ALGORITHMS FOR FAIR PUBLIC AND PRIVATE RESOURCE ALLOCATION

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF COMPUTER SCIENCE
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Benjamin Plaut
August 2021

Abstract

This thesis explores algorithms for resource allocation, with an emphasis on fairness. "Fairness" is an extremely complex concept and our goal in this thesis is not to design an algorithm that is "truly fair" (if such a thing even exists). Rather, we consider a variety of specific objectives inspired by fairness, and analyze their properties in different resource allocation models.

The thesis is split into *private* resource allocation and *public* resource allocation. Private resource allocation handles items such as food and cars, where each individual receives a separate bundle of resources, with the key assumption being that their happiness (or "utility") depends only on what they receive, and not on what others receive. This process typically occurs through decentralized markets. In contrast, in public resource allocation, a centralized government makes a single decision that affects all citizens: for example, how to allocate a city budget, or whether to build a communal pool. Markets and governments are two of the most fundamental institutions in our society, and so any improvement to these systems can have far-reaching benefits.

For both of these settings, we take on the role of the system designer (or "social planner"). Our goal is to design an allocation method (or "mechanism") to distribute these resources in a "good" way, where "good" can have many interpretations, including but not limited to fairness concerns.

We consider two primary approaches to defining "good": *axiomatic* and *welfarist*. An axiom states a desirable property the outcome. For example, a common fairness axiom for private resource allocation is *envy-freeness*: no individual should prefer another individual's bundle to her own. In contrast, the welfarist approach is predicated on a designated *welfare function*, which assigns a single number to each possible outcome, with larger numbers being "better". Our goal then becomes to choose an outcome which maximizes the welfare function. Different welfare functions represent different priorities: for example, maximizing the sum of utilities focuses on overall societal happiness, and maximizing the minimum utility focuses on equality of happiness across individuals.

For all of these settings – private or public, axiomatic or welfarist – we provide allocation mechanisms that are provably "good" in some way. That said, each of our mechanisms has its own drawbacks, and whether or not these mechanisms are truly "fair" depends on a myriad of logistical, cultural, and philosophical factors. However, we hope that this thesis serves as an important step along the never-ending path to improve resource allocation in our society.

This thesis is dedicated to my brother Daniel

Acknowledgments

First and foremost, I would like to thank my advisor Ashish Goel. Ashish has been incredibly supportive, with regards to both research and my personal life, in times of both joy and struggle. With Ashish's mentorship, I was able to make a research contribution that I'm really proud of, much of which makes up this thesis. Also, I'm not sure how many people know this, but Ashish is wickedly funny, making our research meetings a joy.

I would also like to thank Tim Roughgarden. As my first co-author during my PhD, Tim taught me how to effectively structure my research and engage in the research community. His mentorship early on in my PhD was crucial to my growth and success.

Next, I would like to thank John Dickerson and Tuomas Sandholm for their research mentorship during my undergraduate degree, which set me up to succeed as a PhD student.

More generally, this thesis would not have been possible without my excellent collaborators. My co-authors, in reverse alphabetical order (because why are we always stuck with alphabetical order), are: Glen Weyl, Ameya Velingker, Tim Roughgarden, Kostas Kollias, Nicole Immorlica, Reyna Hulett, Sreenivas Gollapudi, Ashish Goel, and Nikhil Garg. I would also like to thank my thesis committee: Aviad Rubinstein, Tim Roughgarden, Omer Reingold, Ashish Goel, and Mohammad Akbarpour.

To Cardinal Swing: thank you for being an amazing community and helping me grow as a dancer, a leader, and a person.

To my friends: Emily Zapinski, Jennifer Yin, Ananya Rajgarhia, Suhaas Reddy, Julia Olivieri, Katherine Lee, Rosemond Ho, and many more who are too numerous to list. Thank you for all of the laughs, tears, and derps.

To Kat: thank you for helping me with my slides, listening to me explain my research in way too much detail, and supporting me through both the highs and the lows. This thesis would not have been the same without you.

To my parents: thank you for instilling in me a love of learning from an early age, and for supporting my personal and academic development at all stages of my life. I am where I am today because of you.

Finally, to Tux, Baci, Carson, and Bates: the best pets anyone could ask for, even if you might just be in it for the food.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Resource allocation is one of the fundamental underpinnings of society. Every purchase that is made, every allotment of government capital – these are instances of resource allocation. When these systems are robust and fair, society functions smoothly and its citizens are taken care of. When these systems are exploitable and unfair, society starts to crumble and inequity runs rampant. The design of "good" resource allocation mechanisms is a monumental task, and requires collaboration between a myriad of disciplines. In this thesis, we make substantial contributions to this effort from a theoretical economics and computer science perspective.

1.1 What is "fairness"?

The title of this thesis includes the word "fair", a complex and weighty term. Is equality of happiness "fair"? Is lack of envy "fair"? Is equal opportunity to exert one's preferences "fair"? The answers to these questions depend on a nuanced tapestry of logistical, cultural, and philosophical factors which we do not address in this thesis. Consequently, it would be inappropriate to claim that any of our mechanisms are truly "fair". Instead, we focus on the specific properties they satisfy: equality of happiness, envy-freeness, etc.

So why is this word included in the title? Because "fairness" is the overarching inspiration for this work. Fairness is like happiness: it is difficult to define what happiness is, and sometimes difficult to even know if we are happy. However, that does not preclude the pursuit of happiness. In the same vein, this thesis is about the pursuit of fairness.[1]

1.2 A general resource allocation framework

All of the resource allocation models used in this thesis (both public and private) fall within the same general framework. There is a fixed set of agents $N = \{1, \dots, n\}$, and a fixed set of resources M, where $m = |M|$. There is also a fixed set χ of feasible outcomes. Each model we consider will

[1] It is also worth noting that "fair" is a standard term in some research areas, e.g., fair division. That is not the primary reason we include it in the title, however.

have different assumptions on χ, and each chapter of the thesis will explicitly state the applicable assumptions.

Each agent i has a *utility function* u_i, where $u_i(\mathbf{x})$ is a real number indicating how much agent i "likes" $\mathbf{x}$. In general, we assume that agents are selfish and wish to maximize their utility $u_i(\mathbf{x})$. Similarly to χ, different chapters of the thesis will make different assumptions on what u_i can be.

For private resource allocation, each element $\mathbf{x} \in \chi$ is a $n \times m$ matrix, specifying how much of resource $j \in M$ is allocated to agent $i \in N$. Since each agent's utility depends only on what she receives, we can write $u_i(\mathbf{x}) = u_i(x_i)$, where x_i is the *bundle* given to agent i. An allocation is feasible only if the total amount allocated of each good j does not exceed the total supply of good j. Formally,

$$\sum_{i \in N} x_{ij} \leq s_j \quad \forall j \in M$$

where x_{ij} is the amount of good j in i's bundle, and s_j is the supply of good j. In general, we use i and k to refer to agents, and j and ℓ to refer to resources.

We also consider *divisible* vs *indivisible* items: a loaf of bread can be split among two, or three, or ten agents, but a car must go entirely to a single agent and thus is *indivisible*. When items are indivisible, this adds a further constraint on χ: namely any $\mathbf{x} \in \chi$ must have only integer entries. When items are divisible, each x_{ij} can be any real number.

For public resource allocation, each "resource" is a public issue that the agents must decide upon. In this case, the outcome for each issue is a real number, and the overall outcome is a vector of length m, where the jth entry is the outcome for issue j. In one chapter, we restrict issues to be binary (e.g., yes or no, which we present as 0 or 1), but in another chapter, we allow the entries to be any real number.

1.3 Objectives

If we simply let each agent take whatever she wants, that may not lead to a good outcome for society overall. Thus our general goal is to design a resource allocation mechanism that has provably "good" outcomes, *even when* agents selfishly try to maximize their utility.

In this section, we define the various notions of "good" that we consider in this thesis. As discussed in the abstract, these are divided into *welfarist* objectives, and *axiomatic* objectives.

1.3.1 Welfarist objectives

A *welfare function* Φ assigns a real number to each feasible outcome $\mathbf{x}$, with higher numbers indicating that the allocation is "better" for society. The welfare function typically only depends on the agent utilities, and not directly on the allocation. Once we select a welfare function Φ, our goal then becomes to design a mechanism such that the outcomes are guaranteed to maximize Φ.

But how should we choose a welfare function? There are many different possibilities, each representing different priorities. The most common welfare functions are

$$\text{utilitarian welfare: } \Phi(\mathbf{x}) = \sum_{i \in N} u_i(\mathbf{x})$$
$$\text{egalitarian welfare: } \Phi(\mathbf{x}) = \min_{i \in N} u_i(\mathbf{x})$$
$$\text{Nash welfare: } \Phi(\mathbf{x}) = \prod_{i \in N} u_i(\mathbf{x})$$

Utilitarian welfare cares only about the overall happiness of society, regardless of equality across agents. We think of this as the most *efficient* outcome. In contrast, egalitarian welfare[2] cares only about equality: maximizing the minimum utility means that an optimal solution will give all agents the same utility. Nash welfare is somewhere in between these two.

All of these can be generalized by a *constant elasticity of substitution (CES)* welfare function, parametrized by a real number $\rho \in (-\infty, 1]$:

$$\Phi_\rho(\mathbf{x}) = \Big(\sum_{i \in N} u_i(\mathbf{x})^\rho \Big)^{1/\rho}$$

For $\rho = 1$, we recover utilitarian welfare. The limit as $\rho \to -\infty$ yields egalitarian welfare [149, 160, 161], whereas $\rho \to 0$ yields Nash welfare [108, 129]. The closer ρ gets to $-\infty$, the more the social planner cares about individual equality (egalitarian welfare being the extreme case of this), and the closer ρ gets to 1, the more the social planner cares about efficiency, or *efficiency* (utilitarian welfare being the extreme case of this). For this reason, ρ is called the *inequality aversion* parameter, and this family of welfare functions is thought to exhibit an *equality/efficiency* tradeoff.

These welfare functions were originally proposed by Atkinson [8]; indeed, his motivation was to measure the level of inequality in a society. Despite being extremely influential in the traditional economics literature (see [53] for a survey), the CES welfare function has received almost no attention in the computational economics community.[3]

In general, different values of ρ will lead to different optimal outcomes. Thus whenever we say that an outcome maximizes CES welfare, we always mean with respect to a particular value of ρ. When the value of ρ is clear from context, we omit this and just say that the omit maximizes CES welfare.

1.3.2 Axiomatic objectives

The welfarist approach is based on maximizing a given function of the agent utilities. An alternative approach is to posit certain properties – or *axioms* – that the outcome should satisfy. The axioms should be easy to verify if the agent utilities are known. This is in contrast to welfarism, there is generally no way to verify that $\mathbf{x}$ is optimal without solving the corresponding optimization problem.

[2] This is also known as max-min welfare or Rawlsian welfare.

[3] To our knowledge, only one computational economics paper outside of our work has studied CES welfare in any context: [7].

Axioms are most commonly considered in the context of fairness. For private resource allocation, perhaps the most popular is *envy-freeness*, which states that no agent should prefer another agent's bundle of resource to her own. Formally, $\mathbf{x}$ is envy-free if for all agents i and k,

$$u_i(x_i) \geq u_i(x_k)$$

A related property is *proportionality*, which states that each agent's utility for her own bundle should be at least $1/n$ times her utility for her favorite possible outcome. Formally,

$$u_i(x_i) \geq \frac{1}{n} \max_{\mathbf{y} \in \chi} u_i(\mathbf{y})$$

It is typically assumed that the utilities are monotonic, i.e., getting more resources cannot decrease an agent's utility. Under that assumption, $\max_{\mathbf{y} \in \chi} u_i(\mathbf{y})$ is simply agent i's utility for getting all of the resources.

There are also variants of each of these properties that we will discuss when they become relevant later in the thesis. Depending on the specific problem (i.e., the assumptions on χ and the utilities), achieving envy-freeness and/or proportionality may be trivial, an interesting challenge, or impossible.

Finally, note that neither of these axioms make much sense in the context of public resource allocation: when all agents get the same "bundle", how do we interpret envy-freeness? Designing axioms for public resource allocation is a vibrant area of research, one that we consider in Chapter 8 of this thesis.

1.3.3 Intersection of welfarist and axiomatic approaches

For many welfare functions, any optimal solution is guaranteed to satisfy certain axioms. Indeed, one of our mechanisms for satisfying (a variant of) envy-freeness will involve maximizing a (non-traditional) welfare function. Furthermore, any CES welfare function is guaranteed to satisfy the following axioms:

1. Monotonicity: if one agent's valuation increases while all others are unchanged, the welfare function should prefer the new allocation.
2. Anonymity: the welfare function should treat all agents the same,
3. Continuity: the welfare function should be continuous.[4].
4. Independence of common scale: scaling all agent valuations by the same factor should not affect which allocations have better welfare than others.
5. Independence of unconcerned agents: when comparing the welfare of two allocations, the comparison should not depend on agents who have the same valuation in both allocations.

[4] A slightly weaker version of continuity is often used: if an allocation $\mathbf{x}$ is strictly preferred to an allocation $\mathbf{y}$, there should be neighborhoods $N(\mathbf{x})$ and $N(\mathbf{y})$ such that every $\mathbf{x}' \in N(\mathbf{x})$ is preferred to every $\mathbf{y}' \in N(\mathbf{y})$. This weaker version only requires a welfare *ordering* and does not require that this ordering be expressed by a function. However, any such ordering which also satisfies the rest of our axioms is indeed representable by a welfare function [60], and so both sets of axioms end up specifying the same set of welfare functions/orderings.

6. The Pigou-Dalton principle: when choosing between equally efficient allocations, the welfare function should prefer more equitable allocations [56, 142].

Furthermore, disregarding monotonic transformations of the welfare function (which of course do not affect which allocations have better welfare than others), the set of welfare functions satisfying these axioms is exactly the set of CES welfare functions with $\rho \in (-\infty, 0) \cup (0, 1]$, including Nash welfare [128].[5] This axiomatic characterization shows that we are not just focusing on an arbitrary class of welfare functions: CES welfare functions are arguably the most reasonable welfare functions.

1.3.4 Motivations behind CES welfare and envy-freeness

CES welfare and envy-freeness have distinctly different flavors, beyond one being an axiom and the other a welfare objective. CES welfare aims to implement a particular equality/efficiency tradeoff for a (potentially large) group of agents as a whole. In contrast, envy-freeness is based on interpersonal comparison: no agent should be treated poorly *in comparison to other agents*.[6] Notably, CES welfare only considers the utility each agent has for her own bundle, whereas envy-freeness considers each agent's utility for every agent's bundle.

It is worth noting that the empty allocation is trivially envy-free: all agents are treated the same, and they are all treated "badly" (i.e., by withholding resources). We typically avoid this degenerate case by requiring that all goods be allocated, but the general point still stands. CES welfare does not suffer from this issue, but that does not make it a "better" objective: it is simply different.

We suggest that although CES welfare and envy-freeness can both be motivated by fairness, they are attempting to capture different aspects of fairness: CES welfare is closer to "equitable", while envy-freeness is closer to "non-discriminatory". The conditions, if any, under which envy implies discrimination or CES welfare corresponds to equitability are far beyond the scope of this thesis. However, it is still valuable to reflect on the differences in motivation between the objectives we consider.

1.4 Contributions of the thesis

With these concepts in hand – a general resource allocation framework and a discussion of objectives – we can describe the primary contributions of this thesis.

1.4.1 Part I: welfarist private resource allocation

Part I studies market mechanisms with welfarist objectives for private resource allocation. We will provide a thorough introduction to markets in Section 1.5, but we need to provide some quick definitions in order to describe our results. There are two main market models: *fixed-budget*, where each agent has a fixed budget of money to spend and seeks to maximize her utility subject to that

[5] This actually does not include max-min welfare, which satisfies weak monotonicity but not strict monotonicity.

[6] Envy-freeness also makes the most sense for a small group of agents, where each agent can compare their bundle to each other agent.

budget, and *quasilinear*, where each agent can spend as much as she wants, and instead she tries to maximize her value for her bundle minus the cost of the bundle.

For both of these models, the majority of the economics literature assumes linear pricing, i.e., buying twice as much always costs twice as much. It is known that linear pricing equilibria in the fixed-budget model and quasilinear model maximize Nash welfare and utilitarian welfare, respectively [72, 73].

In contrast, Part I focuses on *nonlinear* pricing. The overarching themes of this part of the thesis are nonlinear pricing, CES welfare, and the connection between. The material in this part of the thesis is based on three published papers: [94, 96, 143].

Markets beyond Nash welfare for bandwidth allocation

Chapter 2 uses the fixed-budget model to design nonlinear pricing rules which lead to CES welfare maximization.[7] We show that for *bandwidth allocation utilities*[8], nonlinear pricing allows us to obtain market equilibria which maximize (budget-weighted) CES welfare. Furthermore, these prices take a simple form: $p(x_i) = \sum_{i \in N} q_j x_{ij}^{1-\rho}$. Here $p(x_i)$ is the price of bundle x_i, ρ is the parameter of the CES welfare function to be maximized, and $q_1, \dots, q_m$ are optimal Lagrange multipliers in a convex program for maximizing CES welfare. The structure of $p(x_i)$ yields a simple way to compute these price curves: decide on a parameter ρ, ask the agents for their utilities, and solve the corresponding convex program to obtain $q_1, \dots, q_m$.

Optimal Nash equilibria for bandwidth allocation

The above approach requires agents to truthfully report their utilities, which they may not be willing and/or able to do. Chapter 3 provides a way around this. We consider the same basic problem – CES welfare maximization for bandwidth allocation utilities in the fixed-budget model – but focus on strategic behavior. Inspired by the classic trading post mechanism, we propose a mechanism where agents use their fixed budget to bid on goods, and each good is allocated as a nonlinear function of the bids. Rather than asking agents to report their utilities and then explicitly computing prices ourselves, this bidding mechanism allows prices to arise naturally from agents' behavior. We show that the Nash equilibria of this mechanism are also guaranteed to maximize CES welfare, thereby preserving the welfare guarantees from Chapter 2, while also handling strategic behavior.

Counteracting inequality in markets via convex pricing

We view Chapter 4 as the culmination of this line of work. In this chapter, we are able to obtain similar results as before, but for a much wider range of agent utilities. To do this, we pivot to the quasilinear market model, which is much better suited to the analysis of CES welfare maximization (see Section 1.5.2). We show a similar result to that of Chapter 2: for any problem instance, for

[7] This chapter also considers the more general question of what set of market equilibria is possible for nonlinear pricing, but those results are less crucial to the overarching contribution of the thesis.

[8] In bandwidth allocation, each good represents a link in a network, and each agent wants to transmit data through a fixed path of links. An agent's utility is the minimum bandwidth she receives among all links in her desired path.

any $\rho \in (0,1]$, there exists a pricing rule such that all of the market equilibria maximize CES welfare. Chapter 2 only showed this for bandwidth allocation utilities, but in this chapter, we only assume that utilities are concave, differentiable, and homogeneous of fixed degree. The pricing rule takes a similar form: $p(x_i) = \big(\sum_{j\in M} q_j x_{ij}\big)^{1/\rho}$ where ρ determines the CES welfare function to be optimized, and $q_1,\dots,q_m$ are optimal Lagrange multipliers from a convex program for maximizing CES welfare.[9] This pricing rule is convex for $\rho \in (0,1]$, and we discuss connections to real-world convex pricing: in particular, for water.

Whereas the result from Chapter 2 feels quite specialized, we feel that this result is general enough to indicate a fundamental connection between CES welfare and nonlinear pricing, mirroring the relationship between linear pricing and utilitarian welfare (in the quasilinear model) or Nash welfare (in the fixed-budget model).

Comparison of our results in the fixed-budget and quasilinear models

In the fixed-budget model, we show that for any $\rho \in (-\infty,1)$, the pricing rule $p(x_i) = \sum_{j\in M} q_j x_{ij}^{1-\rho}$ leads to (budget-weighted) CES welfare maximization for bandwidth allocation utilities (Chapter 2). In the quasilinear model, we show that for any $\rho \in (0,1]$, the pricing rule $p(x_i) = \big(\sum_{j\in M} q_j x_{ij}\big)^{1/\rho}$ leads to CES welfare maximization for a much larger set of agent utilities (Chapter 4). Although there are some superficial differences between these results (e.g., the permissible values of ρ, the specific form of the pricing curve), we can unify them to some extent.

Essentially, we show that if we use the quasilinear model to obtain a market equilibrium with pricing rule $p(x_i) = \big(\sum_{j\in M} q_j x_{ij}\big)^{1/\rho}$, and let B_i be the spending of agent i at equilibrium, then $(\mathbf{x}, p)$ also forms a fixed-budget equilibrium for agent budgets $B_1,\dots,B_n$. We can extend this idea to show that maximizing CES welfare with respect to a fixed $\rho \in (0,1]$ is actually equivalent to maximizing budget-weighted CES welfare with respect to $\rho - 1$ (where agent i's budget is her equilibrium spending in the quasilinear model). Formally:

Theorem 4.9.2. *Assume agent valuations are concave, differentiable, and homogeneous of fixed degree, let $\rho \in (0,1]$, let $q_1,\dots,q_m \in \mathbb{R}_{\geq 0}$, and let $p(x_i) = \big(\sum_{j\in M} q_j x_{ij}\big)^{1/\rho}$. Assume $\mathbf{x}$ maximizes CES welfare with respect to ρ, and let $B_i = p(x_i)$. Then $\mathbf{x}$ maximizes budget-weighted CES welfare for $\rho - 1$.*

This does not fully unify the results. However, we feel that the remaining differences are artifacts of the strong assumptions on utilities we make in Chapter 2 in order to obtain our result in the fixed-budget model (see Section 1.5.2). Again, we view the result in the quasilinear model to be the capstone result of this part of the thesis, one which points to a fundamental connection between CES welfare and nonlinear pricing.

[9] This result has the same strategic concern as that of Chapter 2, and unfortunately the trading post approach does not work here. However, we are able to provide a truthful mechanism for the case of a single good.

1.4.2 Part II: axiomatic private resource allocation

In Part II, we turn to axiomatic approaches to private resource allocation. We also focus on indivisible items. Perhaps the most pervasive fairness axiom is envy-freeness, which states that no agent prefers another agent's bundle to her own. Unfortunately, this is not always possible with indivisible goods: consider two agents and a single good. Part II consists of two chapters, each representing a different approach to the fundamental asymmetry of indivisible goods. The mechanisms in Part II are not market mechanisms, but instead use a variety of other techniques. The relevant published papers for Part II are [144] and [145].

A new fairness axiom: envy-freeness up to any good (EFX)

The goal of Chapter 5 is to find a relaxation of envy-freeness which is achievable for indivisible goods. We study a property called *envy-freeness up to any good*, also known as EFX. This axiom permits envy in an allocation, but if agent i envies agent j, removing any good from j's bundle should eliminate the envy.[10] Note that this immediately resolves the case of two agents a single good.

Unfortunately, the question of whether EFX allocations are guaranteed to exist seems extremely difficult. This axiom was proposed by [40], who were unable to resolve the question of guaranteed existence "despite significant effort". Although we too were unable to resolve this question in general, we were able to resolve it in some special cases, as well as prove an exponential lower bound on the number of queries required to compute an EFX allocation. Our work, when originally published in 2018, constituted the first formal results regarding EFX. After significant amount of follow-up work (e.g., [3, 39, 46, 92]), it was recently shown that EFX allocations are guaranteed to exist for three agents with additive valuations [45].

Communication complexity of envy-freeness (and friends)

Chapter 6 takes a different approach. This chapter focuses on full envy-freeness, accepting that envy-free allocations do not always exist. Instead, our goal is to efficiently determine whether or not an envy-free allocations exists in a particular problem instance. We view this through the lens of *communication complexity*, where agents can do as much offline computation as they wish, and seek to minimize the amount of information they must transmit amongst themselves. We study the same questions for proportionality, approximations of envy-freeness, and approximations of proportionality, all of which have the same property that existence is not guaranteed for indivisible items.[11] We consider many different problem parameters, including number of players, randomized vs deterministic complexity, and restrictions of agent utility functions. For every combination of parameters, we complete resolve whether the communication complexity is polynomial or exponential (in the number of goods).

[10] This is simply a thought experiment used in the definition of EFX: the good is not actually removed from j's bundle.

[11] Typically, EFX is considered a "relaxation" and not an approximation. As discussed, guaranteed existence of EFX is an open question.

1.4.3 Part III: public resource allocation

Finally, we consider public resource allocation in Part III. Both chapters in Part III consider *public decision-making*, where each of the m resources is an issue for which the group must make a decision (where the "decision" is a real number).

Like Part I, this part of the thesis has a strong emphasis on markets. Markets for private resource allocation are known to have many desirable properties (see Section 1.5), but have not been studied as thoroughly for public resource allocation. In fact, we can draw a strong parallel between this and Part I: markets with linear pricing are well-studied, so we study nonlinear pricing; markets for private goods are well-studied, so we study public goods. That said, the emphasis of Part I is less on understanding markets for public goods in general, and more on designing mechanisms for public decision-making inspired by markets. Part III is based on material from [88] and [103].

Markets for public decision-making

In Chapter 7, our goal is to adapt traditional private goods markets to the public decision-making setting. We first consider a straightforward adaptation of the fixed-budget model: essentially, we set a price for each issue, give each agent one unit of artificial currency, and allow agents to "buy probability" for their preferred outcome (we focus on binary issues with randomized outcomes). However, the market equilibria for this mechanism can have arbitrarily poor welfare. This motivates a more sophisticated type of pricing. We propose *pairwise pricing*, where for each issue, we assign a price to every pair of agents who disagree. This construction allows us to reduce the problem to an equivalent private goods market, and we use this connection that the resulting equilibria in the public decisions market has optimal Nash welfare. Our reduction also allows certain results for private goods markets to be immediately lifted to the public decision-making setting.

A new fairness axiom for public decision-making: equality of power

In Chapter 8, we consider a market mechanism with a combination of welfarist and axiomatic properties. The fundamental premise of our mechanism is that each agent should have equal opportunity to affect the group outcome. We term this *equality of power*, where we interpret "power" as the externality an agent causes to the overall (utilitarian) welfare. Specifically, given a current proposed outcome $\mathbf{x}$, we ask each agent to propose a new outcome $\mathbf{x}'$ subject to the constraint that the decrease in welfare from $\mathbf{x}$ to $\mathbf{x}'$ is at most some small constant γ. By equalizing γ across agents, we ensure equality of power. Equilibrium occurs when the vectors of the proposed new outcomes cancel out, and we show that there always exists an equal power equilibrium which maximizes utilitarian welfare. There are additional technical details regarding the scaling of individual utilities which we defer to Chapter 8.

Taken together, our results shed light on a variety of notions of fairness in a variety of models, for both public and private resource allocation. Although "true" fairness depends on a myriad of factors which we do not consider, we hope that our work contributes to a greater understanding of how resource allocation can be improved in our society.

1.5 Markets

Before diving into the body of the thesis, we provide a technical overview of markets, which will play an important role throughout the thesis (in particular, Parts I and III).

Markets are one of the oldest mechanisms for distributing (private) resources; indeed, commodity prices were meticulously recorded in ancient Babylon for over 300 years [165, 166]. In a market, buyers and sellers exchange goods according to some sort of pricing system, and *market equilibrium*[12] occurs when the demand of the buyers exactly equals the supply of the sellers. This concept was first studied by Walras in the 1870's [173]. In 1954, Arrow and Debreu showed that under some conditions, a market equilibrium is guaranteed to exist [6].

In modern day, markets continue to be the primary way in which private resources are distributed. Stores set prices for their products, and individuals peruse the wares, making purchases if they wish. Markets operate based on the interaction of two fundamental concepts: prices and utilities. Each individual has some utility for each product, or perhaps each combination of products. When deciding whether to make a purchase, the individual compares their utility for the product with the price: if the price is too high, they will not make a purchase, even if the product is very valuable to them. Importantly, this process is frequently subconscious: the individual simply decides whether the purchase is "worth it" at the given price. If a store notices that a product is not selling well, they may choose to lower the price. Conversely, if a product is frequently selling out, they may raise the price. This process – known as *tâtonnement*, and also proposed by Walras [173] – will ideally converge to a market equilibrium.

In this thesis, we focus on markets without production, and where each agent is either a buyer or a seller (but not both). This is because we are imagining ourselves to be the social planner who initially controls all the resources, so we essentially function as the unique seller.

1.5.1 Formal market models

To formalize market equilibrium, we need the concept of a *demand set*. Given a *pricing rule* p which assigns a real number price to any bundle, agent i's demand set $D_i(p)$ is the set of bundles which maximize her utility given to p. Demands sets are best understood for *linear pricing*, i.e., when $p(x_i) = \sum_{j \in M} q_j x_{ij}$, with q_j being the (constant) price of good j.

There are two standard ways to define the demand set. The first assumes that each agent has a fixed budget B_i, and the cost of her purchase must be at most B_i. Therefore

$$D_i(p) = \underset{x_i \in \mathbb{R}^m_{\geq 0}:\ p(x_i) \leq B_i}{\arg\max} u_i(x_i)$$

Alternatively, the price of the bundle can be directly factored in to the utility. *Quasilinear utilities* assumes that each agent wishes to maximize her value for her bundle minus the price she pays. Formally, if v_i is agent i's *valuation* function, then agent i's quasilinear utility is $u_i(x_i) = v_i(x_i) -$

[12] This is also known as Walrasian equilibrium, competitive equilibrium, and general equilibrium, depending on the context.

$p(x_i)$, and the demand set is defined by

$$D_i(p) = \operatorname*{arg\,max}_{x_i \in \mathbb{R}^m_{\geq 0}} \big(v_i(x_i) - p(x_i)\big)$$

Some chapters of the thesis will use the fixed-budget model, and others use the quasilinear model. Each chapter will explicitly state which model it is using (if it uses either).

Once the demand set is defined, we have the following definition of market equilibrium:

Definition 1.5.1. *Given an allocation* $\mathbf{x}$ *and pricing rule* p*,* $(\mathbf{x}, p)$ *is a* market equilibrium *if both of the following hold:*

1. *Each agent receives a bundle in her demand set:* $x_i \in D_i(p)$ *for all* $i \in N$*.*
2. *The market clears: for all* $j \in M$*,* $\sum_{i \in N} x_{ij} \leq s_j$*, and for all* $j \in M$ *with nonzero cost,* $\sum_{i \in N} x_{ij} = 1$.[13]

In the fixed-budget model, it is more accurate to write that $(\mathbf{x}, \mathbf{B}, p)$ is a market equilibrium, where $\mathbf{B} = (B_1, \dots, B_n)$. However, when it is clear from context (typically, when all budgets are the same), we simply write $(\mathbf{x}, p)$. The quasilinear model has no budgets, so this is a non-issue.

Market equilibria have many nice properties, which we will explore in detail later. They include:

1. The mechanism is *anonymous*: different agents purchasing the same bundle always pay the same price.
2. Any market equilibrium[14], i.e., immediately satisfies envy-freeness: if agent i preferred agent j's bundle, agent i would have simply bought x_j instead.[15]
3. There is a strong sense of *agency*: each individual has total control over their purchase, and can be confident that their selection is optimal (based on whatever criteria they wish). This is in contrast to other mechanisms (including ones in this thesis) which simply hand each agent a bundle and say "here it is".
4. Markets are a familiar concept to the general population, so mechanisms and analysis relating to markets may be more easily explainable.
5. There is a plethora of research into market equilibrium that we can lean on when analyzing these mechanisms. In contrast, the non-market mechanisms we use in this thesis typically must be analyzed from scratch.

Finally, and perhaps most importantly, market equilibria for linear pricing offer strong welfare guarantees. Our goal in some of the chapters of this thesis will be to prove welfare guarantees for nonlinear pricing.

