

COMPUTATIONAL COMPLEXITY OF THE HYLLAND-ZECKHAUSER MECHANISM FOR ONE-SIDED MATCHING MARKETS *

VIJAY V. VAZIRANI[†] AND MIHALIS YANNAKAKIS[‡]

Abstract. In 1979, Hylland and Zeckhauser [26] gave a simple and general mechanism for a one-sided matching market, given cardinal utilities of agents over goods. They use the power of a pricing mechanism, which endows their mechanism with several desirable properties – it produces an allocation that is Pareto optimal and envy-free, and the mechanism is incentive compatible in the large. It therefore 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 characterize the computational complexity of this mechanism.

We present the following results:

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 this problem is not in PPAD.
3. A proof of membership of the problem in the class FIXP; as a corollary we get that an HZ equilibrium can always be expressed via algebraic numbers. For this purpose, we give a new proof of the existence of an HZ equilibrium using Brouwer’s fixed point theorem; the proof of Hylland and Zeckhauser used Kakutani’s fixed point theorem, which is more involved.
4. A proof of membership of the problem of computing an approximate HZ equilibrium in the class PPAD.

Key words. Computational Complexity; Hylland-Zeckhauser Mechanism; One-Sided Matching Markets; Brouwer’s fixed point theorem; PPAD; FIXP; dichotomous utilities

MSC codes. 68Q25, 68W40, 68R10, 68U05

Abstract. In 1979, Hylland and Zeckhauser [26] gave a simple and general mechanism for a one-sided matching market, given cardinal utilities of agents over goods. They use the power of a pricing mechanism, which endows their mechanism with several desirable properties – it produces an allocation that is Pareto optimal and envy-free, and the mechanism is incentive compatible in the large. It therefore 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 characterize the computational complexity of this mechanism.

We present the following results:

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 this problem is not in PPAD.
3. A proof of membership of the problem in the class FIXP; as a corollary we get that an HZ equilibrium can always be expressed via algebraic numbers. For this purpose, we give a new proof of the existence of an HZ equilibrium using Brouwer’s fixed point theorem; the proof of Hylland and Zeckhauser used Kakutani’s fixed point theorem, which is more involved.
4. A proof of membership of the problem of computing an approximate HZ equilibrium in the class PPAD.

In subsequent work [8], the problem of computing an approximate HZ equilibrium was shown to be PPAD-hard, thereby establishing it to be PPAD-complete. We leave open the (difficult) question of determining if computing an exact HZ equilibrium is FIXP-hard. We also give pointers to the substantial body of work on cardinal-utility matching markets which followed [38].

*Submitted to the editors 10/3/2024.

Funding: Vazirani was supported in part by NSF grants CCF-1815901 and CCF-2230414. Yannakakis was supported in part by NSF grants CCF-1703925, CCF-2212233 and CCF-2332922.

[†]University of California, Irvine (vazirani@ics.uci.edu).

[‡]Columbia University (mihalis@cs.columbia.edu).

1. Introduction. A *one-sided matching market* consists of n agents, n indivisible goods and the preferences of the agents for the goods, provided in a suitable form. The problem is to design a mechanism which produces allocations satisfying desirable optimality and fairness properties, such as Pareto optimality and envy-freeness. In addition, the mechanism should have desirable game theoretic properties, such as strategy-proofness. This problem arises in various settings, where individuals have to be assigned to available positions, for example, assignment of students to dormitories, people to committees, workers to tasks, etc.

Since different agents may prefer the same goods, randomization is often needed to achieve fairness. A randomized mechanism yields a *random assignment*, which matches each agent i to good j with some probability $x_{ij} \in [0, 1]$; the probabilities x_{ij} form a doubly stochastic matrix, i.e., a fractional perfect matching in the bipartite graph between agents and goods. In some applications in fact, the goods may represent resources that can be shared or divided (e.g. time-shared housing units), in which case the quantities x_{ij} represent the shares of the agents in the goods.

In a brilliant and by-now classic paper, Hylland and Zeckhauser [26] studied the one-sided matching problem and gave a simple and general mechanism, henceforth called HZ. In their model, the preferences of agents are stated using von Neumann-Morgenstern utilities u_{ij} , specifying the utility of agent i for each good $j \in [n]$ (these are called *cardinal preferences*). The HZ mechanism produces a fractional allocation, also called lotteries, which need to be rounded using randomization, that has several desirable properties: it is ex-ante Pareto optimal, i.e. there is no other random assignment in which some agent has strictly higher expected utility and no agent has lower utility; it is ex-ante envy-free, i.e. no agent prefers the allocation of any other agent [26]; and it is incentive compatible in the large [24].

The HZ mechanism uses the power of a pricing mechanism to produce its allocation: The goods are first rendered divisible by viewing each one as one unit of probability shares, and each agent is assumed to have one dollar of fake money. An *HZ equilibrium* consists of prices for the goods and allocations x_{ij} for the agents so that each agent gets a total of one unit of probability shares across all goods; moreover, this should be a utility-maximizing bundle w.r.t. the prices. Hylland and Zeckhauser showed, using Kakutani's fixed point theorem, that an HZ equilibrium always exists. The equilibrium allocation can be viewed as a fractional perfect matching in the complete bipartite graph over n agents and n goods. Using the Birkhoff-von Neumann Theorem, this fractional allocation can be rounded to an integral one whose *ex ante* utility for each agent is the same as the expected utility of her bundle of probability shares. Furthermore, if each agent is allocated a minimum cost utility-maximizing bundle, then Hylland and Zeckhauser showed that the allocation is Pareto optimal; it is envy-free simply because it was produced by a pricing mechanism where all the agents have the same budget.

Despite the prominent standing of the HZ mechanism in economics and the large body of work it has spawned on pseudo-markets, its computational complexity remains unexplored. The conference version of the current paper, [38], initiated work on filling this gap; see Section 1.1.3 for a summary of work on cardinal-utility matching markets that happened in its wake. We note that even though HZ was discovered forty-five years ago, the significance of matching markets has only grown in recent years, with ever more diverse and impactful ones being launched into our economy, e.g., see [20, 15].

HZ can be viewed as a marriage between fractional perfect matching and a linear Fisher market, both of which admit strongly polynomial time algorithms, and fur-

thermore combinatorial ones (i.e., not requiring an LP or a convex program solver). These facts had enticed researchers over the years to seek an efficient algorithm for it. Observe that the HZ market model differs from a linear Fisher market in only one respect, namely in the former, each agent is allocated exactly one unit of goods and in the latter, there is no such restriction. The underlying reason for the polynomial time solvability of the latter [14] is the property of weak gross substitutability¹. 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-complete² [37]. For the case of HZ, this property is destroyed by the restriction that each agent be allocated exactly one unit of goods, see Example 3.4 and Remark 3.5. Indeed this is a key stumbling block one faces while attempting to obtain a polynomial time algorithm for HZ. The only known method for computing an HZ equilibrium is using an algebraic cell decomposition [2], which requires exploration of $n^{5n^2} > 10^{80}$ cells for a problem of size $n = 5$; note that this number is more than the number of atoms in the known universe!

The problem of computing an HZ equilibrium is a total search problem, where the existence of a solution is proved through a fixed point theorem. Two complexity classes for such problems (e.g. Nash equilibria and market equilibria) are PPAD and FIXP. We note that a crucial requirement for membership in PPAD is that there is always a rational solution if all parameters of the instance are rational numbers. In this paper we show that this is not true for HZ equilibria: we give an example consisting of four agents and goods that has only irrational equilibria (in fact, a unique equilibrium), see Section 5.

This irrationality of solutions suggests that the appropriate class for the problem of computing an HZ equilibrium is the class FIXP. The proof in [26], showing the existence of an equilibrium, uses Kakutani's fixed point theorem and it 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 an HZ equilibrium. We define a suitable Brouwer function (a continuous function from a convex, compact set to itself) which uses elementary arithmetic operations to improve the optimality or the feasibility of the current prices and allocations, in case they do not form an equilibrium. Consequently the only fixed points of the adjustment mechanism are equilibria. This yields our proof of membership in FIXP, see Section 6. A corollary of this proof is that an HZ equilibrium can always be expressed via algebraic numbers. Showing FIXP-hardness remains open.

We define in Section 7 a notion of approximate HZ equilibrium: a set of prices and allocations, where the agent's allocations are only approximately optimal. Approximate HZ equilibria are approximately Pareto optimal and envy-free. Using the same Brouwer function as in the FIXP proof, we show membership of the approximate HZ equilibrium problem in PPAD. In subsequent work, [8] showed PPAD-hardness for the approximate HZ equilibrium problem, hence the problem is PPAD-complete.

We give one positive result for HZ: we give a combinatorial, strongly polynomial time algorithm for the *dichotomous case*, i.e., all utilities are 0/1. This involves melding a bipartite graph perfect matching algorithm with ideas from the combinatorial algorithm of [14] for the linear Fisher market, see Section 4. We also extend this result to solving 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

¹Namely, if you increase the price of one good, the demand of another good cannot decrease.

²Independently, PPAD-hardness, though not PPAD membership, was also established in [12].

different for different agents.

However, this approach did not extend any further, say from the bi-valued utilities case to tri-valued utilities, in particular, to the case of $\{0, \frac{1}{2}, 1\}$ utilities. Indeed, this case appears to be intractable and its status is discussed in Section 8. Note that the case of four-valued utilities is PPAD-hard, as shown in [8].

1.1. Related work.

1.1.1. Past Work on Matching Markets. Matching markets can be classified into two classes based on the type of agents' preferences: *cardinal*, where each agent i specifies a numerical utility u_{ij} for each good j , and *ordinal*, where each agent specifies an ordering of the goods. Whereas ordinal-based matching markets are well-developed from the viewpoint of both theory and practice, the status of cardinal-based matching markets has been in a state a flux for some time; however, very recent developments seem to have restored a semblance of order, see Section 1.1.3. These two types of preferences have their individual pros and cons: Whereas ordinal preferences are easier to elicit, cardinal preferences are more expressive, thereby producing higher quality allocations and leading to significant gain in efficiency, e.g., [27] give a striking example with n agents and goods, in which an allocation under cardinal utilities improves every agent by a factor of $\Theta(\log n)$ over the (coarse) allocation made under ordinal information. Furthermore, it is shown in [18] that no ordinal-based allocation can approximate the social welfare (the sum of the agents' utilities) to a factor better than $\Theta(\sqrt{n})$.

In some applications, cardinal utilities are natural and easy to elicit and contain critical information which is lost after conversion to their (coarse) ordinal counterparts, such as in course allocation, e.g., suppose a student is asked to rate the desirability of being assigned to one of four courses, A, B, C, D on the scale $[0, 100]$ and reports $A(100); B(95); C(2); D(0)$. This translates to the ordinal preference list of $A > B > C > D$. In contrast, cardinal utilities of $A(100); B(5); C(2); D(0)$ are qualitatively different, yet lead to the same ordinal list. Assume further that these are utilities of two different students and course A has only one vacant seat. A mechanism using cardinal utilities has sufficient information to produce the desirable allocation of course B for the first student and course A for the second. On the other hand, under ordinal utilities, the second student has a 50% chance of being assigned course B , leading to a poor allocation³.

Important mechanisms for matching markets under ordinal preferences, together with their properties, are the following. For two-sided matching markets, the Gale-Shapley Deferred Acceptance algorithm finds a stable matching, i.e, a matching which is in the core of this game. The algorithm is incentive compatible for the proposing side. For one-sided matching markets, the famous Top Trading Cycles mechanism is an algorithm for reallocating goods which is Pareto efficient (optimal), strategyproof and core-stable; it was discovered by Gale and reported in [34]. Two prominent randomized mechanisms for the one-sided matching problem with ordinal preferences are Random Priority (also known as Random Serial Dictatorship) and Probabilistic Serial [4]. Random Priority [31] is strategy-proof, ex-post Pareto efficient (but not ex-ante), and is not envy-free. Probabilistic Serial [4] is envy-free, it is 'ordinally efficient' (a notion of efficiency appropriate for ordinal preferences that lies in strength between ex-ante and ex-post Pareto efficiency), and satisfies a weak form of strategy-proofness.

³Today, there is even a company, Cognomos, for course allocation, based on cardinal utilities and a CEEI market.

Both Random Priority and Probabilistic Serial are simple randomized mechanisms that can be computed efficiently (in polynomial time).

Before the conference version of this paper [38], only the following computational result on the HZ mechanism was known: using the algebraic cell decomposition technique of [2], [1] gave a polynomial time algorithm for computing an equilibrium for the case that n is a fixed constant. As stated in the Introduction, this algorithm is far from practical even for small values of n .

An immediate generalization of HZ is its exchange extension in which each agent's initial endowment consists of one unit of goods and after trading, each agent desires one unit of goods. [26] gave an example showing that this model may not even admit an equilibrium. That essentially put an end to the exploration of more general cardinal-utility matching markets, other than two exceptions.

First, Echenique et al. [16], considered a hybrid model that is a convex combination of HZ and exchange HZ, and showed that if the convex combination has a non-zero extent of the first model, then equilibrium exists. [23] defined an ϵ -approximate Arrow-Debreu extension of HZ and, using the previous result, proved existence of equilibrium⁴. The second was work on the dichotomous utilities case: For two-sided matching markets under symmetric⁵ dichotomous utilities, existence of equilibrium was shown in [5] and a polynomial time algorithm for finding an equilibrium follows by applying the methods of Section 4 of the current paper. For non-bipartite matching markets under dichotomous utilities, existence of equilibrium was shown in [33] and a polynomial time algorithm for finding an equilibrium was given in [29].

As stated in the Introduction, HZ led to pseudo-markets for a number of applications, e.g., see [6, 24, 28, 30], which use the notion of *competitive equilibrium under equal incomes (CEEI)*.

1.1.2. Past Work on Complexity Classes PPAD and FIXP. Several papers have established membership and hardness in PPAD and FIXP for equilibrium computation problems in different settings. The quintessential complete problem for PPAD is 2-Nash [13, 9] and that for FIXP is multiplayer Nash equilibrium [17]. For the latter problem, computing an approximate equilibrium is PPAD-complete [13].

For the case of market equilibrium, there are two parallel streams of results in the economics literature: one assumes that an excess demand function is given and the other assumes a specific class of utility functions. [17] 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 [40] and [21], respectively. This was followed by a proof of FIXP-hardness for Arrow-Debreu markets with Leontief and PLC utilities [22]. For the case of Arrow-Debreu markets with CES (constant elasticity of substitution) utility functions, [11] 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 [11]. Computing an (exact or approximate) equilibrium under

⁴Furthermore, they proved that this equilibrium satisfies Pareto optimality, approximate envy-freeness, and approximate weak core stability.

⁵Under symmetric dichotomous utilities, if i and j are on different sides of the bipartition, then the utility of i for j is the same as that of j for i .

separable, piecewise-linear, concave (SPLC) utilities for Arrow-Debreu and Fisher markets is also known to be PPAD-complete [10, 12, 37].

1.1.3. The Current Status of Cardinal-Utility Matching Markets.

In this section, we will summarize the developments that followed [38]. First, Chen et al. [8], proved that computing an approximate HZ equilibrium is PPAD-hard. That naturally raised the question of finding alternative mechanisms, especially because of the ever-increasing impact of matching markets in the economy [20, 15] and the advantage of cardinal-utility matching markets over ordinal ones, see Section 1.1.1. Fortunately, alternatives to pricing-based mechanisms for market models had been explored in the past: [36] gave a Nash-bargaining-based mechanism for the linear Arrow-Debreu model (which was traditionally addressed via the pricing mechanism), and it led to the paper [25].

[25] addressed two issues: the intractability of HZ and the paucity of matching market models, in sharp contrast with General Equilibrium Theory, which had defined and extensively studied several fundamental market models to address a number of specialized and realistic situations. [25] defined a rich collection of Nash-bargaining-based matching market models, not only one-sided but also two-sided, and not only in Fisher but also in the Arrow-Debreu setting. Since the Nash bargaining solution is captured by a convex program, these models can be solved in polynomial time using the ellipsoid algorithm. Since the latter is not time-efficient, [25] demonstrated the practicality of these models by giving very fast implementations, using Frank Wolfe and cutting plane algorithms, which solved very large instances, with 20,000 agents and goods, in minutes on a laptop, even for a two-sided matching market.

The next question was determining whether these models match the nice game-theoretic properties of HZ. The solution to a Nash bargaining game is Pareto optimal; however it is neither envy-free nor incentive compatible. Recently [35] showed that the one-sided linear-utilities Nash-bargaining-based model satisfies envy-freeness within factor two and incentive compatibility within factor two; moreover, both results are tight. Additionally, via a reduction from HZ, they showed that the problem of finding an envy free and Pareto optimal allocation in a one-sided market is PPAD-hard; membership of this problem in PPAD was established by [7]. Hence the properties of this model are almost as good as is feasible.

