Computational Complexity of the Hylland-Zeckhauser Scheme for One-Sided Matching Markets

Vijay V. Vazirani
Department of Computer Science, University of California, Irvine, CA, USA
vazirani@ics.uci.edu

Mihalis Yannakakis
Department of Computer Science, Columbia University, New York, NY, USA
mihalis@cs.columbia.edu

Abstract

In 1979, Hylland and Zeckhauser [23] gave a simple and general scheme for implementing a one-sided matching market using the power of a pricing mechanism. Their method has nice properties – it is incentive compatible in the large and produces an allocation that is Pareto optimal – and hence it provides an attractive, off-the-shelf method for running an application involving such a market. With matching markets becoming ever more prevalent and impactful, it is imperative to finally settle the computational complexity of this scheme.

We present the following partial resolution:

1. A combinatorial, strongly polynomial time algorithm for the dichotomous case, i.e., 0/1 utilities, and more generally, when each agent's utilities come from a bi-valued set.
2. An example that has only irrational equilibria, hence proving that this problem is not in PPAD.
3. A proof of membership of the problem in the class FIXP.
4. A proof of membership of the problem of computing an approximate HZ equilibrium in the class PPAD.

We leave open the (difficult) questions of determining if computing an exact HZ equilibrium is FIXP-hard and an approximate HZ equilibrium is PPAD-hard.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory

Keywords and phrases Hyland-Zeckhauser scheme, one-sided matching markets, mechanism design, dichotomous utilities, PPAD, FIXP

Digital Object Identifier 10.4230/LIPIcs.ITCS.2021.59

Related Version Full version: `https://arxiv.org/pdf/2004.01348`

Funding *Vijay V. Vazirani*: Supported in part by NSF grant CCF-1815901.
Mihalis Yannakakis: Supported in part by NSF grants CCF-1703925 and CCF-1763970.

Acknowledgements We wish to thank Federico Echenique, Jugal Garg, Tung Mai and Thorben Trobst for valuable discussions and Richard Zeckhauser for providing us with the Appendix to his paper [23]. In addition, the first author wishes to thank Simons Institute for running a program on matching markets in Fall 2019; this provided valuable exposure to the topic.

1 Introduction

In a brilliant and by-now classic paper, Hylland and Zeckhauser [23] gave a simple and general scheme for implementing a one-sided matching market using the power of a pricing mechanism[1]. Their method produces an allocation that is Pareto optimal, envy-free [23]

[1] See Remark 5 for a discussion of the advantages of this mechanism.

12th Innovations in Theoretical Computer Science Conference (ITCS 2021).
Editor: James R. Lee; Article No. 59; pp. 59:1–59:19
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

and is incentive compatible in the large [22]. The Hylland-Zeckhauser (HZ) scheme can be viewed as a marriage between fractional perfect matching and a linear Fisher market, both of which admit not only polynomial time algorithms but also combinatorial ones. These facts have enticed numerous researchers over the years to seek an efficient algorithm for the HZ scheme. The significance of this problem has only grown in recent years, with ever more diverse and impactful matching markets being launched into our economy, e.g., see [17].

Our work on resolving this problem started with an encouraging sign, when we obtained a combinatorial, strongly polynomial time algorithm for the *dichotomous case*, in which all utilities are 0/1, by melding a prefect matching algorithm with the combinatorial algorithm of [13] for the linear Fisher market, see Section 4. This algorithm can be extended to solve a more general problem which we call the *bi-valued utilities case*, in which each agent's utilities can take one of only two values, though the two values can be different for different agents. However, this approach did not extend any further, as described in the next section.

One-sided matching markets can be classified along two dimensions: whether the utility functions are cardinal or ordinal, and whether agents have initial endowments or not. Under this classification, the HZ scheme is (cardinal, no endowments). Section 1.2 gives mechanisms for the remaining three possibilities as well as their game-theoretic properties. Ordinal and cardinal utility functions have their individual pros and cons, and neither dominates the other. Whereas the former are easier to elicit, the latter are far more expressive, enabling an agent to not only report if she prefers one good to another but also by how much, thereby producing higher quality allocations as illustrated in Example 1, which is taken from [20].

▶ **Example 1.** ([20]) The instance has three types of goods, T_1, T_2, T_3, and these goods are present in the proportion of (1%, 97%, 2%). Based on their utility functions, the agents are partitioned into two sets A_1 and A_2, where A_1 constitute 1% of the agents and A_2, 99%. The utility functions of agents in A_1 and A_2 for the three types of goods are $(1,\ \epsilon,\ 0)$ and $(1,\ 1-\epsilon,\ 0)$, respectively, for a small number $\epsilon > 0$. The main point is that whereas agents in A_2 marginally prefer T_1 to T_2, those in A_1 overwhelmingly prefer T_1 to T_2. Clearly, the ordinal utilities of all agents in $A_1 \cup A_2$ are the same. Therefore, a mechanism based on such utilities will not be able to make a distinction between the two types of agents. On the other hand, the HZ mechanism, which uses cardinal utilities, will fix the price of goods in T_3 to be zero and those in T_1 and T_2 appropriately so that by-and-large the bundles of A_1 and A_2 consist of goods from T_1 and T_2, respectively.

While studying the dichotomous case of two-sided markets, Bogomolnaia and Moulin [5] called it an "important special case of the bilateral matching problem." Using the Gallai-Edmonds decomposition of a bipartite graph, they gave a mechanism that is Pareto optimal and group strategyproof. They also gave a number of applications of their setting, some of which are natural applications of one-sided markets as well, e.g., housemates distributing rooms, having different features, in a house. Furthermore, they say, "Time sharing is the simplest way to deal fairly with indivisibilities of matching markets: think of a set of workers sharing their time among a set of employers." It turns out that the HZ (fractional) equilibrium allocation is a superior starting point for the problem of designing a randomized time-sharing mechanism; this is discussed in Remark 5 after introducing the HZ model. Roth, Sonmez and Unver [29] extended these results to general graph matching under dichotomous utilities; this setting is applicable to the kidney exchange marketplace.

1.1 The gamut of possibilities

The most useful solution for practical applications would of course have been a combinatorial, polynomial time algorithm for the entire scheme. At the outset, this didn't seem unlikely, especially in view of the existence of such an algorithm for the dichotomous case. Next we considered the generalization of the bi-valued utilities case to tri-valued utilities, in particular, to the case of $\{0, \frac{1}{2}, 1\}$ utilities. However, even this case appears to be intractable.

Underlying the polynomial time solvability of a linear Fisher market is the property of weak gross substitutability[2]. We note that this property is destroyed as soon as one goes to a slightly more general utility function, namely piecewise-linear, concave and separable over goods (SPLC utilities), and this case is PPAD-complete[3] [32]; the class PPAD was introduced in [28]. Since equilibrium allocations for the HZ scheme do not satisfy weak gross substitutability, e.g., see Example 9, we were led us to seek a proof of PPAD-completeness.

A crucial requirement for membership in PPAD is to show that there is always a rational equilibrium if all parameters of the instance are rational numbers. However, even this is not true; we found an example which admits only irrational equilibria, see Section 6. This example consists of four agents and goods, and hence can be viewed as belonging to the four-valued utilities case; see Remark 17 for other intriguing aspects of this example.

The irrationality of solutions suggests that the appropriate class for this problem is the class FIXP, introduced in [16]. The proof in [23], showing the existence of an equilibrium, uses Kakutani's theorem and does not seem to lend itself in any easy way to showing membership in FIXP. For this purpose, we give a new proof of the existence of equilibrium. Our proof defines a suitable Brouwer function which adjusts prices and allocations in case they are not an equilibrium. It uses elementary arithmetic operations that improve their optimality or feasibility of the current prices and allocations. The adjustment scheme is such that the only stable prices and allocations are forced to be equilibria.

