Constant Inapproximability for Fisher Markets

ARGYRIOS DELIGKAS, Royal Holloway, United Kingdom
JOHN FEARNLEY, University of Liverpool, United Kingdom
ALEXANDROS HOLLENDER, University of Oxford, United Kingdom
THEMISTOKLIS MELISSOURGOS, University of Essex, United Kingdom

We study the problem of computing approximate market equilibria in Fisher markets with separable piecewise-linear concave (SPLC) utility functions. In this setting, the problem was only known to be PPAD-complete for inverse-polynomial approximations. We strengthen this result by showing PPAD-hardness for constant approximations. This means that the problem does not admit a polynomial time approximation scheme (PTAS) unless PPAD = P. In fact, we prove that computing any approximation better than 1/11 is PPAD-complete. As a direct byproduct of our main result, we get the same inapproximability bound for Arrow-Debreu exchange markets with SPLC utility functions.

CCS Concepts: • **Theory of computation → Problems, reductions and completeness**; **Market equilibria**; **Exact and approximate computation of equilibria**.

Additional Key Words and Phrases: Fisher markets, competitive equilibrium, PPAD

ACM Reference Format:
Argyrios Deligkas, John Fearnley, Alexandros Hollender, and Themistoklis Melissourgos. 2024. Constant Inapproximability for Fisher Markets. In *Conference on Economics and Computation (EC '24), July 8–11, 2024, New Haven, CT, USA*. ACM, New York, NY, USA, 27 pages. https://doi.org/10.1145/3670865.3673533

1 Introduction

Fisher markets [Brainard and Scarf, 2005] are one of the foundational models that have shaped modern economics. In a Fisher market, every buyer has a fixed budget that they spend on their most favorable bundle of goods according to their utility function. The appealing fact for this type of markets is that when the utilities of the buyers satisfy some *"standard sufficient conditions"*, then a *market equilibrium* is *guaranteed* to exist. In a market equilibrium, the prices and the allocation of the goods to the buyers are such that: (a) every buyer is allocated goods that maximize their utility, and (b) the market *clears*, i.e., the supply of each good is exactly equal to the demand of that good.

When the utility functions of the buyers are linear, a market equilibrium can always be computed in polynomial time [Devanur et al., 2008, Orlin, 2010, Végh, 2012]. However, the problem becomes intractable as soon as one considers a very slight generalization of linear utilities: *additive separable piecewise linear concave* (SPLC) utilities. A buyer with an SPLC utility function has a piecewise linear concave utility for each good, and their utility for a bundle of goods is simply the sum of their utilities for each of the individual goods.

Authors' Contact Information: Argyrios Deligkas, Royal Holloway, Egham, United Kingdom, argyrios.deligkas@rhul.ac.uk; John Fearnley, University of Liverpool, Liverpool, United Kingdom, john.fearnley@liverpool.ac.uk; Alexandros Hollender, University of Oxford, Oxford, United Kingdom, alexandros.hollender@cs.ox.ac.uk; Themistoklis Melissourgos, University of Essex, Colchester, United Kingdom, themistoklis.melissourgos@essex.ac.uk.

EC '24, July 8–11, 2024, New Haven, CT, USA

ACM ISBN 979-8-4007-0704-9/24/07
https://doi.org/10.1145/3670865.3673533

Computing a market equilibrium in Fisher markets with SPLC utilities was shown to be a PPAD-complete problem [Chen and Teng, 2009, Vazirani and Yannakakis, 2011], which means that the problem is unlikely to admit a polynomial-time algorithm unless PPAD = P. This motivates the study of *approximate equilibria* in which the condition that the market clears is replaced with an *approximate clearing* constraint. In an ε-approximate market equilibrium, we first normalize the market so that there is exactly one unit of each good, and then we seek a price vector such that every good ε-clears, meaning that the discrepancy between between the supply and the demand for any good is at most ε.

Vazirani and Yannakakis [2011] actually showed that it is PPAD-complete to find an ε-approximate market equilibrium when ε is inversely polynomial in the size of the market, i.e., the number of buyers and the number of goods. This ruled out fully polynomial-time approximation schemes (FPTAS) for the problem unless PPAD = P. However, the existence of a polynomial-time approximation scheme (PTAS) was not ruled out by this result, and a PTAS may well be good enough to clear a market for most practical purposes.

Our contribution. This paper provides the first PPAD-completeness for ε-approximate market equilibria in Fisher markets for a *constant* ε. Hence, a PTAS for Fisher markets cannot exist unless PPAD = P.

Theorem 1.1. *It is* PPAD*-complete to compute an ε-approximate market equilibrium in Fisher markets with SPLC utilities for any constant $\varepsilon < 1/11$.*

We note that our hardness threshold, $1/11$, is relatively large. If we recall that all goods are normalized to have one unit available, then this result states that it is PPAD-complete to find any approximate market equilibrium in which all of the goods clear to within 9% of their total supply. So this essentially rules out the existence of polynomial-time algorithms that can clear a market in practice.

Technical Overview. We prove our result via a reduction from the Pure-Circuit problem that was recently introduced by Deligkas et al. [2022]. In this problem we are given a circuit with NOT, NAND, and PURIFY gates, and we are required to find a satisfying assignment using the values $\{0, 1, \bot\}$. The NOT gate behaves as usual on the values in $\{0, 1\}$, and has no constraints when the input is $\bot$. The NAND gate is required to output 0 if both inputs are 1, and is required to output 1 if at least one input is 0, and otherwise has no constraints. The PURIFY gate has one input and two outputs. If the input is in $\{0, 1\}$, then that value should be copied to both outputs, while if the input is $\bot$, then at least one output should be in $\{0, 1\}$.

The use of the Pure-Circuit problem as a starting point for our reduction enables us to give a construction that is arguably considerably simpler than previous hardness results for Fisher markets. Prior work has used constructions in which one first sets up a price-regulating market, which essentially gives a set of goods that are able to encode values within a certain range, and then adds extra buyers to simulate a PPAD-hard problem, e.g., finding a Nash equilibrium of a bimatrix game.

In contrast, the market that we construct essentially directly implements the gates of the Pure-Circuit instance. For each variable in the Pure-Circuit we introduce a good, and the price of this good encodes the value of the variable. Specifically we define a *low price* L and a *high price* H with $L < H$, and then a price below L encodes a 0, a price above H encodes a 1, while a price between the two encodes $\bot$.

Then we show that the NOT and NAND gates can be implemented by relatively simple gadgets that each contain two buyers. The PURIFY gate is implemented using two chains of NOT gates, where we carefully tweak the parameters of each of the NOT gates to ensure that the constraints of the PURIFY gate are implemented.

Arrow-Debreu exchange markets. Finally, we also show that our hardness result for Fisher markets implies a new hardness result for Arrow-Debreu exchange markets with SPLC utilities. Specifically, we use a well-known reduction from Fisher markets to exchange markets that preserves hardness for ε-equilibria, and we use this to show that computing an ε-equilibrium in an exchange market is also PPAD-hard for all $\varepsilon < 1/11$.

While hardness for Arrow-Debreu exchange markets was already known for constant ε [Rubinstein, 2018], prior work had only shown this for some extremely small unknown constant, and so that work did not rule out the existence of a polynomial-time algorithm that could clear a market in practice. Our hardness result effectively rules this out unless PPAD=P.

1.1 Related work

The problem of computing market equilibria, for both exact and approximate equilibria, has received significant attention over the years. For Fisher markets, Vazirani and Yannakakis [2011] and Chen and Teng [2009] have established PPAD-hardness for SPLC utilities albeit for a sub-constant ε. On the other hand, polynomial-time algorithms were derived for the cases where the utility functions of the buyers are linear [Devanur et al., 2008, Orlin, 2010, Végh, 2012], homogeneous [Eisenberg, 1961], or weak gross substitutes [Codenotti et al., 2005].

Furthermore, when the number of goods is constant Kakade et al. [2004] gave a PTAS while Devanur and Kannan [2008] gave a polynomial-time algorithm for exact equilibria. In fact, the algorithm of [Devanur and Kannan, 2008] also works when the number of buyers is constant in the SPLC utility setting. For non-separable PLC utilities Garg et al. [2022] derived a fixed parameter approximation scheme that has the number of buyers as a parameter.

Matching markets is another important subclass of general Fisher markets. Alaei et al. [2017] designed a polynomial-time algorithm for markets with a constant number of goods or buyers, while Vazirani and Yannakakis [2020] derived a polynomial algorithm when the buyers have dichotomous utilities. For one-sided matching markets, the most famous problem is the Hylland-Zeckhauser market, for which the existence of an equilibrium was initially established in [Hylland and Zeckhauser, 1979] and was recently simplified by Braverman [2021]. Even more recently, Chen et al. [2022] have established PPAD-completeness for the problem.

There has also been interest in Fisher markets with additional constraints. Birnbaum et al. [2010], Devanur [2004], and Vazirani [2010] considered the case where the utilities of the buyers depend on prices of goods through spending constraints. Jalota et al. [2023] considered additional linear constraints that include matching markets, and they gave a tâtonnement process which was found to converge to a market equilibrium in experiments.

The Arrow-Debreu *exchange* market model [Arrow and Debreu, 1954] is another foundational class of markets. In this setting, the goods are brought to the market by the buyers, who then spend the revenue they get from selling their initial endowments. PPAD-hardness for Arrow-Debreu market equilibria has been established for several different settings [Chen et al., 2009, 2017, Codenotti et al., 2006, Garg et al., 2023] and in fact Rubinstein [2018] showed that it is PPAD-hard to compute an ε-equilibrium in Arrow-Debreu markets for a small unknown constant ε. On the positive side, there are several polynomial-time algorithms for linear utilities [Duan et al., 2016, Duan and Mehlhorn, 2015, Garg and Végh, 2019, Jain, 2007, Ye, 2008], or for buyers with weak gross substitutes utilities [Bei et al., 2019, Codenotti et al., 2005, Garg et al., 2023]. The various types of market models have also been studied for the setting where the items are chores that provide disutility to the agents [Boodaghians et al., 2022, Brânzei and Sandomirskiy, 2024, Chaudhury et al., 2022a,b].

The Pure-Circuit problem was recently introduced in [Deligkas et al., 2022], where it was used to prove strong, improved, PPAD-hardness results for a variety of problems related mainly to approximate Nash equilibria. Since then it was further used in [Deligkas et al., 2023] to prove tight PPAD-hardness for approximate Nash equilibria in graphical games, and in [Ioannidis et al., 2023] to prove improved stronger PPAD-hardness in the problem of clearing financial networks. To the best of knowledge, this is the first time that Pure-Circuit has been used to prove hardness for market equilibria.

2 Preliminaries

2.1 Fisher Markets

Fisher markets. A Fisher market is given by a tuple $(G, B, (e_i)_{i\in B}, (u_i)_{i\in B})$, where:

- G is a set of (divisible) goods. Without loss of generality, we assume that there is one unit of each good available.[1]
- B is a set of buyers.
- For every $i \in B$, $e_i > 0$ is the budget of buyer i.
- For every $i \in B$, $u_i : \mathbb{R}^{|G|}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is the utility function of buyer i. For any allocation $x_i \in \mathbb{R}^{|G|}_{\geq 0}$ of goods to buyer i (where $x_{i,j} \geq 0$ denotes the amount of good j allocated to buyer i), $u_i(x_i)$ denotes the utility derived by the buyer. We assume that the utility functions are separable piecewise-linear concave (SPLC), meaning that $u_i(x_i)$ can be written as $\sum_{j\in G} u_{i,j}(x_{i,j})$, where each $u_{i,j} : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ satisfies
  (1) $u_{i,j}(0) = 0$,
  (2) $u_{i,j}$ is continuous and piecewise-linear,
  (3) $u_{i,j}$ is concave but non-decreasing.

Optimal bundles. Given a price vector $p \in \mathbb{R}^{|G|}_{\geq 0}$, where p_j denotes the price of good j, the set of optimal bundles for buyer i, denoted $\mathrm{OPT}_i(p) \subseteq \mathbb{R}^{|G|}_{\geq 0}$, is the set of optimal solutions of the following optimization problem:

$$\begin{aligned} \max \quad & u_i(x_i) \\ \text{s.t.} \quad & \sum_{j\in G} p_j x_{i,j} \leq e_i \\ & x_{i,j} \geq 0 \quad \forall j \in G. \end{aligned} \tag{1}$$

Note that it is possible that $\mathrm{OPT}_i(p) = \emptyset$, if some good has price 0, and the agent is never satiated with this good.

