Computational Hardness of the Hylland-Zeckhauser Scheme*

Thomas Chen[†] Xi Chen[‡] Binghui Peng[§] Mihalis Yannakakis[¶]

Abstract

We study the complexity of the classic Hylland-Zeckhauser scheme [21] for one-sided matching markets. We show that the problem of finding an ϵ-approximate equilibrium in the HZ scheme is PPAD-hard, and this holds even when ϵ is polynomially small and when each agent has no more than four distinct utility values. Our hardness result, when combined with the PPAD membership result of [29], resolves the approximation complexity of the HZ scheme. We also show that the problem of approximating within a certain constant factor the optimal social welfare (the weight of the matching) achievable by HZ equilibria is NP-hard.

1 Introduction

In a *one-sided matching* problem, there is a set A of n agents and a set G of n goods, and we are given a specification of the preferences of the agents for the goods. (In general there can be n_1 agents, each with their own demand d_i, and n_2 goods, each with its own supply, s_j, where $\sum_i d_i = \sum_j s_j$. It is clear that this setting reduces to the simpler case of equal numbers of agents and goods.) The problem is to find a matching between the agents and the goods (assigning a distinct good to each agent) that has desirable properties, such as Pareto optimality, envy-freeness, incentive compatibility. This situation, where only one side has preferences, arises in many settings, such as assigning students to schools, assigning faculty members to committees, workers to tasks, program committee members to papers, students to courses with limited capacity, etc.

Since many agents may have the same or similar preferences, it is usually not possible to offer everybody their favorite good. So a solution mechanism has to strive to be equitable, satisfy the agents as much as possible, and incentivize them to give their true preferences (i.e., not gain an advantage by lying). Randomization is often useful to meet fairness requirements. A randomized solution mechanism has probability $x_{i,j} \in [0, 1]$ of matching each agent i to each good j; these probabilities form a doubly stochastic matrix, i.e., a fractional perfect matching in the bipartite graph between agents and goods. In some applications, the goods may be divisible, or they may represent tasks or resources that can be shared among agents; in these cases the quantities $x_{i,j}$ represent the shares of the agents in the goods.

There are two main ways of specifying the preferences of each agent $i \in [n]$ for the goods: (1) *cardinal preferences*, where we are given the utility $u_{i,j}$ of agent i for each good $j \in [n]$, or (2) *ordinal preferences*, where we are given the agent's total ordering of the goods. Cardinal preferences allow for a finer specification of the agents' preferences (although they may require more effort to produce them). As a result, they can yield better assignments. Consider for instance the following example from [21]: There are 3 agents and 3 goods. The utilities of agents 1 and 2 for the three goods are 100, 10, 0, while agent 3 has utilities 100, 80, 0. The ordinal preferences of the three agents are the same, so any fair mechanism will not distinguish between them, and will give them each probability 1/3 for each good. The expected utilities of the three agents in this solution is $36\frac{2}{3}, 36\frac{2}{3}, 60$. This solution is not Pareto optimal, i.e., there is another solution where all agents are better off. Agent 3 is assigned good 2, and agents 1, 2 randomly split goods 1 and 3. The expected utilities of the three agents in this solution are $50, 50, 80$.

In 1979 Hylland and Zeckhauser proposed a, by now, classic scheme for the one-sided matching problem under cardinal preferences [21]. The scheme uses a pricing mechanism to produce an assignment of probability shares $\{x_{i,j} | i, j \in [n]\}$ of goods to agents, i.e. a fractional perfect matching, and these are then used in a standard way to generate probabilistically an integral perfect matching. The basic idea is to imagine a market where every agent

*The full version of the paper [5] can be accessed at https://arxiv.org/abs/2107.05746

[†]Columbia Univeristy, tc3039@columbia.edu
[‡]Columbia University, xichen@cs.columbia.edu
[§]Columbia University, bp2601@columbia.edu
[¶]Columbia University, mihalis@cs.columbia.edu

has 1 dollar, and the goal is to find prices for the goods and (fractional) allocations $x_{i,j}$ of goods to the agents, such that the market clears (all goods are sold), while every agent maximizes her utility subject to receiving a bundle of goods of size 1 and cost at most 1. Hylland and Zeckhauser showed that such an equilibrium set of prices and allocations always exists, using Kakutani's fixed point theorem. Note that money here is fictitious; no money changes hands. The only goal is to produce the allocation (the shares $x_{i,j}$) so that it reflects the preferences of the agents. The HZ scheme has several desirable properties: it is Pareto optimal, envy-free [21], and it is incentive compatible in the large [19]. The scheme has been extended and generalized in various ways since then.

Although the HZ scheme has several nice properties, one impediment is that, despite much effort, there is no efficient algorithm known to compute an equilibrium solution. This has remained an open problem till now. In [2], Alaei, Khalilabadi and Tardos gave polynomial-time algorithms for the case that the number of goods or the number of agents is a fixed constant (the case of a constant number of goods can be derived also from [9]). Recently in [29], Vazirani and Yannakakis gave a polynomial-time algorithm for the bi-valued case, where every agent's utilities take only two values. They also gave an example showing that the equilibrium prices and allocations can be inherently irrational. In the general case, they showed that the problem of computing an equilibrium solution is in the class FIXP. Furthermore, computing an ϵ-approximate equilibrium is in the class PPAD, where in an approximate equilibrium an agent may get a slightly suboptimal allocation and may spend $1 + \epsilon$ dollars. They leave open the problem whether computing an exact or approximate equilibrium is complete for the classes.

In this paper we resolve the complexity of computing an approximate equilibrium of the HZ scheme. Our main result is:

Theorem 1.1. (Main) *The problem of computing an ϵ-approximate equilibrium of the HZ scheme is PPAD-complete when $\epsilon = 1/n^c$ for any constant $c > 0$.*

In our construction, every agent has at most 4 different utilities for the goods. Thus, the problem is PPAD-complete even for 4-valued utilities[1]. We leave the 3-valued case open. In the full version [5], we give however a simple example with values in $\{0, 0.5, 1\}$ showing that there can be multiple disconnected equilibria, thus suggesting that usual convex programming methods may not work (at least a convex program will not include all equilibria).

A given instance of the one-sided matching problem may have multiple HZ equilibria. All of them are Pareto optimal, but some may be preferable to others when other criteria are considered. One such criterion is the social welfare, i.e., the total weight of the matching (or the sum of utilities of agents). We study the problem of approximating the optimal social welfare achievable by an HZ equilibrium. We show that this is an NP-hard problem:

Theorem 1.2. *Given an instance of the one-side matching problem and a value w, it is NP-hard to distinguish the case that the maximum social welfare of an HZ equilibrium is at least w from the case that it is at most $(\frac{175}{176} + \epsilon)w$ for any constant $\epsilon > 0$.*

1.1 Proof Overview We give an overview of the proof of Theorem 1.1. To prove the problem of finding an approximate HZ equilibrium is PPAD-hard, we give a polynomial-time reduction from *threshold games*, introduced recently by Papadimitriou and Peng [27]. A threshold game is defined on a directed graph $G = (V, E)$, with a variable $x_v \in [0, 1]$ associated with each node $v \in V$. The equilibrium condition is characterized by a comparison operator: $x_v = 1$ if $\sum_{(u,v)\in E} x_u \le 1/2 - \kappa$; $x_v = 0$ if $\sum_{(u,v)\in E} x_u \ge 1/2 + \kappa$ and x_v can take an arbitrary value in $[0, 1]$, otherwise. The PPAD-hardness of threshold game is proved in [27], and it holds for some positive constant $\kappa > 0$ and for sparse graphs. As we will see later, the use of threshold games significantly simplifies the reduction.

From a high-level view, our reduction follows the general framework of previous hardness results on market equilibria [6, 8, 28, 7]: we use prices of an HZ market to simulate variables x_v in a threshold game and the construction is based on the design of two gadgets: *variable gadgets* for each $v \in V$ (to simulate variables x_v and enforce the equilibrium condition at each node v) and *edge gadgets* for each $e = (u, v) \in E$ (to simulate the action of sending x_u to the sum at x_v in the threshold game). However, a major challenge of working with the

[1]The case of a small number of values is natural. For example, some program committees ask the members to rate their level of interest in the submissions by values in a limited range, e.g. 0-4.

HZ scheme is that it is difficult to characterize the equilibrium behavior of agents in this model, as it is complex, nonlinear, and does not admit a closed form solution. As a consequence, it is hard both, to analyze even small instances, and to synthesize instances with desired characteristics. Below we discuss some of the key ideas behind the construction.

