UNIVERSITY OF CALIFORNIA,
IRVINE

Cardinal-Utility Matching Markets and Online Matching

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Computer Science

by

Thorben Alexander Tröbst

Dissertation Committee:
Distinguished Professor Vijay V. Vazirani, Chair
Distinguished Professor David Eppstein
Associate Professor Milena Mihail
Assistant Professor Ioannis Panageas

2024

TABLE OF CONTENTS

LIST OF FIGURES

Page

LIST OF ALGORITHMS

Page

ACKNOWLEDGMENTS

I would like to thank Vijay for advising me and for inviting me to come to UCI in the first place. His continuous excitement and optimism about research has helped me many times to recover motivation when projects did not work out. He usually believed more in me and my work than I did and provided me with many opportunities that I am grateful for.

Aside from Vijay, I would like to thank my co-authors, including those of papers that did not appear in this thesis: Jannis Blauth, Jugal Garg, Stephan Held, Zhiyi Huang, Milena Mihail, Dirk Müller, Ioannis Panageas, Niklas Schlomberg, Vera Traub, Rajan Udwani, and Jens Vygen. It was a great experience to collaborate with you.

On that note, I am grateful towards the various organizations which have provided funding for my research: UCI and the Donald Bren School of ICS, the NSF (via grants CCF-1815901 and CCF-2230414), and the Hausdorff Center for Mathematics.

I would also like to thank those people who have contributed ideas or with whom I had fruitful discussions about research. It will be impossible to recall all of you, but let me at least mention Asaf Ferber for teaching me a lot about probability theory and probabilistic combinatorics; Federico Echenique, Joseph Root, and Bernard Salanié for fruitful discussions about matching markets; Binghui Peng for insights on PPAD-hardness; and Martin Bullinger for discussions on online matching and dynamics for Nash bargaining.

Finally, I would like to express my gratitude towards my friends and family who contributed in a different but no less meaningful way to my PhD. Foremost, I thank my parents Ines and Holger for supporting me and my interests for so many years as well as for igniting my interest in computers in the first place. On that note, a special thanks also goes to my grandmother Helga for nurturing my interest in science at a young age.

I would like to thank the CS theory group as a whole for every board game night, every evening at Eureka, every hike, and just in general for being an awesome group of people. Special thanks go to Evrim and Hadi for the great times we have had.

This thesis is based on a collection of previously published papers:

- Chapter 2 is based on "One-Sided Matching Markets with Endowments: Equilibria and Algorithms" [62] which was published in JAAMAS and co-authored with Jugal Garg and Vijay V. Vazirani. Copyright has been retained and Springer Nature has granted permission of reuse within this thesis.
- Chapter 3 is based on "Cardinal-Utility Matching Markets: The Quest for Envy-Freeness, Pareto-Optimality, and Efficient Computability" [116] which was accepted to EC 2024 and was co-authored with Vijay V. Vazirani. It will appear as an abstract in the conference proceedings. The full version is currently in journal review. All rights have been retained.

- Chapter 4 is based on "Time-Efficient Nash-Bargaining-Based Matching Market Models" [108] which has been accepted to WINE 2024 and was co-authored with Ioannis Panageas and Vijay V. Vazirani. It will appear in the conference proceedings. Copyright has been retained and Springer Nature has granted permission of reuse within this thesis.
- Chapter 6 is based on "Online Matching with High Probability" [101] which appeared in SAGT 2024 and was co-authored with Milena Mihail. Copyright has been retained and Springer Nature has granted permission of reuse within this thesis.
- Chapter 7 is based on "Almost Tight Bounds for Online Hypergraph Matching" [117] which was published in ORL and was co-authored with Rajan Udwani. Copyright belongs to Elsevier and they have granted permission of reuse within this thesis.

Vijay V. Vazirani directed and supervised research which forms the basis for the dissertation.

VITA

Thorben Alexander Tröbst

EDUCATION

Doctor of Philosophy in Computer Science — **2024**
University of California, Irvine — *Irvine, California, US*

Master of Science in Computer Science — **2022**
University of California, Irvine — *Irvine, California, US*

Master of Science in Mathematics — **2019**
University of Bonn — *Bonn, Germany*

Bachelor of Science in Mathematics — **2017**
University of Bonn — *Bonn, Germany*

WORK EXPERIENCE

Graduate Research Assistant — **2019–2024**
University of California, Irvine — *Irvine, California, US*

Teaching Assistant — **2019–2024**
University of California, Irvine — *Irvine, California, US*

Software Engineering Intern — **2023**
Google — *Sunnyvale, California, US*

Software Engineering Intern — **2022**
Google — *Mountain View, California, US*

Undergraduate / Graduate Research Assistant — **2016–2019**
Research Institute for Discrete Mathematics — *Bonn, Germany*

REFEREED JOURNAL PUBLICATIONS

One-Sided Matching Markets with Endowments: Equilibria and Algorithms **2024**
J. Garg, T. Tröbst, V. V. Vazirani
Autonomous Agents and Multi-Agent Systems

Vehicle Routing with Time-Dependent Travel Times: Theory, Practice, and Benchmarks **2024**
J. Blauth, S. Held, D. Müller, N. Schlomberg, V. Traub, T. Tröbst, and J. Vygen
Discrete Optimization

Almost Tight Bounds for Online Hypergraph Matching **2024**
T. Tröbst and R. Udwani
Operations Research Letters

A Real Polynomial for Bipartite Graph Minimum Weight Perfect Matchings **2023**
T. Tröbst and V. V. Vazirani
Information Processing Letters

A Fast $(2 + 2/7)$**-Approximation for Capacitated Cycle Covering** **2021**
V. Traub and T. Tröbst
Mathematical Programming

REFEREED CONFERENCE PUBLICATIONS

Time-Efficient Algorithms for Nash-Bargaining-Based Matching Market Models **2024**
I. Panageas, T. Tröbst, and V. V. Vazirani
Conference on Web and Internet Economics (WINE)

Online Matching with High Probability **2024**
M. Mihail and T. Tröbst
International Symposium on Algorithmic Game Theory (SAGT)

Cardinal-Utility Matching Markets: The Quest for Envy-Freeness, Pareto-Optimality, and Efficient Computability **2024**
T. Tröbst and V. V. Vazirani
ACM Conference on Economics and Computation (EC)

A Nash-Bargaining-Based Mechanism for One-Sided Matching Markets and Dichotomous Utilities **2023**
J. Garg, T. Tröbst, and V. V. Vazirani
International Conference on Autonomous Agents and Multiagent Systems (AAMAS)

An Arrow-Debreu Extension of the Hylland-Zeckhauser Scheme: Equilibrium Existence and Algorithms **2022**
J. Garg, T. Tröbst, and V. V. Vazirani
International Conference on Autonomous Agents and Multiagent Systems (AAMAS)

A Fast $(2 + 2/7)$**-Approximation for Capacitated Cycle Covering** **2020**
V. Traub and T. Tröbst
International Conference on Integer Programming and Combinatorial Optimization (IPCO)

BOOK CHAPTERS

Applications of Online Matching **2023**
T. Tröbst and Z. Huang
Online and Matching-Based Market Design

ABSTRACT OF THE DISSERTATION

Cardinal-Utility Matching Markets and Online Matching

By

Thorben Alexander Tröbst

Doctor of Philosophy in Computer Science

University of California, Irvine, 2024

Distinguished Professor Vijay V. Vazirani, Chair

In this dissertation, we study several aspects of matching-based market design, i.e. the problem of designing mechanisms that can match agents to goods (or to other agents) while satisfying certain desirable properties such as *fairness* or *efficiency*. Common applications include matching students to schools, doctors to hospitals, riders to drivers, search queries to online advertisements, etc.

Part I deals with cardinal-utility matching markets where agents report non-negative utilities for the goods they are interested in. The Hylland Zeckhauser mechanism satisfies many of the desirable properties that we are interested in, however it has been shown to be PPAD-complete. We design alternative, *polynomial time* mechanisms for several settings. In particular, this dissertation provides substantial support for *Nash-bargaining-based mechanisms* as a computationally tractable alternative to HZ.

Part II deals with online matching where some of the agents or goods arive *online* and must be matched immediately and irrevocably. We provide a novel analysis of the classic RANKING algorithm which shows that it achieves its competitive ratio of $1 - 1/e$ with *high probability* as opposed to just in expectation. In addition, we introduce the study of online matching on hypergraphs where we give asymptotically almost tight upper and lower bounds.

Part I

Cardinal-Utility Matching Markets

Chapter 1

Background

A matching market, broadly speaking, refers to any setting in which agents must be matched to either goods or other agents under certain preferences. The key point here is that each agent gets at most one good. In most cases, matching markets do not involve any form of payments but rather the goal is simply to assign resources efficiently and in a fair way. Classic examples of matching markets include matching students to schools, doctors to hospitals, kidney donors to patients, etc. In 2012, the Nobel prize in economics was awarded to Alvin Roth and Lloyd Shapley for their seminal work on matching markets. For a detailed overview of the area, we refer to [40].

In this part of the dissertation, we will consider the classic setting of a centralized, offline market. We are given a set A of *agents* and a set G of *goods*. For the sake of simplicity, we will assume that $|A| = |G| = n$ with the goal henceforth being to find *perfect* matchings between agents and goods.

Matching markets can be classified according to three dichotomies: cardinal-utility vs ordinal-preferences, two-sided vs one-sided, and endowments vs no endowments. In a *cardinal-utility* matching market, each agent $i \in A$ has a non-negative vector $(u_{ij})_{j \in G}$ of utilities over the goods, whereas in an *ordinal-preferences* matching market, each agent $i \in A$ has some total order $\preceq_i \subseteq G^2$ over the goods. For reasons explained in Section 1.1, this dissertation focuses entirely on cardinal utilities.

In the remainder of this chapter, we will present the most significant preliminary information for the study of cardinal-utility matching markets (CMMs), focusing on the one-sided and no-endowments case. We will make an argument as to why CMMs are of particular interest in Section 1.1, followed by background information on market equilibria (Section 1.2), the celebrated Hylland Zeckhauser mechanism (Section 1.3), and the Nash bargaining game (Section 1.4).

In Chapter 2, we will introduce the setting with endowments while focusing primarily on bivalued / dichotomous utilities where we can give polynomial time algorithms. We will study the natural extension of HZ to such a setting with endowments: the Arrow-Debreu Hylland-Zeckhauser (ADHZ) model. In this setting, we will give a new counterexample that shows that ADHZ equilibria do not exist, even in very restricted conditions. On the other hand, we will introduce the ϵ-approximate ADHZ model and show its existence. Our main result is a polynomial time algorithm to compute such ϵ-approximate ADHZ equilibria under dichotomous utilities. Finally, we will briefly return to the standard, non-endowment model and give a rational convex program for HZ with dichotomous utilities.

In Chapter 3, we will return to the standard setting without endowments, but instead broaden our scope to general (i.e. not necessarily bivalued) utilities. It is known that finding an approximate HZ equilibrium is PPAD-hard [28] and so we will investigate whether there might be an alternative mechanism which is able to attain at least attain some of the desirable properties of HZ in polynomial time, especially envy-freeness (EF) and Pareto-optimality (PO). We will show that finding any EF+PO allocation is already PPAD hard but that Nash bargaining gives 2-approximately EF and PO allocations. Moreover, Nash bargaining also turns out to be 2-approximately incentive compatible. Lastly, we will show that EF+PO allocations generally do not exist in two-sided markets and propose an alternative notion of fairness which does.

In Chapter 4, we will focus on the computation of Nash bargaining solutions. Given the results from the previous chapter, Nash bargaining is a promising alternative for HZ under general utilities. In this chapter, we will give two relatively simple algorithms which are efficient both in theory and in practice in order to find approximate Nash bargaining solutions for a variety of matching market models. The algorithms are based on the multiplicative weights update technique and on conditional gradient descent. We will even consider non-bipartite models.

1.1 Preliminaries

In a matching market, our goal is to match a set A of agents to a set G of goods with $|A| = |G| = n$. As mentioned previously, we assume that agents have some form of preferences over the goods and want to find matchings with certain "desirable" properties. A key question is *how* agents should report their preferences and here there are two distinct approaches: ordinal preferences and cardinal utilities.

Under ordinal preferences, agents report preference lists over the goods, i.e. each agent $i \in A$ has some total order $\preceq_i \subseteq G^2$. Sometimes, ties between goods are also allowed. This model is particularly popular in the matching market literature, going back to the seminal work by Gale and Shapley on the Stable Matching Problem and the classic Deferred Acceptance algorithm [56]. Many prominent mechanisms have been developed for matching markets with ordinal preferences such as Random Priority [103], Probabilistic Serial [17], and Top Trading Cycles [113].

In this dissertation, we will instead consider cardinal utilities. Here, each agent $i \in A$ has a non-negative vector $(u_{ij})_{j\in G}$ of rational utilities over the goods. Cardinal utilities allow the agents to express *how much* they prefer one good over another. For this reason, mechanisms that use cardinal utilities have the potential to be much more *efficient* than those which use ordinal preferences. Our standard measure of efficiency is that of *Pareto-optimality*.

Definition 1.1. *Let $(x_{ij})_{i\in A, j\in G}$ and $(y_{ij})_{i\in A, j\in G}$ be fractional perfect matchings (FPMs). We call y* Pareto-better *than x, if $u_i \cdot y_i \geq u_i \cdot x_i$ for all $i \in A$ and $u_i \cdot y_i > u_i \cdot x_i$ for some $i \in A$ where $u_i \cdot y_i = \sum_{j\in G} u_{ij} y_{ij}$ (and likewise for $u_i \cdot x_i$) is the standard inner product. We call x* Pareto-optimal, *if there is no FPM which is Pareto-better than x.*

Note that Definition 1.1 is phrased in terms of *fractional* perfect matchings even though we are often interested in *integral* perfect matchings. The reason for this is an observation which goes back to Hylland and Zeckhauser [75]. Finding integral matchings which are *fair* in any meaningful sense of the word is generally impossible as standard notions from fair division (e.g. EFX, EF1, maximin share etc.) are only relevant when agents can get multiple goods. For example, if there is a good which all agents agree is more desirable than the other goods, whoever does not get that good will be envious. We have no way of compensating them by giving them more of the less desirable goods.

Instead, the idea is to design mechanisms which find desirable *lotteries* over integral matchings thus using the power of randomness to provide fairness. A lottery over integral matchings is simply a convex combination of integral matchings from which we can sample in polynomial time. By a well-known result of Birkhoff and von Neumann, this is exactly equivalent to being a fractional perfect matching (in bipartite graphs).

Theorem 1.1 (Birkhoff [15], von Neumann [124])**.** *Let $(x_{ij})_{i\in A, j\in G}$ be a fractional perfect matching. Then there exist $y^{(1)}, \dots, y^{(k)}$ integral perfect matchings with $k \in O(n^2)$ and corresponding $\lambda_1, \dots, \lambda_k \geq 0$ with $\lambda_1 + \cdots + \lambda_k = 1$ such that $x = \sum_{i=1}^{k} \lambda_i y^{(i)}$. Moreover, both y and λ can be found in polynomial time.*

For this reason, we will generally only consider fractional perfect matchings in this part of this dissertation. If an integral matching is desired (as is often the case), then one should use the algorithm from the Birkhoff-von-Neumann theorem to run a lottery after the fact. For an example of such a lottery, consider the permit lotteries ran by `recreation.gov` to allocate permits to visit various crowded attractions.[1] Visitors can list various dates on which they wish to visit the attraction and a lottery is used to determine who gets a permit and on which day. This can easily be modelled as a one-sided matching market, though we do not know which algorithm they use to run the lottery.

Note that properties which hold *for the lottery* are usually called *ex ante* properties whereas properties that hold *for the sampled integral matching* are called *ex post*. If not otherwise mentioned, we will always consider ex-ante properties by default. In this nomenclature, Definition 1.1 actually refers to ex-ante Pareto-optimality. An integral perfect matching would be ex-post Pareto-optimal if there was no other *integral* perfect matching which was Pareto-better than it. Note that any integral perfect matching which is sampled from an ex-ante Pareto-optimal lottery is automatically ex-post Pareto-optimal.

Finding a matching, integral or fractional, which is merely Pareto-optimal is not particularly challenging; any max-weight matching would satisfy this condition. The goal is to combine this property with other useful properties such as that of envy-freeness defined below.

Definition 1.2. *Let* $(x_{ij})_{i\in A, j\in G}$ *be a fractional perfect matching. If* $i, i' \in A$ *are agents such that* $u_i \cdot x_i < u_i \cdot x_{i'}$, *i.e. agent* i *strictly prefers the bundle of agent* i' *to their own, then we say that* i envies i'. *If no agent envies any other agent,* x *is* envy-free.

Again, note that we have technically defined *ex-ante* envy-freeness here. As already mentioned, ex-post envy-freeness is typically not a useful notion for matching markets as it generally does not exist and neither do its common generalizations.

[1] `https://www.recreation.gov/lottery/available`

The last major property that we desire in the context of matching markets is *incentive compatibility*. Informally, we do not want agents to be able to lie about their utilities in order to gain an advantage. Note that this is not a property of matchings, fractional or otherwise, and rather a property of *mechanisms*, i.e. the assignment of utilities to matchings. A formal definition is given below.

Definition 1.3. *A* mechanism *is a function M which maps utility profiles $(u_{ij})_{i \in A, j \in G}$ to fractional perfect matchings. We call M* dominant-strategy incentive-compatible *(DISC) or simply* incentive-compatible, *if for all utility profiles u and agents $i \in A$, there does not exist any $(\hat{u}_{ij})_{j \in G}$ such that agent i prefers $M(u_{-i}, \hat{u}_i)$ to $M(u)$ wrt. utilities u. In other words, no agent can misreport utilities so that they gain wrt. their original, true utilities. Here we use the notation $(u_{-i}, \hat{u}_i)$ to refer to the utility profile which agrees with u on all agents except for i and agrees with $\hat{u}$ on agent i.*

Unfortunately, due to a result by Zhou [128], there is no mechanism which is Pareto-optimal, envy-free and incentive-compatible. Indeed, envy-freeness can be relaxed to symmetry (equal agents must get equal utility) in this impossibility result. Achieving incentive-compatibility with cardinal utilities turns out to be a rather challenging task; see Abebe et al. [1] for a weak result in this direction. This brings us back to the question: why focus on cardinal utilities in the first place?

Theorem 1.2 (Immorlica et al. [76])**.** *There are instances with k types of agents and goods such that all agents have the same preference order over the goods but the uniform allocation (i.e. every agent gets an equal fraction of each good) is highly inefficient: a different allocation improves the utility of* every *agent by a factor of $\frac{k}{2}$.*

Proof. Let $i \in \{1, \ldots, k\}$, then there will be 2^k identical agents and 2^{i-1} identical goods labeled with number i in our instance. An agent with label i has utility 1 for goods with labels $1, \ldots, i$ and utility 0 for the remaining goods. Assume that all agents break ties in the same way, if necessary by slightly perturbing all utilities to enforce the same ordering of goods for all agents. Note that there are fewer goods than agents as described so far, so add sufficiently many dummy goods to the instance which all agents have utility 0 for. Again, perturb the utilities as necessary to make sure all agents agree on their ordinal preferences to satisfy the condition of the theorem.

Now let us consider how much utility each agent gets under the uniform allocation. Since there are $k2^k$ agents, each agent gets a $\frac{1}{k2^k}$ fraction of every good. So the utility of an agent with label i is

$$\frac{1}{k2^k} \sum_{j=1}^{i} 2^{i-1} = \frac{1}{k2^k}(2^i - 1).$$

On the other hand, consider the allocation x in which we equally distribute the 2^{i-1} goods with label i to the 2^k agents with label i for every i. Afterwards, we can make sure that the allocation is a fractional perfect matching by filling up with dummy goods. Under this allocation, the utility of an agent with label i is $\frac{2^{i-1}}{2^k}$. Therefore x improves the utility of every agent by a factor of at least $\frac{k}{2}$. □

Note that since the agents in the instances from Theorem 1.2 are ordinally indistinguishable, the uniform allocation is the really the only reasonable allocation for a fair mechanism to chose if we only have access to the ordinal preferences. Moreover, the instances involved do not get too large: for n agents and goods, we can potentially improve each agent by a $\theta(\log n)$ factor. This is a highly limiting result for the efficiency of ordinal mechanisms in matching markets and our main motivation for studying cardinal utilities.

1.2 Market Equilibria

A particularly powerful tool for finding fair and efficient allocations of any kind, matchings and otherwise, is the theory of market equilibria. In this section we will give a short overview of markets with money and their equilibria.

The simplest kind of market with money is the linear Fisher market. Here we are given a set A of agents and a set G of goods as well as non-negative, linear utilities $(u_{ij})_{i\in A, j\in G}$. The goods are perfectly divisible and for each good there is exactly one unit of it in the market. Each agent i has some budget $b_i > 0$ which they wish to spend on the goods. However, there is no matching constraint: agents can get more than one total unit of goods as long as they can afford it with their budget.

Note that the restriction that we have exactly one unit of each good is without loss of generality: we can simply rescale the units to make sure this is the case. We also have not specified what kind of numbers are used for the utilities and budgets. Economists and mathematicians would allow real numbers. Since the main focus of this dissertation is on computational matters, we will generally assume that all numbers involved are rational.

Definition 1.4. *Given A, G, u, and b, a* Fisher market equilibrium *consists of an* allocation $(x_{ij})_{i\in A, j\in G}$ *and non-negative* prices $(p_j)_{j\in G}$ *such that*

1. *No agent overspends, i.e. $p \cdot x_i \leq b_i$ for each $i \in A$.*
2. *Each agent i gets an* optimal bundle*, i.e. x_i maximizes $u_i \cdot x_i$ under the constraint that $p \cdot x_i \leq b_i$.*
3. *The market clears, i.e. $\sum_{i\in A} x_{ij} = 1$ for all $j \in G$.*

We remark that there are several common variants of Definition 1.4. A common requirement is that each agent spends their entire budget. However, with linear utilities this is guaranteed by Definition 1.4, provided that each agent has strictly positive utility for at least one good. Similarly, the market clearing condition is often stated in the form: if a good $j \in G$ has $p_j > 0$, then $\sum_{i \in A} x_{ij} = 1$. However, under linear utilities it is clear that a good cannot actually have price 0 as long as there is at least one agent who has positive utility for it since otherwise said agent would demand an infinite quantity of the good. Throughout this thesis we will typically make these common assumptions implicitly.

Theorem 1.3. *Assuming that each agent has positive utility for at least one good and each good has positive utility for at least one agent, a Fisher market equilibrium always exists.*

The Fisher market model itself goes back to Irving Fisher's 1891 PhD thesis. However, Fisher did not prove that a market equilibrium always exists. Arrow and Debreu [5] were the first to prove the existence of market equilibria, though their model is more general than that of Fisher as we will see later in this section. The proof is non-constructive and relies on Kakutani's fixed-point theorem which is an extension of Brouwer's fixed-point theorem that was discovered 50 years after Fisher's work.

Market equilibria have remarkable properties. First, they exactly characterize the set of Pareto-optimal allocations as per the fundamental theorems of welfare economics. Note that Pareto-optimality is defined analogously as in Definition 1.1; we just allow all allocations now, not just fractional perfect matchings.

Theorem 1.4 (First Welfare Theorem). *Let (x, p) be a Fisher market equilibrium. Then x is Pareto-optimal.*

Theorem 1.5 (Second Welfare Theorem). *Let x be a Pareto-optimal allocation in a linear Fisher market. Then there exist non-negative prices $(p_j)_{j \in G}$ and positive budgets $(b_i)_{i \in A}$ such that (x, p) is a Fisher market equilibrium under budgets b.*

Another simple but remarkable fact is that market equilibria are envy-free, provided that the budget of each agent is the same.[2] Again, the definition of envy-freeness in this setting is the same as Definition 1.2.

Theorem 1.6 (Varian [118]). *Let (x, p) be a Fisher market equilibrium for uniform budgets, i.e. all agents have the same budget. Then x is envy-free.*

Proof. Consider any two agents $i, i' \in A$. i cannot envy i' because $b_i = b_{i'}$ and so agent i *could* have bought the same bundle as i'. Since i gets an *optimal* bundle, we therefore know that $u_i \cdot x_i \geq u_i \cdot x_{i'}$. □

We remark that allocating goods via such a market equilibrium process is not an incentive compatible mechanism; it is possible for agents to manipulate prices by misreporting their utilities. However, the potential for such manipulations becomes smaller as the market becomes larger. This was formalized by Roberts and Postlewaite [110], though the details go beyond the scope of this introduction.

Instead, we will now focus on computational matters. As previously mentioned, the existence of Fisher market equilibria follows from a non-constructive fixed-point argument. Can we *compute* a Fisher market equilibrium? It turns out that the Fisher market equilibria are captured by the so-called Eisenberg-Gale (EG) convex program shown below.

$$
\begin{aligned}
\max \quad & \sum_{i \in A} b_i \log(u_i \cdot x_i) \\
\text{s.t.} \quad & \sum_{i \in A} x_{ij} \leq 1 \quad \forall j \in G, \\
& x_{ij} \geq 0 \quad \forall i \in A, j \in G
\end{aligned}
\tag{1.1}
$$

Theorem 1.7 (Eisenberg [46], Gale [54]). *Let x be an optimal solution to* (1.1). *Then there exist non-negative prices $(p_j)_{j \in G}$ (dual variables) such that (x, p) is a Fisher market equilibrium.*

[2]With differing budgets, they satisfy a kind of weighted envy-freeness.

Theorem 1.7 implies that we can simply solve the EG convex program, for example with the ellipsoid method, in order to get both the allocation (primal) and its prices (dual) making up a Fisher market equilibrium. Of course we cannot in general hope to find an *exact* solution to the convex program, but this method does provide an *approximate* equilibrium in polynomial time for some suitable definition of approximate. But we can do even better: Devanur et al. [36] give a *combinatorial* algorithm, the DPSV algorithm, which finds an *exact* equilibrium in polynomial time. This was later improved by Orlin [106] to find equilibria in strongly polynomial time.

Note that the existence of this algorithm implies in particular that there is always a rational equilibrium which is a rather remarkable fact on its own since this implies that the EG convex program always has at least one rational solution. Convex programs with this rather rare property are called rational convex programs (RCPs) and have a tendency to show up in the study of market equilibria [119].

The linear Fisher market is a rather simple market model and many extensions have been considered in the economics literature. It is common to study non-linear utilities of various types and generally, many of the results from this section go through as long as the agents have concave utility functions. Another common extension is the exchange model by Arrow and Debreu [5].

In the Arrow Debreu market model, we are still given agents A and goods G with non-negative, linear utilities $(u_{ij})_{i\in A, j\in G}$. The difference to the Fisher model is that agents no longer come to the market with money. Instead, the each agent has an initial endowment of goods, i.e. there is an initial allocation $(e_{ij})_{i\in A, j\in G}$ with $\sum_{i\in A} e_{ij} = 1$ for all goods $j \in G$. The idea is that the market will determine non-negative prices $(p_j)_{j\in G}$ and then each agent will sell their initial endowment at market prices and buy an optimal bundle with that money.

Definition 1.5. *Given A, G, u, and e, an* Arrow Debreu market equilibrium *consists of an* allocation $(x_{ij})_{i\in A, j\in G}$ *and non-negative* prices $(p_j)_{j\in G}$ *such that*

1. *No agent overspends, i.e.* $p \cdot x_i \leq p \cdot e_i$ *for each* $i \in A$.

2. *Each agent i gets an* optimal bundle, *i.e.* x_i *maximizes* $u_i \cdot x_i$ *under the constraint that* $p \cdot x_i \leq p \cdot e_i$.

3. *The market clears, i.e.* $\sum_{i\in A} x_{ij} = 1$ *for all* $j \in G$.

The Arrow Debreu market is a generalization of the Fisher market. Given a Fisher market with budgets $(b_i)_{i\in A}$, we can transform this into endowments $e_{ij} := \frac{b_i}{\sum_{i'\in A} b_{i'}}$. Since the prices in the Arrow Debreu market are clearly scale invariant, we can rescale them to make sure that $\sum_{j\in G} p_j = \sum_{i\in A} b_i$ in which case these endowments guarantee that every agent i gets exactly b_i units of money for his endowment, thus recovering the Fisher market setting.

As for the linear Fisher market, there are several slight variations of Definition 1.5. We will typically assume that each agent has positive utility for at least one good and positive endowment of at least one good as well as that each good has positive utility for at least one agent. However, even with these assumptions, equilibria do not always exist and a more technical condition is required.

Definition 1.6. *For a given linear Arrow Debreu market, we can define the* demand graph *which is a directed graph on the set of agents which contains an edge (i, i') whenever there exists a good $j \in G$ such that $e_{i'j} > 0$ and $u_{ij} > 0$, i.e. agent i likes at least one good that initially belongs to i'.*

Theorem 1.8 (Gale [55]). *Consider the demand graph H of a linear Arrow Debreu market. Assume that for every agent i which is a singleton strongly connected component of H, there is at least one good $j \in G$ with $u_{ij} > 0$ and $e_{ij} > 0$. Moreover, assume that each agent has positive utility for at least one good and positive endowment of at least one good as well as that each good has positive utility for at least one agent. Then an Arrow Debreu market equilibrium exists.*

We remark that the original existence result by Arrow and Debreu makes a substantially stronger assumption, namely that each agent has a positive endowment of *every* good [5].

On the computational side, Jain [79] gave a rational convex program whose solutions are Arrow Debreu market equilibria, thus allowing us to find an equilibrium in polynomial time, for example via the ellipsoid method. Duan and Mehlhorn [37] were able to find a combinatorial algorithm using a similar approach as the DPSV algorithm for the Fisher market. Garg and Végh [63] further improved this approach to find an equilibrium in strongly polynomial time.

1.3 The Hylland Zeckhauser Mechanism

Let us now return to matching markets. We are given sets A of agents and G of goods with $|A| = |G| = n$ as well as non-negative linear utilities $(u_{ij})_{i\in A, j\in G}$. The goal is to find a perfect matching of agents and goods with desirable properties. This was precisely the problem studied by Hylland and Zeckhauser in their seminal 1979 work [75].

The first observation that they made (and which we already covered in Section 1.1) is the fact that we can relax the problem to consider *fractional* perfect matchings instead. Each good becomes one unit of perfectly divisible probability shares and each agent must be allocated exactly one unit of probability shares in total. Once probability shares have been allocated, we can run a corresponding lottery via Theorem 1.1.

The main idea is that we can then turn the problem into a market: we simply give each agent an equal amount of income and find a competitive equilibrium. The resulting allocation will be fair and efficient. This alone is not too surprising as Varian [118] had made the same observation for the tradtional fair division setting without the matching constraint. However, it is not obvious that a competitive equilibrium even exists in this setting.

Definition 1.7. *Given A, G, u, and b, an* HZ market equilibrium *consists of an* allocation $(x_{ij})_{i\in A, j\in G}$ *and non-negative* prices $(p_j)_{j\in G}$ *such that*

1. *x is a fractional perfect matching.*
2. *No agent overspends, i.e. $p \cdot x_i \le b_i$ for each $i \in A$.*
3. *Each agent i gets an* optimal bundle, *i.e. x_i maximizes $u_i \cdot x_i$ under the constraint that $p \cdot x_i \le b_i$ and $\sum_{j\in G} x_{ij} = 1$.*
4. *Each agent i gets a* cheapest bundle, *i.e. x_i minimizes $p \cdot x_i$ among all bundles with utility at least $u_i \cdot x_i$ and $\sum_{j\in G} x_{ij} = 1$.*

Almost always, the budgets b will be identical in order to guarantee envy-freeness and unless otherwise stated, one should assume that an HZ equilibrium has equal budgets. In some settings it can make sense to give preferential treatment to certain agents in which case non-uniform budgets can be used.

Definition 1.7 is quite similar to the equilibrium notions for the Fisher market and Arrow Debreu market from Section 1.2. There are two main differences:

1. Instead of the normal market clearing condition, we now additionally require that each agent gets exactly one unit of goods. In line with this extra constraint, optimal bundles must only be optimal among those bundles which contain exactly one unit of goods.
2. Agents must not only get optimal bundles, but *cheapest* optimal bundles. This technical constraint is required in order to guarantee Pareto-optimality despite the fact that agents' utilities can be *satiated*. Unlike in a Fisher market, agents in an HZ market may not always spend their entire budget since they are limited to one unit of goods.

As was the case for the Fisher and Arrow Debreu markets, an HZ equilibrium has many desirable properties.

Theorem 1.9 (Hylland, Zeckhauser [75])**.** *Let (x, p) be an HZ equilibrium. Then x is Pareto-optimal.*

Proof. Assume otherwise, i.e. there is some FPM y which is Pareto-better than x. We will now consider how much each agent would spend under y. Note that since each agent in an HZ equilibrium gets a *cheapest* bundle and since $u_i \cdot y_i \geq u_i \cdot x_i$ (by definition of Pareto-optimality), we know that $p \cdot y_i \geq p \cdot x_i$.

But we also know that there is at least one agent i such that $u_i \cdot y_i > u_i \cdot x_i$. For that agent we must have $p \cdot y_i > p \cdot x_i$ since otherwise x_i is not an optimal bundle for i. We can conclude that $\sum_{i \in A} p \cdot y_i > \sum_{i \in A} p \cdot x_i$.

This is a contradiction since a simple calculation shows that

$$\sum_{i \in A} p \cdot x_i = \sum_{i \in A} \sum_{j \in G} p_j x_{ij} = \sum_{j \in G} p_j \sum_{i \in A} x_{ij} = \sum_{j \in G} p_j$$

by virtue of x being a fractional perfect matching. Likewise, $\sum_{i \in A} p \cdot y_i = \sum_{j \in G} p_j$ since y is a fractional perfect matching as well. □

Theorem 1.10 (Hylland, Zeckhauser [75])**.** *Let (x, p) be an HZ equilibrium with equal budgets. Then x is envy-free.*

Proof. The proof is identical to the proof of Theorem 1.6: agents cannot envy another because they can afford the other agents' bundles. □

Theorem 1.11 (He et al. [69])**.** *The HZ mechanism is incentive compatible in the large.*

As mentioned in Section 1.1, there is no mechanism which is Pareto-optimal, envy-free, and incentive compatible [128]. The notion of "incentive compatible in the large" essentially states that the mechanism becomes incentive compatible as long as there are many copies of agents and goods in the market to prevent individuals from affecting the prices too much. The practical relevance of this result is debatable. However, in practice it is also not obvious how to manipulate the mechanism without knowing the other agents' utilities which may be an additional protection against manipulation. Of course, these properties are only interesting because HZ equilibria are guaranteed to exist which is precisely what Hylland and Zeckhauser showed.

Theorem 1.12 (Hylland, Zeckhauser [75]). *An HZ market equilibrium always exists.*

The proof of Theorem 1.12 is similar to other existence proofs for market equilibria and relies critically on Kakutani's fixed-point theorem. It therefore provides no way of actually computing an HZ equilibrium. Unlike the Fisher and Arrow Debreu markets, there is no rational convex program for HZ market equilibria. This was shown by Vazirani and Yannakakis [122] who in fact provided an instance with rational utilities in which there is a unique *irrational* equilibrium.

The irrationality of HZ equilibria also rules out the possibility that finding an HZ equilibrium could be in the complexity class PPAD which contains various other market equilibrium computation problems. On the other hand, Vazirani and Yannakakis [122] showed that finding an HZ equilibrium is in FIXP and that the problem of finding an *ϵ-approximate*[3] HZ equilibrium is indeed in PPAD. Chen et al. [28] recently showed the corresponding hardness: finding an ϵ-approximate HZ equilibrium is PPAD-hard and therefore likely to be computationally intractible.

On the positive side, Vazirani and Yannakakis [122] were able to give a polynomial time algorithm for HZ equilibria under *bivalued* utilities, i.e. each agent has utilities $u_{ij} \in \{a_i, b_i\}$ where a_i, b_i are two non-negative rational numbers that may be different for each agent.

[3]We will cover approximate equilibrium notions in more detail in the following chapters.

Note that the result of Chen et al. [28] established hardness for four-valued utilities, thus leaving the three-valued case open for now. It is also possible to find equilibria in polynomial time assuming that the total number of agent or good *types* is constant. Alaei et al. [3] provide such an algorithm based on the algebraic cell decomposition technique, though their running times increase at least on the order of k^{5k^2} where k is the number of types. Garg et al. [61] improve on this with an algorithm that scales with k^k. Neither algorithm is practical on realistic instances.

1.4 Nash Bargaining

Aside from market equilibria, another strategy for the fair and efficient allocation of resources is based on Nash bargaining. The idea, which goes back to Nash [105], is to consider a more general situation in which n agents are bargaining about some shared outcome. There is a set of potential outcomes (specified purely in terms of the utilities that the agents accrue) as well as a disagreement point which represents the utility that agents get if bargaining fails. A formal definition is given below.

Definition 1.8. *An n-person* Nash bargaining game *consists of a convex, compact set $\mathcal{U} \subseteq \mathbb{R}^n_{\geq 0}$ of n-dimensional utility vectors and a* disagreement point $c \in \mathcal{U}$. *The game is* feasible *if there exists at least one $u \in \mathcal{U}$ with $u_i > c_i$ for all $i \in [n]$.*

Nash asked the question: which point $u^\star \in \mathcal{U}$, called a Nash bargaining point, do the n players reach if they bargain in a rational way with each other? He came up with a series of axioms that $u^\star$ should satisfy:

1. $u^\star$ should be Pareto-optimal.

2. $u^\star$ should be symmetric, i.e. for any permutation $\pi : [n] \to [n]$, the point $(u^\star_{\pi(i)})_{i\in[n]}$ should be a Nash bargaining point for the Nash bargaining problem in which agents have been permuted by π.

3. $u^\star$ should be invariant under affine transformations, i.e. if $f : \mathbb{R}^n \to \mathbb{R}^n$ is some affine transformation then $f(u^\star)$ is a Nash bargaining point for the Nash bargaining game $(f(\mathcal{U}), f(c))$.

4. $u^\star$ should be independent of irrelevant alternatives, i.e. if there is some convex, compact $\mathcal{U}' \subseteq \mathcal{U}$ and $u^\star \in \mathcal{U}'$ then $u^\star$ is also a Nash bargaining point for $(\mathcal{U}', c)$.

These axioms, particularly the fourth one, are somewhat contentious. Still, under certain assumptions on the bargaining process, one can show that rational agents really do reach a point satisfying these conditions [14]. It turns out that there is always a unique Nash bargaining point and it has a simple characterization.

