \pm 3\sqrt{\epsilon} / \mathbf{x}[b]$; when $\mathbf{x}[b]$ and ϵ are bounded away from 0, this is only a constant factor increase in the error.

Algorithm 1 MULTIPLY($| a, b | c$)

1. $G_\zeta(0 \mid h_0)$

2. **for** each $k \in [1/\sqrt{\epsilon}]$:

(a) $G_\zeta(k\sqrt{\epsilon} \mid \zeta_k), G_<(| a, \zeta_k | \bar{a}_k)$

The vector (\bar{a}_k) is the unary representation of a

$$\sum_{k: X[\bar{a}_k] < \epsilon} \sqrt{\epsilon} = \max_{k: X[\bar{a}_k] < \epsilon} k\sqrt{\epsilon} = \mathbf{x}[a] \pm \sqrt{\epsilon}$$

(b) $G_{\times\zeta}(\sqrt{\epsilon} \mid b \mid d_k)$

The vector (d_k) is simply equal to $b \cdot \sqrt{\epsilon}$ everywhere.

$$\sum_{k: X[\bar{a}_k] < \epsilon} \mathbf{x}[d_k] = \left(\sum_{k: X[\bar{a}_k] < 1-\epsilon} \sqrt{\epsilon} \right) \cdot \mathbf{x}[b] \pm \sqrt{\epsilon} = \mathbf{x}[a] \cdot \mathbf{x}[b] \pm 2\sqrt{\epsilon}$$

(c) $G_-(| d_k, \bar{a}_k | e_k)$

The vector (e_k) is $b \cdot \sqrt{\epsilon}$ only when $(\bar{a}_k) < \epsilon$.

$$\sum_{k: X[\bar{a}_k] < \epsilon} \mathbf{x}[e_k] = \mathbf{x}[a] \cdot \mathbf{x}[b] \pm 3\sqrt{\epsilon}$$

(d) $G_+(| h_{k-1}, e_k | h_k)$

Finally, we sum the e_k 's to get $a \cdot b$

$$\mathbf{x}[h_{1/\sqrt{\epsilon}}] = \mathbf{x}[a] \cdot \mathbf{x}[b] \pm 4\sqrt{\epsilon}$$

3. $G_=(| h_{1/\sqrt{\epsilon}} | c)$

Algorithm 2 DIVIDE($| a, b | c$)

1. $G_\zeta(0 || h_0)$

2. **for** each $k \in [1/\sqrt{\epsilon}]$:

(a) $G_{\times\zeta}(k\sqrt{\epsilon} | b | b_k), G_<(| b_k, a | d_k)$

The vector (d_k) is the unary representation of a/b

$$\left(\sum_{k: X[d_k] > \epsilon} \sqrt{\epsilon} \right) \cdot \mathbf{x}[b] = \left(\max_{k: X[d_k] > \epsilon} k\sqrt{\epsilon} \right) \cdot \mathbf{x}[b] = \mathbf{x}[a] \pm \sqrt{\epsilon}$$

(b) $G_{\times\zeta}(\sqrt{\epsilon} | d_k | e_k)$

The vector (e_k) is a $(\sqrt{\epsilon})$ -scaled version of (d_k)

$$\left(\sum \mathbf{x}[e_k] \right) \cdot \mathbf{x}[b] = \mathbf{x}[a] \pm 2\sqrt{\epsilon}$$

(c) $G_+(| h_{k-1}, e_k | h_k)$

Finally, we sum the e_k 's

$$\mathbf{x}[h_{1/\sqrt{\epsilon}}] \cdot \mathbf{x}[b] = \mathbf{x}[a] \pm 3\sqrt{\epsilon}$$

3. $G_=(| h_{1/\sqrt{\epsilon}} | c)$

Our construction has four parts: (1) equiangle sampling segment, (2) computing the displacement, (3) summing the displacement vectors, and (4) closing the loop. The first part contains a new idea introduced in this paper: using a constant size sample. The second part is a more technical description of the implementation of the Hirsch et al. mapping by approximate gates. The third and fourth parts are essentially exactly the same as in [4]. \square

5.1 Equiangle sampling segment

The first information we require in order to compute the Hirsch et al. mapping $f(\mathbf{y})$ is about the subcube to which \mathbf{y} belongs: is it part of the tube? if so, which are the entrance and exit facets? In order to answer those questions, we extract the binary representation of the cube. Recall that our circuit uses brittle comparators; thus when \mathbf{y} is close to a facet between subcubes, the behaviour of the brittle comparators may be unpredictable. We start with the easy case, where \mathbf{y} is actually far from every facet:

Definition 5. We say that \mathbf{y} is an *interior point* if for every i , $|y_i - 1/2| > \epsilon$; otherwise, we say that \mathbf{y} is a *boundary point*.

A very nice property of the Hirsch et al. construction is that whenever \mathbf{y} is at the intersection of two or more facets, the displacement is the same: $g(\mathbf{y}) = \delta \boldsymbol{\xi}_{2n+2}$. Thus, by the Lipschitz property of g , whenever \mathbf{y} is close to the intersection of two or more facets, the displacement is approximately $\delta \boldsymbol{\xi}_{2n+2}$. For such \mathbf{y} 's, we don't care to which subcube they belong.

Definition 6. We say that \mathbf{y} is a *corner point* if there exist distinct $i, j \in [2n + 2]$ such that $|y_i - 1/2| < \epsilon^{1/4}$ and $|y_j - 1/2| < \epsilon^{1/4}$.

(Notice that \mathbf{y} may be an interior point and a corner point at the same time.)

We still have a hard time handling \mathbf{y} 's which are neither an interior point nor a corner point. To mitigate the effect of such \mathbf{y} 's we use an equiangle averaging scheme. Namely we consider the set:

$$E^\epsilon(\mathbf{y}) = \{\mathbf{y}^l = \mathbf{y} + (6l \cdot \epsilon) \mathbf{1} : 0 \leq l < 1/\sqrt{\epsilon}\}$$

where $\mathbf{1}$ denotes the all-ones vector. Notice that since g is λ -Lipschitz for constant λ , $g(\mathbf{y}^l)$ will be approximately the same for all $\mathbf{y}^l \in E^\epsilon(\mathbf{y})$.

Fact 2. *If any $\mathbf{y}^l \in E^\epsilon(\mathbf{y})$ is not a corner point, then at most one $\mathbf{y}^l \in E^\epsilon(\mathbf{y})$ is a boundary point.*

Proof. For each dimension, at most one element in $E^\epsilon(\mathbf{y})$ can be ϵ -close to the $(1/2)$ facet. Thus if two elements in $E^\epsilon(\mathbf{y})$ are boundary points, it must be because of distinct dimensions - and therefore every \mathbf{y}^l is a corner point. \square

Given input \mathbf{y} , we compute the displacement $g(\cdot)$ separately and in parallel for each $\mathbf{y}^l \in E^\epsilon$, and average at the end. Since at most one \mathbf{y}^l is a boundary point, this will incur an error of at most $\sqrt{\epsilon}$.

In the generalized circuit we can construct E^ϵ using $(1/\sqrt{\epsilon})$ auxiliary nodes and G_ζ and G_+ gates:

$$\mathbf{x}[y_i^l] = \min \{ \mathbf{x}[y_i^0] + (6l \cdot \epsilon), 1 \} \pm 2\epsilon$$

5.2 Computing the displacement

For each $\mathbf{y}^l \in E^\epsilon$, we construct a disjoint circuit that approximates the displacement $g(\mathbf{y}^l)$. In the description of the circuit below we omit the index l .

Lemma 1. *The circuit below $O(\sqrt{\epsilon})$ -approximately simulates the computation of the Hirsch et al. displacement:*

1. *Whenever $(\mathbf{x}[y_i])_{i \in [2n+2]}$ is an interior point,*

$$\mathbf{x}[g_i^+] - \mathbf{x}[g_i^-] = g_i(\mathbf{x}[\mathbf{y}]) \pm O(\epsilon^{1/4})$$

2. *Furthermore, whenever $(\mathbf{x}[y_i])_{i \in [2n+2]}$ is a corner point,*

$$\mathbf{x}[g_{2n+2}^+] - \mathbf{x}[g_{2n+2}^-] = \delta \pm O(\sqrt{\epsilon})$$

and $\forall i < 2n + 2$:

$$\mathbf{x}[g_i^+] - \mathbf{x}[g_i^-] = 0 \pm O(\sqrt{\epsilon})$$

Proof. We construct the circuit in five stages: (1) given \mathbf{y} , we extract \mathbf{b} , that is the binary representation of the corresponding subcube in $\{0, 1\}^{2n+2}$; (2) we then compute whether \mathbf{b} belongs to a path in H , and if so which are the previous and next vertices; (3) we compute the centers of the coordinate systems

corresponding to the entrance and exit facets, and label them \mathbf{z}^{in} and \mathbf{z}^{out} ; (4) we project \mathbf{y} to each facet, and transform this projection to the local polar coordinate systems - $(r^{in}, \mathbf{p}^{in})$; and (5) finally, we use all the information above to compute the displacement $\mathbf{g} = g(\mathbf{y})$.

The correctness of Lemma 1 follows from Claims 1-7.

Extract $\mathbf{b} \in \{0, 1\}^{2n+2}$

Our first step is to extract the binary vector \mathbf{b} which represents the subcube to which \mathbf{y} belongs. In other words we want b_i to be the indicator of $y_i < 1/2$. We do that by adding the following gadgets: $G_\zeta(1/2 \parallel c_{1/2})$ and, for each i , $G_{<}(y_i, c_{1/2} \parallel b_i)$. Observe that now

$$\mathbf{x}[b_i] = \begin{cases} 0 \pm \epsilon & \mathbf{x}[y_i] < \mathbf{x}[c_{1/2}] - \epsilon \\ 1 \pm \epsilon & \mathbf{x}[y_i] > \mathbf{x}[c_{1/2}] + \epsilon \end{cases}$$

Claim 1. If $\mathbf{x}[\mathbf{y}]$ is an interior point, $\mathbf{x}[\mathbf{b}]$ is the correct representation (up to ϵ error) of the corresponding bits in $\{0, 1\}^{2n+2}$.

Neighbors in H

Given $\mathbf{x}[\mathbf{b}]$ we can construct, using G_\wedge 's and G_\neg 's and a polynomial number of unused nodes, the circuits S^H and P^H that give the next and previous vertex visited by our collection of paths, H . The output of each circuit is represented by $2n + 2$ unused nodes $\{P_i^H(\mathbf{b})\}$ and $\{S_i^H(\mathbf{b})\}$.

Recall that H is defined in $\{0, 1\}^{2n+1}$, so the last input bit is simply ignored (inside the picture it is always 0); the last output bit is used as follows. Our convention is that starting points and end points correspond to $P^H(\mathbf{b}) = \mathbf{b}$ and $S^H(\mathbf{b}) = \mathbf{b}$, respectively, and likewise for points that do not belong to any path. An exception to this is the $\mathbf{0}$ starting point, which will correspond to $P^H(\mathbf{0}) = (0^{2n+1}; 1)$: This is in accordance with the Hirsch et al. construction, where the home subcube is constructed as if it continues a path from the subcube above it.

Claim 2. If $\mathbf{x}[\mathbf{b}]$ is an ϵ -approximate binary vector, i.e. $\mathbf{x}[\mathbf{b}] \in ([0, \epsilon] \cup [1 - \epsilon, 1])^{2n+2}$, then $\mathbf{x}[P^H(\mathbf{b})]$ and $\mathbf{x}[S^H(\mathbf{b})]$ correctly represent (up to ϵ error) the previous vertex and next vertex in H .

