informs
https://pubsonline.informs.org/journal/moor

MATHEMATICS OF OPERATIONS RESEARCH
Vol. 00, No. 0, Xxxxx 0000, pp. 000–000
ISSN 0364-765X, EISSN 1526-5471

Submitted to *Mathematics of Operations Research*

# Cardinal-Utility Matching Markets: The Quest for Envy-Freeness, Pareto-Optimality, and Efficient Computability

Thorben Tröbst, Vijay V. Vazirani
Department of Computer Science, University of California, Irvine, t.troebst@uci.edu, vazirani@ics.uci.edu



**Abstract.** Recent insights have left cardinal-utility matching markets in a state of flux: the celebrated pricing-based mechanism for one-sided cardinal-utility matching markets due to Hylland and Zeckhauser [26] (HZ), which had long eluded efficient algorithms, was finally shown to be PPAD-complete (Chen et al. [15], Vazirani and Yannakakis [37]). This raises the question: is there a polynomial time mechanism for one-sided cardinal-utility matching markets which achieves the desirable properties of HZ, i.e. envy-freeness (EF) and Pareto-optimality (PO)?

We show that the problem of finding an EF+PO lottery in a one-sided cardinal-utility matching market is by itself PPAD-complete. However, a $(2+\epsilon)$-approximately envy-free and Pareto-optimal lottery can be found in polynomial time using the Nash-bargaining-based mechanism of Hosseini and Vazirani [25]. This mechanism is also $(2+\epsilon)$-approximately incentive compatible.

We next turn to two-sided cardinal-utility matching markets, for which Bogomolnaia and Moulin [9] had shown that for a symmetric, bipartite two-sided matching market with $\{0,1\}$ utilities, rational EF+PO allocations exist. We prove that both these conditions are essential by giving negative results for an asymmetric $\{0,1\}$ utilities market and a symmetric $\{0,1,2\}$ utilities market. We also prove existence of justified-envy-free and weak Pareto-optimal lotteries.

**Funding:** This research was supported by NSF grant CCF-2230414.
**Key words:** Matching markets, market equilibria, mechanism design

## 1. Introduction

Ever since the invention of the Internet and the rise in mobile computing, one-sided and two-sided matching markets have become an important part of the economy, driving innovation in the field of matching-based market design; see Echenique et al. [19] for a detailed overview. Prominent commercial examples of such markets include Google's AdWords market (matching advertisers to user search queries), ride-hailing services (Uber, Lyft), food delivery (Doordash, Uber Eats), freelancing (Taskrabbit, Upwork), and vacation rentals (Airbnb, VRBO).

Furthermore, there are also many useful matching markets *without* monetary transfers which are the focus of this paper. Markets of this kind arise in a variety of scenarios where payments are impractical or even immoral. Well-known examples include the kidney donor market (matching kidney donors with transplant recipients), the National Resident Matching Program (matching doctors to hospitals), school choice (matching students to primary or secondary schools) and course allocation (matching students to courses).

In a *one-sided matching market* we are given a set $G$ of *goods* and a set $A$ of *agents* with $|A| = |G|$. Agents have preferences over the goods, and each agent is to be assigned exactly one good. The goal is to design a mechanism that takes in the preferences of the agents and finds a perfect matching such that certain desirable properties, most notably efficiency, fairness, and incentive compatibility, are achieved.

Since each agent must get exactly one good, it is generally not possible to assign indivisible goods in a fair way. For this reason, it is customary to study probability distributions—or *lotteries*—over matchings instead and to consider notions of *ex-ante* fairness, efficiency, etc.

Such matching markets can be broadly distinguished into two classes based on the kind of preferences the agents have over the goods. In an *ordinal-utility* matching market, each agent $i$ represents their preferences via a total (or sometimes partial) order $\prec_i \subseteq G^2$, whereas in a *cardinal-utility* matching market, each agent $i$ has a vector $(u_{ij})_{j \in G}$ of rational, nonnegative utilities instead.

Ordinal preferences are more restrictive, making them easier to elicit and also making it easier to design incentive-compatible mechanisms for them. Several well-known mechanisms such as Probabilistic Serial (Bogomolnaia and Moulin [8]) and Random Priority (also known as Random Serial Dictatorship) (Abdulkadiroglu and Sönmez [1]) have been developed in the ordinal setting.

A significant limitation of ordinal preferences is their inability to express the *magnitude* to which an agent likes a certain good, which can lead to a significant loss of efficiency. For a simple example, consider a course allocation problem in which students rank courses they wish to take. Student A ranks course 1 as 10/10 and course 2 as 1/10, while student B ranks course 1 as 10/10 and course 2 as 9/10. If there is only one remaining slot in each course, it makes sense to assign student A to course 1 and student B to course 2 in order to greatly increase the combined social welfare of the students. However, from the perspective of an ordinal mechanism, both students have identical preferences, and the only reasonable solution is to randomly assign the students to the courses.

In the previous example, we can improve the social welfare by using cardinal utilities: assign student A to course 1 and student B to course 2. However, this would come at the cost of envy-freeness: student B envies student A. Immorlica et al. [27] provide an even stronger example of the power of cardinal utilities. There are instances with $n$ agents and goods in which all agents have identical ordinal preferences and, as such, the goods would be evenly distributed among the agents under any reasonable ordinal mechanism.

But, remarkably, another allocation exists which is envy-free and improves the cardinal utilities of *every* agent by a factor of $\theta(\log n)$.

Consequently, cardinal utilities have found several applications in the real world. The popular fair division website Spliddit (`www.spliddit.org`) uses cardinal utilities for all of its tools to allow users to split goods, chores, rent, etc. among themselves in a fair and efficient way (Caragiannis et al. [14]). Several academic conferences within theoretical computer science crucially rely on a cardinal-utility matching market for their review process by asking their committee members to rank papers on a scale (e.g. -20 to 20) in order to assign papers to committee members for review. Cognomos (`www.cognomos.com`) successfully incorporates cardinal utilities to find course allocations for universities in a fair and efficient way (Budish et al. [11]).

When it comes to cardinal utilities, the most notable mechanism is that of Hylland and Zeckhauser [26], which is based on a pricing approach. It results in allocations which are (ex-ante) envy-free (EF) and Pareto-optimal (PO), thus satisfying the most common notions of fairness and efficiency respectively. Later, He et al. [24] showed that the HZ mechanism is asymptotically incentive compatible in a certain sense. Note that no mechanism is EF, PO, and also incentive compatible in the traditional sense (Zhou [40]). We review the HZ mechanism in Section 2.

A core issue with the approach is that *computing* an HZ equilibrium is hard in theory and in practice. Vazirani and Yannakakis [37] recently showed that there are instances in which every HZ equilibrium is irrational. They also showed that the problem of finding an approximate HZ equilibrium is in PPAD and conjectured that it is PPAD-complete. This conjecture was confirmed by Chen et al. [15] who proved the corresponding hardness result.

This motivated the search for alternative mechanisms that can achieve some or all of the desirable properties of HZ while being implementable in polynomial time. Fortunately, an alternative to the pricing-based mechanism for market models had been explored in the past: Vazirani [36] gave a Nash-bargaining-based mechanism for the linear Arrow-Debreu model. Building on this idea, Hosseini and Vazirani [25] recently proposed a Nash bargaining-based mechanism for a one-sided matching market; it is polynomial time computable and Pareto-optimal but not necessarily envy-free. We pose the question: is it possible to compute an envy-free and Pareto-optimal lottery in a one-sided cardinal-utility matching market in polynomial time?

Beyond one-sided matching markets and HZ, there are also *two-sided matching markets* in which we are matching agents to other agents. For example, this would apply to the aforementioned school choice market if we allow the schools to also have preferences over students, e.g. via aptitude scores. Except for a few highly restricted special cases (Bogomolnaia and Moulin [9], Roth et al. [34]), it was not known whether envy-free and Pareto-optimal lotteries even exist in this setting.

### 1.1. Our Contributions

**One-Sided Matching Markets** Our most significant contribution is to resolve the question regarding the complexity of finding EF+PO allocations by showing that the problem is PPAD-hard. Together with a recent result showing membership in PPAD by Caragiannis et al. [13], this shows that this problem is PPAD-complete. See Section 2.1 for a short introduction to the complexity class PPAD.

THEOREM (SECTION 2.3). *The problem of finding an EF+PO allocation in a one-sided cardinal-utility matching market is PPAD-hard.*

Our proof works through a polynomial reduction of approximate HZ to the problem of finding EF+PO allocations which is inspired by the fact that HZ allocations and EF+PO allocations coincide in certain *continuum markets* involving infinitely many agents and goods (Ashlagi and Shi [3]). The key idea is to take an HZ instance and add agents and goods so as to approximate such a continuum market without perturbing the HZ equilibria in the instance too much. However, the fact that this yields a working reduction is nonetheless surprising and requires additional ideas since it was already known that EF+PO allocations need not be approximately HZ, even in markets that converge to a continuum market in the limit (Miralles and Pycia [30]).

Along the way, we will also provide a simple polyhedral proof that there are always rational EF+PO allocations. This of course follows from the PPAD membership proof given by Caragiannis et al. [13] but our argument does not rely on the substantial amount of machinery inherent to proving PPAD membership.

Lastly, we show that the Nash-bargaining-based mechanisms for matching markets introduced by Hosseini and Vazirani [25] satisfy an approximate notion of envy-freeness and incentive compatibility.

THEOREM (SECTION 2.4). *The Nash bargaining solution for one-sided cardinal utility matching markets is 2-envy-free and 2-incentive compatible.*

Together with the algorithms given by Panageas et al. [32], this yields a polynomial time mechanism that is $(2+\epsilon)$-envy-free, $(2+\epsilon)$-incentive compatible and Pareto optimal. We remark that HZ is $(1+\epsilon)$-incentive compatible, however only in a certain asymptotic sense which requires that every agent and good has many copies (He et al. [24]). Our results support the Nash-bargaining-based mechanism as a more practical alternative mechanism for one-sided cardinal-utility matching markets. We believe that this is an important development in the growing field of matching-based market design.

**Two-Sided Matching Markets** For two-sided markets, the only cases in which it was previously known that EF+PO allocations exist are when the utilities are in $\{0,1\}$ and symmetric, i.e. each pair of agents either finds their match mutually agreeable or mutually disagreeable. In Section 3.2, we provide counterexamples that show that both of these conditions are necessary: if agents have $\{0,1,2\}$ utilities or asymmetric $\{0,1\}$ utilities, then EF+PO allocations may not exist.

Given this non-existence result, we give a notion of *justified envy-freeness* (JEF) which is related to—but to the best of our knowledge different from—notions of fractional stability from the stable matching literature. We show existence of rational JEF + weak PO allocations via a limiting argument, an equilibrium notion introduced by Manjunath [29], and similar polyhedral techniques as we used for one-sided markets.

THEOREM (SECTION 3.3). *In any two-sided cardinal-utility matching market, a rational JEF + weak PO allocation always exists.*

The Nash-bargaining-based approach by Hosseini and Vazirani [25] and the efficient algorithms by Panageas et al. [32] also extend to two-sided markets. However, in Section 3.4, we give a counterexample to show that in a Nash bargaining solution, agents can have $\theta(n)$-factor justified envy toward other agents.

**1.2. Related Work** Our work builds on the existing literature surrounding the Hylland Zeckhauser mechanism (Hylland and Zeckhauser [26]) and the complexity of computing HZ equilibria. Alaei et al. [2] give an algorithm to compute HZ equilibria which is based on the algebraic cell decomposition technique (Basu et al. [6]). However, this algorithm needs to enumerate at least $n^{5n^2}$ cells and is thus highly impractical even for small values of $n$. Vazirani and Yannakakis [37] give a polynomial time algorithm that computes HZ equilibria for $\{0, 1\}$ utilities, and more generally, when each agent's utilities come from a bi-valued set. They also show FIXP membership for the problem of computing HZ equilibria and PPAD membership for the problem of computing approximate HZ equilibria. Chen et al. [15] show the corresponding PPAD-hardness result, though it remains open whether finding an exact equilibrium is FIXP-hard.

The notion of envy-freeness comes from fair division where it was originally introduced in the context of dividing a single resource amongst the agents (Foley [22], Varian [35]), a problem that is now referred to as the cake cutting problem. It also features prominently in the literature on fair division of indivisible goods. Since it is generally impossible to achieve envy-freeness with indivisible goods, relaxations such as envy-freeness up to one good (EF1) (Budish [10]) or envy-freeness up to any good (EFX) (Caragiannis et al. [14]) are studied instead.

Cole and Tao [17] recently showed that envy-free and Pareto-optimal lotteries exist in a large class of (one-sided) fair division problems that in particular includes our setting. Building on this and recent results by Filos-Ratsikas et al. [21], Caragiannis et al. [13] established PPAD membership for this class of problems, though they leave open the question of showing PPAD-hardness, which we resolve here. They also show that maximizing social welfare over the set of envy-free lotteries is NP-hard, though their construction relies on a more general problem than the matching markets discussed in this paper.