Next, we define the notion of an approximate equilibrium. This is still required to be a fractional perfect matching on agents and goods; agents' allocations are allowed to be slightly suboptimal and/or their cost is allowed to slightly exceed the budget of 1 dollar. We show that the problem of computing such an approximate equilibrium is in PPAD. This involves relating approximate equilibria to the approximate fixed points of the Brouwer function we defined for our proof of membership in FIXP. We leave open the questions of determining if the computation (to desired accuracy) of an exact HZ equilibrium is FIXP-hard, and if the computation of an approximate HZ equilibrium is PPAD-hard.

1.2 Related work

We first present mechanisms for the remaining three possibilities for the classification of one-sided matching market mechanisms given in the Introduction. The famous Top Trading Cycles mechanism is (ordinal, endowments) [30]; it is efficient, strategyproof and core-stable. Under (ordinal, no endowments) are Random Priority [27], which is strategyproof though not efficient or envy-free, and Probabilistic Serial [4], which is efficient and envy-free but not strategyproof. Under (cardinal, endowments) is ϵ-Approximate ADHZ (for Arrow-Debreu Hylland-Zeckhauser) scheme [20], which satisfies Pareto optimality, approximate envy-freeness and incentive compatibility in the large.

---

[2] Namely, if you increase the price of one good, the demand of another good cannot decrease.
[3] Independently, PPAD-hardness was also established in [10].

We are aware of only the following two computational results on the HZ scheme. Using the algebraic cell decomposition technique of [2], [12] gave a polynomial time algorithm for computing an equilibrium for an Arrow-Debreu market under piecewise-linear, concave (PLC) utilities (not necessarily separable over goods) if the number of goods is fixed. One can see that their algorithm can be adapted to yield a polynomial time algorithm for computing an equilibrium for the HZ scheme if the number of goods is a fixed constant. Extending these methods, [1] gave a polynomial time algorithm for the case that the number of agents is a fixed constant.

There are several results establishing membership and hardness in PPAD and FIXP for equilibria computation problems in different settings. The quintessential complete problem for PPAD is 2-Nash [11, 7] and that for FIXP is multiplayer Nash equilibrium [16]. For the latter problem, computing an approximate equilibrium is PPAD-complete [11].

For the case of market equilibria, in the economics literature, there are two parallel streams of results: one assumes that an excess demand function is given and the other assumes a specific class of utility functions. [16] proved FIXP-completeness of Arrow-Debreu markets whose excess demand functions are algebraic. This result is for the first stream and it does not establish FIXP-completeness of Arrow-Debreu markets under any specific class of utility functions. Results for the second stream include proofs of membership in FIXP for Arrow-Debreu markets under Leontief and piecewise-linear concave (PLC) utility functions in [35] and [18], respectively. This was followed by a proof of FIXP-hardness for Arrow-Debreu markets with Leontief and PLC utilities [19]. For the case of Arrow-Debreu markets with CES (constant elasticity of substitution) utility functions, [9] show membership in FIXP but leave open FIXP-hardness.

For the CES market problem stated above, computing an approximate equilibrium is PPAD-complete, and the same holds more generally for a large class of “non-monotonic” markets [9]. Computing an (exact or approximate) equilibrium under separable, piecewise-linear, concave (SPLC) utilities for Arrow-Debreu and Fisher markets is also known to be PPAD-complete [8, 10, 32].

In recent years, several researchers have proposed Hylland-Zeckhauser-type mechanisms for a number of applications, e.g., see [6, 22, 25, 26]. The basic scheme has also been generalized in several different directions, including two-sided matching markets, adding quantitative constraints, and to the setting in which agents have initial endowments of goods instead of money, see [14, 15].

2 The Hylland-Zeckhauser Scheme

Hylland and Zeckhauser [23] gave a general mechanism for a one-sided matching market using the power of a pricing mechanism. Their formulation is as follows: Let $A = \{1, 2, \dots n\}$ be a set of n agents and $G = \{1, 2, \dots, n\}$ be a set of n indivisible goods. The mechanism will allocate exactly one good to each agent and will have the following two properties:

- The allocation produced is Pareto optimal.
- The mechanism is incentive compatible in the large.

The Hylland-Zeckhauser scheme is a marriage between linear Fisher market and fractional perfect matching. The agents will reveal to the mechanism their desires for the goods by stating their von Neumann-Mogenstern utilities. Let u_{ij} represent the utility of agent i for good j. We will use language from the study of market equilibria to describe the rest of the formulation. For this purpose, we next define the linear Fisher market model.

A *linear Fisher market* consists of a set $A = \{1, 2, \ldots n\}$ of n agents and a set $G = \{1, 2, \ldots, m\}$ of m infinitely divisible goods. By fixing the units for each good, we may assume without loss of generality that there is a unit of each good in the market. Each agent i has money m_i and utility u_{ij} for a unit of good j. If x_{ij}, $1 \le j \le m$ is the *bundle of goods allocated to* i, then the utility accrued by i is $\sum_j u_{ij} x_{ij}$. Each good j is assigned a non-negative price, p_j. Allocations and prices, x and p, are said to form an *equilibrium* if each agent obtains a utility maximizing bundle of goods at prices p and the *market clears*, i.e., each good is fully sold and all money of agents is fully spent.

In order to mold the one-sided market into a linear Fisher market, the HZ scheme renders goods divisible by assuming that there is one unit of probability share of each good. An *allocation* to an agent is a collection of probability shares over the goods. Let x_{ij} be the probability share that agent i receives of good j. Then, $\sum_j u_{ij} x_{ij}$ is the *expected utility* accrued by agent i. Each good j has price $p_j \ge 0$ in this market and each agent has 1 dollar with which it buys probability shares. The entire allocation must form a *fractional perfect matching in the complete bipartite graph* over vertex sets A and G as follows: there is one unit of probability share of each good and the total probability share assigned to each agent also needs to be one unit. Subject to these constraints, each agent should buy a utility maximizing bundle of goods *having the smallest possible cost.* Note that the last condition is not required in the definition of a linear Fisher market equilibrium. It is needed here to guarantee that the allocation obtained is Pareto optimal, see [23] for proof of Pareto optimality. A second departure from the linear Fisher market equilibrium is that in the latter, each agent i must spend her money m_i fully; in the HZ scheme, i need not spend her entire dollar. Since the allocation is required to form a fractional perfect matching, all goods are fully sold. We will define these to be *equilibrium allocation and prices*; we state this formally below after giving some preliminary definitions.

▶ **Definition 2.** *Let x and p denote arbitrary (non-negative) allocations and prices of goods. By* size, cost *and* value *of agent i's bundle we mean*

$$\sum_{j \in G} x_{ij}, \qquad \sum_{j \in G} p_j x_{ij} \qquad \text{and} \qquad \sum_{j \in G} u_{ij} x_{ij},$$

respectively. We will denote these by $\text{size}(i)$, $\text{cost}(i)$ *and* $\text{value}(i)$, *respectively.*

▶ **Definition 3** (Hylland and Zeckhauser [23])**.** *Allocations and prices* (x, p) *form an* equilibrium *for the one-sided matching market stated above if:*

1. *The total probability share of each good j is 1 unit, i.e.,* $\sum_i x_{ij} = 1$.
2. *The size of each agent i's allocation is 1, i.e.,* $\text{size}(i) = 1$.
3. *The cost of the bundle of each agent is at most 1.*
4. *Subject to constraints 2 and 3, each agent i maximizes her expected utility at minimum possible cost, i.e., maximize* $\text{value}(i)$, *subject to* $\text{size}(i) = 1$, $\text{cost}(i) \le 1$, *and lastly,* $\text{cost}(i)$ *is smallest among all utility-maximizing bundles of i.*

Using Kakutani's fixed point theorem, the following is shown:

▶ **Theorem 4** (Hylland and Zeckhauser [23])**.** *Every instance of the one-sided market defined above admits an equilibrium; moreover, the corresponding allocation is Pareto optimal.*

Finally, if this "market" is large enough, no individual agent will be able to improve her allocation by misreporting utilities nor will she be able to manipulate prices. For this reason, the HZ scheme is incentive compatible in the large.