Competitive equilibrium. For any $\varepsilon \geq 0$, an ε-approximate market equilibrium is a price vector p and an allocation vector $x = (x_i)_{i\in B}$ satisfying the following conditions:

(1) For each buyer i, x_i is an optimal bundle at prices p, i.e., $x_i \in \mathrm{OPT}_i(p)$.
(2) For each good j, the market clears approximately up to ε units of good, i.e.,

$$\left| \sum_{i\in B} x_{i,j} - 1 \right| \leq \varepsilon.$$

When $\varepsilon = 0$, this corresponds to an exact market equilibrium.

[1] This can be achieved by a simple normalization, and it simplifies the expression for the clearing constraint below.

u	v
0	1
1	0
$\bot$	$\{0, 1, \bot\}$

NOT gate

u	v	w
1	1	0
0	$\{0, 1, \bot\}$	1
$\{0, 1, \bot\}$	0	1
Else		$\{0, 1, \bot\}$

NAND gate

u	v	w
0	0	0
1	1	1
$\bot$	At least one output in $\{0, 1\}$	

PURIFY gate

Fig. 1. The truth tables of the three gates of Pure-Circuit.

Existence of equilibria. The following condition is sufficient to guarantee the existence of a market equilibrium [Maxfield, 1997, Vazirani and Yannakakis, 2011]:

Sufficient Condition: For every buyer $i \in B$, there exists a good $j \in G$ such that $u_{i,j}$ is a strictly increasing function (i.e., buyer i is never satiated with good j).

Computational problem. Let $\varepsilon \geq 0$. The computational problem of computing an ε-approximate market equilibrium is defined as follows:

Input: A Fisher market $(G, B, (e_i)_{i \in B}, (u_i)_{i \in B})$ with SPLC utilities satisfying the sufficient condition for the existence of equilibria. For each $i \in B$ and $j \in G$, $u_{i,j}$ is explicitly described in the input, i.e., for each linear affine piece we are given the positions and values at its endpoints.
Output: An ε-approximate market equilibrium (p, x).

Given (p, x), the equilibrium conditions can be verified in polynomial time, because, for SPLC utilities, the optimization problem (1) defining $\text{OPT}_i(p)$ can be solved in polynomial time using a simple greedy approach;[2] see, e.g., [Garg et al., 2015]. Together with the existence of solutions guaranteed by the sufficient condition, this puts the problem in the complexity class TFNP of total NP search problems. Prior work [Vazirani and Yannakakis, 2011] has shown that the problem lies in the subclass PPAD of TFNP, even for $\varepsilon = 0$. In particular, exact rational solutions are guaranteed to exist. The problem is known to be PPAD-complete for $\varepsilon = 0$, and also when ε is part of the input and inverse-polynomial with respect to the description of the market [Chen and Teng, 2009, Vazirani and Yannakakis, 2011]. No hardness result is known for any constant $\varepsilon > 0$.

2.2 The Pure-Circuit Problem

The Pure-Circuit problem. We will show our hardness result by reducing from the Pure-Circuit problem, which is known to be PPAD-complete [Deligkas et al., 2022]. An instance of the Pure-Circuit problem is given by a node set $V = [n]$ and a set C of gate-constraints (or just *gates*). Each gate $g \in C$ is of the form $g = (T, u, v, w)$ where $u, v, w \in V$ are distinct nodes, and $T \in \{\text{NOT}, \text{NAND}, \text{PURIFY}\}$ is the type of the gate, with the following interpretation.

- If $T = \text{NOT}$, then u is the input of the gate, and v is its output. (w is unused)
- If $T = \text{NAND}$, then u and v are the inputs of the gate, and w is its output.
- If $T = \text{PURIFY}$, then u is the input of the gate, and v and w are its outputs.

We require that each node is the output of exactly one gate.

[2] In fact, given only prices p, it is possible to check in polynomial time whether there exists an allocation x such that (p, x) is an ε-approximate market equilibrium [Vazirani and Yannakakis, 2011]. So we could have equivalently defined the computational problem to only seek equilibrium prices p.

A solution to instance (V, C) is an assignment $\mathbf{x} : V \longrightarrow \{0, 1, \bot\}$ that satisfies all the gates (see Fig. 1), i.e., for each gate $g = (T, u, v, w) \in C$ we have the following.

- If $T =$ NOT in $g = (T, u, v)$, then $\mathbf{x}$ satisfies

$$\mathbf{x}[u] = 0 \implies \mathbf{x}[v] = 1$$
$$\mathbf{x}[u] = 1 \implies \mathbf{x}[v] = 0.$$

- If $T =$ NAND in $g = (T, u, v, w)$, then $\mathbf{x}$ satisfies

$$\mathbf{x}[u] = \mathbf{x}[v] = 1 \implies \mathbf{x}[w] = 0$$
$$(\mathbf{x}[u] = 0) \vee (\mathbf{x}[v] = 0) \implies \mathbf{x}[w] = 1$$

- If $T =$ PURIFY, then $\mathbf{x}$ satisfies

$$\{\mathbf{x}[v], \mathbf{x}[w]\} \cap \{0, 1\} \neq \emptyset$$
$$\mathbf{x}[u] \in \{0, 1\} \implies \mathbf{x}[v] = \mathbf{x}[w] = \mathbf{x}[u].$$

The structure of a Pure-Circuit instance is captured by its *interaction graph*. This graph is constructed on the vertex set $V = [n]$ by adding a directed edge from node u to node v whenever v is the output of a gate with input u. The total degree of a node is the sum of its in- and out-degrees.

Theorem 2.1 ([Deligkas et al., 2022]). Pure-Circuit *is* PPAD*-complete, even when every node of the interaction graph has in- and out-degree at most 2 and total degree at most 3.*

3 Construction of the Market

Given an instance (V, C) of the Pure-Circuit problem, we will construct, in polynomial time, a Fisher market that simulates that instance. In this section, we describe how the Fisher market will be constructed.

The reference good. To implement our construction, we need a specific good that, in every equilibrium, has a price that is close to 1. We call this good the *reference good*, and denote it as ref.

To ensure that the price of the reference good is close to 1, we use a *reference buyer* called b_{ref} who has a budget of $e_{b_{\text{ref}}} = 1$, and the following utility function.

$$u_{b_{\text{ref}},j}(x) = \begin{cases} x & \text{if } j = \text{ref}, \\ 0 & \text{otherwise.} \end{cases}$$

In other words, the reference buyer desires only the reference good, and will therefore spend all of their money on it.

We will ensure that the total amount of demand on ref from all other buyers in the construction will be significantly smaller than 1. This will ensure that in any approximate equilibrium, the price of ref will be close to 1.

Variable encodings. For each variable in the Pure-Circuit instance, we introduce a good that will encode that variable. The value assigned to each variable will be determined by the price of the corresponding good.

Let $s, a \in \mathbb{R}$ be two positive constants that will be fixed later. Given a price vector p, we define a *high price* $H = s \cdot p_{\text{ref}}$, and a *low price* $L = s \cdot H/a = s^2 \cdot p_{\text{ref}}/a$. The idea is that we will fix $s < 1$, while a will be chosen to be very large. Thus, the high price is a specified fraction of the reference price, while the low price is very close to zero.

Given a price vector p for the Fisher market, we extract an assignment to the variables of the Pure-Circuit instance in the following way.

- If $p_v \geq H$ then $v = 1$.

- If $p_v \leq L$ then $v = 0$.
- Otherwise $v = \bot$.

Note that the values of H and L depend on the price of the reference good. Although we know that the price of the reference good will be close to 1, it will not be exactly 1, and thus H and L will vary according to the particular price that is chosen for the reference good. We will use H_{low} and H_{high} to denote a lower and upper bound for H, and we will give exact values for these bounds later. We likewise use L_{low} and L_{high} to give bounds for L.

Auxiliary buyers. The construction will use many *auxiliary buyers*, whose purpose is to buy a pre-specified amount of a particular good. Given a good $j \in G$ and an amount $r \in [0, 1]$ we define the buyer $b = \text{aux}(j, r)$ in the following way.

- The buyer's budget is $e_b = r \cdot H_{\text{high}}$.
- The buyer's utility for good j is defined to be

$$u_{b,j}(x) = \begin{cases} 2 \cdot s \cdot x & \text{if } x \leq r, \\ 2 \cdot s \cdot r & \text{otherwise.} \end{cases}$$

- The buyer's utility for ref is $u_{b,\text{ref}}(x) = x$.
- The buyer's utility for all other goods is zero.

We will ensure that the price of good j is no larger than H_{high}, which means that the auxiliary buyer always has enough money to buy r units of good j. The utilities have been chosen to ensure that, if the good's price is no larger than H_{high}, the auxiliary buyer always strictly prefers to buy r units of good j before buying the reference good. Moreover, once r units of good j have been bought by the auxiliary buyer, their marginal utility for good j becomes 0, whereas they always have positive marginal utility towards the reference good. Thus, as we will later show formally, in any approximate equilibrium the auxiliary buyer will buy exactly r units of good j, and spend the rest of their money on the reference good.

The interface between variables. As mentioned previously, for each variable in the Pure-Circuit instance, there is a corresponding good whose price encodes the value of that variable. To simulate the gates of the Pure-Circuit instance, we will use buyers that buy a certain proportion of the variable-encoding goods.

To ensure that there is a consistent interface between the gates, we introduce a parameter $t \in (0, 0.5)$ with the intention that, if a gate g uses variable v as an input variable, then the buyers that implement g will buy t units of the good that represents v.

From Theorem 2.1 we have that the Pure-Circuit instance has out-degree at most two, so for each variable-encoding good, we expect at most $2t$ units of the good to be bought by the gates that take this variable as an input. Note however that some goods may not be used as an input to exactly two gates. To address this, for each variable-encoding good j that is used as an input to only one gate, we introduce a buyer $\text{aux}(j, t)$, who will buy t units of that good to top-up the amount that is bought. In the case where j is used as an input to no other gate, we introduce two such buyers instead.

In a perfect world, this would ensure that exactly $2t$ units of each variable-encoding good are bought by the gates for which that good is an input. Unfortunately, as we will later show formally, the gates may actually buy slightly less than t units each, and this will be dealt with in our proofs.

NOT gates. We begin by building a NOT gate. The NOT gate is particularly important because we will also use it to implement the PURIFY gate. For this reason, we implement a parameterized gadget $\mathsf{NOT}(\mathsf{in}, \mathsf{out}, r)$, where in is the input good to the NOT gate, out is the output good, and

$r \in [0, 1]$ is a constant. When the gadget is used to implement a NOT gate in the Pure-Circuit instance, we will set $r = r_{\text{NOT}}$, where $r_{\text{NOT}} \in [0, 1]$ is a fixed parameter.

The gadget will have a buyer named the inverter, and also an auxiliary buyer. The interaction between the buyers and the goods is shown below, where an arrow indicates that the buyer has a non-zero utility function for that good.

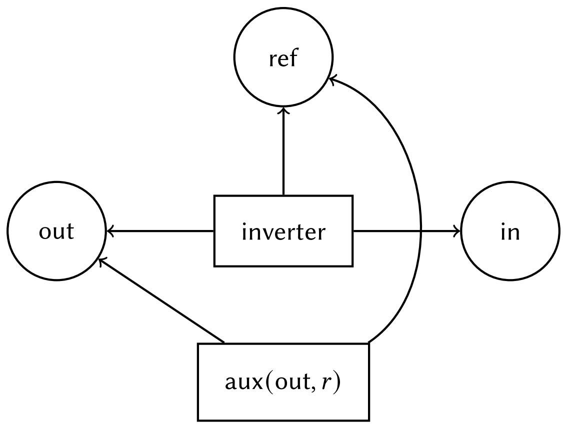

The inverter is specified as follows.

- The budget of the inverter is $e_{\text{inverter}} = t \cdot H_{\text{low}}$.
- The inverter's utility for in is

$$u_{\text{inverter,in}}(x) = \begin{cases} a \cdot x & \text{if } x \leq t, \\ a \cdot t & \text{otherwise.} \end{cases}$$

- The inverter's utility for out is $u_{\text{inverter,out}}(x) = s \cdot x$.
- The inverter's utility for ref is $u_{\text{inverter,ref}}(x) = x$.
- The inverter's utility for all other goods is zero.

Recall that we intend to set $s < 1$ and a to be very large. Thus, the inverter buyer is heavily incentivized to buy t units of the input good before buying anything else. However, once t units of the input good have been bought, the inverter no longer has any interest in buying more units. We call t the *anti-endowment* for the inverter: that buyer must[3] buy a specific amount of the good before being able to spend money elsewhere.

The high-level idea is that if the input variable is 1, then $p_{\text{in}} \geq H \geq H_{\text{low}}$. Since we have fixed the budget of the inverter to be $e_{\text{inverter}} = t \cdot H_{\text{low}}$, this means that the inverter will spend all of their budget buying their anti-endowment, and will have no money left to spend on the output good. Thus, the demand on the output good will be low, and so its price cannot be high.