Markets with a continuum of agents were introduced by Aumann [4]. Zhou [41] showed that in such continuum markets and under locally non-satiating utilities, envy-free and Pareto-optimal allocations coincide with allocations that come from competitive equilibria with equal incomes. In matching markets, the local non-satiation condition is not satisfied, but Ashlagi and Shi [3] show a similar equivalence for HZ. However, Miralles and Pycia [30] show that this holds only in the limit: for "large" markets, EF+PO allocations may not be supported by competitive equilibria from approximately equal incomes.

In order to deal with the intractability of HZ, Hosseini and Vazirani [25] recently proposed an alternative, Nash-bargaining-based mechanism for matching markets. Their approach works for one-sided and two-sided settings with both linear and non-linear utilities. Importantly, they show that their Nash bargaining solutions can be computed very efficiently in practice even on instances with thousands of agents. The idea of operating markets via Nash bargaining instead of pricing goes back to Vazirani [36] who used this approach for the linear Arrow Debreu market. We will introduce the Nash-bargaining-based mechanism in more detail in Section 2.4.

Panageas et al. [32] give algorithms for Nash bargaining in matching markets based on multiplicative weights update and conditional gradient descent which are efficient in practice and provide provable running times bounded by $\mathrm{poly}(n, 1/\epsilon)$. Aziz and Brown [5] show a reduction from HZ to Nash bargaining in the setting with $\{0,1\}$ utilities. They also note that Nash bargaining is not envy-free in general, though, as we will show, it is so in an approximate sense.

For two-sided markets, there have been several attempts at extending the equilibrium notion of Hylland and Zeckhauser. Manjunath [29] as well as Echenique et al. [20] introduce equilibrium notions. In both cases, personalized prices are required, i.e. each agent on one side sees a potentially different set of prices for all other agents on the other side. In Manjunath's equilibrium we will see that we can still get a kind of justified envy-freeness.

Restricted to symmetric $\{0,1\}$ utilities, HZ-like equilibria do exist as shown by Bogomolnaia and Moulin [9] for bipartite markets and Roth et al. [34] for non-bipartite markets. A polynomial time algorithm to compute such equilibria and therefore EF+PO allocations was later given by Li et al. [28].

Beyond this, two-sided markets have been mostly studied under ordinal preferences where stable matching, as introduced by Gale and Shapley [23], is the dominant solution concept. A notable exception is the work by Caragiannis et al. [12] who study the problem of finding a fractional stable matching under cardinal utilities that (approximately) maximizes social welfare.

## 2. One-Sided Matching Markets

In a one-sided matching market we are given a set $A$ of *agents* and a set $G$ of *goods*. We assume that $|A| = |G| = n$ since our goal is to assign exactly one good to each agent (a perfect matching) in a way that satisfies certain desirable properties. In this paper we will focus on *cardinal preferences*, that is each agent $i \in A$ has non-negative utilities $(u_{ij})_{j \in G}$ for every good. The most notable result in the study of cardinal matching markets is the celebrated Hylland-Zeckhauser mechanism (Hylland and Zeckhauser [26]) which we will briefly review here.

Hylland and Zeckhauser note that under cardinal utilities it is possible to reduce the case of *indivisible* goods to the case of *divisible* goods via the Birkhoff-von-Neumann theorem.

THEOREM 1 **(Birkhoff [7], Von Neumann [38]).** *Given a fractional perfect matching $(x_{ij})_{i\in A, j\in G}$, there are $O(n^2)$ integral perfect matchings $y^{(1)}, \dots, y^{(l)}$ and non-negative coefficients $\lambda_1, \dots, \lambda^l$ such that $\sum_{k=1}^{l} \lambda_k = 1$ and $x = \sum_{k=1}^{l} \lambda_k y^{(k)}$. Moreover, both the matchings and the coefficients can be found in polynomial time.*

This means that if we find a fractional perfect matching $x$, we can simply decompose it into a convex combination of integral perfect matchings and run a *lottery* over these matchings. The expected utility of agent $i$ is then $u_i \cdot x_i = \sum_{j\in G} u_{ij} x_{ij}$. Hence, we will—as Hylland and Zeckhauser did—concern ourselves primarily with fractional perfect matchings and linear utilities (which we will also just refer to as *allocations*) and their properties.

Given that we are now in the case of divisible goods with linear utilities, the key insight of Hylland and Zeckhauser is to employ the power of pricing by implementing a *pseudomarket*. Each agent is given one unit of fake currency and we determine prices on the goods as well as an allocation that together form a market equilibrium.

DEFINITION 1. A fractional assignment of goods to agents $(x_{ij})_{i\in A, j\in G}$ and non-negative prices $(p_j)_{j\in G}$ form an *HZ equilibrium* if and only if:

- $x$ is a fractional perfect matching,
- each agent $i$ spends at most their budget, i.e. $\sum_{j\in G} p_j x_{ij} \le 1$, and
- each agent $i$ gets a *cheapest utility-maximizing bundle*, i.e.

$$u_i \cdot x_i = \max \left\{ u_i \cdot y \;\middle|\; \sum_{j\in G} y_j = 1, p \cdot y \le 1 \right\},$$

$$p_j \cdot x_i = \min \left\{ p \cdot y \;\middle|\; \sum_{j\in G} y_j = 1, u_i \cdot y \ge u_i \cdot x_i \right\}.$$

We remark that the condition that the bundle be cheapest among all utility-maximizing bundles is a technical sublety. It is required in order to ensure that the resulting equilibrium allocations are Pareto-optimal.

THEOREM 2 **(Hylland and Zeckhauser [26]).** *An HZ equilibrium always exists. Moreover, if $(x, p)$ is an HZ equilibrium, then $x$ is envy-free and Pareto-optimal.*

Throughout this paper we are always referring to *ex-ante* envy-freeness and Pareto-optimality. The formal definitions are given below.

DEFINITION 2. Given some allocation $x$, agent $i$ has envy towards agent $i'$ if $u_i \cdot x_i < u_i \cdot x_{i'}$. The allocation $x$ is *envy-free* if no agent has envy towards any other agent.

DEFINITION 3. Given two allocations $x$ and $y$, $y$ is Pareto-better than $x$ if $u_i \cdot y_i \ge u_i \cdot x_i$ for all agents $i$ and $u_i \cdot y_i > u_i \cdot x_i$ for some agent $i$. The allocation $x$ is *Pareto-optimal* if no other allocation is Pareto-better than it.

**2.1. The Complexity Class PPAD** The most well-known notion of computational hardness is that of NP-hardness. However, note that when we are considering the computational complexity of total problems, such as finding an HZ equilibrium or an EF+PO allocation, we cannot meaningfully talk about NP-hardness because for these problems, the decision-version is trivial.

By Theorem 2, we know that there always exists an HZ equilibrium (and hence an EF+PO allocation). Problems of this nature, where we know that there always exists a solution and we "merely" have to find it, are called *total search problems*. The analogous class to NP in the space of total search problems is TFNP, which are those problems where solutions can be verified in polynomial time. Unfortunately, it is conjectured that there are no TFNP-complete problems. For this reason, we usually consider subclasses of TFNP such as PPP, PPA, or PLS instead. The class PPAD (*polynomial parity argument in digraphs*) is one of these subclasses of TFNP which contains those total search problems which can be proven to be total via a certain parity argument; see Papadimitriou [33].

A precise definition is outside of the scope of this paper and is not necessary in order to understand our results since we work directly via polynomial reductions. The important thing to note is that many prominent problems in algorithmic game theory and the computation of market equilibria have been shown to be PPAD-complete, most notably the problem of finding a Nash equilibrium [16, 18]. It is strongly believed that PPAD $\neq$ FP, i.e. that PPAD-complete problems are not polynomial time solvable.

**2.2. Rationality** Before discussing matters of computational complexity, we need to discuss whether the allocations that we are interested in even have finite representations. An allocation $x$ is *rational* if $x_{ij} \in \mathbb{Q}$ for all $i, j$. Likewise, we will say that an HZ equilibrium $(x, p)$ is *rational* if both $x$ and $p$ consist only of rational numbers.

In the case of HZ, there are instances which have only irrational equilibria (Vazirani and Yannakakis [37]). In contrast, we will show in this section that there are always rational EF+PO allocations. This follows from the way that PPAD-membership is proven by Caragiannis et al. [13], however our argument is simpler and introduces some basic facts which will be useful in later sections. The core observation is that fractional perfect matchings which are envy-free and Pareto-optimal can be characterized polyhedrally.

Let us start by considering the polytope $P_{\text{PM}}$ of all fractional perfect matchings in the given market.

$$P_{\text{PM}} := \left\{ (x_{ij})_{i\in A, j\in G} \;\middle|\; \begin{array}{lll} \sum_{j\in G} x_{ij} & = 1 & \forall i \in A, \\ \sum_{i\in A} x_{ij} & = 1 & \forall j \in G, \\ x_{ij} & \geq 0 & \forall i \in A, j \in G. \end{array} \right\}$$

It is well-known that Pareto-optimality can be characterized in terms of maximizing along a vector with strictly positive entries (Zadeh [39]). Since agents' utilities are linear and the feasible region is a polytope, one can obtain the corresponding vector in polynomial time using linear programming.

LEMMA 1. *$x^\star \in P_{\text{PM}}$ is Pareto-optimal if and only if there exist positive $(\alpha_i)_{i\in A}$ such that $x^\star$ maximizes $\phi(x) := \sum_{i\in A} \alpha_i u_i \cdot x_i$ over all $x \in P_{\text{PM}}$. Moreover, if $x^\star$ is rational, $\alpha$ can be computed in polynomial time.*

See Appendix A for the proof. Lemma 1 characterizes the Pareto-optimal allocations. Moreover, the envy-free allocations themselves form the polytope $P_{\mathrm{EF}}$ shown below.

$$P_{\mathrm{EF}} := \left\{ (x_{ij})_{i\in A, j\in G} \;\middle|\; \begin{array}{lll} \sum_{j\in G} x_{ij} & = 1 & \forall i \in A, \\ \sum_{i\in A} x_{ij} & = 1 & \forall j \in G, \\ u_i \cdot x_i - u_i \cdot x_{i'} & \geq 0 & \forall i, i' \in A, \\ x_{ij} & \geq 0 & \forall i \in A, j \in G. \end{array} \right\}$$

THEOREM 3. *There is always an EF+PO allocation which is a vertex of $P_{\mathrm{EF}}$ and is thus rational.*

*Proof.* We know that at least one EF+PO allocation $x^\star$ exists since the HZ equilibrium allocation has these properties. By Lemma 1, $x^\star$ maximizes $\phi(x) := \sum_{i\in A} \alpha_i u_i \cdot x_i$ over $P_{\mathrm{PM}}$ for some strictly positive $\alpha$ vector.

Now consider the linear program $\max\{\phi(x) \mid x \in P_{\mathrm{EF}}\}$. $P_{\mathrm{EF}}$ is a polytope so let $x$ be a vertex solution to this LP. Clearly $x$ is envy-free since $x \in P_{\mathrm{EF}}$. But since $x^\star$ is also in $P_{\mathrm{EF}}$ and $P_{\mathrm{EF}} \subseteq P_{\mathrm{PM}}$, we know that $\phi(x) = \phi(x^\star)$. Therefore, by the other direction of Lemma 1, $x$ is Pareto-optimal. $\square$

### 2.3. PPAD-Hardness of Computing EF+PO

We now turn to our main result:

THEOREM 4. *The problem of finding an EF+PO allocation in a one-sided matching market with linear utilities is PPAD-hard.*

Our proof will reduce the problem of finding an approximate HZ equilibrium to that of finding an EF+PO allocation. The former was shown to be PPAD hard recently.

THEOREM 5 **(Chen et al. [15]).** *For any $c > 0$, the problem of finding an $\epsilon$-approximate HZ equilibrium is PPAD-hard for $\epsilon \leq 1/n^c$.*

There are various reasonable notions of $\epsilon$-approximate HZ equilibria which are polynomially equivalent. We will use the following definition. Vazirani and Yannakakis [37] give a proof that this notion is indeed equivalent to the one used by Chen et al. [15].

DEFINITION 4. An assignment $(x_{ij})_{i\in A, j\in G}$ together with non-negative prices $(p_j)_{j\in G}$ are an $\epsilon$-*approximate HZ equilibrium* if and only if

- each agent $i$ satisfies $\sum_{j\in G} x_{ij} \in [1-\epsilon, 1]$,
- each good $j$ satisfies $\sum_{i\in A} x_{ij} \in [1-\epsilon, 1]$,
- each agent $i$ spends at most 1, i.e. $p \cdot x_i \leq 1$,
- each agent $i$ gets a bundle which is at most $\epsilon$ worse than an optimal bundle, i.e.

$$u_i \cdot x_i \geq \max\left\{ u_i \cdot y \;\middle|\; \sum_{j\in G} y_j = 1, p \cdot y \leq 1 \right\} - \epsilon.$$

Note that Chen et al. assume that all utilities lie in $[0, 1]$ and we will do the same for now. Additionally, we remark that there is no requirement that the bundle of an agent be (approximately) cheapest. This condition is necessary to guarantee some form of approximate Pareto-optimality. However, it is not needed for the hardness proof and removing it only makes the theorem stronger.