As stated above, Hylland and Zeckhauser view each agent's allocation as a lottery over goods. In this viewpoint, agents accrue utility in an *expected sense* from their allocations. Once these lotteries are resolved in a manner faithful to the probabilities, an assignment of indivisible goods will result. The latter can be done using the well-known Theorem of Birkhoff [3] and von Neumann [34] which states that any doubly stochastic matrix can be written as a convex combination of permutation matrices, i.e., perfect matchings; moreover, this decomposition can be obtained efficiently. Next, pick one of these perfect matchings from the discrete distribution given by coefficients in the convex combination. As is well known, since the lottery over goods is Pareto optimal *ex ante*, the integral allocation, viewed stochastically, will also be Pareto optimal *ex post*.

Another viewpoint, forwarded by Bogomolnaia and Moulin [5], considers the fractional perfect matching, or equivalently the doubly-stochastic matrix, as the output of the mechanism, i.e., without resorting to randomized rounding. This viewpoint assumes that the agents are going to "time-share" the goods or resources and the doubly-stochastic matrix, which is derived from a market mechanism, provides a "fair" way of doing so.

▶ Remark 5. In their paper studying the dichotomous case of two-sided matching markets, Bogomolnaia and Moulin [5] state that the preferred way of dealing with indivisibilities inherent in matching markets is to resort to time sharing using randomization. Their method builds on the Gallai-Edmonds decomposition of the underlying bipartite graph; this classifies vertices into three categories: disposable, over-demanded and perfectly matched. This is a much more coarse insight into the demand structure of vertices than that obtained via the HZ equilibrium. The latter is the output of a market mechanism in which equilibrium prices reflect the relative importance of goods in an accurate and precise manner, based on the utilities declared by buyers, and equilibrium allocations are as equitable as possible across buyers. Hence the latter yields a more fair and desirable randomized time-sharing mechanism.

3 Properties of Optimal Allocations and Prices

Let p be given prices which are not necessarily equilibrium prices. An optimal bundle for agent i, x_i, is a solution to the following LP, which has two constraints, one for size and one for cost.

$$\max \quad \sum_j x_{ij} u_{ij} \tag{1}$$

$$\text{s.t.} \tag{2}$$

$$\sum_j x_{ij} = 1 \tag{3}$$

$$\sum_j x_{ij} p_j \leq 1 \tag{4}$$

$$\forall j \quad x_{ij} \geq 0 \tag{5}$$

Taking μ_i and α_i to be the dual variables corresponding to the two constraints, we get the dual LP:

$$\min \quad \alpha_i + \mu_i \tag{6}$$

$$\text{s.t.} \tag{7}$$

$$\forall i, j \quad \alpha_i p_j + \mu_i \geq u_{ij} \tag{8}$$

$$\alpha_i \geq 0 \tag{9}$$

Clearly μ_i is unconstrained. μ_i will be called the *offset* on i's utilities. By complementary slackness, if x_{ij} is positive then $\alpha_i p_j = u_{ij} - \mu_i$. All goods j satisfying this equality will be called *optimal goods for agent* i. The rest of the goods, called *suboptimal*, will satisfy $\alpha_i p_j > u_{ij} - \mu_i$. Obviously an optimal bundle for i must contain only optimal goods.

The parameter μ_i plays a crucial role in ensuring that i's optimal bundle satisfies both size and cost constraints. If a single good is an effective way of satisfying both size and cost constraints, then μ_i plays no role and can be set to zero. However, if different goods are better from the viewpoint of size and cost, then μ_i attains the right value so they both become optimal and i buys an appropriate combination. We provide an example below to illustrate this.

▶ **Example 6.** Suppose i has positive utilities for only two goods, j and k, with $u_{ij} = 10$, $u_{ik} = 2$ and their prices are $p_j = 2$, $p_k = 0.1$. Clearly, neither good satisfies both size and cost constraints optimally: good j is better for the size constraint and k is better for the cost constraint. If i buys one unit of good j, she spends 2 dollars, thus exceeding her budget. On the other hand, she can afford to buy 10 units of k, giving her utility of 20; however, she has far exceeded the size constraint. It turns out that her optimal bundle consists of 9/19 units of j and 10/19 units of k; the costs of these two goods being 18/19 and 1/19 dollars, respectively. Clearly, her size and cost constraints are both met exactly. Her total utility from this bundle is 110/19. It is easy to see that $\alpha_i = 80/19$ and $\mu_i = 30/19$, and for these settings of the parameters, both goods are optimal.

We next show that equilibrium prices are invariant under the operation of *scaling* the difference of prices from 1.

▶ **Lemma 7.** *Let p be an equilibrium price vector and fix any $r > 0$. Let p' be such that $\forall j \in G$, $p'_j - 1 = r(p_j - 1)$. Then p' is also an equilibrium price vector.*

Proof. Consider an agent i. Clearly, $\sum_{j \in G} p_j x_{ij} \leq 1$. Now,

$$\sum_{j \in G} p'_j x_{ij} \;=\; \sum_{j \in G} (rp_j - r + 1)x_{ij} \leq 1,$$

where the last inequality follows by using $\sum_{j \in G} x_{ij} = 1$. ◀

Using Lemma 7, it is easy to see that if the allocation x provides optimal bundles to all agents under prices p then it also does so under p'. In the rest of this paper we will enforce that the minimum price of a good is zero, thereby fixing the scale. Observe that the main goal of the Hylland-Zeckhauser scheme is to yield the "correct" allocations to agents; the prices are simply a vehicle in the market mechanism to achieve this. Hence arbitrarily fixing the scale does not change the essential nature of the problem. Moreover, setting the minimum price to zero is standard [23] and can lead to simplifying the equilibrium computation problem as shown in Remark 8.

▶ Remark 8. We remark that on the one hand, the offset μ_i is a key enabler in construing optimal bundles, on the other, it is also a main source of difficulty in computing equilibria for the HZ scheme. We identify here an interesting case in which $\mu_i = 0$ and this difficulty is mitigated. In particular, this holds for all agents in the dichotomous case presented in Section 4. Suppose good j is optimal for agent i, $u_{ij} = 0$ and $p_j = 0$, then it is easy to check that $\mu_i = 0$. If so, the optimal goods for i are simply the maximum bang-per-buck goods; the latter notion is replete in market equilibrium papers, e.g., see [13].

Finally, we extend Example 6 to illustrate that optimal allocations for the Hylland-Zeckhauser model do not satisfy the weak gross substitutes condition in general.

▶ **Example 9.** In Example 6, let us raise the price of k to 0.2 dollars. Then the optimal allocation for i changes to 4/9 units of j and 5/9 units of k. Notice that the demand for j went down from 9/19 to 4/9. One way to understand this change is as follows: Let us start with the old allocation of 10/19 units of k. Clearly, the cost of this allocation of k went up from 1/19 to 2/19, leaving only 17/19 dollars for j. Therefore size of j needs to be reduced to 17/38. However, now the sum of the sizes becomes 37/38, i.e., less than a unit. We wish to increase this to a unit while still keeping cost at a unit. The only way of doing this is to sell some of the more expensive good and use the money to buy the cheaper good. This is the reason for the decrease in demand of j.