[13] We say that good j has nonzero cost for j if there is a bundle x_i such that $x_{i\ell} = 0$ for all $\ell \neq j$, but $p(x_i) > 0$. If $p(x_i) = \sum_{j \in M} q_j x_{ij}$, this is equivalent to $q_j > 0$.

[14] This assumes that the pricing rule is anonymous.

[15] Note that for the fixed-budget model, this assumes that all agents have the same budget.

CES welfare and quasilinear utility

When considering quasilinear utility in this thesis, we will consider a slightly different version of the CES welfare function where we consider each agent's valuation instead of her overall utility:

$$\Phi_\rho(\mathbf{x}) = \Big(\sum_{i\in N} v_i(\mathbf{x})^\rho\Big)^{1/\rho}$$

The reason is that there implicitly exists a seller whose utility is $\sum_{i\in N} p(x_i)$, i.e., the total payment. This conceptually "cancels out" the payment terms in each agent's utility $u_i(x_i) = v_i(x_i) - p(x_i)$, so we do not consider it from a welfare perspective. For simplicity, we use the above definition of CES welfare throughout the thesis, and when not considering quasilinear utility, we simply let $u_i = v_i$.

1.5.2 Comparison of the fixed-budget and quasilinear models

Although these two models share many of the same conceptual messages, their technical implications differ – in particular, the welfare guarantees. For linear pricing, the welfare properties of each model are well-understood (and different):

Theorem 1.5.1 ([72, 73])**.** *For the fixed-budget model and linear pricing, the market equilibrium allocations are exactly the allocations maximizing budget-weighted Nash welfare:* $\Phi_{Nash}(\mathbf{x}, \mathbf{B}) = \prod_{i\in N} u_i(x_i)^{B_i}$.

Theorem 1.5.2 (First Welfare Theorem)**.** *For the quasilinear model and linear pricing, the market equilibrium allocations are exactly the allocations maximizing utilitarian welfare.*[16]

As we discussed earlier, utilitarian welfare is thought be the most efficient, with Nash welfare being something of a compromise between equality and efficiency. So why do these two market models place the linear pricing equilibria at such different points along the equality/efficiency tradeoff?

The key is that the fixed-budget model maximizes *budget-weighted* Nash welfare. If all agent budgets are the same, the outcome may be "equal" in some sense, but if the budgets vary substantially, we lose that claim to equality. Furthermore, consider the following formal connection:

Theorem 4.9.1. *Suppose* $(\mathbf{x}, p)$ *is a market equilibrium in the quasilinear model, and let* $B_i = p(x_i)$ *for each* $i \in N$*. Then* $(\mathbf{x}, \mathbf{B}, p)$ *is a market equilibrium in the fixed-budget model.*

That is, if we use the quasilinear model to establish how much each agent i will spend at equilibrium, and set that to be i's budget in the fixed-budget model, then the two models become equivalent. This of course means that the optimized welfare functions – for linear pricing, utilitarian welfare and budget-weighted Nash welfare – also coincide. This theorem does not only apply to linear pricing, and is our key tool in proving Theorem 4.9.2, which (somewhat) unifies CES welfare maximization in the fixed-budget and quasilinear models. The reader may recall our discussion of this in Section 1.4.1.

[16] This is a special case of the First Welfare Theorem, not the entirety of the First Welfare Theorem.

Advantages of the quasilinear model

As discussed above, these two models are strongly connected, but they still have their own advantages and disadvantages. To understand this, we must consider perhaps the most fundamental difference between the two models: the quasilinear model allows agents to express the *absolute magnitude* of their preferences, and the fixed-budget model does not.[17] Both models are agents to express their relative preferences between goods, i.e., which goods they like more than others. But in the fixed-budget model, each agent will always spend exactly their budget, and they are only deciding how to distribute their budget across goods. Thus if agent 1 values every good twice as much as agent 2, there is no way to express that.

Consequently, the fixed-budget model must treat those two agents the same, so the outcome must not change if an agents scales up her preferences by some factor (this is known as *independence of individual scale*). The only CES welfare function which satisfies this is Nash welfare, the product of utilities. Thus it makes sense that the (linear pricing) equilibria maximize (budget-weighted) Nash welfare in the fixed budget model.

Allowing nonlinear pricing does generalize the model somewhat, but it still leaves agents unable to express the absolute magnitude of their preferences. This means that if we wish to maximize other CES welfare functions – which do not satisfy independence of individual scale – we must make some sort of assumption on the scale of utilities, e.g., every agent's maximum utility is 1.

We do have some results for CES welfare in this model (Chapters 2 and 3), but overall, the quasilinear model is much better suited to maximizing non-Nash CES welfare. In the quasilinear model, agents not only choose how to distribute their budget across goods, they also *choose their budget*. Thus it becomes possible to elicit the absolute magnitude of individual utilities. For this reason, in the quasilinear model, we are able to prove welfare guarantees for a much wider range of agent utilities (Chapter 4).

Advantages of the fixed-budget model

To be honest with the reader, the primary reason that some of our work uses the fixed-budget model is quite mundane: the fixed-budget model is simply used more frequently in our subarea. Thus our earlier private goods work use the fixed-budget model, and our later private goods work switched to the quasilinear model.

That said, there are still advantages of the fixed-budget model. The first is that is that for some application domains – in this thesis, public resource allocation – we are not imagining a traditional market with real money (see Chapter 7). Instead, we are using a market-based framework where each agent is given some amount of "fake money" that they can use to express their preferences. In this case, we can ensure a sense of equality by giving each agent the same amount of fake money.

[17] Note that absolute magnitude of preferences is not meaningful without a unit of comparison: in the case of the quasilinear model, money fills that role. Since agent i wants to maximize $v_i(x_i) - p(x_i)$, we are assuming that $v_i(x_i)$ is expressed in the same units as $p(x_i)$.

Part I

Welfarist Private Resource Allocation

Chapter 2

Markets beyond Nash welfare for bandwidth allocation

In this chapter, we study the allocation of private goods via a market mechanism in the fixed-budget model. We focus on agents with Leontief utilities, a class of utilities which generalizes the bandwidth allocation problem. The majority of the economics and mechanism design literature has focused on *linear* prices, meaning that the cost of a good is proportional to the quantity purchased. Equilibria for linear prices are known to be exactly the maximum Nash welfare allocations.

Price curves allow the cost of a good to be any (increasing) function of the quantity purchased. First, we show that an allocation can be supported by strictly increasing price curves if and only if it is *group-domination-free*. A similar characterization holds for weakly increasing price curves. We use this to show that given any allocation, we can compute strictly (or weakly) increasing price curves that support it (or show that none exist) in polynomial time. These results use a variant of Farkas' Lemma along with a combinatorial argument to construct piecewise linear price curves. For our second main result, we use Lagrangian duality to show that in the bandwidth allocation setting, any allocation maximizing a CES welfare function can be supported by price curves. Taken together, our results show that nonlinear pricing opens up multiple possibilities beyond Nash welfare for market equilibria.

2.1 Introduction

In this chapter, we focus on the fixed-budget market model: each agent i has a set budget B_i, and has no value for leftover money. The simplest version of this model is a *Fisher market*, first proposed in 1892 by Irving Fisher [24, 80]. In Fisher markets, the pricing function p is linear: each good j has a single real-number price p_j, and the cost of a bundle x_i is

$$p(x_i) = \sum_{j \in M} p_j x_{ij}$$

The market equilibria of Fisher markets are well-understood: in particular, they are guaranteed to maximize Nash welfare [72, 73, 104]. However, much less attention has been given to *nonlinear* prices.

There are three motivations behind the work in this chapter. First, in real market economies, prices are often not linear, and depend on the quantity purchased[1]. We refer to prices of this form as *price curves*. For example, "buying in bulk" may allow agents to purchase twice as much of some resource for less than twice the price. In this case, the marginal price decreases as more of the good is purchased. On the other hand, for a scarce resource, a central authority may choose to impose increasing marginal costs to ensure that no single individual can monopolize the resource. Many countries use this type of convex pricing for water in an effort to improve equality of access [175]. A tremendous amount of work has been devoted to understanding the nature of linear prices, despite the pervasiveness of price curves in the real world. This chapter attempts to ask the same fundamental questions of price curves that have been answered for linear prices.

Second, imagine a social planner or mechanism designer who wishes to design a pricing scheme to maximize some welfare function. If the social planner is happy with Nash welfare, then great! They can just use linear pricing. But what if the social planner wishes to maximize a different welfare function? Is it possible that using price curves instead of linear prices allows a wider set of allocations to be equilibria? In particular, are there welfare functions other than Nash welfare such that welfare-maximizing allocations can always be supported by price curves? (We say that an allocation can be *supported* by price curves if there exist prices curves that make that allocation an equilibrium.) Our work answers these questions in the affirmative.

The third motivation involves a more conceptual connection between markets and welfare functions, both of which have been extensively studied in the economics literature. We know that linear-pricing equilibria correspond to maximizing Nash welfare, but does this connection go deeper? Our work hints at an affirmative answer to this question as well.

2.1.1 Leontief utilities

We assume throughout the chapter that agents have Leontief utility functions. An agent with a Leontief utility function desires the goods in fixed proportions, e.g., one unit of CPU for every two units of RAM. We can express agent i's utility as

$$\min_{j \in M : w_{ij} \neq 0} \frac{x_{ij}}{w_{ij}}$$

Recall that M is the set of goods, and x_{ij} is the quantity of good j which agent i receives. When agents have Leontief utilities, a market equilibrium is guaranteed to exist [6].

Leontief utilities exhibit certain convenient properties that other utility functions do not. In particular, such an agent will always purchase her goods exactly in the same proportions, and all

[1] One consequence of this is that there can be an incentive for agents to "team up", which is not the case in linear pricing. For example, it could be cheaper for one person to purchase the resource in bulk and then distribute it, as opposed to each person buying her own: imagine ordering pizza for a party. We do not consider strategic behavior in this chapter; see Section 2.3 for additional discussion.

that changes is how much she purchases. We also assume that each agent has the same amount of money to spend. However, most of our results do extend to the case of unequal budgets, as noted later on.

2.1.2 Bandwidth allocation

Resource allocation with Leontief utilities generalizes the problem of network bandwidth allocation, which is a well-studied area in its own right (for example, the work of Kelly [110] on proportional fairness). In bandwidth allocation, each agent wishes to transmit data along a fixed route of links, and thus desires bandwidth for exactly those links in equal amounts. In our setting, each link corresponds to a good, so agent i has $w_{ij} = 1$ if and only if link j is in her desired route.

In the bandwidth allocation setting, price curves correspond naturally to a signaling mechanism that provides congestion signals (e.g., in the form of a packet mark or drop) and an end-point protocol such as TCP [42] corresponds naturally to agent responses. It has been known that different marking schemes (such as RED and CHOKe [81, 139]) and versions of TCP lead to different objective functions [138], with CES welfare (also known as "α-fairness") being one such objective [20, 126]. However, a market-based understanding was developed only for Nash Welfare, starting with the seminal work of Kelly et al. [110].

2.2 Results

A necessary and sufficient condition for the existence of price curves

Section 2.5 presents our first main result, which concerns the first motivation described above: trying to understand fundamental properties of price curve equilibria. In particular, this section answers the following question: given some allocation, is there a way to tell whether there exist price curves that make this allocation an equilibrium? Furthermore, can such price curves be efficiently computed?

The answer boils down to a property we call *group-domination-freeness*. Roughly, a set of agents $\mathbf{a}$ group-dominates a set of agents $\mathbf{b}$ if these sets are the same size, but for every good j and every threshold $\tau \in \mathbb{R}_{\geq 0}$, the number of agents in $\mathbf{a}$ receiving at least of τ of good j is at least as large as the number of agents in $\mathbf{b}$ receiving at least τ of good j. The formal definition of group domination is given in Section 2.5. An allocation is group-domination-free (GDF) if no group dominates any other group. We show that an allocation can be supported by *strictly* increasing price curves if and only if the allocation is GDF (Theorem 2.5.1)[2]. This characterization results in a polynomial time algorithm to compute the underlying price curves or show that none exist (Theorem 2.5.2). Section 2.8 gives an analogous characterization theorem and polynomial time algorithm for weakly increasing price curves (Theorems 2.8.1 and 2.8.2).

Although the definition of group domination may seem slightly technical, we also demonstrate its relation to the concept of stochastic dominance, and argue that it may in fact be interpreted as

[2] This result extends to the setting of unequal budgets if one instead considers "budget-weighted group-domination-freeness". We elaborate on this in Section 2.5.

a fairness notion. The stochastic dominance interpretation will also suggest that group domination is quite a strong property, and hence group-domination-freeness is a quite a weak assumption.

The proof of these characterization theorems involves the construction of a special matrix we call the *agent-order matrix* A, which is a function of the allocation. We show that existence of strictly increasing price curves is captured by *strongly* positive solutions (every entry of the solution vector is positive) to $A\mathbf{y} = \mathbf{0}$. We relate group-domination-freeness to a property of this matrix, and then invoke a duality theorem equivalent to Farkas' Lemma [141] to complete the proof. The algorithm for computing price curves is a linear program involving the agent-order matrix.

Maximum CES welfare allocations can be supported in bandwidth allocation

Our second main result concerns the second and third motivations: a social planner who wishes to maximize a welfare function other than Nash welfare, and understanding the connection between markets and welfare functions. We know that the maximum Nash welfare allocations can be supported by linear prices. If we allow price curves, are there other welfare functions whose maxima can be supported?

First, we will need some assumption on the agents' weights (recall that w_{ij} denotes agent i' weight for good j). To see this, consider just two agents and one good. Since the agents have the same budget, they must receive equal amounts of the good no matter the price curve. However, if one agent derives less utility per unit of the good, this allocation doesn't maximize any CES welfare function except for Nash welfare[3]. One natural way to handle this is to assume that the agents' weight vectors are normalized in some manner. The bandwidth allocation setting – $w_{ij} \in \{0, 1\}$ for all i and j – constitutes one such possibility (refer to Section 2.1.2 for additional discussion of this setting).

Our second main result is that in the bandwidth allocation setting, the welfare-maximizing allocations for any fixed CES welfare function with $\rho \in (-\infty, 0) \cup (0, 1)$ can be supported by price curves (Theorem 2.6.1). We prove this by writing a convex program to maximize CES welfare, and using duality to construct explicit price curves. Furthermore, these price curves take on a natural form: the cost of buying x of good j is $q_j x^{1-\rho}$, where $q_j \geq 0$ is a constant derived from the dual[4]. This result can be thought of as extending the work on price-based congestion control (pioneered by Kelly et al. [110]) beyond Nash welfare to almost all CES welfare functions.

We also prove a converse of sorts: if an allocation $\mathbf{x}$ can be supported by price curves of the form $q_j x^{1-\rho}$, and the supply is exhausted for every good with nonzero price (i.e., $q_j \neq 0$), then $\mathbf{x}$ is a maximum CES welfare allocation (Theorem 2.6.1). This is analogous to the famous result of Eisenberg and Gale: the linear-pricing equilibrium allocations are exactly the allocations maximizing Nash welfare [72, 73].

One may wonder if Theorem 2.6.1 could be extended to $\rho = 1$, i.e., maximizing the sum of utilities.

---

[3]This example actually holds for a much wider class of utilities, not just Leontief. This is because for a single good, all anyone can do is buy as much of that good as they can.

[4]These results extend to agents with unequal budgets if we instead consider the "budget-weighted CES welfare", i.e., $\left(\sum_{i\in N} B_i u_i^\rho\right)^{1/\rho}$, where B_i is agent i's budget. We discuss this in Section 2.6.4. The price curves will take the exact same form.

	agent 1	agent 2	agent 3
good 1	1	0	1
good 2	0	1	1

Example 2.1: A bandwidth allocation instance where no maximum utilitarian welfare allocation can be supported. The table above gives each agent's weight $w_{ij} \in \{0, 1\}$ for each good. Utilitarian welfare is maximized by giving all of good 1 to agent 1 and all of good 2 to agent 2, leaving agent 3 with nothing. This is impossible to support with price curves, because agent 3 can always buy a nonzero amount of the goods she wants.

Example 2.1 shows that the answer is no, unfortunately. One may also wonder if Theorem 2.6.1 would generalize if we relax our constraint from $w_{ij} \in \{0, 1\}$ to $w_{ij} \in [0, 1]$. The answer is again no; this counterexample is more involved and is given by Theorem 2.9.1 in Section 2.9.

Information required by the social planner

In general, the price curves will depend on agents' preferences, and so the social planner needs to know agents' preferences in order to compute them. This is true of linear-pricing markets as well: the equilibrium prices depend on the utility functions of the agents. For our GDF characterization result, the price curves can have a very complex shape that depends intricately on the specific preferences, unlike linear prices. For this reason, we view the GDF characterization more as a conceptual contribution than an actual mechanism. In contrast, for our CES welfare bandwidth allocation result, the price curves have a very simple shape that is independent of the agents' utility functions (the price of buying x will be $q_j x^{1-\rho}$, where q_j is a Lagrange multiplier corresponding to good j). This structure suggests a simple decentralized primal-dual algorithm similar to the work of Kelly et al. [110], where on each step, every agent updates her (primal) allocation based on the current prices on the links she cares about, and each link updates its (dual) price based on the total flow through that link. We discuss this further in Section 2.6.2.

Additional results

We prove two additional results. First, we consider max-min welfare in Section 2.7, and show that as long as agents' weights are reasonably normalized, allocations with optimal max-min welfare can be supported. Second, in Section 2.8, we give a characterization theorem for weakly increasing price curves (Theorem 2.8.1). Just like strictly increasing price curves, we can compute weakly increasing price curves (or show that none exist) in polynomial time, using a linear program (Theorem 2.8.2).

2.3 Prior work

The study of markets has a long history in the economics literature [173, 169, 6, 24]. Recently, this topic has garnered significant attention in the computer science community as well (see [170] for an algorithmic exposition). The vast majority of the literature has focused on linear prices. Perhaps the most relevant classical result is the Second Welfare Theorem, which states that for any Pareto optimal allocation, there exists a (possibly unequal) redistribution of initial wealth which makes that allocation a (linear-pricing) market equilibrium. In a similar vein, [82] showed that for linear but fully personalized prices (i.e., we can independently assign different prices to different agents for the same good), one can support any Pareto optimal allocation.

The important question, then, is what price curve equilibria offer that these prior results do not. First, in many of societies, a centrally mandated redistribution of wealth is out of the question. Similarly, fully personalized prices mean that we lose any claim to fairness, since agents may be subjected to totally different prices for the same resource. In contrast, price curves do not require a redistribution of wealth, and furthermore are *anonymous*, meaning that all agents are subject to the same pricing scheme. These properties suggest that price curves are more practical, and indeed price curves do appear in practice (as noted previously, the water sector is a good example of this [175]). See Chapter 4 (in particular, Section 4.2.1) for more discussion.

We briefly mention an important property in mechanism design: *strategy-proofness*. A mechanism is strategy-proof if agents can never improve their utility by lying about their preferences. Unfortunately, even in simple settings, the only mechanism for resource allocation that can simultaneously guarantee strategy-proofness and Pareto optimality is *dictatorial*, meaning that one agent receives all of the resources [159]. This is clearly unacceptable, so we sacrifice strategy-proofness in favor of Pareto optimality. Specifically, we assume throughout the chapter that agents always truthfully report their preferences.

The remainder of the chapter is organized as follows. Section 2.4 formally defines the model. Section 2.5 presents our first main result: that for strictly increasing price curves, an allocation can be supported if and only if it is group-domination-free. Section 2.6 gives our second main result: that in the bandwidth allocation setting, every maximum CES welfare allocation can be supported by price curves. Section 2.7 shows that allocations with optimal max-min welfare can be supported by price curves in a wide range of settings. In Section 2.8, we generalize our characterization theorem from Section 2.5 to account for weakly increasing price curves. Section 2.9 provides counterexamples to various claims that one might have hoped to prove. We also discuss in that section why certain other classes of utilities – in particular, linear utilities – are much more difficult to analyze. Finally, Section 2.10 provides some proofs that are omitted from earlier sections.

2.4 Model

We use the basic terminology and notation defined in Chapter 1. We assume goods are divisible, meaning that x_{ij} can be any real number. Thus the main constraint a valid allocation must satisfy is the supply constraint:

$$\sum_{i \in N} x_{ij} \leq s_j \quad \forall j \in M$$

We assume that agents have Leontief utilities:

$$u_i(x_i) = \min_{j \in M : w_{ij} \neq 0} \frac{x_{ij}}{w_{ij}}$$

where w_{ij} denotes the (nonnegative) weight agent i has for good j. For brevity, we will usually write $u_i(x_i) = \min_{j \in M} \frac{x_{ij}}{w_{ij}}$ and leave the $w_{ij} \neq 0$ condition implicit. The same holds for other contexts where we are dividing by a value w_{ij} that may be zero. We assume that agents have Leontief utilities throughout the chapter, and we assume that each agent has nonzero weight for at least one good.

We focus on welfarist objectives in this chapter: specifically, CES welfare:

$$\Phi_\rho(\mathbf{x}) = \left(\sum_{i \in N} u_i(x_i)^\rho \right)^{1/\rho}$$

where ρ is a constant in $(-\infty, 0) \cup (0, 1]$.

In this chapter, we use the fixed-budget market model defined in Section 1.5. Unless otherwise stated, we will assume that all agents have the same budget, and normalize all budgets to 1 without loss of generality. For prices $\mathbf{p} = (p_1, \ldots, p_m) \in \mathbb{R}^m_{\geq 0}$, the *cost* of bundle x_i is $C_\mathbf{p}(x_i) = \sum_{j \in M} x_{ij} p_j$.[5] Bundle x_i is *affordable* for agent i if $C_\mathbf{p}(x_i) \leq 1$. Agent i's *demand set* is the set of her favorite affordable bundles, i.e.,

$$D_i(\mathbf{p}) = \operatorname*{arg\,max}_{x_i \in \mathbb{R}^m_{\geq 0}:\ C_\mathbf{p}(x_i) \leq 1} u_i(x_i).$$

If $p_j > 0$ for all $j \in M$, an agent with Leontief utility will always purchase in exact proportion to her weights: since agent i's utility is determined by $\min_{j \in M} \frac{x_{ij}}{w_{ij}}$, violating these proportions cannot increase her utility. Thus when discussing an arbitrary allocation $\mathbf{x}$, we assume that each bundle x_i is in proportion to agent i's weights: otherwise there is no hope of supporting such an allocation. For brevity, we leave this assumption implicit throughout the chapter, rather than always stating "for an arbitrary allocation $\mathbf{x}$ where each bundle is in proportion to agent i's weights".

The careful reader may note that we are glossing over a detail: if $p_j = 0$ for some good j, agent i can add more of good j to her bundle at no additional cost. This does not affect agent utilities at all, but is technically possible. In order to avoid handling this uninteresting and sometimes messy edge case, we assume throughout the chapter that for agents with Leontief utilities, demand sets and arbitrary allocations are always in exact proportion to agent weights.

Formally, a Fisher market equilibrium $(\mathbf{x}, \mathbf{p})$ is an allocation $\mathbf{x}$ and price vector $\mathbf{p} \in \mathbb{R}^m_{\geq 0}$ such that

1. Each agent receives a bundle in her demand set: $x_i \in D_i(\mathbf{p})$.
2. The market clears: for all $j \in M$, $\sum_{i \in N} x_{ij} \leq s_j$, and if $p_j > 0$, then $\sum_{i \in N} x_{ij} = s_j$.

For a wide class of agent utilities, including Leontief utilities, an equilibrium is guaranteed to exist [6][6]. Furthermore, the equilibrium allocations are the exactly the allocations which maximize

[5]Using the terminology from Section 1.5, $C_\mathbf{p}$ is the *pricing rule*. In this chapter, the pricing rule is fully determined with the prices $p_1, \ldots, p_m$ (or, later, by the price curves $f_1, \ldots, f_m$). For this reason, in this chapter we choose to use use the $C_\mathbf{p}$ and $C_\mathbf{f}$ notation instead of the more general $p(x_i)$.

[6]Specifically, an equilibrium is guaranteed to exist as long agent utilities are continuous, quasi-concave, and non-satiated. The full Arrow-Debreu model also allows for agents to enter to market with goods themselves and not only money; the necessary conditions on utilities are slightly more complex in that setting.

Nash welfare[7]. This is made explicit by the celebrated Eisenberg-Gale convex program [72, 73], and combinatorial approaches to computing market equilibria [61, 104].

2.4.1 Price curves

Our work considers an expanded model where instead of assigning a single price $p_j \in \mathbb{R}_{\geq 0}$ to each good, we assign each good a *price curve* $f_j : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. The function f_j expresses the cost of good j as a function of the quantity purchased. When we say "price curve", we mean a function f_j that is weakly increasing (buying more of a good cannot cost less), normalized ($f_j(0) = 0$), and continuous. Setting $f_j(x) = p_j \cdot x$ for all $j \in M$ and all $x \in \mathbb{R}_{\geq 0}$ yields the Fisher market setting.

Given a vector of price curves $\mathbf{f} = (f_1, \ldots, f_m)$, the cost of a bundle x_i is now $C_{\mathbf{f}}(x_i) = \sum_{j \in M} f_j(x_{ij})$. Although the functions f_j may not be linear, the cost of a bundle is still additive across goods. Each agent's demand set is defined identically to the Fisher market setting: $D_i(\mathbf{f}) = \operatorname*{arg\,max}_{x_i \in \mathbb{R}^m_{\geq 0}:\ C_{\mathbf{f}}(x_i) \leq 1} u_i(x_i)$.

The demand set is intuitively the same as in the Fisher market setting: each agent purchases exactly in proportion to her weights, and buys as much as she can without exceeding her budget. A *price curve equilibrium* $(\mathbf{x}, \mathbf{f})$ is an allocation $\mathbf{x}$ and vector of price curves $\mathbf{f}$ such that

1. Each agent receives a bundle in her demand set: $x_i \in D_i(\mathbf{f})$.
2. The demand does not exceed supply: $\sum_{i \in N} x_{ij} \leq s_j$ for all $j \in M$[8].

We say that price curves $\mathbf{f}$ *support* an allocation $\mathbf{x}$ if $(\mathbf{x}, \mathbf{f})$ is a price curve equilibrium. The first question we address in this chapter is: what allocations $\mathbf{x}$ can be supported by prices curves?

2.5 Group domination

Recall that we require price curves to be continuous and weakly increasing. We wish to theoretically characterize which allocations can be supported by price curves so that we can (1) apply this characterization in our subsequent proofs, and (2) construct an algorithm which can calculate price curves in polynomial time.

The true necessary and sufficient condition for an allocation to be supported by price curves – and an algorithm to compute them – is given in Section 2.8. However, this condition ("locked-agent-freeness") is somewhat unwieldy. Although weakly increasing price curves are sometimes necessary[9],

[7] The conditions for the correspondence between Fisher market equilibria and Nash welfare are slightly stricter than those for market equilibrium existence, but are still quite general. Sufficient criteria were given in [72] and generalized slightly by [104].

[8] For Fisher market equilibria, the second condition also stipulated that whenever $p_j > 0$, $\sum_{j \in M} x_{ij} = s_j$. Without this additional condition, cranking up all prices to infinity would result in trivial equilibria where all agents purchase almost nothing and so would certainly not maximize Nash welfare. Such trivial price curve equilibria do exist under our definition, but since we are not going to make any claims of the form "all price curve equilibria maximize a certain function", there is no issue with allowing these trivial equilibria to exist.

[9] Consider an instance with two agents and two goods, each with supply 1. Let the agents' weights be given by $w_{11} = w_{21} = w_{12} = 1$ and $w_{22} = 0$. Nash welfare is maximized by splitting good 1 evenly between the two agents, and allowing agent 1 to purchase an equal quantity of good 2. This only possible if the price of good 2 is zero: otherwise,

for now we will consider only *strictly* increasing price curves. The corresponding necessary and sufficient condition is the cleaner notion of group-domination-freeness.

2.5.1 Group domination

We have termed the necessary and sufficient condition for the existence of strictly increasing price curves "group-domination-freeness" (GDF). To gain intuition for this condition, consider an allocation $\mathbf{x}$ and agents i, k. We will say that agent i *dominates* agent k if $\forall j \; x_{ij} \geq x_{kj}$ and there exists j for which this inequality is strict. Observe that this would prevent the existence of strictly increasing price curves supporting allocation $\mathbf{x}$ – both agents must spend their full budget (otherwise they could buy more of every good, since price curves are continuous), but agent i buys strictly more than agent k. A similar scenario arises when considering any two weighted sets of agents $\mathbf{a}, \mathbf{b}$. We can represent these weighted groups as vectors with a non-negative weight[10] for each agent, where we require that $\mathbf{a}$ and $\mathbf{b}$ have the same total agent-weight. If for every possible quantity $\tau \in \mathbb{R}_{\geq 0}$ of any good j, considering only agents purchasing at least τ of good j, the weight of the agents in $\mathbf{a}$ is greater than or equal to the weight of agents in $\mathbf{b}$, then $\mathbf{b}$ can never be made to pay more than $\mathbf{a}$. Essentially, for each additional δ of any good, as many agents from $\mathbf{a}$ must purchase that δ as agents from $\mathbf{b}$, so no matter how we price these increments, $\mathbf{b}$ never pays more. If this difference in weights is strict for any (j, τ) pair, that implies some δ increment must cost 0 (because the total expenditure of $\mathbf{a}$ and $\mathbf{b}$ must be equal), violating the requirement that price curves be strictly increasing.

Another way to gain intuition for group domination is by analogy to stochastic dominance. Distribution $\mathbf{a}$ stochastically dominates distribution $\mathbf{b}$ if for every possible payoff $\tau \in \mathbb{R}_{\geq 0}$, the odds of getting at least τ from $\mathbf{a}$ are at least as high as the odds of getting at least τ from $\mathbf{b}$. One consequence of stochastic dominance is that *any* rational agent should prefer $\mathbf{a}$ to $\mathbf{b}$ – there are no trade-offs, $\mathbf{a}$ is simply better than (dominant over) $\mathbf{b}$. In fact, we can directly consider weighted groups of agents as probability distributions. The total weight of each group must be the same – without loss of generality, equal to 1. Consider sampling the allocations x_{ij} for any good j with probability equal to the weight of each agent. The probability distribution $\mathbf{a}$ stochastically dominating $\mathbf{b}$ is exactly equivalent to the weighted group $\mathbf{a}$ group-dominating $\mathbf{b}$. Thus not only does group domination create problems for pricing, it can arguably be considered *unfair*, as $\mathbf{a}$ is in some sense *objectively* better-off than $\mathbf{b}$[11].

The formal definition of this condition is below.

Definition 2.5.1 (**Group-domination-free (GDF)**). *Let* $\mathbf{a} = (a_1, \ldots, a_n)$ *and* $\mathbf{b} = (b_1, \ldots, b_n)$ *be vectors in* $\mathbb{R}^n_{\geq 0}$ *that assign a (possibly zero) weight to each agent, such that* $\sum_{i \in N} a_i = \sum_{i \in N} b_i$. *Then* $\mathbf{a}$ *group-dominates* $\mathbf{b}$ *in an allocation* $\mathbf{x}$ *(denoted* $\mathbf{a} \succ \mathbf{b}$*) if for all* $j \in M$ *and for any threshold*

agent 1 is paying more than agent 2. Recall that the Fisher market equilibrium prices are the dual variables of the convex program for maximizing Nash welfare: thus the price of good 2 being zero corresponds to the fact that the supply constraint for good 2 is not tight in this instance.

[10]Note that this is *not* the same weight as the w_{ij} representing an agent's weight for a good.

[11]See [79] for an introduction to stochastic dominance.