**Overview of the Reduction** The general strategy of the reduction consists of the following five steps.

**Step 1:** We will modify the instance to make sure that all EF+PO allocations are approximate HZ equilibria while making sure that HZ equilibria are not perturbed too much.

**Step 2:** Starting with an EF+PO allocation $x$ in the modified instance, and using a version of the Second Welfare Theorem, we will find prices $p$ and budgets $b$ that make $x$ into a competitive equilibrium.

**Step 3:** We will use the envy-freeness of $x$ to prove that agents with almost-equal utilities have almost-equal budgets in a quantifiable sense.

**Step 4:** Then, we will exploit the structure of our modified instance and the linearity of the agents' utilities to prove that *all* budgets are almost equal which makes $(x, p)$ an approximate HZ-equilibrium.

**Step 5:** Finally, we will transform $(x, p)$ to an approximate HZ equilibrium $(\hat{x}, \hat{p})$ in the original instance, finishing the reduction.

Steps 1, 2, and 5 can be carried out in polynomial time as is required in order to get a polynomial reduction from approximate HZ to EF+PO. Steps 3 and 4 are the crux of the correctness proof. If two agents have *equal* utilities and are *non-satiated*, i.e. they are not getting 1 unit of their maximum utility goods, it is not hard to see that their budgets must be equal. Otherwise, the agent with the smaller budget would necessarily envy the budget with the larger budget. Of course, these conditions are very strong and not typically satisfied between two arbitrary agents in an arbitrary instance; this is the reason a modified instance and additional ideas are needed.

**Step 1: Construction of the Modified Instance** Our modified instance is going to ensure that between any two agents $i$ and $i'$, there is a sequence of agents $i = i^{(0)}, \ldots, i^{(l)} = i'$ such that the utilities of $i^{(t)}$ and $i^{(t+1)}$ are *almost* the same for all $t$. If we can show that $i^{(t)}$ and $i^{(t+1)}$ must have almost the same budget for all $t$, then perhaps we can show that $i$ and $i'$ must have almost the same budget. Moreover, we will ensure that no agent can be satiated. In order to carry out this construction without perturbing approximate HZ equilibria too much, we will need to create many copies of identical agents and identical goods.

DEFINITION 5. If two agents have identical utilities for all goods, we say that they are of the *same type*. Likewise, two goods are of the *same type* if all agents have identical utilities for them.

Fix a positive integer $k \in \mathbb{N}^+$ and some $\epsilon > 0$ such that $k$ is divisible by $n$ and $k \geq \frac{n^3}{\epsilon}$. Then we will create a new instance $I' = (A', G', u')$ as follows.

1. For each good in $G$, we add $k$ identical copies of said good to $G'$. Likewise, for each agent in $A$, we add $k$ identical copies of said agent to $A'$. These copies will allow us to add small amounts of new agents and goods without perturbing the HZ equilibria in the instance.

2. Add $k/n$ identical goods for which every agent has utility 2 which is double of what they can get from goods in $G$. For this reason, we call these *awesome goods* and their limited quantity is going to prevent satiation.
3. For each pair $\{i, i'\}$ of distinct agents in $A$, we create a sequence of *interpolating agents*. Order the *types* of goods in some arbitrary way $\{t_1, \ldots, t_n\}$. Now add up to $\frac{1}{\epsilon}$ agents by starting with the utilities of agent $i$ and slowly increasing / decreasing the utility for $t_1$ goods in steps of $\epsilon$ until we reach the utility that $i'$ has for $t_1$ goods. Repeat this process for $t_2, \ldots, t_n$. The final result of this procedure will be at most $\frac{n}{\epsilon}$ additional agents which slowly interpolate between the utilities of $i$ and $i'$, one type of good at a time. See Figure 1.
4. Finally, add dummy agents to $A'$ until $|A'| = |G'|$. These agents have identical utilities for all goods. Note that we added fewer interpolating agents than awesome goods since $k > \frac{n^3}{\epsilon}$.

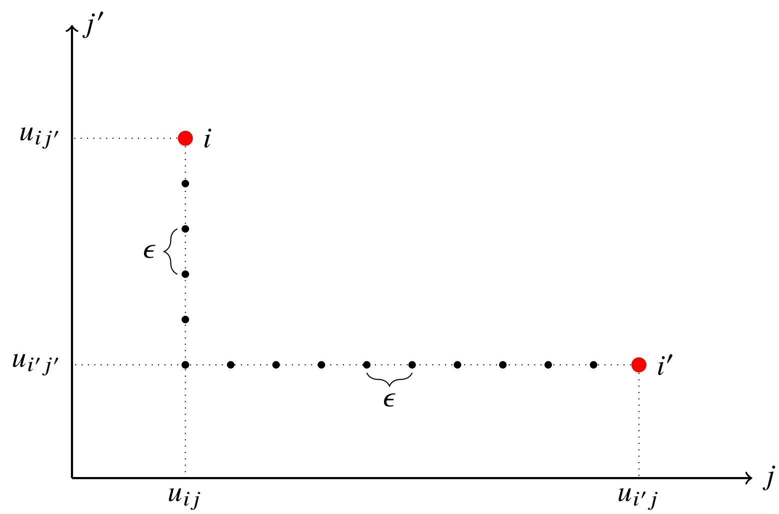


Figure 1. For each pair of agents $i$ and $i'$ (large red dots) we add *interpolating agents* (small black dots) to transition between the utility vector $u_i$ and $u_{i'}$ in small steps. This is done coordinate-wise and this figure depicts an example with only two goods $j$ and $j'$.

Lemma 2. *Let $n' = |A'| = |G'|$ be the number of agents / goods in the modified instance. Then $n' \leq 2kn$.*

*Proof.* We add $kn$ goods through identical copies of the agents in $A$ and $k/n \leq kn$ awesome goods. □

**Step 2: Finding Prices and Budgets** In the following assume that we are given a rational EF+PO allocation $x$ on $I'$ which is encoded with a polynomial number of bits. Our goal will be to construct an approximate HZ solution on $I$. We now carry out Step 2 by finding budgets and prices that make $x$ a competitive equilibrium on $I'$. Recall that by Lemma 1, there exist positive $\alpha_i$ for all $i \in A'$ such that $x$ solves

$$\max \quad \sum_{i \in A'} \alpha_i u_i \cdot x_i$$

$$\begin{aligned} \text{s.t.} \quad & \sum_{j\in G'} x_{ij} = 1 \quad \forall i \in A', \\ & \sum_{i\in A'} x_{ij} = 1 \quad \forall j \in G', \\ & x_{ij} \geq 0 \quad \forall i \in A', j \in G'. \end{aligned}$$

Moreover, we can find such $\alpha_i$ in polynomial time since we obtained them using an LP in the proof of Lemma 1. Consider now an optimal solution $(p,q)$ to the dual.

$$\begin{aligned} \min \quad & \sum_{i\in A'} q_i + \sum_{j\in G'} p_j \\ \text{s.t.} \quad & q_i + p_j \geq \alpha_i u_{ij} \quad \forall i \in A', j \in G' \end{aligned}$$

and define $b_i \coloneqq \alpha_i u_i \cdot x_i - q_i$ to be budget of agent $i$. Note that we may assume that $p, q \geq 0$ since all utilities are non-negative. As shown in Lemma 3, $x$ really is a competitive equilibrium with prices $p$ and budgets $b$.

LEMMA 3. *For every agent $i$, we have that $b_i \geq 0$ and $x_i$ is an optimum solution to*

$$\begin{aligned} \max \quad & u_i \cdot x_i \\ \text{s.t.} \quad & \sum_{j\in G'} x_{ij} \leq 1, \\ & p \cdot x_i \leq b_i, \\ & x_i \geq 0. \end{aligned}$$

*Proof.* First, observe that

$$\sum_{j\in G'} p_j x_{ij} = \sum_{j\in G'} (\alpha_i u_{ij} - q_i) x_{ij} = \alpha_i u_i \cdot x_i - q_i = b_i$$

using complimentary slackness and the fact that $\sum_{j\in G'} x_{ij} = 1$. So $x_i$ is at least feasible and clearly $b_i \geq 0$ since prices are non-negative.

Now take any feasible solution $(y_j)_{j\in G'}$ of the LP. Then

$$\sum_{j\in G'} u_{ij} y_j \leq \sum_{j\in G'} \frac{p_j + q_i}{\alpha_i} y_j \leq \frac{b_i + q_i}{\alpha_i} = u_i \cdot x_i$$

by dual feasibility and the definition of $b_i$. □

**Step 3: Almost Equality of Budgets via Envy-Freeness** Our goal will now be to use envy-freeness in order to show that agents' budgets are approximately equal. First, we need to prove several simple lemmas which ultimately allow us to prove that no agent is satiated (in a quantifiable way).

LEMMA 4. *If $j$ and $j'$ are goods of the same type, then $p_j = p_{j'}$.*

*Proof.* Note that every good is fully matched. Now if the prices were different, then any agent matched to the more expensive good would be violating Lemma 3, leading to a contradiction. □

LEMMA 5. *For any non-dummy agent $i$, $u_i \cdot x_i \leq 1.6$. In particular $i$ is not satiated.*

*Proof.* Since we assume that all the non-awesome goods have utility at most 1, if $i$ were to obtain a total utility of 1.6, it would need to get at least 0.6 units of the utility 2 awesome goods (otherwise $i$'s total utility is less than $0.4 \cdot 1 + 0.6 \cdot 2 = 1.6$). Hence, if $i'$ is any other non-dummy agent, then $u_{i'} \cdot x_i \geq 1.2$ based on just the awesome goods that $i$ has received and the fact that all agents agree which goods are awesome. By the same argument, this means that $i'$ must also have received at least 0.1 units of awesome goods by envy-freeness. Since there are vastly more non-dummy agents than awesome goods, we get a contradiction. □

LEMMA 6. *There exists at least one non-dummy agent $i$ with $b_i > 0$.*

*Proof.* There must be at least one non-dummy agent $i$ who buys a positive fraction of an awesome good. This is because if a non-dummy agent received any amount of an awesome good, this would violate Pareto-optimality since they could swap goods with a non-dummy agent. But since $i$ is not satiated by Lemma 5, the price of said awesome good must be positive and so must the agent's budget. □

In particular, we can rescale all $\alpha$, $p$, $q$, and $b$ so that the maximum budget of any non-dummy agent is exactly 1. In the remainder of this section, we assume that this is the case.

LEMMA 7. *If $i$ and $i'$ are agents of the same type, then $b_i = b_{i'}$.*

*Proof.* Note that by Lemma 5, no agent receives their maximum possible utility. So if $b_i \neq b_{i'}$, assume wlog. that $b_i < b_{i'}$. Then since $i'$ is optimally spending $b_{i'}$ and both agents agree on the utilities of all goods, both agents agree that $i'$ is getting a higher utility bundle than $i$. Thus $i$ would be envious. □

Now that we have established several basic facts about the budgets and bundles of the agents, we will turn to our main objective: show that the budgets are almost equal. As mentioned in our high level plan, we will first show that two agents whose utility vectors are almost equal, must have almost equal budgets. This is done in Lemmas 8 and 9 below.

LEMMA 8. *For any non-dummy agent $i$ we have $\alpha_i \leq 5n^2$.*

*Proof.* Consider an awesome good $j^\star$. By dual feasibility, we know that $p_{j^\star} + q_i \geq \alpha_i u_{ij^\star} = 2\alpha_i$. But on the other hand, note that

$$q_i = \alpha_i u_i \cdot x_i - b_i \leq \alpha_i u_i \cdot x_i \leq 1.6\alpha_i$$

using Lemma 3 and Lemma 5. Comining these inequalities we get $p_{j^\star} \geq 0.4\alpha_i$.

Lastly, we note that the $k/n$ awesome goods can only be sold to the non-dummy agents of which there are at most $2kn$ and each of which has a budget of at most 1 after rescaling. So the price of the awesome goods must be at most $2n^2$ which finishes the proof. □

LEMMA 9. *Let $i, i'$ be two non-dummy agents whose utilities are identical except for the goods of one type where they differ by at most $\epsilon$. Then $|b_i - b_{i'}| \le 5n^2\epsilon$.*

*Proof.* Note that since the $u_i$ and $u_{i'}$ disagree only by epsilon, we have

$$u_i \cdot x_{i'} \ge u_{i'} \cdot x_{i'} - \epsilon \ge u_{i'} \cdot x_i - \epsilon \ge u_i \cdot x_i - 2\epsilon$$

using envy-freeness. In fact, depending on whether $u_i$ or $u_{i'}$ has the higher utility, we can only lose an $\epsilon$ in the first or the last inequality. So we actually get $u_i \cdot x_{i'} \ge u_i \cdot x_i - \epsilon$.