4 Strongly Polynomial Algorithm for Bi-Valued Utilities

In this section, we will study the restriction of the HZ scheme to the bi-valued utilities case, which is defined as follows: for each agent i, we are given a set $\{a_i, b_i\}$, where $0 \leq a_i < b_i$, and the utilities u_{ij}, $\forall j \in G$, are picked from this set. However first, using a perfect matching algorithm and the combinatorial algorithm [13] for linear Fisher markets, we will give a strongly polynomial time algorithm for the dichotomous case, i.e., when all utilities u_{ij} are 0/1. Next we define the notion of equivalence of utility functions and show that the bi-valued utilities case is equivalent to the dichotomous case, thereby extending the dichotomous case algorithm to this case.

We need to clarify that we will not use the main algorithm from [13], which uses the notion of balanced flows and l_2 norm to achieve polynomial running time. Instead, we will use the "simple algorithm" presented in Section 5 in [13]. Although this algorithm is not proven to be efficient, the simplified version we define below, called Simplified DPSV Algorithm, is efficient; in fact it runs in strongly polynomial time, unlike the balanced-flows-based algorithm of [13]. Remark 8 provides an insight into what makes the dichotomous case computationally easier.

We note that recently, [20] gave a *rational convex program (RCP)* for the dichotomous case of HZ, and more recently, [33] made a small modification to our algorithm to obtain a mechanism that is proven to be strategyproof. An RCP, defined in [31], is a nonlinear convex program all of whose parameters are rational numbers and which always admits a rational solution in which the denominators are polynomially bounded. An RCP can be solved exactly in polynomial time using the ellipsoid algorithm and diophantine approximation [21, 24], and therefore directly implies the existence of a polynomial time algorithm for the underlying problem.

Notation. We will denote by $H = (A, G, E)$ be the bipartite graph on vertex sets A and G, and edge set E, with $(i, j) \in E$ iff $u_{ij} = 1$. For $A' \subseteq A$ and $G' \subseteq G$, we will denote by $H[A', G']$ the restriction of H to vertex set $A' \cup G'$. If ν is a matching in H, $\nu \subseteq E$, and $(i, j) \in \nu$ then we will say that $\nu(i) = j$ and $\nu(j) = i$. For any subset $S \subseteq A$ ($S \subseteq G$), $N(S)$ will denote the set of neighbors, in G (A), of vertices in S.

If H has a perfect matching, the matter is straightforward as stated in Steps 1a and 1b; allocations and prices are clearly in equilibrium. For Step 2, we need the following lemma.

▶ **Lemma 10.** *The following hold:*
1. *For any set* $S \subseteq A_2$, $|N(S)| \geq |S|$.
2. *For any set* $S \subseteq G_1$, $|N(S) \cap A_1| \geq |S|$.

Proof.

1. If $|N(S)| < |S|$ then $(G_1 \cup N(S)) \cup (A_2 - S)$ is a smaller vertex cover for H, leading to a contradiction.
2. If $|N(S) \cap A_1| < |S|$ then $(G_1 - S) \cup (A_2 \cup N(S))$ is a smaller vertex cover for H, leading to a contradiction. ◀

The first part of Lemma 10 together with Hall's Theorem implies that a maximum matching in $H[A_2, G_2]$ must match all agents. Therefore in Step 2a, each agent $i \in A_2$ is allocated one unit of a unique good from which it derives utility 1 and having price zero; clearly, this is an optimal bundle of minimum cost for i. The number of goods that will remain unmatched in G_2 at the end of this step is $|G_2| - |A_2|$.

Algorithm 1 Algorithm for the Dichotomous Case.

---

1. If H has a perfect matching, say ν, then do:
   - a. $\forall i \in A$: allocate good $\nu(i)$ to i.
   - b. $\forall j \in G$: $p_j \leftarrow 0$. Go to Step 3.
2. Else do:
   - a. Find a minimum vertex cover in H, say $G_1 \cup A_2$, where $G_1 \subset G$ and $A_2 \subset A$. Let $A_1 = A - A_2$ and $G_2 = G - G_1$.
   - b. Find a maximum matching in $H[A_2, G_2]$, say ν.
   - c. $\forall i \in A_2$: allocate good $\nu(i)$ to i.
   - d. $\forall j \in G_2$: $p_j \leftarrow 0$.
   - e. Run the Simplified DPSV Algorithm on agents A_1 and goods G_1.
   - f. $\forall i \in A_1$: Allocate unmatched goods of G_2 to satisfy the size constraint.
3. Output the allocations and prices computed and Halt.

Allocations are computed for agents in A_1 as follows. First, Step 2e uses the Simplified DPSV Algorithm, which we describe below, to compute equilibrium allocations and prices for the submarket consisting of agents in A_1 and goods in G_1. At the end of this step, the money of each agent in A_1 is exhausted; however, her allocation may not meet the size constraint. To achieve the latter, Step 2f allocates the unmatched zero-priced goods from G_2 to agents in A_1. Clearly, the total deficit in size is $|A_1| - |G_1|$. Since this equals $|G_2| - |A_2|$, the market clears at the end of Step 2f. As shown in Lemma 11, each agent in A_1 also gets an optimal bundle of goods of minimum cost.

Let p be the prices of goods in G_1 at any point in this algorithm. As a consequence of the second part of Lemma 10, the equilibrium price of each good in G_1 will be at least 1. The Simplified DPSV algorithm will initialize prices of goods in G_1 to 1 and declare all goods active. The algorithm will always raise prices of active goods uniformly[4].

For $S \subseteq G_1$ let $p(S)$ denote the sum of the equilibrium prices of goods in S. A key notion from [13] is that of a tight set; set $S \subseteq G_1$ is said to be *tight* if $p(S) = |N(S)|$, the latter being the total money of agents in A_1 who are interested in goods in S. If set S is tight, then the local market consisting of goods in S and agents in $N(S)$ clears. To see this, one needs to use the flow-based procedure given in [13] to show that each agent $i \in N(S)$ can be allocated 1 dollar worth of those goods in S from which it accrues unit utility. Thus equilibrium has been reached for goods in S.

[4] In [13], prices of active goods are raised multiplicatively, which amounts to raising prices of active goods uniformly for our simplified setting.

As the algorithm raises prices of all goods in G_1, at some point a set S will go tight. The algorithm then *freezes* the prices of its goods and removes them from the active set. It then proceeds to raise the prices of currently active goods until another set goes tight, and so on, until all goods in G_1 are frozen.

We can now explain in what sense we need a "simplified" version of the DPSV algorithm. Assume that at some point, $S \subset G_1$ is frozen and goods in $G_1 - S$ are active and their prices are raised. As this happens, agents in $A_1 - N(S)$ start preferring goods in S relative to those in $G_1 - S$. In the general case, at some point, an agent $i \in (A_1 - N(S))$ will prefer a good $j \in S$ as much as her other preferred goods. At this point, edge (i, j) is added to the active graph. As a result, some set $S' \subseteq S$, containing j, will not be tight anymore and will be unfrozen. However, in our setting, the utilities of agents in $(A_1 - N(S))$ for goods in S is zero, and therefore no new edges are introduced and tight sets never become unfrozen. Hence the only events of the Simplified DPSV Algorithm are raising of prices and freezing of sets. Clearly, there will be at most n freezings. One can check details in [13] to see that the steps executed with each freezing run in strongly polynomial time, hence making the Simplified DPSV Algorithm a strongly polynomial time algorithm[5].

▶ **Lemma 11.** *Each agent in A_1 will get an optimal bundle of goods of minimum cost.*

Proof. First note that for agents in A_1, there are no utility 1 goods in G_2 – this follows from the fact that no vertices from $A_1 \cup G_2$ are in the vertex cover picked. Therefore, for $i \in A_1$, an optimum bundle consists of the cheapest way of obtaining one dollar worth of goods from $N\{i\}$, which are in G_1, together with the right amount of zero-priced goods from G_2 to satisfy the size constraint.