$\tau \in \mathbb{R}_{\geq 0}$,

$$\sum_{i \in N:\ x_{ij} \geq \tau} a_i \geq \sum_{i \in N:\ x_{ij} \geq \tau} b_i$$

and there exists a (j, τ) *pair where the inequality is strict.* $\mathbf{x}$ *is group-domination-free if there do not exist vectors* $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n_{\geq 0}$ *such that* $\mathbf{a} \succ \mathbf{b}$ *in* $\mathbf{x}$.[12]

We will also assume without loss of generality that for all $i \in N$, at least one of a_i and b_i is zero, i.e., these are non-overlapping weighted groups: were this not the case, we could define $\mathbf{a}'$ and $\mathbf{b}'$ by $a'_i = a_i - \min(a_i, b_i)$ and $b'_i = b_i - \min(a_i, b_i)$, and we would have $\mathbf{a}' \succ \mathbf{b}'$ if and only if $\mathbf{a} \succ \mathbf{b}$.

Theorem 2.5.1 will show that an allocation can be supported with strictly increasing price curves if and only if it is GDF.

2.5.2 Characterization of allocations supported by strictly increasing price curves

In order to relate the existence of price curves to GDF, first observe that, for agents with Leontief utilities, the conditions for a price curve equilibrium take on a relatively simple form. Recall that by assumption, the allocation to be considered doesn't violate supply, and each agent purchases goods in exact proportion to her weights w_{ij} (see Section 2.4). Then the condition that $x_i \in D_i(\mathbf{f})$ for all i can be captured by Lemma 2.5.1, whose proof appears in Section 2.10. Intuitively, agent i fills up her bundle in proportion to her weights until (a) she reaches her budget and (b) there exists a good where buying more would cost more.

Lemma 2.5.1. *Given price curves* $\mathbf{f}$, $x_i \in D_i(\mathbf{f})$ *if and only if both of the following hold: (a)* $C_{\mathbf{f}}(x_i) = 1$, *and (b) there exists* $j \in M$ *such that for all* $\varepsilon > 0$, $f_j(x_{ij} + \varepsilon w_{ij}) > f_j(x_{ij})$.

We are now almost ready to prove Theorem 2.5.1 relating the existence of price curves to GDF. However, the proof is rather intricate, so we begin by giving an intuitive overview thereof. Throughout, we will refer to the example allocation $\mathbf{x}$ shown in Figure 2.1a to make the argument concrete. (Note that the example allocation shown in the figure implicitly defines a corresponding Leontief utility function for each agent, up to scaling by a constant, since we assume each agent fills up her bundle in exact proportion to her weights w_{ij}.)

We will now use this example to illustrate three key observations regarding the existence of strictly increasing price curves supporting an allocation $\mathbf{x}$: (1) Only the points on the price curves corresponding to agent allocations matter. (2) Only the order of the agents along the price curve for each good, not their absolute allocations, matters. (3) The order of the agents can be captured in an *agent-order matrix* such that weighted column and row sums represent agent costs and group dominations, respectively.

[12] The "-free/-freeness" suffix may remind some readers of the popular fairness notion envy-freeness; this connection is intended. If one agent does envy another, this constitutes an instance of group domination in the allocation, so GDF implies envy-freeness. However, the reverse is not true: for an agent i to envy agent k, k must receive strictly more of every good i cares about; for group domination, the difference need only be strict on one good. All market equilibria are envy-free; GDF is a stronger notion corresponding exactly the the existence of a market equilibrium in this setting.

	good 1	good 2	good 3
agent 1	0.6	0	0.2
agent 2	0.3	0	0
agent 3	0.1	1	0.4
agent 4	0	0	0.4

(a) example allocation **x**

$$\begin{array}{ccc} \text{good 1} & \text{good 2} & \text{good 3} \end{array}$$
$$\left[\begin{array}{ccc:c:cc:c} 1 & 1 & 1 & 0 & 1 & 0 & -1 \\ 1 & 1 & 0 & 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & 1 & 1 & 1 & -1 \\ 0 & 0 & 0 & 0 & 1 & 1 & -1 \end{array}\right]$$

(b) **x** represented as an agent-order matrix A

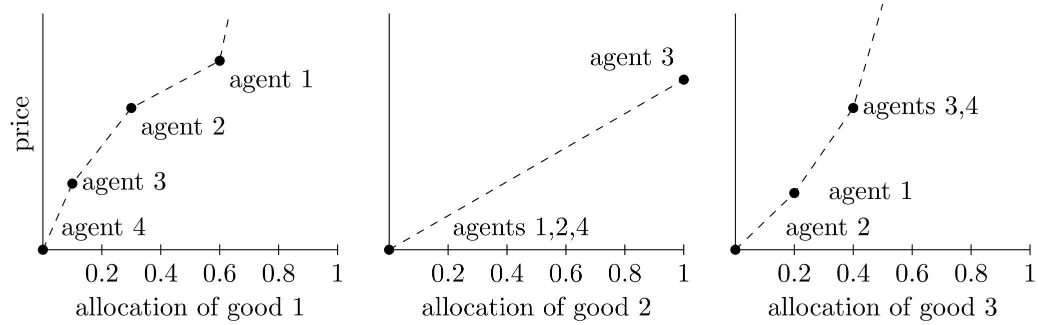

(c) example price curves for allocation **x**

Figure 2.1: An illustrative example allocation and the construction of the corresponding agent-order matrix.

First we address observation (1). Consider the possible price curves shown in Figure 2.1c. Given the price that each agent pays for each good, these are the only points that matter, in the sense that (a) each agent's total cost, which must equal 1, depends only on these points, and (b) an agent must be able to purchase more of a good if the next fixed point on that curve has the same price, and otherwise need not be able to do so, for instance if we make the price curves piece-wise linear as shown. Thus when considering whether price curves are possible, we need only consider the set of prices corresponding to agent allocations.

A similar argument addresses observation (2). As long as we fix the order of points along a price curve, we can change the allocations arbitrarily (assuming they still obey the supply and proportional-purchase assumptions) without changing the prices. Obviously, every agent will still incur a cost of 1, and it will not change whether an agent can purchase more of a good (whether the next point along the curve has the same price).

Finally, we come to the more complicated observation (3). We will first lay out how the agent-order matrix is constructed, then illustrate its connection to both prices and group domination. The matrix will have n rows, one for each agent, and a sub-block for each good, as shown in Figure 2.1b. Within a sub-block, each column will correspond to a non-zero agent allocation (i.e., the non-zero points shown in Figure 2.1c). The entry corresponding to agent i, good j, and allocation threshold $\tau \in \mathbb{R}_{\geq 0}$ will equal 1 if $x_{ij} \geq \tau$ and 0 otherwise. Essentially, this will indicate which agent pays the cost of the first, second, etc. section of each price curve. Additionally, we append a column of -1's to the end of the matrix. To see the connection to prices, consider a vector $\mathbf{y}$ such that $A\mathbf{y} = \mathbf{0}, \mathbf{y} \neq \mathbf{0}$. For instance, Figure 2.2a exhibits such a vector $\mathbf{y}$ for the matrix A shown in Figure 2.1b. $\mathbf{y}$ will represent prices, so we require all the entries to be non-negative, denoted $\mathbf{y} \geq 0$; for strictly increasing

$$\left[\begin{array}{ccc:c:cc:c} 1 & 1 & 1 & 0 & 1 & 0 & -1 \\ 1 & 1 & 0 & 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & 1 & 1 & 1 & -1 \\ 0 & 0 & 0 & 0 & 1 & 1 & -1 \end{array}\right] \times \left[\begin{array}{c} 0 \\ 1 \\ 0 \\ \hdashline 0 \\ \hdashline 0 \\ 1 \\ \hdashline 1 \end{array}\right] = \mathbf{0}$$

(a) column sum of A implies price curves (in this example, weakly increasing)

$$\left[\begin{array}{cccc} 1 & 1 & 1 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ \hdashline 0 & 0 & 1 & 0 \\ \hdashline 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ \hdashline -1 & -1 & -1 & -1 \end{array}\right] \times \left[\begin{array}{c} 1 \\ -1 \\ 0 \\ 0 \end{array}\right] = \left[\begin{array}{c} 0 \\ 0 \\ 1 \\ \hdashline 0 \\ \hdashline 1 \\ 0 \\ \hdashline 0 \end{array}\right]$$

(b) row sum of A implies group domination (precludes strictly increasing price curves)

Figure 2.2: Example row and column sums of the agent-order matrix from Figure 2.1b.

price curves, we require $\mathbf{y}$ to be strongly positive[13], denoted $\mathbf{y} \gg 0$. Specifically, within each block (corresponding to a good j), the first entry represents the cost of increasing from 0 of good j to the first non-zero point on the price curve, the second entry represents the cost of increasing from the first point to the second point, and so on. The last entry in $\mathbf{y}$, which we can assume equals 1 without loss of generality, represents the total cost expended by each agent. Thus $A\mathbf{y} = \mathbf{0}$ ensures that each agent spends exactly 1 unit of money. (Ensuring that condition (b) of Lemma 2.5.1 be met is slightly more complicated. However, for strictly increasing price curves, it is trivially satisfied.)

Thus we can see that the column sums of the agent-order matrix correspond to agent expenditures, where the weight of each column in the sum corresponds to a section of the price curve. *Row* sums, however, correspond to group domination. To see the connection, consider a vector $\mathbf{z}$ such that $A^T\mathbf{z}$ is strictly positive[14], denoted $A^T\mathbf{z} > \mathbf{0}$. For instance, Figure 2.2b exhibits such a vector $\mathbf{z}$ for the matrix A shown in Figure 2.1b. In a given $\mathbf{z}$, the positive entries correspond to the weighted agents in a dominating group $\mathbf{a}$, while the (absolute value of the) negative entries are the weighted agents in group $\mathbf{b}$. Since the last entry of $A^T\mathbf{z}$ must be nonnegative, the total weight of $\mathbf{b}$ is at least as large as that of $\mathbf{a}$. And since $A^T\mathbf{z} > \mathbf{0}$, all the entries are non-negative and at least one other entry must be positive. This means that at every point on a price curve (any j, τ), the weight of group $\mathbf{a}$ purchasing at least τ of good j is at least as much as the weight of group $\mathbf{b}$ purchasing τ, and for some (j, τ) this is strict. Clearly this is equivalent to $\mathbf{a} \succ \mathbf{b}$.

Having constructed the agent-order matrix and related its column and row sums to prices and group domination, respectively, the final step applies a previously-known duality result equivalent to Farkas' Lemma [141], which establishes that valid prices (column sums) exist if and only if group domination (row sums) do *not*. Specifically, we make use of the following result originally due to Stiemke to prove Theorem 2.5.1.

Lemma 2.5.2 (1.6.4 in [167])**.** *For a commutative, ordered field $\mathbb{F}$, A a matrix over $\mathbb{F}$, the following are equivalent: (1) $A\mathbf{y} = \mathbf{0}, \mathbf{y} \gg \mathbf{0}$ has no solution. (2) $A^T\mathbf{z} > \mathbf{0}$ has a solution.*

Theorem 2.5.1. *Let $\mathbf{x}$ be any allocation that obeys the supply constraints and gives at least one agent a nonempty bundle. Then $\mathbf{x}$ be can supported by strictly increasing price curves if and only if*

[13]Recall that a strongly positive vector has every entry greater than 0.

[14]Recall that a strictly positive vector has entries in $\mathbb{R}_{\geq 0}$ with at least one entry non-zero.

$\mathbf{x}$ *is GDF.*

Proof. Recall that an allocation $\mathbf{x}$ can be supported if there exist price curves $\mathbf{f}$ such that $x_i \in D_i(\mathbf{f})\ \forall i \in N$, and $\sum_{i\in N} x_{ij} \leq 1\ \forall j \in M$ (i.e., $\mathbf{x}$ obeys the supply constraints). The latter condition is satisfied by assumption, and by Lemma 2.5.1, for Leontief utilities and strictly increasing price curves, the former condition holds if and only if the cost $C_{\mathbf{f}}(x_i) = 1\ \forall i$.

Let $X_j = \{x_{ij} \mid i \in N\} \setminus \{0\}$ be the set of distinct, non-zero amounts of good j allocated to some agent under $\mathbf{x}$. Label the elements of X_j as $\tau_j^1, \tau_j^2, \ldots, \tau_j^{|X_j|}$ such that $\tau_j^1 < \tau_j^2 < \cdots < \tau_j^{|X_j|}$. Since $f_j(0) = 0$, $f_j(x \notin X_j)$ in some sense doesn't matter – we only require that these "in-between" areas of the price curve don't violate continuity and are strictly increasing. Thus there exist strictly increasing price curves $\mathbf{f}$ supporting $\mathbf{x}$ if and only if there exist functions $f_j' : X_j \to \mathbb{R}_{\geq 0}$ such that $0 < f_j'(\tau_j^1) < f_j'(\tau_j^2) < \ldots < f_j'(\tau_j^{|X_j|})\ \forall j$ and $C_{\mathbf{f}}(x_i) = \sum_j f_j'(x_{ij}) = 1\ \forall i$.

Now we are ready to set up the agent-order matrix $A \in \mathbb{Q}^{n\times(\sum_j |X_j|+1)}$ to which we will apply Lemma 2.5.2. Since each column will represent an allocation point for a specific good (corresponding to its sub-block), we will write the column indices as $\sum_{\ell<j} |X_\ell| + q$, where j indicates the sub-block and $1 \leq q \leq |X_j|$ is the index within that sub-block.

$$A\left[i,\ \sum_{\ell<j} |X_\ell| + q\right] = \begin{cases} -1 & \text{if } j = m+1, q = 1 \text{ (last column)} \\ 0 & \text{if } x_{ij} < \tau_j^q \\ 1 & \text{otherwise} \end{cases}$$

Thus each row of A represents an agent, and each column (except the last) represents one point of the functions $\mathbf{f'}$. Since $\mathbf{x}$ gives at least one agent a nonempty bundle by assumption, A has at least two columns (one allocation point and the column of -1's). We know by Lemma 2.5.2 that $\exists \mathbf{y} \gg \mathbf{0}$ such that $A\mathbf{y} = 0$ if and only if there does not exist a $\mathbf{z}$ such that $A^T\mathbf{z} > \mathbf{0}$. To complete the proof, we will show that the former condition is equivalent to the existence of strictly increasing price curves supporting $\mathbf{x}$, and the latter is equivalent to a group domination.

If $\exists \mathbf{y} \gg \mathbf{0}$ such that $A\mathbf{y} = \mathbf{0}$, we may assume without loss of generality that the last entry in $\mathbf{y}$ is 1. Furthermore, define $f_j'(\tau_j^q) - f_j'(\tau_j^{q-1}) = y_{\sum_{\ell<j} |X_\ell|+q}$ (for convenience, define $f_j'(\tau_j^0) = f_j'(0) = 0$). Clearly $\mathbf{y} \gg \mathbf{0}$ is equivalent to the requirement that $0 < f_j'(\tau_j^1) < \ldots < f_j'(\tau_j^{|X_j|})\ \forall j$. Additionally,

$$C_{\mathbf{f}}(x_i) = \sum_j f_j(x_{ij}) = \sum_j f_j'(x_{ij}) = \sum_j \sum_{q: x_{ij} \geq \tau_j^q} y_{\sum_{\ell<j} |X_\ell|+q} = A_i\mathbf{y} + 1$$

Thus $A\mathbf{y} = \mathbf{0}$ is equivalent to the requirement that $C_{\mathbf{f}}(x_i) = 1\ \forall i$.

Finally, consider $\mathbf{z}$ such that $A^T\mathbf{z} > \mathbf{0}$. This is equivalent to a group domination $\mathbf{a} \succ \mathbf{b}$, where $a_i = z_i$ if $z_i > 0$, $b_i = -z_i$ if $z_i < 0$, and all other entries are 0. Consider the product of the last column of A with $\mathbf{z}$, which equals $-\sum_i z_i \geq 0$. Without loss of generality, we can assume $\sum_i z_i = 0$, and thus $\sum_i a_i = \sum_i b_i$. If this is not true, then $\mathbf{b}$ would have greater weight than $\mathbf{a}$, and decreasing any weight in $\mathbf{b}$ can only increase coordinates of $A^T\mathbf{z}$ or equivalently widen the gap between $\mathbf{a}$ and $\mathbf{b}$ in terms of group domination. Now observe that for any good j and $\tau \in \mathbb{R}_{\geq 0}$, $\sum_{i\in N: x_{ij} \geq \tau}(a_i - b_i)$

is equal to the dot product of column $\sum_{\ell<j} |X_\ell| + q$ of A by $\mathbf{z}$, where q is the largest value such that $\tau_j^q \leq \tau$. This holds because $A\left[i, \sum_{\ell<j} |X_\ell| + q\right]$ is an indicator variable for $x_{ij} \geq \tau_j^q$, and by construction no agent can have an allocation amount between τ_j^q and τ_j^{q+1}. Therefore $A^T\mathbf{z} > \mathbf{0}$ is equivalent to the requirement that $\sum_{i\in N: x_{ij}\geq\tau}(a_i - b_i) \geq 0$ for all (j, τ) and that for some (j, τ) this inequality is strict, i.e., $A^T\mathbf{z} > \mathbf{0}$ is equivalent to $\mathbf{a} \succ \mathbf{b}$. ☐

Remark. Since the matrix A constructed in the proof of Theorem 2.5.1 is over the rationals, we can also assume that the $\mathbf{y}$ or $\mathbf{z}$ obtained via Lemma 2.5.2 are over the rationals. In particular, we can scale $\mathbf{z}$ to obtain $\mathbf{z}' \in \mathbb{Z}^n$ with $A^T z' > 0$. Equivalently, this means that if $\mathbf{a} \succ \mathbf{b}$, we can assume without loss of generality that $a_i, b_i \in \mathbb{Z}$.

This characterization, in addition to allowing us to prove some of our subsequent results, implies that we can compute price curves (or show that they do not exist) for a particular instance in polynomial time. This is exhibited by the following linear program.

Theorem 2.5.2. *Given a set of agents N, goods M, and an allocation $\mathbf{x} \in \mathbb{R}_{\geq 0}^{n\times m}$, let A be the corresponding agent-order matrix. In the following linear program, the optimal objective value is strictly positive if and only if there exist strictly increasing price curves supporting $\mathbf{x}$, in which case $\mathbf{y}$ defines such curves.*

$$\begin{aligned} \max_{\mathbf{y},\eta} \;\; & \eta \\ s.t. \;\; & A\mathbf{y} = \mathbf{0} \\ & y_k \geq \eta \quad \forall k \\ & y_{-1} = 1 \end{aligned}$$

Proof. As per the proof of Theorem 2.5.1, there exist strictly increasing price curves supporting $\mathbf{x}$ if and only if there is a solution to the system $A\mathbf{y} = \mathbf{0}, \mathbf{y} \gg 0$. To turn this into a valid linear program, instead of the strict inequality $y_k > 0$ for each coordinate of $\mathbf{y}$, we write $y_k \geq \eta$ and attempt to maximize η. Furthermore, we restrict the final entry of $\mathbf{y}$ as $\mathbf{y}_{-1} = 1$, since otherwise $\mathbf{y}$ can be scaled arbitrarily. If there is a solution with $\eta > 0$, then $\mathbf{y}$ corresponds to price curves as before, with each entry representing the difference in price between adjacent allocation amounts. These points simply need to be connected, e.g., in a piecewise linear fashion, to constitute valid price curves. ☐

One may wonder if Theorem 2.5.1 generalizes to other classes of utility functions. Unfortunately, the answer in general is no. Example 2.3 gives an instance with linear utilities that is GDF, but cannot be supported by price curves.

In Section 2.7, we will show how the group-domination-freeness concept can be useful for proving that allocations of interest can be supported by price curves: specifically, allocations with optimal (or near optimal) max-min welfare. But first, a word about unequal budgets.

2.5.3 Unequal budgets

It turns out that the characterization theorem of the previous section easily generalizes to agents with unequal budgets. Since price curves are strictly increasing, the only additional requirement for an allocation $\mathbf{x}$ to be supported is that each agent spends her entire budget B_i. In the agent-order matrix, the last column of -1's corresponded to each agent's expenditure, so we simply need to replace -1 with $-B_i$ for each row i.

Following Lemma 2.5.2 with the modified agent-order matrix, the if-and-only-if characterization becomes "budget-weighted group-domination-freeness". A budget-weighted group domination still requires that for all (j, τ),

$$\sum_{i \in N : x_{ij} \geq \tau} a_i \geq \sum_{i \in N : x_{ij} \geq \tau} b_i$$

and that there exists j, τ where the inequality is strict. The only difference is that instead of requiring both groups to have the same total weight, that weight is now scaled by each agent's budget. That is, $\sum_{i \in N} a_i B_i = \sum_{i \in N} b_i B_i$. Note that when $B_i = 1$ for all i, this recovers the definition of group domination.

2.6 CES welfare

In this section, we consider CES welfare functions:

$$\Phi_{CES}(\mathbf{x}) = \Big(\sum_{i \in N} u_i(x_i)^\rho \Big)^{1/\rho}$$

This section contains our second main result: that in the bandwidth setting (i.e., agents have Leontief utilities where $w_{ij} \in \{0, 1\}$ for all $i \in N$, $j \in M$), for any $\rho \in (-\infty, 0) \cup (0, 1)$, any maximum CES welfare allocation can be supported by price curves (Theorem 2.6.1). We present this result this in Section 2.6.1. Next, we discuss why we are optimistic about the possibility of a simple decentralized primal-dual algorithm for computing these price curves, similar to the work of Kelly et al. [110] (Section 2.6.2). We also give a converse of sorts to Theorem 2.6.1 (Section 2.6.3), and briefly discuss the case of unequal budgets (Section 2.6.4). Throughout this section, we let $R_i = \{j \in M : w_{ij} = 1\}$ for brevity.

2.6.1 Main CES welfare result

We now state and prove Theorem 2.6.1. Our proof uses the dual of the convex program for maximizing CES welfare to construct explicit price curves that support a maximum CES welfare allocation. The price curves take the very simple form of $f_j(x) = q_j x^{1-\rho}$ for constants $q_1, \dots, q_m$ that are derived from the dual.

Theorem 2.6.1. *If $w_{ij} \in \{0, 1\}$ for all $i \in N$ and $j \in M$, then for any $\rho \in (-\infty, 0) \cup (0, 1)$, any maximum CES welfare allocation can be supported by price curves of the form $f_j(x) = q_j x^{1-\rho}$ for each $j \in M$.*

Proof. The proof proceeds in three steps.

Step 1: Setting up the convex program. We begin by writing the following program to maximize CES welfare:

$$\max_{\substack{\mathbf{x}\in\mathbb{R}^{n\times m}_{\geq 0},\\ \mathbf{u}=(u_1\dots u_n)\in\mathbb{R}^n_{\geq 0}}} \quad \Big(\sum_{i\in N} u_i^\rho\Big)^{1/\rho}$$

$$s.t.\quad u_i \leq x_{ij} \qquad \forall i\in N, j\in R_i$$

$$\sum_{i\in N} x_{ij} \leq s_j \qquad \forall j \in M$$

We could also have written the first constraint as $u_i w_{ij} \leq x_{ij}$, but since $w_{ij} \in \{0,1\}$, the above formulation is equivalent. Also, the objective $\Big(\sum_{i\in N} u_i^\rho\Big)^{1/\rho}$ is concave for any $\rho \in (-\infty, 0)\cup(0,1)$, so the resulting program is convex.

We can remove the exponent of $1/\rho$ from the objective without affecting the optimal point: the optimal value may be affected, but the optimal solution (i.e., the arg max) will not. When ρ is negative, this changes the program to a minimization program, but this can be handled by adding a factor of $1/\rho$ to the objective.[15] Thus consider a new convex program with objective function $\max\limits_{\mathbf{x}\in\mathbb{R}^{n\times m}_{\geq 0},\mathbf{u}\in\mathbb{R}^n_{\geq 0}} \frac{1}{\rho}\sum_{i\in N} u_i^\rho$, and the same constraints.

Next, we write the Lagrangian of the new program. Let λ_{ij} be the Lagrange multiplier associated with the constraint $u_i \leq x_{ij}$ and let q_j be the Lagrange multiplier associated with the constraint $\sum_{i\in N} x_{ij} \leq s_j$. We will use $\boldsymbol{\lambda}$ and $\mathbf{q}$ to denote the vectors of all such Lagrange multipliers. Then the Lagrangian is given by

$$L(\mathbf{x},\mathbf{u},\boldsymbol{\lambda},\mathbf{q}) = \frac{1}{\rho}\sum_{i\in N} u_i^\rho - \sum_{i\in N}\sum_{j\in R_i} \lambda_{ij}(u_i - x_{ij}) - \sum_{j\in M} q_j\Big(\sum_{i\in N} x_{ij} - s_j\Big)$$

Consider any maximum CES welfare allocation: this corresponds to a point $(\mathbf{x}^*, \mathbf{u}^*)$ which is optimal for the primal. We have strong duality by Slater's condition, so there must exist $\boldsymbol{\lambda}^*$ and $\mathbf{q}^*$ such that $(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*)$ is optimal for L.

Step 2: Using the KKT conditions. The KKT conditions imply that the gradient of L evaluated at $(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*)$ must be zero for every variable with a positive value. Specifically, for each variable y, either $\frac{\partial L}{\partial y} = 0$, or $y = 0$ and $\frac{\partial L}{\partial y} \leq 0$.

First, we have $\frac{\partial L}{\partial u_i}(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*) = {u_i^*}^{\rho-1} - \sum_{j\in R_i} \lambda_{ij}^* = 0$ for every $i \in N$ with $u_i^* > 0$. Suppose $u_i^* = 0$ for some $i \in N$: since $\rho - 1 < 0$, we would have ${u_i^*}^{\rho-1} = \infty$, which contradicts $\frac{\partial L}{\partial y} \leq 0$. Thus $u_i^* > 0$, so $u_i^* = (\sum_{j\in R_i} \lambda_{ij}^*)^{\frac{1}{\rho-1}}$ for all $i \in N$.

Similarly, $\frac{\partial L}{\partial x_{ij}}(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*) = \lambda_{ij}^* - q_j^* = 0$ for every $i\in N$ and $j \in R_i$ with $x_{ij} > 0$. Since $u_i^* > 0$ for all $i \in N$, we must have $x_{ij}^* > 0$ for all $j \in R_i$. Therefore $\lambda_{ij}^* = q_j^*$ for all $i \in N, j \in R_i$, so $u_i^* = \big(\sum_{j\in R_i} q_j^*\big)^{\frac{1}{\rho-1}}$. It will also be helpful to consider $\frac{\partial L}{\partial x_{ij}}(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*)$ for $j \notin R_i$: in this case, we have $\frac{\partial L}{\partial x_{ij}}(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*) = q_j = 0$ whenever $x_{ij}^* > 0$.

[15] We add a factor of $1/\rho$ instead of ρ because this will slightly simplify the analysis.

Next, the KKT conditions also imply that $(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*)$ satisfy *complementary slackness*, meaning that the Lagrange multiplier of any non-tight constraint is equal to 0. We are specifically interested in the constraint $x^*_{ij} \leq u^*_i$ for $j \in R_i$: either $\lambda^*_{ij} = q^*_j = 0$, or

$$x^*_{ij} = u^*_i = \Big(\sum_{j \in R_i} q^*_j\Big)^{\frac{1}{\rho-1}}$$

Step 3: Constructing the price curves. We now use the Lagrange multipliers $\mathbf{q}^*$ to construct explicit price curves. We define $f_j(x)$ by $f_j(x) = q^*_j x^{1-\rho}$. Since $\rho \in (-\infty, 0) \cup (0, 1)$, we have $1-\rho > 0$, so these price curves are in fact increasing. We claim that $(\mathbf{x}^*, \mathbf{f})$ is a price curve equilibrium. To see this, we explicit compute the cost of agent i's bundle x^*_i:

$$C_{\mathbf{f}}(x^*_i) = \sum_{\ell \in M} q^*_\ell (x^*_{i\ell})^{1-\rho} = \sum_{\ell : q_\ell, x_{i\ell} \neq 0} q^*_\ell (x^*_{i\ell})^{1-\rho} = \sum_{\ell : q_\ell, x_{i\ell} \neq 0} q^*_\ell \Big(\sum_{j \in R_i} q^*_j\Big)^{\frac{1-\rho}{\rho-1}} = \sum_{\ell : q_\ell, x_{i\ell} \neq 0} \frac{q^*_\ell}{\sum\limits_{j \in R_i} q^*_j}$$

To show that $C_{\mathbf{f}}(x^*_i) = 1$, we just need to show that $\sum_{j : q_j, x_{ij} \neq 0} q^*_j = \sum_{j \in R_i} q^*_j$. Clearly $\sum_{j : q_j, x_{ij} \neq 0} q^*_j = \sum_{j : x_{ij} \neq 0} q^*_j$. Since $u^*_i > 0$, we have $x_{ij} \neq 0$ for each $j \in R_i$, so $\sum_{j : x_{ij} \neq 0} q^*_j \geq \sum_{j \in R_i} q^*_j$. To show that the reverse inequality holds, it suffices to show that whenever $j \notin R_i$ and $x_{ij} \neq 0$, $q_j = 0$. This is exactly one of the things we showed via the KKT conditions in Step 2.

Thus we have shown that $C_{\mathbf{f}}(x^*_i) = 1$, so x^*_i is affordable to agent i for all $i \in N$. Furthermore, since $u^*_i = \Big(\sum_{j \in R_i} q^*_j\Big)^{\frac{1}{\rho-1}}$ is finite, there must exist $j \in R_i$ with $q^*_j > 0$. Thus there is at least one good $j \in R_i$ such that buying more would cost more money, so by Lemma 2.5.1, x^*_i is in agent i's demand set. We also know that $\sum_{j \in M} x^*_{ij} \leq 1$, since $\mathbf{x}^*$ is a feasible solution to the primal. Therefore $(\mathbf{x}^*, \mathbf{f})$ is a price curve equilibrium. □

To compute the price curves, we only need to know $q_1, \ldots, q_m$. Since these are the optimal Lagrange multipliers of our convex program for computing CES welfare, and that program can be solved in polynomial time, we can compute the price curves in polynomial time.

The structure of the price curve themselves ($f_j(x) = q^*_j x^{1-\rho}$) is also interesting when we consider the interpretation of the parameter ρ: the smaller ρ is, the more we care about agents with small utility. Recall that taking of $\rho \to -\infty$ yields max-min welfare, where we only care about the minimum utility. When $\rho = 1$, we have utilitarian welfare, where we only care about overall efficiency. This roughly corresponds to caring more about agents with higher utility. The limit as $\rho \to 0$ corresponds to Nash welfare, which is a mix of caring about both agents with low utility and those with high utility.

We know that maximum Nash welfare allocations are supported by linear price curves, i.e., those with constant marginal prices. When $\rho < 0$, these marginal prices are increasing, making it easier for agents who are buying less of each good. Since $w_{ij} \in \{0, 1\}$, $u_i(x_i) = x_{ij}$ whenever $w_{ij} \neq 0$, so the agents who are buying less are also the ones with lower utility. Thus price curves of this form for $\rho < 0$ are benefiting the agents with low utility. Furthermore, the smaller ρ is, the faster marginal prices grow, which corresponds to favoring agents with low utility even more. On the other hand, when $\rho > 0$, these marginal prices are decreasing. This favors agents with higher utility, which is consistent with the interpretation of the CES welfare function with $\rho > 0$.

2.6.2 Decentralized primal-dual updates

In this section, we discuss how the simple structure of these price curves suggests a natural decentralized primal-dual algorithm for computing said price curves. As established by Theorem 2.6.1, the price of buying x of good j will be $q_j x^{1-\rho}$. Specifically, for each $j \in M$, q_j be the Lagrange multiplier for the convex program for maximizing CES welfare with respect to that specific value of ρ[16].