Variable gadgets. To simulate variables x_v of a threshold game, our starting point is the following simple sub-market. There are two agents and two goods. Both agents have utility 1 for good 1 and utility $\frac{1}{2\delta^{-1}-1}$ for good 2 ($0 < \delta < 1$ which should be considered as a small constant as discussed below). Moreover, each agent can also choose to buy another zero-priced good, for which it has zero utility. Observe that the set of equilibrium prices for the two goods of the submarket are $(p, 2-p)$ for $p \in [0, \delta]$ (because of the option to buy the zero-priced good). After scaling by $1/\delta$, the price p of good 1 in this sub-market can be used to simulate a variable $x_v \in [0, 1]$ in the threshold game. So we can create such a sub-market M_v for each node v and denote the price of good 1 in M_v by p_v. To finish the reduction, it suffices to create agents that are interested in goods in M_u and M_v, for each edge $e = (u, v) \in E$, such that the total allocation of goods from M_v to them is captured by $-p_u$.[2] The agents created for this task are what we referred to earlier as the edge gadget for e. If achieved, the total allocation of goods of M_v to agents outside would give the desired linear form $-\sum_{(u,v)\in E} p_u$ which is a scaled version of $-\sum_{(u,v)\in E} x_u$. The sub-market M_v can help enforce the equilibrium condition of the threshold game. When $\sum_{(u,v)\in E} x_u$ is too small (so the allocation to agents outside is high due to the negative sign), there is a shortage of goods of M_v, which would lead to $p_v \approx \delta$ and thus, $x_v \approx 1$; on the other hand, if the sum is too large, then there is a surplus of goods in the sub-market M_v, which forces $p_v \approx 0$ and thus, $x_v \approx 0$.

Edge gadgets. The key technical challenge lies in the construction of edge gadgets. Our first attempt is to create an agent who has utility $1/2$ for good 1 in M_u and utility 1 for good 2 in M_v (with price $2 - p_v$). The optimal bundle for this agent, however, is not easy to work with at first sight: for example, the agent is allocated $\frac{1-p_u}{2-p_u-p_v}$ unit of good 2 in M_v. **The first key idea of our reduction is to use first order approximation to simplify a complex function**. That is to say, we set δ to be a sufficiently small constant and apply first-order approximation on the allocation. Ignoring constant factors and constant or lower-order terms, the agent described earlier has $D_u \approx p_u - p_v$ unit of good 1 in M_u and $D_v \approx p_v - p_u$ unit of good 2 in M_v in her optimal bundle. This, however, is far from what we hoped for, as (1) we don't want p_v to appear in D_v, and (2) we want $D_u = 0$. (Notably if we use this agent as our edge gadget for every $e \in E$, then they together are essentially simulating a threshold game over an undirected graph, which admits trivial equilibria.) What about agents with a different set of utilities for goods in M_u and M_v? Perhaps surprisingly, all our attempts fail and there is a fundamental reason for that: it can be shown that, no matter how the utilities are set, the allocation of goods in M_v is always monotonically increasing with p_v, which is undesirable for our purpose.

We circumvent this obstacle using the following three steps. First, we introduce extra goods with a fixed price 2 into the picture (where the fixed price can be enforced easily by creating agents who are interested in these goods only). Second, we replace the current variable gadget with a richer sub-market with three goods, with the price of the new good set to be $\frac{(1+p_v)}{2}$. Agents in our edge gadget can now have access to these new goods which give us a larger design space for their utilities and equilibrium behavior. The last step, which is **the second key idea of our reduction, is to use discrete functions to approximate a continuous function**. This is essential in our proof and we believe it might be of independent interest in reductions of similar settings. Concretely, we design agents that perform comparison operations: there are two possible optimal bundles for each of these agents, and which one it is depends on the sign (in a robust sense) of a certain affine linear form of p_u and p_v. The agent behaves like a step function, which is not useful on its own. However, when combined, one can construct a series of agents by enumerating utilities; these lead to careful cancellations that make sure the total allocation D_u of goods from M_u is 0 as desired.

1.2 Related Work We have already mentioned the most relevant work on the complexity of the Hylland-Zeckhauser scheme. The problem of computing an exact HZ equilibrium is in FIXP, and computing an approximate equilibrium is in PPAD [29]. Polynomial-time algorithms for a fixed number of agents or goods were given in [2]. It has been a longstanding open problem about whether there is a polynomial-time algorithm in the general case.

[2]The negative allocation may look strange but can be achieved (essentially) by offsetting the supply carefully.

The input in the one-sided matching problem is the same as in the classical assignment problem (equivalently, maximum weight perfect matching problem in bipartite graphs). This is one of the most well-studied problems in Operations Research and Computer Science, and several very efficient algorithms have been developed for it over the years. The difference in the one-sided matching problem is that the primary consideration is to produce a solution that has certain desirable fairness and optimality properties for the agents; the goal is not simply the maximization of the weight of the matching. As we show in this paper, computing an HZ solution for the one-sided matching problem is probably computationally harder: it is PPAD-hard to compute any approximate HZ solution, and if we want to maximize the total weight of the matching as a secondary criterion then the problem becomes NP-hard.

In the case of ordinal preferences for the agents in the one-sided matching problem, there are other schemes with nice properties: the Random Priority (also called Random Serial Dictatorship) scheme [1, 25] and the Probabilistic Serial scheme [3]. These have polynomial-time (randomized) algorithms. However, since they are based only on ordinal preferences, they are suboptimal with respect to the agents' utilities, as the earlier simple example shows.

The setting in the HZ scheme is the same as in the linear Fisher market model: the input consists of the utilities $u_{i,j}$ of the agents for the goods, and the problem is to compute equilibrium prices and allocations. The only difference is that when an agent picks her optimal bundle of goods, she must get exactly one unit (in addition to the cost being within the budget of 1 dollar), i.e. the solution must be a (fractional) perfect matching. Although this may seem like a small difference, it has a substantial effect both in the structure of the problem and in its computational complexity: exact solutions may be irrational, and as we show in this paper, finding an approximate solution is PPAD-hard. The linear Fisher model has been studied extensively and there are polynomial-time algorithms for computing equilibria in this model, as well as in the more general Arrow-Debreu model with linear utilities [22, 10, 26, 11, 17].

There is furthermore extensive work on markets with more complex utility functions than linear, such as piecewise linear, Leontief, CES utilities and others, and for many of them it is PPAD-hard or FIXP-hard to compute an approximate or exact equilibrium (e.g. [6, 8, 28, 7, 14, 15]).

Several researchers have proposed Hylland-Zeckhauser-type mechanisms for a number of applications, e.g. see [4, 19, 23, 24]. There are also recent works that have generalized and extended the basic HZ scheme in several directions, for example to two-sided matching markets and to an Arrow-Debreu-type setting where the agents own initial endowments [12, 13, 16]. Note that in the case of initial endowments, an HZ equilibrium may not always exist, so some approximation or slack is needed to ensure existence (see [12, 16]). Motivated by our hardness result, [20] propose a Nash Bargaining model for one-side matching market. Their mechanism is efficient, Pareto optimal and satisfies certain notion of fairness. However, as noted by [20], it does not guarantee envy-freeness and is not incentive compatible in large.

2 Preliminaries

We write $[n]$ to denote $\{1, 2, \ldots, n\}$. Given two integers n and m we use $[n : m]$ to denote integers between n and m, with $[n : m] = \emptyset$ when $m < n$. Given two real number $x, y \in \mathbb{R}$, we use $x = y \pm \epsilon$ to denote $x \in [y - \epsilon, y + \epsilon]$.

2.1 The Hylland-Zeckhauser Scheme We provide a formal description of the Hylland-Zeckhauser scheme for one-sided matching markets [21, 29]. It will be convenient for us to describe it using the language of linear Fisher markets. An *HZ market* M consists of a set $A = [n]$ of n agents and a set $G = [n]$ of n (infinitely) divisible goods. Each agent $i \in A$ has one dollar and there is one unit of each good $j \in G$ in the market. We write $u_{i,j} \in [0, 1]$[3] to denote the utility of one unit of good j to agent i, for each $i \in A$ and $j \in G$. Hence an HZ market M is specified by a positive integer n and utilities $(u_{i,j} : i, j \in [n])$.

Given an HZ market M with n agents and goods, an *HZ equilibrium* [21] consists of an allocation $x = (x_{i,j} : i, j \in [n])$ and a price vector $p = (p_j : j \in [n])$ that are nonnegative and satisfy a list of properties to be described in Definition 2.1. Given x and p, we will refer to $x_i = (x_{i,j} : j \in [n])$ as the *bundle of goods allocated to agent* i. The *cost* of the bundle x_i is given by $\sum_{j\in[n]} p_j x_{i,j}$ and the *value* of x_i to agent i is $\sum_{j\in[n]} u_{i,j} x_{i,j}$. We are ready to define HZ equilibria:

[3]As it will become clear in Definition 2.1, shifting and scaling utilities of agents does not change the set of HZ equilibria. We assume utilities to lie in $[0, 1]$ because we will consider an additive approximation of HZ equilibria in Definition 2.2.