Entrance and exit facets

Let $b_i^{+in} = b_i \wedge \neg P_i^H(\mathbf{b})$, i.e. b_i^{+in} is 1 if the path enters the current subcube via the positive i -th direction; define b_i^{-in} analogously. Let b_i^{in} denote the OR of b_i^{+in} and b_i^{-in} .

The center of the entrance facet is constructed via G_ζ , $G_{\times\zeta}$, G_+ , and G_- according to the formula:

$$z_i^{in} = \begin{cases} 5/12 + (1/6) \cdot b_i & b_i^{in} = 0 \\ 1/2 & b_i^{in} = 1 \end{cases}$$

Construct \mathbf{z}^{out} analogously.

Notice that if we know on which coordinate the path enters, in $\{0, 1\}^{2n+2}$ it has only one possible direction; in the Hirsch et al. hypercube this corresponds to always entering from the center (i.e. from the $y_i = 1/2$ facet). Also, if \mathbf{b} corresponds to a non-trivial starting point $\mathbf{b}^{in} = \mathbf{0}$ and \mathbf{z}^{in} is simply the center of the subcube (and similarly for \mathbf{b}^{out} , \mathbf{z}^{out} when \mathbf{b} is an end point).

Claim 3. If $\mathbf{x}[\mathbf{b}]$, $\mathbf{x}[P^H(\mathbf{b})]$, and $\mathbf{x}[S^H(\mathbf{b})]$ are ϵ -approximate binary vectors, then $\mathbf{x}[\mathbf{z}^{in}]$ and $\mathbf{x}[\mathbf{z}^{out}]$ are 5ϵ -approximations to the centers of the entrance facet and exit facets, respectively.

Max-norm polar coordinates

We are now ready to compute the local max-norm polar coordinates of the projections of \mathbf{y} on the entrance and exit facets.

The max-norm radius is given by

$$r^{in} = \max_{i: b_i^{in}=0} |z_i^{in} - y_i|$$

Finding the maximum of a (length $2n+1$) vector requires some care when on each gate we can incur a constant error, the details are described in Algorithm 3.

The direction (max-norm) unit-vector, \mathbf{p} , is given by

$$p_i^{in} = (z_i^{in} - y_i) / r^{in}$$

Division is computed using DIVIDE, introducing an error of $O(\sqrt{\epsilon}/r^{in})$; this

Algorithm 3 MAX($| a_1, \dots, a_n | b$)

1. $G_\zeta(0 \parallel h_0)$

2. **for** each $k \in [1/\sqrt{\epsilon}]$:

(a) $G_\zeta(k\sqrt{\epsilon} \parallel \zeta_k)$

(b) $G_\zeta(0 \parallel d_{k,0})$

(c) **for** each $i \in [n]$:

i. $G_{<}(\zeta_k, a_i \parallel c_{k,i})$

The vector $(c_{k,i})_k$ is the unary representation of a_i :

$$\forall i \left(\max_{k: \mathbf{x}[c_{k,i}] > \epsilon} k\sqrt{\epsilon} \right) = \mathbf{x}[a_i] \pm \sqrt{\epsilon}$$

ii. $G_{\vee}(d_{k,i-1}, c_{k,i} \parallel d_{k,i})$

The vector $(d_{k,n})$ is the unary representation of $\max a_i$:

$$\left(\max_{k: \mathbf{x}[d_{k,n}] > \epsilon} k\sqrt{\epsilon} \right) = \max \mathbf{x}[a_i] \pm \sqrt{\epsilon}$$

(d) $G_{\times\zeta}(\sqrt{\epsilon} \parallel d_{k,n} \parallel e_k)$

The vector (e_k) is a $(\sqrt{\epsilon})$ -scaled version of (d_k)

$$\left(\sum \mathbf{x}[e_k] \right) = \max \mathbf{x}[a_i] \pm 3\sqrt{\epsilon}$$

(e) $G_+(h_{k-1}, e_k \parallel h_k)$

Finally, we sum the e_k 's

$$\mathbf{x}[h_{1/\sqrt{\epsilon}}] = \max \mathbf{x}[a_i] \pm 4\sqrt{\epsilon}$$

3. $G_=(h_{1/\sqrt{\epsilon}} \parallel b)$

approximation suffices because for $r^{in} < 1/48$, we multiply p_i^{in} by r^{in} when we interpolate. Also, we will use two nodes for each p_i^{in} to represent the positive and negative values. We do the same for $(r^{out}, \mathbf{p}^{out})$.

Claim 4. If $\mathbf{x}[\mathbf{y}]$ is an interior point, then $\mathbf{x}[r^{in}]$ and $\mathbf{x}[r^{out}]$ are $O(\sqrt{\epsilon})$ -approximations to the distances of $\mathbf{x}[\mathbf{y}]$ from $\mathbf{x}[\mathbf{z}^{in}]$ and $\mathbf{x}[\mathbf{z}^{out}]$, respectively. Furthermore, $\mathbf{x}[\mathbf{p}^{in}]$ and $\mathbf{x}[\mathbf{p}^{out}]$ are $O(\sqrt{\epsilon}/\mathbf{x}[r^{in}])$ - and $O(\sqrt{\epsilon}/\mathbf{x}[r^{out}])$ - to the unit-length vectors that point from $\mathbf{x}[\mathbf{y}]$ in the directions of $\mathbf{x}[\mathbf{z}^{in}]$ and $\mathbf{x}[\mathbf{z}^{out}]$, respectively.

The final displacement

Given \mathbf{p}^{in} and \mathbf{b}^{in} , we can compute \mathbf{g}^{in} for the special values of r^{in} . Recall that

$$\mathbf{g}^{in}(\langle r^{in}, \mathbf{p}^{in} \rangle) = \begin{cases} \delta(\mathbf{b}^{+in} - \mathbf{b}^{-in}) & r^{in} = 0 \\ -\delta\mathbf{p}^{in} & r^{in} = 1/48 \\ \delta(\mathbf{b}^{-in} - \mathbf{b}^{+in}) & r^{in} = 1/24 \\ \delta\xi_{2n+2} & r^{in} = 1/12 \end{cases}$$

We use Algorithm 4 to interpolate for intermediate values of r^{in} . We also need to interpolate between \mathbf{g}^{in} and \mathbf{g}^{out} . The ratio at which we interpolate is exactly the ratio between the distance of \mathbf{y} from the entrance and exit facets. We label the positive and negative components of this last interpolation $\mathbf{g}^{(interior)+}$ and $\mathbf{g}^{(interior)-}$, respectively.

When \mathbf{y} is close to both facets, the interpolation may be inaccurate; however, in this case it is a corner point. (We remark that this seems to be the only part of the proof which requires us to set the threshold for a corner point and the final error at $\Theta(\epsilon^{1/4})$ rather than $\Theta(\sqrt{\epsilon})$; this issue may be avoidable by a more careful construction of Algorithms 2 and 4.)

Claim 5. If $\mathbf{x}[\mathbf{y}]$ is an interior point of an intermediate subcube in the tube, and it is not a corner point, then $(\mathbf{x}[\mathbf{g}^{(interior)+}] - \mathbf{x}[\mathbf{g}^{(interior)-}])$ is a $O(\epsilon^{1/4})$ -approximation of the Hirsch et al. displacement $g(\mathbf{x}[\mathbf{y}])$.

Corner points We must ensure that if \mathbf{y} is a corner point, we set $\mathbf{g}^{(corner)+} = \delta\xi_{2n+2}$ and $\mathbf{g}^{(corner)-} = \mathbf{0}$: check over all pairs of coordinates whether $|\mathbf{x}[y_i] - 1/2| < 2\epsilon^{1/4}$ and $|\mathbf{x}[y_j] - 1/2| < 2\epsilon^{1/4}$. Let z be the variable representing the OR of those $\binom{2n+2}{2}$ indicators. Interpolate (e.g. using Algorithm 4) between the

Algorithm 4 INTERPOLATE($a, w_a, b, w_b \mid c$)

1. $G_{\times\zeta}(1/2 \mid w_a \mid w_{a/2})$ and $G_{\times\zeta}(1/2 \mid w_b \mid w_{b/2})$

We divide by 2 before adding in order to stay in $[0, 1]$

2. $G_+(\mid w_{a/2}, w_{b/2} \mid w_{a/2+b/2})$ Add the weights

3. $\text{DIVIDE}(\mid w_{a/2}, w_{a/2+b/2} \mid \bar{w}_a)$ and $\text{DIVIDE}(\mid w_{b/2}, w_{a/2+b/2} \mid \bar{w}_b)$,

\bar{w}_a and \bar{w}_b are the normalized weights:

$$\mathbf{x}[\bar{w}_a] \cdot (\mathbf{x}[w_a] + \mathbf{x}[w_b]) = \mathbf{x}[w_a] \pm (6\sqrt{\epsilon} + 8\epsilon) = \mathbf{x}[w_a] \pm (7\sqrt{\epsilon})$$

4. $\text{MULTIPLY}(\mid \bar{w}_a, a \mid c_a)$ and $\text{MULTIPLY}(\mid \bar{w}_b, b \mid c_b)$

c_a and c_b are the a and b components, respectively, of c :

$$\mathbf{x}[c_a] = \mathbf{x}[w_a] \cdot \mathbf{x}[a] / (\mathbf{x}[w_a] + \mathbf{x}[w_b]) \pm 15\sqrt{\epsilon} / (\mathbf{x}[w_a] + \mathbf{x}[w_b])$$

5. $G_+(\mid c_a, c_b \mid c)$

Finally, c is the interpolation of a and b :

$$\begin{aligned} \mathbf{x}[c] &= (\mathbf{x}[w_a] \cdot \mathbf{x}[a] + \mathbf{x}[w_b] \cdot \mathbf{x}[b]) / (\mathbf{x}[w_a] + \mathbf{x}[w_b]) \\ &\quad \pm 30\sqrt{\epsilon} / (\mathbf{x}[w_a] + \mathbf{x}[w_b]) \end{aligned}$$

$(\mathbf{g}^{(interior)+}, \mathbf{g}^{(interior)-})$ we computed earlier and $\delta\xi_{2n+2}$ with weights z and $-z$. Label the result of the interpolation $(\mathbf{g}^{(tube)+}, \mathbf{g}^{(tube)-})$.

We remark that whenever z is ambiguous, i.e. the second smallest $|\mathbf{x}[y_i] - 1/2|$ is very close to $2\epsilon^{1/4}$, then we cannot predict the value of $\mathbf{x}[z]$; it can take any value in $[0, 1]$. Nevertheless, in this case $\mathbf{x}[\mathbf{y}]$ is not a corner point, thus for most $\mathbf{y}^l \in E^\epsilon$, $\mathbf{x}[\mathbf{y}^l]$ is an interior point. This means that by Claim 5, we would compute the (approximately) correct interior displacement $(\mathbf{x}[\mathbf{g}^{(interior)+}] - \mathbf{x}[\mathbf{g}^{(interior)-}])$. Since $\mathbf{x}[\mathbf{y}]$ is close to a corner point, this value is very close to $\delta\xi_{2n+2} = (\mathbf{x}[\mathbf{g}^{(corner)+}] - \mathbf{x}[\mathbf{g}^{(corner)-}])$. Therefore, although we don't know the value of $\mathbf{x}[z]$, we use it to interpolate between two (approximately) equal vectors - so the result is guaranteed to be (approximately) correct regardless of the value of $\mathbf{x}[z]$.

Claim 6. If $\mathbf{x}[\mathbf{y}]$ is a corner point, then $(\mathbf{x}[\mathbf{g}^{(tube)+}] - \mathbf{x}[\mathbf{g}^{(tube)-}])$ is a $O(\sqrt{\epsilon})$ -approximation of $\delta\xi_{2n+2}$, and thus also a $O(\epsilon^{1/4})$ -approximation of the Hirsch et al. displacement $g(\mathbf{x}[\mathbf{y}])$. Furthermore, Claim 5 continues to hold for $(\mathbf{g}^{(tube)+}, \mathbf{g}^{(tube)-})$.