For bandwidth allocation, each agent's demand set depends only on the prices curves for goods $j \in R_i$, i.e., the goods she cares about. For price curves of the form $f_j(x) = q_j x^{1-\rho}$, this means that each agent's demand given price curves $\mathbf{f}$ depends only on the dual prices q_j for goods $j \in R_i$. This means that given price curves $\mathbf{f}$, agents can update their demands in a decentralized fashion. Furthermore, the price set by each link should depend only on the flow through that link, i.e., $\{x_{ij} : i \in N, j \in R_i\}$. This means that given agent demands, each link can update its dual price q_j in a decentralized way (typically by raising the price if the demand is less than the supply, and increasing the price if the demand exceeds the supply). This suggests a simple decentralized primal-dual algorithm, where on each step, each agent updates her primal allocation x_i in response to the dual prices q_j for $j \in R_i$, and each link updates its dual price in response to primal allocations x_i. This is similar to the work of Kelly et al. [110].

This type of algorithm is also called a *tâtonnement*. One recent approach to tâtonnement makes use of the fact that the equilibrium prices are the Lagrange multipliers in the convex program to maximize Nash welfare, and gives a tâtonnement process that is akin to gradient descent on the dual program [48]. This approach also seems promising for our setting, since $q_1 \dots q_m$ are exactly the Lagrange multipliers in the convex program for maximizing CES welfare. We leave this as an open question.

2.6.3 A converse to Theorem 2.6.1

In this section, we give a converse of sorts to Theorem 2.6.1: if an allocation $\mathbf{x}$ can be supported by price curves $\mathbf{f}$ of the form $f_j(x) = q_j^* x^{1-\rho}$, and the supply is exhausted for any good with nonzero price, then $\mathbf{x}$ must be a maximum CES welfare allocation. The requirement that the supply be exhausted for any good with nonzero price (i.e., $\sum_{i \in N} x_{ij} = s_j$ whenever $q_j \neq 0$) is analogous to the second condition in definition of Fisher market (i.e., standard linear pricing) equilibrium given in Section 2.4.

The proof of Theorem 2.6.1 essentially hinges on the fact that when strong duality holds for a convex program, the KKT conditions are sufficient for optimality. This is analogous to the proof of Theorem 2.6.1, which is based on the fact that the KKT conditions are necessary for optimality. The formal proof appears in Section 2.10.

[16] In general, different values of ρ will lead to different optimal allocations and Lagrange multipliers.

Theorem 2.6.1. *Suppose* $(\mathbf{x}^*, \mathbf{f})$ *is a price curve equilibrium where for all* $j \in M$*,* $f_j(x) = q_j^* x^{1-\rho}$ *for* $\rho \in (-\infty, 1)$ *and nonnegative constants* $q_1^* \dots q_m^*$*. If* $\sum_{i \in N} x_{ij}^* = s_j$ *whenever* $q_j \neq 0$*, then* $\mathbf{x}^*$ *is a maximum CES welfare allocation.*

2.6.4 Unequal budgets

Finally, we address the setting where agents may have different amounts of money to spend. Let B_i be agent i's budget. If we instead consider the *budget-weighted CES welfare* $\Phi_{CES}(\mathbf{x}) = \big(\sum_{i \in N} B_i u_i(x_i)^\rho\big)^{1/\rho}$, then the proof of Theorem 2.6.1 extends directly. Duality tells us that agent i's utility must be $u_i(x_i) = \left(\frac{1}{B_i} \sum_{j \in R_i} q_j^*\right)^{\frac{1}{\rho - 1}}$. By using the same price curve form of $f_j(x) = q_j^* x^{1-\rho}$, we get $C_{\mathbf{f}}(x_i) = \sum_{\ell \in R_i} \frac{B_i q_\ell^*}{\sum\limits_{j \in R_i} q_j^*} = B_i$, so agent i is indeed spending exactly her budget. This can be used to show that any allocation with maximum budget-weighted CES welfare can be supported by price curves.

A social planner may prefer to give the same weight to each agent's utility, even if the budgets are not the same. Unfortunately, allocations with optimal unweighted CES welfare cannot be supported (at least not exactly) when agents have different budgets. To see this, consider two agents with different budgets and a single good: whichever agent has more money must receive a larger portion of the good. But assuming the agents have the same weight for that good (which holds in the bandwidth allocation setting or when weights are normalized somehow), the unweighted CES welfare optimum would give each agent the same amount. This is analogous to the Fisher market setting: the Fisher market equilibria for unequal budgets are exactly the allocations which maximize the budget-weighted Nash welfare.

2.7 Max-min welfare

In this section, we show that under mild assumptions, price curves can support allocations with either optimal max-min welfare, or arbitrarily close to optimal max-min welfare. As before, we assume that agents have Leontief utility functions. Also, we refer to an allocation with optimal max-min welfare as a max-min allocation.

The first thing we observe is that when agent weights are unconstrained in magnitude, there is no hope to support any approximation of max-min welfare. Consider a single good and two agents with weights w_{11} and w_{21} on that good. In this case, each agent i's utility is just x_{i1}/w_{i1}, so the max-min welfare of an allocation $\mathbf{x}$ is $\min(\frac{x_{11}}{w_{11}}, \frac{x_{21}}{w_{21}})$. Now imagine that w_{11} is much larger than w_{21}: agent 1 needs significantly more of the good to achieve the same utility as agent 2. Then any max-min allocation (or even any decent approximation) must give more of the good to agent 1 than agent 2. But since agents have the same budgets, any price curve equilibrium must result in each agent receiving half of the supply of good 1, which is a contradiction.

Thus in order to have any hope of even approximately supporting a max-min allocation, the agent weights must be normalized in some way. Theorem 2.7.1 states that under a quite general normalization assumption, we can support a max-min allocation.

Theorem 2.7.1. *Suppose there exist strictly increasing functions $g_1, \ldots, g_m$ such that for all $i \in N$, $\sum_{j \in M} g_j(w_{ij}) = 1$. Then there exists a max-min allocation that can be supported by price curves.*

Proof. Since the max-min welfare of an allocation is determined by the minimum agent utility, the max-min welfare cannot be improved by making any agent's utility higher than any other. Similarly, since each agent's utility is determined by $\min_{j \in M} x_{ij}/w_{ij}$, the max-min welfare cannot be improved by allocating goods to an agent outside of her desired proportions. Thus there exists a max-min allocation $\mathbf{x}$ where all agents have the same utility u, and where $x_{ij} = u \cdot w_{ij}$ for all $i \in N$ and $j \in M$.

Since GDF is invariant to scaling by constants, this implies that $\mathbf{x}$ is GDF if and only if *the weight vectors themselves* are GDF. That is, $\mathbf{x}$ is GDF if and only if the allocation $\mathbf{x}'$ defined by $x'_{ij} = w_{ij}$ is GDF. One realizes that the assumption of $\sum_{j \in M} g_j(w_{ij}) = 1$ for all $i \in N$ is *literally assuming that there exist (strictly increasing) price curves* that support the allocation $\mathbf{x}'$. Thus $\mathbf{x}'$ is GDF by Theorem 2.5.1, so $\mathbf{x}$ is GDF, which completes the proof. □

One natural corollary of Theorem 2.7.1 is the following:

Corollary 2.7.1.1. *Suppose there exists some $q \geq 1$ so that $\sum_{j \in M} w_{ij}^q = 1$ for all $i \in N$. Then there exists a max-min allocation that can be supported by price curves.*

Theorem 2.7.1 has an interesting conceptual implication. We can think of price curves themselves as a sort of "norm" on the allocation, and any allocation for which there is a "norm" which assigns the same value to each agent's bundle is reasonable enough that it can be supported by price curves. The previous statement can be rephrased as "an allocation can be supported by price curves if and only if there exist price curves which assign the same cost to each agent's bundle", and so is functionally a tautology. Since there exists a max-min allocation which is a constant scaling of the agent weights, this near-tautology carries over.

One final observation is that there are some interesting norms, such as the L_∞ norm, which cannot be written as the sum of increasing functions. In fact, there are cases where no max-min allocation can be supported when agent weights have the same L_∞ norm.[17] Furthermore, the following counterexample falls under the even simpler bandwidth allocation setting: $w_{ij} \in \{0, 1\}$ for all i, j.

Theorem 2.7.2. *There exist instances where $w_{ij} \in \{0, 1\}$ for all $i \in N$ and $j \in M$, but no max-min allocation can be supported.*

Proof. Consider an instance with three agents and two goods, each with supply 1. Let the agent weights be given by the following table:

	agent 1	agent 2	agent 3
good 1	1	0	1
good 2	0	1	1

[17]The L_∞ norm is defined as $\max_{j \in M} w_{ij}$.

The unique max-min allocation is $x_{11} = x_{22} = x_{31} = x_{32} = \frac{1}{2}$. Thus any price curves f_1, f_2 must satisfy $C_{\mathbf{f}}(x_1) = f_1(\frac{1}{2}) = 1$, $C_{\mathbf{f}}(x_2) = f_2(\frac{1}{2}) = 1$. But then $C_{\mathbf{f}}(x_3) = f_1(\frac{1}{2}) + f_2(\frac{1}{2}) = 2$, which is a contradiction. Thus no max-min allocation can be supported. □

The good news is that the L_∞ norm can be approximated to arbitrary precision by L_q norms, leading to the following theorem. We use $\Phi_{MM}(\mathbf{x}) = \min_{i \in N} u_i(x_i)$ to denote the max-min welfare of allocation $\mathbf{x}$.

Theorem 2.7.1. *Suppose that* $\max_{j \in M} w_{ij} = 1$ *for all* $i \in N$*. Then for every* $\epsilon > 0$*, there exists an allocation* $\mathbf{x}$ *that can be supported by price curves where* $\Phi_{MM}(\mathbf{x}) \geq (1 - \epsilon) \max_{\mathbf{x}'} \Phi_{MM}(\mathbf{x}')$.

Proof. Let w'_{ij} be rescaled versions of w_{ij} so that they are L_q-normed for a q to be chosen later. Specifically, let $\alpha_i = (\sum_{j \in M} w_{ij}^q)^{1/q}$, and let $w'_{ij} = w_{ij}/\alpha_i$.

Note that $\sum_{j \in M} w_{ij}'^q = 1$ for all $i \in N$. By Corollary 2.7.1.1, there exists an allocation with optimal max-min welfare with respect to weights w'_{ij} that can be supported by price curves. Let $\mathbf{x}$ be this allocation. Then for all $j \in M$ and all other allocations $\mathbf{x}'$,

$$\min_{i \in N} \frac{x_{ij}}{w'_{ij}} \geq \min_{i \in N} \frac{x'_{ij}}{w'_{ij}}$$

$$\min_{i \in N} \frac{\alpha_i x_{ij}}{w_{ij}} \geq \min_{i \in N} \frac{\alpha_i x'_{ij}}{w_{ij}}$$

$$\min_{i \in N} \alpha_i u_i(x_i) \geq \min_{i \in N} \alpha_i u_i(x'_i)$$

In particular, let $\mathbf{x}^*$ be the allocation maximizing max-min welfare with the respect to the true weights w_{ij}: then $\min_{i \in N} \alpha_i u_i(x_i) \geq \min_{i \in N} \alpha_i u_i(x_i^*)$. Since $u_i(x_i^*) \geq \Phi_{MM}(\mathbf{x}^*)$ by definition, we have $\min_{i \in N} \alpha_i u_i(x_i) \geq \Phi(\mathbf{x}^*) \min_{i \in N} \alpha_i$.

Therefore for all $k \in N$, $\alpha_k u_k(x_k) \geq \Phi(\mathbf{x}^*) \min_{i \in N} \alpha_i$. Therefore

$$u_k(x_k) \geq \Phi(\mathbf{x}^*) \frac{\min_{i \in N} \alpha_i}{\alpha_k}$$

and so

$$\Phi_{MM}(\mathbf{x}) \geq \Phi(\mathbf{x}^*) \frac{\min_{i \in N} \alpha_i}{\max_{i \in N} \alpha_i}$$

It remains to show that there exists $q \geq 1$ such that $\frac{\min_{i \in N} \alpha_i}{\max_{i \in N} \alpha_i} \geq 1 - \epsilon$. This follows from the fact that $\lim_{q \to \infty} \alpha_i = (\sum_{j \in M} w_{ij}^q)^{1/q} = 1$ for all $i \in N$, which completes the proof. □

2.8 Characterization of allocations supported by weakly increasing price curves

In Section 2.5, we showed that an allocation can be supported with *strictly* increasing price curves if and only it is GDF. In this section, we provide the analogous necessary and sufficient condition for the case where any (continuous, weakly increasing) price curves are permitted. This boils down to what we called *locked-agent-freeness* (LAF). LAF is not a particularly interesting condition on its own – though as with GDF it implies a polynomial time algorithm for finding price curves – but it is crucial in allowing us to prove that maximum CES welfare allocations can be supported.

For an allocation $\mathbf{x}$, we wish to determine whether there exist price curves $\mathbf{f}$ such that $(\mathbf{x}, \mathbf{f})$ is a price curve equilibrium. Assuming $\mathbf{x}$ obeys the supply constraints, we just need to determine whether there exist price curves $\mathbf{f}$ such that $x_i \in D_i(\mathbf{f})$ for all $i \in N$.

Recall that $\mathbf{a} \succ \mathbf{b}$ if for all $j \in M$ and $\tau \in \mathbb{R}_{\geq 0}$, $\sum_{i \in N: x_{ij} \geq \tau} (a_i - b_i) \geq 0$, and there exists a (j, τ) pair such that the inequality is strict. As discussed in Section 2.5, this implies that the aggregate spending of $\mathbf{a}$ is at least that of $\mathbf{b}$ for any $\mathbf{f}$, i.e.,

$$\sum_{i \in N} (a_i - b_i) C_{\mathbf{f}}(x_i) \geq 0$$

for any price curves $\mathbf{f}$. Furthermore, we argued that for strictly increasing $\mathbf{f}$, the inequality is strict, so $\mathbf{b}$ cannot be made to pay as much as $\mathbf{a}$. When we allow weakly increasing price curves, $\mathbf{a} \succ \mathbf{b}$ simply implies that, for any marginal price where $\mathbf{a}$ would have to pay strictly more than $\mathbf{b}$, that marginal price must be zero.

We still need to ensure that $x_i \in D_i(\mathbf{f})\ \forall i \in N$, i.e., that every agent spends her full budget and cannot get more utility for free (Lemma 2.5.1). This requirement can be expressed by *locked-agent-freeness.*

Definition 2.8.1 (Locked-agent-free (LAF))**.** *For simplicity, we define two meanings of "locked":*

- *Agent i is locked in an allocation $\mathbf{x}$ if there exists a domination $\mathbf{a} \succ \mathbf{b}$ such that for all $j \in M$ where $x_{ij} > 0$, and all sufficiently small $\varepsilon > 0$, $\mathbf{a} \succ \mathbf{b}$ is strict at $(j, x_{ij} + \varepsilon)$.*
- *The* allocation *is locked if there exists $\mathbf{a} \succ \mathbf{b}$ which is strict at* every *(j, τ) for $\tau \in (0, \max_i x_{ij}]$.*

If nothing is locked in allocation $\mathbf{x}$, we say that $\mathbf{x}$ is locked-agent-free (LAF).

Intuitively, an agent being locked implies that the cost to increase her allocation must be zero, which will violate condition (b) of Lemma 2.5.1. The *allocation* being locked implies that all marginal prices must be zero, and thus all price curves must be identically zero. Clearly, any non-LAF allocation cannot be supported by price curves. Perhaps surprisingly, the opposite directly holds as well, as stated by Theorem 2.8.1.

The proof of Theorem 2.8.1 is similar to the proof of Theorem 2.5.1 for strictly increasing price curves. The main difference is that strictly increasing price curves trivially satisfy condition (b) of Lemma 2.5.1, preventing any agent from getting more utility for free. For weakly increasing price curves, however, we need to add a constraint specifically to ensure that condition is satisfied. Thus in addition to the agent-order matrix, we will define a *marginal-cost matrix* to ensure that no agent has a marginal cost of zero to increase her utility. In order to incorporate this matrix, we use a more general duality result than Lemma 2.5.2 (although still equivalent to Farkas's Lemma [141]), this one due to Motzkin. Recall that $\mathbf{v} > \mathbf{0}$ denotes a strictly positive vector, and $\mathbf{v} \gg \mathbf{0}$ strongly positive.

$$\left[\begin{array}{ccc:c:cc:c} 1 & 1 & 1 & 0 & 1 & 0 & -1 \\ 1 & 1 & 0 & 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & 1 & 1 & 1 & -1 \\ 0 & 0 & 0 & 0 & 1 & 1 & -1 \end{array}\right]$$

(a) $\mathbf{x}$ represented as a agent-order matrix A

$$\left[\begin{array}{ccc:c:cc:c} 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{array}\right]$$

(b) the corresponding marginal-cost matrix C

Figure 2.3: Example construction of the marginal-cost matrix from an agent-order matrix.

Lemma 2.8.1 (1.6.1 in [167])**.** *For matrices A, B, C over $\mathbb{R}$, the following are equivalent.*

1. $A\mathbf{y} = \mathbf{0}, B\mathbf{y} \geq \mathbf{0}, C\mathbf{y} \gg \mathbf{0}$ *has no solution*
2. $A^T\mathbf{u} + B^T\mathbf{v} + C^T\mathbf{w} = \mathbf{0}, \mathbf{v} \geq \mathbf{0}, \mathbf{w} > \mathbf{0}$ *has a solution*

Theorem 2.8.1. *Let $\mathbf{x}$ be an allocation which obeys the supply constraints and gives a nonempty bundle to at least one agent. Then $\mathbf{x}$ can be supported by weakly increasing price curves if and only if it is LAF.*

Proof. Recall that an allocation $\mathbf{x}$ is supported by price curves $\mathbf{f}$ if $x_i \in D_i(\mathbf{f})\ \forall i \in N$, and $\sum_{i\in N} x_{ij} \leq 1\ \forall j \in M$. The latter condition is satisfied by assumption, and by Lemma 2.5.1, for Leontief utilities, the former condition holds if and only if the cost $C_{\mathbf{f}}(x_i) = 1$ and there exists $j \in M$ such that $\forall \varepsilon > 0\ f_j(x_{ij} + \varepsilon w_{ij}) > f_j(x_{ij})$.

As before, let $X_j = \{x_{ij} \mid i \in N\} \setminus \{0\}$ be the set of distinct, non-zero amounts of good j allocated to some agent under $\mathbf{x}$. Label these elements such that $\tau_j^1 < \tau_j^2 < \cdots < \tau_j^{|X_j|}$. Since $f_j(0) = 0$, $f_j(x \notin X_j)$ in some sense doesn't matter – we only require that these "in-between" areas of the price curve are weakly increasing and don't violate continuity. Thus there exist price curves $\mathbf{f}$ supporting $\mathbf{x}$ if and only if there exist functions $f'_j : X_j \to \mathbb{R}_{\geq 0}$ such that

1. for all $j \in M$, $0 \leq f'_j(\tau_j^1) \leq f'_j(\tau_j^2) \leq \cdots \leq f'_j(\tau_j^{|X_j|})$ (weakly increasing)
2. for all $i \in N$, $C_{\mathbf{f}}(x_i) = \sum_j f'_j(x_{ij}) = 1$ (total cost 1)
3. for all $i \in N$, exists $r, j \in M$ such that $f'_j(\tau_j^r = x_{ij} \neq 0) < f'_j(\tau_j^{r+1})$ (positive marginal cost)

Now we are ready to set up the matrices A, B, C (all of width $\sum_j |X_j| + 1$) to which we will apply Lemma 2.8.1. As in the proof of Theorem 2.5.1, A will be the agent-order matrix, and the solution vector $\mathbf{y}$ will represent the marginal prices, with the last entry representing the total cost per agent. Thus, define

$$A\left[i,\ \sum_{\ell<j} |X_\ell| + q\right] = \begin{cases} -1 & \text{if } j = m+1, q = 1 \text{ (last column)} \\ 0 & \text{if } x_{ij} < \tau_j^q \\ 1 & \text{otherwise} \end{cases}$$

Furthermore, let B be the square identity matrix I; this will ensure that the prices are weakly increasing. Finally, we need to define the marginal-cost matrix C. As shown in Figure 2.3, we can create C based only on A: If agent i receives the largest amount of some good (row i has a 1 in the last column of some sub-block), then agent i's row in C is all 1's. Intuitively, we can set the price above $\max_i x_{ij}$ arbitrarily to ensure i has positive marginal cost, so it should be trivial to satisfy $C_i\mathbf{y} > 0$. Otherwise, agent i's row is all zeros, except that within a sub-block if there is a 1 followed by a 0 in row i in A, the position of that 0 becomes a 1 in C. Intuitively, these are the places i would have to buy more of a good to increase her utility. Formally, define

$$C\left[i,\ \sum_{\ell<j}|X_\ell|+q\right] = \begin{cases} 1 & \text{if } \exists j'\ A\left[i, \sum_{\ell\leq j'}|X_\ell|\right] = 1 \\ 1 & \text{if } q\geq 1, A\left[i,\ \sum_{\ell<j}|X_\ell|+q\right] = 0, A\left[i,\ \sum_{\ell<j}|X_\ell|+q-1\right] = 1 \\ 0 & \text{otherwise} \end{cases}$$

Since $\mathbf{x}$ gives at least one agent a nonempty bundle by assumption, A, B, C have at least two columns. We know by Lemma 2.8.1 that $\exists\mathbf{y}$ such that $A\mathbf{y} = \mathbf{0}, B\mathbf{y} \geq \mathbf{0}, C\mathbf{y} \gg \mathbf{0}$ if and only if $\nexists\mathbf{u}, \mathbf{v}, \mathbf{w}$ such that $A^T\mathbf{u} + B^T\mathbf{v} + C^T\mathbf{w} = \mathbf{0}, \mathbf{v} \geq \mathbf{0}, \mathbf{w} > \mathbf{0}$. To complete the proof, we will show that the former condition is equivalent to the existence of weakly increasing price curves supporting $\mathbf{x}$, and the latter is equivalent to either $\mathbf{x}$ or an agent i being locked.

Define $f'_j(\tau_j^q) - f'_j(\tau_j^{q-1}) = y_{\sum_{\ell<j}|X_\ell|+q}$, where for convenience we let $f'_j(\tau_j^0) = f'_j(0) = 0$. Clearly $B\mathbf{y} = \mathbf{y} \geq \mathbf{0}$ is equivalent to the requirement that price curves be weakly increasing. Furthermore, note that $C\mathbf{y} \gg \mathbf{0}$ implies $\mathbf{y} > 0$, so without loss of generality we can assume the last entry of $\mathbf{y}$ is 1. Thus as before, $A\mathbf{y} = 0$ is equivalent to the requirement that every agent's total cost equals 1. Revisiting $C\mathbf{y} \gg \mathbf{0}$, since $\mathbf{y} > \mathbf{0}$ this is trivially satisfied for every row where agent i receives the largest amount of some good – equivalently, agent i's marginal cost can trivially be made positive. Additionally, for all other agents, $C_i\mathbf{y} > 0$ is by definition equivalent to having positive marginal cost. Thus a solution vector $\mathbf{y}$ is equivalent to weakly increasing price curves supporting $\mathbf{x}$.

If no such solution exists, then we have $A^T\mathbf{u} + B^T\mathbf{v} + C^T\mathbf{w} = \mathbf{0}, \mathbf{v} \geq \mathbf{0}, \mathbf{w} > \mathbf{0}$. Rearranging, and since $B = I$, this is equivalent to $A^T\mathbf{u} \geq C^T\mathbf{w}, \mathbf{w} > \mathbf{0}$. Without loss of generality, assume $\mathbf{w}$ is only non-zero on entry i. Furthermore, for all k define $a_k = u_k$ if $u_k > 0$ and $b_k = -u_k$ if $u_k < 0$. Then $A^T\mathbf{u} \geq C^T\mathbf{w}$ is equivalent to $\mathbf{a} \succ \mathbf{b}$ such that the domination is strict wherever C_i is non-zero. If $C_i = \mathbf{1}$, this is equivalent to allocation $\mathbf{x}$ being locked. Otherwise, this is equivalent to agent i being locked. Thus $A^T\mathbf{u} \geq C^T\mathbf{w}, \mathbf{w} > \mathbf{0}$ is equivalent to something being locked in $\mathbf{x}$. □

Finally, we observe that LAF give us the following linear program, which computes price curves (or shows that none exist) in polynomial time.

Theorem 2.8.2. *Given a set of agents N, goods M, and an allocation $\mathbf{x} \in \mathbb{R}_{\geq 0}^{n\times m}$, let A be the corresponding agent-order matrix and C the marginal-cost matrix. In the following linear program, the optimal objective value is strictly positive if and only if there exist strictly increasing price curves supporting $\mathbf{x}$, in which case $\mathbf{y}$ defines such curves.*

$$\max_{\mathbf{y},\eta}\ \eta$$

$$\begin{aligned} s.t. \quad A\mathbf{y} &= \mathbf{0} \\ y_k &\geq 0 \quad \forall k \\ C_i\mathbf{y} &\geq \eta \quad \forall i \\ y_{-1} &= 1 \end{aligned}$$

Proof. As per the proof of Theorem 2.8.1, there exist strictly increasing price curves supporting $\mathbf{x}$ if and only if there is a solution to the system $A\mathbf{y} = \mathbf{0}, \mathbf{y} \geq 0, C\mathbf{y} \gg \mathbf{0}$. To turn this into a valid linear program, we replace the strict inequality $C_i\mathbf{y} > 0$ with $C_i\mathbf{y} \geq \eta$ and attempt to maximize η. Furthermore, we restrict the final entry of $\mathbf{y}$ as $\mathbf{y}_{-1} = 1$, since otherwise $\mathbf{y}$ can be scaled arbitrarily. If there is a solution with $\eta > 0$, then $\mathbf{y}$ corresponds to price curves as before, with each entry representing the difference in price between adjacent allocation amounts. These points simply need to be connected, e.g., in a piecewise linear fashion, to constitute valid price curves. □

2.9 Counterexamples

	agent 1	agent 2
good 1	1	1
good 2	1	0

Example 2.2: An instance where it is necessary to give a price of zero to some goods (which is a form of weakly increasing price curves) in order to support the maximum Nash or CES welfare allocation. Assume each good has supply 1. Nash welfare is maximized by splitting good 1 evenly between the two agents, and allowing agent 1 to purchase an equal quantity of good 2. This only possible if the price of good 2 is zero: otherwise, agent 1 is paying more than agent 2. It can be verified that this same allocation is also the maximum CES welfare allocation for any $\rho \in (-\infty, 0) \cup (0, 1)$. For another interpretation, recall that the Fisher market equilibrium prices are the dual variables of the convex program for maximizing Nash welfare: thus the price of good 2 being zero corresponds to the fact that the supply constraint for good 2 is not tight in this instance.

We showed in Section 2.6 that if $w_{ij} \in \{0, 1\}$ for all $i \in N$ and $j \in M$, then for any $\rho \in (-\infty, 0) \cup (0, 1)$, every maximum CES welfare allocation can be supported by price curves. One natural question is whether this result holds if we only assume that $\max_{j \in M} w_{ij} = 1$ for all $i \in N$. The answer is no, unfortunately, as demonstrated by the following theorem. Theorem 2.9.1 only rules out ρ in the range $(\frac{1}{2}, 1)$, but we conjecture that counterexamples exist for all $\rho \in (-\infty, 0) \cup (0, 1)$.

Theorem 2.9.1. *For agents with Leontief utilities where* $\max_{j \in M} w_{ij} = 1$ *for all* $i \in N$*, for every* $\rho \in (\frac{1}{2}, 1)$*, there exist instances where no maximum CES welfare allocation can be supported by price curves.*

Proof. Consider the following instance with two goods with supply 1, and three agents, whose weights are given by the following table:

	good 1	good 2
agent 1	$1-\varepsilon$	1
agent 2	1	$1-\varepsilon$
agent 3	1	1

Let $\mathbf{x}$ be a maximum CES welfare allocation. For brevity, we write $u_i = u_i(x_i)$. In the proof of Theorem 2.6.1 given in Section 2.6, we used duality to show that for a fixed ρ, any maximum CES welfare allocation $\mathbf{x}$ has the form

$$x_{ij} = w_{ij}\Big(\sum_{j\in M} q_j w_{ij}\Big)^{\frac{1}{\rho-1}}$$

for some constants $q_1, \dots, q_m \in \mathbb{R}_{\geq 0}$. Let $\chi_i = \Big(\sum_{j\in M} q_j w_{ij}\Big)^{\frac{1}{\rho-1}}$. In our case, we have

$$\begin{aligned}
\chi_1 &= \big((1-\varepsilon)q_1 + q_2\big)^{\frac{1}{\rho-1}} \\
\chi_2 &= \big(q_1 + (1-\varepsilon)q_2\big)^{\frac{1}{\rho-1}} \\
\chi_3 &= \big(q_1 + q_2\big)^{\frac{1}{\rho-1}}
\end{aligned}$$

Thus we have $x_{ij} = w_{ij}\chi_i$ for all $i \in N$ and $j \in M$. We proceed by case analysis.

Case 1: $(1-\varepsilon)\chi_1 > \chi_3$. In this case, we have

$$x_{11} = (1-\varepsilon)\chi_1 > \chi_3 = x_{31} \qquad \text{and} \qquad x_{12} = \chi_1 > \chi_3 = x_{32}$$

So $x_{1j} > x_{3j}$ for every good j. Let $\mathbf{a}$ be the vector with $a_1 = 1$ and $a_i = 0$ for $i \neq 1$, and let $\mathbf{b}$ be the vector with $b_3 = 1$ and $b_i = 0$ for $i \neq 3$. Then $\mathbf{a} \succ \mathbf{b}$. Furthermore: the domination is strict at x_{3j} for each good $j \in M$. This means that agent 3 is locked. Therefore by Theorem 2.8.1, the $\mathbf{x}$ cannot be supported by price curves, and we are done.

Case 2: $(1-\varepsilon)\chi_2 > \chi_1$. By a symmetrical argument, we have $x_{2j} > x_{3j}$ for every good j, so agent 3 is again locked, and we are done.

Case 3: $(1-\varepsilon)\chi_1 \leq \chi_3$ and $(1-\varepsilon)\chi_2 \leq \chi_3$. This implies that $(1-\varepsilon)^{\rho-1}\chi_1^{\rho-1} \geq \chi_3^{\rho-1}$ and $(1-\varepsilon)^{\rho-1}\chi_2^{\rho-1} \geq \chi_3^{\rho-1}$. Note that the inequality flipped because $\rho - 1 < 0$. Therefore

$$\begin{aligned}
(1-\varepsilon)^{\rho-1}\chi_1^{\rho-1} + (1-\varepsilon)^{\rho-1}\chi_2^{\rho-1} &\geq 2\chi_3^{\rho-1} \\
(1-\varepsilon)^{\rho-1}\Big((1-\varepsilon)q_1 + q_2\Big) + (1-\varepsilon)^{\rho-1}\Big(q_1 + (1-\varepsilon)q_2\Big) &\geq 2\Big(q_1 + q_2\Big) \\
(1-\varepsilon)^{\rho-1}(2-\varepsilon)\Big(q_1 + q_2\Big) &\geq 2\Big(q_1 + q_2\Big) \\
\ln\Big((1-\varepsilon)^{\rho-1}(2-\varepsilon)\Big) &\geq \ln 2 \\
(\rho - 1)\ln(1-\varepsilon) &\geq \ln 2 - \ln(2-\varepsilon) \\
\rho &\leq 1 + \frac{\ln 2 - \ln(2-\varepsilon)}{\ln(1-\varepsilon)}
\end{aligned}$$

Note that the sign flipped in the last step because $\ln(1-\varepsilon) < 0$.