DEFINITION 2.1. (HZ EQUILIBRIA [21]) *A pair* (x, p)*, where* $x = (x_{i,j} : i, j \in [n]) \in \mathbb{R}^{n\times n}_{\geq 0}$ *and* $p = (p_i : i \in [n]) \in \mathbb{R}^n_{\geq 0}$*, is an* HZ equilibrium *of an HZ market* M *if:*

1. *The total allocation of each good* $j \in [n]$ *is* 1 *unit, i.e.,* $\sum_{i\in[n]} x_{i,j} = 1$.
2. *The total allocation of each agent* $i \in [n]$ *is* 1 *unit, i.e.,* $\sum_{j\in[n]} x_{i,j} = 1$.
3. *The cost of the bundle* x_i *of each agent* $i \in [n]$ *is at most* 1*, i.e.,* $\sum_{j\in[n]} p_j x_{i,j} \leq 1$.
4. *For each* $i \in [n]$*,* x_i *maximizes its value* $\sum_{j\in[n]} u_{i,j}x_{i,j}$ *to agent* i *subject to 2 and 3 above.*[4]

Equivalently the last condition in the definition above can be captured by the following LP:

$$\begin{aligned} \textit{maximize} \quad & \sum_{j\in[n]} u_{i,j}x_{i,j} \\ \textit{s.t.} \quad & \sum_{j\in[n]} x_{i,j} = 1, \ \sum_{j\in[n]} p_j x_{i,j} \leq 1, \text{ and } x_{i,j} \geq 0 \text{ for all } j \in [n]. \end{aligned}$$

Taking μ_i and α_i to be the dual variables, one has the following dual LP that will be useful:

$$\begin{aligned} \textit{minimize} \quad & \alpha_i + \mu_i \\ \textit{s.t.} \quad & \alpha_i \geq 0 \text{ and } \alpha_i p_j + \mu_i \geq u_{i,j}, \text{ for all } j \in [n]. \end{aligned}$$

We will refer to the LP (and its dual LP) above as the LP (or dual LP) for agent i with respect to the price vector p. Let $\mathsf{value}_p(i)$ denote their optimal value. Then it captures the optimal value of any bundle of goods to agent i subject to conditions 2 and 3 in Definition 2.1.

Hylland and Zeckhauser [21] showed that an HZ equilibrium always exists:

THEOREM 2.1. (EXISTENCE [21]) *Every HZ market admits an HZ equilibrium.*

If (x, p) is an equilibrium, then it is easy to see that if we scale the difference of all prices from 1, the resulting price vector p' together with the same allocation x forms also an equilibrium; i.e. for any $r > 0$ with $r \leq \min\{1/(1 - p_j)|p_j < 1\}$, setting $p'_j = 1 + r(p_j - 1)$ for all $j \in [n]$ yields a vector p' such that (x, p') is also an equilibrium (see [29]). The reason is that this scaling does not affect the set of feasible allocations, as can be easily seen, and $\mathsf{value}_{p'}(i) = \mathsf{value}_p(i)$ for all agents $i \in [n]$. That is, price vectors related to each other by this scaling are in a sense equivalent. A consequence of this observation is that we may always assume w.l.o.g. that an equilibrium contains a good with price 0 [21]: If one of the goods has price < 1, then we can always scale the prices so that the minimum price is 0. On the other hand if all prices in an equilibrium are ≥ 1, then all prices must be 1 (the sum of the prices must be $\leq n$, the sum of the agents' budgets), and in this case the cost condition 3 is redundant, and the all-0 vector forms also an equilibrium with the same allocation.

We say that a price vector p is *normalized* if $\min_i p_i = 0$. We will restrict our attention henceforth to normalized price vectors, without always mentioning it explicitly.

Our hardness results hold for the following relaxation studied by Vazirani and Yannakakis [29]:

DEFINITION 2.2. (APPROXIMATE HZ EQUILIBRIA) *Given some* $\epsilon > 0$*, a pair* (x, p)*, where* $x = (x_{i,j} : i, j \in [n]) \in \mathbb{R}^{n\times n}_{\geq 0}$ *and* $p = (p_i : i \in [n]) \in \mathbb{R}^n_{\geq 0}$ *(where* $\min_{i\in[n]} p_i = 0$*)*[5]*, is an* ϵ-approximate HZ equilibrium *of an HZ market* M *if:*

1. *The total allocation of each good* $j \in [n]$ *is* 1 *unit, i.e.,* $\sum_{i\in[n]} x_{i,j} = 1$.

[4]We note that in [21, 29], x_i is required (as a tie-breaking rule) to minimize its cost among all those that maximize the value subject to items 2 and 3. This is needed to ensure Pareto optimality of the equilibrium allocations. However, we do not need this condition for our hardness results, and this only makes the results stronger. So for simplicity, we omit the condition from the definition of exact and approximate equilibria.

[5]The requirement that p be normalized is important in the definition because otherwise condition 3 on the cost has no effect: if (x, p) is any pair that satisfies conditions 1,2,4, then we can always scale p as above to a vector p' where all prices are sufficiently close to 1 so that condition 3 is also satisfied for (x, p').

2. *The total allocation of each agent $i \in [n]$ is 1 unit, i.e., $\sum_{j\in[n]} x_{i,j} = 1$.*

3. *The cost of x_i is at most $1+\epsilon$ for each $i \in [n]$, i.e., $\sum_{j\in[n]} p_j x_{i,j} \le 1+\epsilon$.*

4. *The value $\sum_{j\in[n]} u_{i,j} x_{i,j}$ of x_i to agent i is at least $\mathsf{value}_p(i) - \epsilon$ for each $i \in [n]$.*

An alternative, more relaxed notion of an ϵ-approximate equilibrium, where condition 1 is also relaxed to $|\sum_{i\in[n]} x_{i,j} - 1| \le \epsilon$ for all goods $j \in [n]$, is polynomially equivalent to the above notion [29]. Thus, it follows that computing an ϵ-approximate equilibrium under the more relaxed notion is also PPAD-complete.

2.2 Threshold Games Our PPAD hardness results use *threshold games*, introduced recently by Papadimitriou and Peng [27]. They showed that finding an approximate equilibrium in a threshold game is PPAD-complete.

DEFINITION 2.3. (THRESHOLD GAME [27]) *A threshold game is defined over a directed graph $H = (V, E)$. Each node $v \in V$ represents a player with strategy space $x_v \in [0, 1]$. Let N_v be the set of nodes $u \in V$ with $(u, v) \in E$. Then $x = (x_v : v \in V) \in [0, 1]^V$ is a* κ-approximate equilibrium *if every x_v satisfies*

$$x_v \in \begin{cases} [0,\kappa] & \sum_{u\in N_v} x_u > 0.5+\kappa \\ [1-\kappa,1] & \sum_{u\in N_v} x_u < 0.5-\kappa \\ [0,1] & \sum_{u\in N_v} x_u \in [0.5-\kappa, 0.5+\kappa] \end{cases}$$

THEOREM 2.2. (THEOREM 4.7 OF [27]) *There is a positive constant κ such that the problem of finding a κ-approximate equilibrium in a threshold game is PPAD-hard. This holds even when the in-degree and out-degree of each node is at most 3 in the threshold game.*

3 PPAD-hardness

Our goal in this section is to prove the following theorem:

THEOREM 3.1. *The problem of finding a $(1/n^5)$-approximate HZ equilibrium in an HZ market with n agents and goods is PPAD-hard.*

In the full version [5] (via a padding argument), we give a polynomial-time reduction from the problem of finding a $(1/n^5)$-approximate HZ equilibrium to that of finding a $(1/n^c)$-approximate HZ equilibrium in an HZ market, for any positive constant c. Theorem 1.1 follows by combining the PPAD membership result of [29].

Our plan is as follows. Let $\epsilon = 1/n^5$ throughout this section wherever an HZ market with n agents and goods is concerned. We start with some basic facts about approximate HZ equilibria in Section 3.1 (mainly about how to work with approximately optimal bundles for agents). Then we describe the polynomial-time reduction from threshold games to HZ markets in Section 3.2. Our reduction constructs two types of gadgets, *variable gadgets* and *edge gadgets*, which simulate variables x_v and edges (u, v) in a threshold game, respectively. Using these gadgets, we finish the reduction's correctness proof in Section 3.3; the analysis of these two gadgets is presented afterwards in Section 3.4 and the full version of the paper [5], respectively.

3.1 Basic Facts. Let M be an HZ market with n agents and goods. As it will become clear later, the HZ market we construct in the reduction satisfies $\max_{j\in[n]} u_{i,j} = 1$ for every agent $i \in [n]$. *Hence we assume this is the case in every HZ market discussed in the rest of this section.* In all lemmas of this subsection we assume (x, p) to be an ϵ-approximate HZ equilibrium of M (and skip it in their statements). Recall that prices are normalized: $\min_i p_i = 0$.

We give first an upper bound on the sum of prices:

LEMMA 3.1. $\sum_{j\in[n]} p_j \le 2n$.