Theorem 1.13 (Nash [105])**.** *There always exists a unique point $u^\star$ satisfying axioms 1–4 in any feasible Nash bargaining game, namely the unique maximizer of $\prod_{i\in[n]}(u_i - c_i)$ over all $u \in \mathcal{U}$.*

The product of agents' utilities is often called the *Nash social welfare* and has the very useful property of being log-concave. This is why Nash bargaining is of particular interest to us: after taking the log of the objective, the unique Nash bargaining point can be computed via the following convex program.

$$\begin{aligned} \max \quad & \sum_{i\in[n]} \log(u_i - c_i) \\ \text{s.t.} \quad & u \in \mathcal{U}. \end{aligned} \tag{1.2}$$

Since (1.2) is a convex program, we can find arbitrarily good approximations to the Nash bargaining solution in polynomial time, for example using the ellipsoid method [66]. This "easy" computability is in stark contrast to the computation of market equilibria which often turns out to be intractable. In Chapter 4, we will see that (1.2) can be solved efficiently with very practical algorithms.

Hosseini and Vazirani [70] realized that Nash bargaining might make for a good alternative mechanism to HZ for matching markets. This approach is based on prior work by Vazirani for the Arrow Debreu market [119]. For matching markets, the set $\mathcal{U}$ is the set of all utility vectors that arise from fractional matchings and the disagreement point c is commonly the zero vector, though other disagreement points can be chosen as we will discuss in Chapters 2 and 4. This yields the convex program:

$$\begin{aligned} \max \quad & \sum_{i \in A} \log(u_i \cdot x_i) \\ \text{s.t.} \quad & \sum_{j \in G} x_{ij} \leq 1 \quad \forall i \in A, \\ & \sum_{i \in A} x_{ij} \leq 1 \quad \forall j \in G, \\ & x_{ij} \geq 0 \quad \forall i \in A, j \in G. \end{aligned} \tag{1.3}$$

We remark that the matching constraint ≤ 1 can easily be replaced by a perfect matching constraint $= 1$ since we assume that all utilities are non-negative. The astute reader might notice that (1.3) looks almost the same as (1.1), the Eisenberg Gale convex program for the Fisher market. Indeed, without the matching constraint, Nash bargaining is equivalent to Fisher market equilibria with uniform budgets.

In that sense, Nash bargaining can be seen as a very natural alternative to HZ. Whereas HZ equilibria generalize Fisher equilibria by retaining the same pricing approach in the most natural way, Nash bargaining instead generalizes the Eisenberg Gale convex program in the most natural way. A core claim of this dissertation is that Nash bargaining is an *equally valid* generalization and various Nash-bargaining-based mechanisms will feature prominently in the next three chapters.

Chapter 2

Markets with Endowments

2.1 Introduction

In Chapter 1 we introduced the basic setting of this thesis in which goods should be matched fairly to agents. However, fairness does not always mean that all agents should be treated equally by the mechanism. In this chapter we will investigate a model in which the agents come to the market with initial endowments of goods. This chapter is based on the paper "One-Sided Matching Markets with Endowments: Equilibria and Algorithms" which was joint work with Jugal Garg and Vijay V. Vazirani [62].

The two fundamental market models are the Fisher market and Arrow-Debreu market as discussed in Section 1.2. Recall that in a Fisher market, agents come to the market with money and the goods are assumed to be external. On the other hand, in the Arrow-Debreu market, the agents themselves bring the goods to the market and effectively exchange the goods among themselves. The Arrow-Debreu market is a direct generalization of the Fisher market as discussed in Section 1.2.

The Hylland Zeckhauser mechanism [75] is the defacto standard mechanism for one-sided matching markets *without* any endowments. Fundamentally, HZ can be seen as an extension of the Fisher market model to matchings and it achieves its remarkable fairness and efficiency guarantees due to the power of pricing. The goal in this chapter is to investigate an analogous generalization of the Arrow-Debreu model to matchings.

The Arrow-Debreu setting of one-sided matching markets has several natural applications beyond the Fisher setting, such as allocating students to rooms in a dorm for the next academic year, assuming their current room is their initial endowment. Similarly, school choice, when a student's initial endowment is a seat in a school which they already have. For this reason, the issue of obtaining such an extension of the HZ mechanism, called *ADHZ* in this chapter, was actually already considered by Hylland and Zeckhauser. However, this culminated in an example which inherently does not admit an equilibrium [75].

A recourse to this was given by Echenique, Miralles, and Zhang [42] via their notion of an *α-slack Walrasian equilibrium*. This is a hybrid between the Fisher and Arrow-Debreu settings. Agents have initial endowments of goods and, for some fixed $\alpha \in (0, 1]$, the budget of each agent, for given prices of goods, is $\alpha + (1 - \alpha) \cdot m$, where m is the value for their initial endowment; the agent spends this budget to obtain an optimal bundle of goods. Via a non-trivial proof using Kakutani's fixed point theorem, Echenique et al. proved that an α-slack equilibrium always exists.

2.1.1 Our Contributions

As mentioned, Hylland and Zeckhauser already observed that an ADHZ equilibrium need not always exist. However, even for the regular Arrow-Debreu market, equilibria do not always exist. For an Arrow-Debreu market under linear utilities, Gale [55] defined a *demand graph:* a directed graph on agents with an edge (i, j) if agent i likes a good that agent j

has in her initial endowment. He proved that a sufficiency condition for the existence of equilibrium is that this graph be strongly connected. Hence the question arises: is this a sufficiency condition for equilibrium existence in ADHZ as well? We provide a negative answer to this question: we give an instance of ADHZ whose demand graph is not only strongly connected but also has dichotomous utilities, and yet it does not admit an equilibrium.

Given this, we define an *approximate* equilibrium notion which is closely related to the α-slack equilibria of Echenique et al. [42]. We call this the *ϵ-approximate ADHZ model* and this is the main focus of the chapter. We start by showing existence of ϵ-approximate ADHZ equilibria based on the corresponding existence result for α-slack equilibria.

We prove that the equilibrium in our ϵ-approximate ADHZ model is Pareto optimal, approximately envy-free (in a sense that takes into account the endowments of the agents), and approximately weak core stable. This is analogous to the fairness and efficiency of HZ equilibria that we observed in Section 1.3.

Next, we turn to computational matters. Since the computation of approximate HZ equilibria is PPAD-hard [28], there is little hope in finding ϵ-approximate ADHZ equilibria in polynomial time. For this reason, we focus on the special case of dichotomous utilities. Our results easily extend also to more general bivalued utilities which are well-studied in the literature [18, 122, 12, 60, 39]. Our main result is a combinatorial, polynomial time algorithm for computing ϵ-approximate ADHZ equilibria. This is covered in Section 2.4.

We also study the rationality of equilibria. By extending an example from Vazirani and Yannakakis [122], we show that there are instances on which an ADHZ equilibrium exists but for which all ADHZ equilibria are irrational. However, for dichotomous utilities, we show that there is always a rational α-slack equilibrium. Note that there is also always a rational ϵ-approximate equilibrium but this is a trivial result due to the slack that the ϵ affords us.

Finally, we briefly turn to the classic HZ setting *without* endowments. It was known that HZ with dichotomous utilities always admits rational equilibria [122]. This raises the obvious question: is there a rational convex program for HZ? We answer this question in the positive in Section 2.5.

2.1.2 Related Results

A notable and particularly influential example of a matching market with endowments is the housing market studied by Shapley and Scarf [113]. Here, agents come to the market to trade houses. Each agent has an initial endowment of exactly one house, i.e. is an existing tenant of one house, and has *ordinal* preferences over all the houses in the market. The goal is to find an integral reassignment of the houses to the agents. Shapley and Scarf show that the so-called Top Trading Cycles (TTC) mechanism[1] finds a Pareto-optimal and core-stable allocation. Moreover, it is strategyproof.

An interesting extension of this work is due to Athanassoglou and Sethuraman [7] who consider the same problem but with fractional endowments and fractional reassignments (as we do in this paper). Their mechanism is Pareto-optimal, individually rational, and justified envy-free. Moreover, it is weakly strategyproof. They show that individual rationality, Pareto-optimality, and strategyproofness are incompatible. This negative result was later strengthened by Aziz [10] who showed that the condition of strategyproofness can be relaxed to weak strategyproofness.

[1]Shapley and Scarf attribute the TTC algorithm to Gale.

The core difference between these works and ours is that they rely on ordinal preferences rather than cardinal utilities. As discussed in Section 1.1, mechanisms that have access to cardinal utilities are generally able to achieve much higher degrees of efficiency. However, the cardinal benchmark of Pareto-optimality is much harder to achieve which is why we look towards pricing-based mechanisms like the Hylland Zeckhauser mechanism.

Vazirani and Yannakakis [122] undertook a comprehensive study of the computational complexity of the HZ scheme. They gave a combinatorial polynomial time algorithm for dichotomous utilities and an example which has only irrational equilibria; as a consequence, this problem is not in PPAD. They showed that the problem of computing an exact HZ equilibrium is in the class FIXP and the problem of computing an approximate equilibrium is in PPAD. Chen et al. [28] showed that computing an approximate HZ equilibrium is in fact PPAD-hard. In order to deal with the computational intractability of HZ, a Nash-bargaining-based mechanism was proposed by Hosseini and Vazirani [70]. We will cover this in more detail in the next chapter.

The study of the dichotomous case of matching markets was initiated by Bogomolnaia and Moulin [18]. They studied a two-sided matching market and they called it an "important special case of the bilateral matching problem." Using the Gallai-Edmonds decomposition of a bipartite graph, they gave a mechanism that is Pareto optimal and group strategyproof. They also gave a number of applications of their setting, some of which are natural applications of one-sided markets as well such as housemates distributing rooms in a house. As in the HZ scheme, their mechanism also outputs a doubly-stochastic matrix whose entries represent probability shares of allocations. However, they give another interesting interpretation of this matrix. They say, "time sharing is the simplest way to deal fairly with indivisibilities of matching markets: think of a set of workers sharing their time among a set of employers." Roth, Sönmez and Ünver [111] extended these results to general graph matching under dichotomous utilities; this setting is applicable to the kidney exchange marketplace.

Several researchers have proposed Hylland-Zeckhauser-like mechanisms for a number of other applications, for instance [22, 69, 88, 97]. The basic scheme has also been generalized in several different directions, including two-sided matching markets, adding quantitative constraints, and to some settings in which agents have initial endowments of goods, see [42, 41], though these generalizations often come with severe limitations such as requiring personalized prices. As a result they often do not enjoy the same level of fairness as the traditional HZ mechanism.

2.2 The Hylland-Zeckhauser Mechanism

We will now briefly recall some important facts about Fisher and HZ equilibria; see also Sections 1.2 and 1.3. Both equilibria apply to a setting in which we have a set A of agents and a a set G of goods. Each agent i comes to the market with a budget b_i and has non-negative, rational utilities $(u_{ij})_{j\in G}$ for all the goods in the market. Our goal is to fractionally assign the goods to the agents according to their budgets by determining equilibrium prices.

In the Fisher setting, agents are allowed to get any amount of goods. A formal definition of a Fisher equilibrium is given below.

Definition 2.1. *A* Fisher equilibrium *is a pair* (x, p) *consisting of an* allocation $(x_{ij})_{i\in A, j\in G}$ *and* prices $(p_j)_{j\in G}$ *with the following properties.*

1. *Each agent* i *spends at most their budget, i.e.* $p \cdot x_i \leq b_i$.

2. *Each agent i gets an* optimal bundle*, i.e. a utility maximizing bundle at prices p. Formally:*

$$u_i \cdot x_i = \max \left\{ u_i \cdot y \mid y \in \mathbb{R}^G_{\geq 0}, p \cdot y \leq b_i \right\}.$$

3. *The market clears, i.e. each good with positive price is fully allocated to the agents.*

The set of equilibria of a linear Fisher market is non-empty and corresponds to the set of optimal solutions of the Eisenberg-Gale convex program [47], which is a *rational convex program* (RCP) and in fact it motivated the definition of this concept [119]. Moreover, Fisher equilibria satisfy various nice properties, including equal-type envy-freeness and Pareto optimality.

Definition 2.2. *An allocation is* envy-free *if for any two agents $i, i' \in A$, agent i weakly prefers their allocation to that of i', i.e. $u_i \cdot x_i \geq u_i \cdot x_{i'}$. It is* equal-type envy-free *if the above holds for any two agents with identical budgets.*

The difference between the Fisher setting and the HZ setting is that in the latter, each agent must get exactly one unit of goods. Hence, we also require that $|A| = |G|$ so that a perfect matching can exist. An HZ equilibrium is formally defined below.

Definition 2.3. *A* Hylland-Zeckhauser (HZ) equilibrium *is a pair (x, p) consisting of an* allocation $(x_{ij})_{i\in A, j\in G}$ *and* prices $(p_j)_{j\in G}$ *with the following properties.*

1. *x is a fractional perfect matching, i.e. $\sum_{j\in G} x_{ij} = 1$ for all i and $\sum_{i\in A} x_{ij} = 1$ for all j.*

2. *Each agent i spends at most their budget, i.e. $p \cdot x_i \leq b_i$ (usually $b_i = 1$).*

3. *Each agent i gets an* optimal bundle, *which is defined to be a cheapest utility maximizing bundle:*

$$u_i \cdot x_i = \max \left\{ u_i \cdot y \;\middle|\; y \in \mathbb{R}^G_{\geq 0}, \sum_{j \in G} y_j = 1, p \cdot y \leq b_i \right\},$$
$$p \cdot x_i = \min \left\{ p \cdot y \;\middle|\; y \in \mathbb{R}^G_{\geq 0}, \sum_{j \in G} y_j = 1, u_i \cdot y \geq u_i \cdot x_i \right\}.$$

Like Fisher equilibria, HZ equilibria are Pareto optimal and envy-free (assuming unit budgets). The allocation x found by the HZ mechanism is a fractional perfect matching or a doubly-stochastic matrix. In order to get an integral perfect matching from x, a lottery can be carried out using the Birkhoff-von-Neumann theorem [15, 124].

2.3 The ϵ-Approximate ADHZ Model

In Section 1.2, we also defined the Arrow Debreu (AD) market equilibrium. In this so-called *exchange setting*, we still have agents A and goods G. The difference is that agents have non-negative, rational endowment vectors $(e_{ij})_{j \in G}$ rather than budgets. Here, e_{ij} tells us how much of good j agent i brings to the market. We typically assume that each good is fully owned by the agents, i.e. $\sum_{i \in A} e_{ij} = 1$ for all $j \in G$. A formal definition for the corresponding equilibrium concept is given below.

Definition 2.4. *An* Arrow-Debreu (AD) equilibrium *for a given AD market is a pair (x, p) consisting of an* allocation $(x_{ij})_{i \in A, j \in G}$ *and* prices $(p_j)_{j \in G}$ *with the following properties.*

1. *Each agent spends at most the budget earned from the endowment, i.e. $p \cdot x_i \leq b_i := p \cdot e_i$.*

2. *Each agent i gets an* optimal bundle*:*

$$u_i \cdot x_i = \max \left\{ u_i \cdot y \mid y \in \mathbb{R}^G_{\geq 0}, p \cdot y \leq b_i \right\}.$$

3. *The market clears, i.e. each good with positive price is fully allocated to the agents.*

The AD model generalizes the Fisher model in the sense that any Fisher market can be easily transformed into an AD market by giving each agent a fixed proportion of every good. Clearly, AD equilibria satisfy the condition of individual rationality, defined below, since every agent could always buy back their endowment.

Definition 2.5. *An allocation in an AD market is* individually rational *if for every agent i we have $u_i \cdot x_i \geq u_i \cdot e_i$, i.e. no agent loses utility by participating in the market.*

Individual rationality is a desirable property since otherwise we would be required to incentivize agents to participate in the market in the first place. However, individual rationality fundamentally clashes with envy-freeness. Consider a market consisting of two agents each owning a distinct good. Assume that both agents prefer the good of agent 2 over the good of agent 1, then in any allocation either agent 1 envies agent 2 or agent 2's individual rationality is violated. For this reason we primarily consider a version of equal-type envy-freeness in exchange markets, which demands envy-freeness only for agents with the same initial endowment.

In contrast to Fisher markets, Arrow-Debreu equilibria do not always exist. However, there is a simple necessary and sufficient condition for their existence based on *strong connectivity of demand graph*, due to Gale [55]. See Theorem 1.8. An RCP for this problem was given by Devanur, Garg and Végh [33].

We now turn to the extension of the HZ mechanism to exchange markets. In the *ADHZ market*, we have a set A of *agents* and a set G of *goods* with $|A| = |G| = n$. Each agent i comes with an *endowment* $e_{ij} \geq 0$ of each good j and utilities $u_{ij} \geq 0$. We require that the goods are fully owned by the agents and hence the endowment vector e must be a fractional perfect matching.

Definition 2.6. *An* ADHZ equilibrium *for a given ADHZ market is a pair* (x, p) *consisting of an* allocation $(x_{ij})_{i \in A, j \in G}$ *and* prices $(p_j)_{j \in G}$ *with the following properties.*

1. *x is a fractional perfect matching, i.e.* $\sum_{j \in G} x_{ij} = 1$ *for all i and* $\sum_{i \in A} x_{ij} = 1$ *for all j.*
2. *Each agent spends at most the budget earned from the endowment, i.e.* $p \cdot x \leq b_i := p \cdot e_i$.
3. *Each agent i gets an* optimal bundle, *which is defined to be a cheapest utility maximizing bundle:*

$$u_i \cdot x_i = \max \left\{ u_i \cdot y \;\middle|\; y \in \mathbb{R}^G_{\geq 0}, \sum_{j \in G} y_j = 1, p \cdot y \leq b_i \right\},$$
$$p \cdot x_i = \min \left\{ p \cdot y \;\middle|\; y \in \mathbb{R}^G_{\geq 0}, \sum_{j \in G} y_j = 1, u_i \cdot y \geq u_i \cdot x_i \right\}.$$

This equilibrium notion naturally satisfies the desirable properties that we are looking for.

Theorem 2.1. *ADHZ equilibria are Pareto optimal, individually rational, and equal-type envy-free.*

Proof. Pareto optimality follows from the fact that any ADHZ equilibrium is an HZ equilibrium with certain budgets b. Since any HZ equilibrium is Pareto optimal, we get the same for ADHZ.

Note that the budget of any agent is always enough to buy back their initial endowment. Since they get an optimal bundle, they must get something which they value at least as high as their initial endowment. Thus individual rationality is guaranteed.

If two agents, say 1 and 2, have the same endowment, then their budget will be the same and so agent 1 will never value the 2's bundle higher than their own. Thus ADHZ equilibria are equal-type envy-free. □

In addition, ADHZ equilibria also satisfy the following notion of core-stability.

Definition 2.7. *An allocation x in an ADHZ market is* weakly core-stable *if for any subsets $A' \subseteq A$ and $G' \subseteq G$, there does not exist an allocation $x' \in \mathbb{R}_{\geq 0}^{A' \times G'}$ such that*

- *x' allocates at most one unit of goods to every agent in A',*
- *every good $j \in G'$ is allocated at most to the extent of the endowments of the agents in A', i.e. $\sum_{i \in A'} x'_{ij} \leq \sum_{i \in A'} e_{ij}$, and*
- *every agent in A' receives strictly better utility in x' than in x.*

Theorem 2.2. *ADHZ equilibria are weakly core-stable.*

Proof. Let (x, p) be some ADHZ equilibrium. For the sake of a contradiction, assume that there are $A' \subseteq A$, $G' \subseteq G$ and $x' \in \mathbb{R}_{\geq 0}^{A' \times G'}$ as excluded by the definition of weak core-stability. Now consider the total money spent along allocation x', i.e. the quantity $\sum_{i \in A'} \sum_{j \in G'} p_j x'_{ij}$.

On the one hand we know that only the endowment of the agents in A' is allocated by x'. Thus

$$\sum_{i \in A'} \sum_{j \in G'} p_j x'_{ij} \leq \sum_{i \in A'} \sum_{j \in G'} p_j e_{ij}. \tag{2.1}$$

On the other hand, every agent i receives strictly better utility from x' than from x. But since agents buy optimal bundles in (x, p), this implies that the bundles in x' must be worth more than their budget, i.e.

$$\sum_{j\in G'} p_j x'_{ij} > \sum_{j\in G} p_j e_{ij} \geq \sum_{j\in G'} p_j e_{ij}.$$

Summing this inequality over all $i \in A'$ yields a contradiction to (2.1). □

Like in the case of HZ, equilibrium prices in ADHZ are invariant under the operation of *scaling* the difference of prices from 1, as shown in the following lemma.

Lemma 2.1. *Let p be an equilibrium price vector. For any $r > 0$, let p' be such that $p'_j - 1 = r(p_j - 1)$ for all $j \in G$. Then p' is also an equilibrium price vector.*

Proof. Let x be an equilibrium allocation at prices p. We want to show that (x, p') is also an equilibrium. The core observation is that if $y \in \mathbb{R}^G_{\geq 0}$ satisfies $\sum_{j\in G} y_j = 1$, then

$$\begin{aligned} p' \cdot y &= \sum_{j\in G}(1 + r(p_j - 1))y_j \\ &= 1 + rp \cdot y - r \end{aligned}$$

So if we have another $y' \in \mathbb{R}^G_{\geq 0}$ with $\sum_{j\in G} y'_j = 1$, we can see that $p \cdot y \leq p \cdot y'$ if and only if $p' \cdot y \leq p' \cdot y'$. This implies in particular, that the set of feasible bundles for any agent i (i.e. those that add up to 1 and do not exceed the budget of $p \cdot e_i$ or $p' \cdot e_i$ respectively) is the same. Moreover, if a bundle is cheapest wrt. to p it is also cheapest wrt. to p'. Hence, x consists of optimal bundles for each agent under prices p' and (x, p') is therefore an ADHZ equilibrium. □

Unlike HZ, which always admits an equilibrium, ADHZ has instances which do not admit an equilibrium, as observed by Hylland and Zeckhauser [75]. Below we give a counterexample in which the demand graph is strongly connected and utilities are dichotomous.

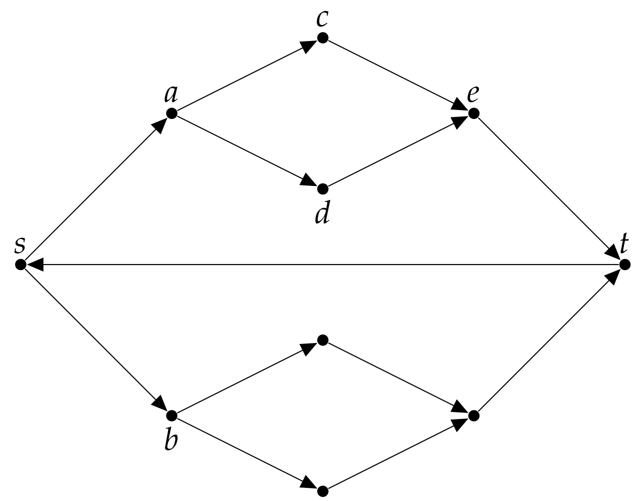

Figure 2.1: The demand graph of an ADHZ market with dichotomous utilities and no equilibrium. Each node represents an agent as well as the good possessed by this agent in their initial endowment. An arrow from i to j represents $u_{ij} = 1$; the rest of the edges have utility 0.

Proposition 2.1. *The ADHZ market with dichotomous utilities in Figure 2.1 does not admit an equilibrium.*

Proof. Assume there is an equilibrium (x, p) in this market. Further, using Lemma 2.1, we can assume that the minimum price is zero. This implies that no agent will buy a zero utility good at a positive price.

Each agent buys a total of one unit of goods and s is the only agent having positive utility for goods a and b. Therefore, at least one of these goods is not fully sold to s and must be sold to an agent deriving zero utility from it. Therefore, this good must have zero price. Without loss of generality, assume $p_a = 0$. Since a has no budget and c and d are desired only by a, $p_c = p_d = 0$, otherwise c and d cannot be sold. For the same reason, $p_e = 0$. Now observe that both agents c and d have a utility 1 edge to a good of price zero, namely e. Therefore, the optimal bundle of both c and d is e. But then e would have to be matched twice which is a contradiction. $\square$

Even if ADHZ equilibria *do* exist, computing them is at least as hard as computing HZ equilibria. This follows from the following reduction.

Proposition 2.2. *Consider an HZ market with unit budgets. Define an ADHZ market by giving every agent as endowment an equal amount of every good. Then every HZ equilibrium in which the prices sum up to n is an ADHZ equilibrium and every ADHZ equilibrium yields an HZ equilibrium by rescaling all prices by $n/\sum_{j\in G} p_j$.*

Vazirani and Yannakakis [122] gave an instance of HZ with four agents and four goods which has one equilibrium in which all agents fully spend their budgets, and allocations and prices are irrational. Since this example satisfies the conditions of Proposition 2.2, we get that the modification of the example of [122], as stated in the Proposition, is an instance for ADHZ having only irrational equilibria.

2.3.1 Existence and Properties of ϵ-Approximate ADHZ Equilibria

Since ADHZ equilibria do not always exist, we study the following approximate equilibrium notion instead.

Definition 2.8. *An ϵ-*approximate ADHZ equilibrium *is an HZ equilibrium (x, p) for a budget vector b with*

$$(1-\epsilon)p \cdot e_i \leq b_i \leq \epsilon + p \cdot e_i$$

for all $i \in A$. We also require that if two agents have the same endowment, then their budget should also be the same.

The additive error term in the upper bound is necessary since otherwise the counterexample from Proposition 2.1 still works. On the other hand, the multiplicative lower bound is useful to get approximate individual rationality. However, one can always find approximate equilibria in which the sum of prices is bounded by n using Lemma 2.1, so we also get

$$p \cdot e_i - \epsilon' \leq b_i \leq p \cdot e_i + \epsilon'$$

where $\epsilon' := n\epsilon$. This implies that we can equivalently define the above notion with additive error terms on both upper and lower bounds.

In our notion of approximate equilibrium, we do not relax the fractional perfect matching constraints or the optimal bundle condition. We only allow the budgets of agents to be slightly different from the money they would normally obtain in an ADHZ market. Hence the step of randomly rounding the equilibrium allocation to an integral perfect matching is the same as in the HZ scheme.

Theorem 2.3. *Any ϵ-approximate ADHZ equilibrium is Pareto optimal, ϵ-approximately individually rational, and equal-type envy-free.*

Proof. Pareto optimality follows just as for the non-approximate ADHZ setting from the fact that an ϵ-approximate ADHZ equilibrium is first and foremost an HZ equilibrium. For approximate individual rationality note that every agent gets a budget of at least $(1 - \epsilon)$ times the cost of their endowment. Hence their utility can decrease by at most a factor of $(1 - \epsilon)$. Equal-type envy-freeness follows immediately from the condition that agents with the same endowment have the same budget. □

One can also define a suitably ϵ-approximate notion of weak core-stability, where instead of demanding that every agent strictly improves in the seceding coalition, we instead require that every agent improves by a factor of more than $\frac{1}{1-\epsilon}$.

Theorem 2.4. *Any ϵ-approximate ADHZ equilibrium is ϵ- approximately weak-core stable.*

Proof. Let (x, p) be an ϵ-approximate ADHZ equilibrium for some budget vector b. Then in order for some other allocation x' to improve agent i's utility by a factor of more than $\frac{1}{1-\epsilon}$, i must spend more than $\frac{b_i}{1-\epsilon}$. But note that $\frac{b_i}{1-\epsilon} \geq p \cdot e_i$. From here the proof is identical to that of Theorem 2.2. □

While approximate equilibrium notions are more amenable to computation, they generally do not lend themselves well to existence proofs. However, our notion of ϵ-approximate ADHZ equilibrium is a slight relaxation of the notion of an α-slack equilibrium introduced in [42].

Definition 2.9. *An α-*slack ADHZ equilibrium *for $\alpha \in (0, 1]$ is an HZ equilibrium (x, p) for a budget vector b in which $b_i = \alpha + (1-\alpha)\sum_{j \in G} p_j e_{ij}$ for all $i \in A$.*

Theorem 2.5 (Theorem 2 in [42]). *In any ADHZ market, α-slack equilibria always exist if $\alpha > 0$.*

Note that any α-slack equilibrium is automatically also an α-approximate equilibrium. Thus we get:

Theorem 2.6. *In any ADHZ market, ϵ-approximate equilibria always exist if $\epsilon > 0$.*

2.4 Algorithm for ϵ-approximate ADHZ under Dichotomous Utilities

Before we can tackle the ADHZ setting, let us first give an algorithm that can compute HZ equilibria with non-uniform budgets. This is an extension of the algorithm presented in [122] and may be of independent interest. In the following, fix some HZ market consisting of n agents and goods with $u_{ij} \in \{0, 1\}$ for all $i \in A$ and $j \in G$. If $u_{ij} = 1$, we will say that i *likes* j (and *dislikes* otherwise). We assume that every agent likes at least one good.

We remark that any HZ equilibrium (x, p) for the utilities u is also an equilibrium for the utilities

$$\tilde{u}_{ij} := \begin{cases} a_i & \text{if } u_{ij} = 1 \\ b_i & \text{if } u_{ij} = 0 \end{cases}$$

where $0 \leq a_i < b_i$ are arbitrary for every agent. This is because

$$\tilde{u}_i \cdot x_i = a_i + (b_i - a_i) u_i \cdot x_i$$

since x is a fractional perfect matching. Hence, utility function $\tilde{u}$ is an affine transformation of utility function u on the space of feasible bundles. This allows us to restrict ourselves to dichotomous utilities instead of the more general bi-valued utilities.

Before we give the actual algorithm, we will first need a characterization of when a price vector can be extended to an HZ equilibrium in terms of certain flow networks.

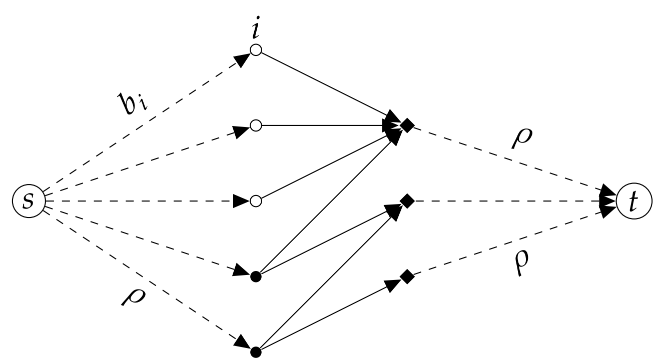

Figure 2.2: Shown is the flow network which corresponds to finding an equilibrium allocation in price class ρ. Filled circles represent agents in $A(\rho)$ with $b_i < \rho$, empty circles are agents in $A(\rho)$ with $b_i \geq \rho$, and diamond vertices are goods in $G(\rho)$. The contiguous edges represent all utility 1 edges and have infinity capacity (utility 0 edges are not part of the network). Dashed edges to empty circle vertices i have capacity b_i whereas the other dashed edges have capacity ρ.

Lemma 2.2. *Let $(p_j)_{j \in G}$ be non-negative prices. For any $\rho \geq 0$, let $G(\rho)$ be the goods which are sold at price ρ and let $A(\rho)$ be those agents for which the cheapest price of any liked good is ρ. Assume that*

- *there is a matching in the utility 1 edges on $A(0) \cup G(0)$ which covers all agents in $A(0)$ and*
- *if $\rho > 0$ is equal to the price of some good, then the flow network shown in Figure 2.2 has a maximum flow of size $\rho|G(\rho)|$.*

Then we can find a fractional perfect matching x which makes (x, p) an HZ equilibrium in polynomial time.

Proof. Allocate every agent in $A(0)$ to some good in $G(0)$ according to the matching which exists by assumption. Let $\rho > 0$, be the price of some good. Then we compute the maximum flow $f^{(\rho)}$ in the flow network from Figure 2.2 and allocate $x_{ij} = f^{(\rho)}_{i,j}/\rho$ for all $i \in A(\rho)$ and $j \in G(\rho)$. Lastly, extend x to a fractional perfect matching by matching the remaining capacity of the agents to the remaining capacity of goods in $G(0)$.

Clearly, no agent exceeds their budget. To see that this yields an HZ equilibrium, note that every agent only spends money on cheapest liked goods and if they do not get allocated entirely to liked goods, then they additionally spend all of their budget. This ensures that every agent gets an optimal bundle. □

Theorem 2.7. *For any rational budget vector b, we can compute an HZ equilibrium in polynomial time.*

Proof. We start in the same way is the algorithm by Vazirani and Yannakakis [122]: by computing a minimum vertex cover in the graph of utility 1 edges, we can partition $A = A_1 \cup A_2$ and $G_1 \cup G_2$ such that

- every agent in A_2 can be matched to a distinct liked good in G_2,
- every agent in A_1 only has liked goods in G_1, and
- for every $S \subseteq G_2$ we have $|N^-(S)| \geq |S|$ where $N^-(S)$ are the agents that have a liked good in S.

Set $p_j = 0$ for all $j \in G_2$ and $p_j = \min_{i \in A_1} b_i$ for all $j \in G_1$. Now we run a DPSV-like [36] algorithm on $A_1 \cup G_1$ to raise prices until certain sets of goods become tight.

For each $i \in A$, let β_i be its *effective* budget at current prices p, that is the minimum of its actual budget b_i and the price of its cheapest liked good. The algorithm will now raise all prices p at the same rate until there is a set $S \subseteq G_1$ which goes *tight* in the sense that $\sum_{i \in \Gamma(S)} \beta_i = \sum_{j \in S} p_j$ where Γ is the collection of agents which have a cheapest liked good in S. At this point, we freeze the prices of the goods in S. If all prices have been frozen we are done. Otherwise, we continue raising all unfrozen prices of goods in G_1.

It is easy to see that if the prices keep rising, eventually each agents' effective budget will be their real budget and so a set must become tight at some point. We will not go into detail here but it is possible to find the next set which will go tight in polynomial time (for example using a parametric max flow algorithm) similar as in DPSV. Finally, since we never unfreeze prices, there will be at most n iterations of the algorithm and hence it runs in polynomial time overall.

We observe that as in the proof of the DPSV algorithm, for any $S \subseteq G_1$, we have that $\sum_{i \in \Gamma(S)} \beta_i \geq \sum_{j \in A} p_j$ and $\sum_{i \in A_1} \beta_i = \sum_{j \in G_1} p_j$. It is then easy to show that this implies that for any price ρ above 0, the corresponding flow network from Figure 2.2 supports a flow of value $\rho|G(\rho)|$ by the max-flow min-cut theorem. Thus we can apply Lemma 2.2 to get an equilibrium allocation. □

Lemma 2.3. *Let b and b' be two budget vectors with $0 \leq b \leq b'$. Assume we are given an HZ equilibrium (x, p) for the budgets b. Then we can compute in polynomial time a new HZ equilibrium (x', p') with $p \leq p'$ for the budgets b'.*

Proof. We will simply run the same algorithm as in the proof of Theorem 2.7, except that this time we start with the prices p. More precisely, we increase the lowest non-zero price until a set goes tight or it becomes equal to the next higher price, then repeat this process until we once again get $\sum_{i \in \Gamma(S)} \beta_i \geq \sum_{j \in A} p_j$ and $\sum_{i \in \Gamma(G_1)} \beta_i = \sum_{j \in G_1} p_j$ where G_1 is now defined as the set of goods with positive prices in (x, p). As in the proof of Theorem 2.7, this will freeze all prices in polynomial time at which point we can use a max-flow min-cut argument to construct the new equilibrium allocation x' in polynomial time. □

Let us now return to the approximate ADHZ setting. Instead of budgets, fix now some fractional perfect matching of endowments $(e_{ij})_{i \in A, j \in G}$.

Theorem 2.8. *An ϵ-approximate ADHZ equilibrium for rational $\epsilon \in (0, 1)$, can be computed in time polynomial in $\frac{1}{\epsilon}$ and n, i.e. by a fully polynomial time approximation scheme.*

Proof. We will iteratively apply Lemma 2.3. Start by setting $b_i^{(1)} := \frac{\epsilon}{2}$ for all $i \in A$ and computing an HZ equilibrium $(x^{(1)}, p^{(1)})$ according to Theorem 2.7. Beginning with $k := 1$, we run the following algorithm.

1. Let $b_i^{(k+1)} := \frac{\epsilon}{2} + (1 - \frac{\epsilon}{2}) \sum_{j \in G} p_j^{(k)} e_{ij}$ for all $i \in A$.
2. Compute a new HZ equilibrium $(x^{(k+1)}, p^{(k+1)})$ for budgets $b^{(k+1)}$ according to Lemma 2.3 using the old equilibrium $(x^{(k)}, p^{(k)})$ as the starting point. Note that since $p^{(k)} \geq p^{(k-1)}$ we always have $b^{(k+1)} \geq b^{(k)}$ and so this is well-defined.
3. Set $k := k + 1$ and go back to step 1.

Note that

$$\sum_{i\in A} b_i^{(k+1)} = \frac{\epsilon}{2}n + \left(1 - \frac{\epsilon}{2}\right)\sum_{j\in G} p_j^{(k)} \quad \le \frac{\epsilon}{2}n + \left(1 - \frac{\epsilon}{2}\right)\sum_{i\in A} b_i^{(k)}$$

and thus

$$\sum_{j\in G} p_j^{(k)} \le \sum_{i\in A} b_i^{(k)} \le n$$

as otherwise we would get $\sum_{i\in A} b_i^{(k+1)} < \sum_{i\in A} b_i^{(k)}$.

Let K be the first iteration such that $p^{(K)} \le \frac{1-\epsilon/2}{1-\epsilon} p^{(K-1)}$. Note that

$$K \le n \log_{\frac{1-\epsilon/2}{1-\epsilon}} \left(\frac{n}{\epsilon}\right) = O\left(\frac{n}{\epsilon}\log\left(\frac{n}{\epsilon}\right)\right)$$

since all non-zero prices are initialized to at least ϵ but are bounded by n. Then $(x^{(K)}, p^{(K)})$ is an ϵ-approximate ADHZ equilibrium with budget vector $b^{(K)}$ because for all $i \in A$ we have

$$\begin{aligned} b_i^{(K)} &= \frac{\epsilon}{2} + \left(1 - \frac{\epsilon}{2}\right) p^{(K-1)} \cdot e_i \\ &\in \left[(1-\epsilon)p^{(K)} \cdot e_i, \epsilon + p^{(K)} \cdot e_i\right]. \end{aligned}$$

Lastly, we note that since the number of iterations is bounded by $O(\frac{n}{\epsilon}\log(\frac{n}{\epsilon}))$ and each iteration runs in polynomial time, the total runtime is polynomial in $\frac{1}{\epsilon}$ and n as claimed. □

2.5 An RCP for the HZ Scheme under Dichotomous Utilities

We will now consider the setting without initial endowments, i.e. the classic HZ setting. Our goal in this section is to give a rational convex program (RCP) for the classic HZ equilibrium. As usual, we will assume without loss of generality that each agent likes at least one good and that each good is liked by at least one agent. Consider now the following convex program.

$$\max \quad \sum_{i \in A} \log(u_i \cdot x_i) \tag{2.2a}$$

$$\text{s.t.} \quad \sum_{i \in A} x_{ij} \leq 1 \quad \forall j \in G, \tag{2.2b}$$

$$\sum_{j \in G} x_{ij} \leq 1 \quad \forall i \in A, \tag{2.2c}$$

$$x_{ij} \geq 0 \quad \forall i \in A, j \in G \tag{2.2d}$$

Note the similarity to the standard Eisenberg Gale RCP for the Fisher market which we introduced in Section 1.2; the only difference is the additional matching constraint. The EG program is an RCP for the Fisher market with arbitrary utilities. For HZ, we have to make the weaker claim that (2.2) is an RCP for dichotomous utilities.