Another question is whether two-sided markets admit envy free and Pareto optimal allocations. [5] gave a positive answer for the case of symmetric dichotomous utilities. However, [35] show that on relaxing either of the conditions, symmetry or dichotomous utilities, such an allocation may not exist⁶. In contrast, the Nash bargaining approach easily yields models for these and more general matching market settings [25], and even for a non-bipartite matching market model [32], which may have applications to roommate sharing. Finally, [32] obtained efficient combinatorial algorithms with proven running times for these models, using the techniques of multiplicative weights update and conditional gradient descent.

A couple of recent papers provide alternative ways of obtaining some of our results: First, [19] prove membership of HZ in FIXP by first developing machinery for proving FIXP membership of problems by constructing a ‘pseudo-gate’ (i.e., a FIXP sub-circuit) that solves linear or convex programs, suitably presented. They then use such a pseudo-gate as a gadget in the construction of a FIXP circuit for various

⁶[35] give examples of two-sided matching markets with symmetric $\{0, 1, 2\}$ utilities, and those with asymmetric dichotomous utilities for which such allocations don’t exist, hence precluding pricing-based mechanisms.

problems, including HZ. In contrast, our proof of membership of HZ in FIXP is more direct, since we construct a circuit for a Brouwer function that adjusts prices and allocations if they do not form an HZ equilibrium. An advantage of our approach is that the same function is used to show membership of the approximate HZ equilibrium problem in PPAD. Second, [23] gave a rational convex program⁷ for the dichotomous case of HZ, thereby implying its polynomial time solvability via an LP-solver. Once again, our approach has advantages: it is more efficient and it gives an insight into the combinatorial structure of the problem.

Organization of the paper:

Section 2 describes the Hylland-Zeckhauser mechanism in more detail and Section 3 gives basic properties of optimal allocations and prices. Section 4 presents our polynomial-time algorithm for the case of bi-valued utilities. Section 5 presents an example instance with four agents and goods whose (unique) equilibrium is irrational. Section 6 shows membership of the problem of computing an HZ equilibrium in the class FIXP. Section 7 defines approximate HZ equilibria and shows that their computation is in the class PPAD. Appendix A proves a sufficient condition for the existence of rational equilibria and uses it to show rationality for the case of three goods.

2. The Hylland-Zeckhauser mechanism . Hylland and Zeckhauser [26] 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 goods⁸. The mechanism will allocate exactly one good to each agent and will have the following properties:

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

The Hylland-Zeckhauser mechanism 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-Morgenstern 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 $A = \{1, 2, \dots, n\}$ of n agents and a set $G = \{1, 2, \dots, m\}$ 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} , $1 \leq j \leq m$ 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 mechanism 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,

⁷i.e., a non-linear convex program which has a rational optimal solution as long as all parameters are rational numbers [36].

⁸The model in the original paper [26] is expressed in terms of assigning n individuals to m jobs, with the requirement that exactly s_j individuals must be assigned to job j , where $\sum_{j \in m} s_j = n$. This model is equivalent to the above formulation, where the n goods correspond to the n job positions.

$\sum_j u_{ij}x_{ij}$ is the *expected utility* accrued by agent i . Each good j has price $p_j \geq 0$ 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. 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 mechanism, 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 2.1. 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

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

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

DEFINITION 2.2. (Hylland and Zeckhauser [26]) 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 $\text{size}(i) = 1$, $\text{cost}(i) \leq 1$, and lastly, $\text{cost}(i)$ is smallest among all utility-maximizing bundles of i .

An allocation (fractional perfect matching) x is *Pareto optimal* if it is not dominated by any other allocation, i.e., there is no allocation y such that $\sum_j u_{ij}y_{ij} \geq \sum_j u_{ij}x_{ij}$ for all agents i , with the inequality strict for at least one agent i . An allocation x is *envy-free* if $\sum_i u_{ij}x_{ij} \geq \sum_j u_{ij}x_{kj}$ for all agents i, k , i.e. no agent i envies the allocation of another agent k .

Using Kakutani's fixed point theorem, Hylland and Zeckhauser showed the following:

THEOREM 2.3. [Hylland and Zeckhauser [26]] Every instance of the one-sided market defined above admits an equilibrium; moreover, the corresponding allocation is Pareto optimal and envy-free.

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 mechanism 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 [39] 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 [26], since the lottery over goods is Pareto optimal *ex ante*, the integral allocation, viewed stochastically, will also be Pareto optimal *ex post*.

A randomized mechanism is called *ex-ante* Pareto optimal (efficient) if, for every instance, its allocation x is Pareto optimal. It is called *ex-post* Pareto optimal if, for every instance, every integral matching that it generates is Pareto optimal (not dominated by any other integral matching). Ex-ante Pareto optimality implies ex-post, but not vice-versa [26]; it is possible that a randomized mechanism is ex-post Pareto optimal but there is another mechanism that yields strictly higher expected utility for all agents on some instances.

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 2.4. 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.

$$(3.1) \quad \max \quad \sum_j x_{ij} u_{ij}$$

$$(3.2) \quad \text{s.t.}$$

$$(3.3) \quad \sum_j x_{ij} = 1$$

$$(3.4) \quad \sum_j x_{ij} p_j \leq 1$$

$$(3.5) \quad \forall j \quad x_{ij} \geq 0$$

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

$$\begin{aligned}
(3.6) \quad & \min \quad \alpha_i + \mu_i \\
(3.7) \quad & \text{s.t.} \\
(3.8) \quad & \forall j \quad \alpha_i p_j + \mu_i \geq u_{ij} \\
(3.9) \quad & \alpha_i \geq 0
\end{aligned}$$

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 3.1. 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 3.2. 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 $\sum_{j \in G} x_{ij} = 1$. \square

Using Lemma 3.2, 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 the condition that the minimum price of a good is zero, thereby fixing the scale. Observe that the main goal of the Hylland-Zeckhauser mechanism 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 [26] and can lead to simplifying the equilibrium computation problem as shown in Remark 3.3.

Remark 3.3. We remark that on the one hand, the offset μ_i is a key enabler in constructing optimal bundles and on the other, it is also a main source of difficulty in computing equilibria for the HZ mechanism. We identify 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 [14].

Finally, we extend Example 3.1 to illustrate that optimal allocations for the Hylland-Zeckhauser model do not satisfy the weak gross substitutes (WGS) condition in general. This is done in Example 3.4 and in Remark 3.5.

EXAMPLE 3.4. *In Example 3.1, 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 .*

Remark 3.5. Let us extract out the main idea behind Example 3.4 in a simple setting where we have taken the liberty of ignoring lower order terms in ϵ . Suppose agent i has positive utilities for goods j and k , with $u_{ij} = 100$, $u_{ik} = \epsilon$, where $\epsilon > 0$ is a very small number. Suppose their prices are $p_j = 100$, $p_k = \epsilon$. Clearly an optimal bundle involves spending almost all the money on good j to maximize utility and a small amount of money on good k to ensure that agent i gets 1 unit of goods. Thus agent i spends (approximately) $1 - \epsilon$ and ϵ on j and k , respectively.

Next, assume that the price of good k is raised to 2ϵ . Now, the optimal bundle involves spending (approximately) $1 - 2\epsilon$ and 2ϵ on j and k , respectively. Hence, on raising the price of k , the demand of j went down.

3.1. 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 critically in Section 6, 7 and to some extent in Section 5.

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. The reason for distinguishing the two types will become clear in Section 6.

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

Recall that good j is optimal for i if⁹ $\alpha_i p_j = u_{ij} - \mu_i$. 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 $\alpha_i > 0$, 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 $\alpha_i p_z < u_{iz} - \mu_i$.

4. Strongly Polynomial Algorithm for Bi-Valued Utilities. In Section 4.1, we will first give a strongly polynomial time algorithm for the *dichotomous case* of HZ, i.e., when all utilities u_{ij} are 0/1. Our algorithm uses some key ideas from the paper of [14], which gave a polynomial time algorithm for the linear-utilities case of Fisher markets; these ideas are summarized in Remark 4.6. We note that in Section 5, [14] presented a “simple algorithm” for this problem; however, it does not run in polynomial time. In Section 8, they enhanced this algorithm with the additional machinery of balanced flows and ℓ_2 norm, and this led to a polynomial time—though not strongly polynomial—algorithm. By exploiting the much simpler structure of the dichotomous case, we managed to fine-tune the “simple algorithm” of [14] to achieve a strongly polynomial algorithm for our problem.

Next, in Section 4.2 we will handle the more general case of *bi-valued utilities*, which is defined as follows: for each agent i , we are given a set of rational numbers $\{a_i, b_i\}$, where $0 \leq a_i < b_i$, and the utilities u_{ij} , $\forall j \in G$, are picked from this set. For this purpose, we will define the notion of *equivalence of utility functions* and use it to reduce the bi-valued utilities case to the dichotomous case.

Notation: We will denote by $H = (A, B, E)$ the bipartite graph on vertex sets A and B , and edge set E , with $(i, j) \in E$ iff $u_{ij} = 1$. For $A' \subseteq A$ and $B' \subseteq B$, we will denote by $H[A', B']$ the restriction of H to vertex set $A' \cup B'$. 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 $S \subseteq A$, $N(S)$ will denote the *set of neighbors*, in B , of vertices in S ; similarly, for any subset $S \subseteq B$, $N(S)$ will denote the set of neighbors, in A , of vertices in S .

4.1. The Dichotomous Case. An instance of the dichotomous case of HZ can be encoded as a bipartite graph $H = (A, B, E)$, where A and B are the set of n agents and n goods, respectively, and for $i \in A, j \in B$, $(i, j) \in E$ if and only if $u_{ij} = 1$. Our algorithm is stated as Algorithm 1.

If H has a perfect matching, computing equilibrium allocations and prices is straightforward, since each agent can be allocated one unit of a unique good from which it derives utility 1 and having price zero; see Steps 1(a) and 1(b) of Algorithm 1. Otherwise in Step 2, we compute a minimum vertex cover for H ; it is necessarily smaller than n . We will need the following lemma.

LEMMA 4.1. *The following hold:*

1. For any set $S \subseteq A_2$, $|N(S)| \geq |S|$.
2. For any set $S \subseteq B_1$, $|N(S) \cap A_1| \geq |S|$.

Proof. 1). If $|N(S)| < |S|$ then $(B_1 \cup N(S)) \cup (A_2 - S)$ is a smaller vertex cover for H , leading to a contradiction.

⁹Note that under this case, optimal goods are not necessarily maximum utility goods; the latter may be suboptimal because their prices are too high.

ALGORITHM 1. Algorithm for the Dichotomous Case

1. If $H = (A, B, E)$ has a perfect matching, say ν , then do:
 - (a) $\forall i \in A$: allocate good $\nu(i)$ to i .
 - (b) $\forall j \in B$: $p_j \leftarrow 0$. Go to Step 3.
2. (a) Find a minimum vertex cover in H , say $(B_1 \cup A_2)$, where $B_1 \subset B$ and $A_2 \subset A$.
 Let $A_1 = A - A_2$ and $B_2 = B - B_1$.
 (b) Find a maximum matching in $H[A_2, B_2]$, say ν .
 (c) $\forall i \in A_2$: allocate good $\nu(i)$ to i ; $\forall j \in B_2$: $p_j \leftarrow 0$.
3. (a) $C \leftarrow A_1$; $D \leftarrow B_1$.
 (b) Consider the subgraph $H[C, D]$.
 (c) Initialization: $p \leftarrow 1$.
 (d) While $D \neq \emptyset$, do:
 - i. Raise p at unit rate.
 - ii. When a set $S \subseteq D$ goes tight, do:
 - A. $\forall j \in S^* : p_j \leftarrow p$.
 - B. $\forall i \in (N(S^*) \cap C) : \text{allocate } 1/p \text{ units of goods from } S^*$.
 - C. $D \leftarrow D - S^*$.
 - D. $C \leftarrow C - N(S^*)$.
- (e) $\forall i \in A_1$: allocate unmatched goods of B_2 to i , to satisfy the size constraint.
4. Output the allocations and prices computed and Halt.

FIG. 2. Algorithm for the Dichotomous Case

546 2). If $|N(S) \cap A_1| < |S|$ then $(B_1 - S) \cup (A_2 \cup N(S))$ is a smaller vertex cover for
 547 H , leading to a contradiction. \square

548 The first part of Lemma 4.1, together with Hall's Theorem, implies that a max-
 549 imum matching in $H[A_2, B_2]$ must match all agents. Therefore in Step 2(a), each
 550 agent $i \in A_2$ is allocated one unit of a unique good, having price zero, from which it
 551 derives utility 1; clearly, this is an optimal bundle of minimum cost for i . The number
 552 of goods that will remain unmatched in B_2 at the end of this step is $|B_2| - |A_2|$.

553 **4.1.1. Computing Allocations and Prices in Subnetwork $H[A_1, B_1]$.** Allo-
 554 cations for agents in A_1 and prices of goods in B_1 are computed in Step 3 of Algorithm
 555 1, in the subnetwork $H[A_1, B_1]$; this is the step which uses ideas from [14], see Remark
 556 4.6. At the end of Step 3(d), each agent in A_1 receives utility 1 goods, worth 1 dollar,
 557 from B_1 . However, if $p_j > 1$, the size of her allocation will be strictly less than one.
 558 To achieve the latter, Step 3(e) allocates the unmatched goods from B_2 , which are
 559 zero-priced, to agents in A_1 . Clearly, the total deficit in size among all agents in A_1
 560 is $|A_1| - |B_1|$. Since this equals $|B_2| - |A_2|$, the market clears at the end of this step.
 561 In Lemma 4.7 we prove that each agent in A_1 gets an optimal bundle of goods of
 562 minimum cost.

563 At any point during the execution of Step 3(d), $C \subseteq A_1$ and $D \subseteq B_1$ and the

algorithm is working on the subnetwork $H[C, D]$. Each good in D has the same price, namely p . We will say that a set $S \subseteq D$ is *tight* if the total worth of goods in S equals the total money possessed by agents in C who desire these goods. The latter set is the neighborhood of S in the subgraph $H[C, D]$, i.e., $N(S) \cap C$. Thus S is tight if $p \cdot |S| = |N(S) \cap C|$.

At the start of Step 3(d), $H[C, D] = H[A_1, B_1]$ and by the second part of Lemma 4.1, every set $S \subseteq D$ satisfies $p \cdot |S| \leq |N(S) \cap A_1|$, where $p = 1$. In Step 3(d)(i), as the algorithm raises p , starting with $p = 1$, at some point a set S will go tight. The price at which a set goes tight and the tight set are given by

$$p^* = \min_{S \subseteq D} \frac{|N(S) \cap A_1|}{|S|} \quad \text{and} \quad S^* = \operatorname{argmin}_{S \subseteq D} \frac{|N(S) \cap A_1|}{|S|},$$

respectively; if several sets go tight, we will assume that S^* represents the maximal tight set¹⁰. In Step 3(d)(ii), the algorithm removes S^* and $N(S^*)$ from D and C , respectively, and attempts to find the next tight set in the rest of the graph. At a general step,

$$p^* = \min_{S \subseteq D} \frac{|N(S) \cap C|}{|S|} \quad \text{and} \quad S^* = \operatorname{argmin}_{S \subseteq D} \frac{|N(S) \cap C|}{|S|}.$$

Next we prove an important monotonicity condition, due to which it suffices to monotonically raise p to find the next tight set. Let S_1, \dots, S_k be the successive sets that go tight. Let $C_1 = A_1$ and in general let $C_{l+1} = A_1 - (S_1 \cup \dots \cup S_l)$.

LEMMA 4.2. For $k > l \geq 1$,

$$\frac{|N(S_{l+1}) \cap C_{l+1}|}{|S_{l+1}|} > \frac{|N(S_l) \cap C_l|}{|S_l|}.$$

Proof. Consider a value of l , with $k > l \geq 1$. Let

$$\frac{|N(S_l) \cap C_l|}{|S_l|} = p.$$

Assume for the sake of contradiction that

$$\frac{|N(S_{l+1}) \cap C_{l+1}|}{|S_{l+1}|} \leq p.$$

Then

$$|N(S_l \cup S_{l+1}) \cap C_l| = |N(S_l) \cap C_l| + |N(S_{l+1}) \cap C_{l+1}| \leq (|S_l| + |S_{l+1}|) \cdot p,$$

contradicting either the minimality of the price at which a set goes tight in the l^{th} iteration or the maximality of the tight set. The lemma follows. \square

COROLLARY 4.3. Suppose the sets S_1, \dots, S_k go tight at prices p_1, \dots, p_k , respectively. Then $1 \leq p_1 < \dots < p_k$.