Proof. Since (x, p) is an ϵ-approximate HZ equilibrium, every good must be sold out and no agent can spend more than $1+\epsilon$. Thus, $\sum_{j\in[n]} p_j \le n(1+\epsilon) < 2n$ using $\epsilon = 1/n^5$. □

Next, we consider an optimal solution (α_i^*, μ_i^*) to the dual LP for agent i and prove the following:

LEMMA 3.2. *$\mu_i^* \geq 0$ and $\alpha_i^* \leq 1$ for every $i \in [n]$.*

Proof. Let ℓ be a good with $p_\ell = 0$. From the dual LP constraints, we have $0 \leq u_{i,\ell} \leq \alpha_i^* p_\ell + \mu_i^* = \mu_i^*$. Moreover, since all utilities are in $[0, 1]$ we have trivially that $\alpha_i^* + \mu_i^* = \mathsf{value}_p(i) \leq 1$. Therefore, we have $\alpha_i^* \leq 1$. ∎

LEMMA 3.3. *If $\mathsf{value}_p(i) \leq 0.9$ then $\alpha_i^* \geq 1/(20n)$ and $\sum_{j\in[n]} p_j x_{i,j} \geq 1 - 20n\epsilon$.*

Proof. Let $u_{i,\ell} = 1$. It follows from Lemma 3.1 that $p_\ell \leq 2n$ and thus,

$$1 = u_{i,\ell} \leq \alpha_i^* p_\ell + \mu_i^* \leq 2n\alpha_i^* + \mu_i^*.$$

On the other hand, $\mathsf{value}_p(i) = \alpha_i^* + \mu_i^* \leq 0.9$. The first part of the lemma follows from adding these two inequalities.

Next, multiplying both sides of the inequalities $\alpha_i^* p_j + \mu_i^* \geq u_{i,j}$ by $x_{i,j}$, summing over all $j \in [n]$, and using $\sum_j x_{i,j} = 1$, we have

$$\sum_{j\in[n]} \alpha_i^* p_j x_{i,j} + \mu_i^* \geq \sum_{j\in[n]} u_{i,j} x_{i,j} \geq \alpha_i^* + \mu_i^* - \epsilon.$$

The second part of the lemma then follows from $\alpha_i^* \geq 1/(20n)$. ∎

Recall that all goods j satisfy $u_{i,j} \leq \alpha_i^* p_j + \mu_i^*$. We say a good j is δ-*suboptimal* for agent i if $u_{i,j} + \delta \leq \alpha_i^* p_j + \mu_i^*$. We show that agent i's good-bundle, x_i, cannot contain significant quantities of suboptimal goods.

LEMMA 3.4. *For every $i \in [n]$, the total allocation in x_i to δ-suboptimal goods is at most $2\epsilon/\delta$.*

Proof. Fix an agent $i \in [n]$. We have $u_{i,j} \leq \alpha_i^* p_j + \mu_i^*$ for all $j \in [n]$, and $u_{i,j} + \delta \leq \alpha_i^* p_j + \mu_i^*$ for δ-suboptimal goods. Let W be the total allocation in x_i to δ-suboptimal goods. Then

$$\sum_{j\in[n]} u_{i,j} x_{i,j} + W\delta \leq \alpha_i^* \sum_{j\in[n]} p_j x_{i,j} + \mu_i^* \sum_{j\in[n]} x_{i,j}.$$

Using the definition of ϵ-approximate HZ equilibria, the LHS is at least

$$\mathsf{value}_p(i) - \epsilon + W\delta = \alpha_i^* + \mu_i^* + W\delta - \epsilon$$

and the RHS is at most $\alpha_i^*(1 + \epsilon) + \mu_i^*$. The lemma follows from $\alpha_i^* \leq 1$ by Lemma 3.2. ∎

We use some of the lemmas above to obtain following corollaries:

COROLLARY 3.1. *Let J be the set of $j \in [n]$ that are* not *δ-suboptimal for i. If $\mathsf{value}_p(i) \leq 0.9$, then*

$$1 - 2\epsilon/\delta \leq \sum_{j\in J} x_{i,j} \leq 1 \quad \textit{and} \quad 1 - 20n\epsilon - \frac{4n\epsilon}{\delta} \leq \sum_{j\in J} p_j x_{i,j} \leq 1 + \epsilon.$$

Proof. The first part follows directly from Lemma 3.4.

The second part follows from Lemma 3.1, Lemma 3.3, and Lemma 3.4. ∎

COROLLARY 3.2. *Let J be the set of goods $j \in [n]$ with $u_{i,j} > 0$. If $\mathsf{value}_p(i) \leq 0.9$, then*

$$1 - 20n\epsilon - \frac{1}{n^2} \leq \sum_{j\in J} p_j x_{i,j} \leq 1 + \epsilon.$$

Proof. The second inequality clearly holds since (x, p) is an ϵ-approximate equilibrium. Suppose that the first inequality does not hold. Then by Lemma 3.3, agent i spends more than $1/n^2$ on zero-utility goods, hence she buys at least an amount $1/2n^3$ of these, since all prices are at most $2n$. Consider a new bundle for i obtained by replacing $1/2n^3$ of the zero-utility goods by a good with utility 1. The cost of the new bundle is still less than 1, i.e. it is a feasible bundle, and the value exceeds that of the original bundle x_i by $1/2n^3 > \epsilon$, contradicting the fact that (x, p) is an ϵ-approximate equilibrium. ∎

Finally we include a simple lemma about the optimal value of an agent:

LEMMA 3.5. *Let $i \in [n]$ and $\ell \in [n]$ with $u_{i,\ell} = 1$. Then $\mathsf{value}_p(i) \geq \min(1, 1/p_\ell)$.*

Proof. If $p_\ell = 0$, then agent i can get value 1 by buying one unit of good ℓ for free.

If $p_\ell > 0$ then there is another good with zero price and thus, agent i can get value $\min(1, 1/p_\ell)$ by buying $\min(1, 1/p_\ell)$ unit of good ℓ and $1 - \min(1, 1/p_\ell)$ unit of a zero price good. ∎

3.2 The Construction. Let $\kappa \in (0,1)$ be the positive constant in Theorem 2.2. Recall that our goal is to give a polynomial-time reduction from the problem of finding a κ-approximate equilibrium in a threshold game (with both in-degree and out-degree at most 3) to that of finding an ϵ-approximate HZ equilibrium in an HZ market with $\epsilon = 1/n^5$.

Let C be a sufficiently large universal constant, and $m = \lceil C/\kappa \rceil$. Let $H = (V, E)$ be a threshold game with $|V| = N$. (Note that N is asymptotically large and should be considered as larger than any function of m.) We write in-deg(v) and out-deg$(v) \le 3$ to denote the in-degree and out-degree of $v \in V$, respectively. We construct an HZ market M_H from H in three steps as described below. This is done by creating *groups* of goods and *groups* of agents, with the guarantee that agents in the same group have the same utility for any good as each other, and that goods in the same group yield the same utility to any agent. We say a group A_i of agents have utility u for a group G_j of goods if all agents in A_i share the same utility u for all goods in G_j. (Intuitively we create a group A_i of agents to simulate an agent with demand and budget $|A_i|$ instead of 1, and a group G_j of goods to simulate a good with a supply of $|G_j|$ units in the market. A technical subtlety though is that in an approximate HZ equilibrium, goods in the same group may not share exactly the same price and agents in the same group may not have exactly the same allocation.)

Step 1: Creating Variable Gadgets. We start with an empty market and create a *variable gadget* for each node $v \in V$ to simulate the variable x_v in the threshold game H. For each node $v \in V$, the variable gadget of v consists of the following three groups of goods and one group of agents:

1. Create three groups of goods $G_{v,1}, G_{v,2}$ and $G_{v,3}$: $G_{v,1}$ has $m^{10} + S_u$ goods, where

$$S_u := (24m^3 + 12m) \cdot \text{out-deg}(u) + (24m^3 + 15m) \cdot \text{in-deg}(u) - 3m,$$

   and $G_{v,2}$ and $G_{v,3}$ both have $2m^{10}$ goods. Let G_v denote the union of $G_{v,1}, G_{v,2}$ and $G_{v,3}$.

2. Create a group A_v of $5m^{10}$ agents. Each agent in A_v has the following utilities for G_v:

$$\frac{1}{2m^2 - 1} \text{ for } G_{v,1}, \quad \frac{m^2+1}{4m^2-2} \text{ for } G_{v,2}, \quad 1 \text{ for } G_{v,3}, \tag{3.1}$$

   and utility 0 for every other good in the market (including those created later).

Looking ahead, we will prove (in Lemma 3.6) that in any ϵ-approximate HZ equilibrium (x, p) of the final HZ market M_H, $p(G_{v,1}), p(G_{v,2})$ and $p(G_{v,3})$ must satisfy

$$0 \le p(G_{v,1}) \lesssim \frac{1}{m^2}, \quad p(G_{v,2}) \approx \frac{1 + p(G_{v,1})}{2} \quad \text{and} \quad p(G_{v,3}) \approx 2 - p(G_{v,1}),$$

where $p(G_{v,\ell})$ denotes the minimum price of goods in $G_{v,\ell}$. Indeed, $p(G_{v,1})$ will be used to simulate the variable x_v in the threshold game H and at the end, we set $x_v \approx m^2 p(G_{v,1})$ for each $v \in V$ to obtain a κ-approximate equilibrium of H.