The resulting right hand side is some real-numbered value, so whenever ρ is greater than that, we obtain a contradiction. Taking the limit as ε goes to 0 shows us that the right hand side may be arbitrarily close to $\frac{1}{2}$. This shows that for any $\rho > \frac{1}{2}$, there exists an $\varepsilon > 0$ such that in the above instance, no maximum CES welfare allocation can be supported by price curves. □

2.9.1 Difficulties in analyzing linear utilities

We assumed throughout the chapter that agents have Leontief utilities. One natural question is whether our results extend to other classes of utilities: in particular, linear utilities. The answer is no, in general. A *linear* utility function is defined by

$$u_i(x_i) = \sum_{j \in M} w_{ij} x_{ij}$$

where w_{ij} is still the weight that agent i has for good j.

Leontief utilities have the very nice property that agents always purchase goods in a fixed proportion. It does not matter exactly how the cost within each bundle was distributed across goods, because each agent will always purchase goods in the same proportions, regardless of the underlying costs. We do not have this luxury with linear utilities. In this setting, the proportions in which each agent purchases goods depend on a complex interaction between her values for the goods, and the price curves. This makes it very difficult to reason about what agents will purchase given a set of price curves. In fact, each agent's optimization problem

$$\underset{x_i \in \mathbb{R}^m_{\geq 0}:\ C_{\mathbf{f}}(x_i) \leq 1}{\arg\max} \quad u_i(x_i)$$

may not even be convex.

Thus in order for $(\mathbf{x}, \mathbf{f})$ to form a price curve equilibrium for linear utilities, a complex set of conditions would need to be satisfied. We note that $C_{\mathbf{f}}(x_i) = 1$ is still necessary, and so GDF is still a necessary condition (for strictly increasing price curves), but it is certainly not sufficient (Example 2.3).

2.10 Omitted proofs

Lemma 2.5.1. *Given price curves $\mathbf{f}$, $x_i \in D_i(\mathbf{f})$ if and only if both of the following hold: (a) $C_{\mathbf{f}}(x_i) = 1$, and (b) there exists $j \in M$ such that for all $\varepsilon > 0$, $f_j(x_{ij} + \varepsilon w_{ij}) > f_j(x_{ij})$.*

Proof. ($\Longleftarrow$) Suppose the above conditions hold, but $x_i \notin D_i(\mathbf{f})$. Then there exists $x_i' \in D_i(\mathbf{f})$ such that $u_i(x_i') = u_i(x_i) + \varepsilon$ for some $\varepsilon > 0$. Since we assume that x_{ij} is proportional to w_{ij}, agent x must receive at least εw_{ij} more of each good j in order to increase her utility by ε. Furthermore,

	agent 1	agent 2
good 1	4	0
good 2	0	4
good 3	1	2
good 4	2	1

Example 2.3: For two agents with linear utilities, group-domination-freeness is not sufficient for the existence of price curves. Consider the instance where the agents' weights are given as above and the available supply of each good is 1. Define $\mathbf{x}$ by $x_{11} = x_{13} = x_{22} = x_{24} = 1$ and $x_{ij} = 0$ otherwise. This allocation is EF and GDF. To see that $\mathbf{x}$ cannot be supported by price curves, let $\tilde{j} = \arg\min_{j\in\{3,4\}} f_j(1)$. If $\tilde{j} = 3$, then the cost of good 3 is at most the cost of good 4, so agent 2 would buy good 3 instead of buying good 4. Similarly, if $\tilde{j} = 4$, then the cost of good 4 is at most the cost of good 3, so agent 1 would buy good 4 that instead of buying good 3.

since price curves are increasing, $f_j(x'_{ij}) \geq f_j(x_{ij})$ for every good j. However, condition (b) of the lemma implies that there exists a good j such that

$$f_j(x'_{ij}) \geq f_j(x_{ij} + \varepsilon w_{ij}) > f_j(x_{ij})$$

and thus

$$C_{\mathbf{f}}(x'_i) = \sum_{j\in M} f_j(x'_{ij}) > \sum_{j\in M} f_j(x_{ij}) = 1$$

which contradicts $x'_i \in D_i(\mathbf{f})$.

($\Longrightarrow$) Now suppose that at least one of the two conditions of the lemma does not hold. If $C_{\mathbf{f}}(x_i) \neq 1$, then either $C_{\mathbf{f}}(x_i) > 1$ and the cost exceeds the budget, or $C_{\mathbf{f}}(x_i) < 1$ so by continuity agent i could purchase more of every good and increase her utility. Either way $x_i \notin D_i(\mathbf{f})$. Thus assume that for every j, there exists an $\varepsilon_j > 0$ such that $f_j(x_{ij} + \varepsilon_j w_{ij}) = f_j(x_{ij})$. Then consider the bundle x'_i defined by $x'_{ij} = x_{ij} + \varepsilon_j w_{ij}$. This bundle has the same cost as x_i, but

$$u_i(x'_i) = \min_{j\in M} \frac{x_{ij} + \varepsilon_j w_{ij}}{w_{ij}} > \min_{j\in M} \frac{x_{ij}}{w_{ij}} = u_i(x_i)$$

contradicting $x_i \in D_i(\mathbf{f})$. □

Theorem 2.6.1. *Suppose $(\mathbf{x}^*, \mathbf{f})$ is a price curve equilibrium where for all $j \in M$, $f_j(x) = q^*_j x^{1-\rho}$ for $\rho \in (-\infty, 1)$ and nonnegative constants $q^*_1 \dots q^*_m$. If $\sum_{i\in N} x^*_{ij} = s_j$ whenever $q_j \neq 0$, then $\mathbf{x}^*$ is a maximum CES welfare allocation.*

Proof. First, for Nash welfare ($\rho = 0$), this is exactly Eisenberg and Gale's result: the linear-pricing equilibrium allocations are exactly the allocations maximizing Nash welfare [72, 73]. Thus for the rest of this proof, we assume $\rho \neq 0$.

The proof follows a duality argument very similar to the proof of Theorem 2.6.1. We use the same convex program for maximizing CES welfare, which, as stated in the proof of Theorem 2.6.1, satisfies strong duality. Suppose $(\mathbf{x}^*, \mathbf{g})$ is a PCE, where $g_j(x) = q^*_j x^{1-\rho}$ for all $j \in M$ for nonnegative constants $q^*_1 \dots q^*_m$. Let $u^*_i = u_i(x^*_i)$ be agent i's utility for $\mathbf{x}^*$, and let $\lambda^*_{ij} = q^*_j$. Let $\mathbf{u}^* = u^*_1 \dots u^*_n$, let $\mathbf{q}^* = q_1^* \dots q_m^*$, and let $\boldsymbol{\lambda}^*$ represent the vector of all λ_{ij}^*'s.

Since our convex program satisfies strong duality, the *Karush-Kuhn-Tucker* (KKT) conditions are both necessary and sufficient for optimality. Specifically, if we can show that $(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*)$ satisfies the KKT conditions, then $(\mathbf{x}^*, \mathbf{u}^*)$ is optimal for the primal. The KKT conditions are primal feasibility, dual feasibility, complementary slackness, and stationarity. Since $\mathbf{x}^*$ is a valid allocation and u_i^* is defined by $u_i^* = u_i(x_i^*)$ for all $i \in N$, primal feasibility of $(\mathbf{x}^*, \mathbf{u}^*)$ immediately follows. Since $q_j^* \geq 0$ for all $j \in M$ by assumption and $\lambda_{ij}^* \geq 0$ by definition, we have dual feasibility as well.

Complementary slackness requires that for every constraint, either the constraint is tight, or the corresponding dual variable is equal to 0. For the supply constraints, we need to show that for all $j \in M$, we have $\sum_{j \in M} x_{ij}^* = s_j$ whenever $q_j^* = 0$. This is satisfied by assumption. For the other constraints, we need to show that for all $i \in N, j \in R_i$, either $\lambda_{ij}^* = 0$ or $x_{ij}^* = u_i^*$. We will show something slightly stronger: either $\lambda_{ij}^* = 0$ or $x_{ij}^* = w_{ij} u_i^*$ (for all $j \in M$, not just in R_i). Since $\lambda_{ij}^* = q_j^*$, we have $\lambda_{ij}^* = 0$ when $q_j^* = 0$. Suppose $q_j^* \neq 0$ and $x_{ij}^* \neq w_{ij} u_i^*$. First, we must have $x_{ij}^* \geq w_{ij} u_i^*$ by the definition of Leontief utility, which implies that $x_{ij}^* > w_{ij} u_i^*$. Also, since $q_j^* \neq 0$, agent i must be spending money on good j; furthermore, she is purchasing more of good j than she needs. Instead, she could purchase $x'_{ij} = w_{ij} u_i^*$ and have some leftover money, which she could use to buy slightly more of every good and increase her utility. This would imply that x_i^* is not in agent i's demand set, which contradicts $(\mathbf{x}^*, \mathbf{f})$ being a price curve equilibrium. Therefore we have $x_{ij}^* = w_{ij} u_i^*$ whenever $q_j \neq 0$, which satisfies the complementary slackness conditions.

For stationarity, we need to show that the gradient of L with respect to $\mathbf{x}$ and $\mathbf{u}$ vanishes at $(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*)$ for every coordinate that is not zero. Specifically, we need to show that for each variable y, either $\frac{\partial L}{\partial y} = 0$, or $y = 0$ and $\frac{\partial L}{\partial y} \leq 0$. First, for $j \in R_i$ we have $\dfrac{\partial L}{\partial x_{ij}}(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^* = \lambda_{ij}^* - q_j^*$; by definition, this is equal to zero. For $j \notin R_i$, we have $\dfrac{\partial L}{\partial x_{ij}}(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*) = -q_j^*$. This is always nonpositive, since $q_j^* \geq 0$. We showed before that $x_{ij}^* = w_{ij} u_i^*$ whenever $q_j \neq 0$. Since $w_{ij} = 0$ for $j \notin R_i$, we have either $x_{ij}^* = 0$ or $q_j^* = 0$; this satisfies the stationarity condition for those variables.

Finally, consider u_i^*: since $(\mathbf{x}^*, \mathbf{f})$ is a price curve equilibrium, everyone must have positive utility (any agent could always buy a nonzero amount of every good, which would give her nonzero utility). Thus we need to show that $\dfrac{\partial L}{\partial u_i}(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*) = {u_i^*}^{\rho - 1} - \sum_{j \in R_i} \lambda_{ij}^* = 0$ for each $i \in N$. Since $(\mathbf{x}^*, \mathbf{f})$ is a price curve equilibrium, each agent must be exhausting her entire budget. Thus $C_{\mathbf{f}}(x_i^*) = 1$ for all $i \in N$, which gives us:

$$C_{\mathbf{f}}(x_i^*) = \sum_{j \in M} f_j(x_{ij}^*) = \sum_{j \in M} q_j^* {x_{ij}^*}^{1-\rho} = 1$$

Furthermore, as argued above, we have $x_{ij}^* = w_{ij} u_i^*$ whenever $q_j \neq 0$. Therefore

$$\sum_{j \in M} q_j^* {x_{ij}^*}^{1-\rho} = \sum_{j : q_j^* \neq 0} q_j^* {x_{ij}^*}^{1-\rho}$$

$$
\begin{aligned}
&= \sum_{j:q_j^* \neq 0} q_j^* (w_{ij} u_i^*)^{1-\rho} \\
&= {u_i^*}^{1-\rho} \sum_{j:q_j^* \neq 0} q_j^* w_{ij}^{1-\rho} \\
&= {u_i^*}^{1-\rho} \sum_{j:q_j^* \neq 0} q_j^* w_{ij}
\end{aligned}
$$

where the last equality is because $w_{ij} \in \{0, 1\}$. Therefore we have

$$
\begin{aligned}
{u_i^*}^{1-\rho} \sum_{j:q_j^* \neq 0} q_j^* w_{ij} &= 1 \\
\sum_{j:q_j^* \neq 0} q_j^* w_{ij} &= {u_i^*}^{\rho-1} \\
{u_i^*}^{\rho-1} - \sum_{j:q_j^* \neq 0} q_j^* w_{ij} &= 0 \\
{u_i^*}^{\rho-1} - \sum_{j \in R_i} q_j^* &= 0 \\
{u_i^*}^{\rho-1} - \sum_{j \in R_i} \lambda_{ij}^* &= 0
\end{aligned}
$$

Therefore $\frac{\partial L}{\partial u_i}(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*) = {u_i^*}^{\rho-1} - \sum_{j \in R_i} \lambda_{ij}^*$ is indeed 0. Thus the KKT conditions are satisfied. Therefore $(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*)$ is optimal for L, which implies that $(\mathbf{x}^*, \mathbf{u}^*)$ is optimal for the primal: in other words, $\mathbf{x}^*$ is a maximum CES welfare allocation. □

2.11 Conclusion

In this chapter, we analyzed price curves in several different settings, focusing on agents with Leontief utilities. Our first main result was that for strictly increasing price curves, an allocation can be supported if and only if it is GDF. We proved this by defining the agent-order matrix, and using duality theorems to show the existence of a strongly positive solution to a particular system of linear equations. Our second main result was that in the bandwidth allocation setting, the maximum CES welfare allocation can be supported by price curves. These price curves took the simple form of $f_j(x) = q_j x^{1-\rho}$. This is contrast to the standard linear pricing setting, where only maximum Nash welfare allocations can be equilibria.

There are many possible directions for future research. The first is the possibility of a simple primal-dual tâtonnement for price curves, as discussed in Section 2.6.2. We think that the approach of [48] seems especially promising in this regard.

A second possible direction is studying price curves for other classes of agent utilities, and in particular, linear utilities. We will discuss in Section 2.9 some of the challenges that linear utilities pose for analyzing price curves, but perhaps everything would fall into place with the right framework.

Last but not least, we are intrigued by the connection between GDF and the agent-order matrix and duality theorems, and we wonder if this connection could be useful for other resource allocation problems as well.

Chapter 3

Optimal Nash equilibria for bandwidth allocation

In bandwidth allocation, competing agents wish to transmit data along paths of links in a network, and each agent's utility is equal to the minimum bandwidth she receives among all links in her desired path. The previous chapter showed the existence of price curve equilibria maximizing CES welfare, but if agents act strategically, these welfare guarantees no longer hold. On the other hand, [30] proposed a mechanism whose Nash equilibria have nearly optimal Nash welfare, but their approach does consider other CES welfare functions.

In this chapter, we achieve the best of both worlds: we give a mechanism parametrized by ρ whose Nash equilibria maximize CES welfare with respect to that same value of ρ. This holds for all CES welfare functions except for $\rho = 1$. Our mechanism is a nonlinear variant of the classic trading post mechanism. We also prove that fully strategyproof mechanisms for this problem are impossible in general, with the exception of max-min welfare.

3.1 Introduction

As discussed in the previous chapter, the bandwidth allocation problem can be defined in two equivalent ways:

1. Agents have Leontief utilities with $w_{ij} \in \{0, 1\}$ for all i, j.

2. Each agent wishes to transmit data across a fixed path in a network, and her utility is equal to the minimum bandwidth she receives among all links in her desired path, i.e., the rate at which she is able to transmit data.

We use exactly the same model as in the previous chapter, with one exception: this chapter focuses on *strategic behavior*. In particular, we will focus on Nash equilibrium. We study this through the lens of *implementation theory*. A mechanism is said to Nash-implement a social choice rule Ψ (for example, Ψ could denote Nash welfare maximization) if every problem instance has least one Nash equilibrium, and every Nash equilibrium outcome is optimal with respect to Ψ. This is similar to saying that the price of anarchy – the ratio of the optimum and the "worst" Nash equilibrium – of the mechanism is 1.[1] In this chapter, we focus on pure Nash equilibria, i.e., we do not consider randomized strategies.

The result of Kelly et al. [110] assumes that agents are not strategic, and thus the Nash equilibria of their mechanism may be poor. In contrast, our augmented trading post mechanism will lead to optimal Nash equilibria, not just for Nash welfare, but for an entire family of welfare functions.

3.1.1 Trading post

Our main tool will be an augmented version of the *trading post* mechanism. In the standard trading post mechanism, each agent i submits a bid $b_{ij} \in \mathbb{R}_{\geq 0}$ on each good j, with the constraint that $\sum_j b_{ij} \leq 1$ for each agent i. Let x_{ij} be the fraction of good j that agent i receives: then trading post's allocation rule is $x_{ij} = \frac{b_{ij}}{\sum_k b_{kj}}$. In words, each agent receives a share of the good proportional to her share of the aggregate bid on that good. The bids consist of "fake money": agents have no value for leftover money.

Trading post has the desirable property that the information requirements are quite light. Each agent's best response only depends on the aggregate bid of the other agents (i.e., $\sum_{k\neq i} b_{kj}$), not on their individual bids. Furthermore, the allocation rule is decentralized in the sense that there is no centralized price computation, and each link j only needs to know the bids $b_{1j}, b_{2j}, \dots b_{nj}$.

However, the vanilla version of trading post also has limitations. First of all, it is not even guaranteed to have a Nash equilibrium for every problem instance.[2] A partial solution to this was proposed by [30]. For every $\varepsilon > 0$, they gave a modified version of trading post (parameterized by ε) that always has a Nash equilibrium, and where every Nash equilibrium attains at least $1 - \varepsilon$ of the maximum possible Nash welfare.[3] In the language of implementation theory, this mechanism Nash-implements a $1 - \varepsilon$ approximation of Nash welfare. In the course of our main result, we will strengthen this to full Nash implementation. It is important to note that their mechanism still uses the linear constraint of $\sum_j b_{ij} \leq 1$; their modification has to do with a minimum allowable bid (see Section 3.1.2 for additional discussion).

In this chapter, we augment the trading post mechanism by allowing nonlinear bid constraints: instead of $\sum_j b_{ij} \leq 1$, we require $\sum_j f_j(b_{ij}) \leq 1$ for each agent i, where each f_j is a nondecreasing function chosen by us ahead of time. Importantly, all agents are still subject to the same bid constraint, and we use the same allocation rule of $x_{ij} = \frac{b_{ij}}{\sum_k b_{kj}}$. This novel augmentation allows us to Nash-implement a wide range welfare functions, as opposed to just Nash welfare. Specifically, we will Nash-implement almost the entire family of CES welfare functions (see Section 3.1.3 for more details). This is our main result.

[1] The price of anarchy [114] concept applies only when Ψ can be written as the maximization of some cardinal function. This is true when Ψ denotes Nash welfare maximization, but is not true in general.

[2] This happens when there is a good that has large enough supply that is not the "rate limiting factor" for any agent; see Sections 3.1.2 and 3.2.1 for additional discussion.

[3] They study *Leontief utilities*, which is a generalization of bandwidth allocation to the setting where agents may desire goods in different proportions.

3.1.2 Related work

Trading post and market games. The trading post mechanism – first proposed by Shapley and Shubik [164], and sometimes called the "Shapley-Shubik game"[4] – is an example of a *strategic market game* (for an overview of strategic market games, see [93]). The study of markets has a long history in the economics literature [6, 24, 169, 173][5], but most of this work assumes that agents are *price-taking*, meaning that they treat the market prices are fixed, and do not behave strategically to affect these prices.[6] A market game, however, treats the agents as strategic players who wish to selfishly maximize their own utility. Trading post does not have explicit prices set by a centralized authority: instead, prices arise implicitly from agents' strategic behavior. In particular, $\sum_k b_{kj}$ – the aggregate bid on good j – functions as the implicit price of good j. Although the trading post mechanism is well-defined for any utility functions, the Nash equilibria are not guaranteed to have many nice properties in general, except in the limit as the number of agents goes to infinity [70] (in this case, the trading post Nash equilibria converge to the price-taking market equilibria).

The paper most relevant to our work is [30], which analyzed the performance of trading post (with a linear bid constraint) with respect to Nash welfare. They showed that for Leontief utilities (which generalize bandwidth allocation), a modified trading post mechanism approximates the Nash welfare arbitrarily well. Specifically, for any $\varepsilon > 0$, they gave a mechanism (parameterized by ε) which achieves a $1-\varepsilon$ Nash welfare approximation: there is at least one Nash equilibrium, and every Nash equilibrium has Nash welfare at least $1-\varepsilon$ times the optimal Nash welfare. Thus the price of anarchy is at most $\frac{1}{1-\varepsilon}$; equivalently, this mechanism Nash-implements a $1-\varepsilon$ approximation of Nash welfare. The reason that they were unable to perfectly implement Nash welfare is because when there is a good with supply much larger than other goods[7], vanilla trading post may not even have a Nash equilibrium. To fix this, they added a minimum allowable bid, and showed that for any $\varepsilon > 0$, there is a minimum bid that gives them a $1-\varepsilon$ Nash implementation. Instead of having a minimum allowable bid, we will add a special bid β, which will allow us to strengthen this to full Nash implementation (see Section 3.2.1).

It is worth noting that [30] also considers a broader class of valuations than Leontief, but for this broader class, only a $1/2$ approximation is achieved. Another recent paper gave a strategyproof mechanism achieving a $1/e \approx .368$ approximation of the optimal Nash welfare [51]. Their $1/e$ approximation guarantee is weaker than the $1/2$ guarantee of [30] (and the $1-\varepsilon$ guarantee for Leontief), but strategyproofness is sometimes more desirable that Nash implementation. Unfortunately, strategyproofness in the bandwidth allocation setting is generally impossible (Theorem 3.5.1).

Price-taking markets. The simplest mathematical model of a price-taking market is the Fisher market, which we discussed earlier in this thesis. Recall that the market equilibria of Fisher markets

[4] A plethora of other names have been applied to this mechanism as well, including the proportional share mechanism [78], the Chinese auction [124], and the Tullock contest in rent seeking [34].

[5] Recently, this topic has garnered significant attention in the computer science community as well (see [170] for an algorithmic exposition).

[6] There is some work treating price-taking market models as strategic games; see e.g., [1, 31, 30].

[7] Specifically, this occurs when a good has price zero. Having a much larger supply than other goods is sufficient but not necessary for this.

are guaranteed to maximize Nash welfare [72, 73], and the equilibrium prices are equal to the optimal Lagrange multipliers in the convex program for maximizing Nash welfare (the Eisenberg-Gale convex program).

In Chapter 2, we extended this model to allow nonlinear prices, where the cost of a good may be any nondecreasing function of the quantity purchased. These functions are called *price curves*. We showed that for bandwidth allocation, for any $\rho \in (-\infty, 1)$, there exist price curves that make every maximum CES welfare allocation a market equilibrium. Furthermore, these prices take a natural form: the cost of purchasing $x \in \mathbb{R}_{\geq 0}$ of good j is $g_j(x) = q_j x^{1-\rho}$, for some nonnegative constants $q_1 \dots q_m$. Interestingly, for $\rho = 0$ – which denotes Nash welfare – this function form reduces to a linear price q_j, and we know that linear pricing maximizes Nash welfare. Furthermore, $q_1 \dots q_m$ are the optimal Lagrange multipliers in the convex program for maximizing CES welfare.

Trading post with linear bid constraints ($\sum_j b_{ij} \leq 1$) can be thought of as a market game equivalent of the Fisher market model: it implements Nash welfare ([30] proved a $1-\varepsilon$ approximation, but we will strengthen this to exact implementation), and the implicit trading post prices (the aggregate bids) are equal to the Fisher market equilibrium prices. Our augmented trading post, with bid constraint $\sum_j f_j(b_{ij}) \leq 1$, can be thought of as a market game equivalent of the price curves model. The augmented trading post mechanism we use to implement CES welfare will use $f_j(b) = b^{1-\rho}$ for each good j, further strengthening this analogy.

Bandwidth allocation. Bandwidth allocation has been studied both with and without monetary payments; we consider the latter setting, following the model of Kelly et al. [110]. Although it has been known that different marking schemes (such as RED and CHOKe [81, 139]) and versions of TCP lead to different objective functions (eg. [138]), a market-based understanding was developed only for Nash Welfare, starting with the pioneering work of Kelly et al. [110]. Furthermore, the market scheme of Kelly et al. is in the price-taking setting; the only strategic market analysis of bandwidth allocation that we are aware of is the $1 - \varepsilon$ approximation of Nash welfare due to [30].

Routing games. A related topic is that of *routing games*. In a routing game, each agent has a fixed source and destination in the network, but chooses which path she uses to get there. Each agent incurs a cost for each link she travels over, and the cost each agent pays is typically nondecreasing function of the total traffic over that link. Each agent wishes to minimize the total cost she incurs by strategically choosing which path to follow. In the standard bandwidth allocation model, each agent has a fixed path, and her goal is to maximize the total amount of flow she is able to send from her source to her destination (which is equal to the minimum bandwidth she receives among links in her path). Instead of choosing which path to follow, each agent's strategy is how she bids (or more generally, how she interacts with the allocation mechanism). For an overview of routing games, see [156].

Although one could consider a model where bandwidth allocation and routing are handled simultaneously (i.e., by allowing agents to choose their paths), that would be less accurate in terms of how the internet actually works: routing (which is handled by IP) and bandwidth allocation (which is handled by TCP) are generally separate problems. Our work is about bandwidth allocation,

	$\rho = -\infty$	$\rho \in (-\infty, 1)$,	$\rho = 1$
Nash-implementable?	✓ (Thm. 3.5.3)	✓ (Thm. 3.4.1)	?
DSE-implementable?	✓ (Thm. 3.5.2)	✗ (Thm. 3.5.1)	✗ (Thm. 3.5.1)

Table 3.1: A summary of our main implementation results. Here $\rho = -\infty$ denotes max-min welfare, $\rho \in (-\infty, 1)$ includes Nash welfare as $\rho = 0$, and $\rho = 1$ denotes utilitarian welfare. DSE stands for "dominant strategy equilibrium". "✓" indicates that the type of implementation specified by the row is possible for the social choice rule specified by the column, while "✗" indicates that we give a counterexample, and "?" indicates an open question.

where pricing-based schemes (like trading post) naturally correspond to signaling mechanisms that indicate which links are congested, and an end-point protocol like TCP [42] can be thought of as agent responses.

Implementation theory. Implementation theory is the study of designing mechanisms whose outcomes coincide with some desirable social choice rule. A social choice rule could be the maximization of a cardinal function, such as a CES welfare function, or something else, such as the set of Pareto optimal allocations. A full survey is outside the scope of this chapter; we direct the interested reader to [123].

The "outcome" of a mechanism is not really well-defined; we need to specify a *solution concept*. The solution concept that we focus on for most of this chapter is Nash equilibrium. Possibly the most crucial result regarding implementation in Nash equilibrium (Nash implementation, for short) is due to Maskin [122], who identified a necessary condition for Nash implementation, and a partial converse. He showed that in a very general environment (much broader than bandwidth allocation), any Nash-implementable social choice rule must satisfy what he calls *monotonicity*. Monotonicity, in combination with a property called *no veto power*, is sufficient for Nash implementation. In Section 3.4.2, we show that CES welfare functions do not satisfy no veto power, and so cannot be Nash-implemented by Maskin's approach.

3.1.3 Our results

Our results fall into two categories, both summarized by Table 3.1.

Nash-implementing CES welfare functions. We view the Nash implementation of CES welfare functions by trading post as our main result (Theorem 3.4.1). For each $\rho \in (-\infty, 1)$, we define an augmented trading post mechanism with a nonlinear bid constraint of $\sum_j b_{ij}^{1-\rho} \leq 1$ for each agent i.[8] We denote this mechanism by $\mathcal{ATP}(\rho)$. We show that $\mathcal{ATP}(\rho)$ has at least one Nash equilibrium, and that all of its Nash equilibria maximize CES welfare.

[8] The reader may notice that for $\rho = 0$ – which corresponds to Nash welfare – this constraint reduces to the standard linear constraint of $\sum_j b_{ij} \leq 1$, which is what we should expect: we know from [30] that trading post with the linear constraint leads to good Nash welfare.

Our result improves that of Chapter 2 by strengthening our price curve equilibrium (which assumes agents are not strategic) to a strategic equilibrium, and improves that of [30] by generalizing from just Nash welfare to all CES welfare functions (except $\rho = 1$) and strengthening their $1 - \varepsilon$ approximation to exact implementation.[9] Furthermore, because the price curve equilibria can be computed in polynomial time (Section 2.6), our Nash equilibria can also be computed in polynomial time.

Our proof makes use of the following results (stated informally):

1. Theorem 3.3.1: Any Nash equilibrium of $\mathcal{ATP}$ can be converted into an "equivalent" price curve equilibrium.

2. Theorem 3.3.2: Any price curve equilibrium can be converted into an "equivalent" Nash equilibrium of $\mathcal{ATP}$.

3. Lemma 3.4.3 (Chapter 2): If $\mathbf{x}$ is a maximum CES welfare allocation, then there exist price curves $\mathbf{g}$ of the form $g_j(x) = q_j x^{1-\rho}$ such that $(\mathbf{x}, \mathbf{g})$ is a price curve equilibrium.

4. Lemma 3.4.4 (Chapter 2): If $(\mathbf{x}, \mathbf{g})$ is a price curve equilibrium and each g_j has the form $g_j(x) = q_j x^{1-\rho}$, then $\mathbf{x}$ is a maximum CES welfare allocation.

Lemmas 3.4.3 and 3.4.4 are slight restatements of results we proved in Chapter 2. Together, they imply that $\mathbf{x}$ is a maximum CES welfare allocation if and only if it is a price curve equilibrium with respect to some price curves $\mathbf{g}$ of the form $g_j(x) = q_j x^{1-\rho}$ (where $q_1 \dots q_m$ are nonnegative constants). Theorems 3.3.1 and 3.3.2 allow us to convert between price curve equilibria and Nash equilibria of $\mathcal{ATP}$, and thus enable us to apply Lemmas 3.4.3 and 3.4.4 to the Nash equilibria of $\mathcal{ATP}(\rho)$. Specifically, Theorem 3.3.1 in combination with Lemma 3.4.4 will show that any Nash equilibrium of $\mathcal{ATP}(\rho)$ maximizes CES welfare, and Theorem 3.3.2 in combination with Lemma 3.4.3 will show that $\mathcal{ATP}(\rho)$ has at least one Nash equilibrium.

Section 3.3 is devoted to proving our reduction between price curve equilibrium and Nash equilibria of trading post: Theorems 3.3.1 and 3.3.2. This reduction is the main tool we use to Nash-implement CES welfare maximization. Section 3.4 then uses this reduction, in combination with Lemmas 3.4.3 and 3.4.4, to prove our main theorem: Theorem 3.4.1.

Our trading post approach breaks down for $\rho = -\infty$ and $\rho = 1$. We are able to Nash-implement $\rho = -\infty$ by a different mechanism (see below), but we were not able to resolve whether $\rho = 1$ is Nash-implementable. We leave this as an open question.

Results for dominant strategy implementation and max-min welfare. A natural question is whether these results can be improved from Nash implementation to implementation in dominant strategy equilibrium (DSE). Section 3.5 shows that the answer is mostly no: for any $\rho \in (-\infty, 1]$, there is no mechanism which DSE-implements CES welfare maximization (Theorem 3.5.1). We do this by showing that there is no strategyproof mechanism for this problem: the revelation principle

[9] It is worth noting that the result of [30] holds for Leontief utilities, a generalization of bandwidth allocation utilities.

tells us that DSE-implementability implies strategyproofness, so impossibility of strategyproofness implies impossibility of DSE implementation.

On the positive side, we show that max-min welfare ($\rho = -\infty$) can in fact be DSE-implemented by a simple revelation mechanism (Theorem 3.5.2). This is actually stronger than strategyproofness: strategyproofness requires truth-telling to be *a* DSE, but does not rule out the possibility of additional dominant strategy equilibria that are not optimal. In contrast, DSE implementation requires *every* DSE to be optimal.