We next describe how to efficiently find the maximal tight set in the graph $H[C, D]$ using flow-based techniques. Define $R(p)$ to be the following network, as a function of p : its vertices are $C \cup D$ together with special vertices s and t , the source and sink.

¹⁰It is easy to see that there is a unique maximal set.

Direct all edges of $H[C, D]$ from D to C and assign them infinite capacity. Connect s to each vertex in D with an edge of capacity p and connect each vertex in C to t with an edge of capacity 1.

A flow in $R(p)$ should be viewed as a flow of value—goods or money—in dollars. The price of good $j \in D$ is p ; recall that there is one unit of each good in the market. Thus in $R(p)$, at most p dollars of flow can go from s to j . Therefore by flow conservation at all vertices other than s and t , from j , goods worth at most p dollars can go to the set of agents i who are adjacent to j ; the latter set is precisely the set of agents who like j . Furthermore, since the edge from agent i to t has capacity 1, in any flow in $R(p)$, the total value of goods going to i is at most 1 dollar, which is the money of agent i .

The next lemma is easy to verify.

LEMMA 4.4. *The following hold:*

1. For any p , $1 \leq p < p^*$, $(s, D \cup C \cup t)$ is the unique min-cut in $R(p)$.
2. For $p = p^*$, the cuts

$$[s, D \cup C \cup t] \quad \text{and} \quad [s \cup S^* \cup (N(S^*) \cap C), (D - S^*) \cup (C - N(S^*)) \cup t]$$

are both min-cuts in $R(p)$, where p^* and S^* are defined above.

3. For any $p > p^*$, $(s, D \cup C \cup t)$ is not a min-cut in $R(p)$. A min-cut in $R(p)$ is of the form

$$[s \cup S \cup (N(S) \cap C), (D - S) \cup (C - N(S)) \cup t].$$

for an appropriate $S \subseteq D$, depending on p , with $S \neq \emptyset$.

By Lemma 4.4, p^* is the largest value of p for which $(s, D \cup C \cup t)$ is a min-cut in $R(p)$. Clearly, p^* is the ratio of two positive integers $\leq n$ and can be found by conducting a binary search on p in the following interval on the real line:

$$\left[1, \frac{|A_1|}{|B_1|}\right].$$

Each step of this binary search requires computation of an s - t min-cut in $R(p)$ which minimizes the s side. The number of iterations needed is $O(\log n)$. Once p^* is found, S^* can be obtained by finding an s - t min-cut in $R(p)$ which maximizes the s side.

LEMMA 4.5. *For goods $j \in S^*$ assign prices $p_j = p^*$. Then the local market consisting of goods in S^* and agents in $N(S^*) \cap C$ clears.*

Proof. Consider the network $R(p^*)$ corresponding to subnetwork $H[(N(S^*) \cap C), S^*]$. By the second of the claims made in Lemma 4.4, any max-flow in $R(p^*)$ has value $|S^*| = p^* \cdot |N(S^*) \cap C|$. Clearly this flow gives a way of distributing goods in S^* among agents in $(N(S^*) \cap C)$ in such a way that the market clears. \square

Remark 4.6. Our setup is much simpler than that of [14]. As a result, we are able to use simplified versions of ideas from that paper. The latter include the network $R(p)$, the notion of tight sets and the raising of p to successively find tight sets.

LEMMA 4.7. *Each agent in A_1 will get an optimal bundle of goods of minimum cost.*

Proof. First note that for an agent $i \in A_1$ and good $j \in B_2$, $(i, j) \notin E$, since the vertex cover picked has no vertices from $A_1 \cup B_2$. Therefore, for $i \in A_1$, the goods she likes are all in B_1 .

Assume that the algorithm finds k tight sets, S_1, \dots, S_k , in that order; the union of these sets is B_1 . Let p_1, p_2, \dots, p_k be the prices of goods in these sets, respectively. By Lemma 4.2, $1 \leq p_1 < p_2 < \dots < p_k$, and for $1 \leq l \leq k$, $p_l = |(N(S_l) \cap C_l)|/|S_l|$. If $i \in (N(S_l) \cap C_l)$, the algorithm will allocate $1/p_l$ amount of goods to i from S_l , costing 1 dollar, as proven in Lemma 4.5.

By definition of neighborhood of sets, if $i \in (N(S_l) \cap C_l)$, then i cannot have edges to S_1, \dots, S_{l-1} but it can have edges to S_{l+1}, \dots, S_k . However, as shown in Lemma 4.2, the goods in the latter sets will have prices exceeding p_l . Therefore, the cheapest goods from which i accrues utility are in S_l , 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 B_2 , which are zero-priced and from which i gets zero utility. Clearly, i gets an optimal bundle of minimum cost. \square

LEMMA 4.8. *Algorithm 1 finds equilibrium prices and allocations for the dichotomous case of HZ. It runs in strongly polynomial time.*

4.2. Reducing Bi-Valued Case to the Dichotomous Case.

DEFINITION 4.9. *Let I be an instance of the HZ mechanism and let the utility function of agent i be $u_i = \{u_{i1}, u_{i2}, \dots, u_{in}\}$. Then $u'_i = \{u'_{i1}, u'_{i2}, \dots, u'_{in}\}$ is equivalent to u_i if it is a positive affine transform of u_i , i.e., if there are two numbers $s > 0$ and h 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.*

LEMMA 4.10. *Let I be an instance of the HZ mechanism 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, $x_i = \{x_{i1}, \dots, x_{in}\}$ is an optimal bundle for i at prices p and hence is a solution to the optimal bundle LP (3.1). 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. \square

Next, let u_i be bi-valued with the two values being $0 \leq a < b$. 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 reduced to an instance of the dichotomous case, with both having the same equilibria. Now using Lemma 4.8 we get:

THEOREM 4.11. *The bi-valued utilities case of HZ admits a rational equilibrium, and there is a strongly polynomial time algorithm for computing equilibrium allocations and prices for it.*

5. An Example Having Only Irrational Equilibria. Our example has 4 agents A_1, \dots, A_4 and 4 goods g_1, \dots, g_4 ¹¹. 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.

Even such a small instance is not easy to analyze. We will 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.

Consider any equilibrium with minimum price 0. We will analyze its properties, and show eventually that they force specific prices and allocations.

LEMMA 5.1. *Equilibrium prices satisfy:*

$$0 = p_1 < p_2 < 1 \quad \text{and} \quad p_3, p_4 > 1.$$

The equilibrium bundle of each agent is of Type D and contains goods having positive utilities only.

Proof. Suppose $p_3 \leq 1$. Then agents A_3 and A_4 will demand 1 unit each of good g_3 , leading to a contradiction. Similarly, if $p_4 \leq 1$ then A_1 and A_2 will demand 1 unit each of g_4 . Therefore, $p_3, p_4 > 1$. Since the maximum utility goods of every agent have price > 1 , all agents spend exactly 1. Therefore, the sum of the prices of the goods is 4.

Suppose $p_2 = 0 \leq p_1$. Then A_1, A_2, A_4 do not buy g_1 , since they prefer g_2 and it is weakly cheaper than g_1 . Therefore A_3 must buy the entire unit of g_1 . Clearly A_1, A_2 do not buy g_3 , since they prefer g_2 . Therefore, the only agent who buys g_3 is A_4 ; however, she cannot afford the entire unit of g_3 since $p_3 > 1$, contradicting market clearing. Therefore $p_2 > 0$ and hence the 0-priced good is g_1 and $p_1 = 0 < p_2$. Furthermore, $p_2 + p_3 + p_4 = 4$.

Next suppose $p_2 \geq 3/4$. Then $p_4 = 4 - (p_2 + p_3) < 9/4$. For both agents A_1 and A_2 , a combination of g_1 and g_4 in proportion 2:1 has a price less than $3/4$ for one unit

¹¹It 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 is given in the Appendix.

and utility 4, and is therefore preferable to g_2 . Hence, A_1, A_2 will not buy any g_2 , and since A_3 does not buy any g_2 either, since she prefers g_1 , it follows that A_4 must buy the entire unit of g_2 . This is possible only if $p_2 = 1$ and A_4 buys nothing else; in particular, she does not buy any g_3 . Clearly, A_1, A_2 do not buy any g_3 since they prefer g_1 . Therefore the entire unit of g_3 must be bought by A_3 , which is impossible because $p_3 > 1$. Hence $p_2 < 3/4$. These facts together with $p_1 = 0 < p_2 < 1 < p_3, p_4$ imply that the agents' bundles are not Type B or C. Therefore they are all of Type D.

Finally we prove that none of the agents will buy an undesirable good (a good with utility 0). For A_1, A_2, A_3 , such a good is dominated by another lower-priced good. Since $p_4 > 1$, A_4 does not buy g_4 . Suppose agent A_4 buys good g_1 . Since she spends 1 dollar, she must also buy g_3 . Therefore we have: $\alpha_4 p_1 + \mu_4 = u_{41} = 0$. Therefore $\mu_4 = 0$. Also $\alpha_4 p_3 + \mu_4 = u_{43} = 5$; therefore $\alpha_4 p_3 = 5$, which implies $\alpha_4 < 5$ since $p_3 > 1$. Furthermore, $\alpha_4 p_2 + \mu_4 \geq u_{42} = 4$, hence $p_2 > 4/5$, which contradicts $p_2 < 3/4$. Therefore, no agent buys any undesirable good. \square

LEMMA 5.2. *One of the agents A_1, A_2 buys all three desirable goods. If A_1 buys g_1, g_2, g_4 , then A_2 buys g_1, g_4 only. If A_2 buys g_1, g_2, g_4 , then A_1 buys g_2, g_4 only.*

Proof. Since all the bundles are of Type D, every bundle has at least two goods; clearly, every good is bought by at least two agents.

Suppose that every agent buys two goods and every good is bought by two agents. If so, one of A_1, A_2 must buy g_1, g_4 and the other must buy g_2, g_4 . Consider the graph with the goods as nodes and an edge joining two nodes if they are bought by the same agent. This graph must be the 4-cycle g_1, g_4, g_2, g_3, g_1 . Therefore for some a , $0 < a < 1$, each agent buys a units of one good and $b = 1 - a$ units of the second good and each good is sold to two agents in the amounts of a and b .

Let $r_i = |1 - p_i|$. Observe that for every edge (g_i, g_j) of the cycle, one price is < 1 and the other price is > 1 , and we have $ap_i + bp_j = 1$. Therefore $ar_i - br_j = 0$, and $\frac{r_i}{r_j} = \frac{b}{a}$. Hence

$$\frac{r_1}{r_4} = \frac{r_4}{r_2} = \frac{r_2}{r_3} = \frac{r_3}{r_1},$$

which implies that all the r_i are equal. Therefore $p_1 = p_2$, contradicting the previous claim that $p_1 < p_2$. Hence at least one of A_1, A_2 will buy all three of her desirable goods.

Suppose that A_1 buys all three desirable goods g_1, g_2, g_4 . Then we have $\alpha_1 p_j + \mu_1 = u_{1j}$ for $j = 1, 2, 4$. Therefore, $(p_4 - p_1)/(p_4 - p_2) = (u_{14} - u_{11})/(u_{14} - u_{12}) = 3/2$. Agent A_2 buys g_4 and at least one of g_1, g_2 . Suppose she buys g_2 . Then $\alpha_2 p_j + \mu_2 = u_{2j}$ for $j = 2, 4$, hence $\alpha_2(p_4 - p_2) = u_{24} - u_{22} = 5$. This implies that $\alpha_2(p_4 - p_1) > 6 = u_{24} - u_{21}$, hence $\alpha_2 p_1 + \mu_2 < u_{21}$, a contradiction. Therefore A_2 does not buy g_2 and she buys g_1 and g_4 only.

Next suppose A_2 buys all three desirable goods g_1, g_2, g_4 . By a similar argument we will prove that A_1 buys only two goods. We have $\alpha_2 p_j + \mu_2 = u_{2j}$ for $j = 1, 2, 4$. Therefore, $(p_4 - p_1)/(p_4 - p_2) = (u_{24} - u_{21})/(u_{24} - u_{22}) = 6/5$. Agent A_1 buys g_4 and at least one of g_1, g_2 . Suppose that she buys g_1 . Then $\alpha_1 p_j + \mu_1 = u_{1j}$ for $j = 1, 4$, hence $\alpha_1(p_4 - p_1) = u_{14} - u_{11} = 6$. This implies that $\alpha_1(p_4 - p_2) > 4 = u_{14} - u_{12}$, hence $\alpha_1 p_2 + \mu_1 < u_{12}$, a contradiction. Therefore, A_1 does not buy g_1 , hence she buys g_2 and g_4 only. \square

THEOREM 5.3. *The instance of Table 1 has a unique equilibrium; the allocations to agents and prices of goods, other than the zero-priced good, are all irrational. The prices are as follows:*

$$p_1 = 0, \quad p_2 = (23 - \sqrt{17})/32, \quad p_3 = (9 + \sqrt{17})/8, \quad p_4 = (69 - 3\sqrt{17})/32.$$

Proof. Let $r_i = |1 - p_i|$. By Lemma 5.1, $r_1 = 1$. We consider the two cases established in Lemma 5.2. We will show that in Case 1 there is a unique equilibrium, while in Case 2 there is no equilibrium.

Case 1. A_1 buys g_1, g_2, g_4 , and A_2 buys g_1, g_4 .

Agent A_3 spends her dollar on goods g_1, g_3 in the proportion $r_3 : r_1$, i.e., $r_3 : 1$. Therefore, $x_{31} = \frac{r_3}{1+r_3}$, $x_{33} = \frac{1}{1+r_3}$. Agent A_4 buys goods g_2, g_3 in the proportion $r_3 : r_2$. Therefore, $x_{42} = \frac{r_3}{r_2+r_3}$, $x_{43} = \frac{r_2}{r_2+r_3}$. Since only agents A_3 and A_4 buy good g_3 , we have $x_{31} = 1 - x_{33} = x_{43}$, and $x_{42} = 1 - x_{43} = x_{33}$. This implies $r_3^2 = r_2 \dots$ (1).

Since agent A_1 buys g_1, g_2, g_4 , we have, $\frac{u_{14}-u_{12}}{u_{12}-u_{11}} = \frac{p_4-p_2}{p_2-p_1}$. Therefore $r_2 + r_4 = 2(1 - r_2) \dots$ (2).

The sum of the prices is equal to 4, therefore $1 + r_2 - r_3 - r_4 = 0 \dots$ (3)

Now we have three equations, (1), (2) and (3), in three unknowns r_2, r_3, r_4 . Using (1) and (2) we can express r_2 and r_4 in terms of r_3 . Letting $r_3 = y$, we have from (1), $r_2 = y^2$, and from (2), $r_4 = 2 - 3r_2 = 2 - 3y^2$. Substituting into (3), we get $4y^2 - y - 1 = 0$.

The only positive solution is $y = \frac{1+\sqrt{17}}{8}$. Therefore,

$$p_1 = 0, \quad p_2 = 1 - r_2 = 1 - y^2 = \frac{23 - \sqrt{17}}{32}, \quad p_3 = 1 + r_3 = 1 + y = \frac{9 + \sqrt{17}}{8},$$

$$p_4 = 1 + r_4 = 3 - 3y^2 = \frac{69 - 3\sqrt{17}}{32}.$$

Once we have the value of y , we get:

$$r_1 = 1, \quad r_2 = y^2 = \frac{9 + \sqrt{17}}{32}, \quad r_3 = y = \frac{1 + \sqrt{17}}{8} \quad \text{and} \quad r_4 = 2 - 3y^2 = \frac{37 - 3\sqrt{17}}{32}.$$

We can compute then the allocations from the r_i . We already expressed the allocations for agents A_3, A_4 in terms of the r_i . Agent A_2 buys goods g_1, g_4 in the proportion $r_4 : r_1$, i.e., $r_4 : 1$. Therefore, $x_{21} = \frac{r_4}{1+r_4}$, $x_{24} = \frac{1}{1+r_4}$. Agent A_1 buys the remaining amount of each good g_1, g_2, g_4 . Thus, the allocations of the agents in terms of the r_i are:

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

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

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

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

We conclude that, if there is an equilibrium in Case 1, then there can be only one and it must have the above prices and allocations.

Conversely, we can verify that the above pair (p, x) is an equilibrium. First we note that all allocations are nonnegative. This is obvious for all the allocations, except for x_{11} , which, after plugging in the values for the r_i 's evaluates to approximately 0.084. Second, note that every good has exactly one unit allocated: for good g_3 this follows from equation (1), and for the other goods it holds because A_1 buys the remaining amounts. Third, every agent buys a total of one unit of goods: this is obvious for agents A_2, A_3, A_4 from the allocations, and for agent A_1 it follows because exactly one unit is sold of each good. Fourth, every agent spends exactly one dollar: this holds for agents A_2, A_3, A_4 because they pay an average price of 1 for their goods, and for agent A_1 it follows from the fact that the total expenditure of the agents, which is equal to the sum of the prices of the goods, is 4 (equation (3)).