Now we can compute

$$\begin{aligned} b_{i'} &= \sum_{j \in G} x_{i'j} p_j \\ &= \sum_{j \in G} x_{i'j}(\alpha_i u_{ij} - q_i) \\ &= \alpha_i u_i \cdot x_i' - q_i \\ &\ge \alpha_i u_i \cdot x' - \epsilon\alpha_i - q_i \\ &= b_i - \epsilon\alpha_i \end{aligned}$$

and using symmetry and Lemma 8 we conclude $|b_i - b_{i'}| \le \epsilon \max\{\alpha_i, \alpha_{i'}\} \le 5n^2\epsilon$. □

Lemma 9 is enough to show that the difference in budgets between "close" agents tends to zero for an inverse-polynomial $\epsilon$. However, between any two distinct agents $i, i' \in A$, it takes us up to $\frac{n^2}{\epsilon}$ agents to interpolate between them and therefore we cannot give any non-trivial bound on the difference in budget between arbitrary agents. It seems as if we have not won anything!

**Step 4: Bounding the Budget Changes for Interpolating Agents** The key argument that makes our construction work is as follows: we are going to show that along any chain of interpolating agents, the budgets cannot change more than $O(n^2)$ many times due to the linearity of the utilities. Before we prove this in full generality, it is insightful to consider a simpler situation in which agents do not have the matching constraint. Without the matching constraint, the optimal thing to do for any agent is to spend their entire budget on whichever goods have the maximum "bang per buck", i.e. those goods $j$ that maximize $\frac{u_{ij}}{p_j}$.

It is not hard to see that when two agents agree on *which* goods are maximum bang per buck, then their budgets must be equal. Otherwise, the agent with the larger budget would be able to buy more of those goods and thus would be envied by the agent with the smaller budget. When we modify the utility of one good, the set of maximum bang per buck goods can only change twice. See Figure 2.

Unfortunately, once we add in the matching constraint which is crucial to our setting, this simple characterization no longer works. The core issue is that with the matching constraint, the optimal bundles of an agent depend not just on the utilities and prices of the goods but also on the budget of the agent. Since our goal is to show that agents have identical budgets, this easily leads to circular reasoning. The way around this is to instead assume that agents have the same optimal bundles for all *potential* budgets.

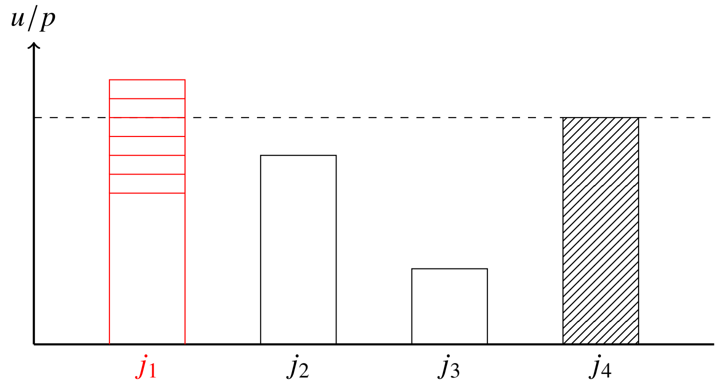


FIGURE 2. Shown is an agent who is interested in goods $j_1$ to $j_4$ which are plotted by their bang per buck. If we change only the utility of good $j_1$ (red) and leave the rest the same, there are only three possible sets of maximum bang per buck goods: $\{j_4\}$, $\{j_1, j_4\}$, and $\{j_1\}$. So along any chain of interpolating agents where we change only the utility for $j_1$ (monotonically), there will be at most two times that the set of maximum bang per buck goods and with it the budget of the agent can change.

DEFINITION 6. For any agent $i \in A'$, define a function $\theta_i(t)$ which maps any $t \geq 0$ to the set of all goods $j \in G$ such that $y_j$ can be positive in an optimum solution to

$$\begin{aligned} \max \quad & u_i \cdot y \\ \text{s.t.} \quad & \sum_{j \in G'} y \leq 1, \\ & p \cdot y \leq t, \\ & y \geq 0. \end{aligned} \tag{1}$$

In other words, $\theta_i(t)$ are the goods which can participate in an optimal bundle for agent $i$ at budget $t$.

LEMMA 10. *Let $i, i' \in A'$ be two agents with $\theta_i = \theta_{i'}$, then $b_i = b_{i'}$.*

*Proof.* Assume otherwise and let $b_i < b_{i'}$ wlog. We will show that $i$ must envy $i'$.

Consider LP (1) with $t = b_{i'}$ which maximizes the utility of agent $i$ but under the higher budget of agent $i'$. We claim that $y = x_{i'}$ is an optimum solution of this LP. To see this, consider the dual as well:

$$\begin{aligned} \min \quad & \mu + \rho b_i \\ \text{s.t.} \quad & \mu + p_j \rho \geq u_{ij}, \\ & \mu, \rho \geq 0. \end{aligned} \tag{2}$$

Now, for any $j$, we know that if $x_{i'j} > 0$, then $j \in \theta_{i'}$ by definition. But since $\theta_{i'} = \theta_i$, this implies that there is some optimal primal solution with $y_j > 0$. By complementary slackness, this implies that $\mu + p_j \rho = u_{ij}$. Therefore, $x_{i'}$ is a feasible solution to the LP which, together with $\mu$ and $\rho$, satisfies the complementary slackness conditions and is therefore optimal.

Finally, since no agent is satiated (Lemma 5), increasing the budget always increases the optimum value of the LP, implying that $u_i \cdot x_i < u_i \cdot x_{i'}$. This contradicts envy-freeness. $\square$

LEMMA 11. *Let $i_1, \ldots, i_m$ be a set of agents such that all agents agree on all utilities except for possibly one type of good. Then $|\{\theta_{i_1}, \ldots, \theta_{i_m}\}| \leq 2n+1$.*

*Proof.* We will give a geometric proof of this fact. First, we will need to understand the behavior of any particular $\theta_i(t)$. We are interested in the goods which can be used in an optimum solution $y$ to (1). By complementary slackness these are the goods for which the corresponding dual constraint is tight in the dual (2).

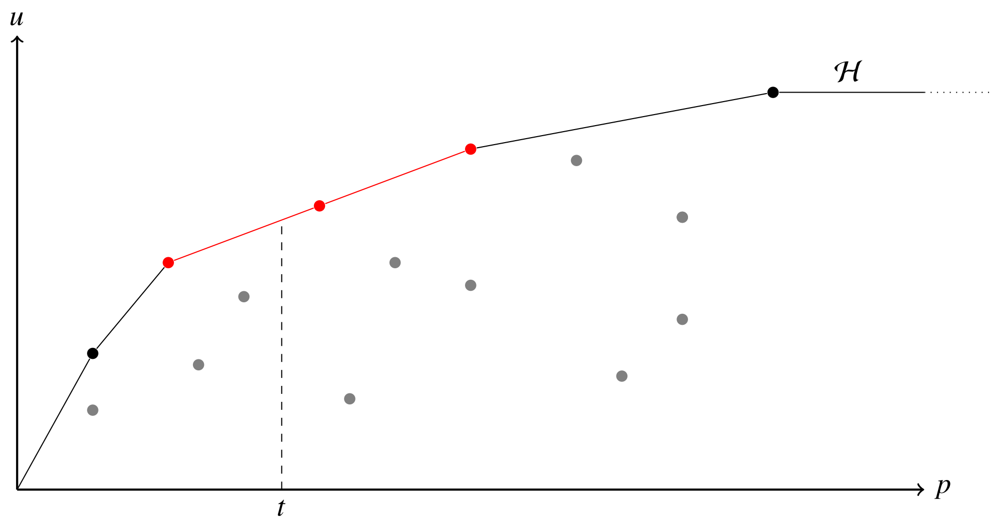


FIGURE 3. Depicted is $\mathcal{H}$ and its relationship to optimal bundles. Each point represents a good or collection of goods with identical price and utility. Gray points are dominated and will never be part of an optimal bundle. Points on $\mathcal{H}$ can be part of an optimal bundle depending on the budget $t$. A typical case is shown in which $\theta_i(t)$ consists of the three red goods that lie on the edge of $\mathcal{H}$ which corresponds to the tight dual constraints at budget $t$.

Now let us interpret this dual geometrically in $\mathbb{R}^2$. The expression $\mu + \rho t$ represents a line in $t$ with non-negative slope. The condition that $\mu + p_j\rho \geq u_{ij}$ means that this line lies above the point $(p_j, u_{ij})$. In other words, the dual objective function for a fixed $t$ is optimized by a line which is as low as possible at $t$ and yet lies above all the points $(p_j, u_{ij})$. This characterizes precisely the upper boundary of the convex hull of the point set

$$\{(0,0)\} \cup \{(p_j, u_{ij}) \mid j \in G'\} \cup \{(\infty, \max_{j \in G'} u_{ij})\}$$

which we will denote by $\mathcal{H}$.

Together with what we already know from complementary slackness, this gives a nice geometric characterization of $\theta_i$. For a given $t$, consider the point $(t, v) \in \mathcal{H}$. If $(t, v)$ is a vertex of the convex hull, i.e. corresponds to $(p_j, u_{ij})$ for some good $j \in G'$, then only this good—or more precisely only goods with identical price and utility—can participate in an optimum bundle. On the other hand, if $(t, v)$ is not a vertex, then it lies on some line $L$ that bounds the convex hull (determined by at least two linearly independent tight dual constraints). $\theta_i(t)$ will then consist of all those goods j such that $(p_j, u_{ij})$ lies on $L$. See Figure 3.

Let us now return to the agents $i_1, \ldots, i_m$ and consider what happens to $\mathcal{H}$ when we shift a single point along the $y$-axis. By the characterization of $\theta$, the only thing we need to know to uniquely determine $\theta$ is which goods lie on $\mathcal{H}$ and out of these which goods are vertices of $\mathcal{H}$. Call this data the *structure* of $\mathcal{H}$.

Let $j$ be the type of good for which the agents have differing utilities. When we remove $j$, we can construct a convex hull $\mathcal{H}'$ on the rest of the goods (corresponding to optimal bundles without type $j$). Finally, observe that the structure of $\mathcal{H}$ only depends on the relationship (below, intersecting, above) which $(p_j, u_{ij})$ has with the at most $n$ lines that bound $\mathcal{H}'$. Since there are only $2n+1$ possible ways in which a point can relate to $n$ lines, this proves the claim. See Figure 4. □

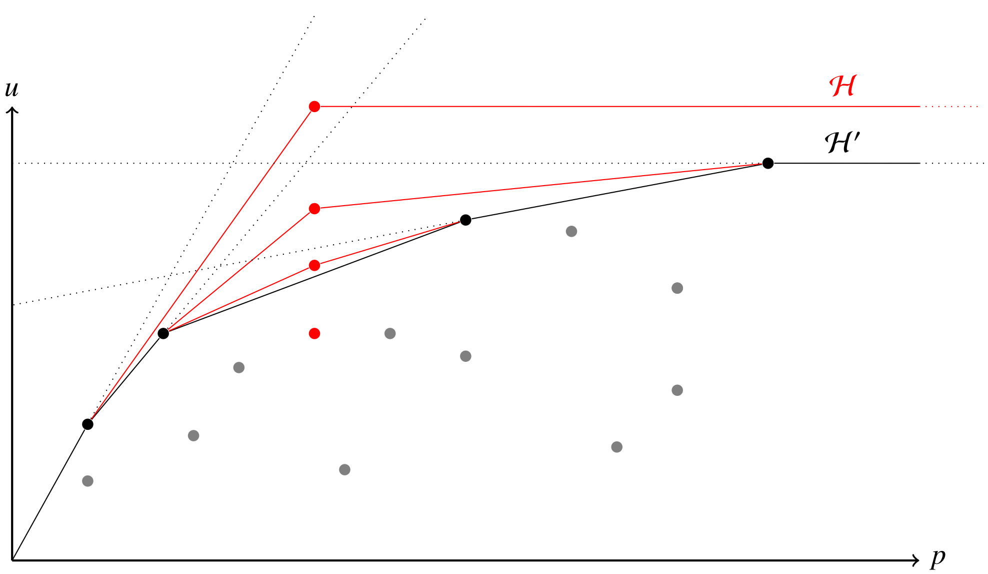


FIGURE 4. Shown are several convex hulls $\mathcal{H}$ (red) as the red good's utility is changed. Note that the structure of $\mathcal{H}$ only changes when we cross one of the bounding lines of $\mathcal{H}$ – the convex hull without the red good.