On the other hand, if the input variable is 0, then $p_{\text{in}} \leq L = s \cdot H/a$. Since we will set a to be very large, this means that the inverter spends almost no money buying their anti-endowment, and so has essentially their entire budget left over. In this scenario, the inverter will have a large amount of money to spend on the output good, which will cause it to have a high price.

The utility of the inverter towards the reference good serves to ensure that the output good's price cannot rise above H. Since $H = s \cdot p_{\text{ref}}$, if $p_{\text{out}} > H$ then the marginal utility of buying the output good would be strictly less than $s/(s \cdot p_{\text{ref}}) = 1/p_{\text{ref}}$, whereas the marginal utility of buying the reference good is $1/p_{\text{ref}}$, and so in this scenario the inverter would spend all of their remaining

[3]Actually, if the price of the output good is very close to zero, then the inverter may prefer to buy the output good before the input good. We will deal with this case separately in our proofs.

money on the reference good, and no money on the output good. We will ensure that the output good fails to clear in this scenario, which will ensure that the price of the output good is capped at H in any approximate equilibrium.

The auxiliary buyer allows us to change how much of the output good can be bought by the inverter, since we can adjust the r parameter to change how much of the output is taken by the auxiliary buyer. The high-level idea here is that by changing r, we can change how the output good's price changes relative to the inverter's remaining budget after buying the anti-endowment. This will be used critically when we use NOT gates to implement the PURIFY gate.

NAND gates. For each NAND gate with inputs in1 and in2, and output out, we use the following construction.

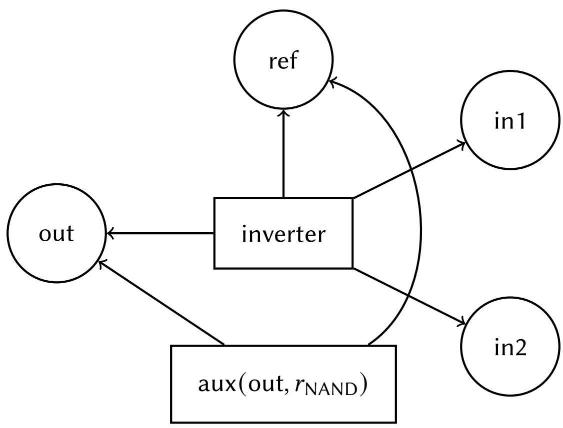

The construction uses two buyers. The auxiliary buyer has a parameter r_{NAND} that will be fixed later. The inverter buyer is slightly different than that of the NOT gate and is specified as follows.

- The budget of the inverter is $e_{\text{inverter}} = 2t \cdot H_{\text{low}}$.
- The inverter's utility for good $j \in \{\text{in1}, \text{in2}\}$ is

$$u_{\text{inverter},j}(x) = \begin{cases} a \cdot x & \text{if } x \leq t, \\ a \cdot t & \text{otherwise.} \end{cases}$$

- The inverter's utility for out is $u_{\text{inverter,out}}(x) = s \cdot x$.
- The inverter's utility for ref is $u_{\text{inverter,ref}}(x) = x$.
- The inverter's utility for all other goods is zero.

This is a straightforward generalization of the NOT gate to two inputs. The inverter now has a budget of $2t \cdot H_{\text{low}}$ and an anti-endowment of t units for both of the input goods. So if both inputs have a high price, the inverter will spend all their money on the input goods, and so they will not be able to increase the price of the output good. If either of the two inputs has a low price, then the inverter will have money left over, and will be able to push the price of the output good higher.

PURIFY gates. For a PURIFY gate with input in and outputs out1 and out2, we use the following construction.

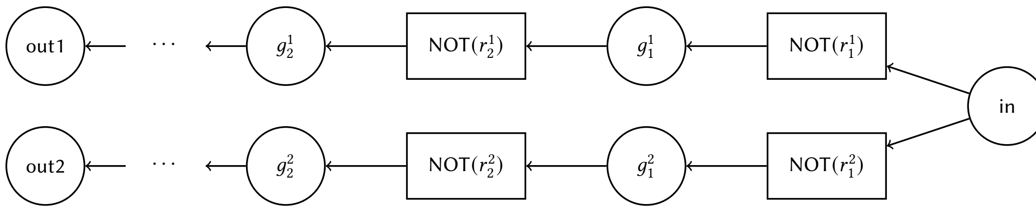

We use two chains of d NOT gates to compute the two outputs, where d is a parameter. To do this, we introduce intermediate goods $\{g_j^1, g_j^2 \ : \ 1 \leq j \leq d-1\}$. To simplify the definition, we use $g_0^1 = g_0^2 = \mathsf{in}$, and we use $g_d^1 = \mathsf{out1}$ and $g_d^2 = \mathsf{out2}$. Then for each $i \in \{1,2\}$, and for each $j \in \{1,2,\ldots,d\}$ we include a gadget $\mathsf{NOT}(g_{j-1}^i,\ g_j^i,\ r_j^i)$ where $r_j^i \in \mathbb{R}$ is a parameter that we will fix later.

We treat each of the NOT gates and each of the intermediate goods as full gates and variables in our instance. This means that each intermediate good j will also have an auxiliary buyer $\mathsf{aux}(j,t)$ that buys t units of the good, to compensate for the fact that the good is only used as an input to one other gate.

The idea is that we will set d to be some large even number. Therefore each chain of NOT gates will have even length. So if the input variable is a 1 or 0, both chains of NOT gates will output that value, as required by the PURIFY constraints. If the input good has a price that is strictly between L and H, meaning that it encodes a $\bot$ value, then one of the two chains is required to output a value in $\{0,1\}$. We will ensure this by carefully selecting values for the parameters r_j^i.

More specifically, the parameters are chosen so that there exists a price p_{in}^* such that, if the price of the input good is at least p_{in}^*, then the prices of the goods in the top chain will increase and we will have that $p_{\mathsf{out1}} = H$. On the other hand, if p_{in} is less than p_{in}^*, then the prices in the bottom chain will decrease and we will have that $p_{\mathsf{out2}} \leq L$. This implies that at least one of the two chains will always output a pure value no matter what the input price is, as required by the PURIFY gate.

Circuit copies. So far we have described a full reduction from Pure-Circuit to a Fisher market, and while this construction is strong enough to give hardness for a constant ε, the bound that we would obtain would be much smaller than $1/11$. The main reason for this is the uncertainty in the price of the reference good. As we mentioned earlier, this price should be close to 1, but as ε increases, our bounds on it get weaker.

To address this, we will introduce k copies of the circuit, where each copy is only required to work when the reference price is within a particular range. Specifically, letting $H_{\min} := s/2$ and $H_{\max} := 2s$, we divide the region $[H_{\min}, H_{\max}]$ into k equally sized non-overlapping regions. Then for each such region $[x,y]$, we build a copy of the circuit setting $H_{\text{low}} = x$ and $H_{\text{high}} = y$. As we will show, the price of the reference good in any ε-equilibrium will be bounded by $p_{\mathsf{ref}}^- := 1/2$ and $p_{\mathsf{ref}}^+ := 2$, and so we will indeed always have that $H \in [H_{\min}, H_{\max}]$.

When we are given an approximate equilibrium of the Fisher market, we first find an interval $[x,y]$ that contains H, and then decode the assignments to the Pure-Circuit instance from the copy that corresponds to that interval, while ignoring all other copies.

The key advantage of this is that each circuit copy can now assume H_{low} and H_{high} are very close together, which then increases the values of ε for which we can show hardness.

The sufficient condition. Finally, we will verify that the sufficient condition for the existence of an equilibrium holds for our construction. Recall that this condition requires that for every buyer i, there exists a good j such that $u_{i,j}$ is a strictly increasing function. There are three types of buyers in our market: inverters, auxiliary buyers, and the reference buyer. All of these buyers have the same utility for the reference good, and so for every buyer i in the market we have $u_{i,\text{ref}}(x) = x$. Therefore the sufficient condition is satisfied.

4 Analysis

Fix any $\varepsilon < 1/11$. The construction described in the previous section uses several parameters. For our proofs, we will fix these parameters to the following values. We first set $\delta := \frac{11}{4} \cdot (1/11 - \varepsilon) > 0$.

Parameter	Value	Description
t	$4/11$	The size of each inverter's anti-endowment
d	$4 \cdot \lceil \ln(4/\delta) \rceil$	The length of the NOT gate chains in a PURIFY gate
k	$110/\delta$	The number of copies of the circuit
s	$1/(20kd\|V\|)$	The inverter's marginal utility for the output good
a	$\max\{2, 4s/\delta\}$	The inverter's marginal utility for the input good
r_{NOT}	$2/11$	The value of r used in the NOT gate
r_{NAND}	$2/11$	The value of r used in the NAND gate
r^1_j, j is odd	0	The r values used in the first NOT chain in a PURIFY gate
r^1_j, j is even	$2/11$	
r^2_j, j is odd	$2/11$	The r values used in the second NOT chain in a PURIFY gate
r^2_j, j is even	0	

For the rest of this section, we will consider an ε-approximate market equilibrium (p, x) of the market. Omitted proofs can be found in the appendix.

General properties of the construction. Before we prove the correctness of each of the individual gates, we first prove some general properties of the construction that will be useful later. We start by considering the reference good. Recall that the reference good was intended to have price close to 1. The following lemma shows this this is indeed the case.

LEMMA 4.1. *We have $p_{\text{ref}} \in [p^-_{\text{ref}}, p^+_{\text{ref}}]$, where $p^-_{\text{ref}} = 1/2$ and $p^+_{\text{ref}} = 2$. In particular, it follows that $H = s \cdot p_{\text{ref}} \in [H_{min}, H_{max}]$, where $H_{min} = s/2$ and $H_{max} = 2s$.*

PROOF. If $p_{\text{ref}} = 0$, then, by construction, the reference buyer will demand an infinite amount of good ref. In particular, $\text{OPT}_{b_{\text{ref}}}(p) = \emptyset$, and we cannot be at an ε-equilibrium. Thus, we must have $p_{\text{ref}} > 0$. In that case, any optimal bundle for buyer b_{ref} will demand exactly $e_{b_{\text{ref}}}/p_{\text{ref}} = 1/p_{\text{ref}}$ units of the reference good. If $p_{\text{ref}} < p^-_{\text{ref}} = 1/2$, then buyer b_{ref} demands $1/p_{\text{ref}} > 2 > 1 + \varepsilon$ units of good. Since there is only one unit of good available, this is a contradiction to the ε-clearing condition. Thus, we must have $p_{\text{ref}} \geq p^-_{\text{ref}}$.

Let $E_{-b_{\text{ref}}}$ denote the sum of budgets of all buyers, except the reference buyer b_{ref}, i.e., $E_{-b_{\text{ref}}} = \sum_{i \in B \setminus \{b_{\text{ref}}\}} e_i$. By construction, the budget of any buyer $i \neq b_{\text{ref}}$ satisfies $e_i \leq H_{\text{max}}$. Indeed, the budget of any such buyer is either $t \cdot H_{\text{low}}$, or $2t \cdot H_{\text{low}}$, or $r \cdot H_{\text{high}}$ for some $r \in [0, 1]$. Furthermore, we have $|B \setminus \{b_{\text{ref}}\}| \leq 4kd|V|$, since there are k copies and in each copy there are at most $d|V|$ goods, and for each such good there are at most four buyers having non-zero utility for it. Thus, we can bound $E_{-b_{\text{ref}}} \leq 4kd|V|H_{\text{max}}$.

We can now proceed to prove the upper bound on the price p_{ref}. The total demand on the reference good is at most $(e_{b_{\text{ref}}} + E_{-b_{\text{ref}}})/p_{\text{ref}}$. This corresponds to the case where all buyers spend all of their budget on the reference good. In order for the reference good to ε-clear, we must thus have that $(e_{b_{\text{ref}}} + E_{-b_{\text{ref}}})/p_{\text{ref}} \geq 1 - \varepsilon$, i.e.,

$$p_{\text{ref}} \leq \frac{e_{b_{\text{ref}}} + E_{-b_{\text{ref}}}}{1 - \varepsilon} = \frac{1}{1 - \varepsilon} + \frac{E_{-b_{\text{ref}}}}{1 - \varepsilon} \leq \frac{1}{1 - \varepsilon} + 5kd|V|H_{\text{max}} \leq \frac{1}{1 - \varepsilon} + 10kd|V|s \leq 2 = p^+_{\text{ref}}$$

where we used $\varepsilon < 1/11$, $E_{-b_{\text{ref}}} \leq 4kd|V|H_{\text{max}}$, $H_{\text{max}} = 2s$, and $s \leq 1/(20kd|V|)$. □

Recall that we have k copies of the circuit, and that each of the circuits is required to work only when p_{ref} is within the given range for that circuit. From now on for the rest of this section we will focus only on the circuit copy that was built for the particular value of p_{ref} in our equilibrium. Observe that by construction we have $H_{\text{low}} \leq H \leq H_{\text{high}}$ in this copy. We can also apply Lemma 4.1 to obtain the following bounds, which will prove useful later. We have

$$H_{\text{high}} - H_{\text{low}} \leq 2s/k \tag{2}$$

and thus

$$1 - 4/k \leq H_{\text{low}}/H_{\text{high}} \leq 1 \tag{3}$$

where we used $H_{\text{high}} \geq H_{\text{min}} = s/2$.