Although every DSE is also a Nash equilibrium, DSE-implementability does *not* imply Nash-implementability [58]. A DSE implementation requires every DSE to be optimal, but there could be Nash equilibria (which are not dominant strategy equilibria) that are not optimal. This means that Theorem 3.5.2 does not imply Nash-implementability of max-min welfare. In fact, our revelation mechanism which DSE-implements max-min welfare is not a Nash implementation: there exist Nash equilibria which are not optimal (see Section 3.5.2 for an example). Our last result of Section 3.5 is that there is a different mechanism which does Nash-implement max-min welfare (Theorem 3.5.3).[10]

The rest of the chapter is structured as follows. Section 3.2 formally defines the models of bandwidth allocation, price curves, trading post, and implementation theory. Section 3.3 presents our reduction between price curves and our augmented trading post mechanism. In Section 3.4, we use this reduction to Nash-implement CES welfare maximization for $\rho \in (-\infty, 1)$. Finally, Section 3.5 handles DSE-implementation and max-min welfare.

3.2 Model

We continue to use the terminology and notation defined in Chapter 1. As in Chapter 2, we assume that goods are divisible and that agents have Leontief utilities. Here we further restrict ourselves to *bandwidth allocation utilities*, which take the form:

$$u_i(x_i) = \min_{j \in R_i} x_{ij}$$

where R_i is the set of links that agent i requires. We assume that $R_i \neq \emptyset$ for all i, i.e., each agent desires at least one good. It will sometimes be useful to consider the Leontief weights w_{ij} where $w_{ij} = 1$ if $j \in R_i$, and 0 otherwise.

We use the same family of CES welfare functions are our objectives:

$$\Phi_\rho(\mathbf{x}) = \left(\sum_{i \in N} u_i(x_i)^\rho \right)^{1/\rho}$$

where ρ is a constant in $(-\infty, 0) \cup (0, 1]$. Recall that the limits as $\rho \to -\infty$ and $\rho \to 0$ yield max-min welfare and Nash welfare, respectively. Throughout the chapter, we will use $\rho = -\infty$ and $\rho = 0$ to

[10] The mechanism for Theorem 3.5.3 is unrelated to trading post: our trading post approach breaks down for both max-min welfare and utilitarian welfare. This is because $g_j(x) = q_j x^{1-\rho}$ is not a valid price curve when $\rho \to -\infty$ or when $\rho = 1$.

denote max-min welfare and Nash welfare (e.g., " This theorem holds for $\rho \in (\infty, 1)$" would include Nash welfare but not max-min welfare).

For $\rho \neq 1$, this function is strictly concave in $u_1(x_1) \dots u_n(x_n)$, so every optimal allocation $\mathbf{x}$ has the same utility vector.[11]

Price curves. We use the same basic definitions of price curves as in Chapter 2. However, in this chapter, we deviate in two small ways:

1. We assume that price curves are either strictly increasing, or identically zero (denoted $g_j \equiv 0$).
2. Our market clearing condition is: $\sum_{i \in N} x_{ij} \leq s_j$ for all $j \in M$, and $\sum_{i \in N} x_{ij} = s_j$ whenever $g_j \not\equiv 0$.

Both of these are for convenience and do not substantially change the technical model. On a high-level, we did not want to make these assumptions in Chapter 2 because we were interested in the general question of which allocations can be market equilibria. In this chapter, we are interested in optimizing a welfare function, so we can restrict ourselves to the type of price curves and allocation that will be useful in that task.

We also emphasize that we are still not considering strategic behavior in the price curve model: rather, we will use the price curve model as a tool for designing and analyzing our primary mechanism (for which we will consider strategic behavior).

3.2.1 The trading post mechanism

In the standard trading post mechanism, each agent i places a *bid* $b_i \in \mathbb{R}^m_{\geq 0}$, where $b_{ij} \in \mathbb{R}_{\geq 0}$ is the amount i bids on good j. Each agent i must obey the constraint $\sum_{j \in M} b_{ij} \leq 1$. We use $\mathbf{b} \in \mathbb{R}^{m \times n}_{\geq 0}$ to represent the matrix of all bids.

Each agent receives a fraction of the good in proportion to the fraction of the total bid on that good. Formally,

$$x_{ij} = \frac{b_{ij}}{\sum_{k \in N} b_{kj}} \cdot s_j$$

As in the Fisher market model, we assume that agents have no value for leftover money. The aggregate bid on good j is $\sum_{k \in N} b_{ij}$, and can be thought of as the "price" of good j: in fact, this analogy will be crucial in our proofs.

We augment the standard trading post mechanism in two ways. The first is necessary in order to ensure the existence of equilibrium when goods have price zero, and the second is to extend this mechanism to implement CES welfare functions beyond Nash welfare.

[11] There could be multiple optimal allocations, however. For example, consider one agent who desires two goods with supply s_1 and $s_2 > s_1$. The agent's optimal utility will be s_1, but we can either allocate the rest of the second good anyway, or leave some unallocated; the utility is unaffected.

Handling goods with price zero

In Fisher markets, it is possible for some goods to have price zero. This occurs when that good is not the "rate-limiting factor", i.e., there is enough of that good for everyone and the supply constraint is not tight. This is a problem for standard trading post: in order to receive any amount of good j, agent i must bid $b_{ij} > 0$. But if the supply constraint is not tight in the Fisher market setting, there will be at least one agent receiving more of the good than they need. Such an agent will decrease their bid so that she is only receiving what she needs. However, this process will continue infinitely, with agents repeatedly decreasing their bids on this good, but never reaching bid 0.

To handle this, we present the following modified allocation rule. We allow an additional special bid of β so that $b_{ij} \in \mathbb{R}_{\geq 0} \cup \{\beta\}$. Conceptually, a bid of 0 indicates that the agent actually does not want the good; bidding β indicates that the agent desires the good, but is hoping to get it for free, so to speak. We treat β as zero in arithmetic, for example, in the constraint $\sum_{j \in M} b_{ij} \leq 1$. Similarly, we interpret $b_{ij} > 0$ to mean $b_{ij} \notin \{0, \beta\}$.

Our modified allocation rule follows this series of steps:

1. If at least one agent bids a positive (i.e., neither 0 nor β) amount on good j, we follow the standard trading post rule: $x_{ij} = \frac{b_{ij}}{\sum_{k \in N} b_{ij}} s_j$.
2. However, if all agents bid 0 either or β on good j, then we allow each agent to have as much good j as they want. Specifically, for any agent i with $b_{ij} = \beta$, let ℓ_i be an arbitrary good with $b_{i\ell_i} > 0$. Then we allocate $x_{ij} = x_{i\ell_i}$. For completeness, if there is no good ℓ with $b_{i\ell} > 0$ (although this will never happen at equilibrium), we set $x_{ij} = 0$. For agents i bidding 0 on good j, we set $x_{ij} = 0$.
3. After following the above steps, for any good ℓ where $\sum_{i \in N} x_{i\ell} > s_\ell$ (violating the supply constraint), for all $i \in N$ bidding β on good ℓ, we set $x_{ij} = 0$ for all $j \in M$ as a penalty. In words, if so many agents try to get good j for free that the supply constraint is violated, they are all penalized by receiving nothing. Not to worry: this will never happen at equilibrium.

This modification will allow us to simulate a good having price zero.

It is important that we allow separate bids of 0 and β. Consider a good j where $b_{kj} \in \{0, \beta\}$ for all $k \in N$. Suppose some agent i does not need good j, and bidding β would cause the supply constraint to be violated and the Step 3 penalty to be invoked. Such an agent can bid 0 on good j, which allows her to still spend no money on this good, without the possibility of invoking the Step 3 penalty.

Allowing nonlinear constraints

It will turn out that trading post with the standard constraint of $\sum_{j \in M} b_{ij} \leq 1$ implements Nash welfare. To implement other CES welfare functions, let $\mathbf{f} = (f_1 \dots f_m)$ be nondecreasing functions from $\mathbb{R}_{\geq 0}$ to $\mathbb{R}_{\geq 0}$. We call $\mathbf{f}$ the *constraint curves*. Like price curves, we assume that each f_j is continuous and normalized. Unlike price curves, we require each f_j to be strictly increasing: $f_j \equiv 0$

is not allowed. Throughout the chapter, we will use $\mathbf{f}, \mathbf{f}'$ to denote constraint curves and $\mathbf{g}, \mathbf{g}'$ to denote price curves.

We define the mechanism $\mathcal{ATP}(\mathbf{f})$ as follows. Given bids $\mathbf{b} = (b_1 \dots b_n) \in \mathbb{R}^{n \times m}_{\geq 0}$, $\mathcal{ATP}(\mathbf{f})$ allocates each good j according to the three-step allocation rule described in the previous section. However, each agent's bid constraint is now

$$\sum_{j \in M} f_j(b_{ij}) \leq 1$$

We can define $C_{\mathbf{f}}(b_i)$ like we defined $C_{\mathbf{g}}(x_i)$ for price curves $\mathbf{g}$ and a bundle x_i. Specifically, $C_{\mathbf{f}}(b_i) = \sum_{j \in M} f_j(b_{ij})$. Thus each agent's constraint is $C_{\mathbf{g}}(x_i) \leq 1$ in the price curves model, and is $C_{\mathbf{f}}(b_i) \leq 1$ in the trading post model.

The most natural case will be when $f_1 \dots f_m$ are all the same function. In particular, let $\mathcal{ATP}(\rho)$ be the mechanism where $f_j(b) = b^{1-\rho}$ for all $j \in M$. In general, we will use $\mathcal{ATP}(\mathbf{f}, \mathbf{b})$ to denote the allocation $\mathbf{x}$ produced by the mechanism $\mathcal{ATP}(\mathbf{f})$ when agents bid $\mathbf{b} \in \mathbb{R}^{n \times m}_{\geq 0}$.

3.2.2 Implementation theory

This section covers only the basic concepts of implementation theory; we direct the reader to [123] for a broad overview of this area.

A *social choice rule* Ψ takes as input a utility profile $\mathbf{u} = u_1 \dots u_n$ and returns a set of "optimal" outcomes. In our case, Ψ will represent maximizing a CES welfare function. Define $\Psi_\rho(\mathbf{u})$ by

$$\Psi_\rho(\mathbf{u}) = \underset{\mathbf{x} \in \mathbb{R}^{n \times m}_{\geq 0}}{\arg\max} \Big(\sum_{i \in N} u_i(x_i)^\rho \Big)^{1/\rho}$$

In general, a social choice rule need not express the maximization of any cardinal function.

Let $\mathcal{C}$ be a solution concept (e.g., Nash equilibrium), H be a mechanism (sometimes called a "game form"), and $H(\mathbf{u})$ be the induced game for utility profile $\mathbf{u}$.[12] Let $\mathcal{C}(H(\mathbf{u}))$ be the set of strategy profiles[13] satisfying $\mathcal{C}$ for that game. For example, if $\mathcal{C}$ denotes Nash equilibrium, then $\mathcal{C}(H(\mathbf{u}))$ would be the set of Nash equilibria of the game $H(\mathbf{u})$. To distinguish between equilibrium strategies (e.g., what agents bid) and equilibrium outcomes (e.g., the resulting allocation), we use $\mathcal{C}_X(H(\mathbf{u}))$ to denote the set of outcomes resulting from strategy profiles satisfying $\mathcal{C}$.

Definition 3.2.1. *A mechanism H $\mathcal{C}$-*implements *a social choice rule Ψ if for any utility profile $\mathbf{u}$,*

$$\emptyset \neq \mathcal{C}_X(H(\mathbf{u})) \subseteq \Psi(\mathbf{u})$$

Using the running example of Nash equilibrium, H Nash-implements Ψ if for any utility profile $\mathbf{u}$, there is at least one Nash equilibrium, and every Nash equilibrium of $H(\mathbf{u})$ results in an outcome

[12] In general, the difference between a game and a mechanism is that the game definition includes the agent utilities, whereas a mechanism does not.

[13] A strategy profile is a list of strategies $S_1 \dots S_n$, where S_i is the strategy played by agent i. For trading post, a strategy is $b_i \in \mathbb{R}^m_{\geq 0}$, and a strategy profile is $\mathbf{b} \in \mathbb{R}^{m \times n}_{\geq 0}$.

that is optimal under Ψ. We denote the set of Nash equilibria of $H(\mathbf{u})$ by $NE(H(\mathbf{u}))$, and the set of outcomes resulting from some Nash equilibrium by $NE_X(H(\mathbf{u}))$. When only a single utility profile $\mathbf{u}$ is under consideration, we will frequently leave $\mathbf{u}$ implicit and write $NE(H)$.

It is worth noting that some of the literature refers to Definition 3.2.1 as *weak implementation*, where *full implementation* requires that $\mathcal{C}_X(H(\mathbf{u})) = \Psi(\mathbf{u})$, i.e., every outcome that is optimal under Ψ should be a Nash equilibrium outcome of $H(\mathbf{u})$. We feel that this distinction is not important in our case, since the utility vector in $\Psi_\rho(\mathbf{u})$ is unique (with the exception of $\rho = 1$, which we do not Nash implement anyway): thus allocations $\mathbf{x} \in \Psi_\rho(\mathbf{u})$ differ only in what they do with leftover supply, i.e., supply that will not affect anyone's utility. If one truly cared about this distinction, our augmented trading post mechanism could be further augmented by allowing each agent another special bid that indicated how much of the leftover supply they wanted. Since these special bids would not affect the utilities, the Nash equilibrium utilities would not be affected, and there would be a combination of leftover supply bids that achieves any maximum CES welfare allocation.[14]

We remind the reader of the following standard definitions:

1. Nash equilibrium: a strategy profile where no agent can strictly improve her utility by unilaterally changing her strategy. We consider only pure Nash equilibria, i.e., we do not allow randomized strategies.

2. Dominant strategy: a strategy that is optimal regardless of what other agents do.

3. Dominant strategy equilibrium (DSE): a strategy profile where each agent plays a dominant strategy.

4. Strategyproofness: A revelation mechanism (i.e., a mechanism that asks each agent to report her utility function) is *strategyproof* if telling the truth is a dominant strategy for every agent.

DSE-implementability implies strategyproofness via the revelation principle[15], but it is *not* generally true that any strategyproof social choice rule is DSE-implementable. Strategyproofness ensures that truth-telling is *a* dominant strategy equilibrium, but there could also be bad equilibria that are not consistent with Ψ.

By definition, every DSE is also a Nash equilibrium. However, it is *not* generally true that DSE-implementability implies Nash-implementability [58]. DSE-implementability requires that every DSE of the mechanism be optimal under Ψ, but the mechanism might have additional Nash equilibria (that are not dominant strategy equilibria) that are not consistent with Ψ. We will need to take both this and the previous paragraph into account when studying DSE implementation.

We now move on to our results, beginning with our reduction between price curves and $\mathcal{ATP}$. This reduction will be the main tool we use to show that $\mathcal{ATP}$ Nash-implement CES welfare maximization.

[14] We would also need to include another penalty step if the leftover supply bids lead to a supply constraint being violated.

[15] See Chapter 9 of [135] for an introduction to the revelation principle.

3.3 Reduction between price curves and augmented trading post

In this section, we show that any equilibrium of our augmented trading post mechanism can be transformed into a price curve equilibrium, and vice versa. Section 3.4 will use this result (along with the existence of price curve equilibria maximizing CES welfare, as shown in Chapter 2) to prove that the $\mathcal{ATP}(\rho)$ mechanism Nash-implements CES welfare maximization.

Section 3.3.1 describes the intuition behind the reduction. Section 3.3.2 presents some useful necessary and sufficient conditions for price curve equilibrium and trading post Nash equilibrium. Section 3.3.3 shows that any trading post Nash equilibrium can be transformed into a price curve equilibrium (Theorem 3.3.1), and Section 3.3.4 shows that any price curve equilibrium can be transformed into a trading post Nash equilibrium (Theorem 3.3.2).

3.3.1 Intuition behind the reduction

First, notice that augmented trading post and price curves have similar-looking constraints: $\sum_{j \in M} f_j(b_{ij}) \leq 1$ and $\sum_{j \in M} g_j(x_{ij}) \leq 1$. If $\mathbf{f} = \mathbf{g}$, these constraints become identical, so b_i is a feasible bid if and only if x_i is a feasible purchase subject to price curves $\mathbf{g}$. Suppose that $(\mathbf{x}, \mathbf{g})$ is a price curve equilibrium. For now, assume each g_j is strictly increasing (the formal proof will also handle the possibility of $g_j \equiv 0$). Let $\mathbf{x}'$ be the outcome of $\mathcal{ATP}(\mathbf{f})$ when agents bid $\mathbf{b}$ (i.e., $\mathbf{x}' = \mathcal{ATP}(\mathbf{f}, \mathbf{b})$), and suppose that $b_{ij} = x_{ij}$ for all i, j: then

$$x'_{ij} = \frac{b_{ij}}{\sum_{k \in N} b_{kj}} s_j = \frac{x_{ij}}{\sum_{k \in N} x_{ij}} s_j = x_{ij}$$

where the last equality uses the fact that $\sum_{k \in N} x_{ij} = s_j$ when $(\mathbf{x}, \mathbf{g})$ is a PCE and $g_j \not\equiv 0$.

Thus the allocation resulting from $\mathcal{ATP}(\mathbf{f})$ under bids $\mathbf{b}$ is in fact $\mathbf{x}$. Furthermore, since $(\mathbf{x}, \mathbf{g})$ is a price curve equilibrium, each agent exhausts her price curve constraint: $C_{\mathbf{g}}(x_i) = 1$. Since $\mathbf{f} = \mathbf{g}$ and $\mathbf{b} = \mathbf{x}$, this implies that $C_{\mathbf{f}}(b_i) = 1$ for all $i \in N$. Furthermore, in any price curve equilibrium with all nonzero prices, each agent should be spending exclusively on goods in her set R_i, and purchasing them in equal amounts. Thus in the trading post outcome $\mathbf{x}'$, each agent i also also spending exclusively on $j \in R_i$ and acquiring them in equal amounts.

We claim that $\mathbf{b}$ is a Nash equilibrium of $\mathcal{ATP}(\mathbf{f})$. Suppose the opposite: then there must exist an agent i and an alternate bid b'_i such that bidding b'_i instead of b_i increases her utility. Thus under b'_i, she receives strictly more of all goods in R_i. But this means that she must be bidding strictly more on each of these goods, which would violate her bid constraint, since $C_{\mathbf{f}}(b_i) = 1$ is already tight. Therefore $\mathbf{b}$ must be a Nash equilibrium of $\mathcal{ATP}(\mathbf{f})$.

The above is an informal proof of one direction of the reduction: transforming price curve equilibria into trading post equilibria. Similarly, if we are given a Nash equilibrium $\mathbf{b}$ of $\mathcal{ATP}(\mathbf{f})$, we can let $\mathbf{g} = \mathbf{f}$ (actually, $\mathbf{g}$ will be a scaled version of $\mathbf{f}$) and $\mathbf{x} = \mathcal{ATP}(\mathbf{f}, \mathbf{b})$, and use the same intuition to show that $(\mathbf{x}, \mathbf{g})$ is a price curve equilibrium.

There are several additional complications. The largest of these is dealing with goods that have price zero in $\mathbf{g}$; indeed, this is the issue that prevents vanilla trading post from implementing Nash welfare maximization [30]. Another difficulty is that in trading post, what you bid depends on others' bids (whereas for price curves, it only depends on $\mathbf{g}$). However, due to the nature of bandwidth allocation utilities, agents will always purchase in proportion to their weights w_{ij}, and the outcomes at equilibrium will correspond. We will end up with the following two theorems:

Theorem 3.3.1. *Let* $\mathbf{f}$ *be constraint curves where each* f_j *is homogenous of degree* α_j *for some* $\alpha_j > 0$*. For bids* $\mathbf{b} \in NE(\mathcal{ATP}(\mathbf{f}))$*, define nonnegative constants* $a_1 \dots a_m$ *by* $a_j = (\sum_{k \in N} b_{kj}/s_j)^{\alpha_j}$*. Define price curves* $\mathbf{g}$ *by*

$$g_j(x) = \begin{cases} 0 & \text{if } b_{ij} \in \{0, \beta\} \ \forall i \in N \\ a_j f_j(x) & \text{otherwise} \end{cases}$$

Let $\mathbf{x} = \mathcal{ATP}(\mathbf{f}, \mathbf{b})$*. Then* $(\mathbf{x}, \mathbf{g})$ *is a price curve equilibrium.*

Theorem 3.3.2. *Let* h *be any constraint curve. Let* $(\mathbf{x}, \mathbf{g})$ *be a price curve equilibrium, and define* $\mathbf{f}$ *and* $\mathbf{b}$ *by*

$$f_j(b) = \begin{cases} h(b) & \text{if } g_j \equiv 0 \\ g_j(b) & \text{otherwise} \end{cases} \qquad b_{ij} = \begin{cases} \beta & \text{if } g_j \equiv 0 \text{ and } j \in R_i \\ 0 & \text{if } g_j \equiv 0 \text{ and } j \notin R_i \\ x_{ij} & \text{otherwise} \end{cases}$$

Then $\mathbf{b}$ *is a Nash equilibrium of* $\mathcal{ATP}(\mathbf{f})$*.*

3.3.2 Equilibrium conditions for price curves and trading post

Recall that $w_{ij} = 1$ if $j \in R_i$, and 0 otherwise. The following lemma for trading post states a useful necessary and sufficient condition for Nash equilibria of $\mathcal{ATP}(\mathbf{f})$.

Lemma 3.3.1. *Let* $\mathbf{x} = \mathcal{ATP}(\mathbf{f}, \mathbf{b})$*. Then* $\mathbf{b} \in NE(\mathcal{ATP}(\mathbf{f}))$ *if and only if all of the following hold:*

1. *For all* $i \in N$*,* $x_{ij} = w_{ij} u_i(x_i)$ *for all* $j \in M$ *where there exists* $k \in N$ *with* $b_{kj} > 0$*.*
2. *For all* $i \in N$*,* $C_{\mathbf{f}}(b_i) = 1$*.*

Proof. ($\implies$) Assume that the two conditions of the lemma are true. First, we claim that $b_{ij} \in \{0, \beta\}$ for all $j \notin R_i$: agent i only spends money on goods in R_i. This is because $w_{ij} u_i(x_i) = 0$ for $j \notin R_i$, but $b_{ij} > 0$ ensures that $x_{ij} > 0$, so $x_{ij} = w_{ij} u_i(x_i)$ would be impossible. Therefore $C_{\mathbf{f}}(b_i) = \sum_{j \in M} f_j(b_{ij}) = \sum_{j \in R_i} f_j(b_{ij}) = 1$.

Now suppose that $\mathbf{b}$ is not a Nash equilibrium: then there exists an agent i and bid b_i' such that $u_i(x_i') > u_i(x_i)$, where $\mathbf{x}'$ is the resulting allocation when agent i bids b_i' and every agent $k \neq i$ still bids b_k. Condition 1 implies that $x_{ij} = u_i(x_i)$ for all $j \in R_i$ with $b_{ij} > 0$ (since $w_{ij} = 1$ for $j \in R_i$).

Since $b_{ij} > 0$ only when $j \in R_i$, we have $x_{ij} = u_i(x_i)$ whenever $b_{ij} > 0$. Thus $u_i(x_i') > x_{ij}$ when $b_{ij} > 0$. Since $x_{ij}' \geq u_i(x_i')$ for all $j \in R_i$, we have $x_{ij}' > x_{ij}$ when $b_{ij} > 0$.

We next claim that $b_{ij}' > b_{ij}$ whenever $b_{ij} > 0$. If there exists $k \neq i$ with $b_{kj} > 0$, then $b_{ij}' > b_{ij}$ is necessary to ensure that $x_{ij}' > x_{ij}$. The only other possibility is that $b_{ij} > 0$, but $b_{ij}' = \beta$, and $b_{kj} \in \{0, \beta\}$ for all $k \neq i$. But in this case, following Step 1 of $\mathcal{ATP}$'s allocation rule, $x_{ij} = s_j$. Then $u_i(x_i) = s_j$. This is the highest utility agent i could ever have, since $u_i(x_i) \leq s_j$ for all $j \in R_i$. This contradicts $u_i(x_i') > u_i(x_i)$. We conclude that $b_{ij}' > b_{ij}$ whenever $b_{ij} > 0$.

Therefore, since each f_j is strictly increasing,

$$\sum_{j \in M} f_j(b_{ij}) = \sum_{j: b_{ij} > 0} f_j(b_{ij}) < \sum_{j: b_{ij} > 0} f_j(b_{ij}') \leq \sum_{j \in M} f_j(b_{ij}')$$

Since $C_{\mathbf{f}}(b_i) = \sum_{j \in M} f_j(b_{ij}) = 1$ by assumption, we have $\sum_{j \in M} f_j(b_{ij}') > 1$. This means that b_i' violates the bid constraint, and so is not a valid bid. Therefore $\mathbf{b}$ is a Nash equilibrium.

($\Longleftarrow$) Suppose that $\mathbf{b}$ is a Nash equilibrium of $\mathcal{ATP}(\mathbf{f})$. If $C_{\mathbf{f}}(b_i) > 1$, b_i violates the supply constraint, so $\mathbf{b}$ cannot be a Nash equilibrium. If $C_{\mathbf{f}}(b_i) < 1$, agent i can improve her utility by bidding slightly more on every good (and thus receiving slightly more of every good). Thus $C_{\mathbf{f}}(b_i) = 1$ must hold.

Suppose $x_{i\ell} \neq w_{i\ell} u_i(x_i)$ for some $\ell \in M$ where there exists $k \in N$ with $b_{k\ell} > 0$. By definition of u_i, $u_i(x_i) w_{i\ell} > x_{i\ell}$ is impossible, so we must have $x_{i\ell} > w_{i\ell} u_i(x_i)$. Consider a new bid b_i' where $b_{ij}' = b_{ij}$ for all $j \neq \ell$, but b_{ij}' is such that $x_\ell = w_{i\ell} u_i(x_i)$ (where $\mathbf{x}'$ is the resulting allocation when i bids b_i' and each $k \neq i$ bids b_k). Thus $b_{i\ell}' < b_{i\ell}$.

By definition of u_i, we have $u_i(x_i') = u_i(x_i)$, but $C_{\mathbf{f}}(b_i') < C_{\mathbf{f}}(b_i) = 1$, since $f_j(b_{ij}') \leq f_j(b_{ij})$ for all $j \in M$, and $f_\ell(b_{i\ell}') < f_\ell(b_{i\ell})$. Thus there must exist a bundle b_i'' with $b_{ij}'' > b_{ij}'$ for all j, but $C_{\mathbf{f}}(b_i'') \leq 1$, i.e., b_i'' obeys the bid constraint. Furthermore, let $\mathbf{x}''$ be the resulting allocation when i bids b_i'' and each $k \neq i$ bids b_k: then $x_{ij}'' > x_{ij}'$ for all $j \in M$. Therefore $u_i(x_i'') > u_i(x_i') = u_i(x_i)$. But this means $\mathbf{b}$ cannot be a Nash equilibrium, which is a contradiction. □

Next, we give an analogous lemma for price curve equilibrium. Note that the last condition in Lemma 3.3.2 is simply one of the conditions in the definition of PCE.

Lemma 3.3.2. *An allocation* $\mathbf{x}$ *and price curves* $\mathbf{g}$ *are a PCE if and only if all of the following hold:*

1. *For all* $i \in N$, $x_{ij} = w_{ij} u_i(x_i)$ *whenever* $g_j \not\equiv 0$.

2. *For all* $i \in N$, $C_{\mathbf{g}}(x_i) = 1$.

3. *For all* $j \in M$, $\sum_{i \in N} x_{ij} \leq s_j$, *and* $\sum_{i \in N} x_{ij} = s_j$ *whenever* $g_j \not\equiv 0$.

Proof. The third condition is simply one of the two conditions in the definition of PCE. The other requirement for PCE is that $x_i \in D_i(\mathbf{g})$ for all $i \in N$, so it suffices to show that $x_i \in D_i(\mathbf{g})$ if and only if $C_{\mathbf{g}}(x_i) = 1$ and $x_{ij} = w_{ij} u_i(x_i)$ whenever $g_j \not\equiv 0$.

($\Longrightarrow$) Suppose that $C_{\mathbf{g}}(x_i) = 1$, and $x_{ij} = w_{ij} u_i(x_i)$ whenever $g_j \not\equiv 0$. We first claim that agent i only spends money on goods in R_i. This is because $w_{ij} u_i(x_i) = 0$ for $j \notin R_i$ (because $w_{ij} = 0$

for $j \notin R_i$), and spending money implies that $g_j \not\equiv 0$ and $x_{ij} > 0$, which makes $x_{ij} = w_{ij} u_i(x_i)$ impossible. Thus $C_{\mathbf{g}}(x_i) = \sum_{j \in R_i} g_j(x_{ij}) = \sum_{j \in R_i : g_j \not\equiv 0} g_j(x_{ij})$.

Now suppose for sake of contradiction that there exists another bundle x_i' that is also affordable, and $u_i(x_i') > u_i(x_i)$. For all $j \in R_i$ with $g_j \not\equiv 0$, we have $x_{ij} = w_{ij} u_i(x_i) = u_i(x_i)$ (because $w_{ij} = 1$ for $j \in R_i$), so $u_i(x_i') > x_{ij}$ for $j \in R_i$, $g_j \not\equiv 0$. Therefore

$$C_{\mathbf{g}}(x_i) = \sum_{j \in R_i : g_j \not\equiv 0} g_j(x_{ij}) < \sum_{j \in R_i : g_j \not\equiv 0} g_j(x_{ij}') \leq \sum_{j \in M} g_j(x_{ij}') = C_{\mathbf{g}}(x_i')$$

Since, $C_{\mathbf{g}}(x_i) = 1$, we have $C_{\mathbf{g}}(x_i') > 1$. But this implies that x_i' is not affordable, which is a contradiction. Therefore $x_i \in D_i(\mathbf{g})$.

($\Longleftarrow$) Suppose $x_i \in D_i(\mathbf{g})$. If $C_{\mathbf{g}}(x_i) > 1$, x_i is not affordable, which is impossible. If $C_g(x_i) < 1$, agent i can improve her utility by purchasing slightly more of every good. Thus $\sum_{j \in M} g_j(x_{ij}) = 1$ must hold.

Suppose $x_{i\ell} \neq w_{i\ell} u_i(x_i)$ for some $\ell \in M$ where $g_j \not\equiv 0$. By definition, $u_i(x_i) w_{i\ell} > x_{i\ell}$ is impossible, so we must have $x_{i\ell} > w_{i\ell} u_i(x_i)$. Consider a bundle x_i' where $x_{ij}' = x_{ij}$ for all $j \neq \ell$, but $x_{i\ell}' = w_{i\ell} u_i(x_i)$. Then $u_i(x_i') = u_i(x_i)$. Furthermore, $g_\ell(x_{i\ell}') < g_\ell(x_{i\ell})$, so $C_{\mathbf{g}}(x_i') < C_{\mathbf{g}}(x_i) \leq 1$. Consider another bundle x_i'' where $x_{ij}'' > x_{ij}'$ for all $j \in M$, but $C_{\mathbf{g}}(x_i'') \leq 1$: this is always possible because each g_j is continuous, and $C_{\mathbf{g}}(x_i') < 1$. Then x_i'' is affordable, but $u_i(x_i'') > u_i(x_i') = u_i(x_i)$. This contradicts $x_i \in D_i(\mathbf{g})$. $\square$

We are now ready to move on to the reduction itself.

3.3.3 Transforming trading post equilibria into price curve equilibria

This direction of the reduction will require an additional mild condition, involving the following definition.