LEMMA 12. *Let $i, i'$ be two non-dummy agents. Then $|b_i - b_{i'}| \le 10\epsilon n^4$.*

*Proof.* Consider the chain of interpolating agents between $i$ and $i'$. There can be at most $n$ types of goods on which $i$ and $i'$ have different utilities. So we can divide these agents into at most $n$ groups inside of which the agents differ only on one good. By Lemma 11, inside each group there are at most $2n+1$ different $\theta$ functions. By Lemma 10, the budgets of agents who have identical $\theta$ functions must be identical. And so there are at most $2n$ opportunities for $\theta$ to change inside each group, totaling to $2n^2$ changes overall. Each of these changes in $\theta$ corresponds to two agents that differ in their utilities by at most $\epsilon$ on one good, thus Lemma 9 applies and we get $|b_i - b_{i'}| \le 2n^2 \cdot 5\epsilon n^2$. □

**Step 5: Contracting to the Original Instance** To finish the proof, let us construct our approximate HZ equilibrium $(\hat{x}, \hat{p})$ on the original instance by contracting the allocation along the copies of goods and agents. For any $i \in A, j \in G$ let $\hat{x}_{ij}$ be the average over all $x_{i'j'}$ where $i'$ are the $k$ identical copies of $i$ and $j'$ are the $k$ identical copies of $j$ in $I'$. The parts of $x$ going to the dummy agents, interpolating goods, and awesome goods are simply dropped.

THEOREM 6. *If $\epsilon \leq \frac{1}{10n^5}$, then $(\hat{x}, \hat{p})$ is a $\frac{3}{n}$-approximate HZ equilibrium in the original instance $I$.*

*Proof.* First, observe that as there are $k$ copies of each agent $i$ and only $k/n$ awesome goods, we have that $\sum_{j \in G} \hat{x}_{ij} = [1 - \frac{1}{n}, 1]$. Likewise, the total number of interpolating and dummy agents is $k/n$ and there are $k$ copies of each good $j$ so $\sum_{i \in A} \hat{x}_{ij} = [1 - \frac{1}{n}, 1]$. This establishes that $\hat{x}$ is an approximately perfect fractional matching.

Moreover, it is clear that no agent overspends as no non-dummy agent spends more than 1 in $I'$ and we have only removed allocations during the contraction.

Finally, we need to show that no agent is far from their optimum bundle. For that, let $y$ be an optimum solution to

$$\begin{aligned} \max \quad & u_i \cdot y \\ \text{s.t.} \quad & \sum_{j \in G} y = 1, \\ & p \cdot y \leq 1, \\ & y. \geq 0 \end{aligned}$$

By Lemma 12, we know that $b_i \geq 1 - \frac{1}{n}$. And so $u_i \cdot x_i \geq (1 - 1/n) u_i \cdot y$ since we could otherwise scale down $y$ and violate Lemma 3. Note that it is important here that $x_i$ was optimal even among bundles that get *at most* one unit of good.

Lastly, we know that $u_i \cdot \hat{x}_i \geq u_i \cdot x_i - \frac{2}{n}$ since the only thing that was lost when contracting were up to $\frac{1}{n}$ awesome goods as mentioned above. Thus

$$u_i \cdot \hat{x}_i \geq (1 - 1/n) u_i \cdot y - \frac{2}{n} \geq u_i \cdot y - \frac{3}{n}$$

finishing the proof. □

*Proof of Theorem 4.* If we choose $\epsilon = \frac{1}{10n^5}$ and $k = 5n^8$, then the constructed instance has at most $20n^9$ agents by Lemma 2. Given a rational EF+PO allocation with polynomial encoding length, we can construct $(\hat{x}, \hat{p})$ as above in polynomial time and get a $\frac{3}{n}$-approximate HZ equilibrium. By Theorem 5, the latter problem is PPAD-hard. □

Lastly, we remark that Theorem 4 can be slightly strengthed to show hardness of computing *approximately* envy-free and Pareto-optimal solutions with inverse polynomial $\epsilon$. Lemmas 9 and 10 require minor modifications for the proof to go through.

**2.4. 2-EF and 2-IC via Nash Bargaining** Now that we have seen that finding EF+PO allocations is PPAD-hard, this raises the question: what is the best that we can actually do in polynomial time? It turns out that Nash bargaining comes to the rescue here. Nash [31] studied the problem of two or more agents bargaining over a common outcome, for example how they should split certain goods among themselves. He showed that there is a unique point that satisfies certain axioms (namely Pareto-optimality, symmetry, invariance under affine transformations of utilities, and independence of irrelevant alternatives) and moreover that this point is characterized as maximizing the product of the agents' utilities, i.e. the Nash social welfare.

In our case, this means that the Nash bargaining solution is given by the solution to

$$\begin{aligned} \max \quad & \prod_{i \in A} u_i \cdot x_i \\ \text{s.t.} \quad & \sum_{j \in G} x_{ij} = 1 \quad \forall i \in A, \\ & \sum_{i \in A} x_{ij} = 1 \quad \forall j \in G, \\ & x_{ij} \geq 0 \quad \forall i \in A, j \in G. \end{aligned} \tag{3}$$

Since the objective function is log-concave, general purpose convex programming techniques can be used to find approximate solutions to this program which is a stark difference to HZ. For this reason, Hosseini and Vazirani [25] proposed Nash bargaining as an alternate solution concept for cardinal-utility matching markets of various kinds. We strengthen the case for Nash bargaining as an HZ alternative by showing that Nash bargaining points are approximately envy-free and approximately incentive compatible.

DEFINITION 7. An allocation $(x_{ij})_{i \in A, j \in G}$ is *$\alpha$-approximately envy-free* or just $\alpha$-EF if for every $i, i' \in A$ we have $u_i \cdot x_i \geq \frac{1}{\alpha} u_i \cdot x_{i'}$. In other words, no agent envies another agent by more than a factor of $\alpha$.

THEOREM 7. *Let $x$ be an optimum solution to* (3). *Then $x$ is 2-EF.*

*Proof.* Assume otherwise, i.e. that there are agents $i, i' \in A$ such that $u_i \cdot x_{i'} = \alpha u_i \cdot x_i$ and $\alpha > 2$. Then we consider what happens when we swap some $\epsilon$-fraction of the bundle that $i$ gets with the bundle that $i'$ gets. This maintains feasibility.

By doing so, the product of the agents' utilities changes by a factor of

$$\frac{(u_i \cdot x_i(1-\epsilon) + \alpha u_i \cdot x_i \epsilon)(u_{i'} \cdot x_{i'}(1-\epsilon) + u_{i'} \cdot x_i \epsilon)}{(u_i \cdot x_i)(u_{i'} \cdot x_{i'})}.$$

We now evaluate the derivative of this expression wrt. to $\epsilon$ at $\epsilon = 0$ and get

$$\frac{(\alpha - 1)(u_i \cdot x_i)(u_{i'} \cdot x_{i'}) + (u_i \cdot x_i)(u_{i'} \cdot x_i - u_{i'} \cdot x_{i'})}{(u_i \cdot x_i)(u_{i'} \cdot x_{i'})} \geq \alpha - 2.$$

But since $\alpha > 2$, this implies the derivative is positive, i.e. for small enough $\epsilon$ the product of the agents' utilities is increasing. This contradicts the fact that $x$ is an optimum solution to (3). □

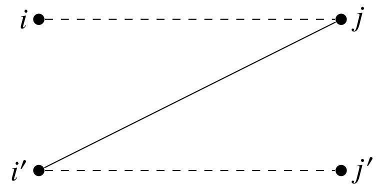


FIGURE 5. Shown is an example instance which demonstrates that 2-EF is tight for Nash bargaining. Dashed edges have utility 1, solid edges have utility 2, and missing edges have utility 0. Clearly both agents prefer $j$ to $j'$. A simple calculation shows that in the Nash bargaining solution, $i$ will get all of $j$ and thus $i'$ will envy $i$ by a factor of 2.

We remark that this bound is tight since Aziz and Brown [5] give an instance in which an agent envies another agent by a factor of 2. See Figure 5.

DEFINITION 8. Consider some mechanism $M$ which maps utility profiles $(u_i)_{i\in A}$ to allocations $(x_i)_{i\in A}$. Then $M$ is called *$\alpha$-incentive compatible* or just $\alpha$-IC if, whenever utilities $u$ and $\hat{u}$ differ only on agent $i$, said agent does not improve by more than a factor of $\alpha$ wrt. to utilities $u$, i.e. $u_i \cdot M(u)_i \geq \frac{1}{\alpha} u_i \cdot M(\hat{u})_i$. This means that no agent stands to gain more than a factor of $\alpha$ by misreporting their utilities.

THEOREM 8. *Any mechanism which maps u to some maximizer of* (3) *is 2-IC.*

*Proof.* The proof of this result is quite similar to the proof of Theorem 7. Consider the original utility profile $u$ and a modified utility profile $\hat{u}$ which differs only on one agent, say agent $l \in A$. Let $x$ be a maximizer of (3) under utilities $u$ and $y$ a maximizer of (3) under utilities $\hat{u}$. Assume that $u_l \cdot y_l = \alpha u_l \cdot x_l$. Our goal is to show that $\alpha \leq 2$.

For small $\epsilon$, we now consider the new allocations $(1-\epsilon)x + \epsilon y$. This allocation cannot increase the product of the utilities $\hat{u}$ compared to $x$ by the maximality of $x$. Thus the derivative wrt. to $\epsilon$ of

$$\prod_{i\in A}(u_i \cdot x_i(1-\epsilon) + u_i \cdot y_i \epsilon)$$

must be non-positive at $\epsilon = 0$. Performing this computation yields

$$\sum_{i\in A} \frac{u_i \cdot y_i - u_i \cdot x_i}{u_i \cdot x_i} \prod_{i'\in A} u_i \cdot x_i \leq 0$$

and therefore

$$\sum_{i\in A\setminus\{l\}} \left(\frac{u_i \cdot y_i}{u_i \cdot x_i} - 1\right) \leq 1 - \frac{u_l y_l}{u_l x_l} = 1 - \alpha.$$

The same argument applies to the allocation $\epsilon x + (1-\epsilon)y$ and the utilities $\hat{u}$ by symmetry, giving the inequality

$$\sum_{i\in A\setminus\{l\}} \left(\frac{\hat{u}_i \cdot x_i}{\hat{u}_i \cdot y_i} - 1\right) \leq 1 - \frac{\hat{u}_l x_l}{\hat{u}_l y_l} \leq 1.$$

Finally note that for all $i \in A \setminus \{l\}$ we have that $u_i = \hat{u}_i$ so after summing up the two inequalities we get:

$$\sum_{i\in A\setminus\{j\}} \left(\frac{u_i \cdot y_i}{u_i \cdot x_i} + \frac{u_i \cdot x_i}{u_i \cdot y_i} - 2\right) \leq 2 - \alpha.$$

By the AM-GM inequality, we know that $\frac{a}{b} + \frac{b}{a} \geq 2$ for all $a, b > 0$ and so this implies that $2 - \alpha \geq 0$ which is precisely what we wanted to show. □

This bound is also tight as shown by the following family of instances. See Figure 6 for an example.

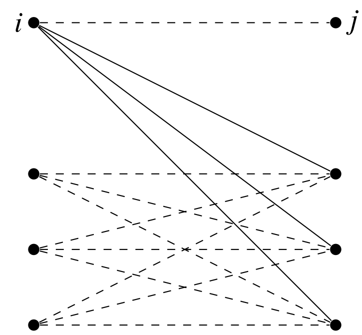


FIGURE 6. Shown is an example instance from the proof of Theorem 9 with $n = 4$. Dashed edges have utility 1, solid edges utility 2, and missing edges have utility 0. Agent $i$ will be fully allocated to good $j$ by Nash bargaining even though they would prefer the "desirable" goods. However, agent $i$ can misrepresent their utilities to look like the other agents therefore get a significant fraction of the desirable goods.

THEOREM 9. *Any mechanism which maps u to some maximizer of* (3) *is not* $(2 - \epsilon)$*-IC for any* $\epsilon > 0$.

*Proof.* Consider the following instance with $n$ agents and $n$ goods. Let there be $n - 1$ desirable goods and one undesirable good. Agent 1 (the agent who will be incentivized to lie) has utility 2 for the desirable goods and utility 1 for the undesirable good whereas all other agents have utility 1 for the desirable goods and utility 0 for the undesirable good. See Figure 6 for $n = 4$.

Let $x$ be the amount that agent 1 is matched to the desirable goods. One can easily see that in optimum solutions to (3), agents with equal utility vectors must have the same overall utility. Otherwise the product of their utilities can be improved. So all other agents must be matched $\frac{n-1-x}{n-1} = 1 - \frac{x}{n-1}$ to the desirable goods. The product of agents' utilities is therefore

$$(2x + (1 - x)) \left(1 - \frac{x}{n-1}\right)^{n-1}$$

and one may check that this is uniquely maximized at $x = 0$. In other words, agent 1 gets nothing from the desirable goods and their utility is 1.