Start/end subcubes and subcubes outside the tube For start/end subcubes (except the home subcube) we use a slightly different (Cartesian) interpolation that yields a fixed point in the center of the subcube, and a displacement of $\delta\xi_{2n+2}$ on all facets but the exit/entrance facet, respectively. For subcubes in the picture but outside the tube, we again set $\mathbf{g} = \delta\xi_{2n+2}$.

Notice that we can infer the type of subcube from the following two-bits vector:

$$T = (\vee_i b_i^{in}, \vee_i b_i^{out})$$

If $T = (0, 0)$, the subcube is outside the tube; when $T = (0, 1)$, we are at a start subcube, while $T = (1, 0)$ corresponds to an end subcube; $T = (1, 1)$ is an intermediate subcube in the tube. Finally, interpolate between the displacement for each type of subcube using T and $-T$; label the result of the interpolation $(\mathbf{g}^+, \mathbf{g}^-)$.

Claim 7. If $\mathbf{x}[\mathbf{y}]$ is either an interior point or a corner point, of any subcube in the slice, then $(\mathbf{x}[\mathbf{g}^+] - \mathbf{x}[\mathbf{g}^-])$ is an $O(\epsilon^{1/4})$ -approximation of the Hirsch et al. displacement $g(\mathbf{x}[\mathbf{y}])$.

□

5.3 Summing the displacement vectors

We are now ready to average over the displacement vectors we computed for each \mathbf{y}^l . Using $G_{\times\zeta}$ and G_+ we have that

$$\mathbf{x}[g_i^+] = \sum_{l=1}^{1/\sqrt{\epsilon}} (\sqrt{\epsilon} \mathbf{x}[g_i^{l+}]) \pm O(\sqrt{\epsilon}) \quad \text{and} \quad \mathbf{x}[g_i^-] = \sum_{l=1}^{1/\sqrt{\epsilon}} (\sqrt{\epsilon} \mathbf{x}[g_i^{l-}]) \pm O(\sqrt{\epsilon})$$

Lemma 2. *For every input $\mathbf{x}[\mathbf{y}]$ and every $i \in [2n+2]$,*

$$\mathbf{x}[g_i^+] - \mathbf{x}[g_i^-] = g_i(\mathbf{x}[\mathbf{y}]) \pm O(\epsilon^{1/4})$$

Proof. By Fact 2, every $\mathbf{y}^l \in E^\epsilon$, except at most one, is either an interior point or a corner point. By Lemma 1, for all those \mathbf{y}^l , $(\mathbf{x}[g_i^{+l}] - \mathbf{x}[g_i^{-l}])$ is at most $O(\epsilon^{1/4})$ -far from the right displacement. The single point which is neither an interior point nor a corner point increases the error by at most $O(\sqrt{\epsilon})$, as does the summation above. Finally, because g is λ -Lipschitz for constant λ , the error added due to the distance between the \mathbf{y}^l 's is again at most $O(\sqrt{\epsilon})$. \square

5.4 Closing the loop

Finally, for each $i \in [2n+2]$: $G_+(\mid y_i^1, g_i^+ \mid y_i')$, $G_+(\mid y_i', g_i^- \mid y_i'')$ and $G_=(\mid y_i'', nil \mid y_i^1)$.

6 ϵ -GCIRCUIT with fan-out 2

In the previous section, we proved that ϵ -GCIRCUIT is PPAD-complete for some constant $\epsilon > 0$. Each generalized circuit gate has fan-in at most 2, which would eventually correspond to a bound on the *incoming* degree of each player in the graphical game. In order to bound the total degree (as well as for the application to A-CEEI), we need to also bound fan-out of each gate.

Theorem 2 (Generalized circuit). *There exists a constant $\epsilon > 0$ such that ϵ -GCIRCUIT with fan-out 2 is PPAD-complete.*

Proof. We present a black-box reduction that converts a general ϵ' -GCIRCUIT instance to an instance of ϵ -GCIRCUIT with fan-out 2, for $\epsilon' = \Theta(\sqrt{\epsilon})$. Daskalakis et al. [8] bound the fan-out of the generalized circuit by introducing a binary tree of $G_=($ gates. Unfortunately, this blows up the error: each $G_=($ gate introduces

Algorithm 5 REAL2UNARY($| a | b_1, \dots, b_{4/\epsilon'}$)

1. $G_=(| a | c_0)$

2. **for** each $k \in [4/\epsilon']$:

(a) $G_=(| c_{k-1} | c_k)$

The c_k 's are simply copies of a , ensuring that each gate has fan-out at most 2:

$$\forall k \quad \mathbf{x}[c_k] = \mathbf{x}[a] \pm (k+1)\epsilon$$

(b) $G_\zeta(k\epsilon'/4 || \zeta_k)$

(c) $G_<(| \zeta_k, c_k | b_k)$

The vector (b_k) is the unary representation of a :

$$\forall i \quad \left(\max_{k: \mathbf{x}[b_k] > \epsilon} k\epsilon'/4 \right) = \mathbf{x}[a] \pm \epsilon'/2$$

an additive error of ϵ , so the increase is proportional to the depth of the tree, i.e. $\Theta(\epsilon \cdot \log n)$. While this was acceptable for [8] who used an exponentially small ϵ , it clearly fails in our setting.

We overcome this obstacle by resorting to logical gates ($G_<$, G_\vee , G_\wedge , and in particular G_-). Recall that the logical gates are at most ϵ' -far from the correct Boolean value. Therefore, concatenating multiple logical gates does not amplify the error. In particular, if any logical gate has a large fan-out, we can distribute its output using a binary tree of G_- gates (we use trees of even depth). When the gate is arithmetic (G_ζ , $G_{\times\zeta}$, $G_=$, G_+ , or G_-), we convert its output to unary representation over $\Theta(\epsilon')$ logical gates (Algorithm 5). Then, we copy the unary representation using trees of G_- gates. Finally, we use $G_{\times\zeta}$ and G_+ gates to convert each copy of the unary representation back to a real value in $[0, 1]$ (Algorithm 6).

It is interesting to note that for constant ϵ and ϵ' , the unary representation has constant size, so the number of new gates is proportional to the original fan-out (i.e. the number of leaves of the binary tree that copies the unary representation). In particular this reduction increases the size of the circuit by a factor of $\Theta(1/\epsilon')$.

Algorithm 6 UNARY2REAL($| b_1, \dots, b_{4/\epsilon'} | a'$)

1. $G_\zeta(0 || d_0)$

2. **for** each $k \in [4/\epsilon']$:

(a) $G_{\times\zeta}(\epsilon'/4 | b_k | c_k)$

The sum of the c_k 's is approximately a :

$$\sum \mathbf{x}[c_k] = \mathbf{x}[a] \pm (\epsilon'/2 + \epsilon \cdot 4/\epsilon')$$

(b) $G_+(| d_{k-1}, c_k | d_k)$

3. $G_=(| d_k | a')$

a' approximately recovers the original a :

$$\mathbf{x}[a'] = \mathbf{x}[a] \pm \epsilon'$$

□

7 From ϵ -GCIRCUIT to ϵ^2 -ANE

The last step in our proof is to reduce the problem of finding an ϵ -approximate solution to the generalized circuit to that of finding an $\Theta(\epsilon^2)$ -ANE. First, we reduce to the problem of finding an ϵ -NE in a (polymatrix, degree 3) graphical game. This reduction is implicit in Daskalakis et al. [8]⁷.

Definition 7. Given a polymatrix game over a degree d graph with k actions for each player, ϵ -NE-POLYMATRIX(d, k) is the problem of finding an ϵ -Nash equilibrium.

Theorem. (Essentially [8]⁸) 2ϵ -GCIRCUIT with fan-out $2 \leq_P \epsilon$ -NE-POLYMATRIX(3, 2)

Daskalakis et al. construct a different gadget for each type of gate. For example, we describe their gadget for the $G_{\times\zeta}$ gate below.

⁷In [8], polymatrix games are called *games with additive utility functions*.

⁸In [8] polymatrix games are called *additive games*.

Lemma 3. ($G_{\times\zeta}$ gadget, [8])

Let v_1, v_2 , and w be players in a graphical game, and suppose that the payoffs of v_2 and w are as follows.

	w plays 0	w plays 1
Payoff for v_2:		
v_2 plays 0	0	1
v_2 plays 1	1	0

Payoffs for w :

	v_1 plays 0	v_1 plays 1
game with v_1:		
w plays 0	0	ζ
w plays 1	0	0
	v_2 plays 0	v_2 plays 1
game with v_2:		
w plays 0	0	0
w plays 1	0	1

Then, in every ϵ -NE $\mathbf{p}[v_2] = \min(\zeta\mathbf{p}[v_1], 1) \pm \epsilon$, where $\mathbf{p}[u]$ denotes the probability that u assigns to strategy 1.

Proof. (Sketch) If $\mathbf{p}[v_2] > \zeta\mathbf{p}[v_1] + \epsilon$, then in every ϵ -NE $\mathbf{p}[w] = 1$, which contradicts $\mathbf{p}[v_2] > \epsilon$. Similarly, if $\mathbf{p}[v_2] < \min(\zeta\mathbf{p}[v_1], 1) - \epsilon$, then $\mathbf{p}[w] = 0$, which yields a contradiction when $\mathbf{p}[v_2] < 1 - \epsilon$. \square

7.1 From $\sqrt{\epsilon}$ -NE to $\Theta(\epsilon)$ -ANE

The reduction above shows hardness for the slightly stronger notion (therefore weaker hardness) of ϵ -Nash equilibrium. Daskalakis et al. [8] show a reduction from $\sqrt{\epsilon}$ -NE to $\Theta(\epsilon)$ -ANE for games with a constant number of players. It is easy to see that the same reduction holds for graphical games with constant degree. We sketch the proof below.

Lemma 4. (Essentially [8]) Given an ϵ -ANE of a graphical game with payoffs in $[0, 1]$ and incoming degree d_{in} , we can construct in polynomial time a $\sqrt{\epsilon} \cdot (\sqrt{\epsilon} + 1 + 4d_{in})$ -NE.

Proof. Let V be the set of players, where each $v \in V$ has utility u^v and action set A^v . Let $\mathbf{x} = (x_a^v) \in \Delta(\times_v A^v)$ be an ϵ -ANE. Let $U_a^v(\mathbf{x}^{-v})$ denote the expected payoff for playing a when all other players play according to \mathbf{x} . Note that

$U_a^v(\mathbf{x}^{-v}) = U_a^v(\mathbf{x}^{\mathcal{N}_{in}(v)})$ depends only on the distributions of the players in the incoming neighborhood of v , which we denote $\mathcal{N}_{in}(v)$. Finally, let $U_{\max}^v(\mathbf{x}^{-v}) = \max_{a \in A^v} U_a^v(\mathbf{x}^{-v})$.

Let $k = k(\epsilon) > 0$ be some large number to be specified later. We construct our new approximate equilibrium by taking, for each player only the strategies that are within ϵk of the optimum:

$$\hat{x}_a^v = \begin{cases} \frac{\hat{x}}{1-z^v} & U_a^v(\mathbf{x}^{-v}) \geq U_{\max}^v(\mathbf{x}^{-v}) - \epsilon k \\ 0 & \text{otherwise} \end{cases}$$

where z^v is the total probability that player v assigns to strategies that are more than ϵk away from the optimum.