The next lemma gives, for each NOT or NAND gate, an upper bound on the amount of the output good that can be allocated to buyers other than the inverter which uses that good as an output.

Lemma 4.2. *For any good $j \in G \setminus \{\text{ref}\}$, let i be the inverter that implements the gate using j as an output variable. If the gate is a* NOT *gate with parameter r, then the total amount of good j allocated to buyers other than i is at most $2t + r$, while if the gate is a* NAND *gate, then the total amount j allocated to buyers other than i is at most $2t + r_{\text{NAND}}$.*

Proof. The only buyers other than buyer i who have non-zero utility toward good j are:

- The auxiliary buyer from the NOT or NAND gate that outputs to good j. A NAND gate auxiliary will buy at most r_{NAND} units of good j, while an auxiliary for a NOT gate with parameter r will buy at most r units of good j.
- The inverters from gates that take good j as an input, who each can buy at most t units of good j.
- Auxiliary buyers that buy at most t units of good j whenever that good is not used as an input to exactly two other gates.

These buyers cannot buy more than the amount specified above because once they have been allocated their specified amount, their marginal utility for good j becomes 0, whereas by the sufficiency condition, all buyers have at least one good with positive marginal utility for which they are never satiated. The same is true for all buyers that have utility 0 for good j.

Hence, for a NOT gate with parameter r, the total amount demanded by buyers other than buyer i is at most $r + 2t$ while for a NAND gate total amount demanded by buyers other than buyer i is at most $r_{\text{NAND}} + 2t$. □

The next lemma states that for each NOT or NAND gate, the inverter must be allocated strictly more than 0 units of the output good.

Lemma 4.3. *For any good $j \in G \setminus \{\text{ref}\}$, let i be the inverter that implements the* NOT *or* NAND *gate that has j as its output variable. Then, buyer i is allocated strictly more than 0 units of good j.*

Next we prove that price of every good other than the reference good can have price at most H.

Lemma 4.4. *For any good $j \in G \setminus \{\text{ref}\}$, $0 < p_j \leq H$.*

The next lemma states that the auxiliary players always do their jobs correctly, meaning that they buy exactly as much of the target good as we have specified.

Lemma 4.5. *For any auxiliary buyer $i \in B$ with target good $j \in G \setminus \{\text{ref}\}$ and mandated amount $r \in [0, 1]$, the buyer is allocated exactly r units of good j, i.e., $x_{i,j} = r$.*

4.1 Bounds on anti-endowment purchases

Recall that our intention is that for each variable-encoding good, the gadgets that take that good as an input should purchase t units of the good as an anti-endowment. As we mentioned earlier, this is unfortunately not the case, and it is possible that less that t units are purchased by each gadget. In this section we formally prove bounds on how much each gadget purchases.

The following lemma shows that if the output good has price strictly greater than L, then the inverter buyer must purchase their full anti-endowment before buying any other good. In particular, since Lemma 4.3 requires the inverter to spend non-zero money on the output good, this means that the inverter must buy t units of all input goods.

Lemma 4.6. *Let j be an inverter in any* NOT *or* NAND *gate. If $p_{\mathsf{out}} > L$ then*

- *if the gate is a* NOT *gate then j must buy t units of* in *before buying any other good; and*
- *if the gate is a* NAND *gate then j must buy t units of* in1 *and t units of* in2 *before buying any other good.*

Proof. If the gate is a NOT gate, then the inverter's marginal utility for in is $a/p_{\mathsf{in}} \geq a/H$, where the second inequality comes from Lemma 4.4. On the other hand, their marginal utility for out is $s/p_{\mathsf{out}} < s/L = a/H$, and their marginal utility for ref is $1/p_{\mathsf{ref}} = s/H$. Since $s < a$, the inverter will buy up to t units of in before buying any other good.

For the case where the gate is a NAND gate, we have that the inverter's marginal utility for in1 is $a/p_{\mathsf{in1}} \geq a/H$, and the inverter's marginal utility for in2 is $a/p_{\mathsf{in2}} \geq a/H$, where in both cases the second inequality comes from Lemma 4.4. We can now use the same argument as above to conclude that the inverter must buy t units of both in1 and in2 before buying any other good. □

The only remaining case is when $p_{\mathsf{out}} \leq L$. In this case the inverter may actually purchase the output good before buying any of the input goods. However, since L is very close to zero, the maximum amount of money that the inverter can spend on the output good without violating the ε-clearing constraint is also very small. This may cause the inverter to buy slightly less than their full anti-endowments. The following pair of lemmas give formal bounds on this for NOT gates and NAND gates, respectively.

Lemma 4.7. *For any* NOT *gate, the inverter's allocation of the input good,* in*, satisfies $x_{\mathsf{inverter,in}} \in [t - 3/k, t]$.*

Lemma 4.8. *For any* NAND *gate, the inverter buyer's allocations of the input goods,* in1 *and* in2*, satisfy $x_{\mathsf{inverter,in1}}, x_{\mathsf{inverter,in2}} \in [t - 5/k, t]$.*

We do not need to treat the PURIFY gates separately, since each PURIFY gate is constructed entirely out of NOT gates. Combining the previous two lemmas gives the following bound.

Lemma 4.9. *For any gate, the total amount of the output good that is allocated to buyers that are not part of the gate's construction lies in $[2\bar{t}, 2t]$, where $\bar{t} := t - 5/k$.*

4.2 Correctness of the gates

We now show that the gate constructions correctly simulate the gates of a Pure-Circuit. In this section we will also use the notation $L_{\mathrm{high}} := s \cdot H_{\mathrm{high}}/a$. Since $H \in [H_{\mathrm{low}}, H_{\mathrm{high}}]$ and $L = s \cdot H/a$, we thus also have $L \leq L_{\mathrm{high}}$.

NOT gates. The constraints of a NOT gate in a Pure-Circuit only require the gate to work when the input is either 0 or 1. So to prove that the NOT gate works, it is sufficient to consider the cases where $p_{\mathsf{in}} \leq L$ and $p_{\mathsf{in}} \geq H$. However, since we use NOT gates to implement PURIFY gates, and since PURIFY gates must output a pure value even when the input is $\perp$, we will need to understand the relationship between the price of the output good and the price of the input good even when $L < p_{\text{in}} < H$. We do this in the following pair of lemmas, which give upper and lower bounds on p_{out} with respect to p_{in}.

LEMMA 4.10. *For any* NOT(in, out, r) *gate with* $r \in [0,1]$ *satisfying* $1 - 2t - r - \varepsilon > 0$, *we have*

$$p_{\text{out}} \leq \max\left(L,\ (H_{low} - p_{\text{in}}) \cdot \frac{t}{1 - 2t - r - \varepsilon}\right).$$

LEMMA 4.11. *For any* NOT(in, out, r) *gate with* $r \in [0,1]$ *satisfying* $1 - 2t - r - \varepsilon > 0$, *we have*

$$p_{\text{out}} \geq \min\left(H,\ (H_{low} - p_{\text{in}}) \cdot \frac{t}{1 - 2\bar{t} - r + \varepsilon}\right).$$

We can now prove that the NOT gate works for pure values by applying the previous two lemmas.

LEMMA 4.12. *For each* NOT *gate with input* in *and output* out, *we have the following.*

- *If* $p_{\text{in}} \geq H$ *then* $p_{\text{out}} \leq L$.
- *If* $p_{\text{in}} \leq L$ *then* $p_{\text{out}} \geq H$.

PROOF. For the first claim we can apply Lemma 4.10 to obtain

$$\begin{aligned} p_{\text{out}} &\leq \max\left(L,\ (H_{\text{low}} - p_{\text{in}}) \cdot \frac{t}{(1 - 2t - r_{\text{NOT}} - \varepsilon)}\right) \\ &\leq \max\left(L,\ (H_{\text{low}} - H) \cdot \frac{t}{(1 - 2t - r_{\text{NOT}} - \varepsilon)}\right) \\ &\leq \max(L, 0) \\ &= L, \end{aligned}$$

where the second inequality uses the fact that $p_{\text{in}} \geq H$, and the third inequality used the fact that $H_{\text{low}} - H < 0$, that $t > 0$, and that $1 - \varepsilon - 2t - r_{\text{NOT}} = 1/11 - \varepsilon > 0$.

For the second claim we can apply Lemma 4.11 to obtain

$$\begin{aligned} p_{\text{out}} &\geq \min\left(H,\ (H_{\text{low}} - p_{\text{in}}) \cdot \frac{t}{1 - 2\bar{t} - r_{\text{NOT}} + \varepsilon}\right) \\ &\geq \min\left(H,\ (H_{\text{low}} - L_{\text{high}}) \cdot \frac{t}{1 - 2\bar{t} - r_{\text{NOT}} + \varepsilon}\right) \end{aligned}$$

where the second inequality uses the fact that $p_{\text{in}} \leq L \leq L_{\text{high}}$.

$$\begin{aligned} (H_{\text{low}} - L_{\text{high}}) \cdot \frac{t}{1 - 2\bar{t} - r_{\text{NOT}} + \varepsilon} &\geq (H - 2s/k - s \cdot H_{\text{high}}/a) \cdot \frac{t}{1 - 2\bar{t} - r_{\text{NOT}} + \varepsilon} \\ &\geq (H - 2s/k - s \cdot (H + 2s/k)/a) \cdot \frac{t}{1 - 2\bar{t} - r_{\text{NOT}} + \varepsilon} \\ &\geq (H - H/k - s \cdot (H + H/k)/a) \cdot \frac{t}{1 - 2\bar{t} - r_{\text{NOT}} + \varepsilon} \\ &\geq (1 - 2/k - 2/k^2) \cdot H \cdot \frac{t}{1 - 2\bar{t} - r_{\text{NOT}} + \varepsilon}. \end{aligned}$$

The second inequality uses the fact that $H_{\text{low}} \geq H - 2s/k$ which arises from Equation (2) and that $L_{\text{high}} = s \cdot H_{\text{high}}/a$ by definition. The third inequality uses the fact that $H_{\text{high}} \leq H + 2s/k$ which again arises from Equation (2). The fourth inequality uses the fact that $H = s \cdot p_{\text{ref}} \leq 2s$ from Lemma 4.1. The fifth inequality uses the fact that $s \leq 1/(20kd|V|) \leq 1/k$ and $a \geq 1$. Since $k \geq 3$, we have that $(1 - 2/k - 2/k^2) > 0$, and so we can continue with the following chain of inequalities on the multiplier of H

$$\begin{aligned}\frac{t \cdot (1 - 2/k - 2/k^2)}{1 - 2\bar{t} - r_{\mathrm{NOT}} + \varepsilon} &\geq \frac{4/11 \cdot (1 - 2/k - 2/k^2)}{1 - 8/11 + 5/k - r_{\mathrm{NOT}} + \varepsilon} \\ &\geq \frac{4/11 \cdot (1 - 2/k - 2/k^2)}{2/11 + 5/k} \\ &\geq 1.\end{aligned}$$

The first inequality used the fact that $\bar{t} \geq t - 5/k$ from Lemma 4.9, the fact that $t = 4/11$ by definition, the second inequality uses the fact that $r_{\mathrm{NOT}} = 2/11$, and the final inequality uses the fact that $k \geq 32$. So we have shown that $p_{\mathrm{out}} \geq \max(H, H)$, and therefore we can conclude that $p_{\mathrm{out}} \geq H$. □

NAND gates. For NAND gates we simply need to verify that the gate works for pure inputs. The following pair of lemmas shows that the NAND gate construction correctly implements the constraints of the NAND gate.