Now agent 1 misreports their utilities as having utility 1 for the desirable good and utility 0 for the undesirable good, i.e. they report the same utilities as all the other agents. But then, by symmetry, this means that agent 1 now gets an equal amount of the desirable goods as all the other agents, i.e. they get $\frac{n-1}{n}$ desirable goods. Thus their actual utility is $2\frac{n-1}{n} + \frac{1}{n}$. Finally, as $n \to \infty$, this implies that any mechanism based on Nash-bargaining cannot be better than 2-IC. □

Panageas et al. [32] give simple, practical algorithms for computing $(1+\epsilon)$-approximate (in the sense that all utilities are within $(1+\epsilon)$ of the Nash bargaining point) Nash bargaining points in $O(\text{poly}(n, 1/\epsilon))$ time and so we get the following corollary.

COROLLARY 1. *There is a $(2+\epsilon)$-EF, PO, $(2+\epsilon)$-IC mechanism for one-sided cardinal-utility matching markets which runs in $O(\text{poly}(n, 1/\epsilon))$ time.*

Finally, note that Hosseini and Vazirani [25] and Panageas et al. [32] also deal with more general settings in which the agents' utilities are not necessarily linear but given by more general (piecewise-linear) concave functions. The above proofs can be adapted to work for non-linear concave utility functions as well, though this is beyond the scope of this paper.

## 3. Two-Sided Matching Markets

A second interesting class of matching markets is that of two-sided markets. A two-sided market is one in which instead of matching goods to agents we match agents to other agents and hence there are preferences from both sides of the market. These markets can be distinguished based on two criteria: whether the underlying graph is bipartite or not and whether the agents' utilities are symmetric or asymmetric.

In a bipartite matching market, we have two sets $A, B$ of agents with $|A| = |B| = n$ and our goal is to match each agent in $A$ to an agent in $B$. Every $i \in A$ has non-negative utilities $u_{ij}$ over $j \in B$ and, likewise, every $j \in B$ has non-negative utilities $w_{ji}$ over $i \in A$. A classic example of this is school choice: students have preferences over schools and schools have preferences over students (e.g. based on test scores). By a slight abuse of notation we use $w_j \cdot x_j$ to mean $\sum_{i \in A} w_{ji} x_{ij}$.

As mentioned in the introduction, there are also non-bipartite matching markets in which we are simply given a set of $2n$ agents and each agent may have utilities over all other agents. In this case one has to be careful with allowing fractional allocations since fractional perfect matchings cannot always be decomposed into integral perfect matchings. Still, these markets are a direct generalization of the bipartite case and so our negative results apply to them as well. In the remainder of this section, we will only consider bipartite two-sided matching markets.

The definitions of Pareto-optimality and envy-freeness extend naturally to this setting.

DEFINITION 9. An allocation in a two-sided matching market is *Pareto-optimal* if there is no way to increase the utility of *any* agent (on either side) without decreasing the utility of another agent (on either side).

DEFINITION 10. An allocation in a two-sided matching market is *envy-free* if no agent prefers another agent's bundle (on their own side) to their own.

Lastly, we will say that a two-sided market has *symmetric* utilities if $u_{ij} = w_{ji}$ for all $i \in A, j \in B$. This is mostly of interest when dealing with $\{0, 1\}$ utilities, in which case a pair of agents is either considered acceptable or not by both parties.

Bogomolnaia and Moulin [9] showed that in the case of a symmetric, bipartite two-sided matching market with $\{0,1\}$ utilities, rational EF+PO allocations exist. Computability is not directly addressed in their paper, though the algorithm by Vazirani and Yannakakis [37] can be adapted for this setting. This result was extended to the non-bipartite case by Roth et al. [34] who proved existence and Li et al. [28] who gave a polynomial time algorithm.

**3.1. Rationality** As was the case for one-sided markets, we can show that if an EF+PO allocation exists, there must be a rational EF+PO allocation. The proofs are essentially identical to those in Section 2.2 so we will not restate them here.

LEMMA 13. *$x^\star \in P_{\mathrm{PM}}$ is Pareto-optimal if and only if there exist positive $(\alpha_i)_{i\in A}$ and $(\beta_j)_{j\in B}$ such that $x^\star$ maximizes $\phi(x) := \sum_{i\in A}\alpha_i u_i\cdot x_i + \sum_{j\in B}\beta_j w_j\cdot x_j$ over all $x\in P_{\mathrm{PM}}$. Moreover, if $x^\star$ is rational, $\alpha$ and $\beta$ can be computed in polynomial time.*

The set of all envy-free allocations is given by the polytope $P_{2\mathrm{EF}}$:

$$P_{2\mathrm{EF}} := \left\{ (x_{ij})_{i\in A, j\in G} \;\middle|\; \begin{array}{lll} \sum_{j\in G} x_{ij} & = 1 & \forall i\in A, \\ \sum_{i\in A} x_{ij} & = 1 & \forall j\in G, \\ u_i\cdot x_i - u_i\cdot x_{i'} & \geq 0 & \forall i,i'\in A, \\ w_j\cdot x_j - w_j\cdot x_{j'} & \geq 0 & \forall j,j'\in B, \\ x_{ij} & \geq 0 & \forall i\in A, j\in G. \end{array} \right\}$$

THEOREM 10. *If an instance of a two-sided bipartite matching market admits an EF+PO allocation, then there is one which is a vertex of $P_{2\mathrm{EF}}$ and is thus rational.*

We will also need the following characterization in Section 3.3. Note that an allocation is weak Pareto-optimal if there is no other allocation that improves on the utility of *every* agent. For a proof see Appendix B.

LEMMA 14. *$x^\star \in P_{\mathrm{PM}}$ is weak Pareto-optimal if and only if there exist non-negative $(\alpha_i)_{i\in A}$ and $(\beta_j)_{j\in B}$ such that $\sum_{i\in A}\alpha_i + \sum_{j\in B}\beta_j > 0$ and $x^\star$ maximizes $\phi(x) := \sum_{i\in A}\alpha_i u_i\cdot x_i + \sum_{j\in B}\beta_j w_j\cdot x_j$ over all $x\in P_{\mathrm{PM}}$.*

**3.2. Non-Existence of EF+PO Solutions** Given that we know that rational EF+PO allocations exist in various matching markets, even two-sided non-bipartite markets with $\{0,1\}$-utilities, an interesting question is whether such allocations exist for any larger classes of instances. We will answer this question in the negative by giving rather limiting counterexamples below.

Per Theorem 10, if an EF+PO allocation exists, then it must be a vertex of the polytope $P_{2\mathrm{EF}}$. Such allocations can often be found heuristically: repeatedly pick random vectors $\alpha\in(0,1]^A$ and $\beta\in(0,1]^B$ and maximize $\sum_{i\in A}\alpha_i u_i\cdot x_i + \sum_{j\in B}\beta_j w_j\cdot x_j$ over $P_{2\mathrm{EF}}$ using an LP solver. This produces a candidate solution $x$ which is Pareto-optimal *among the envy-free allocations*. We can then check whether $x$ is Pareto-optimal *among all solutions* by solving the LP

$$\max \quad \sum_{i\in A} u_i\cdot y_i + \sum_{j\in B} w_j\cdot y_j$$

$$
\begin{aligned}
\text{s.t.} \quad & \sum_{j \in B} y_{ij} = 1 && \forall i \in A, \\
& \sum_{i \in A} y_{ij} = 1 && \forall j \in B, \\
& u_i \cdot y_i \geq u_i \cdot x_i && \forall i \in A, \\
& w_j \cdot y_j \geq w_j \cdot x_j && \forall j \in B, \\
& y_{ij} \geq 0 && \forall i \in A, j \in B.
\end{aligned}
$$

In most instances, this finds an EF+PO allocation relatively quickly. By enumerating small instances we found the examples below which have the fewest positive entries in their utility matrices.

We remark that given the polyhedral nature of the problem, it is possible to design an exact algorithm which can determine in finite time whether an instance has an EF+PO allocation and return it: simply enumerate all vertices of $P_{2\text{EF}}$ and test each one for Pareto-optimality using the LP approach mentioned above. However, this is quite slow in practice due to the exponential number of vertices that $P_{2\text{EF}}$ generally has.

THEOREM 11. *For two-sided matching markets under asymmetric utilities, an EF+PO allocation does not always exist, even for the case of* $\{0,1\}$*-utilities.*

*Proof.* Consider the instance shown in Figure 7a and the Pareto-optimal fractional perfect matching $y$ depicted in that figure. Let $x$ be some allocation in this instance and assume that $x$ is envy-free. We will show that $y$ is strictly Pareto-better than $x$.

First let us show that $x_{24} = \frac{1}{3}$. Note that we must clearly have $x_{24} \geq \frac{1}{3}$ as otherwise $x_{25} > \frac{1}{3}$ or $x_{26} > \frac{1}{3}$ and in those cases agent 4 would envy agent 5 or 6 respectively. On the other hand, assume that $x_{24} = \frac{1}{3} + \epsilon$. Then $u_2 \cdot x_2 \leq \frac{2}{3} - \epsilon$. But then agent 2 envies either agent 1 or agent 3 since among these three, one must get at least $\frac{2}{3}$ of agents 5 and 6. Thus $x_{24} = \frac{1}{3}$ as claimed.

Next, we claim that $x_{14} = \frac{1}{3}$. Again, we clearly have $x_{14} \geq \frac{1}{3}$ as otherwise agent 1 would envy agent 2 or agent 3. But in the other direction, if $x_{14} = \frac{1}{3} + \epsilon$, then $x_{15} + x_{16} = \frac{2}{3} - \epsilon$. By the previous claim, we know that $x_{25} + x_{26} = \frac{2}{3}$ and so $x_{35} + x_{36} = \frac{2}{3} + \epsilon$ which would imply that agent 2 envies agent 3. Thus $x_{14} = \frac{1}{3}$.

Finally, since $x_{24} = \frac{1}{3}$ and $x_{14} = \frac{1}{3}$, we can see that $y$ is Pareto-better than $x$ (regardless of how $x$ assigns the other edges). In particular, $u_1 \cdot y_1 = \frac{2}{3}$ whereas $u_1 \cdot x_1 = \frac{1}{3}$ and $u_i \cdot y_i \geq u_i \cdot x_i$ for all other $i$. □

THEOREM 12. *For two-sided matching markets under symmetric utilities, an EF+PO allocation does not always exist, even in the case of* $\{0,1,2\}$*-utilities.*

*Proof.* Consider the instance shown in Figure 7b together with the depicted Pareto-optimal allocation $y$. Let $x$ be some envy-free allocation. We aim to show that $y$ is Pareto-better than $x$.

First, we can once again see that $x_{24} = \frac{1}{3}$. Note that if $x_{24} < \frac{1}{3}$, then agent 4 will envy agent 5 or agent 6. Vice versa, if $x_{24} > \frac{1}{3}$, then agent 5 or 6 will envy agent 4.


(a) Each arrow represents a utility 1-edge from one side (and utility 0 from the other side).

(b) Dashed edges have utility 1, whereas solid edges have utility 2.

FIGURE 7. Shown are counterexamples for $\{0,1\}$ asymmetric (a) and $\{0,1,2\}$ symmetric utilities. In both cases the edge labels show a Pareto-optimal solution $y$ and edges which are not drawn have utility 0 (assume that $y$ is extended to a fractional perfect matching by filling up with utility 0 edges).

Next, note that $x_{14} = \frac{1}{3}$. Again, we must have $x_{14} \geq \frac{1}{3}$ since otherwise agent 1 would envy agent 2 or agent 3. In the other direction, we cannot have $x_{14} > \frac{1}{3}$ since then agent 2 would envy agent 1 by the previous observation that $x_{24} = \frac{1}{3}$.

Finally, we must have that $x_{25} = x_{26} = \frac{1}{3}$ since otherwise agent 5 would envy agent 6 or vice versa. This determines $x$ on all the edges with positive utility. But now observe that $y$ is Pareto-better than $x$ since $u_1 \cdot y_1 > u_1 \cdot x_1$ and $u_i \cdot y_1 \geq u_i \cdot x_i$ for all other $i$. □

### 3.3. Justified Envy-Freeness

As we have seen in the previous section, in two-sided markets we generally cannot get EF+PO allocations unless we are using symmetric $\{0,1\}$ utilities. Intuitively, the issue is that agents have different entitlements. Consider a market in which an agent $i \in A$ is liked by everyone in $B$ whereas $i' \in A$ is hated by everyone in $B$. It will be difficult to avoid a situation in which $i'$ envies $i$ without sacrificing efficiency.

A way to get around this is to simply formalize this notion of entitlement. In the following, fix some bipartite two-sided matching market with $|A| = |B| = n$ and utilities $u, w$.