Definition 3.3.1. *We say that a function $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is* homogenous of degree $\alpha > 0$ *if for any $b, c \in \mathbb{R}_{\geq 0}$, $f(c \cdot b) = c^\alpha f(b)$.*

Our main result of this section is the following theorem:

Theorem 3.3.1. *Let $\mathbf{f}$ be constraint curves where each f_j is homogenous of degree α_j for some $\alpha_j > 0$. For bids $\mathbf{b} \in NE(\mathcal{ATP}(\mathbf{f}))$, define nonnegative constants $a_1 \dots a_m$ by $a_j = (\sum_{k \in N} b_{kj}/s_j)^{\alpha_j}$. Define price curves $\mathbf{g}$ by*

$$g_j(x) = \begin{cases} 0 & \text{if } b_{ij} \in \{0, \beta\} \ \forall i \in N \\ a_j f_j(x) & \text{otherwise} \end{cases}$$

Let $\mathbf{x} = \mathcal{ATP}(\mathbf{f}, \mathbf{b})$. Then $(\mathbf{x}, \mathbf{g})$ is a price curve equilibrium.

Before proving Theorem 3.3.1, we prove several helpful lemmas (Lemmas 3.3.3 – 3.3.5). Throughout Lemmas 3.3.3 – 3.3.5, we assume $\mathbf{x}, \mathbf{g}$, and $a_1 \dots a_m$ are defined as in Theorem 3.3.1. We also assume that $\mathbf{b} \in NE(\mathcal{ATP}(\mathbf{f}))$. Let $\mathbf{x}'$ be the intermediate allocation after Step 2 of $\mathcal{ATP}$'s allocation rule.

Our first lemma simply states that all agents end up with positive utility.

Lemma 3.3.3. *For all $i \in N$, $u_i(x_i) > 0$.*

Proof. It is always possible for each agent to bid a nonzero amount on each good and obtain nonzero utility. Thus any Nash equilibrium must give each agent nonzero utility. □

The following lemma states that the intermediate allocation after Step 2 is in fact the final allocation.

Lemma 3.3.4. *We have $\mathbf{x} = \mathbf{x}'$.*

Proof. We need to show that Step 3 of $\mathcal{ATP}$'s allocation rule is not invoked. Suppose it were invoked: then there is an agent i who ends up with $x_{ij} = 0$ for all j, and thus $u_i(x_i) = 0$. But this contradicts Lemma 3.3.3. We conclude that $\mathbf{x} = \mathbf{x}'$. □

Lemma 3.3.5 states that under these constraint curves and bids, the bid constraint is equivalent to the price curves constraint.

Lemma 3.3.5. *For all $i \in N$, $C_{\mathbf{g}}(x_i) = C_{\mathbf{f}}(b_i)$.*

Proof. By the allocation rule of $\mathcal{ATP}$, for all $j \in M$ where there exists $k \in N$ with $b_{kj} > 0$, for all $i \in N$, we have

$$x'_{ij} = \frac{b_{ij}}{\sum_{k \in N} b_{kj}} s_j$$

Lemma 3.3.4 implies that $\mathbf{x} = \mathbf{x}'$. Also, since $a_j = (\sum_{k \in N} b_{kj}/s_j)^{\alpha_j}$ we have $s_j / \sum_{k \in N} b_{kj} = a_j^{-1/\alpha_j}$, so

$$x_{ij} = \frac{b_{ij}}{\sum_{k \in N} b_{kj}} s_j = b_{ij} a_j^{-1/\alpha_j}$$

whenever there exists $k \in N$ with $b_{kj} > 0$. By the definition of $\mathbf{g}$, $g_j \not\equiv 0$ if and only if there exists $k \in N$ with $b_{kj} > 0$ (since constraint curves are assumed to be strictly increasing). Therefore

$$\begin{aligned} C_{\mathbf{g}}(x_i) &= \sum_{j \in M} g_j(x_{ij}) \\ &= \sum_{j: g_j \equiv 0} g_j(x_{ij}) + \sum_{j: g_j \not\equiv 0} g_j(x_{ij}) \\ &= \sum_{j: g_j \not\equiv 0} a_j f_j(x_{ij}) \\ &= \sum_{j: g_j \not\equiv 0} a_j f_j(b_{ij} a_j^{-1/\alpha_j}) \\ &= \sum_{j: g_j \not\equiv 0} \frac{a_j}{a_j} f_j(b_{ij}) \end{aligned}$$

$$= \sum_{j:g_j \not\equiv 0} f_j(b_{ij})$$

By the definition of $\mathbf{g}$, $g_j \equiv 0$ is equivalent to $b_{kj} \in \{0, \beta\}$ for all $k \in N$. Thus $\sum_{j:g_j \equiv 0} f_j(b_{ij}) = 0$, so

$$\sum_{j:g_j \not\equiv 0} f_j(b_{ij}) = \sum_{j \in M} f_j(b_{ij})$$
$$= C_{\mathbf{f}}(b_i)$$

as required. □

We are now ready to prove the main result of this section.

Theorem 3.3.1. *Let $\mathbf{f}$ be constraint curves where each f_j is homogenous of degree α_j for some $\alpha_j > 0$. For bids $\mathbf{b} \in NE(\mathcal{ATP}(\mathbf{f}))$, define nonnegative constants $a_1 \dots a_m$ by $a_j = (\sum_{k \in N} b_{kj}/s_j)^{\alpha_j}$. Define price curves $\mathbf{g}$ by*

$$g_j(x) = \begin{cases} 0 & \text{if } b_{ij} \in \{0, \beta\} \ \forall i \in N \\ a_j f_j(x) & \text{otherwise} \end{cases}$$

Let $\mathbf{x} = \mathcal{ATP}(\mathbf{f}, \mathbf{b})$. Then $(\mathbf{x}, \mathbf{g})$ is a price curve equilibrium.

Proof. Since $\mathbf{b} \in NE(\mathcal{ATP}(\mathbf{f}))$, we have $C_{\mathbf{f}}(b_i) = 1$ for all $i \in N$ by Lemma 3.3.1. This implies $C_{\mathbf{g}}(x_i) = 1$ by Lemma 3.3.5. Lemma 3.3.1 also gives us $x_{ij} = w_{ij} u_i(x_i)$ whenever there exists $k \in N$ with $b_{kj} > 0$. As before, $g_j \not\equiv 0$ if and only if there exists $k \in N$ with $b_{kj} > 0$. Therefore $x_{ij} = w_{ij} u_i(x_i)$ whenever $g_j \not\equiv 0$.

Thus in order to apply Lemma 3.3.2, we just need to show that $\sum_{i \in N} x_{ij} \leq s_j$ for all $j \in M$, and that $\sum_{i \in N} x_{ij} = s_j$ whenever $g_j \not\equiv 0$. Since $\mathbf{x}$ is a valid allocation, we immediately have $\sum_{i \in N} x_{ij} \leq s_j$ for all $j \in M$. Consider an arbitrary good j with $g_j \not\equiv 0$: then by the definition of g_j, there exists $k \in N$ with $b_{kj} > 0$. Thus good j is allocated according to Step 1 of $\mathcal{ATP}$'s allocation rule, and we get

$$x'_{ij} = \frac{b_{ij}}{\sum_{k \in N} b_{kj}} \cdot s_j$$

Summing this across agents gives us

$$\sum_{i \in N} x'_{ij} = \sum_{i \in N} \frac{b_{ij}}{\sum_{k \in N} b_{kj}} \cdot s_j = s_j$$

Thus by Lemma 3.3.4, $\sum_{i \in N} x_{ij} = s_j$, as required. Therefore we can apply Lemma 3.3.2 and conclude that $(\mathbf{x}, \mathbf{g})$ is a PCE. □

3.3.4 Transforming price curve equilibria into trading post equilibria

Our main result of this section is the following theorem:

Theorem 3.3.2. *Let h be any constraint curve. Let $(\mathbf{x}, \mathbf{g})$ be a price curve equilibrium, and define $\mathbf{f}$ and $\mathbf{b}$ by*

$$f_j(b) = \begin{cases} h(b) & \text{if } g_j \equiv 0 \\ g_j(b) & \text{otherwise} \end{cases} \qquad b_{ij} = \begin{cases} \beta & \text{if } g_j \equiv 0 \text{ and } j \in R_i \\ 0 & \text{if } g_j \equiv 0 \text{ and } j \notin R_i \\ x_{ij} & \text{otherwise} \end{cases}$$

Then $\mathbf{b}$ is a Nash equilibrium of $\mathcal{ATP}(\mathbf{f})$.

The proof of this theorem is slightly more involved that the proof of Theorem 3.3.1, but the intuition is the same. As before, we prove this theorem via a series of lemmas (Lemmas 3.3.6 – 3.3.11). Let $\mathbf{x}' = \mathcal{ATP}(\mathbf{f}, \mathbf{b})$ be the final allocation resulting from bids $\mathbf{b}$, and let $\mathbf{x}''$ be the allocation resulting from bids $\mathbf{b}$ after Step 2 of $\mathcal{ATP}$'s allocation rule. We use these definitions and assume that $(\mathbf{x}, \mathbf{g})$ is a PCE for the remainder of Section 3.3.4.

As in the other direction of the reduction, our first lemma states that all agents end up with positive utility.

Lemma 3.3.6. *For all $i \in N$, $u_i(x_i) > 0$.*

Proof. Regardless of the price curves, it is always possible for each agent to buy a nonzero amount of each good and obtain nonzero utility. Since $x_i \in D_i(\mathbf{g})$, x_i must give agent i nonzero utility. □

We next claim that for all goods with nonzero price, the intermediate allocation after Step 2 of $\mathcal{ATP}(\mathbf{f})$ is equal to $\mathbf{x}$, the allocation from the price curve equilibrium.

Lemma 3.3.7. *For all $i \in N$, $x''_{ij} = x_{ij}$ whenever $g_j \not\equiv 0$.*

Proof. When $g_j \not\equiv 0$, $b_{ij} = x_{ij}$. Since each good is required by at least one agent, and $u_i(x_i) > 0$ for all i by Lemma 3.3.6, there exists $k \in N$ where $x_{kj} > 0$. Therefore $b_{kj} > 0$, so we follow Step 1 of $\mathcal{ATP}$'s allocation rule. Thus for all $i \in N$,

$$\begin{aligned} x''_{ij} &= \frac{b_{ij}}{\sum_{k \in N} b_{kj}} s_j \\ &= \frac{x_{ij}}{\sum_{k \in N} x_{kj}} s_j \end{aligned}$$

Since $(\mathbf{x}, \mathbf{g})$ is a PCE, Lemma 3.3.2 gives us $\sum_{k \in N} x_{kj} = s_j$ whenever $g_j \not\equiv 0$. Therefore

$$x''_{ij} = \frac{x_{ij}}{\sum_{k \in N} x_{kj}} s_j = x_{ij}$$

as required. □

The next lemma states that for all goods where some agent is bidding a positive amount, every agent's bundle in $\mathbf{x}''$ matches up exactly with her weights and her utility for x_i.

Lemma 3.3.8. *For all $j \in M$ where there exists $k \in N$ with $b_{kj} > 0$, we have $x''_{ij} = w_{ij}u_i(x_i)$.*

Proof. By the definition of $\mathbf{b}$, if $b_{kj} > 0$ for some $k \in N$, then $g_j \not\equiv 0$. Since $(\mathbf{x}, \mathbf{g})$ is a price curve equilibrium, we then have $x_{ij} = w_{ij}u_i(x_i)$ by Lemma 3.3.2. Lemma 3.3.7 gives us $x''_{ij} = x_{ij}$, so $x''_{ij} = w_{ij}u_i(x_i)$. □

Next, we show that each agent's utility for her bundle after Step 2 is equal to her utility for x_i.

Lemma 3.3.9. *For all $i \in N$, $u_i(x''_i) = u_i(x_i)$.*

Proof. It suffices to show that for all $j \in R_i$, $x''_{ij} = u_i(x_i)$.

Case 1: $j \in R_i$ and $g_j \not\equiv 0$. Lemma 3.3.7 implies that $x''_{ij} = x_{ij}$ in this case. Since $(\mathbf{x}, \mathbf{g})$ is a price curve equilibrium, Lemma 3.3.2 implies that $x_{ij} = w_{ij}u_i(x_i)$. Since $w_{ij} = 1$ for $j \in R_i$, $x''_{ij} = u_i(x_i)$, as required.

Case 2: $j \in R_i$ and $g_j \equiv 0$. If $g_j \equiv 0$, the definition of $\mathbf{b}$ implies that all agents bid either β or 0 on j. Thus will be following Step 2 of $\mathcal{ATP}$'s allocation rule. By definition of $\mathbf{b}$, $b_{ij} = \beta$ in this case. Following Step 2 of the $\mathcal{ATP}$ allocation rule, let ℓ_i be a good with $b_{i\ell_i} > 0$: then $x''_{ij} = x''_{i\ell_i}$. Since $b_{i\ell_i} > 0$ by assumption, we have $x''_{i\ell_i} = w_{i\ell_i}u_i(x_i)$ by Lemma 3.3.8. Furthermore, $b_{i\ell_i} > 0$ implies $x''_{i\ell_i} > 0$, so $w_{i\ell_i}u_i(x_i) > 0$. Thus we must have $w_{i\ell_i} = 1$, which implies $x''_{ij} = x''_{i\ell_i} = u_i(x_i)$.

Therefore $x''_{ij} = u_i(x_i)$ for all $j \in R_i$, so $u_i(x''_i) = u_i(x_i)$, as required. □

The next lemma states that the intermediate allocation after Step 2 is equal to the final allocation produced by $\mathcal{ATP}(\mathbf{f}, \mathbf{b})$.

Lemma 3.3.10. *We have $\mathbf{x}' = \mathbf{x}''$.*

Proof. We need to show that the penalty in Step 3 is not invoked. Suppose it is invoked: then there is a good j allocated by Step 2 where $\sum_{i\in N} x''_{ij} > s_j$. For each $i \in N$ bidding β on good j, define ℓ_i as usual: then $x''_{ij} = x''_{i\ell_i}$. Since $b_{i\ell_i} > 0$, Lemma 3.3.8 implies that $x''_{i\ell_i} = w_{i\ell_i}u_i(x_i)$ whenever $b_{ij} = \beta$.

$$s_j < \sum_{i\in N} x''_{ij} = \sum_{i:b_{ij}=\beta} x''_{i\ell_i} = \sum_{i:b_{ij}=\beta} w_{i\ell_i}u_i(x_i)$$

By definition of $\mathbf{b}$, we must have $g_j \equiv 0$: that is the only situation where agents bid β. Furthermore, $b_{ij} = \beta$ if and only if $j \in R_i$. Also using $w_{i\ell_i} \leq 1$ (in reality, $w_{i\ell_i} = 1$ exactly, but we only need the inequality), gives us

$$s_j < \sum_{i:j\in R_i} w_{i\ell_i}u_i(x_i) \leq \sum_{i:j\in R_i} u_i(x_i)$$

Using $w_{ij} = 1$ if and only if $j \in R_i$ then gives us

$$s_j < \sum_{i:j\in R_i} u_i(x_i) = \sum_{i:j\in R_i} w_{ij}u_i(x_i) = \sum_{i\in N} w_{ij}u_i(x_i)$$

By definition of u_i, $x_{ij} \geq w_{ij} u_i(x_i)$ for all $j \in M$. Therefore

$$\sum_{i \in N} x_{ij} \geq \sum_{i \in N} w_{ij} u_i(x_i) > s_j$$

But this implies that $\mathbf{x}$ is not a valid allocation, which is a contradiction. We conclude that Step 3 is not invoked, and thus $u_i(x_i') = u_i(x_i'')$, which is equal to $u_i(x_i)$ by Lemma 3.3.9. □

Next, we show that the price curves constraint and bid constraint coincide.

Lemma 3.3.11. *For all $i \in N$, $C_{\mathbf{f}}(b_i) = C_{\mathbf{g}}(x_i)$.*

Proof. By the definition of $\mathbf{b}$, $b_{ij} \in \{0, \beta\}$ when $g_j \equiv 0$. Therefore:

$$\begin{aligned}
C_{\mathbf{f}}(b_i) &= \sum_{j \in M} f_j(b_{ij}) \\
&= \sum_{j : g_j \not\equiv 0} f_j(b_{ij}) + \sum_{j : g_j \equiv 0} f_j(b_{ij}) \\
&= \sum_{j : g_j \not\equiv 0} f_j(b_{ij}) \\
&= \sum_{j : g_j \not\equiv 0} g_j(x_{ij}) \\
&= \sum_{j \in M} g_j(x_{ij}) \\
&= C_{\mathbf{g}}(x_i)
\end{aligned}$$

□

We are now ready to prove the main result of this section.

Theorem 3.3.2. *Let h be any constraint curve. Let $(\mathbf{x}, \mathbf{g})$ be a price curve equilibrium, and define $\mathbf{f}$ and $\mathbf{b}$ by*

$$f_j(b) = \begin{cases} h(b) & \text{if } g_j \equiv 0 \\ g_j(b) & \text{otherwise} \end{cases} \qquad b_{ij} = \begin{cases} \beta & \text{if } g_j \equiv 0 \text{ and } j \in R_i \\ 0 & \text{if } g_j \equiv 0 \text{ and } j \notin R_i \\ x_{ij} & \text{otherwise} \end{cases}$$

Then $\mathbf{b}$ is a Nash equilibrium of $\mathcal{ATP}(\mathbf{f})$.

Proof. Suppose $(\mathbf{x}, \mathbf{g})$ is a price curve equilibrium. By Lemma 3.3.2, we have $C_{\mathbf{g}}(x_i) = 1$ for all $i \in N$. Thus Lemma 3.3.11 implies that $C_{\mathbf{f}}(b_i) = 1$ as well, which satisfies condition 2 of Lemma 3.3.1.

Lemma 3.3.8 implies that $x_{ij}'' = w_{ij} u_i(x_i)$ whenever there exists $k \in N$ with $b_{kj} > 0$. Combining this with Lemmas 3.3.9 and 3.3.10 gives us $x_{ij}' = w_{ij} u_i(x_i')$ whenever there exists $k \in N$ with $b_{kj} > 0$. This satisfies condition 1 of Lemma 3.3.1. Therefore by Lemma 3.3.1, $\mathbf{b}$ is a Nash equilibrium of $\mathcal{ATP}(\mathbf{f})$.

□

3.4 Nash-implementing CES welfare functions with trading post

In this section, we use the reduction between price curves and augmented trading post to show that for any $\rho \in (-\infty, 1)$, $\mathcal{ATP}(\rho)$ Nash-implements CES welfare maximization. Recall that $\mathcal{ATP}(\rho)$ is the augmented trading post mechanism where $f_j(b) = b^{1-\rho}$ for all $j \in M$. Our key tools will be the reduction from Section 3.3, and pair of lemmas from Chapter 2 regarding price curve equilibria. The final result is Theorem 3.4.1:

Theorem 3.4.1. *For any $\rho \in (-\infty, 1)$, the mechanism $\mathcal{ATP}(\rho)$ Nash-implements the maximum CES welfare social choice rule.*

Before we can prove Theorem 3.4.1, we need one more property: Section 3.4.1 shows that scaling the constraint curves does not affect the set of Nash equilibrium outcomes. We then prove the main theorem in Section 3.4.2.

3.4.1 Nash equilibria of trading post are invariant to scaling of constraint curves

In order to use the reduction from Section 3.3, we would like to set $f_j(b) = g_j(b) = q_j b^{1-\rho}$. However, this would not be a valid mechanism: $q_1 \dots q_m$ depend on the utility profile $\mathbf{u}$, and the mechanism cannot depend on $\mathbf{u}$. In this section, we show that scaling by $q_1 \dots q_m$ does not affect the Nash equilibrium outcomes of $\mathcal{ATP}$. This will allow us to use the mechanism $\mathcal{ATP}(\rho)$ instead, which does not depend on $\mathbf{u}$.

Recall that for the mechanism $\mathcal{ATP}(\mathbf{f})$, $NE(\mathcal{ATP}(\mathbf{f}))$ is the set of Nash equilibrium bids $\mathbf{b}$, and $NE_X(\mathcal{ATP}(\mathbf{f}))$ is set of allocations $\mathbf{x}$ resulting from some $\mathbf{b} \in NE(\mathcal{ATP}(\mathbf{f}))$.

Lemma 3.4.1. *Let $a_1 \dots a_m$ be positive constants and let $\mathbf{f}$ be constraint curves where each f_j is homogenous of degree $\alpha_j > 0$. Define $\mathbf{f}'$ by $f'_j(b) = a_j f_j(b)$. Then $NE_X(\mathcal{ATP}(\mathbf{f})) \subseteq NE_X(\mathcal{ATP}(\mathbf{f}'))$.*

Proof. Let $\mathbf{x}$ be an arbitrary allocation in $NE_X(\mathcal{ATP}(\mathbf{f}))$; we will show that $\mathbf{x} \in NE_X(\mathcal{ATP}(\mathbf{f}'))$. By definition, there exist bids $\mathbf{b} \in NE(\mathcal{ATP}(\mathbf{f}))$ such that $\mathbf{x} = \mathcal{ATP}(\mathbf{f}, \mathbf{b})$. Define $\mathbf{b}'$ by $b'_{ij} = a_j^{-1/\alpha_j} b_{ij}$ when $b_{ij} > 0$ and $b'_{ij} = b_{ij}$ otherwise. We first show that $C_{\mathbf{f}'}(\mathbf{b}') = C_{\mathbf{f}}(\mathbf{b})$:

$$\sum_{j \in M} f'_j(b'_{ij}) = \sum_{j \in M} a_j f_j(a_j^{-1/\alpha_j} b_{ij}) = \sum_{j \in M} a_j (a_j^{-1/\alpha_j})^{\alpha_j} f_j(b_{ij}) = \sum_{j \in M} f_j(b_{ij})$$

Let $\mathbf{x}' = \mathcal{ATP}(\mathbf{f}', \mathbf{b}')$. For any good j where $b'_{kj} > 0$ for some $k \in N$ (and thus also $b_{kj} > 0$),

$$\begin{aligned} x'_{ij} &= \frac{b'_{ij}}{\sum_{k \in N} b'_{kj}} \\ &= \frac{a_j^{-1/\alpha_j} b_{ij}}{\sum_{k \in N} a_j^{-1/\alpha_j} b_{kj}} \end{aligned}$$

$$= \frac{b_{ij}}{\sum_{k \in N} b_{kj}}$$
$$= x_{ij}$$

Thus for any good j where $b'_{kj} > 0$ for some $k \in N$, we have $x'_{ij} = x_{ij}$. For any good j where $b'_{kj} \in \{0, \beta\}$ for all k, we also have $b_{kj} \in \{0, \beta\}$ for all k. Thus in both cases we follow Step 2 of $\mathcal{ATP}$'s allocation rule. Since $x'_{ij} = x_{ij}$ for the good j where $b'_{kj} > 0$ for some k, Step 2 results in $x'_{ij} = x_{ij}$ for goods where $b'_{kj} \in \{0, \beta\}$ for all k. Therefore $\mathbf{x}' = \mathbf{x}$.

This implies that $u_i(x_i) = u_i(x'_i)$ for all $i \in N$. Therefore $x_{ij} = w_{ij} u_i(x_i)$ if and only if $x'_{ij} = w_{ij} u_i(x'_i)$. Thus the conditions of Lemma 3.3.1 hold for $\mathbf{b}, \mathbf{f}$ if and only if they hold for $\mathbf{b}', \mathbf{f}'$. Therefore since $\mathbf{b} \in NE(\mathcal{ATP}(\mathbf{f}))$, we have $\mathbf{b}' \in NE(\mathcal{ATP}(\mathbf{f}'))$, and thus $\mathbf{x} = \mathbf{x}' \in NE_X(\mathcal{ATP}(\mathbf{f}'))$. We conclude that $NE_X(\mathcal{ATP}(\mathbf{f})) \subseteq NE_X(\mathcal{ATP}(\mathbf{f}'))$. □

Lemma 3.4.2. *Let $a_1 \dots a_m$ be positive scalars and let $\mathbf{f}$ be constraint curves where each f_j is homogenous of degree α_j. Define $\mathbf{f}'$ by $f'_j(b) = a_j f_j(b)$. Then $NE_X(\mathcal{ATP}(\mathbf{f})) = NE_X(\mathcal{ATP}(\mathbf{f}'))$.*

Proof. Lemma 3.4.1 gives us $NE_X(\mathcal{ATP}(\mathbf{f})) \subseteq NE_X(\mathcal{ATP}(\mathbf{f}'))$, so it remains only to show that $NE_X(\mathcal{ATP}(\mathbf{f}')) \subseteq NE_X(\mathcal{ATP}(\mathbf{f}))$. We can actually do this by symmetry. Define $a'_1 \dots a'_m$ by $a'_j = 1/a_j$. Then $a'_1 \dots a'_m$ are positive scalars such that $f_j(b) = a'_j f'_j(b)$. Each f'_j is also be homogenous of degree α_j:

$$f'_j(c \cdot b) = a_j f_j(c \cdot b) = c^{\alpha_j} a_j f_j(b) = c^{\alpha_j} f'_j(b)$$

Then we can apply Lemma 3.4.1 with the roles of $\mathbf{f}'$ and $\mathbf{f}$ swapped to give us $NE_X(\mathcal{ATP}(\mathbf{f}')) \subseteq NE_X(\mathcal{ATP}(\mathbf{f}))$, which completes the proof. □

3.4.2 Main theorem

The last tool we need is the following pair of lemmas from Chapter 2, restated here for convenience:

Lemma 3.4.3. *For utility profile $\mathbf{u}$, $\rho \in (-\infty, 1)$, and $\mathbf{x} \in \Psi_\rho(\mathbf{u})$, there exist price curves $\mathbf{g}$ such that $(\mathbf{x}, \mathbf{g})$ is a price curve equilibrium. Furthermore, for each $j \in M$, g_j takes the form $g_j(x) = q_j x^{1-\rho}$ for some nonnegative constants $q_1 \dots q_m$.*

Lemma 3.4.4. *Suppose $\rho \in (-\infty, 1)$ and that price curves $\mathbf{g}$ take the form $g_j(x) = q_j x^{1-\rho}$ for each $j \in M$, for some nonnegative constants $q_1 \dots q_m$. Then if $(\mathbf{x}, \mathbf{g})$ is a PCE, $\mathbf{x} \in \Psi_\rho(\mathbf{u})$.*

Lemma 3.4.3 states that for any maximum CES welfare allocation $\mathbf{x}$, there exist price curves of the form $g_j(x) = q_j x^{1-\rho}$ where $(\mathbf{x}, \mathbf{g})$ is a PCE. Lemma 3.4.4 states the converse: if $\mathbf{g}$ takes the form $g_j(x) = q_j x^{1-\rho}$ for nonnegative constants $q_1 \dots q_m$, and $(\mathbf{x}, \mathbf{g})$ is a PCE, then $\mathbf{x}$ is a maximum CES welfare allocation. Furthermore, $q_1 \dots q_m$ are the Lagrange multipliers of the convex program for maximizing CES welfare, so $q_1 \dots q_m$ can be computed in polynomial time.

We are now finally ready to prove that $\mathcal{ATP}(\rho)$ Nash-implements Ψ_ρ. We will make of the follow results from previous sections:

1. Theorem 3.3.1: Any Nash equilibrium of $\mathcal{ATP}$ can be converted into an "equivalent" price curve equilibrium.
2. Theorem 3.3.2: Any price curve equilibrium can be converted into an "equivalent" Nash equilibrium of $\mathcal{ATP}$.
3. Lemma 3.4.2: The set of Nash equilibrium outcomes of $\mathcal{ATP}$ is invariant to constant scaling of the constraint curves.
4. Lemma 3.4.3: For any maximum CES welfare allocation $\mathbf{x}$, there exist price curves $\mathbf{g}$ of the form $g_j(x) = q_j x^{1-\rho}$ such that $(\mathbf{x}, \mathbf{g})$ is a PCE.
5. Lemma 3.4.4: If there exist price curves $\mathbf{g}$ of the form $g_j(x) = q_j x^{1-\rho}$ such that $(\mathbf{x}, \mathbf{g})$ is a PCE, then $\mathbf{x}$ is a maximum CES welfare allocation.

Recall that for a utility profile $\mathbf{u}$, the induced game of mechanism $\mathcal{ATP}(\rho)$ is denoted by $\mathcal{ATP}(\rho)(\mathbf{u})$. We left $\mathbf{u}$ implicit when dealing with Nash equilibria in previous sections, but we make it explicit here.

Theorem 3.4.1. *For any $\rho \in (-\infty, 1)$, the mechanism $\mathcal{ATP}(\rho)$ Nash-implements the maximum CES welfare social choice rule.*

Proof. We need to show that for any utility profile $\mathbf{u}$, $\emptyset \neq NE_X(\mathcal{ATP}(\rho)(\mathbf{u})) \subseteq \Psi_\rho(\mathbf{u})$: in words, for any $\mathbf{u}$, there is at least one Nash equilibrium, and every Nash equilibrium allocation of $\mathcal{ATP}(\rho)(\mathbf{u})$ is a maximum CES welfare allocation with respect to ρ and $\mathbf{u}$.

Pick any $\mathbf{x}^* \in \Psi_\rho(\mathbf{u})$, and define $q_1 \dots q_m$ and $\mathbf{g}$ as in Lemma 3.4.3. Define $\mathbf{f}$ by $f_j(b) = q_j b^{1-\rho}$ when $q_j \neq 0$, and $f_j(b) = b^{1-\rho}$ when $q_j = 0$. By Lemma 3.4.2, we have $NE_X(\mathcal{ATP}(\rho)(\mathbf{u})) = NE_X(\mathcal{ATP}(\mathbf{f})(\mathbf{u}))$. Thus it suffices to show that $\emptyset \neq NE_X(\mathcal{ATP}(\mathbf{f})(\mathbf{u})) \subseteq \Psi_\rho(\mathbf{u})$.

We first show that $NE_X(\mathcal{ATP}(\mathbf{f})(\mathbf{u})) \neq \emptyset$, i.e., $\mathcal{ATP}(\mathbf{f})(\mathbf{u})$ has at least one Nash equilibrium. By Lemma 3.4.3, $(\mathbf{x}^*, \mathbf{g})$ is a PCE. Since $g_j(x) = q_j x^{1-\rho}$ by Lemma 3.4.3, we have $f_j(b) = g_j(b)$ whenever $q_j \neq 0$ (which is equivalent to $g_j \equiv 0$). When $g_j \equiv 0$, $f_j(b) = b^{1-\rho}$, which is strictly increasing. Thus $\mathbf{f}$ satisfies the requirements of Theorem 3.3.2. If we define $\mathbf{b}$ as a function of $\mathbf{x}^*$ as in Theorem 3.3.2, then by Theorem 3.3.2, $\mathbf{b} \in NE(\mathcal{ATP}(\mathbf{f})(\mathbf{u}))$. Therefore $\mathcal{ATP}(\mathbf{f})(\mathbf{u})$ has at least one Nash equilibrium.