LEMMA 4.13. *For each* NAND *gate with inputs* in1 *and* in2 *and output* out*, if* $p_{\mathrm{in1}} \geq H$ *and* $p_{\mathrm{in2}} \geq H$*, then* $p_{\mathrm{out}} \leq L$*.*

LEMMA 4.14. *For each* NAND *gate with inputs* in1 *and* in2 *and output* out*, if there exists an input good* $j \in \{\mathrm{in1}, \mathrm{in2}\}$ *such that* $p_j \leq L$*, then* $p_{\mathrm{out}} \geq H$*.*

PROOF. Assume, for the sake of contradiction, that $p_{\mathrm{out}} < H$. We have that the marginal utility of buying the reference good is $1/p_{\mathrm{ref}} = s/H$, whereas the marginal utility of the output good is $s/p_{\mathrm{out}} > s/H$. Thus the inverter strictly prefers the output good to the reference good. Since the inverter can never be satiated by the output good, this means that the inverter cannot buy the reference good.

Furthermore, the inverter cannot buy more than t units of either of the input goods, because the inverter's marginal utility becomes 0 once t units have been bought. This means that the inverter can spend at most $t \cdot p_{\mathrm{in1}}$ money on in1 and $t \cdot p_{\mathrm{in2}}$ on in2. The inverter's remaining budget will therefore be at least

$$\begin{aligned}e_{\mathrm{inverter}} - t \cdot p_{\mathrm{in1}} - t \cdot p_{\mathrm{in2}} &= t \cdot (2H_{\mathrm{low}} - p_{\mathrm{in1}} - p_{\mathrm{in2}}) \\ &\geq t \cdot (2H_{\mathrm{low}} - H - L)\end{aligned}$$

money left over after buying all goods other than out, where the final inequality has used the assumption that one of the two input goods j satisfies $p_j \leq L$ while the other has price at most H due to Lemma 4.4. This money must therefore be spent on out. Therefore, the number of units of the output good bought by the inverter will be at least

$$\begin{aligned}t \cdot (2H_{\mathrm{low}} - H - L)/p_{\mathrm{out}} &\geq t \cdot (2H_{\mathrm{low}} - H - L)/H \\ &\geq t \cdot (2H_{\mathrm{low}} - H_{\mathrm{high}} - L_{\mathrm{high}})/H_{\mathrm{high}} \\ &\geq t \cdot (2H_{\mathrm{low}} - H_{\mathrm{high}} - s \cdot H_{\mathrm{high}})/H_{\mathrm{high}} \\ &\geq t \cdot (2H_{\mathrm{low}} - 1.05H_{\mathrm{high}})/H_{\mathrm{high}} \\ &\geq t \cdot (2H_{\mathrm{low}}/H_{\mathrm{high}} - 1.05) \\ &\geq t \cdot (2 \cdot (1 - 4/k) - 1.05) \\ &> 0.75t \\ &= 3/11.\end{aligned}$$

The first inequality uses the fact that $p_{\text{out}} \leq H$ while the second inequality uses the fact that $H \leq H_{\text{high}}$ and $L \leq L_{\text{high}}$. The third inequality uses the fact that $L_{\text{high}} = s \cdot H_{\text{high}}/a$ by definition, and since $a \geq 1$ we therefore have $L_{\text{high}} \leq s \cdot H_{\text{high}}$. The fourth inequality uses the fact that $s \leq 1/(20kd|V|) \leq 1/20$. The sixth inequality uses fact that $H_{\text{low}}/H_{\text{high}} \geq 1 - 4/k$ from Equation Equation (3), and the final inequality uses the fact that $k > 40$ and the fact that $t = 4/11$.

By Lemma 4.9 we have that buyers who are external to the NAND gate will buy at least $2\bar{t}$ units of the output good, and by Lemma 4.5 we have that r_{NAND} units of the output good will be bought by the auxiliary buyer of the NAND gate. Since the good must ε-clear, this means that the number of units of the output good that the inverter can buy is at most

$$\begin{aligned} 1 - 2\bar{t} - r_{\text{NAND}} + \varepsilon &< 1 - 8/11 + 10/k - 2/11 + 1/11 \\ &= 2/11 + 10/k. \\ &< 3/11. \end{aligned}$$

where the first inequality uses the fact that $\bar{t} \geq t - 5/k$ from Lemma 4.9, that $r_{\text{NAND}} = 2/11$, and that $\varepsilon < 1/11$, while the second inequality uses the fact that $k > 110$.

So we have shown that the inverter must buy strictly more than 3/11 units of the output good, but also can buy strictly less than 3/11 units of the output good, so we have arrived at a contradiction. □

PURIFY gates. We now prove the correctness of the PURIFY gates. For the proofs of this section, we will use the following auxiliary notation. Let $A := \frac{1-2t-r-\varepsilon}{t}$, $A' := \frac{1-2t-r'-\varepsilon}{t}$, $B := \frac{1-2\bar{t}-r+\varepsilon}{t}$, $B' := \frac{1-2\bar{t}-r'+\varepsilon}{t}$, where $A, A', B, \in (0,1)$, and $0 < B' = 1 - 4/k - s/a < 1$.

In the following two lemmas and corollary, we prove a relationship between the input price to a chain of NOT gates and the output price.

Lemma 4.15. *Consider a chain of gates* NOT(in, mid, r), NOT(mid, out, r'). *For any $P \in [L, H]$, if $p_{\text{in}} \geq H_{low} \cdot (1 - A) + P \cdot AB'$ then $p_{\text{out}} \geq P$.*

Lemma 4.16. *Consider a chain of gates* NOT(in, mid, r), NOT(mid, out, r'). *For any $P \in [L, H]$, if $p_{\text{in}} \leq H_{low} (1 - B) + P \cdot A'B$ then $p_{\text{out}} \leq P$.*

Corollary 4.17. *For some even $d \geq 2$, consider a chain made by connecting $d/2$ pairs of gates* NOT(g_{j-1}, g_j, r), NOT(g_j, g_{j+1}, r')*, where $j \in \{1, 3, 5, \ldots, d-1\}$. Also, let* in $:= g_0$*, and* out $:= g_d$ *be the input and output of the chain, respectively.*

- *If $p_{\text{in}} \geq H_{low} \cdot \frac{1-A}{1-AB'} + (AB')^{d/2} \cdot \left(H - H_{low} \cdot \frac{1-A}{1-AB'}\right)$, then $p_{\text{out}} \geq H$, and*
- *if $p_{\text{in}} \leq H_{low} \cdot \frac{1-B}{1-A'B} + (A'B)^{d/2} \cdot \left(L - H_{low} \cdot \frac{1-B}{1-A'B}\right)$, then $p_{\text{out}} \leq L$.*

Proof. By construction, the last pair's input serves as the output of the previous pair, and so on, until the second pair's input serves as the output of the first pair. Suppose now that we require the output of the chain to be at least $P \in [L, H]$. Then by repeatedly applying the bound of Lemma 4.15, i.e., substituting P with $H_{\text{low}} \cdot (1 - A) + P \cdot AB'$ for i times, and finally setting $P = H$, we get the first part of the statement.

To prove the first part of the statement, let the lower bound for the input of the $(d/2 - i + 1)$-st pair of NOT gates be P_i, for $i \in \{1, \ldots, d/2\}$. Now set $P_0 = H$, and observe that the first time we substitute (i.e., we consider only the last pair of gates), the lower bound for p_{in} becomes $P_1 := H_{\text{low}} \cdot (1 - A) + P_0 \cdot AB'$ which agrees with the first part of the statement for $d = 2$. We will use this as a base case. Now, given that for some $i \in \{1, \ldots, d/2 - 1\}$ it is $P_i := H_{\text{low}} \cdot \frac{1-A}{1-AB'} + (AB')^i \cdot \left(P_0 - H_{\text{low}} \cdot \frac{1-A}{1-AB'}\right)$, we will prove that if $P_{i+1} = H_{\text{low}} \cdot (1 - A) + P_i \cdot AB'$, then $P_{i+1} = H_{\text{low}} \cdot \frac{1-A}{1-AB'} + (AB')^{i+1} \cdot \left(P_0 - H_{\text{low}} \cdot \frac{1-A}{1-AB'}\right)$.

Indeed, by this substitution we get

$$\begin{aligned}
P_{i+1} &= H_{\text{low}} \cdot (1-A) + \left(H_{\text{low}} \cdot \frac{1-A}{1-AB'} + (AB')^i \cdot \left(P_0 - H_{\text{low}} \cdot \frac{1-A}{1-AB'}\right)\right) \cdot AB' \\
&= H_{\text{low}} \cdot (1-A)\left(1 + \frac{AB'}{1-AB'}\right) + (AB')^i \cdot \left(P_0 - H_{\text{low}} \cdot \frac{1-A}{1-AB'}\right) \cdot AB' \\
&= H_{\text{low}} \cdot \frac{1-A}{1-AB'} + (AB')^{i+1} \cdot \left(P_0 - H_{\text{low}} \cdot \frac{1-A}{1-AB'}\right),
\end{aligned}$$

therefore, by induction on i until $i = d/2$, we complete the first statement of the lemma, for $P_0 = H$.
The second part of the proof is symmetric, using Lemma 4.16 and $P_0 = L$, and we omit it. □

We now prove that the PURIFY gate is correct. Specifically, we consider a PURIFY gate with an input good in and two output goods out1, out2. It consists of two chains of NOT gates; one, called *chain 1*, with input in and output out1, and the other, called *chain 2*, with input in and output out2. Each chain contains $d/2$ pairs of NOT gates, where $d \geq 4\lceil \ln(4/\delta) \rceil$ is an even number. Each chain $i \in \{1, 2\}$ has pairs of gates $\mathsf{NOT}(g^i_{j-1}, g^i_j, r^i_j)$, $\mathsf{NOT}(g^i_j, g^i_{j+1}, r^i_{j+1})$, where the w-th pair, $w \in \{1, \dots, d/2\}$ corresponds to the aforementioned j-th and $(j+1)$-st NOT gates, with $j = 2w - 1$. Also, the first good of both chains is common, i.e., $g^1_0 = g^2_0$. Finally, chain 1 has $\mathsf{in} := g^1_0$, $\mathsf{out1} := g^1_d$, $r^1_j = 0$ for j odd, and $r^1_j = 2/11$ for j even; chain 2 has $\mathsf{in} := g^2_0$, $\mathsf{out2} := g^2_d$, $r^2_j = 2/11$ for j odd, and $r^2_j = 0$ for j even.

Lemma 4.18. *For each* PURIFY *gate we have the following.*

- *If $p_{\mathsf{in}} \geq H$, then $p_{\mathsf{out1}} \geq H$ and $p_{\mathsf{out2}} \geq H$.*
- *If $p_{\mathsf{in}} \leq L$, then $p_{\mathsf{out1}} \leq L$ and $p_{\mathsf{out2}} \leq L$.*
- *If $p_{\mathsf{in}} \in (L, H)$, then at least one of $p_{\mathsf{out1}}, p_{\mathsf{out2}}$ is outside of (L, H).*

5 Inapproximability for Arrow-Debreu Exchange Markets

In this section we show that our inapproximability result also applies to Arrow-Debreu exchange markets.

Exchange markets. An exchange market is given by a tuple $\left(G, B, (w_{i,j})_{i\in B, j\in G}, (u_i)_{i\in B}\right)$, where:

- G is a set of (divisible) goods.
- B is a set of buyers (or traders).
- For every $i \in B$ and $j \in G$, $w_{i,j} \geq 0$ is the endowment of good j owned by buyer i. Without loss of generality, we assume that there is one unit of each good available, i.e., $\sum_{i\in B} w_{i,j} = 1$ for all $j \in G$.
- For every $i \in B$, $u_i : \mathbb{R}^{|G|}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is an SPLC utility function, as defined in the preliminaries.

Optimal bundles. Given a price vector $p \in \mathbb{R}^{|G|}_{\geq 0}$, the set of optimal bundles for buyer i, denoted $\mathrm{OPT}_i(p) \subseteq \mathbb{R}^{|G|}_{\geq 0}$, is the set of optimal solutions of the following optimization problem:

$$\begin{aligned}
\max \quad & u_i(x_i) \\
\text{s.t.} \quad & \sum_{j\in G} p_j x_{i,j} \leq \sum_{j\in G} p_j w_{i,j} \\
& x_{i,j} \geq 0 \quad \forall j \in G.
\end{aligned}$$

In other words, the budget of the buyer is the amount of money obtained by selling its endowment.

Existence of equilibria. The definition of market equilibrium is identical to the one given for Fisher markets in the preliminaries.[4] The following condition is sufficient to guarantee the existence of a market equilibrium for exchange markets [Maxfield, 1997, Vazirani and Yannakakis, 2011]:

Sufficient Condition: The economy graph of the market is strongly connected. This graph is defined on the set of buyers B by introducing a directed edge from buyer i to buyer i' if there exists a good $j \in G$ such that $w_{i,j} > 0$ and $u_{i',j}$ is a strictly increasing function.