Finally, it can be shown that the bundle of every agent is optimal for these prices, using the dual LP and complementary slackness. The dual variables α_i can be calculated as $\alpha_i = \frac{u_{ij}-u_{ik}}{p_j-p_k}$, where g_k, g_j are (any) two goods bought by agent A_i ; the shift $\mu_i = u_{ij} - \alpha_i p_j$ (which is equal to $u_{ik} - \alpha_i p_k$). Thus, for example $\alpha_1 = \frac{u_{12}-u_{11}}{p_2-p_1} = \frac{2}{p_2}$, and $\mu_1 = u_{11} - \alpha_1 p_1 = 2$. Note that $\frac{u_{12}-u_{11}}{p_2-p_1} = \frac{u_{14}-u_{12}}{p_4-p_2} = \frac{u_{14}-u_{11}}{p_4-p_1}$ by equation (2), so it does not matter which goods g_j, g_k in Agent A_1 's bundle are used to calculate α_1 . Also, for each agent A_i , it does not matter which good g_j in her bundle is used to calculate μ_i .

Clearly $\alpha_i \geq 0$ for all i . For all agents A_i and goods g_j in the bundle of A_i , we have $\alpha_i p_j + \mu_i = u_{ij}$, by construction. Furthermore, if good g_j is not in the bundle of A_i then $\alpha_i p_j + \mu_i > u_{ij}$: For agent A_1 and good g_3 , note that g_3 has higher price and lower utility than good g_1 which is in the bundle of A_1 , hence $\alpha_1 p_3 + \mu_1 > \alpha_1 p_1 + \mu_1 = u_{11} > u_{13}$. The same argument applies to agent A_2 and good g_3 , agent A_3 and goods g_2, g_4 (they are both dominated by good g_1 in A_3 's bundle), and to agent A_4 and good g_4 (it is dominated by good g_2). For agent A_2 and good g_2 , note that $\alpha_2 = \frac{u_{24}-u_{21}}{p_4-p_1} = \frac{u_{14}-u_{11}}{p_4-p_1} = \alpha_1$, and $\mu_2 = u_{24} - \alpha_2 p_4 = u_{14} - \alpha_1 p_4 = \mu_1$. Therefore $\alpha_2 p_2 + \mu_2 = \alpha_1 p_2 + \mu_1 = u_{12} = 4 > 3 = u_{22}$. The only remaining case that needs to be checked numerically is agent A_4 and good g_1 . Since $p_1 = 0$ and $u_{41} = 0$, the inequality $\alpha_4 p_1 + \mu_4 > u_{41}$ is equivalent to $\mu_4 > 0$. By construction, $\mu_4 = u_{44} - \alpha_4 p_3 = u_{44} - \frac{u_{43}-u_{42}}{p_3-p_2} p_3 = 5 - \frac{p_3}{p_3-p_2} = \frac{4p_3-5p_2}{p_3-p_2}$. Thus, $\mu_4 > 0$ is equivalent to $4p_3 > 5p_2$, which holds for the above values of p_2, p_3 . Therefore, the values α_i, μ_i satisfy the constraints of the dual LP, and since they and the x_{ij} satisfy clearly also the complementary slackness conditions, it follows that the allocations x_{ij} give optimal bundles to the agents for the prices p_j . Therefore, (x, p) is an equilibrium.

Case 2. A_2 buys g_1, g_2, g_4 , and A_1 buys g_2, g_4 .

We will show that there is no equilibrium in this case. Specifically, we will show that if there is an equilibrium, it must have specific prices and allocations, and we will derive a contradiction.

Consider any equilibrium for Case 2. The allocations for agents A_3, A_4 are the same as in Case 1, i.e., $x_{31} = \frac{r_3}{1+r_3}$, $x_{33} = \frac{1}{1+r_3}$, and $x_{42} = \frac{r_3}{r_2+r_3}$, $x_{43} = \frac{r_2}{r_2+r_3}$. Again we have $x_{31} = x_{43}$ and $x_{42} = x_{33}$, which implies $r_3^2 = r_2 \dots$ (1).

Since agent A_2 buys g_1, g_2, g_4 , we have, $\frac{u_{24}-u_{22}}{u_{22}-u_{21}} = \frac{p_4-p_2}{p_2-p_1}$, therefore $r_2 + r_4 = 5(1 - r_2) \dots$ (2').

The sum of the prices is 4, thus again $1 + r_2 - r_3 - r_4 = 0 \dots$ (3)

We can solve now for r_2, r_3, r_4 . Using (1) and (2') we can express r_2 and r_4 in terms of r_3 . Letting $r_3 = y$, we have from (1), $r_2 = y^2$, and from (2'), $r_4 = 5 - 6r_2 = 5 - 6y^2$. Substituting into (3), we get $7y^2 - y - 4 = 0$.

The only positive solution is $y = \frac{1+\sqrt{113}}{14}$. Therefore,

$$p_1 = 0, \quad p_2 = 1 - r_2 = 1 - y^2 = \frac{41 - \sqrt{113}}{98}, \quad p_3 = 1 + r_3 = 1 + y = \frac{15 + \sqrt{113}}{14},$$

$$p_4 = 1 + r_4 = 6 - 6y^2 = \frac{246 - 6\sqrt{113}}{98}.$$

As in the previous case, the value of y gives:

$$r_1 = 1, \quad r_2 = y^2 = \frac{57 + \sqrt{113}}{98}, \quad r_3 = y = \frac{1 + \sqrt{113}}{14}$$

$$\text{and } r_4 = 5 - 6y^2 = \frac{148 - 6\sqrt{113}}{98}.$$

□

If there is any equilibrium in Case 2, then it must have the above prices. We can compute again the allocations from the r_i . The allocations of A_3 and A_4 are as before. Agent A_1 buys goods g_2, g_4 in the proportion $r_4 : r_2$. Therefore, $x_{12} = \frac{r_4}{r_2+r_4}$, $x_{14} = \frac{r_2}{r_2+r_4}$. Substituting the values of the r_i 's in the expressions for x_{12} and x_{42} , we get $x_{12} = \frac{r_4}{r_2+r_4} \approx 0.554$ and $x_{42} = \frac{r_3}{r_2+r_3} \approx 0.546$. Thus, $x_{12} + x_{42} \approx 1.1 > 1$, i.e., good g_2 is oversold. Therefore, there is no equilibrium in Case 2.

Remark 5.4. 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.

6. 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 [17]. 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 price and allocation 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 prices and allocations, hence proving that Hylland-Zeckhauser is in FIXP.

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: $p_j \leftarrow p_j - r$

FIG. 2. Straight line program for function F_p .

Notation: We will denote the set $\{1, \dots, n\}$ by $[n]$. x_i will denote 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 Algorithm 2 and Algorithm 3, respectively. 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 F : 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, $\text{cost}(i)$ 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, every fixed point of F must satisfy all these conditions and is therefore an equilibrium. Conversely, every equilibrium (p, x) is a fixed point of F .

Intuitively, the function F_p adjusts the prices of goods if they are under- or over-allocated under x , i.e., if condition 1 is violated (step 1 of Algorithm 2) and ensures that the minimum price of a good is 0 (steps 2, 3 of Algorithm 2). The function F_i for an agent i adjusts the agent's allocation x_i if it costs too much under prices p , i.e., if condition 2 is violated (steps 1, 2 of Algorithm 3), or if it is not optimal, i.e. condition 3 is violated (steps 3-7 of Algorithm 3). In the case of a suboptimal allocation x_i , the allocation can be improved by a local transfer of a small enough amount of allocation mass among two or three goods: either from one good to another (from good $k \notin G_i^*$ to good i^* in step 4; from good j to good k in step 5), or from one good to two others (from good k to goods j, l in step 6), or from two goods to a third one (from goods j, l to good k in step 7).

We will prove that 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 , and is left to the reader. The proof for F_i is more delicate and uses a potential function argument based on the

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. $u_{ij} \leq u_{ik}$ 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) $x_{il} := x_{il} - \frac{u_{ik} - u_{ij}}{u_{il} - u_{ij}} d$

FIG. 2. Straight line program for function F_i

913 changes in $\text{value}(i) = \sum_j u_{ij} x_{ij}$ and $\text{cost}(i) = \sum_j p_j x_{ij}$ caused by any change in the
 914 allocation x_i in every step of the algorithm for F_i , as stated in the following lemma.

915 LEMMA 6.1. *Let (p, x) be such that $p \in D_p$, $x \in D_x$. Then, $\text{cost}(i)$ and $\text{value}(i)$
 916 are modified by the steps of F_i as follows.*

- 917 1. *If Steps 1, 2 modify x_i , then the initial cost is > 1 , and steps 1,2 decrease*
 918 *strictly $\text{cost}(i)$.*
- 919 2. *If steps 3, 4 modify x_i , then they increase strictly $\text{value}(i)$ while maintaining*
 920 *$\text{cost}(i) \leq 1$.*
- 921 3. *If anyone of steps 5, 6, 7 modifies x_i , then it increases weakly $\text{value}(i)$ and*
 922 *decreases strictly $\text{cost}(i)$.*

923 *Proof.* For the first part, note that if steps 1,2 modify the allocation x_i , then we
 924 must have $r \sum_k (1 - p_k)_+ > 0$, hence $r > 0$ and $\sum_k (1 - p_k)_+ > 0$. Therefore, the initial
 925 cost $\text{cost}(i) = \sum_j p_j x_{ij} = r + 1 > 1$. The new cost is $\frac{\sum_j p_j x_{ij} + r \sum_j p_j (1 - p_j)_+}{1 + r \sum_j (1 - p_j)_+}$ which
 926 is $< \text{cost}(i)$, because $r \sum_j p_j (1 - p_j)_+ \leq r \sum_j (1 - p_j)_+ < (r \sum_j (1 - p_j)_+) \text{cost}(i)$; the
 927 last inequality is strict because $r \sum_j (1 - p_j)_+ > 0$, and $\text{cost}(i) > 1$.

For the second part, note that if the cost of the allocation, $\text{cost}(i) = \sum_j p_j x_{ij}$ before step 3 is ≥ 1 , then $t = 0$ in line 3, and steps 3, 4 make no change. Suppose that the cost is < 1 , i.e. $t > 0$. For every good $k \notin G_i^*$, if $x_{ik} = 0$ then no change is made for this good. Thus, if steps 3,4 change x_i , then there must be some good(s) $k \notin G_i^*$ with $x_{ik} > 0$. For every such good k , we swap d units of k for i^* , and as a result the value is increased by $d(u_{ii^*} - u_{ik}) > 0$, since $d > 0$ and $u_{ii^*} > u_{ik}$. The cost is increased at most by $d(p_{i^*} - p_k) \leq \frac{t}{n^2}n = \frac{t}{n}$. Hence, over all the goods $k \notin G_i^*$, the cost is increased by less than t , hence it remains < 1 .

For the third part, we consider the following three cases.

- If Step 5 modifies x_i for a pair j, k of goods then we must have $p_j > p_k$ and $x_{ij} > 0$. Since $u_{ij} \leq u_{ik}$, step 5 weakly increases $\text{value}(i)$ and strictly decreases $\text{cost}(i)$.
- If Step 6 kicks in for a triple of goods j, k, l , then the net change in $\text{value}(i)$ is

$$d \frac{u_{il} - u_{ik}}{u_{il} - u_{ij}} u_{ij} + d \frac{u_{ik} - u_{ij}}{u_{il} - u_{ij}} u_{il} - du_{ik} = 0.$$

The net change in $\text{cost}(i)$ is

$$d \frac{u_{il} - u_{ik}}{u_{il} - u_{ij}} p_j + d \frac{u_{ik} - u_{ij}}{u_{il} - u_{ij}} p_l - dp_k = \frac{d\Delta}{u_{il} - u_{ij}},$$

where $\Delta = (u_{ik} - u_{ij})(p_l - p_k) - (u_{il} - u_{ik})(p_l - p_k) < 0$.

- If Step 7 kicks in for a triple of goods j, k, l , then the net change in $\text{value}(i)$ is again 0, and the net change in cost is $\frac{-d\Delta}{u_{il} - u_{ij}} < 0$. \square

COROLLARY 6.2. *If (p, x) is a fixed point of F , then no step of F_i will change x_i .*

Proof. Suppose that some step(s) of F_i change the allocation x_i of fixed point (x, p) , and consider the earliest such step. If it is step 2, then the initial $\text{cost}(i) > 1$, and step 2 decreases strictly the cost. Step 3,4 either do not change the allocation or if they do change it, the new cost is ≤ 1 , i.e. still smaller than the initial one. Steps 5, 6, 7 do not increase the cost, hence the final cost is strictly smaller than the initial. Thus, the final allocation x_i cannot be the same as the initial.

If the earliest step that changes x_i is step 4, then it increases strictly the value and the subsequent steps do not decrease it, hence the final value is strictly higher than the initial. If the earliest modifying step is one of 5, 6, 7, then it decreases strictly the cost, and all other subsequent changes do not increase it. We conclude that no step of F_i can change the allocation x_i of a fixed point. \square

LEMMA 6.3. *If (p, x) is a fixed point of F , as defined in Algorithms 2 and 3, then*

1. $\exists z \in G$ such that $p_z = 0$.
2. $\forall i \in [n]$, $\text{cost}(i) \leq 1$.
3. $\forall j \in [n]$, $\sum_{i \in A} x_{ij} = 1$, i.e. the market clears.

Proof. 1. Steps 2 and 3 of F_p ensure that there is a good with price 0.

2. If for some $i \in [n]$, $\text{cost}(i) > 1$, then Steps 1 and 2 of F_i will modify x_i since $r = \text{cost}(i) - 1 > 0$, and $\sum_k (1 - p_k)_+ > 0$ because some good z has $p_z = 0$.

3. Suppose that there is a good j such that $\sum_i x_{ij} \neq 1$. Since $\sum_j x_{ij} = 1$ for all agents $i \in [n]$, there must be a good k such that $\sum_i x_{ik} < 1$, and another good l such that $\sum_i x_{il} > 1$.

We claim that then $p_k = 0$. Since $\sum_i x_{ik} < 1$, if $p_k > 0$, then line 1 of F_p will strictly decrease p_k , and line 3 certainly does not increase it, contradicting

972 $F_p(p, x) = p$. Thus, $p_k = 0$, the price p_k will stay 0 after line 1, hence $r = 0$
 973 in line 2, and line 3 will not change any prices.

974 On the other hand, we claim that $p_l = n$. Since $\sum_i x_{il} > 1$, if $p_l < n$, then line
 975 1 of F_p will increase strictly p_l , and since line 3 has no effect, this contradicts
 976 $F_p(p, x) = p$.

977 But $\text{cost}(i) = \sum_j p_j x_{ij} \leq 1$ for all $i \in [n]$ implies that $\sum_i \sum_j p_j x_{ij} \leq n$,
 978 which contradicts the fact that $p_l = n$ and $\sum_i x_{il} > 1$, hence $\sum_i p_l x_{il} > n$.

979 **LEMMA 6.4.** *If (p, x) is a fixed point of F , as defined in Algorithms 2 and 3, then*
 980 *x_i is an optimal bundle for i at prices p . Furthermore, $\text{cost}(i)$ is minimum among*
 981 *optimal bundles.*

982 *Proof.* We will consider the following exhaustive list of cases. Each contradiction
 983 is based on applying Corollary 6.2. We will assume that α_i and μ_i are optimal variables
 984 of the dual to i 's optimal bundle LP and that $u = \max_j \{u_{ij}\}$.

985 **Case 1:** Assume that $\text{cost}(i) < 1$. If $B_i \not\subseteq G_i^*$, then Steps 3 and 4 will kick
 986 in, contradicting the fact that (p, x) is a fixed point. Therefore $B_i \subseteq G_i^*$. Clearly,
 987 u is the maximum utility that i can derive from a bundle satisfying $\text{size}(i) = 1$ and
 988 $\text{cost}(i) \leq 1$. Therefore, x_i is an optimal bundle for i . Since step 5 does not modify x_i ,
 989 all goods in B_i must have minimum price among the goods of G_i^* . Therefore, $\text{cost}(i)$
 990 is minimum among the optimal bundles.

991 *Henceforth, we will assume that $\text{cost}(i) = 1$.*

992 **Case 2:** Assume that i derives the same utility from all goods $j \in B_i$ and
 993 $B_i \subseteq G_i^*$. As in the previous case, x_i is an optimal bundle for i and hence each good
 994 in B_i is optimal. Furthermore, again since step 5 does not modify the allocation, as
 995 in Case 1, $\text{cost}(i)$ is minimum among the optimal bundles.