Assume that the algorithm freezes k sets, $S_1, \ldots S_k$, in that order; the union of these sets being G_1. Let $p_1, p_2, \ldots p_k$ be the prices of goods in these sets, respectively. Clearly, successive freezings will be at higher and higher prices and therefore, $1 \leq p_1 < p_2 < \ldots < p_k$, and for $1 \leq j \leq k, p_j = |N(S_j)|/|S_j|$. If $i \in N(S_j)$, the algorithm will allocate $1/p_j$ amount of goods to i from S_j, costing 1 dollar.

By definition of neighborhood of sets, if $i \in N(S_j)$, then i cannot have edges to $S_1, \ldots S_{j-1}$ and can have edges to $S_{j+1}, \ldots, S_k$. Therefore, the cheapest goods from which it accrues unit utility are in S_j, the set from which she gets 1 dollar worth of allocation. The rest of the allocation of i, in order to meet i's size constraint, will be from G_2, which are zero-priced and from which i gets zero utility. Clearly, i gets an optimal bundle of minimum cost. ◀

Since all steps of the algorithm, namely finding a maximum matching, a minimum vertex cover and running the Simplified DPSV Algorithm, can be executed in strongly polynomial time, we get:

▶ **Lemma 12.** *The algorithm given finds equilibrium prices and allocations for the dichotomous case of the Hylland-Zeckhauser scheme. It runs in strongly polynomial time.*

▶ **Definition 13.** *Let I be an instance of the HZ scheme and let the utility function of agent i be $u_i = \{u_{i1}, u_{i12}, \ldots, u_{in}\}$. Then $u'_i = \{u'_{i1}, u'_{i12}, \ldots, u'_{in}\}$ is* equivalent *to u_i if there are two numbers $s > 0$ and $h \geq 0$ such that for $1 \leq j \leq n$, $u'_{ij} = s \cdot u_{ij} + h$. The numbers s and h will be called the* scaling factor *and* shift*, respectively.*

[5] In contrast, in the general case, the number of freezings is not known to be bounded by a polynomial in n, as stated in [13].

▶ **Lemma 14.** *Let I be an instance of the HZ scheme and let the utility function of agent i be u_i. Let u'_i be equivalent to u_i and let I' be the instance obtained by replacing u_i by u'_i in I. Then x and p are equilibrium allocation and prices for I if and only if they are also for I'.*

Proof. Let s and h be the scaling factor and shift that transform u_i to u'_i. By the statement of the lemma, $x_i = \{x_{i1}, \ldots, x_{in}\}$ is an optimal bundle for i at prices p and hence is a solution to the optimal bundle LP (1). The objective function of this LP is

$$\sum_{j=1}^{n} u_{ij} x_{ij}.$$

Next observe that the objective function of the corresponding LP for i under instance I' is

$$\sum_{j=1}^{n} u'_{ij} x_{ij} \;=\; \sum_{j=1}^{n} (s \cdot u_{ij} + h) x_{ij} \;=\; h + s \cdot \sum_{j=1}^{n} u_{ij} x_{ij},$$

where the last equality follows from the fact that $\sum_{j=1}^{n} x_{ij} = 1$. Therefore, the objective function of the second LP is obtained from the first by scaling and shifting. Furthermore, since the constraints of the two LPs are identical, the optimal solutions of the two LPs are the same. Finally, for each $i \in A$: the bundle under allocation x is a minimum cost optimal bundle for I if and only if it is also for I'. The lemma follows. ◀

Next, let u_i be bi-valued with the two values being $0 \le a < b$. Obtain u'_i from u_i by replacing a by 0 and b by 1. Then, u'_i is equivalent to u_i, with the shift and scaling being a and $b - a$, respectively. Therefore the bi-valued instance can be transformed to a unit instance, with both having the same equilibria. Now using Lemma 12 we get:

▶ **Theorem 15.** *There is a strongly polynomial time algorithm for the bi-valued utilities case of the Hylland-Zeckhauser scheme.*

5 Characterizing Optimal Bundles

In this section we give a characterization of optimal bundles for an agent at given prices p which are not necessarily equilibrium prices. This characterization will be used in Section 6.

Notation. For each agent i, let $G^*_i \subseteq G$ denote the set of goods from which i derives maximum utility, i.e., $G^*_i = \arg\max_{j \in G}\{u_{ij}\}$. With respect to an allocation x, let $B_i = \{j \in G \mid x_{ij} > 0\}$, i.e., the set of goods in i's bundle.

We identify the following four types of optimal bundles.

Type A bundles: $\alpha_i = 0$ and $\text{cost}(i) < 1$.

By complementary slackness, optimal goods will satisfy $u_{ij} = \mu_i$ and suboptimal goods will satisfy $u_{ij} < \mu_i$. Hence the set of optimal goods is G^*_i and $B_i \subseteq G^*_i$. Note that the prices of goods in B_i can be arbitrary, as long as $\text{cost}(i) < 1$.

Type B bundles: $\alpha_i = 0$ and $\text{cost}(i) = 1$.

The only difference from the previous type is that $\text{cost}(i)$ is exactly 1.

Type C bundles: $\alpha_i > 0$ and all optimal goods for i have the same utility.

Recall that good j is optimal for i if[6] $\alpha_i p_j = u_{ij} - \mu_i$. Suppose goods j and k are both optimal. Then $u_{ij} = u_{ik}$ and $\alpha_i p_j = u_{ij} - \mu_i = u_{ik} - \mu_i = \alpha_i p_k$, i.e., $p_j = p_k$. Since $\alpha_i > 0$,

[6] Note that under this case, optimal goods are not necessarily maximum utility goods; the latter may be suboptimal because their prices are too high.

by complementary slackness, $\text{cost}(i) = 1$. Further, since $\text{size}(i) = 1$, we get that each optimal good has price 1.

Type D bundles: $\alpha_i > 0$ and not all optimal goods for i have the same utility.

Suppose goods j and k are both optimal and $u_{ij} \neq u_{ik}$. Then $\alpha_i p_j = u_{ij} - \mu_i \neq u_{ik} - \mu_i = \alpha_i p_k$, i.e., $p_j \neq p_k$. Therefore optimal goods have at least two different prices. Since $\alpha_i > 0$, by complementary slackness, $\text{cost}(i) = 1$. Further, since $\text{size}(i) = 1$, there must be an optimal good with price more than 1 and an optimal good with price less than 1. Finally, if good z is suboptimal for i, then $\alpha_i p_z < u_{iz} - \mu_i$.

6 An Example Having Only Irrational Equilibria

Our example has 4 agents $A_1, \ldots, A_4$ and 4 goods $g_1, \ldots, g_4$[7]. The agents' utilities for the goods are given in Table 1, with rows corresponding to agents and columns to goods.

Table 1 Agents' utilities.

	g_1	g_2	g_3	g_4
A_1	2	4	0	8
A_2	2	3	0	8
A_3	2	0	5	0
A_4	0	4	5	0

Thus, agents A_1 and A_2 like, to varying degrees, three goods only, g_1, g_2, g_4, while agents A_3 and A_4 like two goods each, $\{g_1, g_3\}$ and $\{g_2, g_3\}$, respectively. The precise values of the utilities are not that important; the important aspects are: which goods each agent likes, the order between them, and the ratios $\frac{u_{14}-u_{12}}{u_{12}-u_{11}}$ and $\frac{u_{24}-u_{22}}{u_{22}-u_{21}}$. Notice that the latter are unequal.

We show that this example has a unique equilibrium solution with minimum price 0. In this solution, good g_1 has price 0, and all the other goods have positive irrational values. Agents A_1, A_3 and A_4 buy the goods that they like, and A_2 buys g_1 and g_4 only.