The above division is well-defined because for $k > 1$, z^v is bounded away from 1. Moreover, the following claim from [8] formalizes the intuition that when k is sufficiently large, the total weight on actions removed is small, so $\hat{\mathbf{x}}^v$ is close to \mathbf{x}^v :

Claim 8. (Claim 6 of [8])

$$\forall v \in V \quad \sum_{a \in A^v} |\hat{x}_a^v - x_a^v| \leq \frac{2}{k-1}$$

Now, the total change to the expected payoff to player v for each action a , is bounded by the total change in mixed strategies of its *incoming neighbors*:

$$\begin{aligned} |U_a^v(\mathbf{x}^{-v}) - U_a^v(\hat{\mathbf{x}}^{-v})| &= |U_a^v(\mathbf{x}^{\mathcal{N}_{in}(v)}) - U_a^v(\hat{\mathbf{x}}^{\mathcal{N}_{in}(v)})| \\ &\leq \sum_{\mathbf{a} \in A^{\mathcal{N}(v)}} |\hat{x}_{\mathbf{a}}^{\mathcal{N}_{in}(v)} - x_{\mathbf{a}}^{\mathcal{N}_{in}(v)}| \leq \sum_{w \in \mathcal{N}(v)} \sum_{a \in A^w} |\hat{x}_a^w - x_a^w| \leq \frac{2d_{in}}{k-1} \end{aligned}$$

It follows that \hat{x}_a^v is a $(k\epsilon + \frac{4d_{in}}{k-1})$ -NE:

$$\hat{U}_a^v(\mathbf{x}^{-v}) \geq U_a^v(\mathbf{x}^{-v}) - \frac{2d_{in}}{k-1} \geq U_{\max}^v(\mathbf{x}^{-v}) - \epsilon k - \frac{2d_{in}}{k-1} \geq \hat{U}_{\max}^v(\mathbf{x}^{-v}) - \epsilon k - \frac{4d_{in}}{k-1}$$

Finally, take $k = 1 + 1/\sqrt{\epsilon}$ to get that

$$k\epsilon + \frac{4d_{in}}{k-1} \leq \sqrt{\epsilon} \cdot (\sqrt{\epsilon} + 1 + 4d_{in})$$

□

8 ϵ -approximate Bayesian Nash equilibrium

In this section we formally introduce and prove our corollary for ϵ -approximate Bayesian Nash equilibrium in two-player games with incomplete information.

In a game with incomplete information, each player i has a type t_i known only to her, and the players' types $\mathbf{t} = (t_1, t_2)$ are drawn from a joint distribution which is known to everyone. The payoff for player i is a function $u_i(\mathbf{a}, t_i)$ of her own type and all the players' actions.

Definition 8. (e.g. [25])

In a *Bayesian Nash equilibrium*, for every player i and every type t_i , the mixed strategy $x_i(t_i)$ must be a best response in expectation over the other players' types and actions:

$$\mathbb{E}_{\mathbf{t}|t_i} \left[\mathbb{E}_{\mathbf{a} \sim (x_i(t_i), x_{-i}(t_{-i}))} [u_i(\mathbf{a}; t_i)] \right] \geq \max_{x'_i(t_i) \in \Delta A_i} \mathbb{E}_{\mathbf{t}|t_i} \left[\mathbb{E}_{\mathbf{a} \sim (x'_i(t_i), x_{-i}(t_{-i}))} [u_i(\mathbf{a}; t_i)] \right].$$

Similarly, in an ϵ -approximate Bayesian Nash equilibrium, for every player i and every type t_i , the mixed strategy $x_i(t_i)$ must be an ϵ -best response in expectation over the other players' types and actions:

$$\mathbb{E}_{\mathbf{t}|t_i} \left[\mathbb{E}_{\mathbf{a} \sim (x_i(t_i), x_{-i}(t_{-i}))} [u_i(\mathbf{a}; t_i)] \right] \geq \max_{x'_i(t_i) \in \Delta A_i} \mathbb{E}_{\mathbf{t}|t_i} \left[\mathbb{E}_{\mathbf{a} \sim (x'_i(t_i), x_{-i}(t_{-i}))} [u_i(\mathbf{a}; t_i)] \right] - \epsilon.$$

Before we prove our main corollary for incomplete information games, it is helpful to prove the following slightly weaker statement, for two players with many strategies.

Lemma 5. *There exists a constant $\epsilon > 0$, such that given a two-player game with incomplete information where each player has n actions, finding an ϵ -approximate Bayesian Nash equilibrium is PPAD-complete.*

Proof. We reduce from a bipartite polymatrix game, and let the typeset for each of the two players in the incomplete information game correspond to one side of the bipartite graph. The utility of player i on edge (i, j) of the polymatrix game depends on her identity (i), as well as the identity (j) of the vertex on the other side of that edge. We use the types of the incomplete information game, to encode i . We encode j using the strategies of the second player in the incomplete information game.

In more detail, consider a bipartite polymatrix game for which it is PPAD-hard to compute a 4ϵ -approximate Nash equilibrium. Use an affine transformation to change all the payoffs from $[0, 1]$ to $[1/2, 1]$. It is PPAD-hard to find a 2ϵ -approximate Nash equilibrium in the transformed game.

We now construct the two-player incomplete information game: As we hinted before, we let the typeset of each player correspond to the vertices on one side of the bipartite graphical game. Player i has $|T_i|$ types and $2|T_i|$ strategies, where each strategy corresponds to a pair of vertex and strategy for that vertex. If a player plays a strategy whose vertex does not match her type, her payoff is 0. Therefore in every ϵ -approximate Bayesian Nash equilibrium, every player, on every type, plays the two strategies that correspond to her type with probability at least $1 - 2\epsilon$.

Let the joint distribution over types be as follows: pick a random edge in the bipartite graph, and assign the types corresponding to its vertices. Whenever both players play strategies that match their respective types, their payoffs are the payoffs in the (transformed) bimatrix game on that edge. In every ϵ -approximate Bayesian Nash equilibrium, every player, on every type, plays a mixed strategy which is ϵ -best response. Since the other player plays strategies that correspond to the correct vertex with probability at least $1 - 2\epsilon$, the same mixture must be a 2ϵ -best response for the vertex player in the bipartite game. \square

In order to prove the main corollary, we need to reduce the number of actions in the above construction. Observe that we don't need each player to choose an action that uniquely identifies her type. Rather, it suffices to specify which neighbor of the other player's vertex is chosen. This can be done concisely via a coloring of the vertices such that every pair of vertices of distance exactly two have different colors; i.e. a coloring of the square of the polymatrix game's graph. The squared graph has degree $3 \cdot (3 - 1) = 6$, and therefore we can efficiently find a 7-coloring. It suffices for each player to specify one of 7 colors, together with one of 2 strategies for the vertex with that color. Therefore we can encode this choice using only 14 strategies.

Corollary 1. *There exists a constant $\epsilon > 0$, such that given a two-player game with incomplete information where each player has 14 actions, finding an ϵ -approximate Bayesian Nash equilibrium is PPAD-complete.*

9 Relative ϵ -Nash equilibrium

In this section we formally introduce and prove our result for relative ϵ -Nash equilibrium in two-player games.

As we mentioned in Section 3.1, we say that a vector of mixed strategies $\mathbf{x} \in \times_j \Delta A_j$ is a Nash equilibrium if every strategy a_i in the support of x_i is a best response to the actions of the mixed strategies of the rest of the players, x_{-i} ;

$$\mathbb{E}_{a_{-i} \sim x_{-i}} [u_i(a_i, a_{-i})] = \max_{a' \in A_i} \mathbb{E}_{a_{-i} \sim x_{-i}} [u_i(a', a_{-i})].$$

In this section, we study the multiplicative relaxation of this condition. Namely,

Definition 9. We say that \mathbf{x} is a *relative ϵ -Nash equilibrium* (*relative ϵ -NE*; sometimes also *relative ϵ -well supported Nash equilibrium*) if each a_i in the support of x_i is a relative ϵ -best response to x_{-i} . I.e. $\forall a_i \in \text{Supp}(x_i)$,

$$\mathbb{E}_{a_{-i} \sim x_{-i}} [u_i(a_i, a_{-i})] \geq (1 - \epsilon) \max_{a' \in A_i} \mathbb{E}_{a_{-i} \sim x_{-i}} [u_i(a', a_{-i})].$$

Corollary 2 (Relative ϵ -Nash equilibrium). *There exists a constant $\epsilon > 0$ such that finding a relative ϵ -Nash equilibrium in a bimatrix game with positive payoffs is PPAD-complete.*

Proof. We reduce from a bipartite, degree 3, polymatrix game (Theorem 1). The payoff bimatrix is the sum of a main game that simulates the polymatrix game and an imitation game that forces the two players to randomize (approximately) uniformly across all the polymatrix game nodes.

Main game As in the reduction to games of incomplete information, let each player “control” the nodes on one side of the bipartite game graph. Namely, let n be the number nodes on each side of the graph; each player has $2n$ actions, each corresponding to a choice of a node and an action for that node. (We assume without loss of generality that the number of nodes is equal and every node has exactly 3 neighbors.) When the players play strategies that correspond to adjacent nodes in the graphs, they receive the payoffs of the corresponding nodes, scaled by a small positive constant, e.g. $\eta = 0.01$, to fit in $[0, \eta]$. If the nodes played do not share an edge in the bipartite game graph, the utility for both players is zero.

Imitation game We also let the players play a hide-and-seek game with payoffs in $\{0, 1\}$ over a space of n actions. Fix an arbitrary ordering over the nodes on each side of the polymatrix game; action i in the imitation game corresponds to the two rows (respectively, two columns) for node i in the main game. The row player tries to imitate the column player; namely her payoff is 1 if they both play their respective node i , and 0 otherwise. The column player tries to imitate the row player $+1$; her payoff is 1 if the row player plays node i , and she plays node $i+1 \pmod{n}$ (we henceforth drop the \pmod{n} for simplicity of notation).

Structure and value of a relative ϵ -NE

We first claim that in every relative ϵ -NE (x, y) , the two players randomize approximately uniformly across all their nodes. Let $x(i)$ denote the total probability that the row player assigns to node i , and analogously for $y(i)$.

Lemma 6. *In every (x, y) relative ϵ -NE, $x(i), y(i) \in [(1 - 4\eta)/n, (1 + 5\eta)/n] \forall i$.*

Proof. Let $t_R = \max_i x(i)$ and $i_R^* = \arg \max_i x(i)$; we first observe that if $x(j) \leq (1 - 4\eta)t_R < (1 - \epsilon - 3\eta)t_R$, then $y(j+1) = 0$. The column player's payoff from the imitation game, when playing node $j+1$, is at most $x(j)$, while from the main game she can gain at most $x(k) \cdot \eta$ for each row player's node k in the neighborhood of the column player's $j+1$. Since each node has only three neighbors and they all satisfy $x(k) \leq t_R$, her payoff for playing either strategy of node $j+1$ is bounded by $x(j) + 3t_R \cdot \eta < (1 - \epsilon)t_R$. However, she can guarantee a payoff of at least t_R by playing node i_R^* ; therefore by the well-supported condition, $y(j+1) = 0$. Similarly, if we let $t_C = \max_i y(i)$, then whenever $y(j) \leq (1 - 4\eta)t_C$, $x(j) = 0$.

Next, we claim that both players have full support in the imitation game, i.e. $x(i), y(i) > 0 \forall i \in [n]$. Assume by contradiction that this is false, and let, without loss of generality, $x(i) = 0$. Since $x(i) < (1 - 4\eta)/n \leq (1 - 4\eta)t_R$, we also have $y(i+1) = 0$. However, this implies $x(i+1) = 0$, which in turn implies $y(i+2) = 0$, and so forth. By induction we get that $x(i), y(i) = 0 \forall i \in [n]$, which of course cannot be true as their sum is 1.

Therefore, for every $i \in [n]$, we also have that

$$x(i) > (1 - 4\eta)t_R \geq (1 - 4\eta) \mathbf{E}_j x(j) = (1 - 4\eta)/n, \quad (2)$$

and similarly $y(i) > (1 - 4\eta)/n$.