Theorem 2.9. *Under dichotomous utilities, any HZ equilibrium allocation is an optimal solution to* (2.2)*, and every optimal solution of* (2.2) *can be trivially extended to an HZ equilibrium allocation.*

Proof. The proof relies on the KKT conditions of (2.2). Let p_j be the dual variables corresponding to constraint (2.2b) and let α_i be the dual variables corresponding to (2.2c). Note that p and α are non-negative. Then the KKT conditions for an optimal solution of (2.2) are as follows.

1. If $\alpha_i > 0$, then $\sum_{j \in G} x_{ij} = 1$.

2. If $p_j > 0$, then $\sum_{i \in A} x_{ij} = 1$.

3. $u_{ij} \leq u_i \cdot x_i(p_j + \alpha_i)$ for all $i \in A, j \in G$.

4. If $x_{ij} > 0$, then $u_{ij} = u_i \cdot x_i(p_j + \alpha_i)$.

Let us show the forward direction; given an HZ equilibrium (x, p), we wish to show that x is an optimal solution to (2.2). Note that x is a fractional perfect matching between agents and goods and hence satisfies the constraints of the CP. We can assume wlog. that at least one good has price zero.

Let α_i be the unspent budget of agent i, i.e. $\alpha_i := 1 - x_i \cdot p$. We want to show that x together with (α, p) satisfies the KKT conditions of (2.2). Clearly, conditions 1 and 2 are satisfied trivially. Consider now some agent i. We distinguish two cases.

Case 1: there is some $j \in G$ with $u_{ij} = 1$ and $p_j \leq 1$. In this case, $u_i \cdot x_i = 1$ and agent i is exclusively allocated goods j with $u_{ij} = 1$ and which have minimum price among those goods (since i's bundle must be cheapest). Now let $j \in G$ be arbitrary. Clearly if $u_{ij} = 0$, then condition 3 holds trivially and so does 4 (since $x_{ij} = 0$). So assume $u_{ij} = 1$. But then we have $p_j + \alpha_i \geq 1$ with equality when j is a cheapest utility 1 good. This follows because i is exclusively allocated cheapest goods. Finally, this implies that conditions 3 and 4 hold.

Case 2: all $j \in G$ with $u_{ij} = 1$ have $p_j > 1$. In this case, $u_i \cdot x_i = \frac{1}{p_{\min}} < 1$ and $\alpha_i = 0$ where $p_{\min}$ is the minimum price of utility 1 goods. Let $j \in G$ be arbitrary. Once again, if $u_{ij} = 0$, conditions 3 and 4 hold trivially. Otherwise, conditions 3 and 4 hold simply because i can only get allocated to cheapest utility 1 goods in order to satisfy the cheapest bundle condition.

Finally, let us show the reverse direction. Assume that x is an optimal solution to (2.2) and (α, p) are the corresponding dual variables that satisfy the KKT conditions. We call p the prices on the goods and observe that $u_i \cdot x_i > 0$ for all agents i.

Let i, j be arbitrary with $x_{ij} > 0$. Then we can first observe that if $u_{ij} = 0$, we have $p_j = \alpha_i = 0$ by KKT condition 4 and the fact that $u_i \cdot x_i > 0$. And if $u_{ij} = 1$, then conditions 3 and 4 imply together that j is a minimum price utility 1 good for agent i. Hence, each agent is spending their budget optimally.

Next, fix some agent i, multiply the equality in condition 4 by x_{ij}, and sum over all j. This yields

$$u_i \cdot x_i = u_i \cdot x_i \sum_{j \in G} (p_j + \alpha_i) x_{ij}$$

and hence

$$1 = x_i \cdot p + \alpha_i \sum_{j \in G} x_{ij}.$$

If $\alpha_i > 0$, then we know that $\sum_{j \in G} x_{ij} = 1$ and hence $\alpha_i = 1 - x_i \cdot p$. On the other hand, if $\alpha_i = 0$, then we also get $\alpha_i = 1 - x_i \cdot p$. So we know that α_i measures the remaining budget of agent i.

Finally, let S denote the set of agents who get less than one unit of goods, i.e. $S := \{i \in A \mid \sum_{j \in G} x_{ij} < 1\}$, and let T denote the set of partially allocated goods, i.e. $T := \{j \in G \mid \sum_{i \in A} x_{ij} < 1\}$. By condition 2, $p_j = 0$ for each $j \in T$. Observe that there cannot be any utility 1 edge between S and T as otherwise we could increase the objective of (2.2).

Since the number of agents equals the number of goods, the total deficiency of agents in solution x equals the total amount of unallocated goods. Therefore, we can arbitrarily allocate unallocated goods in T to deficient agents in S so as to obtain a fractional perfect matching, say x'. Clearly, (x', p) is still an optimal solution to (2.2) and is an HZ equilibrium since each agent still has a cheapest optimal bundle. □

The proof of Theorem 2.9 shows that for the dichotomous case, the dual of (2.2) yields equilibrium prices. In contrast, for arbitrary utilities, there is no known mathematical construct, no matter how inefficient its computation, that yields equilibrium prices. In a sense, this should not be surprising, since there is a polynomial time algorithm for computing an equilibrium for the dichotomous case [122].

In addition, since the objective function in (2.2) is strictly concave in the utilities $u_i \cdot x_i$, the utility derived by each agent i must be the same in all solutions of (2.2). Hence, we get the following corollary which can be seen as a variant of the well-known *Rural Hospital Theorem*; see [67] for the latter.

Corollary 2.1. *Each agent gets the same utility under all HZ equilibria with dichotomous utilities.*

Finally, let us prove that (2.2) is indeed an RCP by establishing the existence of rational optimal solutions.

Theorem 2.10. *There always exists an optimal solution to* (2.2) *which can be expressed via rational numbers with denominators at most n.*

Proof. Let x be some optimal solution of (2.2) and let (α, p) be the corresponding dual variables. Recall from the proof of Theorem 2.9 that each x_i will be an optimal bundle for agent i with prices p and that $\alpha_i = 1 - x_i \cdot p$. Moreover, we can assume wlog. that x is a fractional perfect matching and hence an HZ equilibrium.

Let $G' \subseteq G$ denote the set of goods with prices strictly greater than 1 and let $A' \subseteq A$ denote the set of agents who are allocated some positive amount of goods from G'. Since agents are getting cheapest optimal bundles, each agent in A' only has utility 1 edges to G'. Moreover, for each $i \in A'$, we know that $\alpha_i = 0$, i.e. agent i must spend all their money.

Consider now the connected components of the bipartite graph (A', G', E) where $E := \{(i, j) \in A' \times G' \mid x_{ij} > 0\}$. Because each agent is only allocated to *cheapest* utility 1 goods, we know that in each connected component C, all goods must have the same price, say $p_C > 1$. But observe that $p_C = \frac{|G(C)|}{|A(C)|}$ where we use $A(C)$ to denote the agents in C and $G(C)$ to denote the goods in C. This is clearly a rational quantity. Moreover, we can reallocate the goods in $G(C)$ to the agents in $A(C)$ such that every agent gets exactly $\frac{1}{p_C}$ of each good so as to make the allocation rational as well.

Now consider the remainder of the agents and goods. Now let G'' be those goods j with $0 < p_j \leq 1$. Let A'' be the agents which get some positive amount from a good in G''. Each agent in A'' will get allocated entirely to goods in G''. Hence $|A''| = |G''|$ and we can reallocate agents in A'' to goods in G'' via an arbitrary perfect matching.

Finally, consider the remaining goods G''' which have price zero. These goods are partially allocated to agents in A' who have zero utility for them but need them to fill up to one unit of goods. Some goods in G''' might also be allocated to lucky agents A''' which have utility 1 for a zero price good. We can reallocate the agents in A''' integrally to utility 1 goods in G''' and distribute the remaining goods in G''' rationally to the agents in A' that need them.

Overall, we end up with a rational allocation in which all denominators are bounded by n and which is still an optimal solution (2.2). □

2.6 Rationality of α-Slack Equilibria under Dichotomous Utilities

If one wishes to compute exact equilibria (if they exist) instead of approximate ones, clearly a necessary condition is that equilibria are always rational. As noted in Section 2.3.1, with general utilities, both HZ and ADHZ may have only irrational equilibria. On the other hand, with $\{0,1\}$-utilities, there are always rational HZ equilibria. In this section we extend this result to α-slack equilibria in the ADHZ setting.

Fix some ADHZ market with $\{0,1\}$-utilities, rational endowment vectors e, and some rational $\alpha > 0$. Our rationality proof will work in two steps: First we show that as a consequence of Theorem 2.7, there always exists a *special* α-slack equilibrium in which prices are minimal in some sense. Then we will show that the price vector of such a special equilibrium is the unique solution to a system of linear equations with rational coefficients, proving rationality of the prices (and hence there also exists a rational allocation).

Definition 2.10. *An HZ equilibrium* (x, p) *is called* special *if*

1. *there is a good* $j \in G$ *with* $p_j = 0$*, and*
2. *for every price* $\rho > 0$ *in* p*, there is an agent* i *whose cheapest liked goods have price* ρ *and whose budget is at most* ρ*, i.e.,* $\rho = \min\{p_j \mid u_{ij} = 1,\ j \in G\}$ *and* $b_i \leq \rho$*.*

Lemma 2.4. *The algorithm described in Theorem 2.7 always computes a special equilibrium.*

Proof. We will show that at any point in the algorithm, that if there is some good of price $\rho > 0$, then there is some $i \in A_1$ such that i's cheapest desirable goods have price ρ and $\beta_i = b_i$. Note that this property holds at the beginning of the algorithm since the prices are set to the minimum budget of an agent in A_1. Furthermore, as prices increase, the number of agents $i \in A_1$ with $\beta_i = b_i$ can only increase.

So the only way in which this property could be lost is at the points where prices are frozen and the remaining prices are increased, thus decreasing the number of cheapest desirable goods for some agents. Let $S \subseteq G_1$ be the set of goods which have been frozen at some point in the algorithm and assume that we have raised prices so that the price ρ of items in $G_1 \setminus S$ is strictly larger than the prices in S. Furthermore, assume for the sake of a contradiction that for all $i \in \Gamma(G_1 \setminus S)$, we have that $\beta_i = \rho < b_i$. But then

$$\sum_{i \in \Gamma(G_1 \setminus S)} \beta_i = \rho \cdot |\Gamma(G_1 \setminus S)| \geq \rho \cdot |G_1 \setminus S| = \sum_{j \in G_1 \setminus S} p_j.$$

This means that $G_1 \setminus S$ would have already been frozen in the algorithm contradicting the fact that ρ is strictly greater than the prices in S. □

Lemma 2.5. *The prices of the HZ equilibrium as computed in Theorem 2.7 depend continuously on the budgets assuming the initial vertex cover is chosen consistently.*

Proof. Let b and b' be two distinct positive budget vectors with $\|b - b'\|_\infty \leq \epsilon$ for some $\epsilon > 0$. Consider running the algorithm on b and b' at the same time, note that initially prices differ by at most ϵ everywhere. Whenever a set S is frozen for the budgets b, all prices in that set must also be frozen for b' soon afterwards since there is at most $n\epsilon$ more budget available (otherwise S would go overtight).

Let p and p' be the prices computed for budgets b and b' respectively. Then we have just observed that $p' \leq p + n\epsilon$ and symmetrically $p \leq p' + n\epsilon$. Thus p depends continuously on b. □

Theorem 2.11. *There exists a special α-slack equilibrium.*

Proof. Let $P := \{p \in \mathbb{R}^G_{\geq 0} \mid \sum_{j \in G} p_j \leq n\}$ be the set of feasible price vectors. Given some $p \in P$, define $f(p)$ to be the prices output by the algorithm from Theorem 2.7 when applied to the budgets

$$b_i := \alpha + (1-\alpha)\sum_{j\in G} p_j e_{ij}$$

for all $i \in A$.

Clearly, f maps P into P and by Lemma 2.5, f is continuous. So by Brouwer's fixed point theorem, it has a fixed point $p^* \in P$. But by definition of f, this fixed point yields an α-slack equilibrium and by Lemma 2.4, this equilibrium is special. □

Lemma 2.6. *Special α-slack equilibria have rational prices.*

Proof. Let (x, p) be a special α-slack equilibrium. Let $0 = \rho_1 \leq \cdots \leq \rho_k$ be the distinct prices in p. For each $\rho_l > 0$, we let $A_m(\rho_l) \subseteq A(\rho_l)$ be those agents whose budget is at most ρ_l and we let $A_s(\rho_l) \subseteq A(\rho_l)$ be the remaining agents whose budget is more than ρ_l.

Since (x, p) is an α-slack equilibrium, we then have that

$$\sum_{i\in A_m(\rho_l)} \left(\alpha + (1-\alpha)\sum_{l'=2}^{k} \rho_{l'} \sum_{j\in G(\rho_{l'})} e_{ij}\right) + \rho_l |A_s(\rho_l)| = \rho_l |G(\rho_l)|. \tag{2.3}$$

where $A_m(\rho_l) \neq \emptyset$ since (x, p) is a special equilibrium. Together with $\rho_1 = 0$, this gives us a system of linear equations with rational coefficients that $(\rho_1, \ldots, \rho_k)$ is a solution to.

Finally, let us show that this system has unique solutions. To see this let there be some other solution vector $(\rho'_1, \ldots, \rho'_k)$. Assume without loss of generality that there is some l with $\rho_l > \rho'_l$ as otherwise we can swap ρ and ρ'. Let l^* be the index maximizing $\frac{\rho_l}{\rho'_l}$ and consider constraint (2.3) for this l^*. But now assuming that ρ' satisfies this constraint, ρ_l cannot satisfy it since, compared to ρ', the right-hand side increases by a factor of $\frac{\rho_l}{\rho'_l}$ whereas the left-hand side increases by strictly less due to the presence of $\sum_{i\in A_m(\rho_{l^*})} \alpha > 0$. □

Theorem 2.12. *There exists a rational α-slack equilibrium.*

Proof. By Theorem 2.11, there always exists a special α-slack equilibrium and by Lemma 2.6, this equilibrium must have rational prices. To get a rational allocation, one can obtain an allocation via a flow network in each price class as shown in Lemma 2.2. Since max-flows in networks with rational weights can always be chosen to be rational, the theorem follows. □

2.7 Discussion

In this chapter, we defined an ϵ-approximate ADHZ model for one-sided matching markets with endowments. We showed that ϵ-approximate ADHZ equilibrium always exists for every $\epsilon > 0$. We strengthened the non-existence of ADHZ equilibrium for the case when the demand graph is not strongly connected and agents have dichotomous utilities. We derived a novel combinatorial polynomial-time algorithm for computing an ϵ-ADHZ equilibrium under dichotomous utilities. Finally, we presented a rational convex program (RCP) for the HZ model under dichotomous utilities, which also implies that the problem is polynomial-time solvable.

Since finding an HZ equilibrium is PPAD-complete [122, 28], there is little hope that we can efficiently find ϵ-approximate ADHZ for general utilities except perhaps for constant ϵ. However, this would generalize a similar problem to find ϵ-approximate HZ equilibria for constant ϵ which also remains open.

One interesting problem is to directly compute an exact α-slack equilibrium with dichotomous utilities. By Section 2.6, we know that there is a rational α-slack equilibrium in this case and it stands to reason that it might be computable in polynomial time.

Finally, it is also interesting to consider what happens when $\alpha \to 0$. By compactness, there exists some limit point of α-slack equilibria as α tends to 0 and this limit point has some interesting properties. See Section 3.3 in the next chapter for an argument along these lines.

Can we compute it?

Chapter 3

The Quest for Envy-Freeness and Pareto-Optimality

3.1 Introduction

We will now turn back to the more classic setting of a cardinal-utility matching market without endowments in which every agent is treated equally. Our goal in this chapter will be to analyze the two most important desirable properties: fairness (as measured by envy-freeness) and efficiency (as measured by Pareto-optimality). In particular, when can we find solutions that meet both criteria? This chapter is based on the paper "Cardinal-Utility Matching Markets: The Quest for Envy-Freeness, Pareto-Optimality, and Efficient Computability" which was joint work with Vijay V. Vazirani [116].

As mentioned in the background chapter, the general objective in a matching market is to find allocations with desirable properties in polynomial time. When it comes to cardinal utilities, the most notable mechanism is that of Hylland and Zeckhauser [75], which is based on a pricing approach. It results in allocations which are (ex-ante) envy-free (EF), i.e.

each agent prefers their own lottery to that of any other agent, and (ex-ante) Pareto-optimal (PO), i.e. it is impossible to improve one agent's expected utility without diminishing the expected utility of some other agent. Later, He et al. [69] showed that the HZ mechanism is also incentive compatible in the large. Note that there is no mechanism which is EF, PO, and also incentive compatible in the traditional sense, i.e. without the "in the large" restriction [128].

However, a core issue with the approach is that *computing* an HZ equilibrium is hard in theory and in practice. Vazirani and Yanakakis [122] 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 the complexity class PPAD and conjectured that it is PPAD-complete. This conjecture was confirmed by Chen et al. [28] who proved the corresponding hardness result.

This motivates the search for alternative mechanisms which can achieve some or all of the desirable properties of HZ while also being implementable in polynomial time. In this chapter, we pose the key question: is it possible to compute an envy-free and Pareto-optimal lottery in a one-sided cardinal-utility matching market in polynomial time using some different mechanism?

Beyond one-sided matching markets and HZ, there are also *two-sided matching markets* in which we are matching agents to other agents. Prominent examples are the kidney donor market or the problem of assigning students to schools if the schools have preferences over the students, e.g. via aptitude scores. Except for a few highly restricted special cases [18, 111], it was not known whether envy-free and Pareto-optimal lotteries even exist in this setting.

3.1.1 Our Contributions

One-Sided Matching Markets

Our most significant contribution is that we resolve the question about the complexity of finding EF+PO allocations by showing that the problem is PPAD-hard. Together with a recent result by Caragiannis et al. [24] this shows that the problem is PPAD-complete.

Theorem (Section 3.2.2)**.** *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; it is inspired by the fact that HZ allocations and EF+PO allocations coincide in certain *continuum markets* involving infinitely many agents and goods [6]. The key idea is to take an HZ instance and add both 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 [102].

On the way we will also give a simple polyhedral proof that there are always rational EF+PO allocations. This of course follows from the PPAD membership proof by Caragiannis et al. [24] 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 [70] satisfy an approximate notion of envy-freeness and incentive compatibility.

Theorem (Section 3.2.3). *The Nash bargaining solution for one-sided cardinal utility matching markets is 2-approximately envy-free and 2-approximately incentive compatible.*

Together with the algorithms given by Panageas et al. [108], this results in a mechanism that is $(2+\epsilon)$-envy-free, $(2+\epsilon)$-incentive compatible and Pareto optimal in polynomial time. We remark that HZ is $(1+\epsilon)$-incentive compatible, however this requires the assumption of a "large market" in which every agent has many copies [69]. Hence, this establishes the Nash-bargaining-based mechanism as a more practical, alternative mechanism for one-sided cardinal-utility matching markets.

Two-Sided Matching Markets

For two-sided markets, the only cases in which it was previously known that EF+PO allocations exist is 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 3.3.2, we provide counterexamples that show that both of these conditions are necessary: if agents have $\{0,1,2\}$ utilities or if they have 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 + weakly PO allocations via a limiting argument, an equilibrium notion introduced by Manjunath [95], and similar polyhedral techniques as we used for one-sided markets.

Theorem (Section 3.3.3). *In any two-sided cardinal-utility matching market, a rational JEF + weakly PO allocation always exists.*

The Nash-bargaining-based approach from [70] and the efficient algorithms from [108] also extend to two-sided markets. However, in Section 3.3.4, we give a counterexample to show that, in a Nash bargaining solution, agents can have $\theta(n)$-factor justified envy towards other agents.

3.1.2 Related Work

Our work builds on the existing literature surrounding the Hylland Zeckhauser mechanism [75] and the complexity of computing HZ equilibria. Alaei et al. [3] give an algorithm to compute HZ equilibria which is based on the algebraic cell decomposition technique [13]. However, this algorithm needs to enumerate at least n^{5n^2} cells. Garg et al. [61] improve on this with an algorithm that requires solving on the order of n^n many linear programs. Both algorithms are highly impractical, even for small instances.

Vazirani and Yannakakis [122] give a polynomial time algorithm that computes HZ equilibria for $\{0, 1\}$ utilities. 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. [28] 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 [52, 118], 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) [22] or envy-freeness up to any good (EFX) [25] are studied instead.

Cole and Tao [31] 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. [50], Caragiannis et al. [24] managed to show 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 thesis.

Markets with a continuum of agents were introduced by Aumann [8]. Zhou [129] 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 [6] show a similar equivalence for HZ. On the other hand, Miralles and Pycia [102] 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 [70] 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 [119] who used this approach for the linear Arrow Debreu market. We will introduce the Nash-bargaining-based mechanism in more detail in Section 3.2.3.

Panageas et al. [108] 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 [11] 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. We cover these results in Chapter 4.

For two-sided markets, there have been several attempts at extending the equilibrium notion of Hylland and Zeckhauser. Manjunath [95] as well as Echenique et al. [42] introduce HZ-like 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 [18] for bipartite markets and Roth et al. [111] for non-bipartite markets. A polynomial time algorithm to compute such equilibria and therefore EF+PO allocations was later given by Li et al. [92].

Beyond this, two-sided markets have been mostly studied under ordinal preferences where stable matching, as introduced by Gale and Shapley [56], is the dominant solution concept. A notable exception is the work by Caragiannis et al. [23] who study the problem of finding a fractional stable matching under cardinal utilities that (approximately) maximizes social welfare.

3.2 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. As usual in this thesis, we will assume cardinal utilities, i.e. 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 [75] which we introduced in Section 1.3.

Recall that the key insight of the HZ mechanism consists in giving agents equal amounts of fake currency and finding a so-called HZ equilibrium. We restate the definition of an HZ equilibrium below, assuming all budgets are 1.

Definition 3.1. *Given A, G, u, an* HZ market equilibrium *consists of an* allocation $(x_{ij})_{i\in A,j\in G}$ *and non-negative* prices $(p_j)_{j\in G}$ *such that*

1. *x is a fractional perfect matching.*
2. *No agent overspends, i.e. $p \cdot x_i \leq 1$ for each $i \in A$.*
3. *Each agent i gets an* optimal bundle*, i.e. x_i maximizes $u_i \cdot x_i$ under the constraint that $p \cdot x_i \leq 1$ and $\sum_{j\in G} x_{ij} = 1$.*
4. *Each agent i gets a* cheapest bundle*, i.e. x_i minimizes $p \cdot x_i$ among all bundles with utility at least $u_i \cdot x_i$ and $\sum_{j\in G} x_{ij} = 1$.*

As mentioned in Section 1.3, HZ equilibria always exist and are both Pareto-optimal and envy-free.

Theorem 3.1 (Hylland, Zeckhauser [75])**.** *An HZ equilibrium always exists. Moreover, if (x, p) is an HZ equilibrium, then x is envy-free and Pareto-optimal.*

3.2.1 Rationality

Unfortunately, there are instances on which the unique Hylland-Zeckhauser equilibrium requires an irrational allocation and prices [122]. In contrast, we will show in this section that there are always EF+PO allocations which are rational. This of course follows from the PPAD membership proof given by Caragiannis et al. [24]. However, our argument is simpler and introduces some basic facts that 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_{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 [126]. 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 3.1. *$x^\star \in P_{\text{PM}}$ 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, α can be computed in polynomial time.*

Proof. 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 α 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:

$$\min \quad \sum_{i \in A} a_i u_i \cdot x_i^\star + \sum_{i \in A} q_i + \sum_{j \in G} p_j \tag{3.1a}$$

$$\text{s.t.} \quad a_i u_{ij} + q_i + p_j \geq u_{ij} \quad \forall i \in A, j \in G, \tag{3.1b}$$

$$a_i \leq 0 \quad \forall i \in A. \tag{3.1c}$$

Then by strong duality

$$\sum_{i \in A} u_i \cdot x_i^\star = \sum_{i \in A} a_i u_i \cdot x_i^\star + \sum_{i \in A} q_i + \sum_{j \in G} p_j. \tag{3.2}$$

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

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

But this follows immediately from the fact that (q, p) is an optimal dual solution to this LP: (3.1b) implies feasibility and (3.2) implies optimality. □

Lemma 3.1 characterizes the Pareto-optimal allocations. Moreover, the envy-free allocations themselves form the polytope P_{EF} shown below.

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

Theorem 3.2. *There is always an EF+PO allocation which is a vertex of P_{EF} and is thus rational.*

Proof. We know that at least one EF+PO allocation $x^\star$ exists since the HZ equilibrium allocation is both envy-free and Pareto-optimal. By Lemma 3.1, $x^\star$ maximizes $\phi(x) := \sum_{i \in A} \alpha_i u_i \cdot x_i$ over P_{PM} for some strictly positive α vector.

Now consider the linear program $\max\{\phi(x) \mid x \in P_{\text{EF}}\}$. P_{EF} is a polytope so let x be a vertex solution to this LP. Clearly x is envy-free since $x \in P_{\text{EF}}$. But since $x^\star$ is also in

P_{EF} and $P_{\text{EF}} \subseteq P_{\text{PM}}$, we know that $\phi(x) = \phi(x^\star)$. Therefore, by the other direction of Lemma 3.1, x is Pareto-optimal. $\square$

3.2.2 PPAD-Hardness of Computing EF+PO

We now turn to our main result:

Theorem 3.3. *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 3.4 (Chen, Chen, Peng, Yannakakis 2022 [28])**.** *For any $c > 0$, the problem of finding an ϵ-approximate HZ equilibrium is PPAD-hard for $\epsilon \leq 1/n^c$.*

There are various reasonable notions of ϵ-approximate HZ equilibria which are polynomially equivalent. We will use the following definition. [1]

Definition 3.2. *An assignment $(x_{ij})_{i\in A, j\in G}$ together with non-negative prices $(p_j)_{j\in G}$ are an* ϵ-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 ϵ worse than an optimal bundle, i.e.*

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

[1] See [122] for a proof that this notion is indeed equivalent to the one that Chen et al. use.

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, we will find prices p and budgets b that make x into a competitive equilibrium using a version of the Second Welfare Theorem.

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 which is why 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 3.3. *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. Note that this 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 with the utilities of agent i and slowly increasing / decreasing the utility for type t_1 goods in steps of ϵ 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 3.1.

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 $k \geq \frac{n^3}{\epsilon}$.

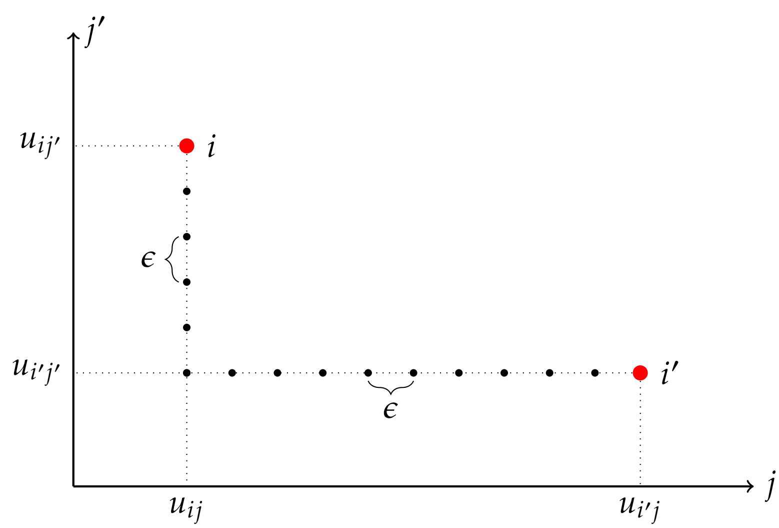

Figure 3.1: For each pair of agents i and i' (large red dots) we add *interpolating agents* (small black dots) to transition between the utility vector u_i and $u_{i'}$ in small steps. This is done coordinate-wise and this figure depicts an example with only two goods j and j'.

Lemma 3.2. *Let $n' = |A'| = |G'|$ be the number of agents and 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 3.1, there exist positive α_i for all $i \in A'$ such that x solves

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

Moreover, we can find such α_i in polynomial time since we obtained them using an LP in the proof of Lemma 3.1. 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 $b_i := \alpha_i u_i \cdot x_i - q_i$ 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 3.3, x really is a competitive equilibrium with prices p and budgets b.

Lemma 3.3. *For every agent i, we have that $b_i \geq 0$ and x_i is an optimal 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, in a quantifiable way, that no agent is satiated.

Lemma 3.4. *If j and j' are goods of the same type, then $p_j = p_{j'}$.*

Proof. Note that every good is fully matched. But if the prices were different, then any agent matched to the more expensive good would be violating Lemma 3.3; they would switch to the cheaper good of the same type. □

Lemma 3.5. *For any non-dummy agent i, $u_i \cdot x_i \leq 1.6$. In particular i is not satiated.*

Proof. If this were not the case, i would need to get at least 0.6 units of an awesome good. But then any other non-dummy agent must get at least 0.1 units of an awesome good by envy-freeness. Since there are many more non-dummy agents than awesome goods, this is a contradiction. □

Lemma 3.6. *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 any 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 3.5, the price of said awesome good must be positive and so must the agent's budget. □

In particular, we can rescale all α, 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 3.7. *If i and i' are agents of the same type, then $b_i = b_{i'}$.*

Proof. Note that by Lemma 3.5, 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 3.8 and 3.9 below.

Lemma 3.8. *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 3.3 and Lemma 3.5. Combining 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 3.9. *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 ϵ. Then $|b_i - b_{i'}| \leq 5n^2\epsilon$.*

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

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

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

Now we can compute

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

and via symmetry and Lemma 3.8 we conclude $|b_i - b_{i'}| \leq \epsilon \max\{\alpha_i, \alpha_{i'}\} \leq 5n^2\epsilon$. □

Lemma 3.9 is enough to show that the difference in budgets between "close" agents tends to zero for an inverse-polynomial ϵ. 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 3.2.

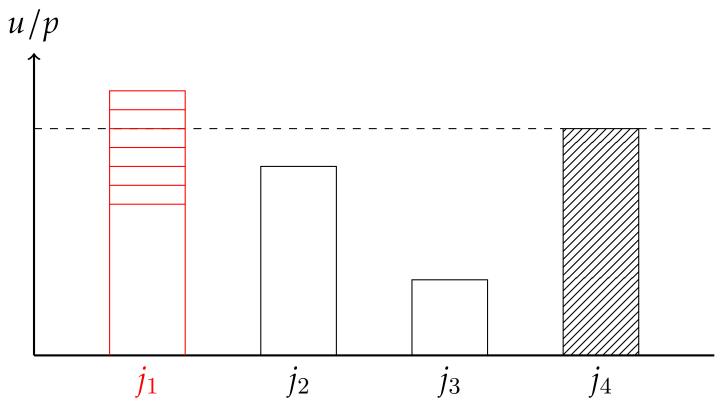

Figure 3.2: Shown is an agent who is interested in goods j_1 to j_4 which are plotted by their bang per buck. If we change only the utility of good j_1 (red) and leave the rest the same, there are only three possible sets of maximum bang per buck goods: $\{j_4\}$, $\{j_1, j_4\}$, and $\{j_1\}$. So along 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.

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.

Definition 3.4. *For any agent i, 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 optimal solution to*

$$\begin{aligned} \max \quad & u_i \cdot y \\ \text{s.t.} \quad & \sum_{j \in G'} y \leq 1, \\ & p \cdot y \leq t, \\ & y \geq 0. \end{aligned} \tag{3.3}$$

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

Lemma 3.10. *Let i, i' 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 (3.3) 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 optimal solution of this LP. To see this, consider the dual as well:

$$\begin{aligned} \min \quad & \mu + \rho b_i \\ \text{s.t.} \quad & \mu + p_j \rho \geq u_{ij}, \\ & \mu, \rho \geq 0. \end{aligned} \tag{3.4}$$

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 μ and ρ, satisfies the complementary slackness conditions and is therefore optimal.

Finally, since no agent is satiated (Lemma 3.5), increasing the budget always increases the optimal value of the LP, implying that $u_i \cdot x_i < u_i \cdot x_{i'}$. This contradicts envy-freeness. □

Lemma 3.11. *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 optimal solution y to (3.3). By complementary slackness these are the goods for which the corresponding dual constraint is tight in the dual (3.4).

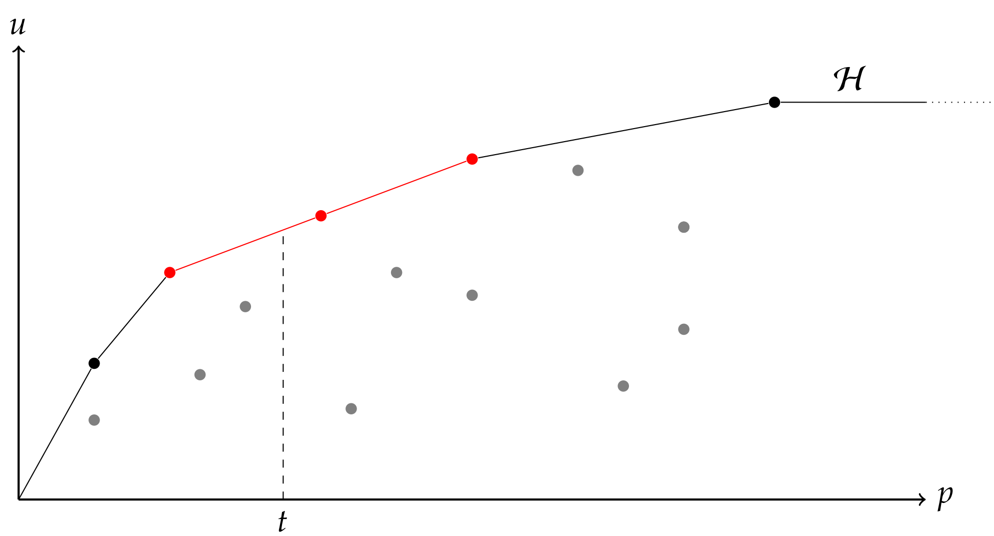

Figure 3.3: Depicted is $\mathcal{H}$ and its relationship to optimal bundles. Each point represents a good or collection of goods with identical price and utility. Gray points are dominated and will never be part of an optimal bundle. Points on $\mathcal{H}$ can be part of an optimal bundle depending 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 $\mathcal{H}$.

And together with what we already know from complementary slackness, this gives a nice geometric characterization of θ_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 optimal 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 3.3.

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 θ, the only thing we need to know to uniquely determine θ 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 3.4. □

Lemma 3.12. *Let i, i' be two non-dummy agents. Then $|b_i - b_{i'}| \leq 5\epsilon n^4$.*

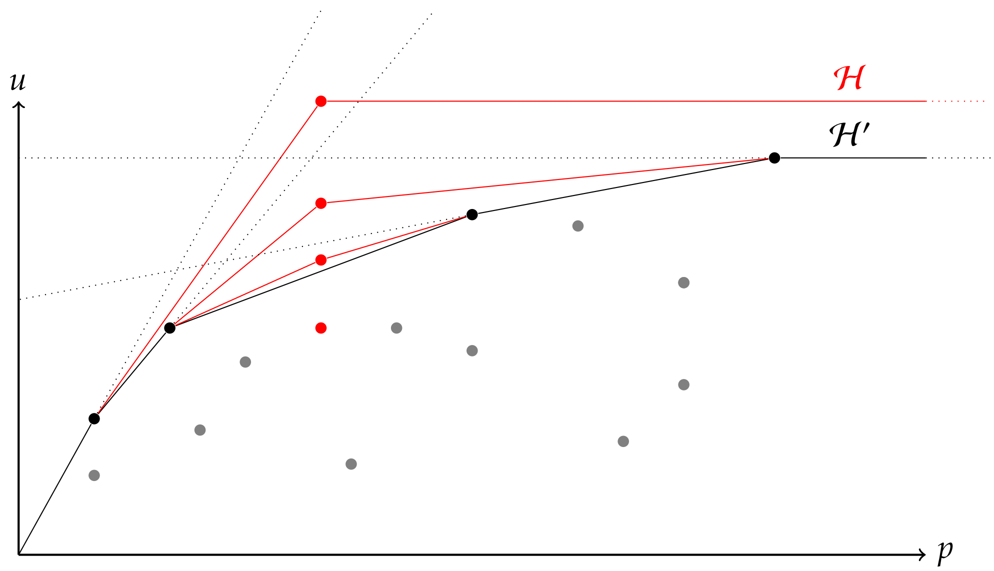

Figure 3.4: Shown are several convex hulls $\mathcal{H}$ (red) as the red good's utility is changed. Note that the structure of $\mathcal{H}$ only changes when we cross one of the bounding lines of $\mathcal{H}'$, the convex hull without the red good.

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 3.11, inside each group there are at most $2n + 1$ different θ functions. By Lemma 3.10, the budgets of agents who have identical θ functions must be identical. And so there are at most $2n$ opportunities for θ to change inside each group, totaling to $2n^2$ changes overall. Each of these changes in θ corresponds to two agents that differ in their utilities by at most ϵ on one good, thus Lemma 3.9 applies and we get $|b_i - b_{i'}| \leq 2n^2 \cdot 5\epsilon n^2$. □

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 agents, and awesome goods are simply dropped.

Theorem 3.5. *If $\epsilon \leq \frac{1}{5n^5}$, then $(\hat{x}, \hat{p})$ is a $\frac{3}{n}$-approximate HZ equilibrium in 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 optimal bundle. For that, let y be an optimal 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 3.12, 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 3.3. 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 3.3. If we choose $\epsilon = \frac{1}{5n^5}$ and $k = 5n^8$, then the constructed instance has at most $10n^9$ agents by Lemma 3.2. 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 3.4, the latter problem is PPAD-hard. □

Lastly, we remark that Theorem 3.3 can be slightly strengthed to show hardness of computing *approximately* envy-free and Pareto-optimal solutions with inverse polynomial ϵ. Lemmas 3.9 and 3.10 require minor modifications for the proof to go through.