996 **Case 3:** Assume that i derives the same utility from all goods $j \in B_i$ and
 997 $B_i \not\subseteq G_i^*$. Let k be a good in B_i and let z be a good having price 0. Each good in B_i
 998 must be a minimum price good having utility u_{ik} , since otherwise Step 5 of F_i will
 999 alter the bundle. Since $\text{cost}(i) = 1$, $\text{size}(i) = 1$ and all goods in B_i have the same
 1000 price, each good in B_i has price 1.

1001 Let l be a good such that $u_{il} > u_{ik}$; observe that any good in G_i^* is such a good.
 1002 We will prove that $p_l > 1 = p_k$. Clearly $u_{iz} < u_{ik}$, since otherwise Step 5 will kick
 1003 in and change the bundle. Hence we have $u_{iz} < u_{ik} < u_{il}$. However, since Step 6 did
 1004 not kick in, $(u_{il} - u_{ik})(p_k - p_z) \leq (u_{ik} - u_{iz})(p_l - p_k)$. Since $(u_{il} - u_{ik})(p_k - p_z) > 0$,
 1005 we get that $(p_l - p_k) > 0$. Therefore $p_l > p_k = 1$. Hence we can conclude that the
 1006 optimal bundle for i at prices p is not a Type A or Type B bundle.

1007 Next, assume for the sake of contradiction that x_i is not an optimal bundle for
 1008 i at prices p ; in particular, this entails that the optimal bundle for i is not Type C.
 1009 Therefore, i 's optimal bundle must be Type D and k is a suboptimal good. As argued
 1010 in Section 3, an optimal Type D bundle must contain a good of price < 1 and a good
 1011 of price > 1 ; let j and l be such goods, respectively. Clearly $u_{iz} < u_{ik} < u_{il}$. Then
 1012 we have,

$$1013 \quad \alpha_i p_j = u_{ij} - \mu_i, \quad \alpha_i p_k > u_{ik} - \mu_i \quad \text{and} \quad \alpha_i p_l = u_{il} - \mu_i$$

1014 Subtracting the first from the second and the second from the third we get

$$1015 \quad \alpha_i (p_k - p_j) > (u_{ik} - u_{ij}) \quad \text{and} \quad \alpha_i (p_l - p_k) < (u_{il} - u_{ik})$$

1016 This gives

$$1017 \quad (u_{il} - u_{ik})(p_k - p_j) - (u_{ik} - u_{ij})(p_l - p_k) > 0.$$

Therefore, Step 6 should kick in, leading to a contradiction. Hence x_i is a Type C optimal bundle. Since all goods in G_i^* have price > 1 , every bundle with cost < 1 is suboptimal, thus x_i has minimum cost among optimal bundles.

Henceforth, we will assume that $\text{cost}(i) = 1$ and $\exists s, t \in B_i$ with $u_{is} < u_{it}$.

Case 4: Assume that the set $\{u_{ij} \mid j \in G\}$ has exactly two elements. Clearly, these utilities must be u_{is} and u_{it} . Now, s must be the zero-priced good, since otherwise Step 5 will kick in. Since $\text{cost}(i) = 1$ and $\text{size}(i) = 1$, $p_t > 1$. Again since Step 5 didn't kick in, s and t must be the cheapest goods having utilities u_{is} and u_{it} . Therefore, x_i is a Type D optimal bundle. It has minimum cost ($=1$) among optimal bundles for the same reason as in case 3.

Case 5: Assume that the set $\{u_{ij} \mid j \in G\}$ has three or more elements. Since $\text{size}(i) = 1$ and $\text{cost}(i) = 1$, $\exists t \in B_i$, s.t. $p_t > 1$. Now, any good having utility u must have price > 1 , since otherwise Step 5 will alter the bundle. Therefore, x_i cannot be a Type A or Type B bundle. Therefore, $\alpha_i > 0$.

Suppose that x_i is not an optimal bundle. Then there are two cases: that the optimal bundle is Type C or Type D. In the first case, let $k \in G$ be an optimal good; $p_k = 1$. Let $j, l \in B_i$ with $p_j < 1 < p_l$ and at least one of j or l is suboptimal. Clearly, $u_{ij} < u_{ik} < u_{il}$, otherwise Step 5 will kick in. Therefore we have

$$\alpha_i p_j \geq u_{ij} - \mu_i, \quad \alpha_i p_k = u_{ik} - \mu_i \quad \text{and} \quad \alpha_i p_l \geq u_{il} - \mu_i,$$

with at least one of the inequalities being strict. Therefore,

$$(u_{ik} - u_{ij})(p_l - p_k) > (u_{il} - u_{ik})(p_k - p_j),$$

and Step 7 should kick in, leading to a contradiction. Hence x_i is a Type C optimal bundle.

Next suppose the optimal bundle is Type D. There are two cases. First, suppose $\exists k \in B_i$ such that k is a suboptimal good for i and there are optimal goods j and l with $u_{ij} < u_{ik} < u_{il}$. Then we have

$$\alpha_i p_j = u_{ij} - \mu_i, \quad \alpha_i p_k > u_{ik} - \mu_i \quad \text{and} \quad \alpha_i p_l = u_{il} - \mu_i$$

As before we get

$$(u_{il} - u_{ik})(p_k - p_j) - (u_{ik} - u_{ij})(p_l - p_k) > 0.$$

Therefore, Step 6 should kick in, leading to a contradiction.

Second, suppose that there is no such good $j \in B_i$. Let v and w be optimal goods with the smallest and largest utilities for i . Then all suboptimal goods in B_i have either less utility than u_{iv} or more utility than u_{iw} . Suppose there are both types of goods, say j and l , respectively. Then Step 7 should kick in with the triple j, v, l . Else there is only one type, say j with $u_j < u_v$. Then $\exists l \in B_i$ with $p_l > 1$. Now, Step 7 should kick in with the triple j, v, l . In the remaining case, $\exists j, l \in B_i$ with $p_j < 1$ and $u_{il} > u_{iw}$. Now, Step 7 should kick in with the triple j, w, l .

The contradictions give us that x_i does not contain a suboptimal good and is hence a Type D optimal bundle. The minimality of the cost holds for the same reason as in Cases 3, 4. \square

Lemmas 6.3 and 6.4 give:

THEOREM 6.5. *The problem of computing an exact equilibrium for the Hylland-Zeckhauser mechanism is in FIXP.*

COROLLARY 6.6. *An HZ equilibrium can always be expressed via algebraic numbers.*

7. 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.

7.1. Definition and Properties. 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. We say that a price vector $p \geq 0$ is *normalized* if $\min_i p_i = 0$.

DEFINITION 7.1. A pair (p, x) of (non-negative) normalized prices p (i.e. $\min_i p_i = 0$) and allocations x 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. (a) 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 $\text{cost}(i) \leq 1$. (b) 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 $v^*(i)$.¹²

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 $\log(1/\epsilon)$.

The requirement in the definition that prices be normalized is important, because if the prices are not restricted then conditions 3 and 4(b) in the definition have no effect: If (p, x) is any pair that satisfies conditions 1, 2 and 4(a) then we can scale the differences of all prices from 1 (as in Lemma 3.2) yielding a vector p' all of whose components are within $\epsilon/2$ of 1, and the resulting pair (p', x) will satisfy also conditions 3 and 4(b): If all prices p'_i are within $\epsilon/2$ of 1, then any unit bundle x_i has cost $\text{cost}'(x_i) = \sum_j x_{ij} p'_j$ also within $\epsilon/2$ of 1, and any two unit bundles have costs within ϵ of each other; therefore conditions 3 and 4(b) will be trivially satisfied. Note also that scaling the differences $p_j - 1$ does not affect conditions 1, 2, and 4(a), since it does not change the optimal bundles and the optimal value $v^*(i)$ of an agent.

We define also a more relaxed version, called a *relaxed ϵ -approximate equilibrium* where the normalization condition is relaxed to $\min_j p_j \leq \epsilon$ and 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 7.2. The problems of computing an ϵ -approximate equilibrium and a relaxed approximate equilibrium are polynomially equivalent.

Proof. Clearly, the relaxed version is no harder than the nonrelaxed version. On the other hand, if we have an algorithm for the relaxed version, then we can compute an ϵ -approximate equilibrium as follows. Given a one-sided market M and a rational $\epsilon > 0$, assume without loss of generality that $\epsilon \leq 1$ and $n \geq 2$. Compute a relaxed δ -approximate equilibrium (p, x) where $\delta = \epsilon/4n$. In this equilibrium the minimum

¹²Condition 4b is included here to be consistent with the definition of HZ equilibrium. In the case of HZ equilibrium, the corresponding condition is needed to ensure Pareto optimality. For approximate equilibria with $\epsilon > 0$, this condition is not necessary (see the proof of Proposition 7.4). Also, the proof of PPAD-hardness in [8] does not need this condition; that is, if we define approximate HZ equilibria without condition 4b, their computation is still PPAD-complete.

price $p_{\min} = \min_j p_j$ may be positive, but is at most δ , and some goods may be oversold or undersold by an amount at most δ .

We compute a new pair (p', x') as follows. For the prices, we normalize them by scaling their difference from 1: Set $p'_j = 1 + (p_j - 1)/(1 - p_{\min})$ for all $j \in G$; then $\min_j p'_j = 0$. For the allocations, we set up a bipartite transportation network N with the nodes corresponding to the agents and the goods that are undersold or oversold. If a good j is oversold then it has an outgoing directed edge (j, i) to every agent i with capacity x_{ij} ; if a good j is undersold then it has an incoming directed edge (i, j) from every agent i with unlimited capacity (we could also set the capacity to 1 or to δ). For each good j that is oversold, the corresponding node is a source with supply $\sum_i x_{ij} - 1$; for each good j that is undersold, the corresponding node is a sink with demand $1 - \sum_i x_{ij}$. Since $\sum_j x_{ij} = 1$ for all agents i , the sum of the supplies over all sources is equal to the sum of the demands over all sinks. We construct a flow f in the network N that ships the excess allocation from the oversold to the undersold goods, and combine f with x to obtain the new allocation x' .

CLAIM 7.3. *We can construct in polynomial time a feasible flow f in N that ships all the supply from the sources to the sinks, where the flow on each edge is at most δ .*

Proof. Perform repeatedly the following action until there is no more any supply and demand left: Pick any source node (good) j with positive supply $\text{supply}(j)$, any outgoing edge (j, i) with positive capacity $\text{cap}(j, i)$, and any sink node (good) j' with positive demand $\text{demand}(j')$, and ship $\min\{\text{supply}(j), \text{cap}(j, i), \text{demand}(j')\}$ amount of flow from j to j' along the edges $(j, i), (i, j')$. Decrease $\text{supply}(j), \text{cap}(j, i)$ and $\text{demand}(j')$ by this amount.

Clearly an invariant of this algorithm is that at all times the total supply is equal to the total demand. Furthermore, for every source node j , the sum of the capacities of the outgoing edges (j, i) is equal to $1 + \text{supply}(j)$. These properties hold initially, and are obviously maintained in every iteration. The properties imply in particular that, if $\text{supply}(j) > 0$ then there is an outgoing edge (j, i) with positive capacity and there is a sink j' with positive $\text{demand}(j')$. Every iteration eliminates either a source, a sink, or an edge (i.e. reduces to 0 respectively its supply, demand or capacity), hence the algorithm terminates in polynomial time.

Since the maximum supply and demand of any node in the initial network is at most δ , the flow on every edge (j, i) and (i, j') is at most δ . \square

Combining this flow f with the allocation x , we obtain a new allocation x' : If a good j was oversold then set $x'_{ij} = x_{ij} - f(j, i)$ for all $i \in A$; if j was undersold then set $x'_{ij} = x_{ij} + f(i, j)$ for all $i \in A$. For the other goods j , $x'_{ij} = x_{ij}$ for all $i \in A$. The new allocation x' satisfies $\sum_j x'_{ij} = 1$ for all agents i (because of the flow conservation at the agent nodes), and $\sum_i x'_{ij} = 1$ for all goods j (because f ships all excess supply from the sources to the sinks). Thus, conditions 1 and 2 in the definition are satisfied.

For condition 4(a), note first that the change in prices by scaling their difference from 1, does not affect the optimal bundles and the optimal value $v^*(i)$ of each agent. The value $\text{value}(i)$ of the allocation x_i is at least $v^*(i) - \delta$. The flow on each edge is at most δ , so every allocation x_{ij} is changed at most by δ . Since all the utilities u_{ij} are in $[0, 1]$, the value of each agent's allocation is changed at most by $n\delta < \epsilon/2$. Therefore, the new value $\text{value}'(i)$ of the allocation x'_i of agent i is at least $v^*(i) - \delta - \epsilon/2 > v^*(i) - \epsilon$.

For conditions 3 and 4b, note first that since $\text{cost}(i) = \sum_j x_{ij} p_j \leq 1 + \delta$ for all agents i , $\sum_i \sum_j x_{ij} p_j = \sum_j p_j \sum_i x_{ij} \leq n(1 + \delta)$. Since $\sum_i x_{ij} \geq 1 - \delta$ for all goods j , it follows that the sum of the prices $\sum_j p_j \leq n(1 + \delta)/(1 - \delta) \leq 2n$. The cost $\text{cost}'(i) = \sum_j x'_{ij} p'_j$ of the new allocation x'_i of agent i with respect to the new prices p' is at most

$\sum_j (x_{ij} + \delta)p_j / (1 - \delta)$ (since $p_{\min} \leq \delta$). Thus $\text{cost}'(i) \leq (\text{cost}(i) + \delta \sum_j p_j) / (1 - \delta) \leq \text{cost}(i) + \delta(1 + 2n) / (1 - \delta) \leq \text{cost}(i) + 3\delta n$. Hence, $\text{cost}'(i) \leq 1 + \delta + 3\delta n \leq 1 + \epsilon$, proving condition 3.

For condition 4b, note that the cost $c^*(i)$ of the optimal bundle for agent i , if it is smaller than 1, will decrease somewhat with respect to the new prices, to $c'^*(i) = 1 + (c^*(i) - 1) / (1 - p_{\min}) \geq c^*(i) - \delta$. Since $\text{cost}(i) \leq c^*(i) + \delta$, it follows that $\text{cost}'(i) \leq \text{cost}(i) + 3\delta n \leq c'^*(i) + 2\delta + 3\delta n \leq c'^*(i) + \epsilon$, proving condition 4b. \square

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.

Approximate equilibria satisfy approximate versions of the envy-freeness and Pareto optimality properties. We say that an allocation x is ϵ -*envy-free* if $\sum_j u_{ij}x_{ij} \geq \sum_j u_{ij}x_{kj} - \epsilon$, for all $i, k \in A$; that is, no agent envies any other agent's bundle more than ϵ . We say that an allocation x is ϵ -*Pareto optimal (efficient)* if there is no other allocation y such that $\sum_j u_{ij}y_{ij} \geq \sum_j u_{ij}x_{ij} + \epsilon$ for all $i \in A$, with the inequality strict for at least one agent i . Clearly, for $\epsilon = 0$ these reduce to the standard notions of envy-freeness and Pareto optimality.

PROPOSITION 7.4. *If (p, x) is an ϵ -approximate equilibrium for a given one-sided matching market, then the allocation x satisfies the following properties:*

1. *It is 2ϵ -envy-free.*
2. *It is 2ϵ -Pareto optimal.*

Proof. For $\epsilon = 0$, these are the well-known properties that every HZ equilibrium is envy-free and Pareto optimal. So, assume $\epsilon > 0$.

1. The bundle x_k of agent k costs at most $1 + \epsilon$. Let x'_k be the unit bundle obtained by taking $x_k / (1 + \epsilon)$ and adding zero-priced goods to get size 1. The cost of x'_k is at most 1 and the value is at least $\sum_j u_{kj}x_{kj} / (1 + \epsilon)$. Therefore, $\sum_j u_{ij}x_{kj} / (1 + \epsilon)$ is at most the value $v^*(i)$ of the optimal bundle for agent i . Hence, $\sum_j u_{ij}x_{ij} \geq \sum_j u_{ij}x_{kj} / (1 + \epsilon) - \epsilon \geq \sum_j u_{ij}x_{kj} - 2\epsilon$, since $u_{ij} \in [0, 1]$ and $\sum_j x_{ij} = 1$.