Step 2: Creating Edge Gadgets. Next we create an *edge gadget* for each edge $e = (u, v) \in E$ to simulate the action of vertex u sending a contribution x_u to the summation at vertex v in the threshold game H (see definition 2.3). For each (directed) edge $e = (u, v) \in E$, the edge gadget of e consists of the following multiple groups of goods and agents (for convenience, we only list goods with positive utilities for each group of agents; every other good has utility 0):

1. Create a group G_e of $32m^5$ goods.
2. Create a group $A_{e,*}$ of $64m^5$ agents. They have utility 1 for G_e.
3. Create a group of $48m^3$ agents $A_{e,1}$. They have utility 1 for $G_{u,3}$ and $1/2$ for $G_{v,1}$.
4. Create m groups $A_{e,2,\ell}$, $\ell \in [m]$, each of 6 agents. They have 1 for G_e and $\ell/(2m^3)$ for $G_{v,1}$.

5. Create m groups $A_{e,3,\ell}$, $\ell \in [m]$, each of 8 agents. They have 1 for G_e and $\ell/(2m^3)$ for $G_{u,1}$.

6. Create m groups $A_{e,4,\ell}$, $\ell \in [2m]$, each of 18 agents. They have 1 for G_e, $\ell/(2m^3)$ for $G_{v,1}$,
$$\frac{1}{4} + \frac{1}{4m^2} + \frac{1}{m^3}$$
for goods in $G_{u,2}$.

For convenience we write A_e to denote the union of groups $A_{e,1}, A_{e,2,\ell}, A_{e,3,\ell}$ and $A_{e,4,\ell}$, for all ℓ.

Step 3: Adding Dummy Goods. So far we have created
$$5m^{10} \cdot |V| + (64m^5 + 48m^3 + 50m) \cdot |E| \tag{3.2}$$
many agents and
$$\sum_{u \in V} (5m^{10} + S_u) + 32m^5 \cdot |E| = (5m^{10} - 3m) \cdot |V| + (32m^5 + 48m^3 + 27m) \cdot |E|$$
many goods. To finish the construction (since the number of goods needs to match that of agents), we create a group of $3m|V| + (32m^5 + 23m)|E|$ dummy goods, which have utility 0 to every agent in the market. This finishes the construction of M_H with n agents and goods, where n is given in (3.2). It is clear that M_H can be built in polynomial time.

Before moving forward, we record a list of simple properties about M_H:

FACT 3.1. *The HZ market M_H satisfies the following properties:*

1. *Every agent in the market has maximum utility* 1*;*

2. *For each node $v \in V$, the number of agents outside of A_v that have a positive utility on at least one group of goods in G_v is at most $288m^3 + 258m = O(m^3)$;*

3. *For each edge $e \in E$, the number of agents outside of $A_{e,*}$ that have a positive utility on G_e is $50m = O(m)$.*

3.3 Proof of Correctness. Let $\epsilon = 1/n^5$. We prove two lemmas about variable gadgets in Section 3.4.

We use $p(G_i)$ to denote the minimum price of goods in a group G_i. The first lemma shows that $p(G_{v,1})$ is between (roughly) 0 and $1/m^2$ and it determines $p(G_{v,2})$ and $p(G_{v,3})$ (approximately).

LEMMA 3.6. *Let (x, p) be an ϵ-approximate HZ equilibrium of M_H. Then $p(G_{v,1})$ satisfies*
$$0 \le p(G_{v,1}) \le \frac{1}{m^2} + O\left(\frac{1}{m^6}\right)$$
for every $v \in V$. Moreover, $p(G_{v,2})$ and $p(G_{v,3})$ satisfy
$$p(G_{v,2}) = \frac{1 + p(G_{v,1})}{2} \pm O\left(\frac{1}{m^7}\right) \quad \text{and} \quad p(G_{v,3}) = 2 - p(G_{v,1}) \pm O\left(\frac{1}{m^7}\right). \tag{3.3}$$

We prove Lemma 3.6 in Section 3.4.

We next show that the variable gadget created for each node $v \in V$ is sensitive to demand from agents outside of A_v. To state the lemma (and the next one), we introduce the following notation: Let G^* be a subset of goods (which could be a group or the union of multiple groups of goods) and A^* be a subset of agents in M_H (which could be a group or the union of multiple groups). We let
$$x^+(G^*, A^*) = \sum_{\substack{i \in A^* \\ j \in G^*: \\ u_{i,j} > 0}} x_{i,j},$$

i.e., the total allocation of G^* to A^* but limited to those goods in G^* with positive utilities to each agent in A^* only. We also write $\overline{A}_v$ to denote all agents in M_H outside of A_v.

The second lemma (which we also prove in Section 3.4) states that if the total allocation of G_v to agents outside of A_v with positive utilities is either more than $S_v + 1$ or less than $S_v - 1$, then $p(G_{v,1})$ must be at one of the two extreme cases accordingly, i.e., either close to 0 or close to $1/m^2$.

Lemma 3.7. *Let (x, p) be an ϵ-approximate HZ equilibrium of M_H. Then for every $v \in V$:*

1. *If $x^+(G_v, \overline{A}_v) \geq S_v + 1$, then we have*

$$p(G_{v,1}) = \frac{1}{m^2} \pm O\left(\frac{1}{m^9}\right);$$

2. *If $x^+(G_v, \overline{A}_v) \leq S_v - 1$, then we have $p(G_{v,1}) \leq O(1/n^2)$.*

Finally we prove the following lemma about M_H, which follows from a detailed analysis of the edge gadgets; the proof can be found in the full version [5]:

Lemma 3.8. *Let (x, p) be an ϵ-approximate HZ equilibrium of M_H. For each $e = (u, v) \in E$,*

$$x^+(G_u, A_e) = 24m^3 + 12m \pm O(1) \quad \text{and} \quad x^+(G_v, A_e) = -6m^3 p(G_{u,1}) + 24m^3 + 15m \pm O(1).$$

We now use these lemmas to prove Theorem 3.1:

Proof. [Proof of Theorem 3.1 assuming Lemmas 3.6, 3.7 and 3.8] Let $H = (V, E)$ be a threshold game, and let (x, p) be an ϵ-approximate HZ equilibrium of M_H. Let $(x_v : v \in V)$ be a profile for H with

$$x_v = \min\big(1, m^2 p(G_{v,1})\big)$$

for each $v \in V$. We prove below that $(x_v : v \in V)$ is a κ-approximate equilibrium of H.

Fix a node $v \in V$. We consider two cases.

1. Case 1: $\sum_{u \in N_v} x_u > 0.5 + \kappa$. In this case, $x^+(G_v, \overline{A}_v)$ is at most

$$\begin{aligned} &\text{out-deg}(v) \cdot \big(24m^3 + 12m + O(1)\big) + \sum_{u \in N_v} \big(24m^3 + 15m - 6m^3 p(G_{u,1}) + O(1)\big) \\ &\quad = S_v + 3m - 6m^3 \sum_{v \in N_v} p(G_{u,1}) + O(1) < S_v - 1. \end{aligned}$$

   It follows from Lemma 3.7 that $p(G_{v,1}) \leq O(1/n^2)$ and thus, $x_v \leq O(m^2/n^2) < \kappa$.

2. Case 2: $\sum_{u \in N_v} x_u < 0.5 - \kappa$. Using $p(G_{u,1}) \leq 1/m^2 + O(1/m^9)$, we have

$$\sum_{u \in N_v} m^2 p(G_{u,1}) < 0.5 - \kappa + O(1/m^7).$$

   Similarly, $x^+(G_v, \overline{A}_v)$ is at least

$$\begin{aligned} &\text{out-deg}(v) \cdot \big(24m^3 + 12m - O(1)\big) + \sum_{u \in N_v} \big(24m^3 + 15m - 6m^3 p(G_{u,1}) - O(1)\big) \\ &\quad = S_v + 3m - 6m^3 \sum_{u \in N_v} p(G_{u,1}) - O(1) > S_v + 1. \end{aligned}$$

   It follows from Lemma 3.7 that $p(G_{v,1}) \geq (1/m^2) - O(1/m^9)$ and thus, $x_v \geq 1 - \kappa$.

This finishes the proof of the theorem. ∎

3.4 Analysis of Variable Gadgets

We prove Lemma 3.6 and Lemma 3.7 in this section. We start with some simple bounds on prices of goods in G_e and $G_{v,3}$, $e \in E$ and $v \in V$:

LEMMA 3.9. *Let (x, p) be an ϵ-approximate HZ equilibrium of M_H. We have $p(G_e) \geq 2(1 - 2\epsilon)$ for every $e \in E$ and $p(G_{v,3}) \geq 5/3$ for every $v \in V$.*

Proof. Fix an $e \in E$. The optimal value of each agent in $A_{e,*}$ is at most $0.5 + \epsilon$; otherwise each of them must receive a bundle with value more than 0.5, which implies that each of them gets more than 0.5 unit of goods in G_e, contradicting with the fact that there are $64m^5$ many agents in $A_{e,*}$ but only $32m^5$ many goods in G_e. On the other hand, the optimal value of each agent in $A_{e,*}$ is at least $\min(1, 1/p(G_e))$ by Lemma 3.5 and thus, $p(G_e) \geq 1/(0.5 + \epsilon) > 2(1 - 2\epsilon)$.