3.2.3 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 [124] studied the problem of two or more agents bargaining over a common outcome, for example how they should split up certain goods amongst themselves. He showed that there is a unique point that satisfies certain axioms and moreover that this point is characterized as maximizing the product of the agents' utilities, i.e. the Nash social welfare. See Section 1.4 for more details.

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

$$
\begin{aligned}
\max \quad & \prod_{i\in A} u_i \cdot x_i \\
\text{s.t.} \quad & \sum_{j\in G} x_{ij} = 1 \quad \forall i \in A, \\
& \sum_{i\in A} x_{ij} = 1 \quad \forall j \in G, \\
& x_{ij} \geq 0 \quad \forall i \in A, j \in G.
\end{aligned}
\tag{3.5}
$$

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 [70] 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 3.5. *An allocation* $(x_{ij})_{i\in A, j\in G}$ *is* α-approximately envy-free *or just* α*-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* α*.*

Theorem 3.6. *Let* x *be an optimal solution to* (3.5)*. 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 ϵ-fraction of the bundle that i gets with an ϵ-fraction of 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 ϵ 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 ϵ the product of the agents' utilities is increasing. This contradicts the fact that x is an optimal solution to (3.5). □

We remark that this bound is tight since Aziz and Brown [11] give an instance in which an agent envies another agent by a factor of 2. See Figure 3.5.

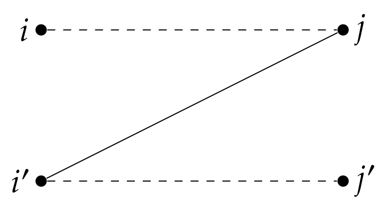

Figure 3.5: Shown is an example instance which demonstrates that 2-EF is tight for Nash bargaining. Dashed edges have utility 1, solid edges have utility 2, and missing edges have utility 0. Clearly both agents prefer j to j'. A simple calculation shows that in the Nash bargaining solution, i will get all of j and thus i' will envy i by a factor of 2.

Definition 3.6. *Consider some mechanism M which maps utility profiles $(u_i)_{i \in A}$ to allocations $(x_i)_{i \in A}$. Then M is called α-*incentive compatible *or just α-IC if, whenever utilities u and $\hat{u}$ differ only on agent i, said agent does not improve by more than a factor of α 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 α by misreporting their utilities.*

Theorem 3.7. *Any mechanism which maps u to some maximizer of* (3.5) *is 2-IC.*

Proof. The proof of this result is quite similar to the proof of Theorem 3.6. 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 (3.5) under utilities u and y a maximizer of (3.5) 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 ϵ, 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 ϵ 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. $\square$

This bound is also tight as shown by the following family of instances. See Figure 3.6 for an example.

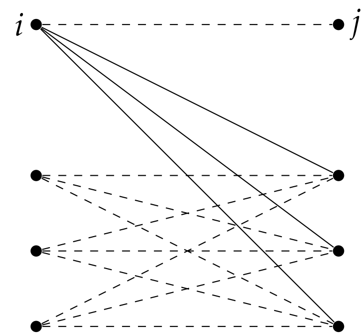

Figure 3.6: Shown is an example instance from the proof of Theorem 3.8 with $n = 4$. Dashed edges have utility 1, solid edges utility 2, and missing edges have utility 0. Agent i will be fully allocated to good j by Nash bargaining even though they would prefer the "desirable" goods. However, agent i can misrepresent their utilities to look like the other agents therefore get a significant fraction of the desirable goods.

Theorem 3.8. *Any mechanism which maps u to some maximizer of* (3.5) *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 3.6 for $n = 4$.

Let x be the amount that agent 1 is matched to the desirable goods. By symmetry[2], 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

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

[2]One can easily see that in optimal solutions to (3.5), agents with equal utility vectors must have the same overall utility. Otherwise the product of their utilities can be improved.

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. [108] give simple, practical algorithms for computing $(1+\epsilon)$-approximate[3] Nash bargaining points in $O(\mathrm{poly}(n, 1/\epsilon))$ time and so we get the following corollary.

Corollary 3.1. *There is a $(2+\epsilon)$-EF, PO, $(2+\epsilon)$-IC mechanism for one-sided cardinal-utility matching markets which runs in $O(\mathrm{poly}(n, 1/\epsilon))$ time.*

Finally, note that both [70] and [108] 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 chapter.

3.3 Two-Sided Matching Markets

A second interesting class of matching markets is that of two-sided markets in which instead of matching goods to agents we match agents to other agents. 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.

[3] Approximate in the sense that all utilities are within $(1+\epsilon)$ of the Nash bargaining point.

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 $\sum_{i \in A} w_{ji} x_{ij}$.

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. For more details on non-bipartite markets, see Chapter 4.

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

Definition 3.7. *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 3.8. *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. See [43] for an interesting viewpoint on symmetric utilities.

Bogomolnaia and Moulin [18] 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 [122] can be adapted for this setting. This result was extended to the non-bipartite case by Roth et al. [111] who proved existence and Li et al. [92] who gave a polynomial time algorithm.

3.3.1 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 3.2.1 so we will not restate them here.

Lemma 3.13. *$x^\star \in P_{\mathrm{PM}}$ 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, α and β can be computed in polynomial time.*

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

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

Theorem 3.9. *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_{2EF} and is thus rational.*

We will also need the following characterization in Section 3.3.3. Note that an allocation is weakly Pareto-optimal if there is no other allocation that improves on the utility of *every* agent.

Lemma 3.14. *$x^\star \in P_{\mathrm{PM}}$ is weakly 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 $x \in P_{\mathrm{PM}}$.*

Proof. The proof is quite similar to the proof of Lemma 3.1. Clearly if $x^\star$ maximizes $\phi(x)$, then it is a weakly Pareto-optimal allocation since any strictly Pareto-better allocation x would satisfy $\phi(x) > \phi(x^\star)$ since at least one α_i or β_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 optimal solution (α, β, 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 optimal 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 optimal dual solution. □

3.3.2 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 3.9, 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

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

In most instances, this finds an EF+PO allocation relatively quickly. By enumerating small instaces 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 $P_{2\mathrm{EF}}$ 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 $P_{2\mathrm{EF}}$ generally has.

Theorem 3.10. *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 3.7a 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 direciton, 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. □

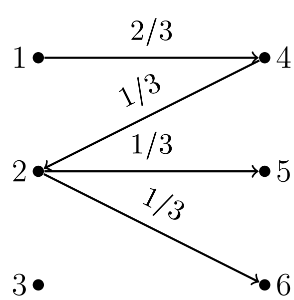

(a) Each arrow represents a utility 1-edge from one side and utility 0 from the other.

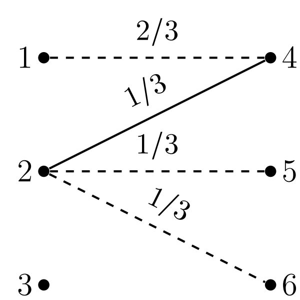

(b) Dashed edges have utility 1, whereas solid edges have utility 2.

Figure 3.7: Shown are counterexamples for $\{0,1\}$ asymetric (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).

Theorem 3.11. *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 3.7b 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.

Next, note that $x_{14} = \frac{1}{3}$. Again, we must have $x_{14} \geq \frac{1}{3}$ since otherwise agent 1 would enyv 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. □

3.3.3 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 3.9. *In an allocation x, agent $i \in A$ has* strongly justified envy *towards $i' \in A$ if $w_{ji} \geq w_{ji'}$ for all $j \in B$ and $u_i \cdot x_i < u_i \cdot x_{i'}$. Strongly justified envy is defined symmetrically for agents in B. An allocation in which there is no strongly justified envy is said to be* weakly justified envy-free (weakly JEF).

Weakly 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 also define a slightly stronger notion of justified envy-freeness.

Definition 3.10. *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, strongly justified envy implies justified envy and therefore JEF implies weakly 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 3.12. *There always exists a rational allocation which is JEF and weakly PO.*

The proof uses a limit argument based on an equilibrium notion introduced by Manjunath [95]. 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 3.11. *A* 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*

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

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

Theorem 3.13 (Manjunath [95])**.** *As long as every agent has a positive endowment of money (i.e. $\omega_i > 0$), a DIP equilibrium always exists.*

Theorem 3.14 (Manjunath [95])**.** *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 3.13 directly but in order to be self-contained, we will give a short proof which uses the existence of DIP equilibria as a black box.

Lemma 3.15. *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. □

Finally, we can use the existence of DIP equilibria to show that JEF and weakly PO allocations exist through another limiting argument.

LEMMA 3.16. *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 $m_i \le \sum_{k \in A \cup B} m_k = \sum_{k \in A \cup B} \omega_k = \epsilon$ and this finishes the proof. By symmetry the same holds for all pairs of agents in B. □

Proof of Theorem 3.12. First, let us show that a JEF and weakly PO allocation exists. By Lemma 3.16, we can pick a sequence $x^{(k)}$ of Pareto-optimal allocations such that

$$u_i \cdot x_i^{(k)} \ge \sum_{\substack{j \in B \\ w_{ji} \ge w_{ji'}}} u_{ij} x_{i'j}^{(k)} - \epsilon_k \tag{3.6}$$

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 weakly Pareto-optimal allocation. Furthermore, it is easy to see that $x^\star$ is justified envy-free by taking the limit over (3.6).

Finally, we can use a similar argument as in the proof of Theorem 3.2 to show that a rational JEF + weakly PO allocation exists as well. Simply pick α, β according to Lemma 3.14 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. □

3.3.4 Justified Envy for Nash Bargaining

As shown in Section 3.2.3, 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 in [108] that shows this not to be the case.

Theorem 3.15. *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'. Agent j has utility 0 for all agents in A. 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 3.8.

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 $u_i \cdot x_i = \frac{1}{n+1}$ but $u_i \cdot x_{i'} = \frac{n}{n+1}$. □

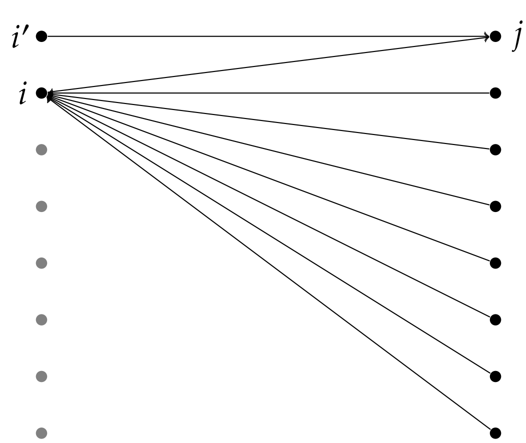

Figure 3.8: Shown is an instance ($n = 8$) with justified envy for Nash bargaining. Agents i and i' compete for j but all agents in $B \setminus \{j\}$ want i so i gets only a small fraction of j. The gray agents are dummy agents and have identical utilities for all agents in B.

3.4 Discussion

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

- Is there a polynomial time algorithm to find α-approximately JEF+PO lotteries in two-sided markets, for any constant α?
- Is Nash bargaining the best we can do for one-sided markets or is there a way to compute an α-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?

Chapter 4

Efficient Algorithms for Nash Bargaining

4.1 Introduction

In the previous chapter we have made a strong case that Nash bargaining is an attractive alternative to HZ since it achieves desirable game-theoretic properties (2-approximate envy-freeness, 2-approximate incentive-compatibility, and Pareto-optimality) while also being computationally tractable. In this chapter we will expand on the last part; we will show that Nash bargaining does not just have polynomial time algorithms in theory by appealing to convex programming theory, but we will actually present simple, practical algorithms with proven running time guarantees. This chapter is based on the paper "Time-Efficient Algorithms for Nash-Bargaining-Based Matching Market Models" which was joint work with Ioannis Panageas and Vijay V. Vazirani [108].

The study of Nash bargaining as an alternative mechanism for cardinal-utility matching markets was initiated by Hosseini and Vazirani [70] based on earlier work by Vazirani [119]. They identified two key shortcomings of the classic pricing-based mechanism due to Hylland and Zeckhauser [75]:

1. There is no known algorithm for HZ which is polynomial time or at least efficient in practice. Note that it was later shown that finding an approximate HZ equilibrium is PPAD-hard [28] and, as we discussed in Chapter 3, even finding *any* EF+PO allocation is already PPAD.

2. The HZ mechanism is difficult to extend to more complex matching market models. We showed in Chapter 2, that there is a reasonable extension to endowments. However, beyond that, the extensions that have been proposed (e.g. by He et al. [69]) tend to require substantial concessions such as individualized prices.

For this reason, Hosseini and Vazirani proposed a different approach, namely Nash bargaining. See Section 1.4 for an overview of Nash bargaining and its application to matching markets. They point out that Nash bargaining is efficient in theory due to being described via a convex program (in constrast, there is no mathematical program that models HZ). For example, this allows for a polynomial time algorithm via the ellipsoid method [66, 123]. Moreover, Nash bargaining is far more flexibile to different settings such as two-sided markets or markets with additional constraints.

Finally, Hosseini and Vazirani perform various computational experiments in which they show that convex programming solvers can be used to solve large matching markets efficiently in practice. They show that markets with up to 10,000 agents and goods are still readily solvable on a laptop. However, they did not attempt to give a worst-case running time bound of their implementation.

While the computability of Nash bargaining-based mechanisms via convex programming is also a significant advantage over HZ, we will further expand on this advantage in this chapter. Our goal will be to show that the simple structure of both the constraints and the objective function allow us to provide simpler algorithms with provable running time guarantees.

4.1.1 Our Contributions

Modeling

In this chapter we will consider three types of cardinal-utility matching markets:

- In one-sided markets, we are matching agents to goods. The agents have utilities for the goods but the goods do not have any preferences for agents.
- In two-sided markets, we are matching agents of one type to agents of another type. Agents have utilities for agents of the other type. See also Section 3.3.
- In non-bipartite markets, we are matching agents to other agents just as in the two-sided markets but we no longer require two distinct types of agents. An example of such a market would be a ride-sharing or roommate market.

One-sided and two-sided Nash-bargaining models were introduced by Hosseini and Vazirani [70]. Non-bipartite markets however are novel to this work. We will also allow endowments in our models; see also Chapter 2. All models are introduced in Section 7.2. Note that in the full version of the paper that this chapter is based on [108], we also consider models with *separable piecewise-linear concave (SPLC)* utilities. This discussion and the related results have been omitted from this chapter for the sake of brevity.

We show that all models satisfy a kind of approximate *equal-share fairness*, though with varying approximation factors. This fact is useful to prove convergence of conditional gradient descent algorithms. Indeed, Hosseini and Vazirani made use of this observation to further speed up their practical algorithms.

Multiplicative Weights Update (MWU)

In Section 4.3, we will introduce our first practical algorithm for Nash-bargaining-based matching market models. The algorithm is based on a multiplicative weights approach and finds an ϵ-approximate Nash bargaining solution in $O(\frac{n^3 \log n}{\epsilon^2})$ time. It yields both allocations and the corresponding dual variables.

The main idea underlying our approach is to reduce the problem of finding an optimal Nash bargaining solution to the problem of finding any feasible solution to a related convex program. A similar idea was used by Fleischer et al. [51] to find equilibria in non-matching markets. However, in our case, the approach is made more complicated by the additional constraints present in our setting.

The MWU algorithm can only deal with one-sided markets though it does allow for endowments and even SPLC utilities. Its analysis depends on a clever use of the KKT conditions of the Nash bargaining convex program together with a rescaling trick. Besides these ideas, the proof follows a relatively standard potential function argument that is common in the analysis of MWU algorithms.

Conditional Gradient Descent (CGD)

In Section 4.4, we will introduce another practical algorithm for Nash-bargaining-based matching markets which is based on conditional gradient descent. It finds ϵ-approximate solutions in $O(\frac{n^3 \kappa^2}{\epsilon})$ time where κ measures the maximum gap between the smallest and largest utilities for any agent. This algorithm is faster than the MWU algorithm for small ϵ, provided that κ is bounded. Moreover, the CGD algorithm is able to deal with two-sided and non-bipartite markets as well.

In order to apply a gradient descent type algorithm to Nash-bargaining in matching markets, we had to overcome two main challenges:

- The objective function of Nash-bargaining is logarithmic and hence neither Lipschitz nor smooth. In order to guarantee convergence, we modify the objective function with an approach similar ot that used by Gao et al. [58] to find Fisher market equilibria. The idea is to extend the objective function by its quadratic extension below a certain point. This approach relies crucially on an equal-share fairness property that Fisher market equilibria have. While Nash-bargaining in a matching market, does *not* satisfy this property, we show that an approximate version of it holds which is enough to carry out the proof.
- Projecting into the feasible regions of our models which are given by matching or flow polytopes, is computationally challenging. For this reason, we use a conditional gradient approach instead which relies on combinatorial matching and flow algorithms that are very efficient in practice. In the case of SPLC utilities, we also had to add an additional "shifting step" to decrease the dependence on the rather large diameter of the feasible region. This last result is covered in the full version of the paper and has been omitted here for the sake of brevity [108].

4.1.2 Related Results

As stated in the introduction, the first paper to suggest the use of a Nash-bargaining-based mechanism for solving a market model which had traditionally been addressed via pricing was due to Vazirani [119]. It builds on the success of the Eisenberg Gale convex program for the linear Fisher market, which shows the equivalence of Nash bargaining and pricing for that market. For the Arrow Debreu market, the natural extension of the EG program does not capture Arrow Debreu equilibria but rather a version of Nash bargaining applied to an exchange market; see Chapter 1 for details on these concepts. Vazirani hence proposed a Nash-bargaining-based approach to exchange markets and gave a combinatorial polynomial time algorithm for Nash-bargaining in this setting.

The more recent work by Hosseini and Vazirani [70] builds on this idea and extends it to matching markets where pricing and Nash bargaining are distinct concepts just as they are for exchange markets. Together with the PPAD-completeness of HZ [122, 28] and the results that we covered in Chapter 3, this makes for a compelling case that Nash bargaining is a practical alternative to HZ in matching markets.

Our models apply to more general settings than just the classic one-sided cardinal-utility matching markets without endowments. Related works on more general matching markets largely focus on extensions of HZ [22, 69, 88, 97]. A notable example is the work by Echenique et al. [42] which defined the notion of an α-slack equilibrium to deal with exchange settings. We build on this in Chapter 2 and define ϵ-approximate ADHZ equilibria which we are able to compute for dichotomous utilities.

Aside from the aforementioned work by Vazirani [119], our results are closely related to the computation of Fisher market equilibria. Several algorithms have been developed for that purpose. Notable examples are the DPSV algorithm [36], which is combinatorial, and an algorithm due to Ye et al. [125] which is based on the interior point method and converges in time $O(\text{poly}(n)\log(1/\epsilon))$.

The multiplicative weights update method (MWU) is a ubiquitous meta-algorithm with numerous applications in different fields [4]. For example, it has been used in max-flow problems [29], discrepancy minimization [91], learning graphical models [84], and even in evolution [27, 100]. It is particularly useful in algorithmic game theory due to its regret-minimizing properties [53, 26], i.e. the time average behavior of MWU leads to (approximate) coarse correlated equilibria (CCE).

Fleischer et al. [51] showed how to apply MWU to the computation of Fisher market equilibria. They give a simple and decentralized algorithm based on the multiplicative weights update method, though its running time is $O(\text{poly}(n)/\epsilon^2)$. Due to the general nature of the technique, their algorithm extends to much more general classes of utilities including some that do not satisfy weak gross substitutability. For our models, a non-trivial extension of the algorithm is required in order to deal with initial endowments and matching constraints.

Gradient descent (GD) is arguably one of the most well-known and effective techniques in optimization. The main reason behind this fact lies in its simplicity and nice properties while only requiring limited information about the objective to the optimized. For convex objectives, one can show that as long as the function is Lipschitz, $O(1/\epsilon^2)$ steps suffice to get an ϵ-approximate solution. Moreover, if the function has a Lipschitz gradient, then $O(1/\epsilon)$ steps suffice. Finally if the function is strongly-convex, one can get an ϵ-approximate optimal solution in $O(\log(1/\epsilon))$ iterations, see [20] for more information. Most recently, gradient descent and its stochastic counterpart have been extensively studied and used for optimizing non-convex objectives with the guarantee of almost-always convergence towards local optima [64, 89, 80]. These results shed light on why GD works well in practice.

Gradient descent has also found numerous applications for computing market equilibria, most notably the recent work of Gao et al. [58], see references therein. They studied first-order methods for the Fisher market by considering gradient descent type algorithms for the various convex programming formulations. They show that one may exploit the fairness of the Fisher market in order to bound the market equilibrium away from the boundary of the feasible region. Using this, they develop a projected gradient descent algorithm that converges in $O(\text{poly}(n)\log(1/\epsilon))$ time and requires only simplex projections. However, the efficiency of this approach depends quite critically on the simple structure of the convex programs for Fisher markets. A more general algorithm of this form would depend on geometric properties of the feasible region and use expensive projections. For these reasons, our work uses a projection-free approach instead.

4.2 Models

In this section we will define and briefly discuss all the Nash-bargaining-based matching market models which we will cover in this paper. They can be distinguished along two axes:

1. Who can be matched with whom? We distinguish one-sided, two-sided, and non-bipartite markets.
2. Do agents bring endowments or are all agents equal?

For each model we will only consider the classic, linear-utility variant. The full version of the paper that this chapter is based on [108] also covers non-linear utilities. The one-sided and two-sided models are due to Hosseini and Vazirani [70]. However, the non-bipartite model is new to this work.

4.2.1 One-Sided Matching Markets

For one-sided matching markets, we have largely covered the setup in Section 1.1. The setting is the same as in the HZ mechanism. In particular, we are given a set A of n agents and a set G of n goods. Each agent i has a non-negative, rational utilitiy vector $(u_{ij})_{j \in G}$ to express their preferences over the goods. We are looking for a fractional perfect matching satisfying desirable properties. In the case that an *integral* perfect matching is desired, one can randomly round the fractional perfect matching using the Birkhoff-von-Neumann procedure which was also covered in Section 1.1.

The Nash-bargaining-based model for such a one-sided market with linear utilities is given by the following convex program which is closely related to the Eisenberg-Gale convex program. See also Sections 1.2 and 1.4 for more background. This is also the version of Nash bargaining that is considered in Chapter 3; it satisfies 2-approximate envy-freeness and 2-approximate incentive-compatibility.

$$\max \quad \sum_{i \in A} \log(u_i \cdot x_i) \tag{4.1a}$$

$$\text{s.t.} \quad \sum_{i \in A} x_{ij} \leq 1 \quad \forall j \in G, \tag{4.1b}$$

$$\sum_{j \in G} x_{ij} \leq 1 \quad \forall i \in A, \tag{4.1c}$$

$$x_{ij} \geq 0 \quad \forall i \in A, j \in G. \tag{4.1d}$$

Note that this version of the convex program uses ≤ 1 instead of $= 1$ constraints. Since the utilities are non-negative, both versions are equivalent. We use this version here because it makes the KKT conditions more convenient.

Next, we consider a market with endowments. Here, every agent i has an initial endowment vector $(e_{ij})_{j \in G}$ which describes how much of good j is initially owned by agent i. We assume that e is rational and a fractional perfect matching. The goal of the mechanism is now to improve every agents' utility over their initial endowment.

Recall from Section 1.4, that Nash bargaining actually has a natural way of dealing with such a setting in the form of *disagreement utilities*. Let us define $c_i = u_i \cdot e_i$ as the utility that agent i has for their initial endowment, then the vector c forms the vector of disagreement utilities, also known as the *disagreement point*. We consider an instance of such a market feasible if there is at least one fractional perfect matching which *strictly* improves the utility of each agent over their disagreement utility. In that case, the Nash bargaining model for such a market is given by the following convex program.

$$\max \quad \sum_{i \in A} \log(u_i \cdot x_i - c_i) \tag{4.2a}$$

$$\text{s.t.} \quad \sum_{i \in A} x_{ij} \leq 1 \quad \forall j \in G, \tag{4.2b}$$

$$\sum_{j \in G} x_{ij} \leq 1 \quad \forall i \in A, \tag{4.2c}$$

$$x_{ij} \geq 0 \quad \forall i \in A, j \in G. \tag{4.2d}$$

Note that this mechanism will always produce allocations which every agent prefers to their initial endowment. Hence the mechanism is *individually rational* and no agent is disincentivized to participate.

With non-negative dual variables p_j and q_i, corresponding to constraints (4.2b) and (4.2c) respectively, the KKT conditions for optimal solutions of (4.2) are given below. By setting $c_i = 0$ for all $i \in A$, we get KKT conditions for program (4.1) as well.

(1) $\forall j \in G : \; p_j > 0 \implies \sum_{i \in A} x_{ij} = 1.$

(2) $\forall i \in A : \; q_i > 0 \implies \sum_{j \in G} x_{ij} = 1.$

(3) $\forall i \in A, \; \forall j \in G : \; p_j + q_i \geq \frac{u_{ij}}{u_i \cdot x_i - c_i}.$

(4) $\forall i \in A, \; \forall j \in G : x_{ij} > 0 \implies p_j + q_i = \frac{u_{ij}}{u_i \cdot x_i - c_i}.$

For $c_i = 0$, this allows us to prove the following approximate equal-share fairness property which guarantees that every agent achieves at least $1/2$ of the utility that they would get under the equal share matching that assigns $\frac{1}{n}$ to all edges. Note that this also follows from

2-approximate envy-freeness (see Chapter 3). However, this result predates the work from that chapter.

Lemma 4.1. *Let x be an optimal solution to* (4.1)*. Then we have $u_i \cdot x_i \geq \frac{1}{2n} \sum_{j \in G} u_{ij}$ for all agents $i \in A$.*

Proof. By the KKT conditions we know that $\frac{u_{ij}}{u_i \cdot x_i} \leq p_j + q_i$ with equality if $x_{ij} > 0$. Moreover, if $p_j > 0$ then $\sum_{i \in A} x_{ij} = 1$ and likewise for q_i. Thus

$$\sum_{j \in G} p_j + \sum_{i \in A} q_i = \sum_{i \in A} \sum_{j \in G} x_{ij}(p_j + q_i) \leq n.$$

In addition, note that $q_i \leq 1$ since otherwise we cannot possibly have $\sum_{j \in G} x_{ij} = 1$ and $\sum_{j \in G} x_{ij}(p_j + q_i) = 1$. Finally, we conclude

$$\begin{aligned} u_i \cdot x_i &= \max \left\{ \frac{u_{ij}}{p_j + q_i} \,\middle|\, j \in G \right\} \\ &\geq \sum_{j \in G} \frac{u_{ij}}{p_j + q_i} \cdot \frac{p_j + q_i}{\sum_{j' \in G} p_{j'} + q_i} \\ &\geq \frac{\sum_{j \in G} u_{ij}}{n q_i + \sum_{j \in G} p_j} \\ &\geq \frac{1}{2n} \sum_{j \in G} u_{ij}. \end{aligned}$$

□

The usefulness of this result is that it bounds the optimal solution away from the boundary of the feasible region. Unfortunately, in the case of endowments, the equal share matching may no longer be feasible and so this notion loses some meaning. We will however give an analogous but weaker bound in Section 4.4, specifically in Lemma 4.9.

4.2.2 Two-Sided Matching Markets

Next, we consider two-sided matching markets. This is the setting that we also covered in Section 3.3 of the previous chapter. Here, we are given equal-sized sets A of agents and J of jobs. Agents have utilities $(u_{ij})_{i\in A, j\in J}$ and jobs have utilities $(w_{ji})_{j\in J, i\in A}$. Both are assumed to be rational and non-negative. As in the one-sided setting, we are looking for a fractional perfect matching as integral ones can be found via the Birkhoff-von-Neumann theorem.

Extending the HZ mechanism to such a two-sided setting is rather challenging. See Section 3.3 and particularly the work by Manjunath [95] which makes an attempt at such a generalization. In constrast, Nash bargaining extends very naturally to a two-sided market. We simply consider the utilities of the jobs equally as important as the utilities of the agents. In that case, the natural Nash-bargaining-based model is given by the following convex program.

$$\max \quad \sum_{i\in A} \log(u_i \cdot x_i) + \sum_{j\in J} \log(w_j \cdot x_j) \tag{4.3a}$$

$$\text{s.t.} \quad \sum_{i\in A} x_{ij} \leq 1 \quad \forall j \in J, \tag{4.3b}$$

$$\sum_{j\in J} x_{ij} \leq 1 \quad \forall i \in A, \tag{4.3c}$$

$$x_{ij} \geq 0 \quad \forall i \in A, j \in J. \tag{4.3d}$$

Note that by a slight abuse of notation, we write $w_j \cdot x_j := \sum_{i\in A} w_{ji} x_{ij}$ here and in the following. In addition, observe that this model reduces to the standard one-sided model if the jobs are indifferent towards the agents.

A key advantage of Nash bargaining is its flexibility. We can easily extend the two-sided model to also include endowments. As before, assume that there is some initial fractional perfect matching $(e_{ij})_{i\in A, j\in J}$ which represents the status quo in the market. Each agent $i \in A$ has a disagreement utility of $c_i := u_i \cdot e_i$ and likewise each job $j \in J$ has a disagreement utility $d_j := w_j \cdot e_j$. The corresponding convex program is given below.

$$\max \quad \sum_{i\in A} \log(u_i \cdot x_i - c_i) + \sum_{j\in J} \log(w_j \cdot x_j - d_j) \tag{4.4a}$$

$$\text{s.t.} \quad \sum_{i\in A} x_{ij} \le 1 \quad \forall j \in J, \tag{4.4b}$$

$$\sum_{j\in J} x_{ij} \le 1 \quad \forall i \in A, \tag{4.4c}$$

$$x_{ij} \ge 0 \quad \forall i \in A, j \in J. \tag{4.4d}$$

The KKT conditions for program (4.4) with non-negative dual variables p and q corresponding to constraints (4.4b) and (4.4c) respectively are:

(1) $\forall j \in J : \ p_j > 0 \implies \sum_{i\in A} x_{ij} = 1.$

(2) $\forall i \in A : \ q_i > 0 \implies \sum_{j\in J} x_{ij} = 1.$

(3) $\forall i \in A, \ \forall j \in J : \ p_j + q_i \ge \frac{u_{ij}}{u_i\cdot x_i - c_i} + \frac{w_{ij}}{w_j\cdot x_j - d_j}.$

(4) $\forall i \in A, \ \forall j \in J : x_{ij} > 0 \implies p_j + q_i = \frac{u_{ij}}{u_i\cdot x_i - c_i} + \frac{w_{ij}}{w_j\cdot x_j - d_j}.$

Nash bargaining does not achieve as strong fairness properties in two-sided markets as it does in one-sided markets (see Section 3.3), and the 2-approximate equal-share fairness does not extend to this setting, even without endowments. However, we once again show a weaker analogue in Section 4.4.

4.2.3 Non-Bipartite Matching Markets

Finally, we will introduce one more setting which was not already investigated by Hosseini and Vazirani [70]: non-bipartite markets. In this setting we simply have a set A of n agents and non-negative, rational utilities $(u_{ij})_{i \in A, j \in A \setminus \{i\}}$; each agent has a utility for every other agent. A feasible solution is now any convex combination of integral matchings in the complete graph over A. This can be represented via a vector $(x_{ij})_{\{i,j\} \subseteq A}$ over all undirected edges in the complete graph on A.

Recall that convex combinations of integral matchings, i.e. points in the matching polytope, are not simply vectors x that satisfy $\sum_{j \in A \setminus \{i\}} x_{ij} \leq 1$ at every $i \in A$. Instead, such a vector must also satisfy $\sum_{\{i,j\} \subseteq B} x_{ij} \leq \frac{|B|-1}{2}$ for every set $B \subseteq A$ of odd cardinality. This is a classic result due to Edmonds [45]. For ease of notation let us denote the collection of all such odd subsets B of A by $\mathcal{O}$. This motivates the convex program:

$$
\begin{aligned}
\max \quad & \sum_{i \in A} \log(u_i \cdot x_i) && && \text{(4.5a)} \\
\text{s.t.} \quad & \sum_{j \in A \setminus \{i\}} x_{ij} \leq 1 && \forall i \in A, && \text{(4.5b)} \\
& \sum_{\{i,j\} \subseteq B} x_{ij} \leq \frac{|B|-1}{2} && \forall B \in \mathcal{O}, && \text{(4.5c)} \\
& x_{ij} \geq 0 && \forall i \in A, j \in A \setminus \{i\}. && \text{(4.5d)}
\end{aligned}
$$

We remark that as in the bipartite case, it is possible to take such a vector and—in a combinatorial, polynomial time manner—decompose it into a convex combination of most n^2 integral matchings. This was shown by Padberg and Wolsey [107]; for a modern proof see [120]. Hence, this allows a similar rounding strategy to the standard bipartite setting if one desires integral allocations.

Lastly, we can once again consider endowments, even for this non-bipartite setting. The extension is routine and results in the following convex program.

$$\max \quad \sum_{i\in A} \log(u_i \cdot x_i - c_i) \tag{4.6a}$$

$$\text{s.t.} \quad \sum_{j\in A\setminus\{i\}} x_{ij} \le 1 \qquad \forall i \in A, \tag{4.6b}$$

$$\sum_{\{i,j\}\subseteq B} x_{ij} \le \frac{|B|-1}{2} \quad \forall B \in \mathcal{O}, \tag{4.6c}$$

$$x_{ij} \ge 0 \qquad \forall i \in A, j \in A \setminus \{i\}. \tag{4.6d}$$

The KKT conditions for the CP (4.6) are given below. The dual variables are $(p_i)_{i\in A}$, corresponding to constraint (4.6b), and $(z_B)_{B\in\mathcal{O}}$, corresponding to constraint (4.6c).

(1) $\forall i \in A: \ p_i > 0 \implies \sum_{j\in A\setminus\{i\}} x_{ij} = 1.$

(2) $\forall B \in \mathcal{O}: \ z_B > 0 \implies \sum_{\{i,j\}\subseteq B} x_{ij} = \frac{|B|-1}{2}.$

(3) $\forall \{i,j\} \subseteq A: \ p_i + p_j + \sum_{B\in\mathcal{O}:\{i,j\}\subseteq B} z_B \ge \frac{u_{ij}}{u_i\cdot x_i - c_i} + \frac{u_{ji}}{u_j \cdot x_j - c_j}.$

(4) $\forall \{i,j\} \subseteq A: \ x_{ij} > 0 \implies p_i + p_j + \sum_{B\in\mathcal{O}:\{i,j\}\subseteq B} z_B = \frac{u_{ij}}{u_i\cdot x_i - c_i} + \frac{u_{ji}}{u_j \cdot x_j - c_j}.$

Once again, we will show a weakened version of equal-share fairness for this setting in Section 4.4 where it will be used to guarantee the convergence of a conditional gradient descent algorithm.

4.3 Multiplicative Weights Update

We will now turn to the first of our two algorithms for Nash bargaining in matching markets. It is based on the classic MWU technique. The algorithm works for the one-sided market with endowments; for a setting with non-linear utilities see the full version of the paper that this chapter is based on [108]. Our goal is hence to prove the following theorem.

Theorem 4.1. *The MWU algorithm (Algorithm 4.1) computes an ϵ-approximate Nash bargaining solution for a one-sided market with endowments, i.e. CP* (4.2)*, in $O\left(\frac{n \log n}{\epsilon^2}\right)$ iterations. Each iteration can be implemented in $O(n^2)$ time.*

In the following, fix some one-sided market with agents A, goods G, $|A| = |G| = n$ and non-negative, rational utilities $(u_{ij})_{i \in A, j \in G}$. Assume that the instance is feasible, i.e. that there is at least one allocation which is better for everyone compared to the disagreement point. There are various possible definitions of an ϵ-approximate Nash bargaining solution. In this section, we will use the following one.

Definition 4.1. *Let $(x_{ij})_{i \in A, j \in G}$ be some positive assignment matrix. We call x an ϵ-*approximate Nash bargaining solution *if*

$$\sum_{i \in A} \log(u_i \cdot x_i - c_i) \geq \sum_{i \in A} \log(u_i \cdot x_i^\star - c_i)$$

where $x^\star$ is an optimal solution to (4.2) *and x is approximately feasible in the sense that $\sum_{i \in A} x_{ij} \leq 1 + \epsilon$ for all $i \in A$ and likewise $\sum_{j \in G} x_{ij} \leq 1 + \epsilon$ for all $j \in G$.*

Note that we would typically want an actually feasible solution. In that case one can simply scale down the solution by $1 + \epsilon$ after the fact.

The general proof strategy consists of two steps:

1. The multiplicative weights technique generally applies to *feasibility* rather than optimization problems. Using the KKT conditions of (4.2), we will establish a feasibility program whose feasible solutions are optimal solutions to (4.2).

2. Then we will show that our algorithm can solve the feasibility program approximately in polynomial time and that this indeed gives an approximate Nash bargaining point.

4.3.1 From Optimization to Feasibility

As noted above, our first goal will be to turn the problem of solving the optimization problem (4.2) into a feasibility program. This program is defined on the allocation $(x_{ij})_{i\in A, j\in G}$ as well as the dual variables $(p_j)_{j\in G}$ and $(q_i)_{i\in A}$.

$$
\begin{aligned}
\mathrm{CP}_i(p,q) &\le u_i \cdot x_i && \forall i \in A,\\
\sum_{i\in A} x_{ij} &\le 1 && \forall j \in G,\\
\sum_{j\in G} x_{ij} &\le 1 && \forall i \in A,\\
\sum_{j\in G} p_j + \sum_{i\in A} q_i &= n + \sum_{i\in A} c_i \min_{j\in G} \frac{p_j+q_i}{u_{ij}} && \forall i \in A, j \in G,\\
x,p,q &\ge 0
\end{aligned}
\tag{4.7}
$$

where

$$
\begin{aligned}
\mathrm{CP}_i(p,q) := \max \quad & u_i \cdot y_i\\
\text{s.t.} \quad & \sum_{j\in G} (p_j+q_i) y_j \le 1 + c_i \min_{j\in G} \frac{p_j+q_i}{u_{ij}},\\
& y \ge 0.
\end{aligned}
$$

From a high level perspective, we would like to express the fact that (x,p,q) satisfy the KKT conditions of CP (4.2) in order to show that x is an optimal solution of (4.2). However, the KKT conditions come with complementarity constraints that are not easily dealt with. But what those complementarity constraints essentially do is to enforce that each agent is getting an optimal bundle of goods given the prices p and q.