2. Suppose that y is an allocation such that $\text{value}(y_i) = \sum_j u_{ij}y_{ij} \geq \text{value}(x_i) + 2\epsilon = \sum_j u_{ij}x_{ij} + 2\epsilon$ for all $i \in A$, with the inequality strict for some $i^* \in A$. Since x is an ϵ -approximate equilibrium, condition 4a implies that $\text{value}(y_i) > v^*(i)$, hence $\text{cost}(y_i) = \sum_j y_{ij}p_j > 1$ for all $i \in A$. Furthermore, we claim that $\text{cost}(y_i) \geq 1 + \epsilon$ for all $i \in A$. If $\text{cost}(y_i) < 1 + \epsilon$, then the unit bundle y'_i obtained by taking $y_i / \text{cost}(y_i)$ and adding zero-priced goods, has cost 1 and value greater than $\text{value}(y_i) / (1 + \epsilon) \geq \text{value}(y_i) - \epsilon \geq \text{value}(x_i) + \epsilon$, contradicting condition 4a for x . Therefore, $\text{cost}(y_i) \geq 1 + \epsilon$ for all $i \in A$, with the inequality strict for i^* .

Since in allocation y , every agent gets a unit bundle and all goods are sold, $\sum_i \text{cost}(y_i) = \sum_j p_j$. Hence, $\sum_j p_j > n(1 + \epsilon)$. This contradicts however the fact that allocation x satisfies by condition 3, $\sum_j p_j = \sum_i \text{cost}(x_i) \leq n(1 + \epsilon)$. \square

7.2. Membership in PPAD. 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 6.

DEFINITION 7.5. *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 $\text{size}(x)$ to denote the number of bits in a binary representation of x .

DEFINITION 7.6. *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 $q(|I| + \text{size}(x))$.*

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 [17] 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 6 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 6, 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 $d \leftarrow \min\{\frac{x_{ij}}{3n^2}, \frac{x_{il}}{3n^2}, ((u_{ik} - u_{ij})(p_l - p_k) - (u_{il} - u_{ik})(p_k - p_j))_+\}$. 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 6 for F hold also for F' .

The family \mathcal{F}' of these functions F'_I is clearly also polynomially computable and polynomially continuous. We shall show that, given an instance I of the matching market problem and a rational $\epsilon > 0$, we can pick a $\delta > 0$ such that $\log(1/\delta)$ is bounded by a polynomial in $|I|$ and $\log(1/\epsilon)$, and every weak δ -approximate fixed point of F'_I is a relaxed ϵ -approximate equilibrium of the market I .

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 $\delta = \epsilon/(n^{10}2^{6m})$. We shall show that every weak δ -fixed point of F'_I is a relaxed ϵ -approximate equilibrium of the matching market I . The proof follows and adapts the proof in Section 6 of the analogous statement for the exact fixed points.

LEMMA 7.7. *If (p, x) is a weak δ -fixed point of F' , then*

1. $\exists z \in G$ such that $p_z \leq \delta$.
2. $\forall i \in [n], \text{cost}(i) \leq 1 + 2n^2\delta$.
3. $\forall j \in [n], 1 - 3n^3\delta \leq \sum_{i \in A} x_{ij} \leq 1 + 3n^2\delta$.

4. $\sum_j p_j < 2n$.

Proof. 1. From Steps 2 and 3 of F_p , there is a good z such that the price of z in the output is 0. Therefore, $p_z \leq \delta$.

2. Suppose for some $i \in [n]$, $\text{cost}(i) > 1$. Then Steps 1 and 2 of F'_i will modify x_i since $r = \text{cost}(i) - 1 > 0$, and $\sum_k (1 - p_k)_+ > 0$ because some good z has $p_z \leq \delta$. The new cost is $\frac{\sum_j p_j x_{ij} + r \sum_j p_j (1 - p_j)_+}{1 + r \sum_j (1 - p_j)_+}$. This is at most $\frac{1 + r + r p_z (1 - p_z)}{1 + r (1 - p_z)} = 1 + \frac{r p_z (2 - p_z)}{1 + r (1 - p_z)} < 1 + 2\delta$. Steps 3, 4 of F'_i will either not change the allocation or if they do change it, the new cost will be less than 1. Steps 5, 6, 7 will not increase the cost. Thus, the final cost will be less than $1 + 2\delta$. Since F'_i changes each coordinate x_{ij} at most by δ and every price p_j is at most n , the total change in the cost is at most $n^2 \delta$. Therefore, the initial cost $\text{cost}(i)$ is at most $1 + 2\delta + n^2 \delta < 1 + 2n^2 \delta$.

3. Suppose that there is a good l such that $\sum_i x_{il} > 1 + 3n^2 \delta$. Since $\sum_j x_{ij} = 1$ for all agents $i \in [n]$, there must be a good k such that $\sum_i x_{ik} < 1 - 3n\delta$. We claim that then $p_k \leq \delta$, and that line 3 of F_p does not change the prices. Since $\sum_i x_{ik} < 1 - 3n\delta$, if $p_k > \delta$, then line 1 of F_p will decrease p_k by more than δ , and line 3 certainly does not increase it, contradicting $\|F_p(p, x) - p\|_\infty \leq \delta$. Thus, $p_k \leq \delta$, the price p_k will become 0 after line 1, hence $r = 0$ in line 2, and line 3 will not change the prices.

On the other hand, we claim that $p_l \geq n - \delta$. Since $\sum_i x_{il} > 1 + 3n^2 \delta$, if $p_l < n - \delta$, then line 1 of F_p will increase p_l by more than δ , and since line 3 does not change the prices, the final value of p_l exceeds the initial value by more than δ , contradicting the assumption that (p, x) is a δ -fixed point.

But $\text{cost}(i) = \sum_j p_j x_{ij} \leq 1 + 2n^2 \delta$ for all $i \in [n]$ implies that $\sum_i \sum_j p_j x_{ij} \leq n(1 + 2n^2 \delta)$, which contradicts the fact that $p_l \geq n - \delta$ and $\sum_i x_{il} > 1 + 3n^2 \delta$, hence $\sum_i p_l x_{il} > (n - \delta)(1 + 3n^2 \delta) > n(1 + 2n^2 \delta)$.

We conclude that $\sum_i x_{il} \leq 1 + 3n^2 \delta$ for all goods l . Since $\sum_i \sum_j x_{ij} = n$, it follows that $\sum_i x_{ij} \geq 1 - 3n^3 \delta$ for all goods j .

4. From part (2), $\sum_j p_j x_{ij} \leq 1 + 2n^2 \delta$ for all agents i , hence $\sum_i \sum_j p_j x_{ij} = \sum_j p_j \sum_i x_{ij} \leq n(1 + 2n^2 \delta)$. From part (3), $\sum_i x_{ij} \geq 1 - 3n^3 \delta$ for all j .

Therefore, $\sum_j p_j \leq \frac{n(1 + 2n^2 \delta)}{1 - 3n^3 \delta} < 2n$. \square

In the case of approximate fixed points, it is possible that multiple steps of F'_i modify the allocation. However, as we will see, because of Lemma 6.1, none of the steps can change the value or the cost by a large amount, because then the other steps cannot reverse the change. Note that if two bundles of an agent differ by at most δ in every coordinate, then their values differ by at most $n\delta$ (because all utilities are in $[0, 1]$), and their costs differ by at most $2n\delta$ (because the sum of the prices is less than $2n$). This holds in particular for the values and the costs of the input and the output allocation of each function F'_i when the input is a weak δ -fixed point.

All the steps of F'_i weakly increase the value of the allocation, except possibly for step 2. Since r in step 1 is $(\text{cost}(i) - 1)_+ \leq 2n^2 \delta$, the changes in each coordinate x_{ij} in step 2 are “small”: From the update formula in step 2, x_{ij} can increase at most by $r \leq 2n^2 \delta$. Thus, the value can increase in step 1 at most by $2n^3 \delta$. On the other hand, coordinate x_{ij} may decrease at most by $x_{ij}(1 - \frac{1}{1 + r \sum_k (1 - p_k)_+}) \leq x_{ij} r n \leq x_{ij} 2n^3 \delta$. Therefore the value can decrease in step 1 also at most by $2n^3 \delta$. As we observed above, the value of the output allocation of F'_i cannot differ from that of the input allocation by more than $n\delta$. Thus, we conclude:

COROLLARY 7.8. *If (p, x) is a weak δ -fixed point, then no step of F'_i changes the value of the allocation by more than $2n^3\delta + n\delta$.*

All steps of F'_i weakly decrease the cost, except possibly for step 4. We show that step 4 does not change the allocation significantly, and thus does not increase the cost very much.

LEMMA 7.9. *Suppose that (p, x) is a weak δ -fixed point of F' . If t in Step 3 of F'_i satisfies $t > 3n^5\delta 2^{2m}$ then the value of x_i is within ϵ of the value $v^*(i)$ of the optimal bundle for agent i under prices p , and the cost of x_i is within ϵ of the minimum cost of an optimal bundle.*

Proof. Steps 1, 2 can decrease the cost at most by $1 + r - \frac{1+r+r \sum_j p_j(1-p_j)_+}{1+r \sum_j (1-p_j)_+} \leq r(1+r) \sum_j (1-p_j)_+ \leq r(1+r)n \leq 3n^3\delta$. Since t in step 3 exceeds $3n^3\delta$, it follows that the cost of the input allocation x_i is not greater than 1. Therefore, steps 1, 2 do not modify x_i .

Let $B'_i = \{j | x_{ij} > 3n^3\delta 2^{2m}\}$. Suppose that there is a good $k \in B'_i - G_i^*$. Then d in step 4 for good k satisfies $d \geq 3n^3\delta 2^{2m}$. The change of the allocation in step 4 increases the value by at least $d(u_{ii^*} - u_{ik}) \geq d/2^{2m} \geq 3n^3\delta$, contradicting Corollary 7.8.

Therefore, $B'_i \subseteq G_i^*$. Let u be the utility for agent i of the goods in G_i^* (the maximum utility). The goods $j \notin G_i^*$ have $x_{ij} \leq 3n^3\delta 2^{2m}$. Therefore the value of x_i is at least $u - 3n^4\delta 2^{2m} > u - \epsilon \geq v^*(i) - \epsilon$.

We show now the claim about the cost. If the min-cost optimal bundle has cost 1, then the claim follows from Lemma 7.7. So assume it has cost < 1 , i.e. it is of type A and consists of goods in G_i^* . Let k be a good in G_i^* with minimum price. The minimum cost of an optimal bundle is p_k .

Step 4 may move some probability mass from goods that are not in G_i^* , hence not in B'_i , to i^* . Since $x_{ij} \leq 3n^3\delta 2^{2m}$ for all $j \notin B'_i$, the total mass moved is at most $3n^4\delta 2^{2m}$, and the cost is increased at most by $3n^5\delta 2^{2m}$. Since the cost of the output allocation of F'_i is within $2n\delta$ of the cost of the input allocation, and all other steps of F'_i weakly decrease the cost, it follows that no step of F'_i can decrease the cost by more than $3n^5\delta 2^{2m} + 2n\delta$.

Let $s = 4n^3\sqrt{\delta} 2^m$, and let $\hat{B}_i = \{j | x_{ij} > s\}$. Clearly, $s > 3n^3\delta 2^{2m}$, and thus $\hat{B}_i \subseteq B'_i \subseteq G_i^*$. We claim that every good $j \in \hat{B}_i$ has price $p_j \leq p_k + s$. If not, then step 5 for the pair j, k will have $d \geq s/3$, and it will decrease the cost by $\frac{d}{n}(p_j - p_k) > \frac{s^2}{3n} > 3n^5\delta 2^{2m} + 2n\delta$, a contradiction. Therefore, $p_j \leq p_k + s$ for all $j \in \hat{B}_i$. The allocation x_i has probability mass at most ns in the goods that are not in \hat{B}_i , and thus their cost is at most n^2s . Therefore the cost of x_i is at most $p_k + s + n^2s < p_k + \epsilon$. \square

We assume henceforth that t in step 3 is at most $t_0 = 3n^5\delta 2^{2m}$. Step 4 increases the cost at most by t and the other steps of F'_i weakly decrease the cost. Since the difference between the final and the initial cost is at most $2n\delta$, we have:

COROLLARY 7.10. *No step of F'_i decreases the cost of the allocation by more than $t_0 + 2n\delta < 4n^5\delta 2^{2m}$.*

We show now the approximate optimality of the agents' bundles in an approximate fixed point.

LEMMA 7.11. *If (p, x) is a weak δ -fixed point of F' , then the value of x_i is within ϵ of the optimal value of a bundle for agent i at prices p , and the cost of x_i is within*

ϵ of the minimum cost among optimal bundles.

Proof. Lemma 7.9 showed the result in the case that t in step 3 satisfies $t > t_0 = 3n^5\delta 2^{2m}$. So assume henceforth that $t \leq t_0$. Thus, the cost of the allocation after step 2 is $\geq 1 - t_0$. The cost of the input allocation x_i is at least as great, and is at most $1 + 2n^2\delta$ by Lemma 7.7. It follows that the cost of the input allocation x_i , as well as the allocations after step 2 and after step 4 are all in the interval $[1 - t_0, 1 + 2n^2\delta]$ (i.e., they are close to 1).

Let x'_i be the allocation after step 4. Let $\text{value}'(i)$ be the value of x'_i and $\text{value}(i)$ the value of x_i . As we observed earlier, steps 1, 2 change the value of the allocation at most by $2n^3\delta$, and step 3, 4 change each x_{ij} at most by t_0/n^2 , hence they increase the value at most by t_0/n . Therefore, $\text{value}(i) - 2n^3\delta - (t_0/n) \geq \text{value}'(i) - \epsilon/2$.

Let $B'_i = \{j | x'_{ij} > s\}$, where $s = 4n^3\sqrt{\delta}2^m$. We start with some useful properties of the goods in B'_i .

CLAIM 7.12. *For every good $j \in B'_i$ and every good k with $u_{ik} \geq u_{ij}$, it holds that $p_j \leq p_k + s$.*

Proof. Suppose the claim is not true and consider step 5 for the pair j, k . We have $d \geq s/3$, and step 5 decreases the cost by $\frac{d}{n}(p_j - p_k) > \frac{s^2}{3n} > 4n^5\delta 2^{2m}$, in contradiction to Corollary 7.10. \square

Thus, every good in B'_i has price that is close to the minimum price among goods with the same or higher utility. On the other hand, goods with strictly higher utility must have distinctly higher price:

CLAIM 7.13. *If $j \in B'_i$ and $p_j \geq 1/2$, then every good l with higher utility $u_{il} > u_{ij}$ has price $p_l > p_j + 2^{-2m-2}$. The same holds also for all goods j in an optimal bundle.*

Proof. Let j be a good in B'_i and let l be another good such that $u_{il} > u_{ij}$. Let z be a good with minimum price. By Lemma 7.7, $p_z \leq \delta$, hence $u_{iz} < u_{ij}$. Consider step 6 for the triple z, j, l . We have $d = \min\{\frac{x_{ij}}{3n^2}, \Delta\}$, where $\Delta = (u_{il} - u_{ij})(p_j - p_z) - (u_{ij} - u_{iz})(p_l - p_j)$. We know that $u_{il} - u_{ij} \geq 2^{-2m}$, $p_j - p_z \geq (1/2) - \delta$. If $p_l - p_j \leq 2^{-2m-2}$ then $\Delta \geq 2^{-2m-3} > 3n^2s$. Thus, $d \geq \frac{s}{3n^2}$, step 6 will modify the allocation and decrease the cost by $\frac{d\Delta}{u_l - u_z} > s^2 > 4n^6\delta 2^{2m}$, contradicting Corollary 7.10. Therefore, $p_l > p_j + 2^{-2m-2}$.

The argument for the case that j is a good in an optimal bundle is similar. If $p_l \leq p_j + 2^{-2m-2}$, then applying step 6 for the triple z, j, l to the optimal bundle will reduce its cost while keeping the same value, and then its value can be increased by further transferring some probability mass from good j to l , contradicting the optimality of the bundle. \square

We will prove now the approximate optimality of the allocations x'_i and x_i . We distinguish cases depending on the type of an optimal bundle for the prices p .

Case 1. *The optimal bundle is of type A or B, i.e., there is a good $k \in G_i^*$ with price $p_k \leq 1$.*

Let p_k be the smallest price of a good in G_i^* ; this is also the minimum cost of an optimal bundle. The value $v^*(i)$ of the optimal bundle is u , the maximum utility of a good. We argue that most of the probability mass of x'_i is allocated to goods in G_i^* . The goods not in B'_i have total size at most ns and cost at most n^2s . The goods in B'_i have price at most $p_k + s$ by Claim 7.12. Since the cost of x'_i is close to 1, B'_i must contain goods with price close to 1, therefore p_k must be close to 1. Specifically, the cost of x'_i is at least $1 - t_0$ and at most $p_k + s + n^2s$, hence $p_k \geq 1 - t_0 - n^2s - s$. The goods in $B'_i \setminus G_i^*$ have price at most $p_k - 2^{-2m-2}$ by

Claim 7.13. If the total size of the goods in $B'_i \setminus G_i^*$ is y , then the cost of x'_i is at most $n^2s + y(p_k - 2^{-2m-2}) + (1 - y - ns)(p_k + s)$. Since the cost is at least $1 - t_0$, it follows that $y \leq 2^{2m+2}(n^2s + t_0) \leq 2^{2m+3}n^2s$. Therefore, the value of x'_i is at least $(1 - ns - y)u \geq u - \epsilon/2$. Hence the value of x_i is at least $v^*(i) - \epsilon$. The cost of x'_i is at most $p_k + s + n^2s < p_k + \epsilon/2$, hence the cost of x_i is less than $p_k + \epsilon$.