Finally, we claim that the maximum probability assigned to any node is at most $t_R, t_C \leq 1/((1 - 4\eta)n) < (1 + 5\eta)/n$. Otherwise, $\mathbf{E}_i x(i) < (1 - 4\eta)t_R$, which in particular implies that for some $i \in [n]$, $x(i) < (1 - 4\eta)t_R$, violating (2) \square

Corollary 6. *In every (x, y) relative ϵ -NE, the expected utilities of both players for playing any strategy is in $[(1 - 4\eta)/n, (1 + 9\eta)/n]$.*

Completing the proof of Corollary 2

Given a relative ϵ -NE (x, y) , we can take, for each node i , the mixed strategy induced by the probabilities $x(i : 1)/x(i)$ and $x(i : 2)/x(i)$ that the row player assigns to each action (respectively, $y(j : 1)/y(j)$ and $y(j : 2)/y(j)$ assigned by the column player). We claim that this strategy profile is an ϵ' -approximate equilibrium for the polymatrix game for some small constant $\epsilon' = \Theta(\epsilon)$ (which is PPAD-hard to find by Theorem 1).

Assume by contradiction that this is not the case. Then without loss of generality there exists a node i which is controlled by the row player, and, under the induced strategy profile, can improve her payoff by ϵ' by deviating from strategy 1 to strategy 2. Therefore in the two player game, the expected payoff to the row player from playing $x(i : 2)$ is greater than the payoff from playing $x(i : 1)$ by at least $\epsilon' \cdot \eta \cdot \frac{1}{n} \cdot (1 - O(\eta))$. (We multiply ϵ' by the scaling factor η and the probability $\frac{1}{n} \cdot (1 - O(\eta))$ of observing any of its neighbors.) Yet, by Corollary 6 the total payoff from playing $x(i : 2)$ is at most $(1 + 9\eta)/n$. Therefore, the row player improves her expected payoff by a factor of $1 + \Omega(\epsilon' \cdot \eta) \geq 1 + 2\epsilon$ when deviating from $x(i : 1)$ to $x(i : 2)$ - but this contradicts our assumption that (x, y) is a relative ϵ -NE. \square

10 Non-monotone markets

In this section we formally introduce our result for non-monotone markets, discuss it, and compare it to the original result of Chen, Paparas, and Yannakakis [5]. A sketch of the proof appears in Appendix A.

Intuitively, a market is monotone if increasing the price of some good, while fixing the rest of the prices, never increases the excess demand for that good. Formally, we have the following definition by Chen et al.:

Definition 10. ([5]) Let M be a market with $k \geq 2$ goods. We say that M is *non-monotone* at price vector $\boldsymbol{\pi}$ if there exist $c > 0$, and a good g_1 such that:

- the excess demand $Z_1(y_1, \dots, y_k)$ is a continuous function (rather than correspondence) over $\mathbf{y} \in B(\boldsymbol{\pi}, c)$;
- $Z_1(\boldsymbol{\pi}) > 0$;
- the partial derivative of $\partial Z_1 / \partial y_1$ exists and is continuous over $B(\boldsymbol{\pi}, c)$;
- and $\partial Z_1 / \partial y_1(\boldsymbol{\pi}) > 0$.

We say that a market M is *non-monotone* if there exists such a rational price vector $\boldsymbol{\pi} \geq \mathbf{0}$, and $Z_1(\boldsymbol{\pi})$ is *moderately computable*; i.e. for any $\gamma > 0$, $Z_1(\boldsymbol{\pi})$ can be approximated to within γ in time polynomial in $1/\gamma$.

In general we want to talk about non-monotone families of utility functions, i.e. ones that support non-monotone markets. Formally,

Definition 11. ([5]) We say that a family \mathcal{U} of utility functions is *non-monotone* if:

- \mathcal{U} is countable;
- if $u: [0, \infty)^k \rightarrow \mathbb{R}$ is in \mathcal{U} , then so is $u'(x_1, \dots, x_m) = a \cdot u(x_{l_1}/b_1, \dots, x_{l_k}/b_k)$ for any indices $l_1, \dots, l_k \in [m]$ and positive (rational) a, b_1, \dots, b_k ;
- $u(\mathbf{x}) = g^*(x_i)$ is in \mathcal{U} for some strictly increasing $g: [0, \infty) \rightarrow \mathbb{R}$; and
- there exists a non-monotone market $M_{\mathcal{U}}$ with utilities from \mathcal{U} .

We need to include one more definition: that of ϵ -tight market equilibrium.

Definition 12. A price vector $\boldsymbol{\pi}$ is an ϵ -tight approximate market equilibrium of M if there exists a $\mathbf{z} \in Z(\boldsymbol{\pi})$ (the excess demand at $\boldsymbol{\pi}$) such that for every good j , $|z_j| \leq \epsilon W_j$, where W_j is the sum of the endowments of good j .

Our main result for non-monotone markets equilibria is now formally defined:

Corollary 3 (Non-monotone markets). *Let \mathcal{U} be any non-monotone family of utility functions. There exists a constant $\epsilon_{\mathcal{U}} > 0$ such that given a market M where the utility of each trader is either linear or taken from \mathcal{U} , finding an $\epsilon_{\mathcal{U}}$ -tight approximate market equilibrium is PPAD-hard.*

10.1 Why are non-monotone markets hard?

Before delving into the details of the construction, we attempt to reach some intuition: why should we expect equilibrium computation to be hard in non-monotone markets? Probably the most intuitive algorithm for finding market equilibrium is via tatonnement: raise the prices of over-demanded goods, and decrease the prices of under-demanded goods. For many markets, the tatonnement process is also computationally efficient [6]. One obvious problem is that when the market is non-monotone, the tatonnement step actually takes us further away from equilibrium. However, the non-monotonicity is only local: if we set the (relative) price of the non-monotone good high enough, even the most enthusiastic traders can only afford a small amount.

The “real” reason that tatonnement fails to converge efficiently for non-monotone markets is a little more subtle. What happens when the demand for the non-monotone good g increases by a factor of $(1 + \delta)$ for some small δ ? The tatonnement increases the price of g , which further increases the demand. Eventually, the price is high enough, and the demand is reduced; but due to the non-monotonicity we may have to increase the price by larger factor, i.e. $(1 + \delta')$ for $\delta' > \delta$. Now, another trader with a positive endowment of g has increased her spending budget by $(1 + \delta')$, further increasing the demand for yet another good (by a larger factor). Thus a small change in the demand for one good may cause a much larger change in the demand for another good. Exploiting this “butterfly effect” lies at the heart of Chen et al.’s construction.

10.2 High-level structure of the proof

Our reduction from polymatrix games to non-monotone markets closely follows the footsteps of [5]. To gain some intuition, consider two goods g_{2i-1} and g_{2i} for each player i , corresponding to her two available strategies (soon each of those goods will become a subset of goods). Let $\pi(g_{2i-1})$ and $\pi(g_{2i})$ denote their corresponding prices; those prices correspond to the probabilities that player i assigns to her respective strategies. For every $i, j \in [n]$, we add a trader who is interested in selling g_{2i-1} and buying g_{2j-1} (and similarly for $(2i, 2j - 1)$, $(2i - 1, 2j)$, and $(2i, 2j)$). This trader has an endowment of g_{2i-1} that is proportional to $P_{2i-1, 2j-1}$, the utility of player j in the bimatrix game with player i , when they both play the first strategy. Qualitatively, if the price $\pi(g_{2i-1})$ is high (player i assigns a high probability to her first strategy), and $P_{2i-1, 2j-1}$ is

high (player j receives a high utility from playing her first strategy to i 's first strategy), then the demand for good g_{2j-1} is high - implying a high price in every approximate market equilibrium (i.e. player j indeed assigns a high probability to her first strategy).

In order to turn this qualitative intuition into a proof we use a more complex construction. The main difficulty comes from the need to amplify the effect of a higher income for one trader on the incomes of other traders. To this end we consider, for each $i \in [n]$, two sequences of goods: $g_{2i-1} = g_{2i-1,0}, g_{2i-1,1}, \dots, g_{2i-1,4t} = h_{2i-1}$ and $g_{2i} = g_{2i,0}, g_{2i,1}, \dots, g_{2i,4t} = h_{2i}$. The trader mentioned in the previous paragraph actually owns $P_{2i-1,2j-1}$ units of good h_{2i-1} ; she is still interested in good g_{2j-1} . Now we construct (Lemma 13) a chain of gadgets that use copies of the non-monotone markets in \mathcal{U} to amplify the small gap guaranteed between $\pi(g_{2j-1})$ and $\pi(g_{2j})$ to a larger gap between $\pi(h_{2j-1})$ and $\pi(h_{2j})$.

Additionally, we want to bound the range that these prices may take. In Lemma 10 we use a price regulating gadget [3, 28] to control the relative prices of $\pi(g_{2i-1,j})$ and $\pi(g_{2i,j})$. In Lemma 12 we show that the sums $\pi_{i,j} = \pi(g_{2i-1,j}) + \pi(g_{2i,j})$ are approximately equal. Finally, in Section A.5 we combine these lemmata to formalize a quantitative version of the qualitative intuition described above.

10.3 Adaptations for constant factor inapproximability

As mentioned in the introduction, Corollary 3 has a weakness in comparison to the results of Chen et al. [5]: it only applies to *tight* approximate market equilibrium.

Maintaining the constant hardness of approximation through most of [5]'s proof is rather straightforward, but there are a few hurdles along the way. To understand the first obstacle, we must consider a subtle issue of normalization. Chen et al. normalize the bimatrix game between every pair of players to have an average value of $1/2$. While this does not change the absolute utility gained from any deviation, the relative utility from deviation is now divided by a factor of $\Theta(n)$. In contrast, in Theorem 1 we prove hardness for a constant ϵ when normalizing with respect to a constant degree (3), i.e. each player participates in only a constant number of bimatrix games. We overcome this difficulty by using a different normalization: only edges (i.e. bimatrix games) that belong to the game graph will have an average utility of $1/2$, while the utilities on other

edges remains 0. Since we proved hardness for a degree 3 graphical game, the normalization only costs us a constant factor.

More serious complications arise when trying to prove [5]’s Lemma 3⁹ for a constant ϵ . This lemma says that certain prices (in fact, these are sums of prices), denoted $\pi_{i,0}$ for $i \in [n]$, are approximately equal. A key step in the proof of [5] is to show, roughly, that in every $\epsilon(n)$ -approximate market equilibrium,

$$\pi_{i,0} \geq \frac{1}{n} \sum_{j \in [n]} \pi_{j,0} - O(\epsilon(n))$$

When $\epsilon(n)$ is polynomially small, this immediately implies that $\min_{i \in [n]} \pi_{i,0}$ is within $O(\epsilon(n))$ of the average, and therefore it must also be that $\max_{i \in [n]} \pi_{i,0}$ is within $O(n \cdot \epsilon(n))$ of the average. When taking a larger $\epsilon(n)$, this reasoning breaks. The first modification we make to overcome this obstacle, is to require $\epsilon(n)$ -tight approximate market equilibrium. This gives a two-sided bound:

$$\left| \pi_{i,0} - \frac{1}{n} \sum_{j \in [n]} \pi_{j,0} \right| = O(\epsilon(n)) \quad (3)$$

A second issue that arises in the same inequality, is that with our new normalization, which depends on the game graph G , we can only prove that $\pi_{i,0}$ approximates the values of its neighbors, denoted $\mathcal{N}_G(i)$. In other words, (3) becomes

$$\left| \pi_{i,0} - \frac{1}{|\mathcal{N}_G(i)|} \sum_{j \in \mathcal{N}_G(i)} \pi_{j,0} \right| = O(\epsilon(n)) \quad (4)$$

In order to relate the value of $\pi_{i,0}$ to the corresponding values of the neighboring vertices, $\pi_{j,0}$ ’s, we consider T consecutive applications of (4): $\pi_{i,0}$ is $O(T \cdot \epsilon)$ -close to the expectation over $\pi_{j,0}$ where j is taken from the distribution of a T -steps random walk on G starting from i . For example, if G is a constant degree expander, the random walk converges in $O(\log n)$ steps, yielding a $(1/\log n)$ -inapproximability result.