For a simplified setting, consider the case in which $c_i = 0$ for all $i \in A$. In this case, $\mathrm{CP}_i(p, q)$ gives us the maximum utility that agent i can obtain under prices p on the goods with a budget of $1 - q_i$. Our feasibility program then reduces to the condition that the sum of all prices is n and each agent gets an optimal bundle with the given prices and budgets.

With the introduction of endowments, the program gets more complicated. $\mathrm{CP}_i(p, q)$ now measures the optimal bundle that agent i could get if they were given a budget of $1 - q_i + d_i$ where d_i is the budget needed to buy back a bundle of utility c_i, their disagreement utility. Note that it is necessary that each agent gets a budget that is large enough to buy back to their disagreement utility. Otherwise, the resulting allocation would not be individually rational.

Theorem 4.2. *Let x be a Nash bargaining solution, i.e. an optimal solution to CP* (4.2)*, then there exist non-negative dual variables p, q such that (x, p, q) is a feasible solution for* (4.7)*. Conversely, if (x, p, q) is a feasible solution to* (4.7)*, then x is an optimal solution for* (4.2).

Proof. We will start with the first half of the theorem. Let x be a Nash bargaining solution. Then there are non-negative dual variables p, q that satisfy the KKT conditions which we restate below for convenience.

(1) $\forall j \in G: \; p_j > 0 \implies \sum_{i \in A} x_{ij} = 1$.

(2) $\forall i \in A: \; q_i > 0 \implies \sum_{j \in G} x_{ij} = 1$.

(3) $\forall i \in A, \, \forall j \in G: \; p_j + q_i \geq \frac{u_{ij}}{u_i \cdot x_i - c_i}$.

(4) $\forall i \in A, \, \forall j \in G: x_{ij} > 0 \implies p_j + q_i = \frac{u_{ij}}{u_i \cdot x_i - c_i}$.

Our goal is to show that (x, p, q) is a feasible solution for (4.7). Observe that by definition of CP_i, we get that

$$\begin{aligned}
\mathrm{CP}_i(p,q) &= \left(\max_{j\in G} \frac{u_{ij}}{p_j + q_i}\right)\left(1 + c_i \min_{j\in G} \frac{p_j + q_i}{u_{ij}}\right) \\
&= \left(\frac{1}{\min_{j\in G} \frac{p_j+q_i}{u_{ij}}}\right)\left(1 + c_i \min_{j\in G} \frac{p_j + q_i}{u_{ij}}\right) \\
&= c_i + \frac{1}{\min_{j\in G} \frac{p_j+q_i}{u_{ij}}} \\
&= c_i + \max_{j\in G} \frac{u_{ij}}{p_j + q_i}.
\end{aligned}$$

From the KKT conditions, we can deduce that $u_i \cdot x_i \geq c_i + \frac{u_{ij}}{p_j+q_i}$ for all $i \in A$ and $j \in G$. Accordingly, we get $u_i \cdot x_i \geq \mathrm{CP}_i(p,q)$ as is required by our feasibility program.

Finally, we use the KKT conditions to compute

$$\begin{aligned}
\sum_{j\in G} p_j + \sum_{i\in A} q_i &= \sum_{j\in G} p_j \sum_{i\in A} x_{ij} + \sum_{i\in A} q_i \sum_{j\in G} x_{ij} \\
&= \sum_{i\in A} \sum_{j\in G} x_{ij}(p_j + q_i) \\
&= \sum_{i\in A} \sum_{j\in G} x_{ij} \frac{u_{ij}}{u_i \cdot x_i - c_i} \\
&= \sum_{i\in A} \frac{u_i \cdot x_i}{u_i \cdot x_i - c_i} \\
&= n + \sum_{i\in A} \frac{c_i}{u_i \cdot x_i - c_i} \\
&= n + \sum_{i\in A} c_i \min_{j\in G} \frac{p_j + q_i}{u_{ij}}
\end{aligned}$$

and hence prove the first half of the theorem.

For the converse direction let (x, p, q) be a feasible point of (4.7). Our goal is to show that x is a Nash bargaining point. Clearly x is at least a fractional matching. Moreover, it is not hard to see that $\mathrm{CP}_i(p, q) > c_i$ and hence x is a feasible point of (4.2). It remains to show that x, p, q satisfy the KKT conditions stated above.

The first observation is that since $u_i \cdot x_i \geq \mathrm{CP}_i(p, q)$ holds by assumption, we have

$$\sum_{j \in G} (p_j + q_i) x_{ij} \geq 1 + c_i \min_{j \in G} \frac{p_j + q_i}{u_{ij}}. \tag{4.8}$$

This is because $\mathrm{CP}_i(p, q)$ represents the maximum utility one can get with prices $p_j + q_i$ and a budget of $1 + c_i \min_{j \in G} \frac{p_j + q_i}{u_{ij}}$. So since the bundle x_i provides this much utility, it must cost at least this budget. Otherwise, $\mathrm{CP}_i(p, q)$ would be larger.

Now sum the above inequality over all $i \in A$ to obtain:

$$\sum_{i \in A} \sum_{j \in G} (p_j + q_i) x_{ij} \geq n + \sum_{i \in A} c_i \min_{j \in G} \frac{p_j + q_i}{u_{ij}}. \tag{4.9}$$

Since (x, p, q) is a solution to (4.7), we also have

$$\begin{aligned} n + \sum_{i \in A} c_i \min_{j \in G} \frac{p_j + q_i}{u_{ij}} &= \sum_{j \in G} p_j + \sum_{i \in A} q_i \\ &\geq \sum_{j \in G} p_j \sum_{i \in A} x_{ij} + \sum_{i \in A} q_i \sum_{j \in G} x_{ij} \\ &= \sum_{i \in A} \sum_{j \in G} (p_j + q_i) x_{ij}. \end{aligned} \tag{4.10}$$

The inequality comes from the fact $\sum_{i \in A} x_{ij} \leq 1$ for all $j \in G$ and $\sum_{j \in G} x_{ij} \leq 1$ for all $i \in A$.

So the inequalitites in (4.9) and (4.10) must actually be equalities. We can also deduce from the tight inequality in (4.9) that $\sum_{j \in G} x_{ij} = 1$ whenever $q_i > 0$ and $\sum_{i \in A} x_{ij} = 1$ whenever $p_j > 0$. This takes care of the first two KKT conditions.

Because (4.9) is an equality, the same must hold for (4.8) for all $i \in A$. This allows us to deduce that x_i is actually a feasible solution for the program that defines $\mathrm{CP}_i(p, q)$ and thus $u_i \cdot x_i = \mathrm{CP}_i(p, q)$. But this implies:

$$\begin{aligned} u_i \cdot x_i &= \mathrm{CP}_i(p, q) \\ &= \left(\max_{j \in G} \frac{u_{ij}}{p_j + q_i}\right)\left(1 + c_i \min_{j \in G} \frac{p_j + q_i}{u_{ij}}\right) \\ &= c_i + \max_{j \in G} \frac{u_{ij}}{p_j + q_i} \\ &= c_i + \frac{u_{ij}}{p_j + q_i} \text{ for all } j \in G \text{ s.t. } x_{ij} > 0. \end{aligned}$$

We conclude that $u_i \cdot x_i \geq c_i + \frac{u_{ij}}{p_j + q_i}$, with equality only if $x_{ij} > 0$. This establishes the remaining two KKT conditions and hence finishes the proof. □

4.3.2 Main Analysis

By Theorem 4.2, we know that finding a feasible solution to (4.7) is equivalent to finding a Nash bargaining point. Our approach is now to find an approximate feasible solution to (4.7). Finally, we will show that this approximate feasible solution is also an approximate Nash bargaining point in a certain sense. Our MWU algorithm is given below; see Algorithm 4.1.

On a high level the algorithm works as follows. Recall that the prices must satisfy

$$\sum_{j \in G} p_j + \sum_{i \in A} q_i = n + \sum_{i \in A} c_i \min_{j \in G} \frac{p_j + q_i}{u_{ij}} \tag{4.11}$$

in order to be a feasible solution to (4.7). During the algorithm, our prices p, q may temporarily lose this property and become *unscaled prices* $\tilde{p}, \tilde{q}$.

In Algorithm 4.1, we run T phases where T is to be determined later. In each phase we do the following:

1. First, we rescale the unscaled prices $\tilde{p}^{(t)}, \tilde{q}^{(t)}$ that were generated in the previous phase so that they satisfy (4.11).
2. Next, each agent i independently buys an optimal bundle at these prices, i.e. they solve for CP_i.
3. Then, we determine what the maximum amount is that any agent or good is allocated. Call this $\sigma^{(t)}$.
4. Finally, we make a multiplicative update to the prices in the standard way according to how over- or underdemanded each agent / good is. The resulting prices are the unscaled prices for the next iteration.

At the end, we return a weighted average of the allocations and prices. Our goal in this section will be to show that $(\overline{x}, \overline{p}, \overline{q})$ is an approximate feasible solution to (4.7) and hence that $\overline{x}$ is an approximate Nash bargaining point.

Before we can actually analyze Algorithm 4.1, we need to show that it is well-defined, i.e. that the rescaling step can actually be carried out. This is captured by the following lemma.

Lemma 4.2. *Assume that the instance is feasible, i.e. that there exists at least one allocation which improves the utility of every agent over their disagreement utility. Given $\tilde{p}, \tilde{q}$, we can always rescale them to p, q so that $\sum_{j \in G} p_j + \sum_{i \in A} q_i = n + \sum_{i \in A} c_i \min_{j \in G} \frac{p_j + q_i}{u_{ij}}$.*

Proof. Start by defining

$$\lambda := \sum_{j \in G} \tilde{p}_j + \sum_{i \in A} \tilde{q}_i - \sum_{i \in A} c_i \min_{j \in G} \frac{\tilde{p}_j + \tilde{q}_i}{u_{ij}}.$$

Algorithm 4.1: MULTIPLICATIVE WEIGHTS UPDATE

Initialize prices $\tilde{p}^{(0)} \leftarrow \mathbf{1}, \tilde{q}^{(0)} \leftarrow \mathbf{1}$

for $t = 1$ *to* T **do**

  Rescale $\tilde{p}^{(t)}, \tilde{q}^{(t)}$ by a common factor $a^{(t)}$ to $p^{(t)}, q^{(t)}$ such that $\sum_{j\in G} p_j^{(t)} + \sum_{i\in A} q_i^{(t)} = n + \sum_{i\in A} c_i \min_{j\in G} \frac{p_j^{(t)}+q_i^{(t)}}{u_{ij}}$.

  for $i = 1$ *to* n **do**

    $x_i^{(t)} \leftarrow \arg\max_y \left\{u_i \cdot y : \sum_{j\in G} \left(p_j^{(t)} + q_i^{(t)}\right) y_j \le 1 + c_i \min_{j\in G} \frac{p_j^{(t)}+q_i^{(t)}}{u_{ij}}\right\}$.

  Compute demand vectors $d^{(t)} \leftarrow \sum_{i\in A} x_i^{(t)}, h^{(t)} \leftarrow \sum_{j\in G} x_j^{(t)}$.

  $\sigma^{(t)} \leftarrow \min\left(\frac{1}{\max_{j\in G} d_j^{(t)}}, \frac{1}{\max_{i\in A} h_i^{(t)}}\right)$

  for $j = 1$ *to* n **do**

    $\tilde{p}_j^{(t+1)} \leftarrow \tilde{p}_j^{(t)}(1 + \epsilon\sigma^{(t)} d_j^{(t)})$

  for $i = 1$ *to* n **do**

    $\tilde{q}_i^{(t+1)} \leftarrow \tilde{q}_i^{(t)}(1 + \epsilon\sigma^{(t)} h_i^{(t)})$

return $\overline{x} \leftarrow \frac{\sum_{t=1}^T \sigma^{(t)} x^{(t)}}{\sum_{t=1}^T \sigma^{(t)}}, \overline{p} \leftarrow \frac{\sum_{t=1}^T \sigma^{(t)} p^{(t)}}{\sum_{t=1}^T \sigma^{(t)}}, \overline{q} \leftarrow \frac{\sum_{t=1}^T \sigma^{(t)} q^{(t)}}{\sum_{t=1}^T \sigma^{(t)}}$.

Provided that $\lambda > 0$, we can set $p := \frac{n}{\lambda}\tilde{p}$ and $q = \frac{n}{\lambda}\tilde{q}$ and the lemma follows. So our goal is now to show that $\lambda > 0$.

Now let us use the feasibility: we know that there is some fractional matching $(y_{ij})_{i\in A, j\in G}$ which satisfies $u_i \cdot y_i > c_i$ for all $i \in A$. In particular, we have $\max_{j\in G} u_{ij} \ge u_i \cdot y_i > c_i$. Now compute:

$$\begin{aligned}
\sum_{i\in A} c_i \min_{j\in G} \frac{\tilde{p}_j + \tilde{q}_i}{u_{ij}} &\le \sum_{i\in A} c_i \min_{j\in G} \frac{\tilde{p}_j + \tilde{q}_i}{u_i \cdot y_i} \\
&\le \sum_{i\in A} c_i \frac{1}{n} \sum_{j\in G} \frac{\tilde{p}_j + \tilde{q}_i}{u_i \cdot y_i} \\
&< \frac{1}{n} \sum_{i\in A} \sum_{j\in G} (\tilde{p}_j + \tilde{q}_i) = \sum_{j\in G} \tilde{p}_j + \sum_{i\in A} \tilde{q}_i.
\end{aligned}$$

Hence, $\lambda > 0$ as desired. □

In order to bound the rate of convergence of Algorithm 4.1, we use a potential function argument; see [4] for more details on this classic technique. The potential function in iteration t is simply the sum of the unscaled prices, i.e. $\Phi(t) \coloneqq \sum_{j \in G} \tilde{p}_j^{(t)} + \sum_{i \in A} \tilde{q}_i^{(t)}$. We will give upper and lower bounds on $\Phi(T)$ where T is the total number of iterations. This allows us to ultimately show the desired $O(1/\epsilon^2)$ convergence rate.

Lemma 4.3. *For any $T \in \mathbb{N}$, we have $\Phi(T) \leq 2n \cdot \exp\left(\epsilon \sum_{t=1}^{T-1} \sigma^{(t)}\right)$ where $\sigma^{(t)}$ is the reciprocal maximum demand in iteration t; see Algorithm 4.1.*

Proof. Pick some time index t. Then we compute

$$\begin{aligned}
\frac{\Phi(t+1)}{\Phi(t)} &= \frac{\sum_{j \in G} \tilde{p}_j^{(t+1)} + \sum_{i \in A} \tilde{q}_i^{(t+1)}}{\Phi(t)} \\
&= \frac{\sum_{j \in G} (1 + \epsilon \sigma^{(t)} d_j^{(t)}) \tilde{p}_j^{(t)} + \sum_{i \in A} + (1 + \epsilon \sigma^{(t)} h_i^{(t)}) \tilde{q}_i^{(t)}}{\Phi(t)} \\
&= 1 + \epsilon \sigma^{(t)} \frac{\sum_{j \in G} d_j^{(t)} \tilde{p}_j^{(t)} + \sum_{i \in A} h_i^{(t)} \tilde{q}_i^{(t)}}{\Phi(t)}
\end{aligned}$$

by making use of the definition of the udpates of $\tilde{p}, \tilde{q}$.

Next, we factor out the rescaling factor $a^{(t)}$ and apply the definitions of d and h. This allows us to bound:

$$\begin{aligned}
\frac{\sum_{j \in G} d_j^{(t)} \tilde{p}_j^{(t)} + \sum_{i \in A} h_i^{(t)} \tilde{q}_i^{(t)}}{\Phi(t)} &= \frac{\sum_{i \in A} \sum_{j \in G} (p_j^{(t)} + q_i^{(t)}) x_{ij}^{(t)}}{a^{(t)} \Phi(t)} \\
&\leq \frac{\sum_{i \in A} \left(1 + c_i \min_{j \in G} \frac{p_j^{(t)} + q_i^{(t)}}{u_{ij}}\right)}{a^{(t)} \Phi(t)} \\
&= \frac{n + \sum_{i \in A} c_i \min_{j \in G} \frac{p_j^{(t)} + q_i^{(t)}}{u_{ij}}}{a^{(t)} \Phi(t)}
\end{aligned}$$

The crucial inequality follows directly from the definition of $x_i^{(t)}$ in Algorithm 4.1.

Finally, observe that because of the rescaling, we have that

$$\frac{n + \sum_{i \in A} c_i \min_{j \in G} \frac{p_j^{(t)} + q_i^{(t)}}{u_{ij}}}{a^{(t)}\Phi(t)} = \frac{\sum_{j \in G} p_j^{(t)} + \sum_{i \in A} q_i^{(t)}}{a^{(t)}\Phi(t)} = 1$$

and thus

$$\frac{\Phi(t+1)}{\Phi(t)} \leq 1 + \epsilon\sigma^{(t)} \leq 1 + \exp(\epsilon\sigma^{(t)}). \tag{4.12}$$

Finally, we can bound $\Phi(T)$ by expressing it as a telescopic product over (4.12). The resulting bound is

$$\Phi(T) \leq \Phi(1) \cdot \exp\left(\epsilon \sum_{t=1}^{T-1} \sigma^{(t)}\right) = 2n \cdot \exp\left(\epsilon \sum_{t=1}^{T-1} \sigma^{(t)}\right).$$

□

Lemma 4.4. *We have*

$$\Phi(T) \geq \exp\left(\epsilon(1-\epsilon)\max\left(\max_{j \in G} \sum_{t=1}^{T} \sigma^{(t)} d_j^{(t)}, \max_{i \in A} \sum_{t=1}^{T} \sigma^{(t)} h_i^{(t)}\right)\right).$$

Proof. Note that by the definition of Φ and the way the prices are updated, we know that

$$\begin{aligned}\Phi(T) &= \sum_{j \in G} \tilde{p}_j^{(T)} + \sum_{i \in A} \tilde{q}_i^{(T)} \\ &= \sum_{j \in G} \prod_{t=1}^{T-1} (1 + \epsilon\sigma^{(t)} d_j^{(t)}) + \sum_{i \in A} \prod_{t=1}^{T-1} (1 + \epsilon\sigma^{(t)} h_i^{(t)}).\end{aligned}$$

Finally, we use the identity $e^{\mu(1-y)} \le 1 + \mu$ for $0 < \mu \le y < 1$ to compute:

$$\begin{aligned}
&\sum_{j\in G}\prod_{t=1}^{T-1}(1+\epsilon\sigma^{(t)}d_j^{(t)}) + \sum_{i\in A}\prod_{t=1}^{T-1}(1+\epsilon\sigma^{(t)}h_i^{(t)}) \\
&\ge \sum_{j\in G}\exp\left(\epsilon(1-\epsilon)\sum_{t=1}^{T-1}\sigma^{(t)}d_j^{(t)}\right) + \sum_{i\in A}\exp\left(\epsilon(1-\epsilon)\sum_{t=1}^{T-1}\sigma^{(t)}h_i^{(t)}\right) \\
&\ge \max\left(\max_{j\in G}\exp\left(\epsilon(1-\epsilon)\sum_{t=1}^{T-1}\sigma^{(t)}d_j^{(t)}\right), \max_{i\in A}\exp\left(\epsilon(1-\epsilon)\sum_{t=1}^{T-1}\sigma^{(t)}h_i^{(t)}\right)\right) \\
&= \exp\left(\epsilon(1-\epsilon)\max\left(\max_{j\in G}\sum_{t=1}^{T}\sigma^{(t)}d_j^{(t)}, \max_{i\in A}\sum_{t=1}^{T}\sigma^{(t)}h_i^{(t)}\right)\right). \qquad \square
\end{aligned}$$

By combining Lemmas 4.3 and 4.4, we can conclude that

$$\epsilon(1-\epsilon)\max\left(\max_{j\in G}\sum_{t=1}^{T}\sigma^{(t)}d_j^{(t)}, \max_{i\in A}\sum_{t=1}^{T}\sigma^{(t)}h_i^{(t)}\right) \le \ln(2n) + \epsilon\sum_{t=1}^{T-1}\sigma^{(t)}. \tag{4.13}$$

Let us now use this fact to show that the averaged output $(\overline{x}, \overline{p}, \overline{q})$ is indeed an approximately feasible solutin of (4.7).

Lemma 4.5. *If* $T \ge \frac{2n\log_{1+\epsilon}(2n)}{\epsilon} = O(\frac{n\log n}{\epsilon^2})$*, then we have*

$$\begin{aligned}
\mathrm{CP}_i(\overline{p}, \overline{q}) &\le u_i \cdot \overline{x}_i && \forall i \in A, \\
\sum_{i\in A}\overline{x}_{ij} &\le \frac{1}{1-2\epsilon} = 1 + O(\epsilon) && \forall j \in G, \\
\sum_{j\in G}\overline{x}_{ij} &\le \frac{1}{1-2\epsilon} = 1 + O(\epsilon) && \forall i \in A, \\
\sum_{j\in G}\overline{p}_j + \sum_{i\in A}\overline{q}_i &= n + \sum_{i\in A} c_i \min_{j\in G}\frac{\overline{p}_j + \overline{q}_i}{u_{ij}} && \forall i \in A, j \in G, \\
\overline{x}, \overline{p}, \overline{q} &\ge 0.
\end{aligned}$$

In other words, $(\overline{x}, \overline{p}, \overline{q})$ is a feasible solution to (4.7)*, except for an $O(\epsilon)$ violation of the matching constraints.*

Proof. Non-negativity is trivial and the fact that

$$\sum_{j\in G} \overline{p}_j + \sum_{i\in A} \overline{q}_i = n + \sum_{i\in A} c_i \min_{j\in G} \frac{\overline{p}_j + \overline{q}_i}{u_{ij}}$$

holds for all $i \in A, j \in G$ comes from the fact that we always rescale prices so that this is satisfied.

Now let us show that $\mathrm{CP}_i(\overline{p}, \overline{q}) \le u_i \cdot \overline{x}_i$ holds. Clearly we have that $\mathrm{CP}_i(p^{(t)}, q^{(t)}) \le u_i \cdot x_i^{(t)}$ for every individual iteration t. In fact, this holds with equality by the choice of $x_i^{(t)}$. The reason why the inequality extends also to the averages is simply because $\mathrm{CP}_i(p, q)$ is a convex function. This is because we can write $\mathrm{CP}_i(p, q) = c_i + \max_{j\in G} \left\{ \frac{u_{ij}}{p_j+q_i} \right\}$ and taking the maximum over a collection of convex functions is convex. So we can take the $\sigma^{(t)}$-weighted average over the inequalities $\mathrm{CP}_i(p^{(t)}, q^{(t)}) \le u_i \cdot x_i^{(t)}$ to get $\mathrm{CP}_i(\overline{p}, \overline{q}) \le u_i \cdot \overline{x}_i$ by Jensen's inequality.

Finally, it remains to show that we approximately satisfy the matching constraints which requires a more substantial calculation. Note that $\sum_{i\in A} \overline{x}_{ij}$ can also be expressed as $\sum_{t=1}^{T} \sigma^{(t)} d_j^{(t)}$ and likewise $\sum_{j\in G} \overline{x}_{ij}$ can be expressed as $\sum_{t=1}^{T} \sigma^{(t)} h_i^{(t)}$. This follows directly from the definitions in Algorithm 4.1.

Using inequality (4.13), we get

$$\begin{aligned} \max\left(\max_{j\in G} \sum_{t=1}^{T} \sigma^{(t)} d_j^{(t)}, \max_{i\in A} \sum_{t=1}^{T} \sigma^{(t)} h_i^{(t)} \right) &\le \frac{\ln(2n) + \epsilon \sum_{t=1}^{T} \sigma^{(t)}}{\epsilon(1-\epsilon) \sum_{t=1}^{T} \sigma^{(t)}} \\ &= \frac{1}{1-\epsilon} + \frac{\ln(2n)}{\epsilon(1-\epsilon) \sum_{t=1}^{T} \sigma^{(t)}}. \end{aligned} \tag{4.14}$$

Observe that at every iteration t, there exists at least one $j \in G$ or $i \in A$ so that $\tilde{p}_j^{(t)}$ or $\tilde{q}_i^{(t)}$ increases by a factor of $(1+\epsilon)$. Hence after T iterations it follows via the pigeonhole principle that

$$\max\left(\max_{j\in G}\tilde{p}_j^{(T+1)}, \max_{i\in A}\tilde{q}_i^{(T+1)}\right) \geq (1+\epsilon)^{T/2n} \geq (2n)^{1/\epsilon},$$

where the last inequality holds by setting $T \geq \frac{2n\log_{1+\epsilon}(2n)}{\epsilon}$.

This now lets us compute:

$$\begin{aligned}
&\max\left(\max_{j\in G}\frac{\sum_{t=1}^{T}\sigma^{(t)}d_j^{(t)}}{\sum_{t=1}^{T}\sigma^{(t)}}, \max_{i\in A}\frac{\sum_{t=1}^{T}\sigma^{(t)}h_i^{(t)}}{\sum_{t=1}^{T}\sigma^{(t)}}\right)\\
&\quad= \max\left(\max_{j\in G}\frac{\ln\prod_{t=1}^{T}\exp\left(\epsilon\sigma^{(t)}d_j^{(t)}\right)}{\epsilon\sum_{t=1}^{T}\sigma^{(t)}}, \max_{i\in A}\frac{\ln\prod_{t=1}^{T}\exp\left(\epsilon\sigma^{(t)}h_i^{(t)}\right)}{\epsilon\sum_{t=1}^{T}\sigma^{(t)}}\right)\\
&\quad\geq \max\left(\max_{j\in G}\frac{\ln\prod_{t=1}^{T}(1+\epsilon\sigma^{(t)}d_j^{(t)})}{\epsilon\sum_{t=1}^{T}\sigma^{(t)}}, \max_{i\in A}\frac{\ln\prod_{t=1}^{T}(1+\epsilon\sigma^{(t)}h_i^{(t)})}{\epsilon\sum_{t=1}^{T}\sigma^{(t)}}\right)\\
&\quad= \max\left(\max_{j\in G}\frac{\ln\tilde{p}_j^{(T+1)}}{\epsilon\sum_{t=1}^{T}\sigma^{(t)}}, \max_{i\in A}\frac{\ln\tilde{q}_i^{(T+1)}}{\epsilon\sum_{t=1}^{T}\sigma^{(t)}}\right)\\
&\quad= \frac{1}{\epsilon\sum_{t=1}^{T}\sigma^{(t)}}\ln\left(\max\left(\max_{j\in G}\tilde{p}_j^{(T+1)}, \max_{i\in A}\tilde{q}_i^{(T+1)}\right)\right)\\
&\quad\geq \frac{\ln 2n}{\epsilon^2\sum_{t=1}^{T}\sigma^{(t)}}.
\end{aligned} \tag{4.15}$$

Finally, combining (4.14) and (4.15) yields

$$\max\left(\max_{j\in G}\sum_{i\in A}\overline{x}_{ij}, \max_{i\in A}\sum_{j\in G}\overline{x}_{ij}\right) \leq \frac{1}{1-\epsilon} + \frac{\epsilon}{1-\epsilon}\cdot\max\left(\max_{j\in G}\sum_{i\in A}\overline{x}_{ij}, \max_{i\in A}\sum_{j\in G}\overline{x}_{ij}\right)$$

or equivalently

$$\max\left(\max_{j\in G}\sum_{i\in A}\overline{x}_{ij},\max_{i\in A}\sum_{j\in G}\overline{x}_{ij}\right)\leq\frac{1}{1-2\epsilon}=1+O(\epsilon).\qquad\square$$

Proof of Theorem 4.1. By Lemma 4.5, we know that Algorithm 4.1 converges to an ϵ-approximately feasible solution of (4.7) in $O(\frac{n\log n}{\epsilon})$ many iterations. It is not hard to see that each iteration can be implemented in $O(n^2)$ time. It only remains to show that the output allocation $\overline{x}$ is indeed also an ϵ-approximate Nash bargaining point.

Let $x^\star$ be an actual Nash bargaining point, i.e. a maximizer of (4.2). We claim that

$$\prod_{i\in A}(u_i\cdot x_i^\star-c_i)\leq\prod_{i\in A}(u_i\cdot\overline{x}_i-c_i).$$

To see this, first observe that for any $i\in A$, we have

$$\frac{u_i\cdot x_i^\star-c_i}{u_i\cdot\overline{x}_i-c_i}\leq\frac{u_i\cdot x_i^\star-c_i}{\mathrm{CP}_i(\overline{p},\overline{q})}\leq\sum_{j\in G}(\overline{p}_j+\overline{q}_i)x_{ij}^\star-c_i\min_{j\in G}\frac{\overline{p}_j+\overline{q}_i}{u_{ij}}.$$

In essence, this is because $\mathrm{CP}_i(\overline{p},\overline{q})$ measures the optimal budget wrt. to prices $\overline{p},\overline{q}$ and so the utility of $x^\star$ can be bounded in terms of the cost of this bundle using the prices $\overline{p},\overline{q}$. Finally, we can use the AM-GM inequality to show

$$\begin{aligned}\prod_{i\in A}\frac{u_i\cdot x_i^\star-c_i}{u_i\cdot\overline{x}_i-c_i}&\leq\left(\frac{1}{n}\sum_{i\in A}\left(\sum_{j\in G}(\overline{p}_j+\overline{q}_i)x_{ij}^\star-c_i\min_{j\in G}\frac{\overline{p}_j+\overline{q}_i}{u_{ij}}\right)\right)^{\frac{1}{n}}\\&\leq\left(\frac{1}{n}\left(\sum_{j\in G}\overline{p}_j+\sum_{i\in A}\overline{q}_i-\sum_{i\in A}c_i\min_{j\in G}\frac{\overline{p}_j+\overline{q}_i}{u_{ij}}\right)\right)^{\frac{1}{n}}\\&=1.\end{aligned}\qquad\square$$

4.4 Conditional Gradient Descent

Since the problem of finding Nash-bargaining points can be captured by a convex program, we can leverage techniques from convex optimization such as gradient descent to compute rapidly converging approximations. In this section we will provide a conditional gradient descent algorithm which is relatively simple and projection free while converging at a rate of $O(1/\epsilon)$. This convergence is faster than the MWU algorithm from Section 4.3, however this comes at the cost of more expensive iterations and a dependence on the gap between the largest and smallest utilities.

We will split our analysis in to three parts. First, we will cover the simplest case: one-sided markets without endowments. Then we will extend our algorithm to non-bipartite markets; this also includes the two-sided setting as a special case. Lastly, we cover the extension to endowments. For endowments, we will only explicitly handle one-sided markets, though the results can easily be extended to the other settings as well.

4.4.1 One-Sided Markets without Endowments

Let us now consider the simplest setting of a one-sided market with without any endowments. We are given a set A of agents and G of goods with $|A| = |G| = n$ and utilities u_{ij} for all $i \in A$ and $j \in G$. Our goal is to solve the convex program (4.1) that we introduced in Section 4.2.

In the following we will assume that the utilities have been rescaled so that $\max\{u_{ij} \mid j \in G\} = 1$ for all i and $\sum_{j \in G} u_{ij} \geq \frac{1}{\kappa}$ for some value κ. In particular, the objective function is 1-strongly concave with respect to $u_i \cdot x_i$ but not x_{ij}. Note that since $u_i \cdot x_i$ is not bounded from below, the objective function is neither Lipschitz nor smooth.

Recall also that we have shown in Lemma 4.1 that

$$u_i \cdot x_i = \sum_{j \in G} u_{ij} x_{ij} \geq \frac{1}{2n} \sum_{j \in G} u_{ij} \geq \frac{1}{2\kappa n}.$$

for any optimal solution x. This is crucial to fix the lack of smoothness; we could restrict ourselves to the problem

$$\begin{aligned} \max \quad & \sum_{i \in A} \log(u_i \cdot x_i) \\ \text{s.t.} \quad & u_i \cdot x_i \geq \frac{1}{2\kappa n} \quad && \forall i \in A, \\ & \sum_{i \in A} x_{ij} \leq 1 && \forall j \in G, \\ & \sum_{j \in G} x_{ij} \leq 1 && \forall i \in A, \\ & \quad x \geq 0. \end{aligned}$$

without changing the optimal solution. However, while general purpose projected gradient descent algorithms can achieve $O(\log(1/\epsilon))$ convergence rates on this modified problem, we wish to exploit the combinatorial structure of the solution space. Therefore we will modify the objective function instead.

Definition 4.2. *The* quadratic extension of the logarithm *at* $x_0 > 0$ *is given by*

$$\eta(x; x_0) := \begin{cases} \log(x_0) + (x - x_0)\frac{1}{x_0} - \frac{1}{2}(x - x_0)^2 \frac{1}{x_0^2} & \text{if } x \leq x_0, \\ \log(x) & \text{otherwise.} \end{cases}$$

Note that $\eta(x; x_0)$ is $\frac{1}{x_0^2}$-smooth everywhere, i.e. its gradient is $\frac{1}{x_0^2}$-Lipschitz. This motivates the following convex program as an alternative to (4.1).

$$
\begin{aligned}
\max \quad & \sum_{i \in A} \eta\left(u_i \cdot x_i; \frac{1}{2\kappa n}\right) \\
\text{s.t.} \quad & \sum_{i \in A} x_{ij} \leq 1 \quad \forall j \in G, \\
& \sum_{j \in G} x_{ij} \leq 1 \quad \forall i \in A, \\
& x \geq 0.
\end{aligned} \tag{4.16}
$$

Lemma 4.6. *Every optimal solution to* (4.16) *is also optimal for* (4.1).

Proof. First let x be the optimal solution for (4.1). By Lemma 4.1, the objective function of the two programs agrees up to the first-order at x and thus x is also optimal for (4.16). But now any other optimal solution x' for the modified problem must satisfy $u_i \cdot x'_i = u_i \cdot x_i$ for all i since the objective is strictly concave in u_i. Thus it is also an optimal solution for (4.1). □

In the following let

$$
\begin{aligned}
\phi(x) &:= \sum_{i \in A} \log(u_i \cdot x_i), \\
\psi(x) &:= \sum_{i \in A} \eta\left(u_i \cdot x_i; \frac{1}{2\kappa n}\right)
\end{aligned}
$$

be the objective functions of (4.1) and (4.16) respectively. We need to show that Lemma 4.6 can be extended to hold for *approximate* solutions as well. In the following, we use the notation $u(x)$ to represent the vector $(u_i \cdot x_i)_{i \in A}$.

Lemma 4.7. *Let x be an arbitrary allocation and $x^\star$ an optimal one. Let $\psi(x^\star) - \psi(x) \le \epsilon$ for some $\epsilon > 0$. Then $||u(x) - u(x^\star)||^2 \le 2\epsilon$. Moreover, if $\delta \in (0, 1/2)$, and $\epsilon = \min\left\{\delta, \frac{\delta^{2/3}}{8n^2\kappa^2}\right\}$, then $\phi(x^\star) - \phi(x) \le O(1)\delta$.*

Proof. The first part follows directly from the fact that ψ is 1-strongly concave over the feasible region. Now observe that by Taylor's theorem we have

$$\eta\left(u_i \cdot x_i; \frac{1}{2n\kappa}\right) - \log(u_i \cdot x_i) \le \frac{2}{3}\frac{\left|u_i \cdot x_i - \frac{1}{2n\kappa}\right|^3}{(u_i \cdot x_i)^3}.$$

But because $u_i \cdot x_i^\star \ge \frac{1}{2n\kappa}$ and $||u(x) - u(x^\star)||^2 \le 2\epsilon$, we know that

$$\left|u_i \cdot x_i - \frac{1}{2n\kappa}\right| \le |u_i \cdot x_i - u_i \cdot x_i^\star|$$

and

$$u_i \cdot x_i \ge \frac{1}{2n\kappa} - \sqrt{2\epsilon} \ge (1 - \delta^{1/3})\frac{1}{2n\kappa}.$$

Finally, we compute

$$\begin{aligned}
\phi(x^\star) - \phi(x) &= \psi(x^\star) - \psi(x) + \psi(x) - \phi(x) \\
&\le \delta + \sum_{i \in A}\left(\eta\left(u_i \cdot x_i; \frac{1}{2n\kappa}\right) - \log(u_i \cdot x_i)\right) \\
&\le \delta + O(1)n^3\kappa^3 \sum_{i \in A} |u_i \cdot x_i - u_i \cdot x_i^\star|^3 \\
&\le \delta + O(1)n^3\kappa^3 ||u(x) - u(x^\star)||^{3/2} \\
&\le O(1)\delta.
\end{aligned}$$

□

These two lemmas together imply that it suffices to solve (4.16). Approximate solutions to this modified program are also approximate solutions for (4.1), both wrt. the Euclidean norm on the utility vectors and the original objective function. Finding an approximate solution in an efficient way way can be achieved using the conditional gradient method over the matching polytope; see Algorithm 4.2. Note that the gradient of ψ is easily computable.

Algorithm 4.2: CONDITIONAL GRADIENT DESCENT

$x^{(0)} \equiv 0$ **for** $t \leftarrow 1, \ldots, T$ **do**

  for $i \in A, j \in G$ **do**

    $w_{ij}^{(t)} \leftarrow \partial_{ij}\psi(x^{(t-1)})$

  $y^{(t)} \leftarrow$ max-weight matching with weights $w^{(t)}$ $x^{(t)} \leftarrow \left(1 - \frac{2}{t+1}\right) x^{(t-1)} + \frac{2}{t+1} y^{(t)}$

return $x^{(T)}$

THEOREM 4.3. *Algorithm 4.2 returns some x with $\psi(x^*) - \psi(x) \leq \epsilon$ in $O\left(\frac{n^3\kappa^2}{\epsilon}\right)$ many iterations. Each iteration can be implemented in $O(n^3)$ time.*

Proof. The conditional gradient algorithm converges in $O\left(\frac{D^2 L}{\epsilon}\right)$ iterations where D is the diameter of the polytope that is being optimized over and L is the smoothness of the objective function. For a modern proof of this fact, see for example [77]. In this case, the diameter of the matching polytope is $2\sqrt{n}$ and $L = 4n^2\kappa^2$ since $\eta\left(u_i; \frac{1}{2n\kappa}\right)$ is clearly $4n^2\kappa^2$-smooth. The amount of work in each iteration is $O(n^2)$ except for the computation of the max weight matching which can be done in $O(n^3)$ using the Hungarian method. □

We remark that it is in principle possible to achieve faster convergence rates for conditional gradient type algorithms by leveraging strong concavity in addition to smoothness [59]. Note that while ψ is strongly concave in the utilities, it is unfortunately not strongly concave in the allocation x.