We assume henceforth that the minimum price of a good in G_i^* is > 1 , thus the optimal bundle is of type C or D and has cost=1. The claim of the lemma about the cost thus holds by Lemma 7.7, and we only need to prove the claim about the value.

Case 2. The optimal bundle is of type C.

Thus the optimal bundle has cost 1, and contains goods with the same utility, $v^*(i)$, and price 1. Let k be an optimal good. All the goods of B'_i with utility strictly smaller than u_{ik} have price $\leq 1 - 2^{-2m-2}$ (by Claim 7.13). Let y be the total size of these goods. The goods of B'_i with utility u_{ik} have price at most $1 + s$ (by Claim 7.12). Suppose that B'_i does not have any goods with utility $> u_{ik}$. Then the cost of x'_i is at most $n^2s + y(1 - 2^{-2m-2}) + (1 - y - ns)(1 + s)$. Since the cost is at least $1 - t_0$, it follows that $y \leq 2^{2m+3}n^2s$. Therefore, the value of x'_i is at least $(1 - ns - y)u_{ik} \geq u_{ik} - \epsilon/2$, from which it follows that $\text{value}(i) \geq v^*(i) - \epsilon$.

We assume thus that $y > 2^{2m+3}n^2s$, which means that there are goods in B'_i with utility $> u_{ik}$, and there are also goods in B'_i with utility $< u_{ik}$ (since $y > 0$). Let $L = \{j \in B'_i | u_{ij} < u_{ik}\}$, $R = \{l \in B'_i | u_{il} > u_{ik}\}$. By Claim 7.13, every good $j \in L$ has price $p_j < p_k - 2^{-2m-2} = 1 - 2^{-2m-2}$, and every good $l \in R$ has price $p_l > 1 - 2^{-2m-2}$.

For any good $j \in L$ and any good $l \in R$, consider Step 7 of F'_i for the triple of goods j, k, l . Let α_i, μ_i be the optimal dual values. We have:

$$\alpha_i p_j = g_j + u_{ij} - \mu_i, \quad \alpha_i p_k = \alpha_i = u_{ik} - \mu_i, \quad \alpha_i p_l = g_l + u_{il} - \mu_i$$

where $g_j, g_l \geq 0$. We have $\alpha_i(1 - p_j) = u_{ik} - u_{ij} - g_j \leq 1$ and $1 - p_j \geq 2^{-2m-2}$ (by Claim 7.13), hence $\alpha_i \leq 2^{2m+2}$. The quantity $\Delta = (u_{ik} - u_{ij})(p_l - p_k) - (u_{il} - u_{ik})(p_k - p_j)$ is equal to $g_l(p_k - p_j) + g_j(p_l - p_k)$. If there is a $l \in R$ such that $g_l \geq n^2s2^{2m+2}$, and we let j be any element of L , then the quantity Δ for the triple j, k, l is at least n^2s , thus $d \geq \frac{s}{3n^2}$, and step 7 will decrease the cost by $\frac{d\Delta}{u_{il} - u_{ij}} \geq \frac{s^2}{3} > 4n^6\delta2^{2m}$, contradicting Corollary 7.10. Similarly, if there is a $j \in L$ such that $g_j \geq n^2s2^{2m+2}$, and we take l to be any element of R , the quantity Δ for the triple j, k, l will be at least n^2s , leading to the same contradiction. We conclude that $g_j < n^2s2^{2m+2}$ for all $j \in L \cup R$. Note that for all $j \in S = \{j \in B'_i | u_{ij} = u_{ik}\}$, we have $p_j \leq p_k + s = 1 + s$, hence $\alpha_i p_j \leq u_{ij} - \mu_i + \alpha_i s \leq u_{ij} - \mu_i + s2^{2m+2}$, i.e. $g_j \leq s2^{2m+2}$.

Thus, for all $j \in B'_i$, we have $\alpha_i p_j \leq u_{ij} - \mu_i + n^2s2^{2m+2}$. Multiplying each equation by x_{ij} and summing over all $j \in B'_i$ we get that $\alpha_i \sum_{j \in B'_i} x_{ij} p_j \leq \sum_{j \in B'_i} x_{ij} u_{ij} - \mu_i \sum_{j \in B'_i} x_{ij} + n^2s2^{2m+2} \sum_{j \in B'_i} x_{ij}$. The left hand side is α_i times the cost of x'_i , except for the goods that are not in B'_i , hence it is at least $\alpha_i(1 - t_0 - n^2s)$. We have also $\sum_{j \in B'_i} x_{ij} \geq 1 - ns$. The value of x'_i is at least $\sum_{j \in B'_i} x_{ij} u_{ij}$ and the optimal value $v^*(i)$ is equal to $\alpha_i + \mu_i$. Thus the difference between $v^*(i)$ and the value of x'_i is $v^*(i) - \text{value}'(i) \leq \alpha_i(t_0 + n^2s) + n^2s2^{2m+2} + \mu_i ns \leq 2^{2m+2}(t_0 + n^2s) + n^2s2^{2m+2} + ns < \epsilon/2$. It follows then as before that the initial value $\text{value}(i) > v^*(i) - \epsilon$.

Case 3. The optimal bundle is of type D.

The optimal bundle has cost 1 and contains some good l with price > 1 and some good j with price < 1 . Clearly $u_{ij} < u_{il}$. Let α_i, μ_i be again the optimal dual values. We have

$$\alpha_i p_j = u_{ij} - \mu_i, \quad \alpha_i p_l = u_{il} - \mu_i, \quad \alpha_i = v^*(i) - \mu_i$$

Note that $\alpha_i = \frac{u_{il}-u_{ij}}{p_l-p_j} < 2^{2m+2}$, since $p_l - p_j > 2^{-2m-2}$ by Claim 7.13. For every other good k , we have $\alpha_i p_k = u_{ik} - \mu_i + g_k$, where $g_k \geq 0$. We say that k is a *near-optimal* good if $g_k \leq n^2 s 2^{2m+2}$, and k is *very suboptimal* if $g_k > n^2 s 2^{2m+2}$. Note that if a good $k \in B'_i$ has equal utility to u_{ij} or u_{il} then its price is within s of p_j or p_l respectively (by Claim 7.12), hence $g_k \leq \alpha_i s \leq s 2^{2m+2}$, i.e. k is near-optimal.

CLAIM 7.14. *Let y be the total size of the very suboptimal goods in B'_i . If $y \leq 2^{2m+3} n^2 s$, then $\text{value}'(i) \geq v^*(i) - \epsilon/2$ and $\text{value}(i) \geq v^*(i) - \epsilon$.*

Proof. Let N_i be the set of near-optimal goods of B'_i . For every $k \in N_i$ we have $\alpha_i p_k \leq u_{ik} - \mu_i + g$, where $g = n^2 s 2^{2m+2}$. Multiplying each equation by x_{ik} and summing up over all $k \in N_i$, we get

$$\alpha_i \sum_{k \in N_i} p_k x_{ik} \leq \sum_{k \in N_i} u_{ik} x_{ik} - \mu_i \sum_{k \in N_i} x_{ik} + g \sum_{k \in N_i} x_{ik}$$

The total size of the goods in N_i is $\sum_{k \in N_i} x_{ik} \geq 1 - ns - y$. Their cost, $\sum_{k \in N_i} p_k x_{ik}$ is at least $1 - t_0 - n^2 s - ny$. Since $\text{value}'(i) \geq \sum_{k \in N_i} u_{ik} x_{ik}$ and $v^*(i) = \alpha_i + \mu_i$, we have:

$$v^*(i) - \text{value}'(i) \leq \alpha_i(t_0 + n^2 s + ny) + \mu_i(ns + y) + g$$

Since $\alpha_i \leq 2^{2m+2}$, $\mu_i \leq 1$, and from the assumed upper bounds on y and g , we conclude that $v^*(i) - \text{value}'(i) \leq \epsilon/2$. This implies as before that $\text{value}(i) \geq v^*(i) - \epsilon$. \square

Thus, assume that $y > 2^{2m+3} n^2 s$. This means in particular that B'_i contains some very suboptimal goods. We distinguish cases depending on how their utility compares to the utilities u_{ij}, u_{il} of the goods j, l in the optimal bundle. We will derive in each case a contradiction.

Subcase 1. There is a very suboptimal good $k \in B'_i$ such that $u_{ij} < u_{ik} < u_{il}$. Consider step 6 for the triple j, k, l . The quantity $\Delta = (u_{il} - u_{ik})(p_k - p_j) - (u_{ik} - u_{ij})(p_l - p_k)$ is equal to $g_k(p_l - p_j)$. We have $g_k \geq n^2 s 2^{2m+2}$ and $p_l - p_j \geq 2^{-2m-2}$ (by Claim 7.13), thus $\Delta \geq n^2 s$. Therefore the parameter d in step 6 is $d \geq s/3n^2$, and step 6 decreases the cost by $\frac{d\Delta}{u_{il}-u_{ij}} \geq s^2/3 > 4n^6 \delta 2^{2m}$, contradicting Corollary 7.10.

Subcase 2. There is a very suboptimal good $h \in B'_i$ such that $u_{ih} > u_{il}$. If all the goods in B'_i have utility $\geq u_{il}$, then the value of x'_i is $\text{value}'(i) \geq (1 - ns)u_{il} \geq u_{il} - \epsilon/2 \geq v^*(i) - \epsilon/2$, and the result follows. Thus, assume that B'_i has a good k with $u_{ik} < u_{il}$. Consider step 7 for the triple k, l, h . The quantity $\Delta = (u_{il} - u_{ik})(p_h - p_l) - (u_{ih} - u_{il})(p_l - p_k)$ is equal to $g_h(p_l - p_k) + g_k(p_h - p_l)$. Since $g_h \geq n^2 s 2^{2m+2}$ and $p_l - p_k \geq 2^{-2m-2}$ (by Claim 7.13), it follows that $\Delta \geq n^2 s$. Thus, $d \geq s/3n^2$, and step 7 will decrease the cost again by $\frac{d\Delta}{u_{il}-u_{ij}} \geq s^2/3$, contradicting Corollary 7.10.

Subcase 3. All very suboptimal goods k of B'_i have $u_{ik} < u_{ij}$. Note that then all very suboptimal goods k have price $p_k \leq u_{ij} - 2^{-2m-2} < 1 - 2^{-2m-2}$ by Claim 7.13. We claim that B'_i must contain a good h with utility $> u_{ij}$. For, if all goods in B'_i have utility $\leq u_{ij}$, then they all have price $\leq p_j + s < 1 + s$, and then we can argue as in Case 2 that the total size of the goods of B'_i with price $\leq 1 - 2^{-2m-2}$ must be at most $2^{2m+3} n^2 s$, contradicting the fact that the size y of the very suboptimal goods of B'_i is more than $2^{2m+3} n^2 s$.

Thus, let h be a good of B_i with utility $u_{ih} > u_{ij}$, and consider Step 7 for the triple k, j, h . The quantity $\Delta = (u_{ij} - u_{ik})(p_h - p_j) - (u_{ih} - u_{ij})(p_j - p_k)$ is equal to $g_h(p_j - p_k) + g_k(p_h - p_j)$. Since $g_k \geq n^2 s 2^{2m+2}$ and $p_h - p_j \geq 2^{-2m-2}$, it follows that $\Delta \geq n^2 s$. Thus, $d \geq s/3n^2$, and step 7 will decrease the cost again by $\frac{d\Delta}{u_{il}-u_{ij}} \geq s^2/3$, contradicting Corollary 7.10. \square

8. Discussion. A major open question is proving a stronger hardness result for the computation, to desired precision, of an exact HZ equilibrium. We conjecture that this problem is FIXP-complete; resolving this remains a challenging open problem.

We have given a strongly polynomial algorithm for computing an HZ equilibrium for the case of bi-valued utilities and [8] have shown that the approximate HZ equilibrium problem is PPAD-hard even when all agents have 4-valued utilities. That leaves the case of tri-valued utilities. We believe that this case also has instances with only irrational equilibria; perhaps even for utilities $\{0, \frac{1}{2}, 1\}$. Finding such an example or proving rationality is non-trivial and we leave it as an open problem. Furthermore, it is possible that even this case is intractable to solve exactly or approximately.

We also leave the open question of finding other special cases, besides the bi-valued case, for which an exact or approximate HZ equilibrium is easy to compute.

9. 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 [26]. We also thank the reviewers for their constructive comments.

Vazirani was supported in part by NSF grants CCF-1815901 and CCF-2230414; Yannakakis was supported in part by NSF grants CCF-1703925, CCF-2212233 and CCF-2332922.

Appendix A. Rationality and structure of HZ equilibria.

In this section we give a simple sufficient condition for the existence of a rational HZ equilibrium, and use it to show that instances with three goods have a rational equilibrium.

Let I be a general instance with a set A of agents and a set G of goods. We will allow here the goods to have a positive, integer supply s_j that is possibly greater than 1 (as in the model in the original HZ paper [26]), where the sum $\sum_{j \in G} s_j$ of the supplies is equal to the number of agents. As is well known, this model can be reduced to the model where each good has a supply of one unit, by creating s_j copies of each good j . Since this replication increases however the number of goods, and we want to show rationality more generally for the case of three goods and any number of agents, we will allow here general integer supplies $s_j > 0$ for the goods.

Consider an HZ equilibrium $q^* = (p^*, x^*)$ for I . For each agent i , let $B_i = \{j \in G \mid x_{ij}^* > 0\}$ be the bundle of items that i buys in q^* . We call the family of bundles $\{B_i \mid i \in A\}$ the *structure* of the equilibrium q^* . We will give a set of constraints that characterize equilibria with a given structure.

Given the structure of an equilibrium q^* we can infer some properties of the equilibrium. If $B_i \subseteq G_i^*$ (i.e. all goods in the bundle of agent i have maximum utility for i) then they must all have the same price, the price must be at most 1, and all other goods in G_i^* must have equal or higher price. On the other hand, if B_i is not a subset of G_i^* then $\text{cost}^*(i) = \sum_j p_j^* x_{ij}^* = 1$ in the solution q^* and the corresponding dual variable $\alpha_i^* > 0$ (agent i is of type C or D). Let $A_0 = \{i \in A \mid B_i \subseteq G_i^*\}$.

Let A_1 be the set of agents $i \notin A_0$ such that all goods j in B_i have the same utility u_{ij} for i (but not the maximum utility since $i \notin A_0$). For an agent $i \in A_1$, all the goods in B_i must have the same price in the equilibrium q^* , and the price must be exactly 1 since $\text{cost}^*(i) = 1$. For agents $i \in A_0 \cup A_1$, we know their average utility $v_i^* = \sum_j u_{ij} x_{ij}^*$ in q^* : it is the common utility of the goods in B_i .

Let $A_2 = A \setminus (A_0 \cup A_1)$. Given the structure $\{B_i \mid i \in A\}$ of an HZ equilibrium q^* , we construct a system (C) of constraints in the variables $\{p_j \mid j \in G\}, \{x_{ij} \mid i \in A, j \in G\}, \{v_i \mid i \in A_2\}, \{\beta_i \mid i \in A_1 \cup A_2\}$. The variables $v_i, i \in A_2$ represent the average

utilities of the agents, and the variables $\beta_i, i \in A_1 \cup A_2$ represent the reciprocals $1/\alpha_i$ of the dual variables. Note that $\alpha_i^* > 0$ for $i \notin A_0$, hence the reciprocals exist.

Instead of using the dual variables α_i, μ_i and the dual constraints $\alpha_i p_j + \mu_i \geq u_{ij}$, with equality for all $j \in B_i$ (see eq. (8) in Section 3 and the complementary slackness condition), it is more helpful here to use the equivalent constraint $p_j - 1 \geq \beta_i(u_{ij} - v_i)$, with equality for all $j \in B_i$. To see the equivalence, note that if we multiply by x_{ij} the equation $\alpha_i p_j + \mu_i = u_{ij}$ for $j \in B_i$ and sum over all $j \in B_i$, we obtain $\alpha_i \sum_j p_j x_{ij} + \mu_i = \sum_j u_{ij} x_{ij}$ using the fact that $x_{ij} = 0$ for $j \notin B_i$ and $\sum_j x_{ij} = 1$. Since $\text{cost}(i) = \sum_j p_j x_{ij} = 1$, it follows that $\alpha_i + \mu_i = v_i$. Hence $\alpha_i p_j + \mu_i \geq u_{ij}$ implies $p_j - 1 \geq \beta_i(u_{ij} - v_i)$. The other direction is similar.