Specifically, we show the following:

▶ **Theorem 16.** *The stated instance has a unique equilibrium. In this equilibrium, the allocations to agents and prices of goods, other than the zero-priced good, are all irrational numbers. The prices are as follows:*
$p_1 = 0, \quad p_2 = (23 - \sqrt{17})/32, \quad p_3 = (9 + \sqrt{17})/8, \quad p_4 = (69 - 3\sqrt{17})/32.$

Letting $r_i = |1 - p_i|$, the allocations x_{ij} of each good g_j to each agent A_i are as follows:

$$A_1: \quad x_{11} = 1 - \frac{r_3}{1+r_3} - \frac{r_4}{1+r_4}, \qquad x_{12} = \frac{r_2}{r_2+r_3}, \qquad x_{13} = 0, \qquad x_{14} = \frac{r_4}{1+r_4}$$

$$A_2: \quad x_{21} = \frac{r_4}{1+r_4}, \qquad x_{22} = 0, \qquad x_{23} = 0, \qquad x_{24} = \frac{1}{1+r_4}$$

$$A_3: \quad x_{31} = \frac{r_3}{1+r_3}, \qquad x_{32} = 0, \qquad x_{33} = \frac{1}{1+r_3}, \qquad x_{34} = 0$$

$$A_4: \quad x_{41} = 0, \qquad x_{42} = \frac{r_3}{r_2+r_3}, \qquad x_{43} = \frac{r_2}{r_2+r_3}, \qquad x_{44} = 0$$

[7] It can be shown, by analyzing relations in the bipartite graph on agents and goods with edges corresponding to non-zero allocations, that any instance with 3 agents and 3 goods and rational utilities has a rational equilibrium.

The proof of the theorem is given in the full paper. Even for a small instance like this one, the analysis of the HZ equilibria is not simple. We outline here the main steps. Consider any equilibrium with minimum price 0. We first analyze qualitatively the prices of the goods with respect to 0 and 1: We show that the zero-priced good must be good 1; good 2 must have price strictly between 0 and 1, and goods 3 and 4 must have price strictly greater than 1. Next we characterize qualitatively which goods are bought in non-zero quantity by each agent: We show that every agent buys only goods that she likes. Agents A_3 and A_4 buy both goods that they like, but only one of agents A_1, A_2 buys all three goods that she likes. If A_1 buys g_1, g_2, g_4, then A_2 buys only g_1, g_4. If A_2 buys g_1, g_2, g_4, then A_1 buys only g_2, g_4. Finally, we analyze quantitatively each of these two cases. In each case, we show that the prices satisfy a system of (nonlinear) equations, and the system has a unique nonnegative solution. Furthermore, the prices determine uniquely the allocation. In the first case where A_1 buys all three goods that she likes, the unique solution and the corresponding allocations are as given in the Theorem. In the second case where agent A_2 buys all three goods, we show that there is no equilibrium: the unique prices that satisfy the equations yield allocations that violate the market clearance conditions (a good is oversold). The analysis uses heavily the primal-dual complementary slackness equations. We refer to the full paper for the details.

▶ Remark 17. Observe that in the equilibrium, the allocations of all four agents are irrational even though each one of them spends their dollar completely and the allocations form a fractional perfect matching, i.e., add up to 1 for each good and each agent.

7 Membership of Exact Equilibrium in FIXP

In this section, we will prove that the problem of computing an HZ equilibrium lies in the class FIXP, which was introduced in [16]. This is the class of problems that can be cast, in polynomial time, as the problem of computing a fixed point of an algebraic Brouwer function. Recall that basic complexity classes, such as P, NP, NC and #P, are defined via machine models. For the class FIXP, the role of "machine model" is played by one of the following: a straight line program, an algebraic formula, or a circuit; further it must use the standard arithmetic operations of +, - * /, min and max. We will establish membership in FIXP using straight line programs. Such a program should satisfy the following:

1. The program does not have any conditional statements, such as if ... then ... else.
2. It uses the standard arithmetic operations of +, - * /, min and max.
3. It never attempts to divide by zero.

A *total problem* is one which always has a solution, e.g., Nash equilibrium and Hylland-Zeckhauser equilibrium. A total problem is in FIXP if there is a polynomial time algorithm which given an instance I of length $|I| = n$, outputs a polynomial sized straight line program which computes a function F_I on a closed, convex, real-valued domain $D(n)$ satisfying: each fixed point of F_I is a solution to instance I.

Let p and x denote the allocation and price variables. We will give a function F over these variables and a closed, compact, real-valued domain D for F. The function will be specified by a polynomial length straight line program using the algebraic operations of $+, -, *, /$, min and max, hence guaranteeing that it is continuous. We will prove that all fixed points of F are equilibrium allocations and prices, hence proving that Hylland-Zeckhauser is in FIXP.

Notation. We denote the set $\{1, \ldots . n\}$ by $[n]$. x_i denotes agent i's bundle. For each agent i, choose one good from G_i^* and denote it by i^*. If e is an expression, we will use $(e)_+$ as a shorthand for $\max\{0, e\}$.

Domain $D = D_p \times D_x$, where D_p and D_x are the domains for p and x, respectively, with $D_p = \{p \mid \forall j \in [n], p_j \in [0, n]\}$ and $D_x = \{x \mid \forall i \in [n], \ \sum_{j \in G} x_{ij} = 1, \text{ and } \forall i, j \in [n], x_{ij} \geq 0\}$.

Let $(p', x') = F(p, x)$. (p, x) can be viewed as being composed of $n+1$ vectors of variables, namely p and for each $i \in [n]$, x_i. Similarly, we will view F as being composed of $n + 1$ functions, F_p and for each $i \in [n]$, F_i, where $p' = F_p(p, x)$ and for each $i \in [n]$, $x'_i = F_i(p, x)$. The straight line programs for F_p and F_i are given in the figures below.

It is easy to see that if F_i alters a bundle, the new bundle still remains in the domain; in particular, $\forall i \in [n], \ \text{size}(i) = 1$. Similarly, it is easy to see that the output of F_p is in the domain D_p.

Requirements on equilibria. Observe that (p, x) will be an equilibrium for the market if, in addition to the conditions imposed by the domain, it satisfies the following:

1. $\forall j \in [n], \ \sum_{i \in A} x_{ij} = 1$.
2. $\forall i \in [n], \ \text{cost}(i) \leq 1$.
3. $\forall i \in [n]$, x_i is an optimal bundle for i. Furthermore, cost(i) is minimum over all optimal bundles.

Function F has been constructed in such a way that if any of these conditions is not satisfied by (p, x), then $F(p, x) \neq (p, x)$, i.e., (p, x) is not a fixed point of F. Equivalently, we show that every fixed point of F must satisfy all these conditions and is therefore an equilibrium. The converse also holds. That is, we show:

▶ **Lemma 18.** *Every fixed point (x, p) of F is an equilibrium of the matching market. Conversely, every equilibrium (p, x), where some good has price 0, is a fixed point of F.*

Algorithm 2 Straight line program for function F_p.

1. For all $j \in [n]$ do: $p_j \leftarrow \min\{n, \max\{0, p_j + \sum_{i \in A} x_{ij} - 1\}\}$
2. $r \leftarrow \min_{j \in [n]}\{p_j\}$
3. For all $j \in [n]$ do: $p_j \leftarrow p_j - r$

The proof of correctness for the function F is nontrivial, and is given in the full paper. The proof makes essential use of the characterization of the optimal bundles from Section 5. We first show that F has the following key property: if (p, x) is a fixed point, then no step of F will change (p, x), i.e., it couldn't be that some step(s) of F change (p, x) and some other step(s) change it back, restoring it to (p, x). This is easy to check for F_p. The proof for F_i is more delicate and uses a potential function argument based on the changes in value$(i) = \sum_j u_{ij} x_{ij}$ and cost$(i) = \sum_j p_j x_{ij}$ caused by any change in the allocation x_i in every step of the algorithm for F_i. Given the key property, we then show that if a pair (x, p) does not satisfy one of the equilibrium requirements, then some step of F will change (x, p), hence (x, p) is not a fixed point. We refer to the full paper for the details of the proofs.