Nevertheless, ψ can be written as $g(Ux)$ where g is strongly concave and Ux is the linear transformation of allocations to utilities. There are more complex variants of conditional gradient methods, for example those which involve taking "away steps" [78], which can be shown to converge in $O(\log(1/\epsilon))$ phases even in this more general setting. However, these algorithms depend on difficult to compute Hoffman-type constants of U relative to the matching polytope for which there is no known polynomial bound in the instance parameters.

4.4.2 Non-Bipartite Markets without Endowments

The result from the previous section can be extended to the non-bipartite setting. This also includes the two-sided setting as a special case. Recall that we have a set of n agents A with non-negative utilities u_{ij} for all $i, j \in A$. The goal is then to solve the convex program (4.5).

Assume that the utilities have been rescaled so that $\max\{u_{ij} \mid j \in A\} = 1$ for all i and $\sum_{j \in A} u_{ij} \geq \frac{1}{\kappa}$ for some value κ. Since it is possible to efficiently optimize over the matching polytope using combinatorial methods, the primary ingredient is to once again bound the utilities away from 0. Algorithm 4.2 does not have to be modified; one simply has to use the appropriate matching algorithm and surrogate objective function.

Lemma 4.8. *Let x be an optimal solution to (4.5), then for all agents i we have*

$$u_i \cdot x_i \geq \frac{1}{2n^2} \sum_{j \in A \setminus \{i\}} u_{ij} \geq \frac{1}{2\kappa n^2}.$$

Proof. The proof will be similar to the proof of Lemma 4.1, however the more complicated KKT conditions will yield a weaker bound. So let p and z be the optimal dual variables for our solution of (4.5).

Using the KKT conditions (see Section 4.2), we know that

$$\frac{u_{ij}}{u_i \cdot x_i} + \frac{u_{ji}}{u_j \cdot x_j} \leq p_i + p_j + \sum_{B \in \mathcal{O} : \{i,j\} \subseteq B} z_B$$

for all edges $\{i, j\} \subseteq A$ with equality when $x_{ij} > 0$. We use the notation $\mathcal{O}$ to denote the collection of odd subsets of A. Therefore

$$\begin{aligned}
\sum_{i \in A} p_i + \sum_{B \in \mathcal{O}} \frac{|B| - 1}{2} z_B &= \sum_{\{i,j\} \subseteq A} \left(p_i + p_j + \sum_{B \in \mathcal{O} : \{i,j\} \subseteq B} z_B \right) x_{ij} \\
&= \sum_{\{i,j\} \subseteq A} \left(\frac{u_{ij} x_{ij}}{u_i \cdot x_i} + \frac{u_{ji} x_{ij}}{u_j \cdot x_j} \right) \\
&= n.
\end{aligned}$$

In addition, we can observe from the KKT conditions that

$$\begin{aligned}
u_i \cdot x_i &= \max \left\{ \frac{u_{ij}}{p_i + p_j + \sum_{B \in \mathcal{O} : \{i,j\} \subseteq B} z_B - u_{ji} / (u_j \cdot x_j)} \;\middle|\; j \in A \setminus \{i\} \right\} \\
&\geq \max \left\{ \frac{u_{ij}}{p_i + p_j + \sum_{B \in \mathcal{O} : \{i,j\} \subseteq B} z_B} \;\middle|\; j \in A \right\} \\
&\geq \frac{\sum_{j \in A} u_{ij}}{n p_i + \sum_{j \in A \setminus \{i\}} p_j + \sum_{B \in \mathcal{O} : i \in B} (|B| - 1) z_B} \\
&\geq \frac{1}{2n^2} \sum_{j \in A \setminus \{i\}} u_{ij}.
\end{aligned}$$

□

This bound may initially seem weak when compared to the $\frac{1}{2n} \sum_{j \in G} u_{ij}$ bound from Lemma 4.1. However, this is actually tight up to a constant factor. Consider for some parameter $\ell \in \mathbb{N}$, an instance that has $2\ell + 1$ agents. Agents $1, \ldots, \ell$ have utility 1 for agent $2\ell + 1$ and 0 everywhere else. Agents $\ell + 1, \ldots, 2\ell$ are indifferent, i.e have utility 1 for all agents. Finally, agent $2\ell + 1$ has utility 1 for agents $\ell + 1, \ldots, 2\ell$. One can then show that in the optimal allocation x, agent $2\ell + 1$ will only get $\frac{1}{\ell+1}$ units from their utility 1 edges. Therefore

$$\frac{\sum_{j \in A} u_{ij}}{u_i \cdot x_i} = \ell(\ell + 1) \approx \frac{n^2}{4}.$$

However, this utility bound is still adequate in order to show convergence of the conditional gradient method with the modified objective function

$$\psi(x) \coloneqq \sum_{i \in A} \eta \left(u_i \cdot x_i ; \frac{1}{2n^2\kappa} \right).$$

Theorem 4.4. *Algorithm 4.2 returns some x with $\psi(x^\star) - \psi(x) \leq \epsilon$ in $O\left(\frac{n^5\kappa^2}{\epsilon}\right)$ many iterations where $x^\star$ is an optimal solution to* (4.5)*. Each iteration can be implemented in $O(n^3)$ time.*

Proof. The difference in the number of iterations compared to Theorem 4.3 comes from the fact that ψ is now $4n^4\kappa^2$-smooth. Each iteration can still be implemented in $O(n^3)$ time by using a weighted matching algorithm, now for non-bipartite graphs. □

4.4.3 Extension to Endowments

Finally, let us consider the case in which agents come with preexisting endowments, i.e. each agent i has some disagreement utility c_i and the goal is then to solve the convex program (4.2). We will only consider the one-sided model here since the extension to two-sided and non-bipartite markets works much the same.

The added difficulty of this setting is that, in general, a feasible solution may not exist. We thus assume that there exists a feasible solution $(\hat{x}, \delta)$ of the LP

$$\begin{aligned}
\max \quad & \delta \\
\text{s.t.} \quad & u_i \cdot x_i \geq (1+\delta) c_i && \forall i \in A, \\
& \sum_{i \in A} x_{ij} \leq 1 && \forall j \in G, \\
& \sum_{j \in G} x_{ij} \leq 1 && \forall i \in A, \\
& x \geq 0, \\
& \delta \geq 0.
\end{aligned} \tag{4.17}$$

with $\delta > 0$. We call this δ the *feasibility gap* of the instance.

In practice there are two likely scenarios as to how one might obtain δ. Either one solves the above linear program to find the optimal δ or the disagreement utilities c_i are defined as some constant fraction of the agents' utilities over their initial endowments. More precisely, assume that each agent i comes to the market with endowments e_{ij} over the goods j. Then one may simply define $c_i := \frac{1}{1+\delta} u_i \cdot e_i$. The market will then guarantee that no agent gets worse by a factor of more than $(1 + \delta)$ which is an approximate notion of individual rationality.

Regardless of how one obtains a feasible solution to (4.17), we may use it to derive a similar lower bound as in Lemma 4.1. Once again, assume that the utilities have been rescaled so that $\max\{u_{ij} \mid j \in G\} = 1$ for all i and $\sum_{j \in G} u_{ij} \geq \frac{1}{\kappa}$ for some $\kappa > 0$.

Lemma 4.9. *Let x be an optimal solution to* (4.2)*, then for all agents i we have*

$$u_i \cdot x_i \geq c_i + \frac{1}{2n^2(1+1/\delta)} \sum_{j \in G} u_{ij} \geq c_i + \frac{1}{2n^2(1+1/\delta)\kappa}.$$

Proof. Let p_j and q_i as always be optimal dual variables for the matching constraints in (4.2). Then by the KKT conditions we have that $\frac{u_{ij}}{u_i \cdot x_i - c_i} \leq p_j + q_i$ with equality if $x_{ij} > 0$.

From this complementarity we can deduce that

$$u_i \cdot x_i - c_i = \max \left\{ \frac{u_{ij}}{p_j + q_i} \,\middle|\, j \in G \right\}$$

and

$$\frac{u_i \cdot x_i}{u_i \cdot x_i - c_i} = \sum_{j \in G} x_{ij}(p_j + q_i).$$

Together this implies that

$$\sum_{j \in G} x_{ij}(p_j + q_i) = 1 + c_i \min \left\{ \frac{p_j + q_i}{u_{ij}} \,\middle|\, j \in G \right\}.$$

Now we may use our feasible solution $(\hat{x}, \delta)$ to (4.17) in order to bound

$$\begin{aligned}
\sum_{j \in G} p_j + \sum_{i \in A} q_i &= \sum_{i \in A} \sum_{j \in G} x_{ij}(p_j + q_i) \\
&= n + \sum_{i \in A} c_i \min \left\{ \frac{p_j + q_i}{u_{ij}} \,\middle|\, j \in G \right\} \\
&\leq n + \frac{1}{1+\delta} \sum_{i \in A} \sum_{j \in G} u_{ij} \hat{x}_{ij} \min \left\{ \frac{p_{j'} + q_i}{u_{ij'}} \,\middle|\, j' \in G \right\} \\
&\leq n + \frac{1}{1+\delta} \sum_{i \in A} \sum_{j \in G} \hat{x}_{ij}(p_j + q_i) \\
&\leq n + \frac{1}{1+\delta} \left(\sum_{j \in G} p_j + \sum_{i \in A} q_i \right)
\end{aligned}$$

which implies in particular that $\sum_{j \in G} p_j + \sum_{i \in A} q_i \leq \left(1 + \frac{1}{\delta}\right) n$.

Finally, using the same idea as in Lemma 4.1, we get

$$
\begin{aligned}
u_i \cdot x_i &= c_i + \max\left\{ \frac{u_{ij}}{p_j + q_i} \;\middle|\; j \in G \right\} \\
&\geq c_i + \frac{\sum_{j \in G} u_{ij}}{n q_i + \sum_{j \in G} p_j} \\
&\geq c_i + \frac{1}{2n^2(1 + 1/\delta)} \sum_{j \in G} u_{ij}.
\end{aligned}
$$

□

The modified convex program will thus use the objective function

$$
\psi(x) := \sum_{i \in A} \eta\left(u_i \cdot x_i - c_i; \frac{1}{2n^2(1 + 1/\delta)\kappa}\right)
$$

and we can now use Algorithm 4.2 to solve the modified convex program with the new gradients

$$
\partial_{ij}\psi(x) = \begin{cases} \frac{u_{ij}}{u_i \cdot x_i} & \text{if } u_i \cdot x_i \geq \frac{1}{2n^2(1+1/\delta)\kappa} \\ 4n^4(1 + 1/\delta)^2\kappa^2 u_i \cdot x_i u_{ij} & \text{otherwise.} \end{cases}
$$

Theorem 4.5. *Algorithm 4.2 returns x with $\psi(x^\star) - \psi(x) \leq \epsilon$ in $O\left(\frac{n^5(1+1/\delta)^2\kappa^2}{\epsilon}\right)$ many iterations where $x^\star$ is an optimal solution to* (4.2)*. Each iteration can be implemented in $O(n^3)$ time.*

Proof. The only difference to Theorem 4.3 is that the objective function ψ is now $4n^4(1 + 1/\delta)^2\kappa^2$-smooth instead of $4n^2\kappa^2$ as before. □

We remark that an analogous result to Lemma 4.7 holds also in this setting, i.e. convergence in the modified objective ψ implies convergence of the agents' utilities in the Euclidean norm. Thus, if one chooses ϵ on the order of $O\left(\frac{1}{n\sqrt{(1+1/\delta)\kappa}}\right)$, then one is guaranteed that $u_i \cdot x_i \geq c_i$ for all agents i.

4.5 Discussion

In this chapter, we have introduced two different algorithms for Nash bargaining in cardinal-utility matching markets which are practical, simple, and have good theoretical running time guarantees. Our work leaves three interesting open problems for future research:

- Note that our multiplicative weights approach only applies to one-sided and bipartite two-sided markets. Can it be extended to non-bipartite markets?
- Since our work focuses on the theoretical guarantees, it would also be interesting to see an actual practical comparison of high quality implementations of these techniques. Note that the work by Hosseini and Vazirani [70] provides an implementation of conditional gradient descent which also uses some of the insights from this chapter. However, they did not implement the MWU algorithm to compare against.
- Both of our algorithms extend techniques that can be used for Fisher market equilibria. There are other techniques such as tâtonnement [104] and proportional response dynamics [127] which also converge towards such equilibria. Can they be extended to our setting as well?

Part II

Online Matching

Chapter 5

Background

5.1 Introduction

In the *Online Bipartite Matching Problem* (OBMP), we are given a bipartite graph $G = (S, B, E)$ consisting of a set B of *buyers* and a set S of *sellers* (or, alternatively, *goods*). The buyers arrive online, one by one and in adversarial order, and whenever a buyer $i \in B$ arrives, the set $N(i)$ of its neighbors among the sellers is revealed. At that point we must match i immediately and irrevocably to a currently unmatched neighbor or leave i unmatched forever. Our goal is to maximize the total size of the matching.

The problem was introduced in the seminal work by Karp, Vazirani, and Vazirani [82] who proposed the Ranking algorithm as solution to it and showed that it is $(1 - 1/e)$-competitive in expectation (see Sections 5.2 and 5.3 for details). Over the years, online matching problems have received a large amount of interest due to the vast number of applications created by the internet and mobile computing. Online advertising alone poses the AdWords Problem [99, 74, 121] that lies at the heart of a multi-billion dollar market. Another interesting application, the Fully Online Matching Problem [71, 73], came about due to the rise of ride-sharing and/or ride-hailing apps such as Uber and Lyft where riders and drivers come online and need to be matched almost instantaneously while minimizing some function of latency and distance traveled to the rider. For a more complete overview of online matching and its place in matching-based market design, we refer to [98] and [40].

In the remainder of this chapter we will give an overview of some of the core concepts and definitions surrounding the Online Bipartite Matching Problem, including the notion of competitiveness and the RANKING algorithm. Then, in the following chapters, we will present several novel results on online matching.

In Chapter 6 we will extend the classic $(1 - 1/e)$-competitiveness result for RANKING and show that it holds *with high probability* as opposed to just *in expectation*. As we shall see, the proof is a remarkably natural and beautiful application of the method of bounded differences. We also extend this result to several other variants of online matching which will be defined later.

In Chapter 7 we will look at generalizations of online matching on *hypergraphs*, i.e. where the online vertices must be matched to sets (hyperedges) of offline vertices. We will give new upper and lower bounds while focusing on the case in which the hypergraph has large rank.

5.2 Competitiveness

Before we can study online matching, we need to fix some constraints and understand what our actual objective is. It is in general impossible to perfectly solve the OBMP, i.e. we cannot guarantee that we always find the largest possible matching in the underlying graph G. When faced with hard problems in computer science, we often turn to *approximations*.

However, in the case of online matching it is not entirely clear with respect to what solution the approximation guarantee should be measured.

The most common metric is α-competitiveness which measures the quality of a proposed online algorithm relative to an oracle with perfect information of the arrival of future vertices. In other words, this means that our benchmark is the maximum matching in the underlying graph G. A formal definition is given below.

Definition 5.1. *An algorithm $\mathcal{A}$ is α-competitive if for any instance $G = (S, B, E)$ of the Online Bipartite Matching Problem and any possible arrival order of the buyers B, $\mathcal{A}$ outputs a matching of size at least $\alpha \cdot \mathrm{OPT}$ where OPT is the size of a maximum matching in G.*

It is not immediately obvious that α-competitiveness for any positive α can even be achieved. However, in the case of the OBMP this follows from some classic results about matchings. Recall that a matching M is called *maximal* if we cannot add any other edge to it without intersecting an edge in M. Recall also that the size of any *maximal* matching is at least half the size of any *maximum* matching. We call an algorithm $\mathcal{A}$ *greedy* if it matches buyers whenever they have at least one unmatched neighbor when they arrive. Observe that any greedy algorithm always produces a maximal matching and is thus $\frac{1}{2}$-competitive. In fact, one can show that for the OBMP there is never a reason not to be greedy since the most that we can gain from not matching a vertex now is to match one more vertex later.

Theorem 5.1 (Karp et al. 1990 [82]). *No deterministic algorithm can be α-competitive for the OBMP for $\alpha > \frac{1}{2}$.*

Proof. Let $\mathcal{A}$ be some candidate algorithm. Consider the example instance shown in Figure 5.1 which clearly has $\mathrm{OPT} = 2$. Let i_1 arrive first.

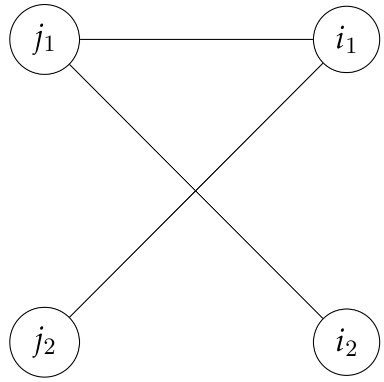

Figure 5.1: Shown is an example demonstrating the upper bound of $\frac{1}{2}$-competitiveness for deterministic online algorithms for the OBMP.

First, consider what happens if $\mathcal{A}$ does not match i_1 at all. In that case, consider a different instance in which i_2 does not exist. Here $\mathrm{OPT} = 1$ and $\mathcal{A}$ gets nothing so the competitive ratio of $\mathcal{A}$ would be 0.

Next, if $\mathcal{A}$ matches i_1 to j_1, then it cannot later match i_2 and so the competitive ratio will be at most $\frac{1}{2}$. The same applies of $\mathcal{A}$ matches i_1 to j_2: simply consider a different instance in which i_2 has an edge to j_2 instead of j_1. The key point here is that all three mentioned instances are indistinguishable from the perspective of $\mathcal{A}$ when i_1 arrives and since $\mathcal{A}$ is deterministic, it must make the same decision in all of them.

Thus the overall competitive ratio of $\mathcal{A}$ can be at most $\frac{1}{2}$. □

A key observation made by Karp et al. [82] was that while *deterministic* aglorithms cannot be more than $\frac{1}{2}$-competitive, *randomized* algorithms can. Let us first restate our definition of competitiveness for randomized algorithms.

Definition 5.2. *A randomized algorithm $\mathcal{A}$ is α-competitive (in expectation) if for any instance $G = (S, B, E)$ of the Online Bipartite Matching Problem and any possible arrival order of the buyers B, the* expected *size of the matching produced by $\mathcal{A}$ is at least $\alpha \cdot \mathrm{OPT}$ where OPT is the size of a maximum matching in G.*

Now we can see that we can in fact beat the $\frac{1}{2}$-competitiveness bound in the example from Figure 5.1. When i_1 arrives, we match it uniformly at random to either j_1 or j_2. The expected size of the matching is then $\frac{3}{2}$ which is $\frac{3}{4}\mathrm{OPT}$.

The key question is how to extend this result to arbitrary instances. The most straightforward generalization is to always match buyers to available sellers uniformly and independently at random. This algorithm is called Random and while it is slightly better than deterministic algorithms for small instances, its competitive ratio is unfortunately still just $\frac{1}{2}$ in general [82].

5.3 Ranking

Nonetheless, it is possible to improve on the $\frac{1}{2}$-competitiveness through a clever use of randomness in the form of the Ranking algorithm (Algorithm 5.1). The idea and namesake behind the algorithm is to start by computing a uniformly random *ranking* on the sellers. When a buyer arrives, they are matched to the highest ranked seller available.

Algorithm 5.1: Ranking (with permutation)

```
Sample a uniformly random permutation π on S.
for each buyer i who arrives do
    Match i to the first unmatched buyer in N(i) wrt. to π.
```

The reason why Ranking is better than Random is not obvious. Intuitively, it has to do with the fact that Ranking has a self-correcting nature: it is more likely to match to a seller who has had fewer opportunities to be matched in the past. This is because a seller who has had many opportunities to be matched is likely worse in the sampled ranking since otherwise it would have *already* been matched.

It turns out that this feature of Ranking boosts its competitive ratio to $1 - \frac{1}{e} \approx 0.63$. However, the original proof of this fact due to Karp et al. [82] is rather complicated. We will present the primal-dual analysis of Ranking which is based on work by Devanur et al. [35] and Eden et al. [44]. Chapter 6 and Chapter 7 will employ similar techniques.

The first step is to cleverly reframe the sampling of the ranking of the sellers. Instead of sampling a random permutation we instead sample *prices*. Almost any distribution works here and leads to the same algorithm but for the analysis, a particular choice is optimal. See Algorithm 5.2.

Algorithm 5.2: Ranking (with prices)

for *each seller j* **do**
    Sample $y_j \in [0, 1]$ uniformly at random.
    Set $p_j := e^{y_j - 1}$.
for *each buyer i who arrives* **do**
    Match i to the unmatched buyer $j \in N(i)$ minimizing p_j.

Next, we follow the terminology by Eden et al. [44]. We call p_j the *price* of seller j. We introduce dual variables $(r_j)_{j \in S}$ (revenues) and $(u_i)_{i \in B}$ (utilities). When buyer i is matched to seller j, we set $r_j = p_j$ and $u_i = 1 - p_j$. The revenues of sellers who are never matched are set to 0 and likewise the utilities of buyers who are not matched is set to 0 as well. Observe that the total amount of revenue and utility is exactly the size of the matching produced by Ranking.

Lemma 5.1. *Let $(i, j) \in E$ be arbitrary and fix all prices except for p_j. Let $u^\star$ be the utility of buyer i under the Ranking if seller j were removed from the list of sellers. Then:*

1. *No matter what p_j is, $u_i \geq u^\star$.*

2. *If $p_j < 1 - u^\star$, then j will be matched.*

Proof (Sketch). For Claim 1, the idea is to show that adding a new seller only improves the utilities of all the buyers. This is called the *monotonicity* of the *posted price mechanism* and can be shown via induction. For details we refer to [44], though we will give a similar argument later in Chapter 6.

Claim 2 follows from the observation that $1 - u^\star$ is the price of the seller that i would have been matched to with j missing. So if $p_j < 1 - u^\star$ and j has not been matched by the time that i arrives, then j would just get matched to i. □

Lemma 5.2. *Let $(i, j) \in E$ be arbitrary. Then $\mathbb{E}[u_i + r_j] \geq 1 - \frac{1}{e}$.*

Proof. Fix all prices except for p_j and once again let $u^\star$ be the utility of i if j were removed from the instance. Part 1 of Lemma 5.1 implies that $\mathbb{E}[u_i] \geq u^\star$ where the expectation is now only over p_j.

Now let $y^\star$ be such that $1 - e^{y^\star - 1} = u^\star$. Then by part 2 of Lemma 5.1, we know that if $y_j < y^\star$, j will be matched since in that case $p_j < 1 - u^\star$. Therefore

$$\mathbb{E}[r_j] \geq \int_0^{y^\star} e^{y_j - 1}\, \mathrm{d}y_j = e^{y^\star - 1} - \frac{1}{e} = 1 - u^\star - \frac{1}{e}.$$

Summing up both of these bounds yields $\mathbb{E}[r_j + u_i] \geq 1 - \frac{1}{e}$ as claimed. □

Theorem 5.2. *Ranking is $(1 - \frac{1}{e})$-competitive.*

Proof. Let M be some maximum matching. As already noted, the size of the matching produced by Ranking is given by $\sum_{j \in S} r_j + \sum_{i \in B} u_i$. Now simply observe that by Lemma 5.2, we have

$$\mathbb{E}\left[\sum_{j \in S} r_j + \sum_{i \in B} u_i\right] \geq \sum_{(i,j) \in M} \mathbb{E}[r_j + u_i] \geq \sum_{(i,j) \in M} \left(1 - \frac{1}{e}\right) = \left(1 - \frac{1}{e}\right) \mathrm{OPT}.$$ □

Ranking has been massively influential in the online matching literature and there are many variants of the algorithm for various related problems, some of which we will discuss in Chapter 6. However, it is not the only popular algorithmic technique in the field. The other two notable algorithms are Water Filling / Balance and the rounding algorithms based on Online Correlated Selection (OCS) [48].

5.4 Water Filling and Balance

Ranking is the most prevalent algorithmic technique for unweighted or vertex weighted, *integral* variants of online matching. However, if the matching is allowed to be *fractional*, the Water Filling algorithm is a popular alternative which, unlike Ranking, can be extended to work with edge weights as well [34]. The idea goes back to Kalyanasundaram and Pruhs [81] who developed the Balance algorithm that we will discuss later.

For now, we will consider the Fractional Online Bipartite Matching Problem (FOBMP). The setup is the same, i.e. we have an underlying graph $G = (S, B, E)$ with sellers and buyers. Buyers arrive online in adverserial order and must be matched immediately and irrevocably. However, we are now allowed to match buyers fractionally to their neighbors when they arrive. The goal is thus to determine non-negative $(y_{ij})_{(i,j)\in E}$ such that $\sum_{j\in S} y_{ij} \leq 1$ for all $i \in B$ and $\sum_{i\in B} y_{ij} \leq 1$ for all $j \in S$.

Recall that the size of the maximum integral matching and the size of the maximum fractional matching are the same in bipartite graphs. So the offline optimum OPT does not change between the integral and fractional variants of the problem. Therefore, the FOBMP is in general an easier problem than the OBMP.

The challenge in online matching in general is that when a buyer arrives, we do not know which of the sellers are still needed for future buyers. The idea behind Water Filling is quite natural: we simply try to be as conservative as possible. Whenever a buyer i arrives, we start matching it to its neighbor with the lowest *fill level*, i.e. the amount that they are currently matched. This way we delay filling up a seller completely by as much as possible. If multiple neighbors of i have the same fill level, we will fill them up at the same rate.

Another way to think of this is in terms of a continuous process. As long as an arriving buyer is not fully matched, we match infinitesimal amounts $\mathrm{d}t$ to whichever seller is least matched. Formally, this is Algorithm 5.3.

Algorithm 5.3: WATER FILLING

for $e \in E$ **do**
    Initialize $y_e := 0$.
for *each buyer i who arrives* **do**
    while *i is not fully matched and has a neighbor who is not fully matched* **do**
        Let $j \in N(i)$ be a neighbor minimizing $\sum_{i' \in N(j)} y_{i'j}$.
        Increase y_{ij} by some infinitesimal $\mathrm{d}t$.

At this point, we should briefly discuss the issue of computational complexity. In online algorithms in general, there is often little focus on running times since the main challenge is usually to achieve the best possible competitive ratio, regardless of running time. In the case of RANKING, it is clear that the algorithm can be implemented very efficiently. However, Algorithm 5.3 describes a process which by its very nature of using infinitesimals cannot be directly implemented like this at all.

It is not too hard to see that Algorithm 5.3 can be turned into a "proper" algorithm; the updates that need to be applied to y for each buyer's arrival can be computed in polynomial time. However, for more elaborate variants of WATER FILLING, it is not always trivial to come up with an algorithm that can be implemented, let alone in polynomial time. In that case, one can approximate the process by replacing $\mathrm{d}t$ with some small Δt instead.

We will now give a short proof, based on an analysis by Devanur et al. [34], that Water Filling is $(1 - \frac{1}{e})$-competitive for the FOBMP. This might initially seem disappointing. After all, the FOBMP is supposed to be *easier* than the OBMP and yet Water Filling does not perform any better than Ranking! Note though that Water Filling is *deterministic* whereas Ranking is *randomized*. Indeed, we showed in Section 5.2 that deterministic algorithms for the OBMP cannot be more than $\frac{1}{2}$-competitive. More importantly, Water Filling generalizes better to more complex settings which is what we will make use of in Chapter 7.

Just like for the analysis of Ranking, we will use the terminology of Eden et al. [44]. For each seller $j \in S$ we will have a revenue $r_j \geq 0$ and likewise for each buyer $i \in B$ we will have a utility $u_i \geq 0$. The instantaneous price p_j of a seller is $e^{x_j - 1}$ where $x_j = \sum_{i \in N(j)} y_{ij}$ is the fill level. When i is matched to j by $\mathrm{d}t$, we increase r_j by $p_j \mathrm{d}t$ and u_i by $(1 - p_j)\mathrm{d}t$. Note how we still maintain the invariant that the size of the matching produced by the algorithm is always equal to $\sum_{j \in S} r_j + \sum_{i \in B} u_i$.

Lemma 5.3. *At the end of the algorithm, we have $u_i + r_j \geq 1 - \frac{1}{e}$ for any $(i, j) \in E$.*

Proof. Let x_j be the final fill level of seller j. Note that

$$r_j = \int_0^{x_j} e^{t-1} \, \mathrm{d}t = e^{x_j - 1} - \frac{1}{e}.$$

Now consider what happened when i was getting matched during the algorithm. Since the final fill level of j is x_j, i would have gotten matched exclusively to sellers with a fill level of at most x_j. Thus for each $\mathrm{d}t$ its utility would have increased by at least $1 - e^{x_j - 1}$. Therefore:

$$u_i \geq \int_0^1 (1 - e^{x_j - 1}) \, \mathrm{d}t = 1 - e^{x_j - 1}.$$

Finally, summing up these two inequalities yields $r_j + u_i \geq 1 - \frac{1}{e}$ as claimed. □

Theorem 5.3. *Water Filling is* $(1 - \frac{1}{e})$*-competitive.*

Proof. The proof is now essentially identical to the proof of Theorem 5.2. □

Lastly, Water Filling can also be used in settings where the matching is not fractional but in which the sellers can be matched many times. If each seller can be matched exactly b times, this is called the Online Bipartite b-Matching Problem and essentially approximates the FOBMP. This is actually where the idea originated in the form of the Balance algorithm; see Algorithm 5.4.

Algorithm 5.4: Balance

```
for each buyer i who arrives do
    Match i to a neighbor which has been matched the least (if possible).
```

Kalyanasundaram and Pruhs [81] showed that this algorithm is $(1 - (1 + \frac{1}{b})^{-b})$-competitive which approaches $1 - \frac{1}{e}$ as $b \to \infty$. The modern analysis is quite similar to that of Water Filling with the price of a seller being $(1 + \frac{1}{b})^{x_j - b}$ if they have been matched x_j many times.

5.5 Upper Bounds via Yao's Principle

We have seen that Ranking is $(1 - \frac{1}{e})$-competitive for the OBMP. But can we do any better? This question was asked and answered already in the original Ranking paper [82] and it turns out the answer is: no, we cannot. In this last section we will see a common upper bounding technique used in online algorithms which we will make use of in Chapter 7. It is based on the following useful lemma.

Lemma 5.4 (Yao's Principle). *Let α be the competitive ratio of any randomized algorithm for the OBMP. Let Δ be a distribution over instances of the OBMP consisting of both the graph $G = (S, B, E)$ and the arrival order of the buyers. Finally, let β be the expected competitive ratio of the best deterministic online algorithm on the distribution Δ. Then $\alpha \leq \beta$.*

Yao's principle comes from a game theoretic view on the competitive ratio. We can think of the competitive ratio as being determined by a zero-sum game between two players: the algorithm player and the instance player. The algorithm player tries to create the best possible online algorithm whereas the instance player tries to generate the most challenging instance. By the well-known minimax principle due to von Neumann, as long as the players are allowed to use mixed strategies (i.e. randomized algorithms and distributions of instances), the order of the players does not matter. Thus, Yao's principle follows.

Theorem 5.4 (Karp et al. 1990 [82]). *There is no $(1 + \frac{1}{e} + \epsilon)$-competitive algorithm for the OBMP for any $\epsilon > 0$.*

Proof. Clearly we are going to use Yao's principle, i.e. we need to come up with a distribution over difficult instances and show that no deterministic algorithm does well on them. Let us fix some large $n \in \mathbb{N}$ to determine the size of the instances, i.e. there will be n buyers and n sellers. The instances are generated as follows.

For the first arriving buyer, we add edges to all sellers. For the second arriving buyer, we connect it to $n - 1$ sellers which are chosen uniformly at random out of the n total sellers. In general, for the i'th arriving buyer, we connect it to a subset of $n - i - 1$ sellers which is chosen uniformly among the $n - i$ sellers that the $(i - 1)$'th buyer had access to.

Now it is easy to see that $\mathrm{OPT} = n$, i.e. there is a perfect matching. Each buyer can be matched to the unique seller which is no longer available to the next arriving buyer.

Consider some optimal deterministic algorithm $\mathcal{A}$. As noted previously, we may assume without loss of generality that $\mathcal{A}$ is greedy.

We are going think of $\mathcal{A}$ as a stochastic process with time $t \geq 0$. At each integral t, a buyer arrives and $\mathcal{A}$ matches it if possible. We keep track of two random variables: $x(t)$ is simply the number of buyers who have not yet been considered by $\mathcal{A}$ and $y(t)$ is the number of sellers which are still eligible to be matched. By "eligible" we refer to the sellers which have not been matched previously and for which there will still be edges to other buyers later.

Note that $\Delta x(t) = x(t) - x(t+1) = -1$ and $x(0) = n$. Likewise, $y(0) = n$. As long as $y(t) \geq 1$, it too decreases by at least one in each time step since $\mathcal{A}$ is greedy and will match a seller if possible. However, it *may* decrease by 2. Consider a buyer i which arrives at some time t. By construction, there is a unique seller $j \in N(i)$ which is not a neighbor of the buyers that arrive after i. If j is currently unmatched at time t and $\mathcal{A}$ does not match i to j, then j becomes ineligible in addition to the seller that i is actually matched to.

Crucially, due to the construction of our instances, the set of $y(t)$ sellers which is eligible at time t is equally likely to be any subset of size $y(t)$ of the $x(t)$ sellers that *could* be eligible. Moreover, the probability that $\mathcal{A}$ does not pick the one special seller which is going to become ineligible anyways is $\frac{y(t)-1}{y(t)}$. Together, this yields

$$\mathbb{E}[\Delta y(t)] = \mathbb{E}[y(t) - y(t+1)] = -1 - \frac{y(t)}{x(t)}\frac{y(t)-1}{y(t)}.$$

This allows us to set up the following stochastic difference equation:

$$\frac{\mathbb{E}[\Delta y(t)]}{\mathbb{E}[\Delta x(t)]} = 1 + \frac{y(t)-1}{x(t)}$$

By Kurtz' theorem [87], this is closely approximated by the solution to the corresponding differential equation

$$\frac{\mathrm{d}y}{\mathrm{d}x} = 1 + \frac{y(t) - 1}{x(t)}$$

with probability tending to 1 as $n \to \infty$. Its solution with the given boundary conditions of $x(0) = y(0) = n$ is

$$y = 1 + x\left(\frac{n-1}{n} + \ln\frac{x}{n}\right). \tag{5.1}$$

Finally, by solving (5.1), we can obtain that if $y = 1$, i.e. there is only one remaining eligible seller, we have $x = \frac{n}{e} + o(n)$. This implies that the competitive ratio of $\mathcal{A}$ on our distribution of instances is $1 - \frac{1}{e} - o(1)$ and by Yao's principle, the theorem follows. □

Lastly, we remark that the above proof can be easily extended to work for fractional algorithms as well: the optimum competitive ratio of the FOBMP is also $1 - \frac{1}{e}$.

Chapter 6

High Probability Guarantees

6.1 Introduction

In Chapter 5, we introduced the Online Bipartite Matching Problem and the Ranking algorithm which achieves the optimal competitive ratio of $1 - \frac{1}{e}$. To be more precise, we showed that this bound holds *in expectation* over the randomness used by Ranking. In this chapter, we will turn towards proving competitiveness *with high probability* instead. This chapter is based on the paper "Online Matching with High Probability" which was joint work with Milena Mihail [101].

For many online matching problems, there are extensions of Ranking which achieve competitive ratios of $1 - \frac{1}{e}$ or at the very least strictly greater than $\frac{1}{2}$. Often, these are best-known for their respective problems. However, all results on Ranking-like online matching algorithms in the literature only establish the competitive ratio *in expectation* without guaranteeing any form of concentration beyond the trivial bounds that follow from Markov's inequality.

Under the restriction that the graph is d-regular, Cohen and Wajc [30] proposed the MARKING algorithm for Online Bipartite Matching and showed that it has a competitive ratio of $1 - O(\sqrt{\log d}/\sqrt{d})$ in expectation and $1 - O(\log n/\sqrt{d})$ with high probability. They remark that this is the first high probability guarantee $> 1/2$ for Online Bipartite Matching, though only in this restricted setting. Accordingly, the analysis presented in this chapter provides the first such bound without additional assumptions on the problem instances.

The analysis of concentration bounds for randomized algorithms goes back to the 1970s with classic results such as the second moment bound for QUICKSORT [112]. See [38] for an extensive overview of the field. However, it has remained the case that in the analysis of algorithms, results are usually quantified in terms of expected solution quality only.

In some sense this is due to the well-known fact that, as a consequence of standard Chernoff bounds, any randomized algorithm which is good in expectation can be boosted to be good with probability $1 - \frac{1}{n}$ by simply repeating it $O(\log n)$ many times. But it is precisely in the case of online algorithms where this argument fails due to the fact that online algorithms, by definition, cannot be repeated. Despite this, the literature on high probability bounds for online algorithms is relatively sparse (for some exceptions, see e.g. [85, 90]). Given the impact that RANKING has had over the last 30 years, it is quite remarkable that such a fundamental aspect of it had been left unanswered.

6.1.1 Our Results and Techniques

Our results concern RANKING type algorithms in three different settings: the classic Online Bipartite Matching Problem (see Section 6.2), the Fully Online Matching Problem inspired by ride-sharing (see Section 6.3) and the Online Vertex-Weighted Bipartite Matching Problem inspired by the internet advertising markets (see Section 6.4).

In Section 6.2 we will show the following result, complementing the $\left(1-\frac{1}{e}\right)$-competitiveness result of Ranking for the Online Bipartite Matching Problem (Theorem 5.2).

Theorem 6.1. *Let $G = (S, B, E)$ be an instance of the Online Bipartite Matching Problem which admits a matching of size n. Then for any $\alpha > 0$ and any arrival order,*

$$\mathbb{P}\left[|M| < \left(1 - \frac{1}{e} - \alpha\right) n\right] < e^{-2\alpha^2 n}$$

where M is the random variable denoting the matching generated by Ranking.

The key technical ingredient for this result is a bounded differences property of the random variable $|M|$ (see Lemma 6.2). We prove this via structural properties of matchings (see Lemma 6.3) similar, to ones which have been used in previous analyses of Ranking [65, 16]. Together with McDiarmid's inequality shown below (a consequence of Azuma's inequality), this gives rise to a particularly natural proof of Theorem 6.1.