Theorem 5.1. *It is* PPAD*-complete to compute an ε-approximate market equilibrium in Arrow-Debreu exchange markets with SPLC utilities for any constant $\varepsilon < 1/11$.*

Proof. To prove this result we use a simple folklore reduction from Fisher markets to exchange markets. Fix any $\varepsilon < 1/11$ and let $(G, B, (e_i)_{i\in B}, (u_i)_{i\in B})$ denote a Fisher market that satisfies the sufficient condition for Fisher markets. We construct the corresponding exchange market $(G, B, (w_{i,j})_{i\in B, j\in G}, (u_i)_{i\in B})$, where the endowments are given by

$$w_{i,j} := \frac{e_i}{\sum_{k\in B} e_k}$$

for all $i \in B$ and $j \in G$. Note that have $\sum_{i\in B} w_{i,j} = 1$ for all $j \in G$, as desired. Furthermore, it is easy to see that the exchange market satisfies the sufficient condition, because we have $w_{i,j} > 0$ for all i, j and moreover, by the sufficient condition for the Fisher market, for any $i \in B$ there exists a good $j \in G$ such that $u_{i,j}$ is a strictly increasing function. Thus, the economy graph is strongly connected.

Now consider any ε-approximate market equilibrium (p, x) of the exchange market. It is easy to see that the prices are invariant to scaling, so without loss of generality we can assume that

$$\sum_{j\in G} p_j = \sum_{i\in B} e_i.$$

As a result the budget available to buyer i in the exchange market at prices p is

$$\sum_{j\in G} p_j w_{i,j} = \sum_{j\in G} p_j \frac{e_i}{\sum_{k\in B} e_k} = e_i \frac{\sum_{j\in G} p_j}{\sum_{k\in B} e_k} = e_i.$$

But this is exactly the budget of buyer i in the Fisher market, so it follows that bundle x_i is also optimal for buyer i in the Fisher market. Finally, since $\sum_{i\in B} w_{i,j} = 1$ for all $j \in G$, the ε-clearing constraint in the exchange market and Fisher market are identical. It follows that (p, x) is also an ε-approximate market equilibrium for the Fisher market. Thus, finding an ε-approximate market equilibrium (p, x) of the exchange market is PPAD-hard. Furthermore, membership in PPAD is known from [Vazirani and Yannakakis, 2011]. □

Acknowledgments

Argyrios Deligkas was supported by EPSRC Grant EP/X039862/1 "NAfANE: New Approaches for Approximate Nash Equilibria". John Fearnley was supported by EPSRC Grant EP/W014750/1 "New Techniques for Resolving Boundary Problems in Total Search".

References

Saeed Alaei, Pooya Jalaly Khalilabadi, and Eva Tardos. 2017. Computing equilibrium in matching markets. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. 245–261.

Kenneth J. Arrow and Gerard Debreu. 1954. Existence of an Equilibrium for a Competitive Economy. *Econometrica* 22, 3 (1954), 265–290. https://doi.org/10.2307/1907353

[4]This is sometimes called an ε-tight market equilibrium; see the discussion in [Chen et al., 2017].

Xiaohui Bei, Jugal Garg, and Martin Hoefer. 2019. Ascending-price algorithms for unknown markets. *ACM Transactions on Algorithms (TALG)* 15, 3 (2019), 1–33.

Benjamin Birnbaum, Nikhil R Devanur, and Lin Xiao. 2010. New convex programs and distributed algorithms for Fisher markets with linear and spending constraint utilities. *Unpublished manuscript* (2010).

Shant Boodaghians, Bhaskar Ray Chaudhury, and Ruta Mehta. 2022. Polynomial Time Algorithms to Find an Approximate Competitive Equilibrium for Chores. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2285–2302. https://doi.org/10.1137/1.9781611977073.92

William C Brainard and Herbert E Scarf. 2005. How to compute equilibrium prices in 1891. *American Journal of Economics and Sociology* 64, 1 (2005), 57–83.

Mark Braverman. 2021. Optimization-friendly generic mechanisms without money. *arXiv preprint arXiv:2106.07752* (2021).

Simina Brânzei and Fedor Sandomirskiy. 2024. Algorithms for Competitive Division of Chores. *Mathematics of Operations Research* 49, 1 (2024), 398–429. https://doi.org/10.1287/moor.2023.1361

Bhaskar Ray Chaudhury, Jugal Garg, Peter McGlaughlin, and Ruta Mehta. 2022a. Competitive Equilibrium with Chores: Combinatorial Algorithm and Hardness. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*. 1106–1107. https://doi.org/10.1145/3490486.3538255

Bhaskar Ray Chaudhury, Jugal Garg, Peter McGlaughlin, and Ruta Mehta. 2022b. On the Existence of Competitive Equilibrium with Chores. In *Proceedings of the 13th Innovations in Theoretical Computer Science Conference (ITCS)*. 41:1–41:13. https://doi.org/10.4230/LIPICS.ITCS.2022.41

Thomas Chen, Xi Chen, Binghui Peng, and Mihalis Yannakakis. 2022. Computational Hardness of the Hylland-Zeckhauser Scheme. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2253–2268.

Xi Chen, Decheng Dai, Ye Du, and Shang-Hua Teng. 2009. Settling the Complexity of Arrow-Debreu Equilibria in Markets with Additively Separable Utilities. In *Proceedings of the 50th IEEE Symposium on Foundations of Computer Science (FOCS)*. 273–282. https://doi.org/10.1109/FOCS.2009.29

Xi Chen, Dimitris Paparas, and Mihalis Yannakakis. 2017. The Complexity of Non-Monotone Markets. *J. ACM* 64, 3 (2017), 20:1–20:56. https://doi.org/10.1145/3064810

Xi Chen and Shang-Hua Teng. 2009. Spending is not easier than trading: on the computational equivalence of Fisher and Arrow-Debreu equilibria. In *International Symposium on Algorithms and Computation*. Springer, 647–656.

Bruno Codenotti, Sriram V Pemmaraju, and Kasturi R Varadarajan. 2005. On the polynomial time computation of equilibria for certain exchange economies. In *SODA*, Vol. 5. 72–81.

Bruno Codenotti, Amin Saberi, Kasturi Varadarajan, and Yinyu Ye. 2006. Leontief economies encode nonzero sum two-player games. In *SODA*, Vol. 6. 659–667.

Argyrios Deligkas, John Fearnley, Alexandros Hollender, and Themistoklis Melissourgos. 2022. Pure-Circuit: Strong Inapproximability for PPAD. In *Proceedings of the 63rd Symposium on Foundations of Computer Science (FOCS)*. 159–170. https://doi.org/10.1109/FOCS54457.2022.00022

Argyrios Deligkas, John Fearnley, Alexandros Hollender, and Themistoklis Melissourgos. 2023. Tight inapproximability for graphical games. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 5600–5607.

Nikhil R Devanur. 2004. The spending constraint model for market equilibrium: Algorithmic, existence and uniqueness results. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*. 519–528.

Nikhil R Devanur and Ravi Kannan. 2008. Market equilibria in polynomial time for fixed number of goods or agents. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, 45–53.

Nikhil R Devanur, Christos H Papadimitriou, Amin Saberi, and Vijay V Vazirani. 2008. Market equilibrium via a primal–dual algorithm for a convex program. *Journal of the ACM (JACM)* 55, 5 (2008), 1–18.

Ran Duan, Jugal Garg, and Kurt Mehlhorn. 2016. An improved combinatorial polynomial algorithm for the linear Arrow-Debreu market. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on discrete algorithms*. SIAM, 90–106.

Ran Duan and Kurt Mehlhorn. 2015. A combinatorial polynomial algorithm for the linear Arrow–Debreu market. *Information and Computation* 243 (2015), 112–132.

Edmund Eisenberg. 1961. Aggregation of utility functions. *Management Science* 7, 4 (1961), 337–350.

Jugal Garg, Edin Husić, and László A Végh. 2023. An auction algorithm for market equilibrium with weak gross substitute demands. *ACM Transactions on Economics and Computation* 11, 3-4 (2023), 1–24.

Jugal Garg, Ruta Mehta, Milind Sohoni, and Vijay V. Vazirani. 2015. A Complementary Pivot Algorithm for Market Equilibrium under Separable, Piecewise-Linear Concave Utilities. *SIAM J. Comput.* 44, 6 (2015), 1820–1847. https://doi.org/10.1137/140971002

Jugal Garg, Yixin Tao, and László A Végh. 2022. Approximating Equilibrium under Constrained Piecewise Linear Concave Utilities with Applications to Matching Markets. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2269–2284.

Jugal Garg and László A Végh. 2019. A strongly polynomial algorithm for linear exchange markets. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. 54–65.

Aanund Hylland and Richard Zeckhauser. 1979. The efficient allocation of individuals to positions. *Journal of Political economy* 87, 2 (1979), 293–314.

Stavros D Ioannidis, Bart de Keijzer, and Carmine Ventre. 2023. Clearing Financial Networks with Derivatives: From Intractability to Algorithms. *arXiv preprint arXiv:2312.05139* (2023).

Kamal Jain. 2007. A polynomial time algorithm for computing an Arrow–Debreu market equilibrium for linear utilities. *SIAM J. Comput.* 37, 1 (2007), 303–318.

Devansh Jalota, Marco Pavone, Qi Qi, and Yinyu Ye. 2023. Fisher markets with linear constraints: Equilibrium properties and efficient distributed algorithms. *Games and Economic Behavior* 141 (2023), 223–260.

Sham M Kakade, Michael Kearns, and Luis E Ortiz. 2004. Graphical economics. In *Learning Theory: 17th Annual Conference on Learning Theory, COLT 2004, Banff, Canada, July 1-4, 2004. Proceedings 17*. Springer, 17–32.

Robert R. Maxfield. 1997. General equilibrium and the theory of directed graphs. *Journal of Mathematical Economics* 27, 1 (1997), 23–51. https://doi.org/10.1016/0304-4068(95)00763-6

James B Orlin. 2010. Improved algorithms for computing Fisher's market clearing prices: Computing Fisher's market clearing prices. In *Proceedings of the forty-second ACM symposium on Theory of computing*. 291–300.

Aviad Rubinstein. 2018. Inapproximability of Nash equilibrium. *SIAM J. Comput.* 47, 3 (2018), 917–959. https://doi.org/10.1137/15M1039274

Vijay V Vazirani. 2010. Spending constraint utilities with applications to the adwords market. *Mathematics of Operations Research* 35, 2 (2010), 458–478.

Vijay V. Vazirani and Mihalis Yannakakis. 2011. Market equilibrium under separable, piecewise-linear, concave utilities. *J. ACM* 58, 3 (2011), 1–25.

Vijay V Vazirani and Mihalis Yannakakis. 2020. Computational complexity of the Hylland-Zeckhauser scheme for one-sided matching markets. *arXiv preprint arXiv:2004.01348* (2020).

László A Végh. 2012. Strongly polynomial algorithm for a class of minimum-cost flow problems with separable convex objectives. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. 27–40.

Yinyu Ye. 2008. A path to the Arrow–Debreu competitive market equilibrium. *Mathematical Programming* 111, 1-2 (2008), 315–348.

A Omitted proofs from Section 4

Proof of Lemma 4.3. From Lemma 4.2 we have that buyers other than buyer i can buy at most

$$\max\left(\{r_{\text{NOT}}, r_{\text{NAND}}\} \cup \{r_q^p \;:\; 1 \leq p \leq 2, 1 \leq q \leq d\}\right) + 2t$$
$$\geq 2/11 + 8/11 = 10/11$$

units of good j, where we have used the fact that $t = 4/11$, that $r_{\text{NOT}} = r_{\text{NAND}} = 2/11$, and that each $r_q^p \leq 2\varepsilon < 2/11$. Since $\varepsilon < 1/11$ we have that $10/11 < 1 - \varepsilon$, so if $x_{i,j} = 0$, then good j does not ε-clear, and so we are not in an ε-equilibrium. □

Proof of Lemma 4.4. We first prove that $p_j > 0$. For the sake of contradiction, suppose that some good $j \in G \setminus \{\text{ref}\}$ has $p_j = 0$. Recall that in the Pure-Circuit instance, each vertex is an output of exactly one gate. As a result, in our construction, for every good $j \in G \setminus \{\text{ref}\}$ there is a buyer i, namely the inverter of the gate whose output is j, who is never satiated with j. Since the price of good j is 0, this buyer will demand an infinite amount of good j, making $\text{OPT}_i(p) = \emptyset$. This implies that we cannot be in an ε-equilibrium.