Thus, we have:

▶ **Theorem 19.** *The problem of computing an exact equilibrium for the Hylland-Zeckhauser scheme is in FIXP.*

Algorithm 3 Straight line program for function F_i.

1. $r \leftarrow (\sum_j p_j x_{ij} - 1)_+$.
2. For all $j \in [n]$ do: $x_{ij} \leftarrow \frac{x_{ij} + r \cdot (1-p_j)_+}{1 + r \cdot \sum_k (1-p_k)_+}$
3. $t \leftarrow (1 - \sum_j p_j x_{ij})_+$
4. For all $k \notin G_i^*$ do:
   a. $d \leftarrow \min\{x_{ik}, \frac{t}{n^2}\}$
   b. $x_{ik} \leftarrow x_{ik} - d$
   c. $x_{ii^*} := x_{ii^*} + d$
5. For all pairs j, k of goods s.t. $u_{ij} \le u_{ik}$ do:
   a. $d \leftarrow \min\{x_{ij}, (p_j - p_k)_+\}$
   b. $x_{ij} \leftarrow x_{ij} - d/n$
   c. $x_{ik} \leftarrow x_{ik} + d/n$
6. For all triples j, k, l of goods such that $u_{ij} < u_{ik} < u_{il}$ do:
   a. $d \leftarrow \min\{x_{ik}, ((u_{il} - u_{ik})(p_k - p_j) - (u_{ik} - u_{ij})(p_l - p_k))_+\}$
   b. $x_{ik} \leftarrow x_{ik} - d$
   c. $x_{ij} \leftarrow x_{ij} + \frac{u_{il} - u_{ik}}{u_{il} - u_{ij}} d$
   d. $x_{il} \leftarrow x_{il} + \frac{u_{ik} - u_{ij}}{u_{il} - u_{ij}} d$
7. For all triples j, k, l of goods such that $u_{ij} < u_{ik} < u_{il}$ do:
   a. $d := \min(x_{ij}, x_{il}, ((u_{ik} - u_{ij})(p_l - p_k) - (u_{il} - u_{ik})(p_k - p_j))_+)$
   b. $x_{ik} := x_{ik} + d$
   c. $x_{ij} := x_{ij} - \frac{u_{il} - u_{ik}}{u_{il} - u_{ij}} d$
   d. $x_{il} := x_{il} - \frac{u_{ik} - u_{ij}}{u_{il} - u_{ij}} d$

8 Membership of Approximate Equilibrium in PPAD

In this section we define approximate equilibria, and show that the problem of computing an approximate equilibrium is in PPAD.

First let us scale the utilities of all the agents so that they lie in $[0, 1]$. This can be done clearly without loss of generality without changing the equilibria.

▶ **Definition 20.** *A pair (p, x) of (non-negative) prices and allocations is an* ϵ-approximate equilibrium *for a given one-sided market if:*

1. *The total probability share of each good j is 1 unit, i.e., $\sum_i x_{ij} = 1$.*
2. *The size of each agent i's allocation is 1, i.e., $\text{size}(i) = 1$.*
3. *The cost of the allocation of each agent is at most $1 + \epsilon$.*
4. *The value of the allocation of each agent i is at least $v^*(i) - \epsilon$ where $v^*(i)$ is the value of the optimal bundle for agent i under prices p, i.e. the optimal value of the program: maximize $\text{value}(i)$, subject to $\text{size}(i) = 1$ and $\text{cost}(i) \le 1$. Furthermore, we require that the cost of the allocation x_i is at most $c^*(i) - \epsilon$, where $c^*(i)$ is the minimum cost of a bundle for agent i that has the maximum value $v^*(i)$.*

The corresponding computational problem is: Given a one-sided matching market M and a rational $\epsilon > 0$ (in binary as usual), compute an ϵ-approximate equilibrium for M. Polynomial time in this context means time that is polynomial in the encoding size of the market M and $\log(1/\epsilon)$. We define also a more relaxed version, called a *relaxed ϵ-approximate equilibrium* where condition 1 is relaxed to $|\sum_i x_{ij} - 1| \leq \epsilon$ for all goods j. It is easy to see that the two versions are polynomially equivalent, i.e., if one can be solved in polynomial time then so can the other.

▶ **Proposition 21.** *The problems of computing an ϵ-approximate equilibrium and a relaxed approximate equilibrium are polynomially equivalent.*

Note however, that in general an ϵ-approximate equilibrium may not be close to an actual equilibrium of the matching market. This phenomenon is similar to the case of market equilibria for the standard exchange markets and to the case of Nash equilibria for games.

We will show membership of the approximate equilibrium problem in PPAD by showing that a relaxed approximate equilibrium can be obtained from an approximate fixed point of a variant of the function F defined in Section 7.

▶ **Definition 22.** *A* weak ϵ-approximate fixed point *of a function F (or* weak ϵ-fixed point *for short) is a point x such that $||F(x) - x||_\infty \leq \epsilon$.*

Let $\mathcal{F}$ be a family of functions, where each function F_I in $\mathcal{F}$ corresponds to an instance I of a problem (in our case a one-sided matching market) that is encoded as usual by a string. The function F_I maps a domain D_I, to itself. We assume that D_I is a polytope defined by a set of linear inequalities with rational coefficients which can be computed from I in polynomial time; this clearly holds for our problem. We use $|I|$ to denote the length of the encoding of an instance I (i.e., the length of the string). If x is a rational vector, we use $size(x)$ to denote the number of bits in a binary representation of x.

▶ **Definition 23.** *A family $\mathcal{F}$ of functions is* polynomially computable *if there is a polynomial q and an algorithm that, given the string encoding I of a function $F_I \in \mathcal{F}$ and a rational point $x \in D_I$, computes $F_I(x)$ in time $q(|I| + size(x))$.*
A family $\mathcal{F}$ of functions is polynomially continuous *if there is a polynomial q such that for every $F_I \in \mathcal{F}$ and every rational $\epsilon > 0$ there is a rational δ such that $\log(1/\delta) \leq q(|I|+\log(1/\epsilon))$ and such that $||x - y||_\infty \leq \delta$ implies $||F_I(x) - F_I(y)||_\infty \leq \epsilon$ for all $x, y \in D_I$.*

It was shown in [16] that, if a family of functions is polynomially computable and polynomially continuous, then the corresponding weak approximate fixed point problem (given I and rational $\delta > 0$, compute a weak δ-approximate fixed point of F_I) is in PPAD. The family $\mathcal{F}$ of functions for the online matching market problem defined in Section 7 is obviously polynomially computable. It is easy to check also that it is polynomially continuous.

We will use a variant F' of the function F of Section 7, where the functions F_i for the allocations are modified as follows. Step 5 for all pairs j, k of goods, and steps 6, and 7 for all triples j, k, l are applied all independently in parallel to the allocation that results after step 4. In order for the allocation to remain feasible (i.e. have $x_{ij} \geq 0$ for all i, j), we change line 5a in F'_i to $d \leftarrow \min\{\frac{x_{ij}}{3}, (p_j - p_k)_+\}$, change line 6a to $d \leftarrow \min\{\frac{x_{ik}}{3n^2}, ((u_{il} - u_{ik})(p_k - p_j) - (u_{ik} - u_{ij})(p_l - p_k))_+\}$, and we change line 7a to $d \leftarrow \min\{\frac{x_{ij}}{3n^2}, \frac{x_{il}}{3n^2}, ((u_{ik} - u_{ij})(p_l - p_k) - (u_{il} - u_{ik})(p_k - pj))_+\}$. In this way, a coordinate x_{ij} can be decreased by the operations of step 5 for all pairs j, k at most by $x_{ij}/3$ in total, and the same is true for the total decrease from the operations of steps 6 and 7 for all triples involving good j; therefore, the coordinates x_{ij} remain nonnegative. The function for the prices remains the same as before. All the properties shown in Section 7 for F hold also for F'. The family $\mathcal{F}'$ of these functions F'_I is clearly also polynomially computable and polynomially continuous.