Lemma 6.1 (McDiarmid's Inequality [96]). *Let $c_1, \ldots, c_n \in \mathbb{R}_+$ and consider some function $f : [0,1]^n \to \mathbb{R}$ satisfying*

$$|f(x_1, \ldots, x_{i-1}, x_i', x_{i+1}, \ldots, x_n) - f(x_1, \ldots, x_n)| \le c_i$$

for all $x \in [0,1]^n$, $i \in [n]$ and $x_i' \in [0,1]$. Moreover let Δ^n be the uniform distribution on $[0,1]^n$. Then for all $t > 0$, we have

$$\mathbb{P}_{x \sim \Delta^n}\left[f(x) < \mathbb{E}_{y \sim \Delta^n}[f(y)] - t\right] < e^{-\frac{2t^2}{\sum_{i=1}^n c_i^2}}.$$

We want to contrast this technique briefly with two related results. The analysis of Marking by Cohen and Wajc [30] uses the d-regularity of the graph in an essential way. They are able to show directly that the probabilities that the offline vertices are unmatched are negatively correlated and apply a Chernoff bound. In fact, they even show that the probability that any given offline vertex is matched goes to 1 as $d \to \infty$ which is certainly not the case for Ranking.

The technique by Komm et al. [85] can be used to show concentration bounds for several problems which are loosely related to online matching such as the Online k-Server Problem. Their key idea is to use a repeating strategy where any existing randomized algorithm is used and simply restarted periodically when certain conditions are met. This improves the expectation guarantee of the original algorithm to a high probability guarantee similar to typical re-running technique for non-online algorithms. However, this only works if one can indeed cheaply restart the algorithm without harming the analysis which is the case in the Online k-Server Problem but not in the Online Bipartite Matching Problem.

In Section 6.3 we will define the Fully Online Matching Problem and the natural extension of Ranking for this setting. We remark that we allow for non-bipartite graphs here and we will give a similar concentration bound as in Theorem 6.1.

Theorem 6.2. *Let G be an instance of the Fully Online Matching Problem which admits a matching of size n. Then for any $\alpha > 0$,*

$$\mathbb{P}\left[|M| < (\rho - \alpha)\, n\right] < e^{-\alpha^2 n}$$

where M is the random variable denoting the matching generated by Ranking and ρ is the competitive ratio of Ranking for this setting.

We remark that by [71], we know $\rho > 0.521$ and for the special case where G is bipartite, we have $\rho = W(1) \approx 0.567$.

In Section 6.4 we will consider the Online Vertex-Weighted Bipartite Matching Problem. In this setting, a generalization of Ranking was shown to be $(1 - \frac{1}{e})$-competitive by Aggarwal et al. [2]. We will modify this algorithm to show the following.

Theorem 6.3. *For any $\alpha > 0$, there exists a variant of Ranking such that for any instance $G = (S, B, E)$ with weights $w : S \to \mathbb{R}_+$ of the Online Vertex-Weighted Bipartite Matching, any arrival order of B and any matching M^*,*

$$\mathbb{P}\left[w(M) < \left(1 - \frac{1}{e} - \alpha\right) w(M^*)\right] < e^{-\frac{1}{50}\alpha^4 \frac{w(M^*)^2}{||w||_2^2}}$$

where M denotes the matching generated by Ranking and

$$w(M) := \sum_{\{i,j\} \in M} w_j.$$

Lastly, we argue that this bound also applies to the Online Single-Valued Bipartite Matching Problem which is a variant of the vertex-weighted problem in which goods can be matched multiple times.

6.2 Online Bipartite Matching

It is common to analyze Ranking by replacing the sampling of the permutation π (see Algorithm 5.1) by independently sampling prices (see Algorithm 5.2). Note that the precise distribution of the prices does not matter for the actual performance of the algorithm. In the following we will simply sample an independent, uniform $x_j \in [0, 1]$ for every $j \in S$ called the *rank* of j. Then, sorting S by the values of x_j yields a uniformly random permutation. Formally, this is Algorithm 6.1.

Algorithm 6.1: Ranking (with uniform ranks)

for $j \in S$ **do**
  Sample a uniformly random $x_j \in [0, 1]$.

for *each buyer i who arrives* **do**
  Match i to an unmatched $j \in N(i)$ minimizing x_j.

In the following, consider a fixed graph $G = (S, B, E)$ with a fixed arrival order. Assume that $|S| = n$ and that G has a matching of size n. We define a function $f : [0, 1]^S \to \mathbb{R}$ by letting $f(y)$ be the size of the matching M generated by Algorithm 6.1 if $x_j = y_j$ for all $j \in S$. Our goal will then be to show the following lemma. It is a different perspective on a structural property that appears under various forms in the online matching literature (e.g. Lemma 2 in [16]).

Lemma 6.2 (Bounded Differences)**.** *Let $x \in [0, 1]^S$, $j^\star \in S$ and $\theta \in [0, 1]$ be arbitrary. Define x'_j to be θ if $j = j^\star$ and x_j otherwise. Then $|f(x) - f(x')| \le 1$.*

Note that Lemma 6.2 implies Theorem 6.1 via McDiarmid's inequality (Lemma 6.1). Specifically, by applying McDiarmid to the function f with $c \equiv 1$ we get

$$\begin{aligned}\mathbb{P}\left[|M| < \left(1 - \frac{1}{e} - \alpha\right) n\right] &\le \mathbb{P}_{x\sim\Delta^S}\left[f(x) < \mathbb{E}_{y\sim\Delta^S}[f(y)] - \alpha n\right] \\ &\le e^{-2\alpha^2 n}\end{aligned}$$

where we used that $(1 - \frac{1}{e})n \le \mathbb{E}_{y\sim\Delta^S}[f(y)]$ since Ranking is $(1 - \frac{1}{e})$-competitive. It remains to prove Lemma 6.2.

Lemma 6.3. *Let $j \in S$, then we can define the graph G_{-j} which contains all vertices of G except for j. For some fixed values of $x \in [0, 1]^S$, we let M be the matching produced by Ranking in G and let M_{-j} be the matching produced by Ranking in G_{-j}. Then $|M_{-j}| \le |M| \le |M_{-j}| + 1$.*

Proof. For any buyers $i, i' \in B$, let $N^{(i)}(i')$ be the set of neighbors of i' in G which are unmatched by the time that i arrives in the run of Ranking with the fixed ranks x. Likewise, let $N^{(i)}_{-j}(i')$ be the set of unmatched neighbors of i' in the run of Ranking on G_{-j} when i arrives. We claim that for all $i \in B$ there exists some $j' \in S$ such that for all $i' \in B$ we have $N^{(i)}(i') = N^{(i)}_{-j}(i')$ or $N^{(i)}(i') = N^{(i)}_{-j}(i') \cup \{j'\}$.

Let us show this claim via induction on $i \in B$ in order of arrival. Note that when the first buyer arrives, this holds for $j' = j$ because we have removed only j' from the graph and nobody has been matched yet. Now assume that the statement holds when i arrives, we need to see that it still holds after i has been matched. Clearly, if i gets matched to the same vertex in G and in G_{-j}, then the inductive step follows trivially.

So now assume that i gets matched to different vertices in G and in G_{-j}. By the inductive hypothesis this can only happen if i gets matched to j' in G and it gets matched to some other j'' (potentially $j'' = \perp$, i.e. it is not matched at all) in G_{-j}. But then $N^{(i+1)}(i') = N^{(i+1)}_{-j}(i')$ or $N^{(i+1)}(i') = N^{(i+1)}_{-j}(i') \cup \{j''\}$ for all $i' \in B$. Thus the claim holds by induction.

Finally, let us see that the claim implies the lemma. First note that since i always has more unmatched neighbors in G than in G_{-j}, we have $|M| \geq |M_{-j}|$. But on the other hand, if at some time in the algorithm i is matched to j' in G and not matched at all in G_{-j}, then we have that $N^{(i+1)}(i') = N^{(i+1)}_{-j}(i')$ for all $i' \in B$. Thus the two runs will be identical from that point onward and $|M| = |M_{-j}| + 1$. □

Finally, we can show that this implies the bounded differences property of f that we claimed in Lemma 6.2.

Proof of Lemma 6.2. By Lemma 6.3 we know that removing a good from the graph can decrease the size of the matching computed by Ranking by at most one, assuming that the values of the x_j are fixed. But of course if we are removing $j^\star \in S$, the matching $M_{-j^\star}$

computed by Ranking in $G_{-j^\star}$ does not depend on the value of $x_{j^\star}$ or $x'_{j^\star}$. So we have

$$|M_{-j^\star}| \leq f(x) \leq |M_{-j^\star}| + 1$$

and

$$|M_{-j^\star}| \leq f(x') \leq |M_{-j^\star}| + 1$$

which implies $|f(x) - f(x')| \leq 1$ as claimed. □

As we have already seen, this is enough to prove Theorem 6.2 in the case where $|S| = n$. To prove the general case we can use a simple reduction. In particular, assuming that there is a matching M of size n but $|S| > n$, let S_M be the goods covered by M and let $G_M = (S_M, B, E)$. We have seen in Lemma 6.3 that for any fixed $x \in [0,1]^S$, Ranking will produce a matching in G that is not smaller than the matching it produces in G_M when run with x restricted to S_M. Therefore, Theorem 6.2 on G_M implies Theorem 6.2 on G which establishes the general case.

6.3 Fully Online Matching

In the Fully Online Matching Problem we have a not necessarily bipartite graph G the vertices of which arrive and depart online in adversarial order. When a vertex arrives, it reveals all of its edges to vertices that have already arrived. By the time it departs, its entire neighborhood is guaranteed to have been revealed.

This problem was introduced by Huang et al. [71] and is motivated by ride-sharing. Each vertex represents a rider who, upon arrival, is willing to wait only for a certain amount of time. Two riders can only be matched if the time that they spend on the platform overlaps, even in the offline solution. This additional condition allows Huang et al. to show that the generalization of Ranking shown in Algorithm 6.2 is 0.521-competitive in general and 0.567-competitive on bipartite graphs.

Algorithm 6.2: Fully Online Ranking

---

for *vertex i who arrives* **do**
  Sample a uniformly random $x_i \in [0, 1]$.
for *vertex i who departs* **do**
  Match i to an unmatched $j \in N(i)$ minimizing x_j.

In order to show a concentration bound, we can apply similar techniques as in Section 6.2. Let $G = (V, E)$ be a graph which admits a perfect matching of size n. Then let $f : [0, 1]^V \to \mathbb{R}$ once again represent the size of the matching generated by Algorithm 6.2 when given the x_i values. The corresponding bounded differences condition then becomes:

Lemma 6.4 (Bounded Differences)**.** *Let $x \in [0, 1]^V$, $i^\star \in V$ and $\theta \in [0, 1]$ be arbitrary. Define x'_i to be θ if $i = i^\star$ and x_i otherwise. Then $|f(x) - f(x')| \le 1$.*

This implies Theorem 6.2 as before though note that this time we will lose a factor of 2 since we now have $2n$ variables. We remark that this follows directly from Lemma 2.3 in [71] but for completeness we will give a short proof sketch.

Lemma 6.5. *Using the notation from Lemma 6.3, we have $|M_{-j}| \le |M| \le |M_{-j}| + 1$ for any $j \in V$ and fixed values of $x \in [0, 1]^V$.*

Proof. As in the proof of Lemma 6.3, let $N^{(i)}(i')$ (or $N^{(i)}_{-j}(i')$) be the set of neighbors of i' in G (or G_{-j}) which is unmatched by the time that i departs in the run of Fully Online Ranking with the fixed values of x. We claim that for all $i \in V$, there exists some $j' \in V$ such that for all $i' \in V$, we have $N^{(i)}(i') = N^{(i)}_{-j}(i')$ or $N^{(i)}(i') = N^{(i)}_{-j}(i') \cup \{j'\}$.

This claim follows via an almost identical induction as in Lemma 6.3. Then, since i always has more unmatched neighbors in G than in G_{-j}, we have $|M| \geq |M_{-j}|$. And if at some time in the algorithm, i is matched to j' in G and not matched at all in G_{-j}, then we have that $N^{(i+1)}(i') = N_{-j}^{(i+1)}(i')$ for all $i' \in V$. Thus the two runs will the identical from that point onward and $|M| = |M_{-j}| + 1$. □

Since Lemma 6.5 implies Lemma 6.4, this yields Theorem 6.2 for graphs which contain a perfect matching. But as in Section 6.2, we may drop this condition by reducing a graph G with a matching M to the subgraph induced by the vertices covered by M. Adding the vertices back in only increases the performance of Fully Online Ranking by Lemma 6.5.

6.4 Online Vertex-Weighted Bipartite Matching

In this section, we will consider a weighted extension of the Online Bipartite Matching Problem which has been inspired by online advertising markets. In the Online Vertex-Weighted Bipartite Matching Problem, we have a bipartite graph $G = (S, B, E)$ with vertex weights $w : S \to \mathbb{R}_+$ on the offline vertices. Here S represents the advertisers and B represents website impressions or search queries which should get matched to ads from the advertisers. The vertices B arrive online in adversarial order and should get matched to a neighbor j such that the total weight of the matched vertices in S is maximized. This problem can be seen as a special case of the AdWords Problem which instead imposes edge-weights and budgets on the offline vertices.

Perhaps somewhat surprisingly, it took 20 years for Ranking to be extended for the unweighted to the vertex-weighted setting by Aggarwal et al. [2]. This is because in the presence of weights, it is no longer enough to pick a uniformly random permutation over the offline vertices. Instead, one has to skew the permutation so that higher weight vertices are more likely to appear first. This is done elegantly in Algorithm 6.3 by ordering the vertices not by their x_j but rather by the careful chosen quantity $w_j(1 - e^{x_j - 1})$.

Note the similarity to Algorithm 5.2. In the unweighted case, we employed a very similar pricing strategy to *analyze* the algorithm. However, this was not the original analysis of Karp et al. [82] so this generalization was not obvious at the time.

Algorithm 6.3: VERTEX-WEIGHTED RANKING

---

1 **for** $j \in S$ **do**
2   Sample a uniformly random $x_j \in [0, 1]$.
3 **for** *each buyer i who arrives* **do**
4   Match i to an unmatched $j \in N(i)$ maximizing $w_j \left(1 - e^{x_j - 1}\right)$.

Unfortunately, Algorithm 6.3 does not lend itself to an analysis via the method of bounded differences. This is because a vertex with small weight, which should have little impact on the total weight of the matching, can sometimes be chosen over a vertex with much larger weight. See the example shown in Figure 6.1.

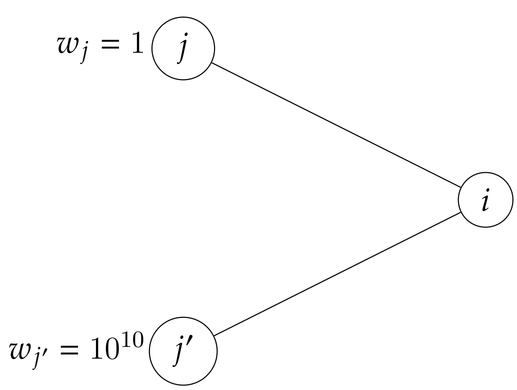

Figure 6.1: Shown is a simple instance in which the value of x_j can have a large impact on the final matching despite the fact that w_j is small. If $x_{j'} \gg 1 - 10^{-10}$, i will choose j in line 4 for sufficiently small values of x_j.

In particular, the problem lies with the fact that $w_j(1 - e^{x_j - 1})$ can get arbitrarily close to 0 if x_j gets close to 1. We will overcome this problem by changing the function slightly. For any $\epsilon > 0$, we consider ϵ-Ranking as shown in Algorithm 6.4.

Algorithm 6.4: ϵ-Ranking

```
for j ∈ S do
    Sample a uniformly random x_j ∈ [0, 1].
for each buyer i who arrives do
    Match i to an unmatched j ∈ N(i) maximizing w_j (1 − e^{x_j − 1 − ϵ}).
```

In the following, fix some instance $G = (S, B, E)$ with vertex-weights w and some $\epsilon > 0$. Then we let $f : [0, 1]^S \to \mathbb{R}$ represent the total weight of the matching generated by Algorithm 6.4 with fixed samples x_j. We will show that ϵ-Ranking is still $(1 - \frac{1}{e} - \epsilon)$-competitive while also allowing us to give a concentration bound.

To give a concise proof of the $(1 - \frac{1}{e} - \epsilon)$-competitiveness, we will folllow a similar approach as we did to prove the competitiveness of Ranking in Chapter 5. We will use the primal-dual analysis due to Devanur et al. [35] with the economic viewpoint by Eden et al. [44].

As usual we associate random variables r_j with all $j \in S$ and u_i with $i \in B$. The value $w_j e^{x_j - 1 - \epsilon}$ represents the *price* of j and whenever a match between i and j is made, this is a *sale*. We will then set r_j (the *revenue*) to be $w_j e^{x_j - 1 - \epsilon}$ and u_i (the *utility*) to be $w_j e^{x_j - 1 - \epsilon}$. If a vertex is never matched, its revenue or utility respectively will be zero.

Lemma 6.6. *Using the notation from Lemma 6.3, we have that for all $j \in S$ and fixed samples x,*

$$w(M_{-j}) - \frac{2}{\epsilon} w_j \leq w(M) \leq w(M_{-j}) + w_j.$$

Additionally, for any $i \in B$, its utility u_i in the run on G will be no less than in the run on G_{-j}.

Proof. For any buyers $i, i' \in B$, let $N^{(i)}(i')$ be the set of neighbors of i' in G which are unmatched by the time that i arrives in the run of Algorithm 6.4 with the fixed values of x. Likewise, let $N_{-j}^{(i)}(i')$ be the set of unmatched neighbors of i' in the run of ϵ-Ranking on G_{-j} when i arrives. We claim that for all $i \in B$ there exists some $j' \in S$ such that

$$w_{j'}(1 - e^{x_{j'}-1-\epsilon}) \leq w_j(1 - e^{x_j-1-\epsilon})$$

and for all $i' \in B$, we have $N^{(i)}(i') = N_{-j}^{(i)}(i')$ or $N^{(i)}(i') = N_{-j}^{(i)}(i') \cup \{j'\}$.

This claim is almost the same as in the proof of Lemma 6.3 and may likewise be shown via induction. Note that the extra condition on $w_{j'}$ holds at the beginning where $j' = j$ and whenever i matches to j', it frees up a vertex j'' with

$$w_{j''}(1 - e^{x_{j''}-1-\epsilon}) \leq w_{j'}(1 - e^{x_{j'}-1-\epsilon})$$

due to the fact that j' was picked over j'' in line 4. If i was not even matched in G_{-j}, we can simply set $j' = j$ for the induction.

Now note that since $N_{-j}^{(i)}(i) \subseteq N^{(i)}(i)$ for all $i \in B$, we always maximize over a larger set in line 4. Thus the utility of i will be no smaller in the run on G compared to the run on G_{-j}.

On the other hand, let $T \subseteq S$ be the set of goods matched in the run on G and let $T_{-j} \subseteq S \setminus \{j\}$ be the set of goods matched in the run on G_{-j}. Then we observe that $T \setminus T_{-j} \subseteq \{j\}$ because for all $j' \neq j$, if j' gets matched to i in M, then either $j' \in N_{-j}^{(i)}(i)$, implying that i will match to j' in M_{-j} , or j' was already matched to some other vertex. In both cases, if $j' \in T$ then $j' \in T_{-j}$. This implies that $w(M) \leq w(M_{-j}) + w_j$.

We also have that $|T_{-j} \setminus T| \leq 1$. Simply imagine a buyer $i^\star$ that arrives after all other buyers and has edges to all goods. Then by the claim, there exists some $j' \in S$ such that

$$(S \setminus \{j\}) \setminus T_{-j} = N_{-j}^{(i^\star)}(i^\star) \subseteq N^{(i^\star)}(i^\star) \cup \{j'\} = S \setminus (T \cup \{j'\})$$

and so $T_{-j} \subseteq T \cup \{j'\}$. This implies that $w(M) \geq w(M_{-j}) - w_{j'}$.

Finally, we also know by the claim that $w_{j'}(1 - e^{x_{j'}-1-\epsilon}) \leq w_j(1 - e^{x_j-1-\epsilon})$ which implies

$$w_{j'} \leq \frac{1}{1 - e^{-\epsilon}} w_j \leq \frac{1}{\left(1 - \frac{1}{e}\right)\epsilon} w_j \leq \frac{2}{\epsilon} w_j.$$

Thus we have shown $w(M_{-j}) - \frac{2}{\epsilon} w_j \leq w(M) \leq w(M_{-j}) + w_j$ as required. □

Lemma 6.7 (Bounded Differences). *Let $x \in [0,1]^S$, $j^\star \in S$ and $\theta \in [0,1]$ be arbitrary. Define x'_j to be θ if $j = j^\star$ and x_j otherwise. Then $|f(x) - f(x')| \leq \left(1 + \frac{2}{\epsilon}\right) w_{j^\star}$.*

Proof. As in the proof of Lemma 6.2, we can simply remove $j^\star$ and apply Lemma 6.6. Then

$$w(M_{-j^\star}) - \frac{2}{\epsilon} w_{j^\star} \leq f(x) \leq w(M_{-j^\star}) + w_{j^\star},$$
$$w(M_{-j^\star}) - \frac{2}{\epsilon} w_{j^\star} \leq f(x') \leq w(M_{-j^\star}) + w_{j^\star}$$

which implies the result. □

Lemma 6.8. *For any $\{i, j\} \in E$, we have $\mathbb{E}[r_j + u_i] \geq (1 - \frac{1}{e} - \epsilon) w_j$.*

Proof. Fix all ranks x except for x_j. Then we can define u^* to be the utility of i when ϵ-Ranking is ran on G_{-j}. By Lemma 6.6, we know that $u_i \geq u^*$, regardless of the value of x_j.

On the other hand, if x_j is small enough that $w_j(1 - e^{x_j-1-\epsilon}) > u^*$, then j will definitely get matched because if j is not yet matched by the time that i arrives, then clearly j will be chosen in line 4 of the algorithm and so it gets matched to i. Now if u^* is very small, this may be the case for all values of x_j and in that case

$$\mathbb{E}[r_j \mid x_{-j}] \geq \int_0^1 w_j e^{t-1-\epsilon}\,\mathrm{d}t = \left(1 - \frac{1}{e}\right) e^{-\epsilon} w_j \geq \left(1 - \frac{1}{e} - \epsilon\right) w_j.$$

Otherwise there will be some value $z \in [0,1]$ such that $w_j(1 - e^{z-1-\epsilon}) = u^*$ and then we can compute

$$\mathbb{E}[r_j \mid x_{-j}] \geq \int_0^z w_j e^{t-1-\epsilon}\,\mathrm{d}t = \left(1 - \frac{1}{e}\right) w_j - u^*.$$

But clearly, in both cases we have

$$\mathbb{E}[r_j + u_i \mid x_{-j}] \geq \mathbb{E}[r_j \mid x_{-j}] + u^* \geq \left(1 - \frac{1}{e} - \epsilon\right) w_j$$

and so in particular $\mathbb{E}[r_j + u_i] \geq (1 - \frac{1}{e} - \epsilon) w_j$ as claimed. □

Lemma 6.9. *ϵ-Ranking is $(1 - \frac{1}{e} - \epsilon)$-competitive.*

Proof. Let M^* be a maximum weight matching and let M be the matching output by ϵ-Ranking. Notice that every time we match an edge in the algorithm, we increase $\sum_{j \in S} r_j + \sum_{i \in B} u_i$ by exactly the weight of the edge. Thus by Lemma 6.8,

$$\begin{aligned} \mathbb{E}[w(M)] = \mathbb{E}\left[\sum_{j \in S} r_j + \sum_{i \in B} u_i\right] &\geq \sum_{\{i,j\} \in M^*} \mathbb{E}[r_j + u_i] \\ &\geq \sum_{\{i,j\} \in M^*} \left(1 - \frac{1}{e} - \epsilon\right) w_j = \left(1 - \frac{1}{e} - \epsilon\right) w(M^*) \end{aligned}$$

and therefore ϵ-Ranking is $(1 - \frac{1}{e} - \epsilon)$-competitive. □

Finally, we have the tools necessary to show Theorem 6.3 by combining Lemma 6.7 with Lemma 6.9.

Proof of Theorem 6.3. Given some $\alpha > 0$, we consider the algorithm $\frac{\alpha}{2}$-Ranking which we know to be $(1 - \frac{1}{e} - \frac{\alpha}{2})$-competitive by Lemma 6.9. We apply Lemma 6.1 (McDiarmid's inequality) with Lemma 6.7 (bounded differences). This gives us

$$\mathbb{P}\left[w(M) < \left(1 - \frac{1}{e} - \alpha\right) w(M^*)\right] < e^{-2\frac{\alpha^2}{2}\frac{w(M^*)^2}{(1+4/\alpha)^2||w||_2^2}}$$
$$\leq e^{-\frac{\alpha^4}{50}\frac{w(M^*)^2}{||w||_2^2}}$$

where we use that $\alpha < 1$ since otherwise the bound holds trivially. □

The results of this section may also be extended to a generalization of the Online Vertex-Weighted Bipartite Matching Problem which is called the Online Single-Valued Bipartite Matching Problem. The setup is almost identical in that we still have a bipartite graph $G = (S, B, E)$ with vertex weights $w : S \to \mathbb{R}_+$ on the offline vertices. However, now each offline vertex j also has a capacity $c_j \in \mathbb{N}$ that represents how often it is allowed to be matched.

Clearly, Theorem 6.3 can be extended to this setting by simply creating c_j many copies of each offline vertex j. This can be done implicitly and in a capacity-oblivious way by sampling a new x_j every time j is matched during the Ranking (or ϵ-Ranking) algorithm.

Recently, Vazirani [121] showed that this "resampling" is in fact not necessary, i.e. that the same value of x_j can be used for every copy of j while still achieving $(1-\frac{1}{e})$-competitiveness of Ranking; see Algorithm 6.5.

Algorithm 6.5: SINGLE-VALUED RANKING

for $j \in S$ **do**
    Sample a uniformly random $x_j \in [0, 1]$.
for *each buyer i who arrives* **do**
    Match i to a $j \in N(i)$ which has been matched less than c_j times, maximizing $w_j \left(1 - e^{x_j - 1}\right)$.

The main benefit of Algorithm 6.5 is that it uses fewer random bits than running RANKING on the reduced instance with c_j many copies of each offline vertex j. However, it will accordingly be less tightly concentrated which leads to a version of Theorem 6.3 in which the bound depends not on $||w||_2^2$ but rather on $\sum_j (c_j w_j)^2$.

6.5 Discussion

We have shown that RANKING and its many variants achieve their competitive ratios with high probability rather than just in expectation. This leaves several interesting open problems. The first is to show a concentration bound for the original weighted version of RANKING rather than ϵ-RANKING. As mentioned, the bounded differences approach fails due to the large influence that vertices with small weight can have on the matching. However, this should happen rarely and so a more fine-grained analysis may be able to overcome this challenge.

A second interesting prospect is to consider the AdWords problem. Vazirani [121] showed that a variant of RANKING can be used for AdWords with small bids under the assumption of the so-called *no-surpassing property* which tends to hold in practice though the bound is once again given in terms of expectation. An advantage of this approach over the classic MSVV algorithm [99] is that RANKING does not need to know about the budgets. It may be possible to show a concentration bound for this algorithm as well. However, this is made more challenging by the fact that the setting is edge-weighted.

Chapter 7

Hypergraph Matching

7.1 Introduction

In the previous chapters we have discussed several variants of the Online Bipartite Matching Problem, such as vertex weights or the fully online setting. We also discussed the Online Bipartite b-Matching Problem in which offline vertices can be matched multiple times. In this chapter, we will look at a different generalization in which vertices are matched to entire sets of vertices at once. This chapter is based on the paper "Almost Tight Bounds for Online Hypergraph Matching" which was joint work with Rajan Udwani [117].

A hypergraph $G = (V, E)$ consists of a set V of vertices and a set E of *hyperedges*.[1] Each hyperedge $e \in E$ is a non-empty subset of V. The maximum size of any edge is called the *rank* of the hypergraph. Note that a hypergraph of rank 2 is just an undirected graph in the standard sense. A matching in a hypergraph is simply a collection of vertex-disjoint hyperedges.

[1]Technically, this is an undirected hypergraph. We will only consider undirected hypergraphs in this chapter.

In Section 7.2, we will introduce several slightly different models of the Online Hypergraph Matching Problem (OHMP), in which hyperedges (or vertices) arrive online in adverserial order and must be matched immediately and irrevocably. Applications for hypergraph matching arise in a variety of settings such as revenue management in airlines [114, 115], combinatorial auctions [83, 86], and ridesharing [109].

In the rest of the chapter, our goal will be to give lower and upper bounds on the competitive ratios for the Online Hypergraph Matching Problem, particularly in the case of large rank. A notable difference between OBMP and the OHMP is the issue of efficient computability. In most online matching problems, it is possible to compute the offline optimal solution (i.e. the comparison point for the competitive ratio) in polynomial time using standard, combinatorial algorithms. However, it is NP hard to find a $\frac{\Omega(\log k)}{k}$-approximate maximum hypergraph matching [68].

We will consider both integral and fractional variants of the OHMP. In the integral case, any greedy algorithm is $\frac{1}{k}$-competitive and this is optimal among deterministic algorithms, similar to the results that we have seen in the bipartite case in Chapter 5. Due to the hardness result mentioned above, no *polynomial time* online algorithm can have a competitive ratio better than $\frac{\Omega(\log k)}{k}$ unless $\mathrm{P} = \mathrm{PN}$. This raises the question whether there is any, potentially randozimed and exponential time, online algorithm which is substantially better than $\frac{1}{k}$-competitive. In Section 7.3, we partially resolve this question in the negative by showing that there is no online algorithm which is $\frac{2+\epsilon}{k}$-competitive for any $\epsilon > 0$ and large k, thus leaving only a constant factor gap.

In the fractional case, there is no computational hardness as fractional hypergraph matchings (which we will define in Section 7.2) can be computed via a linear program. Moreover, this is a special case of the Online Packing Problem for which Buchbinder et al. [21] give a $\Omega\left(\frac{1}{\log k}\right)$-competitive algorithm. In the special case of the fractional OHMP, we give a simpler algorithm which is $\frac{1-o(1)}{\ln k}$-competitive. Moreover, we will show a matching upper bound showing that no online algorithm is better than $\frac{1}{\ln k}$-competitive.

7.2 Models

There are two reasonable variants of online hypergraph matching: vertex arrival and edge arrival. We will start by discussing the vertex arrival model since this is more closely inspired by the OBMP though in the rest of this chapter we will focus on edge arrival.

In the vertex arrival model, we are given a hypergraph $G = (V, E)$ where $V = B \cup S$ can be divided into buyers (online vertices) and sellers (offline vertices). Every hyperedge contains exactly one buyer and at least one seller. Buyers arrive one by one in adversarial order and reveal the hyperedges incident to them. At that point we must immediately and irrevocably match at most one of these hyperedges with the goal of maximizing the total number of matched hyperedges.

Clearly, this model generalizes the OBMP if the rank of the hypergraph is 2. One may add the additional restriction that the hypergraph be k-partite, i.e. that $V = B \cup S_1 \cup \cdots \cup S_{k-1}$ such that every edge contains exactly one element from B and each S_l. However, it is not immediately clear if this helps in any way. The issue is that finding a maximum hypergraph matching is NP-hard even in a 3-partite hypergraph; this problem is usually called the 3D Matching Problem.

As already mentioned, we will largely focus on edge arrival. In the edge arrival model, all vertices are offline and instead edges arrive one by one at which point we only need to decide whether we want to match them or not. Note that vertex arrival can simply be reduced to edge arrival by letting all the incident edges to an online vertex arrive in some arbitrary order.

However, in the case of large rank k, edge arrival and vertex arrival are essentially the same. This is because we can turn an edge arrival instance of rank k into a vertex arrival instance of rank $k+1$: simply add one unique vertex for each arriving edge. Since the focus of this chapter is on large rank, we will restrict ourselves to the edge arrival model as it is conceptually simpler.

We will also consider the fractional variant of the OHMP. In that case, the optimal offline solution is a solution to the following linear program.

$$\max \quad \sum_{e \in E} x_e \tag{7.1a}$$

$$\text{s.t.} \quad \sum_{\substack{e \in E \\ v \in e}} x_e \leq 1 \quad \forall v \in V, \tag{7.1b}$$

$$x_e \geq 0 \quad \forall e \in E \tag{7.1c}$$

Each decision variable $x_e \in [0,1]$ in the LP captures the fraction of edge e included in the matching. The matching constraint (7.1b) enforces the total fraction of edges incident on a vertex $v \in V$ to be at most 1. An online algorithm can match an arbitrary fraction of each arriving edge, subject to the same constraint as the LP. We remark that unlike in the bipartite matching case, a feasible solution of the LP is not in general a convex combination of integral hypergraph matchings.

Finally, note that by adding dummy vertices, we can assume wlog. that every edge has *exactly* k offline vertices, making the instance k-uniform.

7.2.1 Related Work

Perhaps closest to our setting is the work of Buchbinder and Naor [21] on online packing. They considered an online packing problem that generalizes the fractional online hypergraph matching problem studied here and gave a $O(\log k)$ competitive algorithm. A special case of the online packing problem was considered earlier in [9], in the context of online routing.

Another closely related line of work is on the problem of network revenue management [114, 115]. This is a stochastic arrival setting where seats in flights are offline resources allocated to sequentially arriving customers. A customer with multi-stop itinerary requires a seat on each flight in the itinerary. Recently, Ma et al. [94] gave a $\frac{1}{k+1}$ algorithm for network revenue management.

Another stream of work has focused on hypergraph matching from the perspective of ridesharing. Pavone et al. [109] introduced a hypergraph matching problem with deadlines to capture applications in ridesharing. Their model and results are incomparable to ours. Lowalekar et al. [93] consider a model inspired by ridesharing but with a stochastic arrival sequence. Finally, several papers [83, 86] consider related settings in combinatorial auctions that correspond to online hypergraph matching with stochastic arrivals.

As mentioned previously, for the standard (offline) maximum hypergraph matching problem, Hazan et al. [68] showed that unless $\mathrm{P} = \mathrm{NP}$, no polynomial time algorithm can find a $\frac{\Omega(log k)}{k}$-approximate maximum hypergraph matching. The best-known approximation algorithm is due to Cygan [32] who provides a factor $\frac{3}{k}$ approximation.

Lastly, we want to mention the recent work of Borst et al. [19] which appeared after the results in this chapter were announced but before they were published. They consider the integral, vertex-arrival variant of the online hypergraph matching and give an optimal algorithm in the case of $k = 3$ which turns out to have a competitive ratio of $(e-1)/(e+1) \approx 0.4621$. In addition, they give an algorithm which beats the greedy algorithm for larger k under the assumption that the degree of online vertices is bounded.

7.3 Integral Matchings

In the following, fix some $k \geq 2$. We will focus on the OHMP in the edge arrival model. We start with a simple observation which is essentially folklore.

Theorem 7.1. *There is a $\frac{1}{k}$-competitive algorithm for the OHMP. This is the best possible competitive ratio for deterministic algorithms.*

Proof. Consider the greedy online algorithm that includes an arriving edge e in the matching if it can be added to the matching, i.e. if it is disjoint with all previously included edges. Let $e_1, \ldots, e_\ell$ be the set of edges included in an offline optimum solution and let $V_o \subseteq V$ denote the set of offline vertices that are covered by the edges chosen in the online algorithm. For each edge e_i, at least one of the vertices that it intersects must be included in V_o. Thus, $|V_o| \geq \frac{1}{k}\ell\, k$ and the online algorithm picks at least $\frac{1}{k}\ell$ hyperedges.

To see that this is the best possible competitive ratio, consider the following instance. First, an edge e with k vertices arrives. Now, for every $v \in e$, another edge arrives which contains v and $k-1$ "fresh" vertices. Clearly $\mathrm{OPT} = k$ but this requires one to *not* match e. Now consider some deterministic algorithm $\mathcal{A}$. If $\mathcal{A}$ matches e, then its competitive ratio is no better than $\frac{1}{k}$. But if $\mathcal{A}$ does not match e, then consider another instance in which no edge arrives after e: $\mathcal{A}$ would have a competitive ratio of 0 in that case. □

In our first result, we show that even a randomized and possibly exponential time online algorithm cannot achieve a much better competitive ratio for this problem. To show this, we will use Yao's principle, as restated below (recall Section 5.5).

Lemma 7.1 (Yao's Principle)**.** *Let α be the best competitive ratio of any randomized algorithm. Let β be the competitive ratio of the best deterministic algorithm against some fixed distribution of instances. Then $\alpha \leq \beta$.*

Before we get to our main result, we will first give a slightly weaker result that serves both as a warm up and as a gadget for the main result.

Theorem 7.2. *For even k, there does not exist a $\frac{4+\epsilon}{k}$ competitive algorithm for the k-uniform online hypergraph matching problem for any $\epsilon > 0$.*

Proof. Using Yao's principle, we will construct a distribution of instances with even k where $\mathrm{OPT} = \frac{k}{2}$ but the best deterministic online algorithm can only achieve an expected matching size of 2.

For any given even value k, the overall (random) instance G_k will consist of $\frac{k}{2}$ "red" edges and $\frac{k}{2}$ "blue" edges constructed in $\frac{k}{2}$ phases. In each phase, there will be one red and one blue edge which look indistinguishable to any online algorithm. The idea is that if the algorithm ever picks a blue edge, it will be locked out of future edges, thus limiting the expected matching size. See Figure 7.1 for an example of the construction. The construction of the red and blue edge in each phase proceeds as follows:

1. Let A be a set of vertices which intersects every previous blue edge exactly once. Create a new edge e_1 which consists of A and $k - |A|$ many new vertices that have not been in any edges yet. Let e_1 arrive in the instance.

2. Now let A' be a second set of vertices which intersects every previous blue edge *and* e_1 exactly once. Create a new edge e_2 which consists of A' and $k - |A'|$ many new vertices. Let e_2 arrive in the instance.

3. Randomly let one of $\{e_1, e_2\}$ be red and the other blue with equal probability.

Note that the sets A and A' can always be found because each edge contains k vertices and we have $k/2$ phases. The crucial property of this construction is that each blue edge intersects all future edges whereas each red edge is disjoint from all future edges. In particular, the $\frac{k}{2}$ red edges form a maximum size matching, i.e. $\mathrm{OPT} = \frac{k}{2}$.