DEFINITION 11. In an allocation $x$, agent $i \in A$ has *strong justified envy* towards $i' \in A$ if $w_{ji} \geq w_{ji'}$ for all $j \in B$ and $u_i \cdot x_i < u_i \cdot x_{i'}$. Strong justified envy is defined symmetrically for agents in $B$. An allocation in which there is no strong justified envy is said to be *weak justified envy-free (weak JEF)*.

Weak justified envy-freeness is a reasonable notion in many settings. For example, in school choice a student who scores higher on all relevant tests should not envy a student who scores lower. However, it is somewhat unsatisfying that $i$ is only justified in their envy of $i'$ when *all* agents prefer $i$ to $i'$, even agents that $i$ does not care about. For this reason, we define a stronger notion of justified envy-freeness.

DEFINITION 12. In an allocation $x$, agent $i \in A$ has *justified envy* towards $i' \in A$ if

$$u_i \cdot x_i < \sum_{\substack{j \in B \\ w_{ji} \geq w_{ji'}}} u_{ij} x_{i'j}$$

and likewise for agents in $B$. An allocation in which there is no justified envy is *justified envy-free (JEF)*.

Clearly, strong justified envy implies justified envy and therefore JEF implies weak JEF. We remark that in the case of an integral matching, being JEF is equivalent to being a stable matching. We will show the following.

THEOREM 13. *There always exists a rational allocation which is JEF and weak PO.*

The proof uses a limit argument based on an equilibrium notion introduced by Manjunath [29]. This equilibrium is conceptually similar to an HZ equilibrium with three crucial differences:

- While each agent is endowed with some amount of fake currency, the value of this currency is not normalized. Instead there is a price $p_m$ that determines the "price of money".
- Prices are double-indexed, i.e. an agent in $B$ may have different prices for every agent in $A$.
- Prices can be negative. They effectively represent transfers between the two sides of agents.

We do not need the full generality of the equilibrium notion of Manjunath and will give a slightly simplified definition assuming linear utilities. Each agent $i \in A$ (and likewise for agents in $B$) has some initial endowment $\omega_i > 0$ of "money" and they will receive not just an allocation $(x_{ij})_{j \in B}$ but also some money $m_i \geq 0$. We assume that their utility is given by $u_i \cdot x_i + m_i$. Likewise for the agents in $B$.

DEFINITION 13. An *double-indexed price (DIP) equilibrium* consists of an assignment $(x_{ij})_{i \in A, j \in B}$, money assignments $(m_k)_{k \in A \cup B}$, individualized prices $(p_{ij})_{i \in A, j \in B}$ and $(q_{ji})_{j \in B, i \in A}$, and the price of money $p_m$ satisfying:

1. $x$ is a fractional matching (but not necessarily perfect).
2. The money is redistributed exactly, i.e. $\sum_{i \in A \cup B} \omega_i = \sum_{i \in A \cup B} m_i$.
3. Each agent $i \in A$ (and likewise for agents in $B$) receives an optimal bundle in the sense that $(x_i, m_i)$ maximizes

$$\begin{aligned} \max \quad & u_i \cdot x_i + m_i \\ \text{s.t.} \quad & \sum_{j \in B} x_{ij} \leq 1, \\ & p_i \cdot x_i + p_m m_i \leq p_m \omega_i, \\ & x_{ij} \geq 0 \quad \forall j \in B. \end{aligned}$$

4. $p_{ij} = -q_{ji}$ for all $i \in A, j \in B$.

THEOREM 14 (**Manjunath [29]**). *As long as every agent has a positive endowment of money (i.e. $\omega_i > 0$), a DIP equilibrium always exists.*

THEOREM 15 (**Manjunath [29]**). *The allocation in a DIP equilibrium is Pareto-optimal.*

We require the allocation to be a fractional perfect matching. It is possible to modify the proof of Theorem 14 directly but in order to be self-contained, we will give a short proof which uses Theorem 14 as a black box.

LEMMA 15. *As long as every agent has a positive endowment of money (i.e. $\omega_i > 0$), a DIP equilibrium in which $x$ is a fractional perfect matching always exists.*

*Proof.* For each $k \in \mathbb{N}^+$, consider a modified instance in which every zero utility is replaced by $\frac{1}{k}$. Each of these instances has some DIP equilibrium and clearly in each of these equilibria, the allocation must be a fractional perfect matching since otherwise this would immediately violate Pareto-optimality.

Since the prices are scale invariant, we can rescale them so that the maximum price is bounded by 1. Then both allocations, money assignments, and prices are bounded so by compactness one can find a convergent subsequence of these DIP equilibria. The limiting point is a DIP equilibrium in the original instance with an allocation which is a fractional perfect matching. □

LEMMA 16. *If $\omega_i = \frac{\epsilon}{2n}$ for all $i \in A \cup B$, and $(x, m, p, q, p_m)$ is a DIP equilibrium for these budgets, then for all $i, i' \in A$ (and likewise for agents in $B$) we have*

$$u_i \cdot x_i \geq \sum_{\substack{j \in B \\ w_{ji} \geq w_{ji'}}} u_{ij} x_{i'j} - \epsilon.$$

*Proof.* Let $i, i' \in A$. Consider $j \in B$ with $x_{i'j} > 0$ and $w_{ji} \geq w_{ji'}$. Then we can see that $p_{ij} \leq p_{i'j}$. If this were not the case, then since $q_{ji} = -p_{ij}$ and $q_{ji'} = -p_{i'j}$, we would have $q_{ji} < q_{ji'}$ and thus $j$ could redistribute some of their bundle from $i'$ to $i$ decreasing their total expenditure without decreasing their utility. This is a contradiction to the fact that $j$ gets an optimal bundle since they could then increase $m_j$ to get a strictly better bundle.

This means that

$$\sum_{\substack{j \in B \\ w_{ji} \geq w_{ji'}}} p_{ij} x_{i'j} \leq \sum_{\substack{j \in B \\ w_{ji} \geq w_{ji'}}} p_{i'j} x_{i'j} \leq p_m(\omega_{i'} - m_{i'}) \leq p_m \omega_{i'} = p_m \omega_i$$

where we used that all agents have equal endowments of money in the last equality. But since $i$ maximizes their utility among all bundles which cost at most $p_m \omega_i$, this implies that

$$u_i \cdot x_i + m_i \geq \sum_{\substack{j \in B \\ w_{ji} \geq w_{ji'}}} u_{ij} x_{i'j}.$$

Finally note that $m_i \leq \sum_{k \in A \cup B} m_k = \sum_{k \in A \cup B} \omega_k = \epsilon$ and this finishes the proof. By symmetry the same holds for all pairs of agents in $B$. □

*Proof of Theorem 13.* First, let us show that a JEF and weak PO allocation exists. By Lemma 16, we can pick a sequence $x^{(k)}$ of Pareto-optimal allocations such that

$$u_i \cdot x_i^{(k)} \geq \sum_{\substack{j \in B \\ w_{ji} \geq w_{ji'}}} u_{ij} x_{i'j}^{(k)} - \epsilon_k$$

for all $i, i' \in A$ (likewise for agents in $B$) and $\epsilon_k \to 0$. Since the set of all fractional perfect matchings is compact, we can find a convergent subsequence. Without loss of generality, assume that $x^{(k)}$ converges to some $x^\star$. Clearly $x^\star$ is itself a fractional perfect matching.

The limit point of a sequence of Pareto-optimal allocations is a weak Pareto-optimal allocation. Furthermore, it is easy to see that $x^\star$ is justified envy-free.

Lastly, we can use a similar argument as in the proof of Theorem 3 to show that a rational JEF + weak PO allocation exists as well. Simply pick $\alpha, \beta$ according to Lemma 14 and then find a vertex solution which maximizes $\sum_{i \in A} \alpha_i u_i \cdot x_i + \sum_{j \in B} \beta_j w_j \cdot x_j$ over the polytope of all justified envy-free allocations. □

**3.4. Justified Envy for Nash Bargaining** As shown in Section 2.4, Nash bargaining yields an approximately envy-free and Pareto-optimal allocation in the case of one-sided matching markets. One might reasonably conjecture that it achieves approximately *justified* envy-freeness in the two-sided setting. We give a counterexample below based on a similar example due to Panageas et al. [32] that shows this not to be the case.

THEOREM 16. *There are instances on $n$ vertices such that in the Nash bargaining solution $x$, there are agents $i, i' \in A$ such that all agents in $B$ prefer $i$ to $i'$ and yet $u_i \cdot x_i = \frac{1}{n} u_i \cdot x_{i'}$.*

*Proof.* Our instance has three special agents: $i, i' \in A$ and $j \in B$. All agents in $B \setminus \{j\}$ have utility 1 for $i$ but 0 for everyone else in $A$, including $i'$. Agents $i$ and $i'$ both have utility 1 for agent $j$ and utility 0 for all other agents. The agents in $A \setminus \{i, i'\}$ are dummy agents and have identical utility for all agents in $B$. See Figure 8.

Consider a Nash bargaining solution $x$. The agents in $B \setminus \{j\}$ must all be allocated an equal amount of agent $i$, since otherwise we could increase the product of the agents' utilities by making them equal. Let $y$ be this amount. Then we must have $x_{ij} = 1 - (n-1)y$ and $x_{i'j} = (n-1)y$. Therefore $y$ must maximize

$$(1 - (n-1)y) \cdot (n-1)y \cdot y^{n-1}$$

which implies that $y = \frac{n}{n^2-1}$. Then we can compute that $u_i \cdot x_i = \frac{1}{n+1}$ whereas $u_i \cdot x_{i'} = \frac{n}{n+1}$. □

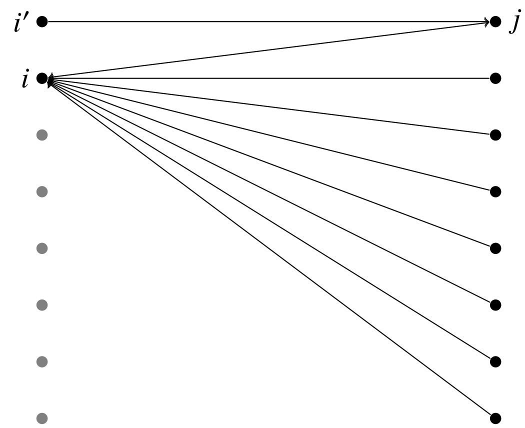


FIGURE 8. Shown is an instance ($n = 8$) with strong justified envy for Nash bargaining. Agents $i$ and $i'$ compete for $j$ but all agents in $B \setminus \{j\}$ want $i$ so $i$ gets only a small fraction of $j$. The gray agents are dummy agents and have identical utilities for all agents in $B$.

## 4. Conclusion

We have resolved the question of whether we can obtain polynomial time mechanisms which give EF+PO lotteries in cardinal-utility matching markets: we can not unless FP = PPAD. However, this leaves several interesting open questions:

- Is there a polynomial time algorithm to find $\alpha$-approximately JEF+PO lotteries in two-sided markets, for any constant $\alpha$?
- Is Nash bargaining the best we can do for one-sided markets or is there a way to compute an $\alpha$-envy-free and Pareto-optimal lottery for $\alpha < 2$ in polynomial time?
- Is there a way to compute an envy-free lottery in polynomial time which satisfies some relaxed notion of Pareto-optimality?

**Acknowledgements** The authors would like to thank Bernard Salanié for mentioning continuum markets to us. We would also like to thank Binghui Peng, Federico Echenique, and Joseph Root for valuable discussions.

## References

[1] Abdulkadiroglu A, Sönmez T (1998) Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66(3):689–701, ISSN 00129682, 14680262, URL `http://www.jstor.org/stable/2998580`.

[2] Alaei S, Jalaly Khalilabadi P, Tardos E (2017) Computing equilibrium in matching markets. *Proceedings of the 2017 ACM Conference on Economics and Computation*, 245–261, EC '17 (New York, NY, USA: Association for Computing Machinery), ISBN 9781450345279, URL `http://dx.doi.org/10.1145/3033274.3085150`.

[3] Ashlagi I, Shi P (2016) Optimal allocation without money: An engineering approach. *Management Science* 62(4):1078–1097, URL `https://EconPapers.repec.org/RePEc:inm:ormnsc:v:62:y:2016:i:4:p:1078-1097`.

[4] Aumann RJ (1964) Markets with a continuum of traders. *Econometrica* 32(1/2):39–50, ISSN 00129682, 14680262, URL `http://www.jstor.org/stable/1913732`.

[5] Aziz H, Brown E (2020) Random assignment under bi-valued utilities: Analyzing hylland-zeckhauser, nash-bargaining, and other rules. *arXiv preprint arXiv:2006.15747* .

[6] Basu S, Pollack R, Roy M (1995) 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* .

[7] Birkhoff G (1946) Tres observaciones sobre el algebra lineal. *Univ. Nac. Tucuman, Ser. A* 5:147–154.