Let I be a given instance of the matching market problem. Every utility u_{ij} is a rational number, without loss of generality in $[0, 1]$, which is given as the ratio of two integers represented in binary. Let m be the maximum number of bits needed to represent a utility. Note that every nonzero u_{ij} is at least $1/2^m$ and the difference between any two unequal utilities is at least $1/2^{2m}$. Given a positive rational ϵ (wlog in $[0, 1]$), let $\delta = \epsilon/(n^{10}2^{6m})$. Note that $\log(1/\delta)$ is bounded by a polynomial in $|I|$ and $\log(1/\epsilon)$. We show the following:

▶ **Lemma 24.** *Every weak δ-approximate fixed point of F'_I is a relaxed ϵ-approximate equilibrium of the market I.*

The proof follows and adapts the proof in Section 7 of the analogous statement for the exact fixed points. The proof is rather involved; we refer to the full paper for the details. As a consequence of Lemma 24 and Proposition 21 we have:

▶ **Theorem 25.** *The problem of computing an ϵ-approximate equilibrium of a given matching market is in PPAD.*

9 Discussion

As stated in the Introduction, a conclusive proof of intractability of the HZ scheme, via either a proof of FIXP-hardness for exact equilibrium or PPAD-hardness for approximate equilibrium, has eluded us. One of the difficulties is the following: Optimal bundles of agents in an HZ equilibrium may include zero-utility zero-priced goods as "fillers" to satisfy the size constraint, e.g., observe their use in our Algorithm for the case of dichotomous utilities. In this easy setting, we knew which were the "filler" goods. However, when faced with a complex instance of HZ, we don't a priori know which zero-utility goods will be used as "fillers". Therefore, even though an agent may have very few positive utility goods, other goods are also in play, thereby giving no "control" on the equilibrium outcome.

We propose exploring the following avenue, in addition to the usual ones, for arriving at evidence of intractability: Relax the notion of polynomial time reducibility suitably and obtain a weaker result than FIXP-hardness or PPAD-hardness.

Other open problems related to our work are: obtain efficient algorithms for computing approximate equilibria, suitably defined, and identify other special cases, besides the bi-valued case, for which equilibrium is easy to compute. Additionally, generalizations and variants of the HZ scheme deserve attention, most importantly to two-sided matching markets [14].

Encouraged by success on the bi-valued utilities case, we considered its generalization to the tri-valued utilities case, in particular, $\{0, \frac{1}{2}, 1\}$ utilities. We believe even this case has instances with only irrational equilibria. Finding such an example or proving rationality is non-trivial and we leave it as an open problem. Furthermore, it will not be surprising if even this case is intractable; resolving this is a challenging open problem.

References

1. Saeed Alaei, Pooya Jalaly Khalilabadi, and Eva Tardos. Computing equilibrium in matching markets. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 245–261, 2017.
2. Saugata Basu, Richard Pollack, and MF Roy. A new algorithm to find a point in every cell defined by a family of polynomials. *Quantifier Elimination and Cylindrical Algebraic Decomposition, B. Caviness and J. Johnson eds., Springer-Verlag, to appear*, 1995.
3. Garrett Birkhoff. Tres observaciones sobre el algebra lineal. *Univ. Nac. Tucuman, Ser. A*, 5:147–154, 1946.

4 Anna Bogomolnaia and Hervé Moulin. A new solution to the random assignment problem. *Journal of Economic theory*, 100(2):295–328, 2001.

5 Anna Bogomolnaia and Hervé Moulin. Random matching under dichotomous preferences. *Econometrica*, 72(1):257–279, 2004.

6 Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

7 Xi Chen, Decheng Dai, Ye Du, and Shang-Hua 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*, pages 273–282. IEEE, 2009.

8 Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM (JACM)*, 56(3):1–57, 2009.

9 Xi Chen, Dimitris Paparas, and Mihalis Yannakakis. The complexity of non-monotone markets. *J. ACM*, 64(3):20:1–20:56, 2017. `doi:10.1145/3064810`.

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

11 Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

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

13 Nikhil R Devanur, Christos H Papadimitriou, Amin Saberi, and Vijay V Vazirani. Market equilibrium via a primal–dual algorithm for a convex program. *Journal of the ACM (JACM)*, 55(5):22, 2008.

14 Federico Echenique, Antonio Miralles, and Jun Zhang. Constrained pseudo-market equilibrium. *arXiv preprint*, 2019. `arXiv:1909.05986`.

15 Federico Echenique, Antonio Miralles, and Jun Zhang. Fairness and efficiency for probabilistic allocations with endowments. *arXiv preprint*, 2019. `arXiv:1908.04336`.

16 K. Etessami and M. Yannakakis. On the complexity of Nash equilibria and other fixed points. *SIAM J. Comput.*, 39(6):2531–2597, 2010.

17 Simons Institute for the Theory of Computing. Online and matching-based market design, 2019. URL: `https://simons.berkeley.edu/programs/market2019`.

18 Jugal Garg, Ruta Mehta, and Vijay V. Vazirani. Dichotomies in equilibrium computation and membership of PLC markets in FIXP. *Theory of Computing*, 12(1):1–25, 2016.

19 Jugal Garg, Ruta Mehta, Vijay V Vazirani, and Sadra 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*, pages 890–901, 2017.

20 Jugal Garg, Thorben Tröbst, and Vijay V Vazirani. An Arrow-Debreu extension of the Hylland-Zeckhauser scheme: Equilibrium existence and algorithms. *arXiv preprint*, 2020. `arXiv:2009.10320`.

21 Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2. Springer Science & Business Media, 2012.

22 Yinghua He, Antonio Miralles, Marek Pycia, and Jianye Yan. A pseudo-market approach to allocation with priorities. *American Economic Journal: Microeconomics*, 10(3):272–314, 2018.

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

24 Kamal Jain. A polynomial time algorithm for computing an Arrow–Debreu market equilibrium for linear utilities. *SIAM Journal on Computing*, 37(1):303–318, 2007. `doi:10.1137/S0097539705447384`.

25 Phuong Le. Competitive equilibrium in the random assignment problem. *International Journal of Economic Theory*, 13(4):369–385, 2017.

26 Andy McLennan. Efficient disposal equilibria of pseudomarkets. In *Workshop on Game Theory*, volume 4, page 8, 2018.

27 Hervé Moulin. Fair division in the age of internet. *Annual Review of Economics*, 2018.

28 C. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *JCSS*, 48(3):498–532, 1994.

29 Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. Pairwise kidney exchange. *Journal of Economic theory*, 125(2):151–188, 2005.

30 Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of mathematical economics*, 1(1):23–37, 1974.

31 V. V. Vazirani. The notion of a rational convex program, and an algorithm for the Arrow-Debreu Nash bargaining game. *JACM*, 59(2), 2012.

32 V. V. Vazirani and M. Yannakakis. Market equilibria under separable, piecewise-linear, concave utilities. *JACM*, 58(3), 2011.

33 Vijay V. Vazirani. Efficient, strategyproof mechanisms for one-sided matching markets. Manuscript, 2020.

34 John Von Neumann. A certain zero-sum two-person game equivalent to the optimal assignment problem. *Contributions to the Theory of Games*, 2(0):5–12, 1953.

35 Mihalis Yannanakis. Unpublished, 2013.