Now consider some deterministic online algorithm $\mathcal{A}$. Let α_i be the probability that $\mathcal{A}$ matches the red edge in phase i and let β_i be the probability that $\mathcal{A}$ matches the blue edge in phase i. Clearly, since the red and blue edges are determined independently and uniformly at random, we must have $\alpha_i = \beta_i$. Moreover, since at most one blue edge can be picked, we know $\alpha_1 + \cdots + \alpha_{k/2} \leq 1$. Thus the expected size of the matching generated by $\mathcal{A}$ is at most

$$\alpha_1 + \cdots + \alpha_{k/2} + \beta_1 + \cdots + \beta_{k/2} \leq 2. \qquad \square$$

Theorem 7.3. *If k is a power of two, then here does not exist a $\frac{2+\epsilon}{k}$ competitive algorithm for the online hypergraph matching problem for any $\epsilon > 0$.*

Proof. We will use induction to create a distribution over graphs H_k for powers of two k, with the following properties:

1. There are k red and k blue edges.

2. The edges appear in k phases, each of which consists of one red and one blue edge where the color is chosen uniformly and independently at random.

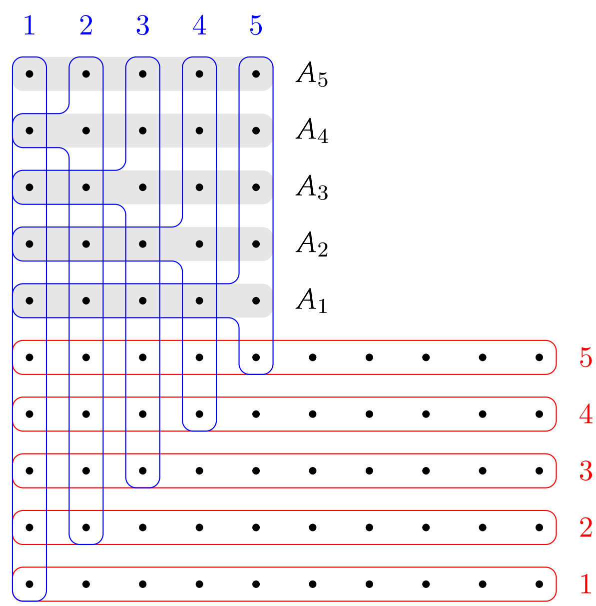

Figure 7.1: Shown is gadget G_{10} proving that a competitive ratio of $\frac{4}{k} + \epsilon$ is imposisble for $k = 10$. The numbers indicate in which phase each edge was added. The lightly shaded areas represent the vertex sets $A_1, \ldots, A_5$ which are useful for the construction of H_k.

3. Every blue edge intersects all future edges.

4. Every red edge is disjoint from all future edges.

H_1 is trivial to construct. We will just have a single vertex which is simultaneously in both a red and blue singleton edge. Suppose that we can construct $H_{k/2}$.

Now in order to construct H_k, we first employ the $\frac{k}{2}$ phases of G_k. After this we can construct $\frac{k}{2}$ disjoint sets $A_1, \ldots, A_{k/2}$ of $\frac{k}{2}$ vertices such that each A_i intersects all blue edges and none of the red edges. See again Figure 7.1. Now, for the remaining $k/2$ phases, we recursively employ the distribution $H_{k/2}$ as follows. In phase $i + \frac{k}{2}$, we extend the the two edges from phase i of $H_{k/2}$ by the set A_i to form sets of size k. This gives us the two edges of rank k for phase $i + \frac{k}{2}$.

Finally, one may check that properties 1 through 4 are satisfied by induction. Thus, we may conclude the proof similar to the proof of Theorem 7.2: the optimum solution picks all k red edges whereas any deterministic online algorithm can only get an expected value of 2 since at most one blue edge can be picked and red and blue edges in each phase are indistinguishable for the online algorithm. □

7.4 Fractional Matchings

Inspired by the Water Filling (or Balance) algorithm of [81, 99], designed for the OBMP and its variants, we propose the following online algorithm for the *fractional* OHMP in the edge arrival model.

Algorithm 7.1: Hypergraph Water Filling

For each $e \in E$, set $y_e := 0$.
for *each edge e which arrives* **do**
    Increase y_e continuously as long as $\sum_{v \in e} (k \ln k)^{x_v - 1} \leq 1$ where $x_v := \sum_{f \in E : v \in f} y_f$ is the *fill level* of v.

The final value of variable y_e in the algorithm is the fraction of edge e that is included in the matching. At any moment, variable x_v in the algorithm captures the total fraction of edges incident on vertex v that have been included in the matching. In other words, the value of x_v is the fraction of v that has already been matched by the algorithm. In order to preserve resources for future edges, Algorithm 7.1 stops matching an edge when the value of $\sum_{v \in e} (k \ln k)^{x_v - 1}$ reaches 1. For illustration, consider the following scenarios at the arrival of edge e incident on resources $\{1, \ldots, k\}$.

1. $x_i = 0$ for all $i \in [k]$ on arrival of e. Then, the algorithm will match a $\frac{\ln(\ln k)}{\ln(k) + \ln(\ln k)}$ fraction of the edge before stopping.

2. $x_i = \frac{1}{2}$ for all $i \leq \sqrt{k \ln k}$ and $x_i = 0$ otherwise. Then, for sufficiently large values of k, the algorithm does not match any fraction of edge e.

We would like to note that Algorithm 7.1 constructs a fractional matching by augmenting the primal solution y, whereas the online packing algorithm of [21] augments the dual solution. We show that Algorithm 7.1 achieves the best possible competitive ratio guarantee for large k.

Theorem 7.4. *Algorithm 7.1 is* $\frac{1-o(1)}{\ln(k)}$*-competitive for the fractional OHMP with edge arrivals.*

Proof. Given a hypergraph $G = (V, E)$, let ALG denote the total size of the fractional hypergraph matching obtained by Hypergraph Water Filling and let OPT denote the value of the optimal fractional offline solution. We use a primal-dual approach inspired by [21, 35] to prove the result. Note that the dual of (7.1) is given by:

$$\begin{aligned} \min \quad & \sum_{v \in V} r_v \\ \text{s.t.} \quad & \sum_{v \in e} r_v \geq 1 \quad \forall e \in E, \\ & r_v \geq 0 \quad v \in V. \end{aligned}$$

It suffices to find non-negative $(r_v)_{v \in V}$ which satisfy

$$\sum_{v \in V} r_v \leq \mathrm{ALG}, \tag{7.2}$$

$$\sum_{v \in e} r_v \geq \frac{1 - \frac{1}{\ln k}}{\ln k + \ln(\ln k)} = \frac{1 - o(1)}{\ln k} \quad \forall e \in E. \tag{7.3}$$

From there, an approximate complementary slackness argument can be used to get the $\frac{1-o(1)}{\ln k}$ bound on the competitive ratio; see the proof of Theorem 5.2 from the background chapter.

We set the dual variables using the following procedure. In the beginning, all variables r_i are set to 0. When Algorithm 7.1 is matching edge e in line 3 by some infinitesimal amount $\mathrm{d}t$, we increase r_v by $(k \ln k)^{x_v - 1}\mathrm{d}t$ for all $v \in e$. Note that by the condition in line 3, we know that (7.2) holds at the end of the algorithm. It remains to show that (7.3) is also satisfied. Fix an arbitrary edge $e \in E$ and consider the following two cases.

Case 1: Let $e \in E$ be arbitrary and let x_v be the final fill levels of the vertices $v \in e$. If $x_v = 1$ for any v, we know that

$$r_v = \int_0^1 (k \ln k)^{t-1}\,\mathrm{d}t = \frac{1 - 1/(k \ln k)}{\ln k + \ln(\ln k)} \geq \frac{1 - 1/\ln k}{\ln(k) + \ln(\ln k)},$$

so the this one vertex is already enough for (7.3).

Case 2: Otherwise, since all $x_v < 1$ at the end of the algorithm, we must have that $P := \sum_{v \in e} (k \ln k)^{x_v - 1} \geq 1$. But in that case, we can compute:

$$\begin{aligned} \sum_{v \in e} r_v &\geq \sum_{v \in e} \int_0^{x_v} (k \ln k)^{t-1}\,\mathrm{d}t \\ &= \frac{P - 1/\ln k}{\ln k + \ln(\ln k)} \end{aligned}$$

which shows the claim and thus the theorem. □

Next, we show that this bound is asymptotically tight, i.e. that no online algorithm can beat the performance of Algorithm 1 for large k.

Theorem 7.5. *For any $\epsilon > 0$ and k large enough, there does not exist any online algorithm which is $\frac{1+\epsilon}{\ln(k)}$-competitive for the fractional OHMP with edge arrivals.*

Proof. Let $\mathcal{A}$ be some algorithm for the fractional online hypergraph matching problem. We can assume wlog. that $\mathcal{A}$ is deterministic. This is because if $\mathcal{A}$ is randomized, we may

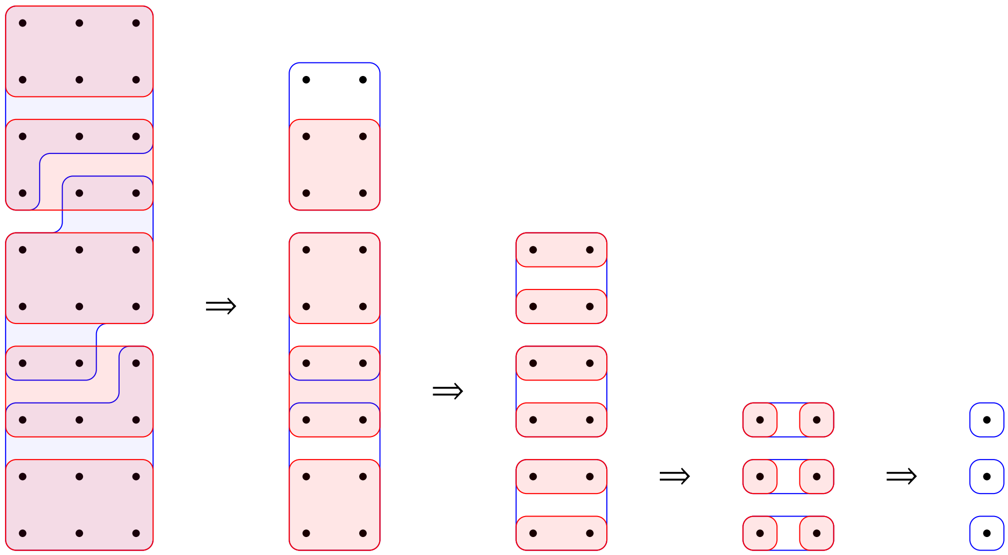

Figure 7.2: Shown is the upper-bounding construction with $k = 10$, $l = 3$, $\delta = 0.5$. In each step we replace the blue edges with as many red edges of $\frac{1}{1+\delta}$ times the size as possible. Then we pick the l red edges that the algorithm puts the most weight on, make those the new blue edges and repeat until only singleton edges are left.

create another algorithm $\mathcal{A}'$ which simply fractionally allocates every edge e with the expected value of $\mathcal{A}$. Then $\mathcal{A}'$ is a deterministic algorithm that performs just as well as $\mathcal{A}$.

Fix some large $l \in \mathbb{N}$ and small $\delta > 0$. We will now construct an instance in which every hyperedge has size at most k. Some will have size strictly less than k but if a k-uniform instance is desired, we can simply fill up with dummy vertices. The instance is created according to the following procedure (see Figure 7.2).

1. Set $m \leftarrow k$ and let l disjoint edges of size k arrive. Let U be the set of all vertices in these l edges.
2. If $m = 0$, stop. Otherwise, set $m \leftarrow \lfloor \frac{m}{1+\delta} \rfloor$.
3. Partition U into as many disjoint edges of size m as possible and let these arrive.
4. Let $e_1, \ldots, e_l$ be the l of these edges that $\mathcal{A}$ matches the most.
5. Update $U \leftarrow e_1 \cup \cdots \cup e_l$ and go back to step 2.

Now let α be the competitive ratio of $\mathcal{A}$. Our first observation is that steps 2–5 execute $(1-o(1))\log_{1+\delta}(k)$ many times as $k \to \infty$. Moreover, in each iteration, we cover l blue edges with

$$\left\lfloor \frac{lm}{\left\lfloor \frac{m}{1+\delta} \right\rfloor} \right\rfloor \geq (1+\delta)l - 1$$

red edges. The optimal solution can thus be increased by at least $\delta l - 1$ by shifting weight from the blue edges to the red edges. Overall, this yields $\mathrm{OPT} \geq (1-o(1))\log_{1+\delta}(k)(\delta l - 1)$ and therefore $\mathrm{ALG} \geq \alpha(1-o(1))\log_{1+\delta}(k)(\delta l - 1)$.

Let $E^\star \subseteq E$ be the set of edges which are picked at various points in step 4 and let y be the fractional matching constructed by $\mathcal{A}$. Then because these edges are always the l most covered edges we know that

$$y(E^\star) \geq \min_{m \geq 1} \frac{l}{\left\lfloor \frac{lm}{\left\lfloor \frac{m}{1+\delta} \right\rfloor} \right\rfloor} \mathrm{ALG} \geq \frac{1}{1+\delta-\frac{1}{l}} \mathrm{ALG}.$$

Lastly, we know that that all edges in $E^\star$ overlap in l vertices, namely the l vertices that are contained in the final iteration of the loop. This implies that $y(E^\star) \leq l$. Combining these inequalities, we thus get

$$\begin{aligned} \alpha &\leq \frac{\left(1+\delta-\frac{1}{l}\right) l}{(1-o(1))\log_{1+\delta}(k)(\delta l - 1)} \\ &= \frac{((1+\delta)l - 1)\ln(1+\delta)}{(1-o(1))(\delta l - 1)} \cdot \frac{1}{\ln(k)} \end{aligned}$$

Finally, observe that for small δ, large l and large k, we get $\alpha < \frac{1+\epsilon}{\ln(k)}$ as claimed. □

7.5 Edge Weights

We will now consider a variant of the fractional OHMP which allows for non-negative edge weights $(w_e)_{e \in E}$. It is quite easy to show that no online algorithm can achieve a bounded competitive ratio in the presence of edge weights, not even for the OBMP. Whenever we match any edge we risk that there may be a later edge including the same vertex of arbitrarily larger value. For this reason, Feldman et al. [49] introduced the *free disposal condition* in the context of online matching.

The idea behind free disposal is that in many applications of online matching (such as online advertising), it makes sense to allow us to drop edges from the matching with no penalty. Of course we do not count such dropped edges towards our objective function and we do not allow us to go back and re-match them later.

We will now give an algorithm for the fractional OHMP with edge arrivals and edge weights which is essentially a generalization of Algorithm 7.1 using ideas inspired by the work of Devanur et al. [34] on free disposal.

Algorithm 7.2: Hypergraph Weighted Water Filling

For each $e \in E$, let $y_e := 0$.
For each $v \in V$, let $f_v(t) := \sum_{e: v \in e, w_e \geq t} y_e$ for all $t \geq 0$.
for *each edge e which arrives* **do**
    while $\sum_{v \in e} \int_0^{w_e} (k \ln(k))^{f_v(t)-1} \, \mathrm{d}t \leq w_e$ **do**
        for $v \in e$ *with* $x_v = 1$ **do**
            Let e_v^- be a minimum weight edge with $v \in e_v^-$ and $y_{e_i^-} > 0$.
            $y_{e_v^-} \leftarrow y_{e_v^-} - \mathrm{d}s$
        $y_e \leftarrow y_e + \mathrm{d}s$

Theorem 7.6. *For any $\epsilon > 0$ and k large enough, Algorithm 7.2 is $\frac{1-\epsilon}{\ln k}$-competitive for online fractional weighted hypergraph matching problem with free disposal.*

Proof. The proof will use a similar primal-dual approach as Theorem 7.4 with non-negative dual variables r_v for all $v \in V$. The dual in question is given by the following linear program.

$$
\begin{aligned}
\min \quad & \sum_{v \in V} r_v \\
\text{s.t.} \quad & \sum_{v \in e} r_v \geq w_e \quad \forall e \in E, \\
& \quad r_v \geq 0 \qquad v \in V.
\end{aligned}
$$

Once again, when Algorithm 7.2 is matching an edge e by some amount $\mathrm{d}s$, we increase the dual by at most an according amount. Note that the total increase in the matching is $w_e - \sum_{v \in e} w_{e_v^-}$ where $w_{e_v^-} := 0$ if $x_v < 1$.

We increase each r_v by $\int_{w_{e_v^-}}^{w_e} (k \ln k)^{f_v(t)-1} \, \mathrm{d}t \mathrm{d}s$. By definition of f_v, we know that $f_v(t) = 1$ for all $t < w_{e_v^-}$ and thus

$$
\begin{aligned}
& \sum_{v \in e} \int_{w_{e_v^-}}^{w_e} (k \ln k)^{f_v(t)-1} \, \mathrm{d}t \\
&= \sum_{v \in e} \int_0^{w_e} (k \ln k)^{f_v(t)-1} \, \mathrm{d}t - \sum_{v \in e} \int_0^{w_{e_v^-}} (k \ln k)^{f_v(t)-1} \, \mathrm{d}t \\
&\leq w_e - \sum_{v \in e} w_{e_v^-}
\end{aligned}
$$

using the condition in line 4 of the algorithm. This implies that at the end of the algorithm, we have $\sum_{v \in V} r_v \leq \mathrm{ALG}$. It remains to show that for any $e \in E$, we get

$$
\sum_{v \in e} r_v \geq w_e \frac{1 - 1/\ln k}{\ln k + \ln(\ln k)}.
$$

Let f_v be the step function defined in Algorithm 7.2 at the end of the algorithm. Then the total r_v collected by each $v \in V$ satisfies

$$\begin{aligned} r_v &= \int_0^\infty \int_0^{f_v(t)} (k \ln k)^{s-1} \, \mathrm{d}s \, \mathrm{d}t \\ &= \int_0^\infty \frac{\sum_{v \in e} (k \ln k)^{f_v(t)-1} - 1/\ln k}{\ln k + \ln(\ln k)} \, \mathrm{d}t \end{aligned}$$

Now let $P(t) := \sum_{v \in e} (k \ln k)^{f_v(t)-1}$, then this means that

$$\begin{aligned} \sum_{v \in e} r_v &= \int_0^\infty \frac{P(t) - 1/\ln k}{\ln k + \ln(\ln k)} \, \mathrm{d}t \\ & \int_0^{w_e} \frac{P(t) - 1/\ln k}{\ln k + \ln(\ln k)} \, \mathrm{d}t \\ &= \frac{\int_0^{w_e} P(t) \, \mathrm{d}t - w_e/\ln k}{\ln k + \ln(\ln k)}. \end{aligned}$$

Finally, by the condition in line 4 of the algorithm and the fact that $f_v(t)$ only increases during the algorithm for all t, we know that $\int_0^{w_e} P(t)\mathrm{d}t \geq w_e$ which establishes the claim and thus the theorem via the same complimentary slackness argument as in the proof of Theorem 7.4. □

7.6 Discussion

In this paper we have given a tight asymptotic bound for the fractional k-uniform hypergraph matching problem and an almost tight bound for the integral variant. This leaves room for some interesting directions for future research.

A major open problem is to beat the $\frac{1}{k}$ lower bound in the integral setting. In fact, just recently Gamlath et al. [57] showed that for $k = 2$, no algorithm beats $\frac{1}{2}$, even if the underlying graph is bipartite. However, their construction in fact shows this result for the fractional setting and where we know how to beat $\frac{1}{k}$ for large k. It thus remains open whether $\frac{1}{k} + \epsilon$ is achievable for *any* k.

In fact, for small k, one may also explicitly distinguish between edge arrival and vertex arrival models as mentioned in Section 7.2 or even the fully-online arrival model of Huang et al. [72]. To the best of our knowledge, the only result here is the recent one by Borst et al. [19] who managed to get an optimal algorithm for the $k = 3$ case under vertex arrivals.

Finally, we remark that even in the fractional setting, exactly tight bounds are only known for $k = 2$ and finding a tight non-asymptotic result remains open.

Bibliography

[1] R. Abebe, R. Cole, V. Gkatzelis, and J. D. Hartline. A truthful cardinal mechanism for one-sided matching. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2096–2113. SIAM, 2020.

[2] G. Aggarwal, G. Goel, C. Karande, and A. Mehta. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1253–1264. SIAM, 2011.

[3] S. Alaei, P. Jalaly Khalilabadi, and E. Tardos. Computing equilibrium in matching markets. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, page 245–261, New York, NY, USA, 2017. Association for Computing Machinery.

[4] S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of computing*, 8(1):121–164, 2012.

[5] K. J. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica: Journal of the Econometric Society*, pages 265–290, 1954.

[6] I. Ashlagi and P. Shi. Optimal allocation without money: An engineering approach. *Management Science*, 62(4):1078–1097, 2016.

[7] S. Athanassoglou and J. Sethuraman. House allocation with fractional endowments. *International Journal of Game Theory*, 40:481–513, 2011.

[8] R. J. Aumann. Markets with a continuum of traders. *Econometrica*, 32(1/2):39–50, 1964.

[9] B. Awerbuch, Y. Azar, and S. Plotkin. Throughput-competitive on-line routing. In *Proceedings of 1993 IEEE 34th Annual Foundations of Computer Science*, pages 32–40. IEEE, 1993.

[10] H. Aziz. An impossibility result for housing markets with fractional endowments. *Operations Research Letters*, 46(5):473–475, 2018.

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

[12] M. Babaioff, T. Ezra, and U. Feige. Fair and truthful mechanisms for dichotomous valuations. In *Proc. 35th Conf. Artif. Intell. (AAAI)*, pages 5119–5126, 2021.

[13] S. Basu, R. Pollack, and M. Roy. A new algorithm to find a point in every cell defined by a family of polynomials. *Quantifier Elimination and Cylindrical Algebraic Decomposition, B. Caviness and J. Johnson eds., Springer-Verlag, to appear*, 1995.

[14] K. Binmore, A. Rubinstein, and A. Wolinsky. The nash bargaining solution in economic modelling. *The RAND Journal of Economics*, pages 176–188, 1986.

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

[16] B. Birnbaum and C. Mathieu. On-line bipartite matching made simple. *SIGACT News*, 39(1):80–87, 2008.

[17] A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic theory*, 100(2):295–328, 2001.

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

[19] S. Borst, D. Kashaev, and Z. K. Koh. Online matching on 3-uniform hypergraphs. *arXiv preprint arXiv:2402.13227*, 2024.

[20] S. Bubeck. Convex optimization: Algorithms and complexity. *Found. Trends Mach. Learn.*, 8(3-4):231–357, 2015.

[21] N. Buchbinder and J. Naor. Online primal-dual algorithms for covering and packing. *Mathematics of Operations Research*, 34(2):270–286, 2009.

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

[23] I. Caragiannis, A. Filos-Ratsikas, P. Kanellopoulos, and R. Vaish. Stable fractional matchings. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, page 21–39, New York, NY, USA, 2019. Association for Computing Machinery.

[24] I. Caragiannis, K. A. Hansen, and N. Rathi. On the complexity of pareto-optimal and envy-free lotteries. *arXiv preprint arXiv:2307.12605*, 2023.

[25] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum nash welfare. *ACM Trans. Econ. Comput.*, 7(3), sep 2019.

[26] N. Cesa-Bianchi and G. Lugoisi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

[27] E. Chastain, A. Livnat, C. H. Papadimitriou, and U. V. Vazirani. Algorithms, games, and evolution. *Proc. Natl. Acad. Sci. USA*, 111(29):10620–10623, 2014.

[28] T. Chen, X. Chen, B. Peng, and M. Yannakakis. Computational hardness of the hylland-zeckhauser scheme∗. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2253–2268. SIAM, 2022.

[29] P. F. Christiano, J. A. Kelner, A. Madry, D. A. Spielman, and S. Teng. Electrical flows, laplacian systems, and faster approximation of maximum flow in undirected graphs. In L. Fortnow and S. P. Vadhan, editors, *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 273–282. ACM, 2011.

[30] I. R. Cohen and D. Wajc. Randomized online matching in regular graphs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 960–979. SIAM, 2018.

[31] R. Cole and Y. Tao. On the existence of pareto efficient and envy-free allocations. *Journal of Economic Theory*, 193:105207, 2021.

[32] M. Cygan. Improved approximation for 3-dimensional matching via bounded path-width local search. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 509–518. IEEE, 2013.

[33] N. Devanur, J. Garg, and L. Végh. A rational convex program for linear Arrow-Debreu markets. *ACM Trans. Econom. Comput.*, 5(1):6:1–6:13, 2016.

[34] N. R. Devanur, Z. Huang, N. Korula, V. S. Mirrokni, and Q. Yan. Whole-page optimization and submodular welfare maximization with online bidders. *ACM Transactions on Economics and Computation (TEAC)*, 4(3):1–20, 2016.

[35] N. R. Devanur, K. Jain, and R. D. Kleinberg. Randomized primal-dual analysis of ranking for online bipartite matching. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete algorithms*, pages 101–107. SIAM, 2013.

[36] N. R. Devanur, C. H. Papadimitriou, A. Saberi, and V. V. Vazirani. Market equilibrium via a primal–dual algorithm for a convex program. *J. ACM*, 55(5), 2008.

[37] R. Duan and K. Mehlhorn. A combinatorial polynomial algorithm for the linear arrow–debreu market. *Information and Computation*, 243:112 – 132, 2015. 40th International Colloquium on Automata, Languages and Programming (ICALP 2013).

[38] D. P. Dubhashi and A. Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009.

[39] S. Ebadian, D. Peters, and N. Shah. How to fairly allocate easy and difficult chores. In *Proc. 21st Conf. Auton. Agents and Multi-Agent Systems (AAMAS)*, 2022.

[40] F. Echenique, N. Immorlica, and V. Vazirani, editors. *Online and Matching-Based Market Design*. Cambridge University Press, 2023.

[41] F. Echenique, A. Miralles, and J. Zhang. Fairness and efficiency for probabilistic allocations with endowments. *arXiv preprint arXiv:1908.04336*, 2019.

[42] F. Echenique, A. Miralles, and J. Zhang. Constrained pseudo-market equilibrium. *American Economic Review*, 111(11):3699–3732, November 2021.

[43] F. Echenique, J. Root, and F. Sandomirskiy. Stable matching as transportation, 2024.

[44] A. Eden, M. Feldman, A. Fiat, and K. Segal. An economics-based analysis of RANKING for online bipartite matching. In H. V. Le and V. King, editors, *4th Symposium on Simplicity in Algorithms, SOSA 2021*, pages 107–110. SIAM, 2021.

[45] J. Edmonds. Maximum matching and a polyhedron with 0,1-vertices. *Journal of Research of the National Bureau of Standards Section B Mathematics and Mathematical Physics*, page 125, 1965.

[46] E. Eisenberg. Aggregation of utility functions. *Management Science*, 7(4):337–350, 1961.

[47] E. Eisenberg and D. Gale. Consensus of subjective probabilities: the Pari-Mutuel method. *The Annals of Mathematical Statistics*, 30:165–168, 1959.

[48] M. Fahrbach, Z. Huang, R. Tao, and M. Zadimorghaddam. Edge-weighted online bipartite matching. In *Proceedings of the 61th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, 2020.

[49] J. Feldman, N. Korula, V. Mirrokni, S. Muthukrishnan, and M. Pál. Online ad assignment with free disposal. In *International workshop on internet and network economics*, pages 374–385. Springer, 2009.

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

[51] L. Fleischer, R. Garg, S. Kapoor, R. Khandekar, and A. Saberi. A fast and simple algorithm for computing market equilibria. In *International Workshop on Internet and Network Economics*, pages 19–30. Springer, 2008.

[52] D. K. Foley. Resource allocation and the public sector.(1967). 1967.

[53] D. Fudenberg and D. K. Levine. *The Theory of Learning in Games*. MIT Press Books. The MIT Press, 1998.

[54] D. Gale. *The theory of linear economic models*. University of Chicago press, 1989.

[55] D. Gale. The linear exchange model. *Journal of Mathematical Economics*, 3(2):205–209, l976.

[56] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

[57] B. Gamlath, M. Kapralov, A. Maggiori, O. Svensson, and D. Wajc. Online matching with general arrivals. pages 26–37, 11 2019.

[58] Y. Gao and C. Kroer. First-order methods for large-scale market equilibrium computation. *Advances in Neural Information Processing Systems*, 33, 2020.

[59] D. Garber and E. Hazan. A linearly convergent variant of the conditional gradient algorithm under strong convexity, with applications to online and stochastic optimization. *SIAM Journal on Optimization*, 26(3):1493–1528, 2016.

[60] J. Garg, A. Murhekar, and J. Qin. Fair and efficient allocations of chores under bivalued preferences. In *Proc. 36th Conf. Artif. Intell. (AAAI)*, 2022.

[61] J. Garg, Y. Tao, and L. A. Végh. *Approximating Equilibrium under Constrained Piecewise Linear Concave Utilities with Applications to Matching Markets*, pages 2269–2284.

[62] J. Garg, T. Tröbst, and V. V. Vazirani. One-sided matching markets with endowments: Equilibria and algorithms. *Autonomous Agents and Multi-Agent Systems*, 38(2):40, 2024.

[63] J. Garg and L. A. Végh. A strongly polynomial algorithm for linear exchange markets. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 54–65, 2019.

[64] R. Ge, F. Huang, C. Jin, and Y. Yuan. Escaping from saddle points - online stochastic gradient for tensor decomposition. In P. Grünwald, E. Hazan, and S. Kale, editors, *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, volume 40 of *JMLR Workshop and Conference Proceedings*, pages 797–842. JMLR.org, 2015.

[65] G. Goel and A. Mehta. Online budgeted matching in random input models with applications to AdWords. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 982–991, 2008.

[66] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2. Springer Science & Business Media, 2012.

[67] D. Gusfield and R. W. Irving. *The stable marriage problem: structure and algorithms*. MIT press, 1989.

[68] E. Hazan, S. Safra, and O. Schwartz. On the complexity of approximating k-set packing. *computational complexity*, 15(1):20–39, 2006.

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

[70] M. Hosseini and V. V. Vazirani. Nash-bargaining-based models for matching markets: One-sided and two-sided; fisher and arrow-debreu. In M. Braverman, editor, *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA*, volume 215 of *LIPIcs*, pages 86:1–86:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[71] Z. Huang, N. Kang, Z. G. Tang, X. Wu, Y. Zhang, and X. Zhu. How to match when all vertices arrive online. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 17–29, 2018.

[72] Z. Huang, N. Kang, Z. G. Tang, X. Wu, Y. Zhang, and X. Zhu. Fully online matching. *J. ACM*, 67(3), may 2020.

[73] Z. Huang, B. Peng, Z. G. Tang, R. Tao, X. Wu, and Y. Zhang. Tight competitive ratios of classic matching algorithms in the fully online model. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2875–2886. SIAM, 2019.

[74] Z. Huang, Q. Zhang, and Y. Zhang. Adwords in a panorama. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1416–1426, Los Alamitos, CA, USA, nov 2020. IEEE Computer Society.

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

[76] N. Immorlica, B. Lucier, G. Weyl, and J. Mollner. Approximate efficiency in matching markets. In *International Conference on Web and Internet Economics*, pages 252–265. Springer, 2017.

[77] M. Jaggi. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In S. Dasgupta and D. McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 427–435, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.

[78] M. Jaggi and S. Lacoste-Julien. On the global linear convergence of frank-wolfe optimization variants. *Advances in Neural Information Processing Systems*, 28, 2015.

[79] K. Jain. A polynomial time algorithm for computing an arrow–debreu market equilibrium for linear utilities. *SIAM Journal on Computing*, 37(1):303–318, 2007.

[80] C. Jin, P. Netrapalli, R. Ge, S. M. Kakade, and M. I. Jordan. On nonconvex optimization for machine learning: Gradients, stochasticity, and saddle points. *J. ACM*, 68(2):11:1–11:29, 2021.

[81] B. Kalyanasundaram and K. R. Pruhs. An optimal deterministic algorithm for online b-matching. *Theoretical Computer Science*, 233(1-2):319–325, 2000.

[82] R. M. Karp, U. V. Vazirani, and V. V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing*, pages 352–358, 1990.

[83] T. Kesselheim, K. Radke, A. Tönnis, and B. Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *European symposium on algorithms*, pages 589–600. Springer, 2013.

[84] A. R. Klivans and R. Meka. Learning graphical models using multiplicative weights. In C. Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 343–354. IEEE Computer Society, 2017.

[85] D. Komm, R. Královic, R. Královic, and T. Mömke. Randomized Online Algorithms with High Probability Guarantees. In E. W. Mayr and N. Portier, editors, *31st International Symposium on Theoretical Aspects of Computer Science (STACS 2014)*, volume 25 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 470–481, Dagstuhl, Germany, 2014. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[86] N. Korula and M. Pál. Algorithms for secretary problems on graphs and hypergraphs. In *International Colloquium on Automata, Languages, and Programming*, pages 508–520. Springer, 2009.

[87] T. G. Kurtz. Solutions of ordinary differential equations as limits of pure jump markov processes. *Journal of Applied Probability*, 7(1):49–58, 1970.

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

[89] J. D. Lee, I. Panageas, G. Piliouras, M. Simchowitz, M. I. Jordan, and B. Recht. First-order methods almost always avoid strict saddle points. *Math. Program.*, 176(1-2):311–337, 2019.

[90] S. Leonardi, A. Marchetti-Spaccamela, A. Presciutti, and A. Rosén. On-line randomized call control revisited. *SIAM Journal on Computing*, 31(1):86–112, 2001.

[91] A. Levy, H. Ramadas, and T. Rothvoss. Deterministic discrepancy minimization via the multiplicative weight update method. In F. Eisenbrand and J. Könemann, editors, *Integer Programming and Combinatorial Optimization - 19th International Conference, IPCO 2017, Waterloo, ON, Canada, June 26-28, 2017, Proceedings*, volume 10328 of *Lecture Notes in Computer Science*, pages 380–391. Springer, 2017.

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

[93] M. Lowalekar, P. Varakantham, and P. Jaillet. Competitive ratios for online multi-capacity ridesharing. *arXiv preprint arXiv:2009.07925*, 2020.

[94] Y. Ma, P. Rusmevichientong, M. Sumida, and H. Topaloglu. An approximation algorithm for network revenue management under nonstationary arrivals. *Operations Research*, 68(3):834–855, 2020.

[95] V. Manjunath. Fractional matching markets. *Games and Economic Behavior*, 100:321–336, 2016.

[96] C. McDiarmid. *On the method of bounded differences*, page 148–188. London Mathematical Society Lecture Note Series. Cambridge University Press, 1989.

[97] A. McLennan. Efficient disposal equilibria of pseudomarkets. In *Workshop on Game Theory*, volume 4, page 8, 2018.

[98] A. Mehta et al. Online matching and ad allocation. *Foundations and Trends® in Theoretical Computer Science*, 8(4):265–368, 2013.

[99] A. Mehta, A. Saberi, U. Vazirani, and V. V. Vazirani. Adwords and generalized online matching. *Journal of the ACM (JACM)*, 54(5):22–es, 2007.

[100] R. Mehta, I. Panageas, and G. Piliouras. Natural selection as an inhibitor of genetic diversity: Multiplicative weights updates algorithm and a conjecture of haploid genetics [working paper abstract]. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS 2015, Rehovot, Israel, January 11-13, 2015*, page 73. ACM, 2015.

[101] M. Mihail and T. Tröbst. Online matching with high probability. In G. Schäfer and C. Ventre, editors, *Algorithmic Game Theory*, pages 21–34, Cham, 2024. Springer Nature Switzerland.

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

[103] H. Moulin. Fair division in the age of internet. *Annual Review of Economics*, 2018.

[104] T. Nan, Y. Gao, and C. Kroer. On the convergence of tâtonnement for linear fisher markets. *arXiv preprint arXiv:2406.12526*, 2024.

[105] J. Nash. Two-person cooperative games. *Econometrica: Journal of the Econometric Society*, pages 128–140, 1953.

[106] J. B. Orlin. Improved algorithms for computing fisher's market clearing prices: Computing fisher's market clearing prices. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 291–300, 2010.

[107] M. W. Padberg and L. A. Wolsey. Fractional covers for forests and matchings. *Mathematical Programming*, 29(1):1–14, 1984.

[108] I. Panageas, T. Tröbst, and V. V. Vazirani. Time-efficient algorithms for nash-bargaining-based matching market models. In *International Conference on Web and Internet Economics*. Springer, 2024.

[109] M. Pavone, A. Saberi, M. Schiffer, and M. W. Tsao. Online hypergraph matching with delays. *Operations Research*, 2022.

[110] D. J. Roberts and A. Postlewaite. The incentives for price-taking behavior in large exchange economies. *Econometrica: journal of the Econometric Society*, pages 115–127, 1976.

[111] A. E. Roth, T. Sönmez, and M. U. Ünver. Pairwise kidney exchange. *Journal of Economic theory*, 125(2):151–188, 2005.

[112] R. Sedgewick. *Quicksort*. PhD thesis, Stanford University, May 1975.

[113] L. Shapley and H. Scarf. On cores and indivisibility. *Journal of mathematical economics*, 1(1):23–37, 1974.

[114] R. W. Simpson. *Using network flow techniques to find shadow prices for market demands and seat inventory control*. MIT Press, Cambridge, MA, 1989.

[115] K. Talluri and G. Van Ryzin. An analysis of bid-price controls for network revenue management. *Management science*, 44(11-part-1):1577–1593, 1998.

[116] T. Tröbst and V. V. Vazirani. Cardinal-utility matching markets: The quest for envy-freeness, pareto-optimality, and efficient computability. *arXiv preprint arXiv:2402.08851*, 2024.

[117] T. Tröbst and R. Udwani. Almost tight bounds for online hypergraph matching. *Operations Research Letters*, 55:107143, 2024.

[118] H. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974.

[119] V. V. Vazirani. 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, 2012.

[120] V. V. Vazirani. An extension of the Birkhoff-von Neumann Theorem to non-bipartite graphs. *arXiv preprint arXiv:2010.05984*, 2020.

[121] V. V. Vazirani. Towards a Practical, Budget-Oblivious Algorithm for the Adwords Problem Under Small Bids. In P. Bouyer and S. Srinivasan, editors, *43rd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2023)*, volume 284 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 21:1–21:14, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[122] V. V. Vazirani and M. Yannakakis. Computational complexity of the Hylland-Zeckhauser mechanism for one-sided matching markets. In *Innovations in Theoretical Computer Science*, 2021.

[123] N. Vishnoi. *Algorithms for Convex Optimization*. Cambridge University Press, 2021. To appear.

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

[125] Y. Ye. A path to the arrow–debreu competitive market equilibrium. *Mathematical Programming*, 111(1):315–348, 2008.

[126] L. Zadeh. Optimality and non-scalar-valued performance criteria. *IEEE Transactions on Automatic Control*, 8(1):59–60, 1963.

[127] L. Zhang. Proportional response dynamics in the fisher market. *Theoretical Computer Science*, 412(24):2691–2698, 2011. Selected Papers from 36th International Colloquium on Automata, Languages and Programming (ICALP 2009).

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

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