[8] Bogomolnaia A, Moulin H (2001) A new solution to the random assignment problem. *J. Econ. Theory* 100:295–328, URL `https://api.semanticscholar.org/CorpusID:14413046`.

[9] Bogomolnaia A, Moulin H (2004) Random matching under dichotomous preferences. *Econometrica* 72(1):257–279.

[10] Budish E (2011) The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6):1061–1103, ISSN 00223808, 1537534X, URL `http://www.jstor.org/stable/10.1086/664613`.

[11] Budish E, Cachon GP, Kessler JB, Othman A (2017) Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Operations Research* 65(2):314–336, URL `http://dx.doi.org/10.1287/opre.2016.1544`.

[12] Caragiannis I, Filos-Ratsikas A, Kanellopoulos P, Vaish R (2019) Stable fractional matchings. *Proceedings of the 2019 ACM Conference on Economics and Computation*, 21–39, EC '19 (New York, NY, USA: Association for Computing Machinery), ISBN 9781450367929, URL `http://dx.doi.org/10.1145/3328526.3329637`.

[13] Caragiannis I, Hansen KA, Rathi N (2023) On the complexity of pareto-optimal and envy-free lotteries. *arXiv preprint arXiv:2307.12605* .

[14] Caragiannis I, Kurokawa D, Moulin H, Procaccia AD, Shah N, Wang J (2019) The unreasonable fairness of maximum nash welfare. *ACM Trans. Econ. Comput.* 7(3), ISSN 2167-8375, URL `http://dx.doi.org/10.1145/3355902`.

[15] Chen T, Chen X, Peng B, Yannakakis M (2022) Computational hardness of the Hylland-Zeckhauser scheme. *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2253–2268 (SIAM).

[16] Chen X, Deng X (2006) Settling the complexity of two-player nash equilibrium. *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, 261–272, URL `http://dx.doi.org/10.1109/FOCS.2006.69`.

[17] Cole R, Tao Y (2021) On the existence of pareto efficient and envy-free allocations. *Journal of Economic Theory* 193:105207, ISSN 0022-0531, URL `http://dx.doi.org/https://doi.org/10.1016/j.jet.2021.105207`.

[18] Daskalakis C, Goldberg PW, Papadimitriou CH (2009) The complexity of computing a nash equilibrium. *SIAM Journal on Computing* 39(1):195–259, URL `http://dx.doi.org/10.1137/070699652`.

[19] Echenique F, Immorlica N, Vazirani V, eds. (2023) *Online and Matching-Based Market Design* (Cambridge University Press), ISBN 9781108831994.

[20] Echenique F, Miralles A, Zhang J (2021) Constrained pseudo-market equilibrium. *American Economic Review* 111(11):3699–3732, URL `http://dx.doi.org/10.1257/aer.20201769`.

[21] Filos-Ratsikas A, Hansen KA, Høgh K, Hollender A (2023) Ppad-membership for problems with exact rational solutions: A general approach via convex optimization.

[22] Foley DK (1967) Resource allocation and the public sector.(1967) .

[23] Gale D, Shapley LS (1962) College admissions and the stability of marriage. *The American Mathematical Monthly* 69(1):9–15, ISSN 00029890, 19300972, URL `http://www.jstor.org/stable/2312726`.

[24] He Y, Miralles A, Pycia M, Yan J (2018) A pseudo-market approach to allocation with priorities. *American Economic Journal: Microeconomics* 10(3):272–314.

[25] Hosseini M, Vazirani VV (2022) Nash-bargaining-based models for matching markets: One-sided and two-sided; fisher and arrow-debreu. Braverman M, ed., *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA*, volume 215 of *LIPIcs*, 86:1–86:20 (Schloss Dagstuhl - Leibniz-Zentrum für Informatik), URL `http://dx.doi.org/10.4230/LIPICS.ITCS.2022.86`.

[26] Hylland A, Zeckhauser R (1979) The efficient allocation of individuals to positions. *Journal of Political economy* 87(2):293–314.

[27] Immorlica N, Lucier B, Weyl G, Mollner J (2017) Approximate efficiency in matching markets. R Devanur N, Lu P, eds., *Web and Internet Economics*, 252–265 (Cham: Springer International Publishing), ISBN 978-3-319-71924-5.

[28] Li J, Liu Y, Huang L, Tang P (2014) Egalitarian pairwise kidney exchange: fast algorithms via linear programming and parametric flow. *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, 445–452.

[29] Manjunath V (2016) Fractional matching markets. *Games and Economic Behavior* 100:321–336, ISSN 0899-8256, URL http://dx.doi.org/https://doi.org/10.1016/j.geb.2016.10.006.

[30] Miralles A, Pycia M (2016) Large vs. continuum assignment economies: Efficiency and envy-freeness.

[31] Nash J (1950) The bargaining problem. *Econometrica* 18(2):155–162.

[32] Panageas I, Tröbst T, Vazirani VV (2021) Combinatorial algorithms for matching markets via nash bargaining: One-sided, two-sided and non-bipartite. *arXiv preprint arXiv:2106.02024* .

[33] Papadimitriou CH (1994) On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences* 48(3):498–532, ISSN 0022-0000, URL http://dx.doi.org/https://doi.org/10.1016/S0022-0000(05)80063-7.

[34] Roth AE, Sönmez T, Ünver MU (2005) Pairwise kidney exchange. *Journal of Economic theory* 125(2):151–188.

[35] Varian H (1974) Equity, envy, and efficiency. *Journal of Economic Theory* 9(1):63–91, URL https://EconPapers.repec.org/RePEc:eee:jetheo:v:9:y:1974:i:1:p:63-91.

[36] Vazirani VV (2012) The notion of a rational convex program, and an algorithm for the arrow-debreu nash bargaining game. *Journal of the ACM (JACM)* 59(2):1–36.

[37] Vazirani VV, Yannakakis M (2021) Computational complexity of the Hylland-Zeckhauser mechanism for one-sided matching markets. *Innovations in Theoretical Computer Science*.

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

[39] Zadeh L (1963) Optimality and non-scalar-valued performance criteria. *IEEE Transactions on Automatic Control* 8(1):59–60, URL http://dx.doi.org/10.1109/TAC.1963.1105511.

[40] Zhou L (1990) On a conjecture by gale about one-sided matching problems. *Journal of Economic Theory* 52(1):123–135.

[41] Zhou L (1992) Strictly fair allocations in large exchange economies. *Journal of Economic Theory* 57(1):158–175.

## Appendix A: Characterization of Pareto-Optimality

*Proof of Lemma 1.* Clearly if $x^\star$ maximizes $\phi(x)$, then it is a Pareto-optimal allocation since any Pareto-better allocation $x$ would satisfy $\phi(x) > \phi(x^\star)$ since $\alpha$ is strictly positive.

For the other direction, note that by Pareto-optimality, $x^\star$ is a maximizer of the linear program:

$$\begin{aligned} \max \quad & \sum_{i \in A} u_i \cdot x_i \\ \text{s.t.} \quad & u_i \cdot x_i \geq u_i \cdot x_i^\star && \forall i \in A, \\ & \sum_{j \in G} x_{ij} = 1 && \forall i \in A, \\ & \sum_{i \in A} x_{ij} = 1 && \forall j \in G, \\ & x_{ij} \geq 0 && \forall i \in A, j \in G. \end{aligned}$$

Consider a solution $(a, q, p)$ to the dual program:

$$\min \quad \sum_{i \in A} a_i u_i \cdot x_i^\star + \sum_{i \in A} q_i + \sum_{j \in G} p_j \tag{4a}$$

$$\text{s.t.} \quad a_i u_{ij} + q_i + p_j \geq u_{ij} \quad \forall i \in A, j \in G, \tag{4b}$$

$$a_i \leq 0 \quad \forall i \in A. \tag{4c}$$

Then by strong duality

$$\sum_{i \in A} u_i \cdot x_i^\star = \sum_{i \in A} a_i u_i \cdot x_i^\star + \sum_{i \in A} q_i + \sum_{j \in G} p_j. \tag{5}$$

Define $\alpha_i := 1 - a_i$. Then clearly $\alpha_i > 0$ for all $i$ since $a_i \leq 0$. Now we want to show that $x^\star$ is a maximizer of

$$\begin{aligned} \max \quad & \sum_{i \in A} \alpha_i u_i \cdot x_i \\ \text{s.t.} \quad & \sum_{i \in G} x_{ij} = 1 && \forall i \in A, \\ & \sum_{j \in A} x_{ij} = 1 && \forall j \in G, \\ & x_{ij} \geq 0 && \forall i \in A, j \in G. \end{aligned}$$

But this follows immediately from the fact that $(q, p)$ is an optimal dual solution to this LP: (4b) implies feasibility and (5) implies optimality. $\square$

## Appendix B: Characterization of Weak Pareto-Optimality

*Proof of Lemma 14.* The proof is quite similar to the proof of Lemma 1. Clearly if $x^\star$ maximizes $\phi(x)$, then it is a Pareto-optimal allocation since any strong Pareto-better allocation $x$ would satisfy $\phi(x) > \phi(x^\star)$ since at least one $\alpha_i$ or $\beta_j$ is positive.

For the other direction, note that by weak Pareto-optimality, $(x^\star, 0)$ is a maximizer of the linear program:

$$\begin{aligned} \max \quad & t \\ \text{s.t.} \quad & u_i \cdot x_i - t \geq u_i \cdot x_i^\star && \forall i \in A, \\ & w_j \cdot x_j - t \geq w_j \cdot x_j^\star && \forall j \in B, \\ & \sum_{j \in B} x_{ij} = 1 && \forall i \in A, \\ & \sum_{i \in A} x_{ij} = 1 && \forall j \in B, \\ & x_{ij} \geq 0 && \forall i \in A, j \in B, \\ & t \geq 0. \end{aligned}$$

Consider an optimum solution $(\alpha, \beta, p, q)$ to the dual program:

$$\begin{aligned} \min \quad & \sum_{i \in A} q_i + \sum_{j \in B} p_j - \sum_{i \in A} \alpha_i u_i \cdot x_i - \sum_{j \ B} \beta_j w_j \cdot x_j \\ \text{s.t.} \quad & q_i + p_j - \alpha_i u_{ij} - \beta_j w_{ji} \geq 0 \quad \forall i \in A, j \in B, \\ & \sum_{i \in A} \alpha_i + \sum_{j \in B} \beta_j \geq 1, \\ & \alpha_i \geq 0 \quad \forall i \in A, \\ & \beta_i \geq 0 \quad \forall j \in B \end{aligned}$$

Then we may see that $x^\star$ is an optimum solution to

$$\begin{aligned} \max \quad & \sum_{i \in A} \alpha_i u_i \cdot x_i + \sum_{j \in B} \beta_j w_j \cdot x_j \\ \text{s.t.} \quad & \sum_{i \in B} x_{ij} = 1 \quad \forall i \in A, \\ & \sum_{j \in A} x_{ij} = 1 \quad \forall j \in B, \\ & x_{ij} \geq 0 \quad \forall i \in A, j \in B. \end{aligned}$$

since $(p, q)$ gives an optimum dual solution. $\square$

# Quest for EF+PO: Review Response

Thorben Tröbst, Vijay V. Vazirani

May 18, 2025

We thank both reviewers for their feedback and reviews of our paper. We have addressed all the comments in the reviews; see responses to the individual comments below.

## Reviewer 1

1. Improved the confusing period in definitions 2 / 3.

2. Added a new section that gives some background on PPAD (Section 2.1).

3. Corrected the indices in the first LP of Step 2.

4. The 0.6 comes from the fact that if agent $i$ were to get less than 0.6 awesome goods, its total utility is less than $0.4 \cdot 1 + 0.6 \cdot 2 = 1.6$ since the non-awesome goods only give utility 1 at most. The exact number here is arbitrary, the point is that no agent can get fully satiated, i.e. get a utility 2 bundle since this would necessarily cause envy as there are very few utility 2 goods.

   We have expanded the proof of the lemma to make it more clear.

5. This includes the new agents, i.e. $A'$, since we want to apply this along the interpolating agents.

6. Yes, looks like we lost a factor of 2 there. We've updated Lemma 12 (and as a consequence Theorem 6 and the proof of Theorem 4) with the corrected constant factors.

7. By "two-sided" in this context, we are referring to the fact that each edge now has preferences associated with it from both sides. In a one-sided market, there is a clear distinction into agents and goods: agents have preferences over goods but goods have no preferences over agents. But in a two-sided market, we are matching agents to other agents and each has preferences over the others.

   We have added a clarifying remark to the beginning of section 3.

8. Fixed typo

9. Fixed typo

## Reviewer 2

1. We have added a section on PPAD (Section 2.1).

2. We call an equilibrium rational if it can be entirely expressed with rational numbers. We have added an explanation to the beginning of Section 2.2 on rationality.