We now prove that $p_j \leq H$. So, for the sake of contradiction, suppose that $p_j > H$. By construction, as mentioned above, good j is the output good of some gate, and in particular, an inverter uses this good as an output. That inverter's marginal utility for ref is $1/p_{\text{ref}}$, and for good j the marginal utility is $s/p_j < s/H = 1/p_{\text{ref}}$. Moreover, the inverter is never satiated with ref, meaning that they will demand zero units of good j. We can now use Lemma 4.3 to obtain a contradiction. □

Proof of Lemma 4.5. We first argue that $x_{i,j} \leq r$. This follows from the fact that, once buyer i is allocated r units of good j, their marginal utility for good j is 0, whereas their marginal utility for the reference good is $1/p_{\text{ref}} > 0$. Thus buyer j strictly prefers the reference good to good j, and so cannot be allocated more than r units of good j.

Next we argue that $x_{i,j} \geq r$. If buyer i has been allocated strictly less than r units of good j, then their marginal utility for good j is $2s/p_j$. From Lemma 4.4, we have that $p_j \leq H$, so their marginal utility for good j is at least $2s/H$. On the other hand, their marginal utility for reference good is $1/p_{\text{ref}} = s/H$, and their marginal utility for all other goods is 0. Thus buyer i strictly prefers good j to any other good. Moreover, since $e_i = r \cdot H_{\text{high}} \geq r \cdot H$, buyer i has enough money to buy r units of good j. Thus any optimal bundle must allocate at least r units of good j to buyer i. □

Proof of Lemma 4.7. We first consider the case where $p_{\text{out}} > L$. By Lemma 4.6, we have that the inverter must buy t units of out before buying any other good. There are two sub-cases.

- If $e_{\text{inverter}} \geq t \cdot p_{\text{in}}$ then the inverter has enough money to buy t units of in, and does so.
- If $e_{\text{inverter}} < t \cdot p_{\text{in}}$ then the inverter does not have enough money to buy t units of in. This means that they spend all of their money on in, and therefore they they will demand zero units of out, which contradicts Lemma 4.3.

We now consider the case where $0 < p_{\text{out}} \leq L$. First note that the marginal utility of out is $s/p_{\text{out}} > s/L = a/H$, whereas the marginal utility of ref is $1/p_{\text{ref}} = s/H$. Since $a > s$, and since the inverter is never satiated by out, this means that the inverter cannot spend anything on ref. Due to the ε-clearing constraint, we know that the inverter can buy at most $1 + \varepsilon$ units of out, and so the inverter will spend at most $(1 + \varepsilon) \cdot L$ money on out. All remaining money must be spent on in, so the inverter will spend at least $e_{\text{inverter}} - (1 + \varepsilon) \cdot L$ money on in. Therefore, since $p_{\text{in}} \leq H$ by

Lemma 4.4, the inverter will buy at least

$$\begin{aligned}
&(e_{\text{inverter}} - (1+\varepsilon)\cdot L)/H \\
&= t \cdot H_{\text{low}}/H - (1+\varepsilon)\cdot s/a \\
&\geq t \cdot H_{\text{low}}/H_{\text{high}} - (1+\varepsilon)\cdot s/a \\
&\geq t \cdot (1 - 4/k) - (1+\varepsilon)\cdot s/a \\
&\geq t - 4t/k - (1+\varepsilon)\cdot s/a \\
&\geq t - 2/k - (1+\varepsilon)\cdot s/a \\
&\geq t - 2/k - 2\cdot(1/20k) \\
&\geq t - 3/k
\end{aligned}$$

units of in. Here we used that $t \leq 1/2$, that $s \leq 1/(20kd|V|) \leq 1/20k$, that $a \geq 1$, and that $1+\varepsilon \leq 2$. □

Proof of Lemma 4.8. This proof is very similar to the proof of Lemma 4.7, but now we must account for the fact that the inverter has two input goods.

We first consider the case where $p_{\text{out}} > L$. By Lemma 4.6, the inverter must buy t units of in1 and t units of in2 before buying any other good. There are now two sub-cases.

- If $e_{\text{inverter}} \geq t \cdot p_{\text{in1}} + t \cdot p_{\text{in2}}$ then the inverter has enough money to buy t units of in1 and t units of in2, and does so.
- If $e_{\text{inverter}} < t \cdot p_{\text{in1}} + t \cdot p_{\text{in2}}$ then the inverter does not have enough money to buy t units of in1 and t units of in2. This means that they spend all of their money on the input goods, and therefore they they will demand zero units of out, which contradicts Lemma 4.3.

We now consider the case where $0 < p_{\text{out}} \leq L$. First note that the marginal utility of out is $s/p_{\text{out}} > s/L = a/H$, whereas the marginal utility of ref is $1/p_{\text{ref}} = s/H$. Since $a > s$, and since the inverter is never satiated by out, this means that the inverter cannot spend anything on ref. Due to the ε-clearing constraint, we know that the inverter can buy at most $1+\varepsilon$ units of out, and so the inverter will spend at most $(1+\varepsilon)\cdot L$ money on out.

All remaining money must be spent on in1 and in2 so the inverter will spend at least $e_{\text{inverter}} - (1+\varepsilon)\cdot L$ money on these two goods. Note also that the inverter cannot buy more than t units of either input good, since their marginal utility for that input becomes 0 once they have t units. Thus, for any input good $j \in \{\text{in1}, \text{in2}\}$, in the worst case the inverter buys t units of the other input and pays price H for those units, and then spends the rest of their money on good j and pays price H for those units as well. So the inverter will buy at least

$$\begin{aligned}
(e_{\text{inverter}} - t\cdot H - (1+\varepsilon)\cdot L)/H &= 2t \cdot H_{\text{low}}/H - t - (1+\varepsilon)\cdot s/a \\
&\geq 2t \cdot (1-4/k) - t - 2\cdot(1/20k) \\
&\geq t - 8t/k - 1/10k \\
&\geq t - 4/k - 1/10k \\
&\geq t - 5/k
\end{aligned}$$

units of good j. Here we used that $t \leq 1/2$, that $s \leq 1/(20kd|V|) \leq 1/20k$, that $a \geq 1$, and that $1+\varepsilon \leq 2$. □

Proof of Lemma 4.9. By construction, for each good, there are exactly two buyers that are not part of the gate's construction that are interested in buying that good. Those buyers are either inverter buyers that implement a NOT or NAND gate, or are auxiliary buyers that top-up the amount bought in the case where the variable encoded by that good is the input to fewer than two other gates. By Lemma 4.7 and Lemma 4.8 the inverter buyers will buy some amount in the range $[t - 5/k, t]$, while by Lemma 4.5 the auxiliary buyers will buy exactly t units of the good. Thus the total amount of the good bought by these buyers lies in the range $[2\bar{t}, 2t]$. □

Proof of Lemma 4.10. We begin by considering the case where $p_{\text{out}} > L$. In this case, by Lemma 4.6, which works for NOT gates with any value of r, we have that the inverter of the NOT gate must buy t units of the input good before buying any other good. This means that the inverter can spend at most

$$e_{\text{inverter}} - t \cdot p_{\text{in}} = t \cdot (H_{\text{low}} - p_{\text{in}})$$

money on goods other than the input good.

By Lemma 4.2, buyers other than the inverter demand at most $2t + r$ units of the output good. Thus, to ensure that the output good ε-clears, we require that at least $1 - \varepsilon - 2t - r$ units of the output good are bought by the inverter, and so the inverter must spend at least $p_{\text{out}} \cdot (1 - \varepsilon - 2t - r)$ money on the output good.

For this to be possible we must have

$$t \cdot (H_{\text{low}} - p_{\text{in}}) \geq p_{\text{out}} \cdot (1 - \varepsilon - 2t - r),$$

and therefore

$$p_{\text{out}} \leq (H_{\text{low}} - p_{\text{in}}) \cdot \frac{t}{1 - \varepsilon - 2t - r}$$

since $1 - \varepsilon - 2t - r > 0$. So we have shown that either $p_{\text{out}} \leq L$, or that the bound above holds. So we can conclude that

$$p_{\text{out}} \leq \max\left(L,\ (H_{\text{low}} - p_{\text{in}}) \cdot \frac{t}{1 - \varepsilon - 2t - r}\right).$$

□

Proof of Lemma 4.11. We first consider the case where $p_{\text{out}} < H$. In this case, the marginal utility of buying the reference good is $1/p_{\text{ref}} = s/H$, whereas the marginal utility of the output good is $s/p_{\text{out}} > s/H$. Thus the inverter strictly prefers the output good to the reference good. Since the inverter can never be satiated by the output good, this means that the inverter cannot buy the reference good.

Furthermore, the inverter cannot buy more than t units of the input good, because the inverter's marginal utility becomes 0 at that point. This means that the inverter can spend at most $t \cdot p_{\text{in}}$ money on the input good, and will have at least

$$e_{\text{inverter}} - t \cdot p_{\text{in}} = t \cdot (H_{\text{low}} - p_{\text{in}})$$

money left over after buying all goods other than out. This money must therefore be spent on out.

By Lemma 4.9 we have that buyers who are not part of the NOT gate will buy at least $2\bar{t}$ units of the output good, and by Lemma 4.5 we have that r units of the output good will be bought by the auxiliary buyer of the NOT gate. Since the good must ε-clear, this means that the inverter can buy at most $1 - 2\bar{t} - r + \varepsilon$ units of the output good.

Since $t \cdot (H_{\text{low}} - p_{\text{in}})$ money must be spent on the output good, and since at most $1 - 2\bar{t} - r + \varepsilon$ units of that good can be bought, we have

$$t \cdot (H_{\text{low}} - p_{\text{in}}) \leq (1 - 2\bar{t} - r + \varepsilon) \cdot p_{\text{out}},$$

because otherwise there would be money left over that is not spent on any good, which cannot happen in an optimal allocation. Rearranging this gives

$$p_{\text{out}} \geq (H_{\text{low}} - p_{\text{in}}) \cdot \frac{t}{1 - 2\bar{t} - r + \varepsilon}$$

since $1 - 2\bar{t} - r + \varepsilon \geq 1 - 2t - r - \varepsilon > 0$.

So we have shown that we either have $p_{\text{out}} \geq H$, or we have the inequality given above. Hence we have

$$p_{\text{out}} \geq \min\left(H,\ (H_{\text{low}} - p_{\text{in}}) \cdot \frac{t}{1 - 2\bar{t} - r + \varepsilon}\right).$$

□

Proof of Lemma 4.13. Assume, for the sake of contradiction, that $p_{\text{out}} > L$. Then by Lemma 4.6 we have that the inverter must buy t units of in1 and t units of in2 before buying any other good. The inverter's remaining budget after buying t units of both of the inputs goods is

$$\begin{aligned}
&e_{\text{inverter}} - t \cdot p_{\text{in1}} - t \cdot p_{\text{in2}} \\
&= 2t \cdot H_{\text{low}} - t \cdot p_{\text{in1}} - t \cdot p_{\text{in2}} \\
&\leq 2t \cdot H_{\text{low}} - t \cdot H - t \cdot H \\
&\leq 0.
\end{aligned}$$

Hence the inverter cannot spend any money on the output good, which contradicts Lemma 4.3. □

Proof of Lemma 4.15. Consider the gate $\mathsf{NOT}(\text{in}, \text{mid}, r)$ with $p_{\text{in}} \geq H_{\text{low}} \cdot (1 - A) + P \cdot AB'$. This implies that $H_{\text{low}} - p_{\text{in}} \leq H_{\text{low}} \cdot A - P \cdot AB'$, or equivalently, $(H_{\text{low}} - p_{\text{in}})/A \leq H_{\text{low}} - P \cdot B'$. Now notice that $H_{\text{low}} - P \cdot B' \geq H_{\text{low}} - P \cdot \left(\frac{H_{\text{low}}}{H_{\text{high}}} - L/H\right) \geq L$, where the two inequalities come from the fact that $B' = 1 - 4/k - s/a \leq \frac{H_{\text{low}}}{H_{\text{high}}} - \frac{L}{H}$ and $P \leq H \leq H_{\text{high}}$, respectively. Therefore, Lemma 4.10 implies that $p_{\text{mid}} \leq H_{\text{low}} - P \cdot B'$.

Now consider the gate $\mathsf{NOT}(\text{mid}, \text{out}, r')$. From above, we have $H_{\text{low}} - p_{\text{mid}} \geq P \cdot B'$, which implies $(H_{\text{low}} - p_{\text{mid}})/B' \geq P$. Since $P \leq H$, Lemma 4.11 implies that $p_{\text{out}} \geq P$. □

Proof of Lemma 4.16. Consider the gate $\mathsf{NOT}(\text{in}, \text{mid}, r)$ with $p_{\text{in}} \leq H_{\text{low}}\,(1 - B) + P \cdot A'B$. This implies that $H_{\text{low}} - p_{\text{in}} \geq H_{\text{low}} \cdot B - P \cdot A'B$, or equivalently, $(H_{\text{low}} - p_{\text{in}})/B \geq H_{\text{low}} - P \cdot A'$. Notice that $H_{\text{low}} - P \cdot A' \leq H_{\text{low}} \leq H$. Therefore, Lemma 4.11 implies that $p_{\text{mid}} \geq H_{\text{low}} - P \cdot A'$.