Below is the system (C) of constraints for the given structure. In the constraints (4), (4') below we use v_i also for agents $i \in A_1$, but in this case they are not variables but fixed rational constants (the common utility of the goods in B_i for $i \in A_1$).

- (1) $\sum_i x_{ij} = s_j$ for all $j \in G$;
- (2) $\sum_j x_{ij} = 1$ for all $i \in A$;
- (3) $\sum_j u_{ij} x_{ij} = v_i$ for all $i \in A_2$;
- (4) $p_j - 1 = \beta_i(u_{ij} - v_i)$ for all $i \in A_1 \cup A_2, j \in B_i$;
- (4') $p_j - 1 \geq \beta_i(u_{ij} - v_i)$ for all $i \in A_1 \cup A_2, j \notin B_i$;
- (5) $p_j \leq 1$ for all $j \in \cup\{B_i | i \in A_0\}$;
- (6) $p_j \leq p_k$ for all $i \in A_0, j \in B_i, k \in G_i^*$;
- (7) $x_{ij} = 0$ for all $i \in A, j \notin B_i$;
- (7') $x_{ij} > 0$ for all $i \in A, j \in B_i$;
- (8) $\beta_i > 0$ for all $i \in A_1 \cup A_2$;

PROPOSITION A.1. *There is an HZ equilibrium with the given structure if and only if there is a solution to the constraint system (C).*

Proof. By our previous discussion, the parameters p, x, v, β of any HZ equilibrium q^* with the given structure satisfy the constraints (C).

Conversely, consider any solution (p, x, v, β) to the set (C) of constraints. Note that if we multiply by the same positive constant the differences $p_j - 1$ of all prices from 1, and the values β_i for all $i \in A_1 \cup A_2$ we obtain another solution to (C). Thus, we can assume without loss of generality that all prices in the solution are non-negative by scaling if necessary their difference from 1 and the β_i . From constraints (1) and (2) of (C), every agent consumes exactly one unit and there are exactly s_j units of each good j consumed. Multiplying (4) by x_{ij} and summing over all $j \in B_i$, using (2), (7) and (3), we have that every agent $i \in A_1 \cup A_2$ satisfies $\sum_j p_j x_{ij} - 1 = \beta_i(\sum_j u_{ij} x_{ij} - v_i) = 0$, hence $\text{cost}(i) = \sum_j p_j x_{ij} = 1$. For agents $i \in A_0$, (2), (5) and (7) imply that $\text{cost}(i) = \sum_j p_j x_{ij} \leq 1$. Thus, the allocations x under prices p are feasible for all agents.

For agents $i \in A_1 \cup A_2$, (4) and (4') imply that allocation x is optimal, because any other feasible allocation y satisfies $0 \geq \sum_j p_j y_{ij} - 1 \geq \beta_i(\sum_j u_{ij} y_{ij} - v_i)$, hence $v_i \geq \sum_j u_{ij} y_{ij}$. For agents $i \in A_0$, optimality is implied from $B_i \subseteq G_i^*$; furthermore,

the minimality of $\text{cost}(i)$ among optimal allocations follows from (6). Therefore, every solution to the above set (C) of constraints yields a HZ equilibrium; the equilibrium has the given structure by (7), (7'). \square

The set (C) of constraints is not linear because of the constraints (4) and (4') for agents $i \in A_2$. Note that (4), (4') for agents $i \in A_1$ are linear because v_i is a constant for these agents, not a variable. Thus, if $A_2 = \emptyset$, then (C) is a system of linear inequalities with rational coefficients; hence it has a rational solution, since we assumed the existence of a HZ equilibrium with the given structure.

So, assume $A_2 \neq \emptyset$. We will give below a sufficient condition for (C) to have a rational solution. Form the bipartite graph with node set A_2 in one part and $\cup\{B_i | i \in A_2\}$ in the other part, and with edge set $\{(i, j) | i \in A_2, j \in B_i\}$. Merge any two good-nodes j, j' if there exists an agent $i \in A$ such that B_i contains both j, j' and $u_{ij} = u_{ij'}$. Note that any solution of (C) (and any equilibrium) must have $p_j = p_{j'}$ by constraints (4), (6).

Let H be the resulting bipartite graph after performing the above merging process. The graph H has node set $A_2 \cup G'$ where every node v of G' represents a set of goods that must have the same price. For each agent $i \in A_2$, let $N(i)$ denote the neighborhood (set of adjacent nodes) of node i in H . Observe that $|N(i)| \geq 2$ for all $i \in A_2$. To see this, recall that since $i \in A_2$, there are at least two goods $j, j' \in B_i$ such that $u_{ij} \neq u_{ij'}$. By constraint (4), any solution to (C) must have $p_j \neq p_{j'}$. If j and j' were merged in H , then there would be no solution to (C) , contradicting our assumption that there is an HZ equilibrium with the given structure.

LEMMA A.2. *Suppose that there is an HZ equilibrium with the given structure. If some agent $i \in A_2$ has $N(i) = G'$, then (C) has rational solution, and hence there is a rational HZ equilibrium with the given structure.*

Proof. Suppose that there is an HZ equilibrium with the given structure, and that $N(i) = G'$. We can assume without loss of generality that the value of β_i in the corresponding solution q^* of (C) is any arbitrary positive constant, by scaling if necessary all the quantities $\beta_k, k \in A_1 \cup A_2$ and $p_j - 1, j \in G$. So assume that $\beta_i = 1$ in the solution q^* . Constraints (4) imply that for every pair of goods $j, j' \in B_i$, their prices in q^* satisfy $p_j - p_{j'} = u_{ij} - u_{ij'}$, a rational number.

Consider any other agent $k \in A_2$ and two goods $l, l' \in B_k$ with $u_{kl} \neq u_{kl'}$. Since $N(i) = G'$, there are goods $j, j' \in B_i$ that are represented in H by the same nodes respectively as l, l' , hence $p_l = p_j$ and $p_{l'} = p_{j'}$. From constraints (4) for agent k , we have $p_l - p_{l'} = \beta_k(u_{kl} - u_{kl'})$, hence $\beta_k = (p_l - p_{l'}) / (u_{kl} - u_{kl'}) = (p_j - p_{j'}) / (u_{kl} - u_{kl'}) = (u_{ij} - u_{ij'}) / (u_{kl} - u_{kl'})$.

Thus, we can determine the values of the variables $\beta_k, k \in A_2$ in q^* , and they are all rational. We can substitute these values into the system (C) , making the constraint system linear, and solve the system to obtain a rational solution. \square

Remark. It can be shown that the same conclusion holds more generally if there is an agent $i \in A_2$ that can reach all nodes in H using the following search algorithm:
 1. Initialize the set R of reached agents to $\{i\}$ and the set S of reached good-nodes to $N(i)$.
 2. While there is an agent $k \in A_2 \setminus R$ such that $|N(k) \cap S| \geq 2$, add k to R and add $N(k)$ to S .

Using similar arguments as in Lemma A.2, we can show that if the search reaches all the nodes of H then we can determine the values of all $\beta_k, k \in A_2$ and they are all rational. We can then substitute these values in the system (C) , and solve the

resulting linear system to obtain a rational equilibrium. We omit the proof of this fact, as Lemma A.2 suffices for the next corollary.

COROLLARY A.3. *Every instance with three goods (and any number of agents) has a rational equilibrium.*

Proof. Let I be an instance with three goods and q^* a HZ equilibrium. Given the structure of q^* , let A_0, A_1, A_2 be the partition of the agents defined above. If $A_2 = \emptyset$ then there is a rational equilibrium. So assume $A_2 \neq \emptyset$, and let H be the bipartite graph defined above with node set $A_2 \cup G'$. Recall that $|N(i)| \geq 2$ for every $i \in A_2$, hence $|G'| \geq 2$. If $N(i) = G'$ for some $i \in A_2$ then the result follows from Lemma A.2.

Thus assume $N(i) \neq G'$ for all $i \in A_2$. This implies that $|G'| = 3$ and $|N(i)| = 2$ for all $i \in A_2$. This means in particular that $G' = G$, there is no merging of nodes, hence every agent i in $A_0 \cup A_1$ has $|B_i| = 1$, and consumes one unit of one good. Let s'_j be the remaining supplies of the three goods that are consumed by agents in A_2 . Every agent in A_2 has unequal utilities for the two goods in his bundle, hence the two goods in his bundle have unequal prices. If every pair of goods is the bundle of some agent, then the three goods must have distinct prices, which implies that an agent whose bundle is the pair of goods with the two lowest prices spends less than an agent whose bundle is the pair with the two highest prices, contradicting the fact that $\text{cost}(i) = 1$ for all agents $i \in A_2$. Therefore, one pair of goods, say $\{1, 2\}$, is not the bundle of any agent. Let n_1 be the number of agents with bundle $\{1, 3\}$ and n_2 the number of agents with bundle $\{2, 3\}$. Let $d_i = |p_j - 1|$ for $j = 1, 2, 3$.

Consider an agent i with bundle $B_i = \{1, 3\}$. Since $\text{cost}(i) = p_1 x_{i1} + p_3 x_{i3} = 1$ and $x_{i1} + x_{i3} = 1$, we have $x_{i1} = \frac{d_3}{d_1 + d_3}$ and $x_{i3} = \frac{d_1}{d_1 + d_3}$. Hence all the n_1 agents i with bundle $B_i = \{1, 3\}$ consume the same amount $x_{i1} = \frac{d_3}{d_1 + d_3}$ of good 1, and since no other agent in A_2 consumes good 1, we have $x_{i1} = s'_1/n_1$, and $x_{i3} = 1 - x_{i1} = 1 - s'_1/n_1$; both are rational numbers. Similarly, all n_2 agents with bundle $B_i = \{2, 3\}$ consume the same amount s'_2/n_2 of good 2 and $1 - s'_2/n_2$ of good 3, also rational amounts. If we set the minimum price to 0, and solve for the other two prices, they will be also rational numbers. \square

REFERENCES

- [1] 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.
- [2] S. BASU, R. POLLACK, AND M. 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] G. BIRKHOFF, *Tres observaciones sobre el algebra lineal*, Univ. Nac. Tucuman, Ser. A, 5 (1946), pp. 147–154.
- [4] A. BOGOMOLNAIA AND H. MOULIN, *A new solution to the random assignment problem*, Journal of Economic theory, 100 (2001), pp. 295–328.
- [5] A. BOGOMOLNAIA AND H. MOULIN, *Random matching under dichotomous preferences*, Econometrica, 72 (2004), pp. 257–279.
- [6] E. BUDISH, *The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes*, Journal of Political Economy, 119 (2011), pp. 1061–1103.
- [7] I. CARAGIANNIS, K. A. HANSEN, AND N. RATHI, *On the complexity of pareto-optimal and envy-free lotteries*, arXiv preprint arXiv:2307.12605, (2023).
- [8] T. CHEN, X. CHEN, B. PENG, AND M. YANNAKAKIS, *Computational hardness of the Hylland-Zeckhauser mechanism*, in Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA, SIAM, 2022, pp. 2253–2268.
- [9] 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.
- [10] X. CHEN, X. DENG, AND S.-H. TENG, *Settling the complexity of computing two-player Nash equilibria*, Journal of the ACM (JACM), 56 (2009), pp. 1–57.
- [11] X. CHEN, D. PAPARAS, AND M. YANNAKAKIS, *The complexity of non-monotone markets*, J. ACM, 64 (2017), pp. 20:1–20:56, <https://doi.org/10.1145/3064810>.
- [12] 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.
- [13] C. DASKALAKIS, P. W. GOLDBERG, AND C. H. PAPADIMITRIOU, *The complexity of computing a Nash equilibrium*, SIAM Journal on Computing, 39 (2009), pp. 195–259.
- [14] 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.
- [15] F. ECHENIQUE, N. IMMORLICA, AND V. V. VAZIRANI, eds., *Online and Matching-Based Market Design*, Cambridge University Press, 2023.
- [16] F. ECHENIQUE, A. MIRALLES, AND J. ZHANG, *Constrained pseudo-market equilibrium*, arXiv preprint arXiv:1909.05986, (2019).
- [17] K. ETESSAMI AND M. YANNAKAKIS, *On the complexity of Nash equilibria and other fixed points*, SIAM J. Comput., 39(6) (2010), pp. 2531–2597.
- [18] A. FILOS-RATSIKAS, S. K. S. FREDERIKSEN, AND J. ZHANG, *Social welfare in one-sided matchings: Random priority and beyond*, in Proc. 7th Intl. Symposium on Algorithmic Game Theory, 2014, pp. 1–12.
- [19] A. FILOS-RATSIKAS, K. A. HANSEN, K. HØGH, AND A. HOLLENDER, *FIXP-membership via convex optimization: Games, cakes, and markets*, in 62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022, IEEE, 2021, pp. 827–838.
- [20] S. I. FOR THE THEORY OF COMPUTING, *Online and matching-based market design*, 2019. <https://simons.berkeley.edu/programs/market2019>.
- [21] J. GARG, R. MEHTA, AND V. V. VAZIRANI, *Dichotomies in equilibrium computation and membership of PLC markets in FIXP*, Theory of Computing, 12 (2016), pp. 1–25.
- [22] 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.
- [23] J. GARG, T. TRÖBST, AND V. V. VAZIRANI, *One-sided matching markets with endowments: Equilibria and algorithms*, in 21st International Conference on Autonomous Agents and Multiagent Systems, AAMAS, 2022, pp. 463–471.
- [24] 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.
- [25] M. HOSSEINI AND V. V. VAZIRANI, *Nash-bargaining-based models for matching markets: One-sided and two-sided; fisher and arrow-debreu*, in Proceedings of the 2022 Conference on Innovations in Theoretical Computer Science, ITCS, 2022.
- [26] A. HYLLAND AND R. ZECKHAUSER, *The efficient allocation of individuals to positions*, Journal of Political economy, 87 (1979), pp. 293–314.
- [27] N. IMMORLICA, B. LUCIER, G. WEYL, AND J. MOLLNER, *Approximate efficiency in matching markets*, in International Conference on Web and Internet Economics, Springer, 2017, pp. 252–265.
- [28] P. LE, *Competitive equilibrium in the random assignment problem*, International Journal of Economic Theory, 13 (2017), pp. 369–385.
- [29] J. LI, Y. LIU, L. HUANG, AND P. TANG, *Egalitarian pairwise kidney exchange: fast algorithms via linear programming and parametric flow*, in Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems, 2014, pp. 445–452.
- [30] A. MCLENNAN, *Efficient disposal equilibria of pseudomarkets*, in Workshop on Game Theory, vol. 4, 2018, p. 8.
- [31] H. MOULIN, *Fair division in the age of internet*, Annual Review of Economics, (2018).
- [32] I. PANAGEAS, T. TRÖBST, AND V. V. VAZIRANI, *Time-efficient algorithms for Nash bargaining-based matching market models*, Submitted to WINE, (2024).
- [33] A. E. ROTH, T. SÖNMEZ, AND M. U. ÜNVER, *Pairwise kidney exchange*, Journal of Economic theory, 125 (2005), pp. 151–188.
- [34] L. SHAPLEY AND H. SCARF, *On cores and indivisibility*, Journal of mathematical economics, 1 (1974), pp. 23–37.
- [35] T. TROEBST AND V. V. VAZIRANI, *Cardinal-utility matching markets: The quest for envy-*

- 1722 *freeness, pareto-optimality, and efficient computability*, in Proceedings of the 2024 ACM
 1723 Conference on Economics and Computation, 2024.
- 1724 [36] V. V. VAZIRANI, *The notion of a rational convex program, and an algorithm for the Arrow-*
 1725 *Debreu Nash bargaining game*, J. Assoc. Comput. Mach., 59(2) (2012).
- 1726 [37] V. V. VAZIRANI AND M. YANNAKAKIS, *Market equilibria under separable, piecewise-linear, con-*
 1727 *cave utilities*, J. Assoc. Comput. Mach., 58(3) (2011).
- 1728 [38] V. V. VAZIRANI AND M. YANNAKAKIS, *Computational complexity of the Hylland-Zeckhauser*
 1729 *mechanism for one-sided matching markets*, in 12th Innovations in Theoretical Computer
 1730 Science Conference, ITCS, vol. 185 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum für
 1731 Informatik, 2021, pp. 59:1–59:19.
- 1732 [39] J. VON NEUMANN, *A certain zero-sum two-person game equivalent to the optimal assignment*
 1733 *problem*, Contributions to the Theory of Games, 2 (1953), pp. 5–12.
- 1734 [40] M. YANNAKAKIS, *Unpublished*, 2013.