⁹Compare with our Lemma 12. The reader may also want to refer to Lemma 6 in the full version of [5].

Achieving constant hardness

Finally, in order to achieve hardness for a constant ϵ , we want a graph with constant mixing time - and this clearly cannot be done with a constant degree¹⁰. Instead, in Section A.2 we construct a normalized game whose graph has a constant mixing time, each vertex has degree $O(\sqrt{n})$, and yet approximating Nash equilibrium is hard for a constant ϵ . In short, we take n copies of the original n -player game (our new game has n^2 players). For any pair of players that play a (non-trivial) bimatrix game in the original game, we have a copy of the same bimatrix game between all $\binom{n}{2}$ pairs of their respective copies. We also add a trivial bimatrix game between every pair of players that belong to the same copy of the original game. In Section A.2 we argue that these newly added trivial edges are only a constant fraction of all edges in the new game graph, yet this graph has a constant mixing time.

References

- [1] Yakov Babichenko. Query complexity of approximate nash equilibria. In *STOC*, pages 535–544, 2014.
- [2] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061 – 1103, 2011.
- [3] Xi Chen, Decheng Dai, Ye Du, and Shang-Hua Teng. Settling the Complexity of Arrow-Debreu Equilibria in Markets with Additively Separable Utilities. In *FOCS*, pages 273–282, 2009.
- [4] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *J. ACM*, 56(3), 2009.
- [5] Xi Chen, Dimitris Paparas, and Mihalis Yannakakis. The complexity of non-monotone markets. In *STOC*, pages 181–190, 2013. Preprint of full version is available at <http://arxiv.org/abs/1211.4918v1>.

¹⁰In fact, it seems that a graph where the random walks starting from any pair of neighbors converge in constant time would suffice. We do not know whether such graphs can be constructed with constant degree.

- [6] Bruno Codenotti, Benton McCune, and Kasturi Varadarajan. Market equilibrium via the excess demand function. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*, STOC '05, pages 74–83, New York, NY, USA, 2005. ACM.
- [7] Constantinos Daskalakis. On the complexity of approximating a nash equilibrium. *ACM Transactions on Algorithms*, 9(3):23, 2013.
- [8] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *Commun. ACM*, 52(2):89–97, 2009.
- [9] Argyrios Deligkas, John Fearnley, Rahul Savani, and Paul G. Spirakis. Computing Approximate Nash Equilibria in Polymatrix Games. In *Web and Internet Economics - 10th International Conference, WINE 2014, Beijing, China, December 14-17, 2014. Proceedings*, pages 58–71, 2014.
- [10] D.K. Foley. Resource Allocation and the Public Sector. *Yale Economic Essays*, 7(1):45–98, 1967.
- [11] Paul W. Goldberg and Aaron Roth. Bounds for the query complexity of approximate equilibria. In *EC*, pages 639–656, 2014.
- [12] Sergiu Hart and Noam Nisan. The query complexity of correlated equilibria. *CoRR*, abs/1305.4874, 2013.
- [13] Sébastien Hémon, Michel de Rougemont, and Miklos Santha. Approximate Nash Equilibria for Multi-player Games. In *SAGT*, pages 267–278, 2008.
- [14] Michael D. Hirsch, Christos H. Papadimitriou, and Stephen A. Vavasis. Exponential lower bounds for finding brouwer fix points. *J. Complexity*, 5(4):379–416, 1989.
- [15] Michael Kearns. *Graphical games*, chapter 7, pages 159–180. Cambridge University Press, 2007.
- [16] Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. In *EC*, pages 36–41, 2003.
- [17] R. R. Maxfield. General equilibrium and the theory of directed graphs. *Journal of Mathematical Economics*, 27(1):23–51, 1997.

- [18] John Nash. Non-cooperative games. *The Annals of Mathematics*, 54:286–295, 1951.
- [19] Abraham Othman. *Course Match User Manual*, 1.3 edition, 2014.
- [20] Abraham Othman, Christos H. Papadimitriou, and Aviad Rubinfeld. The complexity of fairness through equilibrium. In *EC*, pages 209–226, 2014.
- [21] Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.*, 48(3):498–532, 1994.
- [22] Aviad Rubinfeld. Computational complexity of approximate nash equilibrium in large games. *CoRR*, abs/1405.0524, 2014.
- [23] Grant Schoenebeck and Salil P. Vadhan. The computational complexity of nash equilibria in concisely represented games. *TOCT*, 4(2):4, 2012.
- [24] Eran Shmaya. Brouwer Implies Nash Implies Brouwer. <http://theoryclass.wordpress.com/2012/01/05/brouwer-implies-nash-implies-brouwer/>, 2012.
- [25] Satinder P. Singh, Vishal Soni, and Michael P. Wellman. Computing approximate bayes-nash equilibria in tree-games of incomplete information. In *EC*, pages 81–90, 2004.
- [26] W. Thomson and H.R. Varian. Theories of justice based on symmetry. *Social Goals and Social Organizations: Essays in Memory of Elisha Pazner*, 1985.
- [27] H. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974.
- [28] Vijay V. Vazirani and Mihalis Yannakakis. Market equilibrium under separable, piecewise-linear, concave utilities. *J. ACM*, 58(3):10, 2011.

A Non-monotone markets: proof of inapproximability

In this section we prove our main inapproximability result for non-monotone markets:

Corollary 3 (Non-monotone markets). *Let \mathcal{U} be any non-monotone family of utility functions. There exists a constant $\epsilon_{\mathcal{U}} > 0$ such that given a market M where the utility of each trader is either linear or taken from \mathcal{U} , finding an $\epsilon_{\mathcal{U}}$ -tight approximate market equilibrium is PPAD-hard.*

A.1 Normalized polymatrix games

We identify n -player, 2-strategy polymatrix graphical games with $2n \times 2n$ matrices by letting the (i, j) -th 2×2 block correspond to the payoff matrix of player j in the bimatrix game with player i .

Given a game \mathcal{G} , let \mathbf{P}' be the $2n \times 2n$ induced payoff matrix. We normalize \mathbf{P}' as follows:

$$P_{2i,2j-1} = \begin{cases} 1/(2\Delta) + (P'_{2i,2j-1} - P'_{2i,2j}) / (2\Delta) & (i, j) \in E \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

where E is the edge set for the graphical game¹¹ and Δ is the maximum degree. We define $P_{2i,2j}$, $P_{2i-1,2j-1}$, $P_{2i-1,2j}$ analogously. Notice that \mathbf{P} and \mathbf{P}' have the same ϵ -Nash equilibrium, up to the normalization by the degree Δ . In particular finding an (ϵ/Δ) -Nash equilibrium in \mathbf{P} continues to be PPAD-complete.

Observe that in this formulation, finding an ϵ -Nash equilibrium is equivalent to finding a vector $\mathbf{x} \in [0, 1]^{2n}$ s.t. $x_{2i-1} + x_{2i} = 1$ and

$$\begin{aligned} \mathbf{x}^\top \cdot \mathbf{P}_{2i-1} > \mathbf{x}^\top \cdot \mathbf{P}_{2i} + \epsilon &\implies x_{2i} = 0 \\ \mathbf{x}^\top \cdot \mathbf{P}_{2i-1} < \mathbf{x}^\top \cdot \mathbf{P}_{2i} - \epsilon &\implies x_{2i-1} = 0 \end{aligned}$$

A.2 Games on graphs with a constant mixing time

Given the correspondence defined above between n -player games and $2n \times 2n$ matrices, we see that the structure of the game graph plays a non-trivial role

¹¹Notice that this definition allows self-loops in the game graph.

in the construction. In particular, adding trivial edges between vertices, i.e. adding zero-utility bimatrix games between players, has no effect on the utility of the players, but changes the corresponding normalized matrix. For reasons that will become clear much later in the proof, we would like our game graph to have a constant mixing time.

Indeed, a trivial candidate with very fast mixing is the complete graph. However, such a blowup in the degree would dilute our inapproximability factor in the *normalized game*. Instead, we consider n copies (v^1, \dots, v^n) of each player $v \in V$ in the original game. If players u and v play a bimatrix game $\mathcal{G}_{u,v}$ in the original game \mathcal{G} , then for every $i, j \in [n]$, we construct the same bimatrix game $\mathcal{G}_{u,v}$ between u^i and v^j . Our game graph now consists of n^2 vertices, each with degree¹² $3n$. Finally, within each copy V^i , we add trivial edges between all the vertices not otherwise connected (including self-loops). Normalize this game using (5). We use $\bar{\mathcal{G}}$ to denote the new game and \bar{G} for the new game graph; we henceforth let $\Delta = 4n - 3$ denote the degree. In the next two lemmata we show that this game satisfies the two properties we need: finding an ϵ -Nash equilibrium of $\bar{\mathcal{G}}$ is PPAD-complete, and the mixing time of \bar{G} is constant.

Lemma 7. *Given an ϵ -Nash equilibrium in $\bar{\mathcal{G}}$, we can (efficiently) construct a $(4\epsilon/3)$ -Nash equilibrium for \mathcal{G} .*

Proof. For each player v , we take the average of the mixed strategies of v^1, \dots, v^n . The utility of v is the same as the average of utilities of v^1, \dots, v^n , and if v has a $(4\epsilon/3)$ -improving deviation, then at least one of the copies v^i has an ϵ -improving deviation. (The $(4/3)$ factor comes from the change in the degree.) \square

Lemma 8. *Let $\pi_{v^i, T}$ be the distribution of a random walk on \bar{G} after T steps, starting from v^i , and let π^* be the uniform distribution on the vertices of \bar{G} . Then*

$$\|\pi_{v^i, T} - \pi^*\|_1 \leq \left(\frac{1}{4}\right)^{T/2} + \left(\frac{3}{4}\right)^{T/2} = 2^{-O(T)}$$

Proof. At each step of the random walk, there is a constant probability (greater than $3/4$) of walking on a non-trivial edge, which takes us to another (uniformly random) copy of the original game; thereafter the copy of the game remains uniformly random. Similarly, at each step there is a constant probability (greater

¹²Theorem 1 promises a graphical game of degree at most 3. It is not hard to extend to a 3-regular graph game with only a constant loss in the approximation factor.

than $1/4$) of moving to a vertex within the same copy (again, uniformly random). Thus conditioned on having walked on a non-trivial edge, and then on an edge within the same copy, the distribution is uniform. Since all vertices have the same degree, this is also the stationary distribution, and we never leave it. \square

For simplicity, in the following we redefine n to be the size of \overline{G} (and hence $\Delta \approx 4\sqrt{n}$).

A.3 Construction

Let N be a sufficiently large constant, and let $t = \log N$. Note that N depends on the parameters of the non-monotone market in \mathcal{U} , but not on the size n of our construction. We use the notation $O_N(\cdot)$ to denote the asymptotic behaviour when N goes to infinity (but arbitrarily slower than n). We prove that it is PPAD-hard to find an η -tight approximate market clearing equilibrium for $\eta = N^{-8\epsilon}$, where ϵ is the inapproximability factor from Lemma 7.

For each vertex $i \in [n]$ we construct a series of $4t + 1$ gadgets $\mathcal{R}_{i,j}$, for $j \in [0 : 4t]$. Each gadget is composed of:

Main goods $g_{2i-1,j}$ and $g_{2i,j}$ are the main goods in the reduction. They are used to encode the weights assigned to strategies x_{2i-1} and x_{2i} , respectively.