Next fix a $v \in V$. With a similar argument, the optimal value of each agent in A_v is at most

$$\frac{1}{5m^{10}} \cdot \left(\frac{m^{10} + S_v}{2m^2 - 1} + 2m^{10} \cdot \frac{m^2 + 1}{4m^2 - 2} + 2m^{10} \right) + \epsilon < 3/5$$

when m is sufficiently large. On the other hand, by Lemma 3.5 the optimal value of each agent in A_v is at least $\min(1, 1/p(G_{v,3}))$ and thus, $p(G_{v,3}) \geq 5/3$. ∎

From this we can show that every agent in M_H has optimal value at most 0.9:

LEMMA 3.10. *Let (x, p) be an ϵ-approximate HZ equilibrium of M_H. Then every agent in M_H has optimal value (with respect to p) at most* 0.9.

Proof. As shown in the previous lemma, the optimal value of each agent in a group A_v of a variable gadget is at most $3/5$, and the optimal value of each agent in a group $A_{e,*}$ of an edge gadget is at most $0.5 + \epsilon$. The claim for the agents in the groups A_e follows from the prices of the goods in $G_{u,3}$ and G_e, which are the goods that have utility 1 for these agents (the other goods have utility $1/2$ or less). ∎

This allows us to apply lemmas in Section 3.1. It immediately leads to the following corollary:

COROLLARY 3.3. *For every group G_j of goods, the maximum price in G_j is at most $p(G_j) + 1/n^2$.*

Proof. Assume for a contradiction that there is a good in G_j with price at least $p(G_i) + 1/n^2$. Then for each agent i in the market, we have $\alpha_i^* \geq 1/(20n)$ by Lemma 3.3 and thus this good is $\Omega(1/n^3)$-suboptimal (by comparing with the good in G_j with price $p(G_j)$). Hence its allocation to agent i is $O(1/n^2)$, and the total allocation of this good in x is $O(1/n)$, a contradiction. ∎

Before proving Lemma 3.6 we show that $p(G_e)$ is very close to 2:

LEMMA 3.11. *For every edge $e \in E$ we have $p(G_e) = 2 \pm O(1/m^4)$.*

Proof. We have by Lemma 3.9 that $p(G_e) \geq 2 - O(\epsilon)$. For the upper bound note that by Lemma 3.9 and Lemma 3.3, goods in G_e are $\Omega(1/n)$-suboptimal to agents with zero utility so their total allocation to such agents is $O(n^2\epsilon)$ by Lemma 3.4. By Fact 3.1 the total allocation of G_e to agents outside $A_{e,*}$ with a positive utility is $O(m)$ and thus, the rest of $32m^5 - O(m)$ units of G_e are allocated to agents in $A_{e,*}$. So

$$\big(32m^5 - O(m)\big)\, p(G_e) \leq 64m^5(1 + \epsilon),$$

which implies that $p(G_e) \leq 2 + O(1/m^4)$. This finishes the proof of the lemma. ∎

We are now ready to prove Lemma 3.6:

Proof. [Proof of Lemma 3.6] Fixing any node $v \in V$, we let q_ℓ denote $p(G_{v,\ell})$ and y_ℓ to denote the total allocation of $G_{v,\ell}$ to agents in A_v in x, for each $\ell \in \{1, 2, 3\}$. We also write u_ℓ to denote the utility of $G_{v,\ell}$ to agents in A_v given in (3.1). We start by showing that most goods in G_v go to A_v.

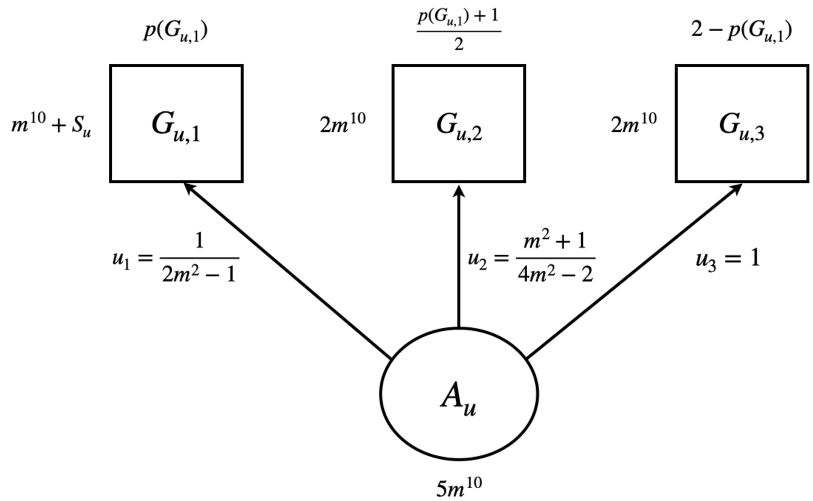

Figure 1: The variable gadget.

LEMMA 3.12. *We have* $y_1 \geq m^{10} - O(m^3)$ *and* $y_2, y_3 \geq 2m^{10} - O(m^3)$.

Proof. Let α^* and μ^* be an optimal solution to the dual LP of agents in A_v. Then $\alpha^* q_\ell + \mu^* \geq u_\ell$ for each ℓ. We consider the following two cases.

First we consider the case when $\mu^* \geq u_1/2 = \Omega(1/m^2)$. This implies that goods outside of G_v are $\Omega(1/m^2)$-suboptimal for A_v and thus, by Lemma 3.4, the total allocation of them to agents in A_v is $O(m^{10}) \cdot O(m^2\epsilon) < 1$. As a result, $y_1 + y_2 + y_3 \geq 5m^{10} - 1$ from which the claim follows.

Next consider the case when $\mu^* < u_1/2$. By Lemma 3.2 ($\alpha^* \leq 1$) we have $q_\ell \geq \Omega(1/m^2)$ for every ℓ. This implies that agents with zero utilities to G_v can be allocated only $n \cdot O(nm^2\epsilon) < 1$ units of G_v given that they are $\Omega(1/nm^2)$-suboptimal by Lemma 3.3. On the other hand, by Fact 3.1 the allocation to agents outside A_v with positive utilities for G_v is at most $O(m^3)$. So the rest of G_v must be allocated to A_v and the claim follows. ∎

Now that we have $y_\ell \geq m^{10} - O(m^3)$ for all $\ell \in \{1, 2, 3\}$, we proceed to prove (3.3). Let (α^*, μ^*) denote an optimal solution to the dual LP for A_v. By Lemma 3.4 and taking $\delta = 20\epsilon$, we have

$$u_\ell \leq \alpha^* q_\ell + \mu^* \leq u_\ell + \delta, \quad \text{for all } \ell \in \{1, 2, 3\}. \tag{3.4}$$

If this were not the case (i.e. the second inequality is violated for some ℓ), then goods in $G_{v,\ell}$ are δ-suboptimal to A_v and their total allocation to agents in A_v can be no more than $5m^{10} \cdot 2\epsilon/\delta = m^{10}/2$, a contradiction.

Combining (3.4) and $u_2 = (3u_1 + u_3)/4$, we have

$$\alpha^* \left(\frac{3q_1 + q_3}{4}\right) + \mu^* - \delta \leq \alpha^* q_2 + \mu^* \leq \alpha^* \left(\frac{3q_1 + q_3}{4}\right) + \mu^* + \delta.$$

Using $\alpha^* \geq 1/(20n)$ from Lemma 3.3, we have

$$q_2 = \frac{3q_1 + q_3}{4} \pm O(n\epsilon). \tag{3.5}$$

Next, using Corollary 3.2 and Corollary 3.3 we have

$$5m^{10}(1 - O(1/n^2)) \leq q_1 y_1 + q_2 y_2 + q_3 y_3 \leq 5m^{10}(1 + \epsilon + 1/n^2).$$

Plugging in $y_1 = m^{10} \pm O(m^3)$ and $y_2, y_3 = 2m^{10} \pm O(m^3)$ and (3.5), we have $q_1 + q_3 = 2 \pm O(1/m^7)$. Together with (3.5) again we obtain

$$q_3 = 2 - q_1 \pm O(1/m^7) \quad \text{and} \quad q_2 = (1 + q_1)/2 \pm O(1/m^7).$$

Finally we give an upper bound on q_1. We first note that $q_1 < q_3$; otherwise goods in $G_{v,1}$ are $\Omega(1)$-suboptimal to agents in A_v, contradicting with $y_1 = m^{10} \pm O(m^3)$. Using (3.4) we have

$$\mu^* \le \frac{u_1 q_3 - u_3 q_1 + O(\delta)}{q_3 - q_1}.$$

But when $q_1 \ge 1/m^2 + 1/m^6$ (and thus, $q_3 \le 2 - 1/m^2$), the nominator of the RHS is

$$u_1 q_3 - u_3 q_1 \le \frac{1}{2m^2 - 1} \cdot \left(2 - \frac{1}{m^2}\right) - 1 \cdot \left(\frac{1}{m^2} + \frac{1}{m^6}\right) \le -\frac{1}{m^6}.$$

So we have $\mu^* < 0$, in contradiction with $\mu^* \ge 0$ by Lemma 3.2. ◻

Next we prove Lemma 3.7:

Proof. [Proof of Lemma 3.7] We use the same notation from the proof of the last lemma.