Now consider the gate $\mathsf{NOT}(\text{mid}, \text{out}, r')$. As shown above, we have $H_{\text{low}} - p_{\text{mid}} \leq P \cdot A'$, which means that $(H_{\text{low}} - p_{\text{mid}})/A' \leq P$. Since $P \geq L$, Lemma 4.10 implies that $p_{\text{out}} \leq P$. □

Proof of Lemma 4.18. First, we will need to calculate the quantities A, B, A', B' for each chain, where we add a subscript $i \in \{1, 2\}$ to indicate the chain they refer to. For ease of presentation, we define $\delta := \frac{11}{4} \cdot (1/11 - \varepsilon) > 0$, and notice that $\bar{t} \geq t - 5/k \geq t - \delta/22$. By construction, for chain 1, without loss of generality we can use only the first pair of NOT gates (i.e., $j \in \{1, 2\}$), and by the specified r_j^i above, we get

$$\begin{aligned}
A_1 &:= \frac{1 - 2t - r_1^1 - \varepsilon}{t} = \frac{3/11 - \varepsilon}{4/11} \\
B_1 &:= \frac{1 - 2\bar{t} - r_1^1 + \varepsilon}{t} \in \left[\frac{3/11 + \varepsilon}{4/11}, \frac{3/11 + \varepsilon}{4/11} + \delta\right] \\
A_1' &:= \frac{1 - 2t - r_2^1 - \varepsilon}{t} = \frac{1/11 - \varepsilon}{4/11} \\
B_1' &:= \frac{1 - 2\bar{t} - r_2^1 + \varepsilon}{t} \in \left[\frac{1/11 + \varepsilon}{4/11}, \frac{1/11 + \varepsilon}{4/11} + \delta\right]
\end{aligned}$$

Similarly, for chain 2 we have

$$\begin{aligned}
A_2 &:= \frac{1-2t-r_1^2-\varepsilon}{t} = \frac{1/11-\varepsilon}{4/11} \\
B_2 &:= \frac{1-2\bar{t}-r_1^2+\varepsilon}{t} \in \left[\frac{1/11+\varepsilon}{4/11}, \frac{1/11+\varepsilon}{4/11}+\delta\right] \\
A_2' &:= \frac{1-2t-r_2^2-\varepsilon}{t} = \frac{3/11-\varepsilon}{4/11} \\
B_2' &:= \frac{1-2\bar{t}-r_2^2+\varepsilon}{t} \in \left[\frac{3/11+\varepsilon}{4/11}, \frac{3/11+\varepsilon}{4/11}+\delta\right]
\end{aligned}$$

Finally, notice that $B_2 \in [A_1 - 2\delta, A_1 - \delta]$, and $B_1' \in [A_2' - 2\delta, A_2' - \delta]$.

Corollary 4.17 implies that for chain 1 there is an interval (R_1^L, R_1^U) of p_{in} prices for which p_{out} is not guaranteed to be at most L or at least H. In particular,

$$R_1^L := H_{\text{low}} \cdot \frac{1-B_1}{1-A_1'B_1} + (A_1'B_1)^{d/2} \cdot \left(L - H_{\text{low}} \cdot \frac{1-B_1}{1-A_1'B_1}\right),$$

and

$$R_1^U := H_{\text{low}} \cdot \frac{1-A_1}{1-A_1B_1'} + (A_1B_1')^{d/2} \cdot \left(H - H_{\text{low}} \cdot \frac{1-A_1}{1-A_1B_1'}\right)$$

Similarly, the corollary implies that, for chain 2, the respective interval is (R_2^L, R_2^U) with

$$R_2^L := H_{\text{low}} \cdot \frac{1-B_2}{1-A_2'B_2} + (A_2'B_2)^{d/2} \cdot \left(L - H_{\text{low}} \cdot \frac{1-B_2}{1-A_2'B_2}\right),$$

and

$$R_2^U := H_{\text{low}} \cdot \frac{1-A_2}{1-A_2B_2'} + (A_2B_2')^{d/2} \cdot \left(H - H_{\text{low}} \cdot \frac{1-A_2}{1-A_2B_2'}\right)$$

We will now show that

$$R_1^L < R_1^U \le R_2^L < R_2^U.$$

First, by Corollary 4.17, it is immediate that $R_1^L < R_1^U$. Otherwise, $p_{\text{in}} = R_1^L$ would result into $p_{\text{out1}} \le L$, but because $R_1^L \ge R_1^U$, we would have $p_{\text{out1}} \ge H$. This implies $H \le L = s \cdot H/a$, a contradiction, since $s < 1/2 < a$. Similarly, by Corollary 4.17, it is straightforward that $R_2^L < R_2^U$.

Now we will show that $R_1^U \leq R_2^L$. We have

$$
\begin{aligned}
R_1^U &= H_{\text{low}} \cdot \frac{1 - A_1}{1 - A_1 B_1'} + (A_1 B_1')^{d/2} \cdot \left(H - H_{\text{low}} \cdot \frac{1 - A_1}{1 - A_1 B_1'} \right) \\
&= H_{\text{low}} \cdot \frac{1 - A_1}{1 - A_1 B_1'} \left(1 - (A_1 B_1')^{d/2} \right) + H \cdot (A_1 B_1')^{d/2} \\
&\leq H_{\text{low}} \cdot \frac{1 - A_1}{1 - (B_2 + 2\delta) B_1'} \left(1 - (A_1 B_1')^{d/2} \right) + H \cdot (A_1 B_1')^{d/2} && (A_1 \leq B_2 + 2\delta) \\
&\leq H_{\text{low}} \cdot \frac{1 - B_2 - \delta}{1 - B_2 (A_2' - \delta) - 2B_1' \delta)} \left(1 - (A_1 B_1')^{d/2} \right) + H \cdot (A_1 B_1')^{d/2} && (B_1' \leq A_2' - \delta) \\
&= H_{\text{low}} \cdot \frac{1 - B_2 - \delta}{1 - A_2' B_2 - \delta B_1'} \left(1 - (A_1 B_1')^{d/2} \right) + H \cdot (A_1 B_1')^{d/2} && (\delta, B_1' = B_2 > 0) \\
&\leq H_{\text{low}} \cdot \frac{1 - B_2 - \delta}{1 - A_2' B_2 - \delta B_1'} \left(1 - (1/8)^{d/2} \right) + H \cdot (3/8)^{d/2} && \left(\frac{1}{2} \cdot \frac{1}{4} \leq A_1 B_1' \leq \frac{3}{4} \cdot \frac{1}{2} \right) \\
&\leq H_{\text{low}} \cdot \frac{1 - B_2}{1 - A_2' B_2 - \delta B_1'} - H_{\text{low}} \cdot \frac{\delta}{1 - A_2' B_2 - \delta B_1'} \left(1 - (1/8)^{d/2} \right) + H \cdot (3/8)^{d/2} \\
&\leq H_{\text{low}} \cdot \frac{1 - B_2}{1 - A_2' B_2} \left(1 + \frac{\delta}{1 - A_2' B_2 - \delta B_1'} \right) - H_{\text{low}} \cdot \frac{\delta}{1 - A_2' B_2 - \delta B_1'} + \left(H_{\text{low}} \cdot \frac{1}{1 - A_2' B_2} + H \right) \cdot (3/8)^{d/2} \\
&\leq H_{\text{low}} \cdot \frac{1 - B_2}{1 - A_2' B_2} - H_{\text{low}} \cdot \frac{\delta}{1 - A_2' B_2 - \delta B_1'} \left(1 - \frac{1 - B_2}{1 - A_2' B_2} \right) + \left(H_{\text{low}} \cdot \frac{1}{1 - A_2' B_2 - \delta B_1'} + H \right) \cdot (3/8)^{d/2} \\
&\leq H_{\text{low}} \cdot \frac{1 - B_2}{1 - A_2' B_2} - H_{\text{low}} \cdot \frac{1 - B_2}{1 - A_2' B_2} (3/8)^{d/2} && (4) \\
&\leq H_{\text{low}} \cdot \frac{1 - B_2}{1 - A_2' B_2} \left(1 - (3/8)^{d/2} \right) \\
&\leq H_{\text{low}} \cdot \frac{1 - B_2}{1 - A_2' B_2} \left(1 - (A_2' B_2)^{d/2} \right) + L \cdot (A_2' B_2)^{d/2} && \left(A_2' B_2 \leq \frac{3}{4} \cdot \frac{1}{2}, \text{ and } L \geq 0 \right) \\
&= H_{\text{low}} \cdot \frac{1 - B_2}{1 - A_2' B_2} + (A_2' B_2)^{d/2} \cdot \left(L - H_{\text{low}} \cdot \frac{1 - B_2}{1 - A_2' B_2} \right) \\
&= R_2^L,
\end{aligned}
$$

where Inequality (4) holds due to our choice of

$$
d \geq 4 \lceil \ln(4/\delta) \rceil \geq \frac{2}{\ln(8/3)} \cdot \ln \left(\frac{1 + \frac{1 - B_2}{1 - A_2' B_2} + H/H_{\text{low}} + \delta}{\delta \cdot \frac{1 - A_2' B_2}{B_2 (1 - A_2')}} \right),
$$

where $\frac{1 - B_2}{1 - A_2' B_2} \leq 1$ (since $B_2, A_2' \in (0, 1)$), $\frac{1 - A_2' B_2}{B_2 (1 - A_2')} \geq 2$, $H/H_{\text{low}} \leq 2$, and $\delta \leq 1$.

Now we are ready to prove the first part of the lemma's statement. Suppose $p_{\mathsf{in}} \geq H$. Notice that

$$\begin{aligned}
H &= H\left(1-(A_2B_2')^{d/2}\right) + H\cdot (A_2B_2')^{d/2} \\
&\geq H\cdot \frac{1-A_2}{1-A_2B_2'}\left(1-(A_2B_2')^{d/2}\right) + H\cdot (A_2B_2')^{d/2} && (0 < A_2, B_2' < 1)\\
&\geq H_{\text{low}}\cdot \frac{1-A_2}{1-A_2B_2'}\left(1-(A_2B_2')^{d/2}\right) + H\cdot (A_2B_2')^{d/2} \\
&= H_{\text{low}}\cdot \frac{1-A_2}{1-A_2B_2'} + (A_2B_2')^{d/2}\cdot\left(H - H_{\text{low}}\cdot\frac{1-A_2}{1-A_2B_2'}\right)\\
&= R_2^U.
\end{aligned}$$

Therefore, $p_{\mathsf{in}} \geq R_2^U > R_1^U$, and as a result, $p_{\mathsf{out1}} \geq H$, and $p_{\mathsf{out2}} \geq H$.

Moving on to the proof of the second part of the statement, suppose $p_{\mathsf{in}} \leq L$. We have

$$\begin{aligned}
L &= L\left(1-(A_1'B_1)^{d/2}\right) + L\cdot (A_1'B_1)^{d/2}\\
&\leq \frac{s\delta}{2}\cdot\left(1-(A_1'B_1)^{d/2}\right) + L\cdot (A_1'B_1)^{d/2} && (L \leq s\delta/2)\\
&\leq H_{\text{low}}\cdot\frac{1-B_1}{1-A_1'B_1}\left(1-(A_1'B_1)^{d/2}\right) + L\cdot (A_1'B_1)^{d/2} && (H_{\text{low}} \geq H_{\min} \geq s/2, 1-B_1=\delta, A_1', B_1 \geq 0)\\
&= H_{\text{low}}\cdot\frac{1-B_1}{1-A_1'B_1} + (A_1'B_1)^{d/2}\cdot\left(L - H_{\text{low}}\cdot\frac{1-B_1}{1-A_1'B_1}\right)\\
&= R_1^L.
\end{aligned}$$

So, $p_{\mathsf{in}} \leq R_1^L < R_2^L$, therefore, $p_{\mathsf{out1}} \leq L$, and $p_{\mathsf{out2}} \leq L$.

Now, for the third part of the statement, suppose that $p_{\mathsf{in}} \in (L, H)$. If $p_{\mathsf{in}} \leq R_2^L$, then by Corollary 4.17 we have that $p_{\mathsf{out2}} \leq L$. If $p_{\mathsf{in}} > R_2^L$, then $p_{\mathsf{in}} > R_1^U$ (since we showed that $R_2^L \geq R_1^U$), and again from Corollary 4.17, $p_{\mathsf{out1}} \geq H$. Therefore, if $p_{\mathsf{in}} \in (L, H)$, then it cannot be that both p_{out1} and p_{out2} are in (L, H). ☐