Non-monotone gadget For each $j \in [4t]$, we include additional goods $s_{i,j,3}, \dots, s_{i,j,k}$ and a non-monotone gadget

$$\text{NM}(\mu, \gamma, g_{2i-1,j}, g_{2i,j}, s_{i,j,3}, \dots, s_{i,j,k})$$

This means that we scale the non-monotone market guaranteed to exist in \mathcal{U} according to parameters γ and μ such that when all the prices are approximately the same, the excess demand of $g_{2i-1,j}$ increases linearly with its price. Formally, we have the following lemma by Chen et al.

Lemma 9. *(Lemma 3.1 of [5]) There exist two (not necessarily rational) positive constants c and d with the following property. Given any $\gamma > 0$, one can build a market M_γ with utilities from \mathcal{U} and goods $g_{2i-1,j}, g_{2i,j}, s_{i,j,3}, \dots, s_{i,j,k}$ in polynomial time in $1/\gamma$ such that:*

Let $f_{\gamma,\mu}(x)$ denote the excess demand function of $g_{2i-1,j}$ when the price of g_{2i-1} is $1+x$, and the prices of all other $k-1$ goods are $1-x$. Then $f_{\gamma,\mu}$

is well defined over $[-c, c]$ with $|f_{\gamma, \mu}(0)| \leq \mu\gamma$ and its derivative $f'_{\gamma, \mu}(0) = d > 0$. For any $x \in [-c, c]$, $f_{\gamma, \mu}(x)$ also satisfies

$$|f_{\gamma, \mu}(x) - f_{\gamma, \mu}(0) - \mu dx| \leq |\mu x/D|, \text{ where } D = \max\{20, 20/d\}.$$

Finally, we would like to set $\mu = \Delta/d$; in particular, this would imply that $f_{\gamma, \mu}(x) \approx \Delta x$. However, as mentioned above, d may be irrational. Instead, let d^* be a positive rational constant that satisfies

$$1 - 1/D \leq d^* \cdot d \leq 1$$

We set the parameters $\mu = d^* \Delta$ and $\gamma = 1/N^6$.

Price regulating gadget For $j \in [4t]$, we include a price regulating gadget

$$\mathbf{PR}(\tau, \alpha_j, g_{2i-1,j}, g_{2i,j}, s_{i,j,3}, \dots, s_{i,j,k}),$$

whereas for $j = 0$, we don't have goods $s_{i,0,3}, \dots, s_{i,0,k}$, and simply include the gadget

$$\mathbf{PR}(\tau, \alpha_0, g_{2i-1,0}, g_{2i,0}).$$

The parameters are set to $\tau = N\Delta$ and $\alpha_i = 2^i/N^5$. Notice that $\alpha_0 = N^{-5}$ and $\alpha_{4t} = N^{-1} = \beta$.

This gadget ensures that in any approximate equilibrium, the price ratio $\pi(g_{2i-1,j})/\pi(g_{2i,j})$ is always in the range $\left[\frac{1-\alpha_j}{1+\alpha_j}, \frac{1+\alpha_j}{1-\alpha_j}\right]$. Furthermore, within each gadget $\mathcal{R}_{i,j}$, the prices of all the goods besides $g_{2i-1,j}$ are exactly equal:

$$\pi(g_{2i,j}) = \pi(s_{i,j,3}) = \dots = \pi(s_{i,j,k}).$$

More specifically, we have two traders T_1 and T_2 with endowments $(k-1)\tau$ of $g_{2i-1,j}$, for T_1 and τ of each of the other goods for T_2 . The utilities are defined as

$$\begin{aligned} u_1 &= (1 + \alpha_j) x(g_{2i-1,j}) + (1 - \alpha_j) \left(x(g_{2i,j}) + \sum_{l=3}^k x(s_{i,j,l}) \right) \\ u_2 &= (1 - \alpha_j) x(g_{2i-1,j}) + (1 + \alpha_j) \left(x(g_{2i,j}) + \sum_{l=3}^k x(s_{i,j,l}) \right). \end{aligned}$$

In particular, T_1 and T_2 do not trade whenever $\pi(g_{2i-1,j})/\pi(g_{2i,j}) \in \left(\frac{1-\alpha_j}{1+\alpha_j}, \frac{1+\alpha_j}{1-\alpha_j}\right)$.

Auxiliary goods For $j = 0$, we also include an auxiliary good AUX_i . Its eventual purpose is to disentangle the price of $g_{2i-1,0}$ and $g_{2i,0}$ from the utility that the actions of player i causes to other players.

Single-minded traders graph

We connect the groups of goods ($\mathcal{R}_{i,j}$'s) using the following single-minded traders. We use $(w, g_1 : g_2)$ to denote a trader with endowment w of good g_1 who only wants good g_2 . Similarly, we use $(w, g_1, g_2 : g_3)$ to denote a trader who has an endowment w of each of g_1 and g_2 , and only wants g_3 .

1. For each $i \in [n]$ and $j \in [0 : 4t - 1]$, we add two traders from $\mathcal{R}_{i,j}$ to $\mathcal{R}_{i,j+1}$: $(\Delta, g_{2i-1,j} : g_{2i-1,j+1})$ and $(\Delta, g_{2i,j} : g_{2i,j+1})$. These traders help propagate price discrepancies from $g_{i,0}$ to $g_{i,4t}$.
2. Recall that we use g_i as short for $g_{i,0}$ and h_i for $g_{i,4t}$. For each pair $(i, j) \in E$ we add the following four traders: $(\Delta P_{2i-1,2j-1}, h_{2i-1} : g_{2j-1})$, $(\Delta P_{2i,2j-1}, h_{2i} : g_{2j-1})$, $(\Delta P_{2i-1,2j}, h_{2i-1} : g_{2j})$, $(\Delta P_{2i,2j}, h_{2i} : g_{2j})$. Since \mathbf{P} is normalized, we have

$$\Delta P_{2i-1,2j-1} + \Delta P_{2i-1,2j} = \Delta P_{2i,2j-1} + \Delta P_{2i,2j} = 1$$

These traders will enforce the approximate Nash equilibrium.

3. Connect the auxiliary goods: We let

$$r_{2j-1} = 2\Delta - \Delta \sum_{i \in \mathcal{N}(j)} (P_{2i-1,2j-1} + P_{2i,2j-1}) > 0$$

$$r_{2j} = 2\Delta - \Delta \sum_{i \in \mathcal{N}(j)} (P_{2i-1,2j} + P_{2i,2j}) > 0$$

note that $r_{2j-1} + r_{2j} = 2\Delta$.

We add the following traders: $((1 - \beta) r_{2j-1}, \text{AUX}_j : g_{2j-1})$, $((1 - \beta) r_{2j}, \text{AUX}_j : g_{2j})$, and $((1 - \beta) \Delta, g_{2j-1}, g_{2j} : \text{AUX}_j)$.

Notice that the economy graph is strongly connected (because G is strongly connected); therefore an equilibrium always exists [17]. The supplies and demands for each good are summarized in Table 1.

A.4 Structure of a market equilibrium

We now prove some properties that every η -tight approximate equilibrium π must satisfy. Recall that $\eta = N^{-8}\epsilon$, where ϵ is the inapproximability factor for the polymatrix game.

We begin with the application of the price regulating markets:

Lemma 10. *For every $i \in [n]$ and $j \in [0 : 4t]$,*

$$\frac{1 - \alpha_j}{1 + \alpha_j} \leq \frac{\pi(g_{2i-1}, j)}{\pi(g_{2i}, j)} \leq \frac{1 + \alpha_j}{1 - \alpha_j}$$

and

$$\pi(g_{2i}, j) = \pi(s_{i,j,3}) = \cdots = \pi(s_{i,j,k})$$

Proof. Follows from the construction of the price regulating markets. For more details see the proof¹³ of Lemma 6 in the full version of [5], or previous works that use similar gadgets [3, 28]. \square

We henceforth use $\pi_{i,j}$ to denote the sum of the (i, j) -th main goods: $\pi_{i,j} = \pi(g_{2i-1,j}) + \pi(g_{2i,j})$.

Lemma 11.

$$(1 - O_N(\eta)) \pi_{i,0}/2 \leq \pi(\text{AUX}_i) \leq (1 + O_N(\eta)) \pi_{i,0}/2$$

Proof. The total supply of AUX_i is $2(1 - \beta)\Delta$, yet the demand from the single-minded trader $((1 - \beta)\Delta, g_{2i-1}, g_{2i} : \text{AUX}_i)$ is $(1 - \beta)\Delta \frac{\pi_{i,0}}{\pi(\text{AUX}_i)}$. (For the upper bound we use the fact that π is a *tight* approximate market equilibrium.) \square

We are now ready to prove that the cost of every pair of main goods is approximately the same. Let $\delta = N^2\eta$.

¹³In their statement, Chen et al. require an ϵ -additively approximate equilibrium, and for a much smaller ϵ . However their proof continues to hold with our parameters.

Table 1: Goods and traders

<u>good</u> [total supply]	supplied by	demanded by
<u>g_{2i-1}, g_{2i}</u> [$N\Delta(1 + o_N(1))$]	<u>PR trader;</u> <u>$(\Delta, g_{2i-1} : g_{2i-1,1}),$</u> <u>$(\Delta, g_{2i} : g_{2i,1});$</u> <u>$((1 - \beta)\Delta, g_{2i-1}, g_{2i} : \text{AUX}_i)$</u>	<u>PR traders;</u> <u>$(\Delta P_{2j-1,2i-1}, h_{2j-1} : g_{2i-1}),$</u> <u>$\vdots$</u> <u>$(\Delta P_{2j,2i}, h_{2j} : g_{2i});$</u> <u>$((1 - \beta)r_{2i-1}, \text{AUX}_i : g_{2i-1}),$</u> <u>$((1 - \beta)r_{2i}, \text{AUX}_i : g_{2i})$</u>
<u>h_{2i-1}</u> [$(k - 1)N\Delta(1 + o_N(1))$]	<u>PR trader;</u> <u>NM traders;</u> <u>$(\Delta P_{2i-1,2j-1}, h_{2i-1} : g_{2j-1}),$</u> <u>$(\Delta P_{2i-1,2j}, h_{2i-1} : g_{2j})$</u>	<u>PR traders;</u> <u>NM traders;</u> <u>$(\Delta, g_{2i-1,4t-1} : h_{2i-1})$</u>
<u>h_{2i}</u> [$N\Delta(1 + o_N(1))$]	<u>PR trader;</u> <u>NM traders;</u> <u>$(\Delta P_{2i,2j-1}, h_{2i} : g_{2j-1}),$</u> <u>$(\Delta P_{2i,2j}, h_{2i} : g_{2j})$</u>	<u>PR traders;</u> <u>NM traders;</u> <u>$(\Delta, g_{2i,4t-1} : h_{2i})$</u>
<u>$g_{2i-1,j}$</u> [$(k - 1)N\Delta(1 + o_N(1))$]	<u>PR trader;</u> <u>NM traders;</u> <u>$(\Delta, g_{2i-1,j} : g_{2i-1,j+1})$</u>	<u>PR traders;</u> <u>NM traders;</u> <u>$(\Delta, g_{2i-1,j-1} : g_{2i-1,j})$</u>
<u>$g_{2i,j}$</u> [$N\Delta(1 + o_N(1))$]	<u>PR trader;</u> <u>NM traders;</u> <u>$(\Delta, g_{2i,j} : g_{2i,j+1})$</u>	<u>PR traders;</u> <u>NM traders;</u> <u>$(\Delta, g_{2i,j-1} : g_{2i,j})$</u>
<u>$s_{i,j,l}$</u> [$N\Delta(1 + o_N(1))$]	<u>PR trader;</u> <u>NM traders;</u>	<u>PR traders;</u> <u>NM traders;</u>
<u>AUX_i</u> [$N\Delta(1 + o_N(1))$]	<u>$((1 - \beta)r_{2i-1}, \text{AUX}_i : g_{2i-1}),$</u> <u>$((1 - \beta)r_{2i}, \text{AUX}_i : g_{2i})$</u>	<u>$((1 - \beta)\Delta, g_{2i-1}, g_{2i} : \text{AUX}_i)$</u>

Lemma 12. *Let $\pi_{\max} = \max_{i,j} \pi_{i,j}$ and $\pi_{\min} = \min_{i,j} \pi_{i,j}$, then*

$$\pi_{\max}/\pi_{\min} \leq 1 + O_N(\delta).$$

Proof. The proof of this lemma is the main obstacle which requires the tightness assumption of the market equilibrium, as well as our bound on the mixing time from Lemma 8.