First given that q_2 and q_3 are $\Omega(1)$, the total allocation of $G_{v,2}$ to agents with zero utility on $G_{v,2}$ is at most $n \cdot O(n\epsilon)$ using Lemma 3.4; the same applies to $G_{v,3}$.

Suppose that $x^+(G_v, A_v) \le S_v - 1$. Because G_v contains $5m^{10} + S_v$ goods while A_v contains only $5m^{10}$ agents, for G_v to be fully sold out, the total allocation of $G_{v,1}$ to agents with zero utility on $G_{v,1}$ must be $1 - o_n(1)$. This implies that $q_1 \le 1/n^2$ since otherwise, the total allocation for $G_{v,1}$ is at most $n \cdot O(n^3\epsilon) = o_n(1)$, using $\epsilon = 1/n^5$.

Next, suppose $x^+(G_v, \overline{A}_v) \ge S_v + 1$. Given that there are $5m^{10} + S_v$ goods in G_v and $5m^{10}$ agents in A_v, there must be an agent in A_v who is allocated at least $1/(5m^{10})$-unit of goods outside of G_v (for which it has zero utility). Since such goods are μ^*-suboptimal, we have $\mu^* \le 5m^{10}\epsilon$. On the other hand, recall (3.4) with $\delta = 20\epsilon$. We have $\alpha^*(q_1 + q_3) + 2\mu^* = u_1 + u_3 \pm 2\delta$ and thus,

$$\alpha^* = \frac{u_1 + u_3}{2}\left(1 \pm O\left(\frac{1}{m^7}\right)\right)$$

using Lemma 3.6. Then $q_1 = (u_1 - \mu^* \pm \delta)/\alpha^* = 1/m^2 \pm O(1/m^9)$. ◻

4 Hardness of Approximating Optimal Social Welfare

In this section we study the problem of approximating the optimal social welfare (defined as the total utility of all agents) achievable by an HZ equilibrium. For this purpose we study the following gap problem for a constant $\rho < 1$: the input is an HZ market M together with a parameter SW, and it is promised that the optimal social welfare achievable by an exact HZ equilibrium of M is either at least SW or at most $\rho \cdot \mathsf{SW}$. The goal is to tell which case it is. We show that there is no polynomial-time algorithm for the gap problem when $\rho > 175/176$, assuming $\mathsf{NP} \ne \mathsf{P}$.

THEOREM 4.1. *Assuming* $\mathsf{NP} \ne \mathsf{P}$, *for any constant* $\epsilon > 0$, *there is no polynomial-time algorithm for the gap problem when* $\rho = (175/176) + \epsilon$.

4.1 Construction.

We reduce from MAX 3SAT, which is hard to approximate better than 7/8 [18]: Given a 3SAT instance, it is NP-hard to distinguish the case that the formula is satisfiable from the case that every truth assignment satisfies at most a fraction $\frac{7}{8} + \epsilon$ of the clauses, for any $\epsilon > 0$. Given a 3SAT instance with m clauses and n variables, we construct the following HZ market. Throughout the proof, we fix $K = m^3$.

Creating Variable Gadget. We first introduce the variable gadget. For convenience, we only list non zero utilities. For each $i \in [n]$

1. Create three groups of goods $G_{i,1}$, $G_{i,2}$, $G_{i,3}$, and $|G_{i,1}| = K$, $|G_{i,2}| = 2K$ and $|G_{i,3}| = K$.
2. Create two groups of agents $A_{i,1}$, $A_{i,2}$, and $|A_{i,1}| = |A_{i,2}| = 2K$.
3. Agents in $A_{i,1}$ have utility $\frac{1}{2K^2}$ for $G_{i,1}$, $\frac{1}{K^2}$ for $G_{i,2}$. Agents in $A_{i,2}$ have utility $\frac{1}{2K^2}$ for $G_{i,3}$, $\frac{1}{K^2}$ for $G_{i,2}$.

In an (exact) HZ equilibrium, all goods within a group have the same price. We use $p(G_{i,\ell})$ to denote the price, $\ell \in [3]$.

Creating Clause Gadget. We next construct clause gadgets. For each $j \in [m]$,

1. Create a group G_j of K goods
2. Create a group $A_{j,*}$ of $2K$ agents, who have utility $1/K^2$ for G_j.
3. Create an agent A_j with utility 1 for G_j. It has utility $5/6$ for $G_{i,1}$ if the j-th clause contains x_i and utility $5/6$ for $G_{i,3}$ if the j-th clause contains $\overline{x}_i$.

Adding Dummy Goods Thus far, we have described $4Kn + Km$ goods and $4Kn + (2K+1)m$ agents. We add $(K+1)m$ extra dummy goods that have zero utilities for all agents. In a normalized (exact) HZ equilibrium, these goods have zero price.

4.2 Proof of Correctness.

We provide the proof of completeness and soundness separately.

Completeness Given a 3SAT instance that has a satisfying assignment, we construct a HZ equilibrium with social welfare at least $11m/12 - O(1/m^2)$. Fix a satisfying assignment.

We assign the j-th clause to the $\phi(j)$-th variable, if the latter satisfies the clause. If there are multiple such variables, we choose an arbitrary one. We set $\ell(j) = 1$ if the j-th clause contains $x_{\phi(j)}$, otherwise $\ell(j) = 3$. Let s_i be the total number of clauses assigned to the i-th variable. The equilibrium prices are as follows.

1. The price of dummy goods is 0.
2. The price of G_j is $p(G_j) = \frac{2K+1}{K}$, $j \in [m]$.
3. For variable gadget $i \in [n]$, if $x_i = 1$, then $(p(G_{i,1}), p(G_{i,3}), p(G_{i,2})) = (0, \frac{8}{5}, \frac{4}{5})$, otherwise, we have $(p(G_{i,1}), p(G_{i,2}), p(G_{i,3})) = (\frac{4}{5}, \frac{8}{5}, 0)$.

Next, we specify the equilibrium allocation.

1. Agents of $A_{j,*}$ take $\frac{2K^2}{2K+1}$ of G_j and $\frac{2K^2+2K}{2K+1}$ of dummy goods, $j \in [m]$.
2. Agent A_j takes $\frac{K}{2K+1}$ of G_j and $\frac{K+1}{2K+1}$ of $G_{\phi(j),\ell(j)}$, $j \in [m]$.
3. If $x_i = 1$, then agents in $A_{i,1}$ obtain $\frac{5K}{4}$ of $G_{i,2}$, $\frac{3K}{4}$ of $G_{i,1}$; agents in $A_{i,2}$ obtain $\frac{3K}{4}$ of $G_{i,2}$, K of $G_{i,3}$, $\frac{K}{4} - s_j \cdot \frac{K+1}{2K+1}$ of $G_{i,1}$ and $s_j \cdot \frac{K+1}{2K+1}$ of the dummy good, $i \in [n]$. If $x_i = 0$, then we define the allocation symmetrically, switching the groups of agents $A_{i,1}$ and $A_{i,2}$, and the groups of goods $G_{i,1}$ and $G_{i,3}$.

One can verify that this is indeed a HZ equilibrium. Agent A_j has utility $\frac{K}{2K+1} + \frac{K+1}{2K+1} \cdot \frac{5}{6} = \frac{11K+5}{12K+6}$, and hence, the social welfare is at least $m \cdot \frac{11K+5}{12K+6} \geq \frac{11}{12}m - O(1/m^2)$.

Soundness Consider any normalized HZ equilibrium (x, p). We first characterize the equilibrium behaviour of variable gadgets. In an (exact) HZ equilibrium, we say a variable gadget is *vacant* if no agents outside of the gadget purchase goods inside the gadget, and we call other gadgets *non-vacant*. Loosely speaking, only non-vacant gadgets are of interest, as vacant gadgets do not interact with the rest of market and their utility is negligible.

We give a detailed analysis of variable gadgets to prove the following lemma in the full version [5]:

Lemma 4.1. *For any $i \in [n]$, suppose the i-th variable gadget is non-vacant. Then the equilibrium price is one of the following three cases.*

$$p(G_{i,1}) = \frac{4}{5} \pm O\left(\frac{1}{m^2}\right), \quad p(G_{i,2}) = \frac{8}{5} \pm O\left(\frac{1}{m^2}\right), \quad p(G_{i,3}) = 0 \pm O\left(\frac{1}{m^2}\right) \tag{4.6}$$

$$\textit{or} \quad p(G_{i,1}) = 0 \pm O\left(\frac{1}{m^2}\right), \quad p(G_{i,2}) = \frac{8}{5} \pm O\left(\frac{1}{m^2}\right), \quad p(G_{i,3}) = \frac{4}{5} \pm O\left(\frac{1}{m^2}\right) \tag{4.7}$$

$$\textit{or} \quad p(G_{i,1}) = \frac{2}{3} \pm O\left(\frac{1}{m^2}\right), \quad p(G_{i,2}) = \frac{4}{3} \pm O\left(\frac{1}{m^2}\right), \quad p(G_{i,3}) = \frac{2}{3} \pm O\left(\frac{1}{m^2}\right) \tag{4.8}$$

The following lemma follows a similar argument of Lemma 3.11, we omit the proof:

LEMMA 4.2. *The price of goods G_j satisfies $p(G_j) = 2 + O(1/m^2)$, $j \in [m]$.*

We are now ready to wrap up the proof of soundness. Given an equilibrium (x, p) that (approximately) maximizes the social welfare, we look at each non-vacant variable gadget. Based on the three cases stated in Lemma 4.1, we extract the i-th variable to be 1 if Eq. (4.7) holds and 0 if Eq. (4.6) holds. We do nothing for the case of Eq. (4.8) and those vacant variables (gadgets).