Recall that by Lemma 10, the prices of all the goods in each gadget $\mathcal{R}_{i,j}$ are approximately equal. Thus, using our bound on the clearing error, we have that for each $(i, j) \in [n] \times [0 : 4t]$,

$$|\text{total spent on } \mathcal{R}_{i,j} - \text{total worth of } \mathcal{R}_{i,j}| \leq O_N(\eta \cdot kN\Delta) \pi_{i,j} = O_N(\eta \cdot N\Delta) \pi_{i,j} \quad (6)$$

By Walras' Law, the traders within each $\mathcal{R}_{i,j}$ (i.e. the price regulating and non-monotone gadgets) contribute the same to both quantities in (6). Similarly, by Lemma 11, the auxiliary traders contribute $O_N(\eta\Delta\pi_{i,0})$ -approximately the same (for $j = 0$). Therefore the money spent on $\mathcal{R}_{i,j}$ by the single minded traders is approximately the same as the total worth of endowments in $\mathcal{R}_{i,j}$ of single minded traders:

- For each $(i, j) \in [n] \times [4t]$, the restriction of (6) to the single-minded traders gives

$$|\Delta\pi_{i,j-1} - \Delta\pi_{i,j}| = O_N(\eta \cdot N\Delta) \pi_{i,j} \quad (7)$$

- Similarly, for each group $\mathcal{R}_{i,0}$, we have

$$\left| \sum_{l \in \mathcal{N}(i)} \pi_{l,4t} - \Delta\pi_{i,0} \right| = O_N(\eta \cdot N\Delta) \pi_{i,0} \quad (8)$$

Applying (7) inductively, we have that for any $i \in [n]$ and for any $j, l \in [0 : 4t]$,

$$|\pi_{i,j} - \pi_{i,l}| = O_N(\eta \cdot Nt) \pi_{i,l}$$

Combining, with (8) we have,

$$\left| \pi_{i,0} - \frac{1}{\Delta} \sum_{l \in \mathcal{N}_{\mathcal{G}}(i)} \pi_{l,0} \right| = O(\eta \cdot Nt) \pi_{i,0}$$

Thus for each i , $\pi_{i,0}$ is $O(\eta \cdot Nt)$ -approximately equal to the average of its neighbors in \overline{G} . Repeating this argument T times, we have that $\pi_{i,0}$ is $O_N(T\eta \cdot Nt)$ -approximately equal to the expectation over a T -step random walk in \overline{G} starting from i . By Lemma 8, after $T = O(\log \delta)$ steps the random walk δ -approximately converges to the uniform distribution, and we have

$$\begin{aligned} \left| \pi_{i,0} - \frac{1}{n} \sum_{l \in [n]} \pi_{l,0} \right| &= O_N(T\eta \cdot Nt) \frac{1}{n} \sum_{l \in [n]} \pi_{l,0} + \delta \max_{l \in [n]} \pi_{l,0} \\ &= O_N(\delta) \max_{l \in [n]} \pi_{l,0} \end{aligned}$$

□

Finally, we have the following lemma which describes the action of the non-monotone gadgets.

Lemma 13. (*Lemma 6 of [5]*)

$$\begin{aligned} \frac{1 + \alpha_{j-1}}{\pi(g_{2i-1,j-1})} = \frac{1 - \alpha_{j-1}}{\pi(g_{2i,j-1})} &\implies \frac{1 + \alpha_j}{\pi(g_{2i-1,j})} = \frac{1 - \alpha_j}{\pi(g_{2i,j})} \text{ and} \\ \frac{1 - \alpha_{j-1}}{\pi(g_{2i-1,j-1})} = \frac{1 + \alpha_{j-1}}{\pi(g_{2i,j-1})} &\implies \frac{1 - \alpha_j}{\pi(g_{2i-1,j})} = \frac{1 + \alpha_j}{\pi(g_{2i,j})} \end{aligned}$$

Proof. The demand for $g_{2i-1,j}$ and $g_{2i,j}$ comes from three sources: the single-minded traders, $(\Delta, g_{2i-1,j-1} : g_{2i-1,j})$ and $(\Delta, g_{2i,j-1} : g_{2i,j})$; the non-monotone gadget; and the price regulating gadget. Assume without loss of generality that the first premise holds, i.e. $\frac{1+\alpha_{j-1}}{\pi(g_{2i-1,j-1})} = \frac{1-\alpha_{j-1}}{\pi(g_{2i,j-1})}$. When the prices of $g_{2i-1,j}$ and $g_{2i,j}$ are equal, the demand from $(\Delta, g_{2i-1,j-1} : g_{2i-1,j})$ is larger since she has more income from $g_{2i-1,j-1}$. In order to account for this difference, $\pi(g_{2i-1,j})$ must be higher - but then the demand from the traders in the non-monotone market increases. Thus we further have to increase $\pi(g_{2i-1,j})$, until we reach the threshold of the price regulating traders: $(1 + \alpha_j) / (1 - \alpha_j)$.

Formally, normalize π such that $\pi_{i,j} = \pi(g_{2i-1,j}) + \pi(g_{2i,j}) = 2$. Thus by Lemma 12, $\pi_{i,j-1}$ is also $O_N(\delta)$ -close to 2. Let $f(x)$ denote the excess demand from the traders in the non-monotone gadget when $\pi(g_{2i-1,j}) = 1 + x$ and $\pi(g_{2i-1,j}) = \pi(s_{i,j,3}) = \dots = \pi(s_{i,j,k}) = 1 - x$ (recall from Lemma 10 that the latter prices are always equal to each other). By Lemma 9, we have that

$|f(0)| \leq \mu\gamma$, and for all $x \in [-c, c]$:

$$|f(x) - f(0) - \mu dx| \leq |\mu x/D|.$$

Now, let $\pi(g_{2i-1,j-1}) = 1 + y$; notice that by Lemma 12, $y = \alpha_{j-1} \pm O_N(\delta)$. Let $h(x, y)$ excess demand from all traders besides the two that belong to the price regulating gadget. Then,

$$h(x, y) = f(x) + \frac{\Delta(1+y)}{1+x} - \Delta = f(x) - \frac{\Delta x}{1+x} + \frac{\Delta y}{1+x}$$

For small x , we show that $f(x) \approx \Delta x/(1+x)$. More precisely,

$$\begin{aligned} |f(x) - \Delta x/(1+x)| &\leq |f(x) - \mu dx| + |\mu dx - \Delta x| + \Delta |x - x/(1+x)| \\ &\leq |f(0)| + 2|\mu x/D| + 2\Delta x^2 \\ &\leq \Delta \cdot (d^* \cdot \gamma + 2x/20 + 2x^2) \\ &\leq \frac{\Delta y}{3} \end{aligned}$$

where the first inequality follows from the triangle inequality; the second follows by application of Lemma 9 for the first difference and the definitions of μ and d^* for the second; the third inequality applies the Lemma 9 again; finally the last inequality holds because for sufficiently large constant N , the parameters γ and x are sufficiently small.

Therefore, the excess demand must be balanced by the demand from the price regulating traders, implying that indeed $\frac{1+\alpha_j}{\pi(g_{2i-1,j})} = \frac{1-\alpha_j}{\pi(g_{2i,j})}$. \square

A.5 From market equilibrium to Nash equilibrium

To complete the proof of Corollary 3, we must construct an ϵ -Nash equilibrium from any η -tight approximate market equilibrium.

For each $i \in [n]$, let $\theta_i = (\pi(h_{2i-1}) + \pi(h_{2i}))/2$. We define

$$x_{2i-1} = \frac{\pi(h_{2i-1})/\theta_i - (1-\beta)}{2\beta} \text{ and } x_{2i} = \frac{\pi(h_{2i})/\theta_i - (1-\beta)}{2\beta} \quad (9)$$

Observe that $x_{2i-1} + x_{2i} = 1$.

Suppose that

$$\mathbf{x}^\top \cdot \mathbf{P}_1 \geq \mathbf{x}^\top \cdot \mathbf{P}_2 + \epsilon$$

We show that this forces $x_1 = 1$ and $x_2 = 0$; by the discussion in Section A.1 this implies that \mathbf{x} is indeed an ϵ -Nash equilibrium.

The following traders spend money on g_1 :

1. For each $i \in \mathcal{N}_{\overline{G}}(1)$, there is a $(\Delta P_{i,1}, h_i; g_1)$ trader. The total money these traders spend on g_1 is

$$\sum \Delta P_{i,1} \cdot \pi(h_1) = \sum \Delta P_{i,1} (1 - \beta + 2\beta \cdot x_i) \theta_{\lceil i/2 \rceil}$$

2. For each $i \in \mathcal{N}_{\overline{G}}(2)$, there is a $(\Delta P_{i,2}, h_i; g_2)$ trader. The total money these traders spend on g_2 is

$$\sum \Delta P_{i,2} \cdot \pi(h_2) = \sum \Delta P_{i,2} (1 - \beta + 2\beta \cdot x_i) \theta_{\lceil i/2 \rceil}$$

3. $((1 - \beta) r_1, \text{AUX}_1 : g_1)$ and $((1 - \beta) r_2, \text{AUX}_1 : g_2)$ traders

Let M_1 be the total amount that these traders spend on g_1 . Then

$$M_1 = \sum_{i \in \mathcal{N}_{\overline{G}}(1)} \Delta P_{i,1} (1 - \beta + 2\beta \cdot x_i) \theta_{\lceil i/2 \rceil} + (1 - \beta) r_1 \pi(\text{AUX}_1)$$

Normalizing the prices such that $\frac{1}{n} \sum_{i=1}^n \theta_i = 1/\Delta$, this means that

$$M_1 \geq 2(1 - \beta) + 2\beta \mathbf{x}^\top \cdot \mathbf{P}_1 - O_N(\delta)$$

Similarly,

$$M_2 \leq 2(1 - \beta) + 2\beta \mathbf{x}^\top \cdot \mathbf{P}_2 + O_N(\delta)$$

Therefore,

$$M_1 \geq M_2 + 2\beta\epsilon - O_N(\delta) = M_2 + \Theta_N(\beta\epsilon)$$

so the difference between the demands for g_1 and g_2 from these traders is

$$\frac{M_1}{\pi(g_1)} - \frac{M_2}{\pi(g_2)} \geq \frac{M_2 + \Theta_N(\beta\epsilon)}{\pi(g_1)} - \frac{M_2(1 + \alpha_0)}{\pi(g_1)(1 - \alpha_0)} = \Theta_N(\beta\epsilon)$$

Thus the price regulating traders T_1 and T_2 must have different demands for g_1 and g_2 - but this can only happen when

$$\frac{1 + \alpha_0}{\pi(g_1)} = \frac{1 - \alpha_0}{\pi(g_2)}$$

Therefore, by consecutive applications of Lemma 13,

$$\frac{1 + \beta}{\pi(h_1)} = \frac{1 - \beta}{\pi(h_2)}$$

Finally, by (9) this implies that $x_1 = 1$ and $x_2 = 0$.

□