The total utility of all agents in $A_{i,\ell}$, $i \in [n], \ell \in [2]$, and all agents in $A_{j,*}$, $j \in [m]$ is at most $O(1/m^2)$. We focus on the utility of agents A_j, $j \in [m]$. If the j-th clause is satisfied, then one of the 5/6 utility goods has zero price, and one can see that the utility is (at most)

$$\left(\frac{1}{2} \pm O\left(\frac{1}{m^2}\right)\right) \cdot 1 + \left(\frac{1}{2} \pm O\left(\frac{1}{m^2}\right)\right) \cdot \frac{5}{6} = \frac{11}{12} \pm O\left(\frac{1}{m^2}\right).$$

On the other hand, if the j-th clause is not satisfied, we still don't need to consider the vacant gadgets (as there is no interactions), and the 5/6 utility goods have price at least $(2/3) \pm O(1/m^2)$. Hence the utility is at most

$$\left(\frac{1}{4} \pm O\left(\frac{1}{m^2}\right)\right) \cdot 1 + \left(\frac{3}{4} \pm O\left(\frac{1}{m^2}\right)\right) \cdot \frac{5}{6} = \frac{7}{8} \pm O\left(\frac{1}{m^2}\right).$$

Thus, if the truth assignment satisfies at most $(\frac{7}{8} + \epsilon)m$ clauses then the social welfare is at most

$$\left(\frac{7}{8} + \epsilon\right) m \cdot \left(\frac{11}{12} + O\left(\frac{1}{m^2}\right)\right) + \left(\frac{1}{8} - \epsilon\right) m \cdot \left(\frac{7}{8} + O\left(\frac{1}{m^2}\right)\right) + O\left(\frac{1}{m}\right) = \frac{175}{192}m + \frac{1}{24}\epsilon m + O\left(\frac{1}{m}\right).$$

From [18], it is NP-hard to distinguish the case that all clauses can be satisfied (in which case there is an equilibrium with social welfare $\frac{11}{12}m - O(1/m^2)$) from the case that at most $(\frac{7}{8} + \epsilon)m$ clauses can be satisfied (in which case the maximum social welfare is at most $\frac{175}{192}m + \frac{1}{24}\epsilon m + O(1/m)$). The theorem follows.

The construction can be easily modified, if desired, so that all utilities are in $[0, 1]$, and every agent has minimum utility 0 and maximum utility 1.

5 Discussion

In this paper we resolved the complexity of computing an approximate equilibrium in the Hylland-Zeckhauser scheme for one-sided matching markets: we showed that the problem is PPAD-complete, and this holds even for inverse polynomial approximation and four-valued utilities. We leave open the complexity of exact equilibria, in particular whether the problem is FIXP-complete. Another open question is whether the PPAD-hardness of the approximation problem holds also for 3-valued utilities.

Acknowledgement

Work supported by NSF grants CCF-1703925, CCF-1763970, IIS-1838154, CCF-2106429 and CCF-2107187.

References

[1] A. Abdulkadiroglou and T. Sonmez, *Random serial dictatorship and the core from random endowments in house allocation problems*, Econometrica, 66 (1998), pp. 689–702.

[2] S. Alaei, P. Jalaly Khalilabadi, and E. Tardos, *Computing equilibrium in matching markets*, in Proceedings of the 2017 ACM Conference on Economics and Computation, 2017, pp. 245–261.

[3] A. Bogomolnaia and H. Moulin, *A new solution to the random assignment problem*, Journal of Economic theory, 100 (2001), pp. 295–328.

[4] E. Budish, *The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes*, Journal of Political Economy, 119 (2011), pp. 1061–1103.

[5] T. Chen, X. Chen, B. Peng, and M. Yannakakis, *Computational hardness of the hylland-zeckhauser scheme*, CoRR, abs/2107.05746 (2021).

[6] X. Chen, D. Dai, Y. Du, and S.-H. Teng, *Settling the complexity of Arrow-Debreu equilibria in markets with additively separable utilities*, in 2009 50th Annual IEEE Symposium on Foundations of Computer Science, IEEE, 2009, pp. 273–282.

[7] X. Chen, D. Paparas, and M. Yannakakis, *The complexity of non-monotone markets*, J. ACM, 64 (2017), pp. 20:1–20:56.
[8] X. Chen and S.-H. Teng, *Spending is not easier than trading: on the computational equivalence of Fisher and Arrow-Debreu equilibria*, in International Symposium on Algorithms and Computation, Springer, 2009, pp. 647–656.
[9] N. R. Devanur and R. Kannan, *Market equilibria in polynomial time for fixed number of goods or agents*, in 2008 49th Annual IEEE Symposium on Foundations of Computer Science, IEEE, 2008, pp. 45–53.
[10] N. R. Devanur, C. H. Papadimitriou, A. Saberi, and V. V. Vazirani, *Market equilibrium via a primal–dual algorithm for a convex program*, Journal of the ACM (JACM), 55 (2008), p. 22.
[11] R. Duan and K. Mehlhorn, *A combinatorial polynomial algorithm for the linear arrow–debreu market*, Information and Computation, 243 (2015), pp. 112–132.
[12] F. Echenique, A. Miralles, and J. Zhang, *Constrained pseudo-market equilibrium*, arXiv preprint arXiv:1909.05986, (2019).
[13] ———, *Fairness and efficiency for probabilistic allocations with endowments*, arXiv preprint arXiv:1908.04336, (2019).
[14] K. Etessami and M. Yannakakis, *On the complexity of Nash equilibria and other fixed points*, SIAM Journal on Computing, 39 (2010), pp. 2531–2597.
[15] J. Garg, R. Mehta, V. V. Vazirani, and S. Yazdanbod, *Settling the complexity of Leontief and PLC exchange markets under exact and approximate equilibria*, in Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, 2017, pp. 890–901.
[16] J. Garg, T. Tröbst, and V. V. Vazirani, *An arrow-debreu extension of the Hylland-Zeckhauser scheme: Equilibrium existence and algorithms*, arXiv preprint arXiv:2009.10320, (2020).
[17] J. Garg and L. A. Végh, *A strongly polynomial algorithm for linear exchange markets*, in Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, 2019, pp. 54–65.
[18] J. Håstad, *Some optimal inapproximability results*, Journal of the ACM (JACM), 48 (2001), pp. 798–859.
[19] Y. He, A. Miralles, M. Pycia, and J. Yan, *A pseudo-market approach to allocation with priorities*, American Economic Journal: Microeconomics, 10 (2018), pp. 272–314.
[20] M. Hosseini and V. V. Vazirani, *Nash-bargaining-based models for matching markets, with implementations and experimental results*, arXiv preprint arXiv:2105.10704, (2021).
[21] A. Hylland and R. Zeckhauser, *The efficient allocation of individuals to positions*, Journal of Political economy, 87 (1979), pp. 293–314.
[22] K. Jain, *A polynomial time algorithm for computing an Arrow-Debreu market equilibrium for linear utilities*, SIAM J. Comput., 37 (2007), pp. 303–318.
[23] P. Le, *Competitive equilibrium in the random assignment problem*, International Journal of Economic Theory, 13 (2017), pp. 369–385.
[24] A. McLennan, *Efficient disposal equilibria of pseudomarkets*, in Workshop on Game Theory, vol. 4, 2018, p. 8.
[25] H. Moulin, *Fair division in the age of internet*, Annual Review of Economics, (2018).
[26] J. Orlin, *Improved algorithms for computing Fisher's market clearing prices*, in Proc. 42nd ACM Symp. Theory of Computing, 2010, pp. 291–300.
[27] C. Papadimitriou and B. Peng, *Public goods games in directed networks*, in Proceedings of the 22nd ACM Conference on Electronic Commerce, 2021.
[28] V. V. Vazirani and M. Yannakakis, *Market equilibria under separable, piecewise-linear, concave utilities*, J. Assoc. Comput. Mach., 58(3) (2011).
[29] V. V. Vazirani and M. Yannakakis, *Computational complexity of the Hylland-Zeckhauser scheme for one-sided matching markets*, in 12th Innovations in Theoretical Computer Science Conference (ITCS 2021), Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.