It remains to show that $NE_X(\mathcal{ATP}(\mathbf{f})(\mathbf{u})) \subseteq \Psi_\rho(\mathbf{u})$, i.e., every Nash equilibrium outcome of $\mathcal{ATP}(\mathbf{f})(\mathbf{u})$ is a maximum CES welfare allocation. Consider an arbitrary $\mathbf{x} \in NE_X(\mathcal{ATP}(\mathbf{f})(\mathbf{u}))$. Then there exists $\mathbf{b} \in NE(\mathcal{ATP}(\mathbf{f})(\mathbf{u}))$ such that $\mathbf{x} = \mathcal{ATP}(\mathbf{f}, \mathbf{b})$. Noting that each f_j is homogenous of degree $1-\rho$, define $\mathbf{g}'$ as a function of $\mathbf{f}$ and $a_1 \dots a_m$ as a function of $\mathbf{b}$ as in Theorem 3.3.1:

$$a_j = \Big(\frac{\sum_{k \in N} b_{kj}}{s_j}\Big)^{1-\rho} \quad \text{and} \quad g'_j(x) = \begin{cases} 0 & \text{if } b_{ij} \in \{0, \beta\} \ \forall i \in N \\ a_j f_j(x) & \text{otherwise} \end{cases}$$

By Theorem 3.3.1, $(\mathbf{x}, \mathbf{g}')$ is a PCE. Furthermore, we can write each g'_j as $g'_j(x) = q'_j x^{1-\rho}$ for nonnegative constants $q'_1 \dots q'_m$. Therefore by Lemma 3.4.4, $\mathbf{x} \in \Psi_\rho(\mathbf{u})$.

Thus we have shown that $\mathbf{x} \in \Psi_\rho(\mathbf{u})$ for all $\mathbf{x} \in NE_X(\mathcal{ATP}(\mathbf{f})(\mathbf{u}))$, so $NE_X(\mathcal{ATP}(\rho)(\mathbf{u})) \subseteq \Psi_\rho(\mathbf{u})$. Since $NE_X(\mathcal{ATP}(\rho)(\mathbf{u})) = NE_X(\mathcal{ATP}(\mathbf{f})(\mathbf{u}))$, we conclude that $\emptyset \neq NE_X(\mathcal{ATP}(\rho)(\mathbf{u})) \subseteq \Psi_\rho(\mathbf{u})$. □

Finally, we note that a Nash equilibrium $\mathbf{b} \in NE(\mathcal{ATP}(\rho)(\mathbf{u}))$ can be computed in polynomial time. Since $q_1 \dots q_m$ are the Lagrange multipliers of the convex program for maximizing CES welfare, they can be computed in polynomial time. Then Theorem 3.3.2 can be applied to obtain $\mathbf{b}' \in NE(\mathcal{ATP}(\mathbf{f})(\mathbf{u}))$, and finally Lemma 3.4.2 yields an equivalent $\mathbf{b} \in NE(\mathcal{ATP}(\rho)(\mathbf{u}))$.

Maskin's approach and no veto power

As discussed in Section 3.1.2, Maskin proved that in a very general environment, any social choice rule satisfying monotonicity and no veto power is Nash-implementable [122]. We briefly show that bandwidth allocation does not satisfy no veto power for any $\rho \in (-\infty, 1]$, and thus is not conducive to Maskin's approach.

Definition 3.4.1. *A social choice rule* Ψ *satisfies* no veto power *if whenever there exists an allocation* $\mathbf{x}$ *where for all* $i \in N$ *except at most 1,* $u_i(x_i) \geq u_i(y_i)$ *for all allocations* $\mathbf{y}$*, we have* $\mathbf{x} \in \Psi(\mathbf{u})$.

In words, if there is a single allocation that everyone (except at most one agent) agrees is their favorite, then that allocation should be optimal under Ψ (the last agent should not be able to "veto" this allocation). In general, agents will not agree on a favorite allocation: each agent would like to receive all of the resources herself. However, when agents' R_i sets are pairwise disjoint, it is possible for all agents to agree on a favorite allocation.

Consider an instance with n agents and n goods, each with supply 1. For all $i \in N$, let $R_i = \{i\}$: each agent just desires a single good. Consider the allocation $\mathbf{x}$ where for all $i \in \{1 \dots n-1\}$, $x_{ii} = 1$, but $x_{nn} = 0$ (and $x_{ij} = 0$ otherwise). For agents $1 \dots n-1$, this is the most utility they can possibly get, so this satisfies the precondition of Definition 3.4.1. However, for any $\rho \in (-\infty, 1]$, $\mathbf{x} \notin \Psi_\rho(\mathbf{u})$, because the CES welfare can be improved by increasing x_{nn}. Specifically, for every ρ, the unique optimal CES allocation has $x_{ii} = 1$ for all $i \in \{1 \dots n\}$.

3.5 Dominant strategy implementation, strategyproofness, and max-min welfare

In Section 3.4, we showed that for every $\rho \in (-\infty, 1)$, CES welfare maximization is Nash-implementable. A natural question to ask is whether this result can be improved to dominant strategy equilibrium implementation (DSE implementation, for short). In Section 3.5.1, we show that for every $\rho \in (-\infty, 1]$, Ψ_ρ is not DSE-implementable (Theorem 3.5.1). In contrast, Section 3.5.2 shows that max-min welfare ($\rho = -\infty$) is in fact DSE-implementable by a simple revelation mechanism (Theorem 3.5.2). Finally, Section 3.5.3 uses a more complex mechanism to show that max-min welfare is also Nash-implementable (Theorem 3.5.3).

Review of relevant concepts. An important property related to DSE implementation is strategyproofness. Recall that a mechanism is strategyproof when honestly reporting one's preferences is always a dominant strategy. As discussed in Section 3.2.2, DSE-implementability implies strategyproofness by the revelation principle, but the converse is not necessarily true: strategyproofness ensures that truth-telling is *a* dominant strategy equilibrium, but there could also be bad dominant strategy equilibria. For our positive result DSE result (Theorem 3.5.2), we will show that our mechanism is strategyproof, and also that there are no bad dominant strategy equilibria. For our negative DSE result (Theorem 3.5.1), we show that the social choice rule in question is not strategyproof, which implies that it is not DSE-implementable.

Furthermore, as also discussed in Section 3.2.2, DSE-implementability does *not* imply Nash-implementability: DSE-implementability requires every DSE to be consistent with Ψ, but the mechanism might have additional (non-DSE) Nash equilibria that are not consistent with Ψ. In fact, our DSE implementation of max-min welfare is not a Nash implementation: it may have Nash equilibria that are not optimal (see Section 3.5.2). In Section 3.5.3, we give a more complex mechanism that does Nash-implement max-min welfare (Theorem 3.5.3).

We briefly discuss a subtlety relating to uniqueness (and lack thereof). In a sense, all strategyproof mechanisms that implement a social choice rule Ψ are the same: they all ask agents to report their utility functions $u_1 \dots u_n$, then compute an outcome $\mathbf{x} \in \Psi(\mathbf{u})$[16]. However, if $\Psi(\mathbf{u})$ contains multiple elements (i.e., there are multiple optimal allocations), it may matter which is chosen. If leftover supply is allocated arbitrarily, it can be hard to reason about the optimal allocation under different utility profiles. Furthermore, not even the optimal vector of agent utilities is unique for max-min welfare and utilitarian welfare (although it is for $\rho \in (-\infty, 1)$).

Consequently, for both of our positive results (Theorems 3.5.2 and 3.5.3), we will specify our mechanism such that it selects a unique allocation for each utility profile $\mathbf{u}$ (even when they are multiple optimal allocations). For our negative result (Theorem 3.5.1), we will give an instance where an agent lying makes her utility in *every* new optimal allocation strictly larger than her utility in *every* optimal allocation under a truthful report.

3.5.1 For all $\rho \in (-\infty, 1]$, CES welfare maximization is not DSE-implementable

To show impossibility of DSE implementation, it is sufficient to show impossibility of strategyproofness. Our counterexample will be the following instance with 5 agents and 7 goods, where each row is an agent, each column is a good, and the cell in the ith row and jth column gives w_{ij}:

	g1	g2	g3	g4	g5	g6	g7
agent 1	1	1	0	0	0	0	1
agent 2	0	0	1	1	0	0	1
agent 3	0	0	0	0	1	1	1
agent 4	1	0	1	0	1	0	0
agent 5	0	1	0	1	0	1	0

[16]This is in the setting where no payments are involved, like this chapter. If payments are allowed, these mechanisms can of course differ in what agents are asked to pay.

Let the supply of good 7 be 2, and let all other goods have supply 1. Notice that agents 1, 2, and 3 all conflict on good 7, but otherwise are not in competition. Agents 4 and 5 are not in competition with each other, but each conflicts with each of agents 1, 2, and 3. Let $\mathbf{u}$ denote this utility profile, and $\mathbf{u}'$ denote the utility profile where $R'_4 = \{1, 3, 5, 7\}$ instead of $R_4 = \{1, 3, 5\}$, and all other utilities are unchanged. We will claim that under utility profile $\mathbf{u}$, agent 4 can increase her utility by misreporting R'_4 instead of R_4.

We will prove this using two main lemmas. Lemma 3.5.1 states that when agent 4 truthfully reports R_4, her utility is strictly less than 1/2. Lemma 3.5.3 states that when agent 4 lies and reports R'_4 instead, her utility is at least 1/2 (Lemma 3.5.2 is a tool used in the proof of Lemma 3.5.3). Note that each lemma is referring to agent 4's *true* utility function u_4.

Lemma 3.5.1. *For every $\rho \in (-\infty, 1]$, every $\mathbf{x} \in \Psi_\rho(\mathbf{u})$ has $u_4(x_4) < 1/2$.*

Proof. For $\rho = 0$, an optimal Nash welfare allocation can be computed explicitly, and any such allocation $\mathbf{x}$ will have $u_4(x_4) < 1/2$. Recall that although the optimal allocation may not be unique, the optimal utility vector is, since Nash welfare is strictly concave.

Let $A = \{1, 2, 3\}$ and $B = \{4, 5\}$. For $\rho \in (-\infty, 0) \cup (0, 1]$, we write the following convex program for maximizing CES welfare:

$$\begin{aligned} \max_{u_1,u_2,u_3,u_4,u_5 \in \mathbb{R}_{\geq 0}} \quad & (u_1^\rho + u_2^\rho + u_3^\rho + u_4^\rho + u_5^\rho)^{1/\rho} \\ s.t. \quad & u_i + u_k \leq 1 \quad \forall i \in A, k \in B \\ & u_1 + u_2 + u_3 \leq 2 \end{aligned}$$

We are using u_i as a variable in the convex program, but we will reserve $u_i(x_i)$ to denote agent i's utility for $\mathbf{x} \in \Psi_\rho(\mathbf{u})$.

By construction, for every $i \in A$ and $k \in B$, there is a good such j such that $x_{ij} + x_{kj} \leq 1$ where $w_{ij} = w_{kj} = 1$. This implies that for any $\mathbf{x} \in \Psi_\rho(\mathbf{u})$, $u_i(x_i) + u_k(x_k) \leq 1$ for all $i \in A$ and $k \in B$. Similarly, the supply constraint of good 7 implies that $u_1(x_1) + u_2(x_2) + u_3(x_3) \leq 2$. Furthermore, any such allocation is indeed feasible: simply let $x_{ij} = w_{ij} u_i(x_i)$ for all $i \in N$. This means that the set of possible utilities for feasible allocations is equal to the set of feasible solutions $u_1 \dots u_5$ to the above convex program. Thus this program correctly maximizes CES welfare. This implies that for every $\mathbf{x} \in \Psi_\rho(\mathbf{u})$, there exists an optimal solution to the above convex program $u_1^* \dots u_5^*$ such that $u_i(x_i) = u_i^*$ for all $i \in N$. We proceed by case analysis.

Case 1: $u_1^* + u_2^* + u_3^* = 2$. In this case, one of those three agents must have utility at least 2/3. Since agent 4 is in competition with each of those agents for a good with supply 1, this implies that $u_4^* = 1/3 < 1/2$.

Case 2: $u_1^* + u_2^* + u_3^* \neq 2$. Because of the convex program's second constraint, $u_1^* + u_2^* + u_3^* > 2$ is not feasible, so we must have $u_1^* + u_2^* + u_3^* < 2$. In this case, the convex program above reduces to:

$$\max_{u_1,u_2,u_3,u_4,u_5 \in \mathbb{R}_{\geq 0}} \quad (u_1^\rho + u_2^\rho + u_3^\rho + u_4^\rho + u_5^\rho)^{1/\rho}$$

$$s.t.\ \ u_i + u_k \leq 1 \ \ \forall i \in A, k \in B$$

We next claim that there exists u_A^* and u_B^* such that $u_i^* = u_A^*$ for all $i \in A$, and $u_k^* = u_B^*$ for all $i \in B$. Suppose there exists $i, i' \in A$ such that $u_i^* < u_{i'}^*$. Then we could increase u_1^* while still obeying the constraints of this program, and increase $\frac{1}{\rho}(u_1^\rho + u_2^\rho + u_3^\rho + u_4^\rho + u_5^\rho)$. That would imply that $u_1^* \ldots u_5^*$ could not be optimal. Furthermore, an identical argument applies to the agents in **b**.

Therefore there exists u_A^* and u_B^* such that $u_i^* = u_A^*$ for all $i \in A$, and $u_k^* = u_B^*$ for all $i \in B$, so we can further rewrite the convex program as

$$\max_{u_A, u_B \in \mathbb{R}_{\geq 0}} \ (3u_A^\rho + 2u_B^\rho)^{1/\rho}$$
$$s.t.\ \ u_A + u_B \leq 1$$

where (u_A^*, u_B^*) is the optimal solution of this program. Clearly we must have $u_B^* = 1 - u_A^*$. Furthermore, we claim that we can change the objective function from $(3u_A^\rho + 2u_B^\rho)^{1/\rho}$ to $\frac{1}{\rho}(3u_A^\rho + 2u_B^\rho)$, and that this changes the optimal value of the program, but does not change the optimal solution, i.e., the argmax. This because when $\rho > 0$, maximizing $(3u_A^\rho + 2u_B^\rho)^{1/\rho}$ is equivalent to maximizing $3u_A^\rho + 2u_B^\rho$, and when $\rho < 0$, maximizing $(3u_A^\rho + 2u_B^\rho)^{1/\rho}$ is equivalent to minimizing $3u_A^\rho + 2u_B^\rho$, which is equivalent to maximizing $\frac{1}{\rho}(3u_A^\rho + 2u_B^\rho)$. Thus our new convex program is

$$\max_{u_A \in [0,1]} \ \frac{1}{\rho}(3u_A^\rho + 2(1 - u_A)^\rho)$$

This is a program we can analyze. For $\rho = 1$, the objective function becomes $u_A + 2$, so we immediately have $u_A^* = 1$ and thus $u_4^* = u_B^* = 0$. For $\rho < 1$, we take the derivative with respect to u_A should be 0 when evaluated at u_A^*:

$$\begin{aligned}
3{u_A^*}^{\rho-1} - 2(1 - u_A^*)^{\rho-1} &= 0 \\
3{u_A^*}^{\rho-1} &= 2(1 - u_A^*)^{\rho-1} \\
3^{\frac{1}{\rho-1}} u_A^* &= 2^{\frac{1}{\rho-1}}(1 - u_A^*) \\
(3^{\frac{1}{\rho-1}} + 2^{\frac{1}{\rho-1}}) u_A^* &= 2^{\frac{1}{\rho-1}} \\
u_A^* &= \frac{2^{\frac{1}{\rho-1}}}{3^{\frac{1}{\rho-1}} + 2^{\frac{1}{\rho-1}}} \\
u_A^* &= \frac{1}{(3/2)^{\frac{1}{\rho-1}} + 1}
\end{aligned}$$

Since $\rho < 1$, $\rho - 1$ is negative, so $\frac{1}{\rho-1}$ is negative. Then since $3/2 > 1$, $(3/2)^{\frac{1}{\rho-1}} < 1$. Altogether, this implies that

$$u_1^* = u_2^* = u_3^* = u_A^* > 1/2 \quad \text{and} \quad u_4^* = u_5^* = u_B^* < 1/2$$

as required. $\square$

The following lemma is a standard property of strictly concave and differentiable functions: it essentially states that any such function is bounded above by any tangent line. This lemma is sometimes called the “Rooftop Theorem”. To avoid confusion with $\mathbf{u}'$ and $\mathbf{x}'$, we use Dh to denote the derivative of h, instead of h' (this is Euler’s notation for derivatives).

Lemma 3.5.2. *Let $h : \mathbb{R} \to \mathbb{R}$ be strictly concave and differentiable. Then for all $a, b \in \mathbb{R}$ where $a \neq b$, $h(a) < h(b) + (Dh(b))(a - b)$.*

Lemma 3.5.3. *For any $\rho \in (-\infty, 1]$, for any $\mathbf{x}' \in \Psi_\rho(\mathbf{u}')$, $u_4(x'_4) \geq 1/2$.*

Proof. It suffices to show that $u'_4(x'_4) \geq 1/2$: since $R_4 \subset R'_4$, we have

$$u_4(x'_4) = \min_{j \in R_4} x'_{4j} \geq \min_{j \in R'_4} x'_{4j} = u'_4(x'_4)$$

For $\rho = 1$, the set of optimal allocations can be computed explicitly to see that for all $\mathbf{x}' \in \Psi_{\rho=1}(\mathbf{u}')$, $u'_i(x'_i) = 1/2$ for all $i \in N$.

We now use this to show that for any $\rho \in (-\infty, 1)$, for any $\mathbf{x}' \in \Psi_\rho(\mathbf{u}')$, $u'_i(x'_i) = 1/2$ for all $i \in N$. Intuitively, the larger ρ is, the more we care about efficiency and the less we care about fairness. But if the most efficient solution (i.e., the optimal allocation for $\rho = 1$) coincides with the most fair solution (i.e., having all utilities equal), then no matter how much we care about efficiency vs fairness, we should get the same outcome.

Let $\mathbf{x}^*$ be any allocation in $\Psi_{\rho=1}(\mathbf{u}')$: then $u'_i(x^*_i) = 1/2$ for all $i \in N$. Fix a $\rho \in (-\infty, 1)$ and let $h(a) = \frac{1}{\rho} a^\rho$ if $\rho \neq 0$, and $h(a) = \log(a)$ if $\rho = 1$. In both of these h is strictly concave and differentiable. Consider any allocation $\mathbf{y}$ where for some $i \in N$, $u_i(y_i) \neq 1/2$. For brevity, let $u'^*_i = u'_i(x^*_i)$ and $u'^y_i = u'_i(y_i)$.

For all such i, Lemma 3.5.2 implies that $h(u'^y_i) < h(u'^*_i) + (Dh(u'^*_i))(u'^y_i - u'^*_i)$ for $u'^*_i \neq u'^y_i$. When $u'^y_i = 1/2 = u'^*_i$, we have $h(u'^y_i) = h(u'^*_i) + (Dh(u'^*_i))(u'^y_i - u'^*_i) = 0$. Thus for $\rho \neq 0$, we have

$$\begin{aligned}
\sum_{i \in N} \frac{1}{\rho} (u'^y_i)^\rho &= \sum_{i \in N} h(u'^y_i) \\
&< \sum_{i \in N} \Big(h(u'^*_i) + (Dh(u'^*_i))(u'^y_i - u'^*_i) \Big) \\
&= \sum_{i \in N} \Big(\frac{1}{\rho} (u'^*_i)^\rho + (Dh(1/2))(u'^y_i - u'^*_i) \Big) \\
&= \frac{1}{\rho} \sum_{i \in N} (u'^*_i)^\rho + \sum_{i \in N} (Dh(1/2))(u'^y_i - u'^*_i) \\
&= \frac{1}{\rho} \sum_{i \in N} (u'^*_i)^\rho + (Dh(1/2)) \Big(\sum_{i \in N} u'^y_i - \sum_{i \in N} u'^*_i \Big)
\end{aligned}$$

Since $\mathbf{x}^* \in \Psi_{\rho=1}(\mathbf{u}')$, $\sum_{i \in N} u'^*_i \geq \sum_{i \in N} u'^y_i$. Therefore $(Dh(1/2))\big(\sum_{i \in N} u'^y_i - \sum_{i \in N} u'^*_i\big) \leq 0$, so

$$\frac{1}{\rho} \sum_{i \in N} (u'^y_i)^\rho < \frac{1}{\rho} \sum_{i \in N} (u'^*_i)^\rho$$

As before, this implies that $\left(\sum_{i \in N} (u_i'^y)^\rho\right)^{1/\rho} < \left(\sum_{i \in N} (u_i'^*)^\rho\right)^{1/\rho}$. The analysis for $\rho = 0$ (i.e., Nash welfare) is the same, except we end up with $\sum_{i \in N} \log(u_i'^y) < \sum_{i \in N} \log(u_i'^*)$ instead, which implies $\prod_{i \in N} u_i'^y < \prod_{i \in N} u_i'^*$.

Thus for any allocation $\mathbf{y}$ where there exists $i \in N$ with $u_i'(y_i) \neq 1/2$, the CES welfare of $\mathbf{x}^*$ is better than the CES welfare of $\mathbf{y}$. This implies that for any $\rho \in (-\infty, 1]$, any $\mathbf{x}' \in \Psi_\rho(\mathbf{u}')$ must have $u_i'(x_i') = 1/2$ for all $i \in N$. □

Theorem 3.5.1. *For all $\rho \in (-\infty, 1]$, Ψ_ρ is not DSE-implementable.*

Proof. Similar to the proof of Theorem 3.5.2, it suffices to show that Ψ_ρ cannot be computed in a strategyproof mechanism. Suppose there were a strategyproof mechanism H: then for utility profile $\mathbf{u}$, H must return an allocation $\mathbf{x} \in \Psi_\rho(\mathbf{u})$, and for utility profile $\mathbf{u}'$, H must return an allocation $\mathbf{x}' \in \Psi_\rho(\mathbf{u}')$. By Lemmas 3.5.1 and 3.5.3, we have $u_4(x_4) < 1/2$ and $u_4(x_4') \geq 1/2$. If agent 4 reports $R_4' = \{1, 3, 5, 7\}$ instead of $R_4 = \{1, 3, 5\}$, she alters the utility profile from $\mathbf{u}$ to $\mathbf{u}'$, which resulting in her receiving a bundle with higher utility. Therefore H is not strategyproof. □

3.5.2 Maxmin welfare is DSE-implementable

We will claim that Mechanism 1 DSE-implements max-min welfare. We are using u_i as a variable in the convex program in Step 2, but we will reserve $u_i(x_i)$ for denoting agent i's utility for a bundle x_i. We could have used $u_i \geq \gamma$ and $u_i w_{ij} \leq x_{ij}$ for our first two constraints, but requiring $u_i = \gamma$ and $u_i w_{ij} = x_{ij}$ ensures a unique solution (and does not affect the optimal value).

Mechanism 1 A simple revelation mechanism which DSE-implements max-min welfare.

---

1. Ask each agent i report her set of desired goods R_i (which fully specifies her utility function u_i and her weights $w_{i1} \ldots w_{im}$).
2. Let $(\mathbf{x}^*, \mathbf{u}^*)$ be an optimal solution the following convex program:

$$\begin{aligned} \max_{\substack{\mathbf{x} \in \mathbb{R}_{\geq 0}^{n \times m}, \\ \mathbf{u} = (u_1 \ldots u_n) \in \mathbb{R}_{\geq 0}^n}} \quad & \gamma \\ s.t. \quad & u_i = \gamma && \forall i \in N \\ & u_i w_{ij} = x_{ij} && \forall i \in N, j \in M \\ & \sum_{i \in N} x_{ij} \leq s_j && \forall j \in M \end{aligned}$$

3. Return the allocation $\mathbf{x}^*$.

We say that an allocation $\mathbf{x}$ is *max-min-optimal* if the minimum utility in $\mathbf{x}$ is the largest possible minimum utility among all valid allocations. Formally, $\min_{i \in N} u_i(x_i) = \max_{\mathbf{x}'} \min_{i \in N} u_i(x_i')$.

Lemma 3.5.4. *Assume agents truthfully report their desired sets of goods. Then Mechanism 1 correctly computes a max-min-optimal allocation.*

Proof. Let $\mathbf{x}$ be the allocation returned by Mechanism 1, and suppose there exists $\mathbf{y}$ such that $\min_{k \in N} u_k(y_k) > \min_{k \in N} u_k(x_k)$. Consider the allocation $\mathbf{x}'$ where $x'_{ij} = w_{ij} \min_{k \in N} u_k(y_k)$ for all i, j. Then $u_i(x'_i) = \min_{k \in N} u_k(y_k)$ for all $i \in N$. Furthermore, $y_{ij} \geq w_{ij} u_i(y_i)$ by definition of $u_i(y_i)$, and $u_i(y_i) \geq u_i(x'_i)$, so

$$y_{ij} \geq w_{ij} u_i(y_i) \geq w_{ij} u_i(x'_i) = x'_{ij}$$

Thus since $\mathbf{y}$ is a valid allocation, so is $\mathbf{x}'$. Therefore $\mathbf{x}'$ is feasible for our convex program, and $\min_{i \in N} u_i(x'_i) = \min_{i \in N} u_i(y_i) > \min_{i \in N} u_i(x_i)$, so $\mathbf{x}$ could not have been an optimal solution to our convex program. □

Lemma 3.5.5. *Mechanism 1 is strategyproof.*

Proof. Let $R_1 \dots R_n$ be the true desired sets of goods. Suppose for sake of contradiction that there exists an instance where an agent i can increase her utility by reporting some $R'_i \neq R_i$. Let x_i and x'_i be agent i's bundles when she reports R_i and R'_i, respectively (assuming other agents make the same reports in both cases). Due to the constraint $u_i w_{ij} = x_{ij}$, our mechanism will set $x'_{ij} = 0$ for all $j \notin R'_i$. If $R'_i \subsetneq R_i$, then there exists a $j \in R_i$ where $j \notin R'_i$. Thus $x'_{ij} = 0$, which implies that $u_i(x'_i) = 0$, since $j \in R_i$.

Suppose $R_i \subsetneq R'_i$, and let $w'_{i1} \dots w'_{im}$ be the weights associated with R'_i. In this case, there exists a $j \in R'_i \setminus R_i$, so $w'_{ij} = 1$ and $w_{ij} = 0$. We claim that any utility vector $u_1 \dots u_n$ that is feasible in the original program (when agent i reports R_i) is also feasible in the new program (when agent i reports R'_i). Let $x_{kj} = u_k w_{kj}$ and $x'_{kj} = u_k w'_{kj}$. Since $w'_{kj} \geq w_{kj}$ for all k, j, we have $x'_{kj} \geq x_{kj}$ for all $j \in M$, so $\sum_{k \in N} x_{kj} \leq \sum_{k \in N} x'_{kj} \leq s_j$. Thus if $\mathbf{x}'$ and $u_1 \dots u_n$ are feasible together, so are $\mathbf{x}$ and $u_1 \dots u_n$. This means that the optimal value of the new program is at most the optimal value of the original program: the objective functions are the same, and the feasible set for the new program is a subset of that of the original program. Since each agent's utility is equal to the objective value of the convex program, this means that agent i's utility when she reports R'_i cannot improve.

Thus we have shown that reporting $R'_i \neq R_i$ cannot improve agent i's utility. We conclude that this mechanism is strategyproof. □

Theorem 3.5.2. *Mechanism 1 DSE-implements max-min welfare.*

Proof. Lemma 3.5.5 implies that there is at least one DSE: in particular, truthful revelation is a DSE. Thus it remains only to show that there are no "bad" dominant strategy equilibria, i.e., every DSE results in a max-min-optimal allocation.

We claim that in any DSE, the vector of utilities is the same as in the truthful DSE, which we know has optimal max-min welfare by Lemma 3.5.4. Since this is a revelation mechanism, each agent just reports a utility function u_i. Let $\mathbf{u} = u_1 \dots u_n$ be the true utility profile, and let $\mathbf{u}' = u'_1 \dots u'_n$ be an arbitrary DSE. Since the mechanism is strategyproof, we know that $\mathbf{u}$ is also a DSE. Thus for every agent i, u_i and u'_i are both dominant strategies (it is possible that $u_i = u'_i$).

For each $r \in \{1 \dots n+1\}$, define another utility profile $\mathbf{u}^r$ where each agent $i \in \{1 \dots r-1\}$ reports u_i and each agent $i \in \{r \dots n\}$ reports u'_i. Suppose every agent is reporting according to $\mathbf{u}^r$. If agent r switches from reporting u'_r to truthfully reporting u_r, she alters the utility profile from $\mathbf{u}^r$ to $\mathbf{u}^{r+1}$. Let $\mathbf{x}^r$ and $\mathbf{x}^{r+1}$ be the resulting allocations for reported utility profiles $\mathbf{u}^r$ and $\mathbf{u}^{r+1}$, respectively. Since reporting u_r and reporting u'_r are both dominant strategies for agent r, she must be indifferent between $\mathbf{x}^r$ and $\mathbf{x}^{r+1}$ (according to her true utility function, u_r). Formally, $u_r(x_r^r) = u_r(x_r^{r+1})$.

Next, by the definition of the mechanism, each agent has the same utility for her resulting bundle (according to the utility function she reports). Let γ^r be every agent's utility for the allocation $\mathbf{x}^r$, according to her reported utility function u_i^r: $u_i^r(x_i^r) = \gamma^r$. Our next claim is that $u_i(x_i^r) = \gamma^r$, i.e., each agent's true utility for $\mathbf{x}^r$ is γ^r. As before, we know that $R_i \subseteq R'_i$ for each agent i: reporting $R'_i \subsetneq R_i$ always results in getting zero utility. This means that $R_i \subseteq R_i^r$. Furthermore, the convex program ensures that each agent i receives the same amount of every good in her reported set R_i^r. Thus we have

$$u_i^r(x_i^r) = \min_{j \in R_i^r} x_{ij}^r = \min_{j \in R_i} x_{ij}^r = u_i(x_i^r)$$

Therefore $u_i(x_i^r) = \gamma^r$, i.e., each agent's true utility for $\mathbf{x}^r$ is γ^r. In particular, $u_r(x_r^r) = \gamma^r$ and $u_r(x_r^{r+1}) = \gamma^{r+1}$.

We showed above that $u_r(x_r^r) = u_r(x_r^{r+1})$, so we now have $\gamma^{r+1} = \gamma^r$ for all r. This implies $\gamma^1 = \gamma^{n+1}$. Thus each agent's true utility for $\mathbf{x}^1$ (which is γ^1) is the same as each agent's true utility for x^{n+1} (which is γ^{n+1}). By definition, $\mathbf{x}^{n+1}$ is the resulting allocation when each agent truthfully reports u_i, and $\mathbf{x}^1$ is the resulting allocation when each agent reports u'_i. Thus we have shown that each agent's utility is the same in these two allocations.

Since $\mathbf{u}'$ was an arbitrary DSE, we have shown that in any DSE, every agent's utility is the same as in the truthful outcome. Therefore the outcome of any DSE is a max-min-optimal allocation. □

The above mechanism does not Nash-implement max-min welfare

In this section, we show that our DSE implementation of max-min welfare is not a Nash implementation, i.e., there may be Nash equilibria that are not optimal. Consider an instance with n agents and n goods, where each agent i's true set of desired goods is $R_i = \{i\}$. Assume each good has supply 1. The unique max-min-optimal allocation has $x_{ii} = 1$ for all $i \in N$ and $x_{ij} = 0$ for $j \neq i$, i.e., it gives the entirety of each good to the unique agent who desires it. This results in each agent having utility 1, and thus max-min welfare of 1.

Now consider the strategy profile where each agent i reports that she desires every good, i.e., reports M. The resulting allocation will give each agent exactly $1/n$ of each good, resulting in each agent's utility (according to her true utility function) being $1/n$. We claim that this is a Nash equilibrium. Suppose an agent i reports R'_i instead of M. If $R'_i = \emptyset$, agent i receives nothing, so that cannot increase her utility. Thus let j be any good in R'_i. Since the other $n-1$ agents are also reporting that they desire good j, our mechanism would divide j evenly across all the agents, resulting in agent i receiving $x_{ij} = 1/n$. Since our mechanism gives each agent equal the same quantity of each good in their reported set, agent i does not receive more than $1/n$ of any good, so her utility is at most $1/n$ (in fact, it will be exactly $1/n$). Thus agent i cannot improve her utility by bidding some $R'_i \neq M$, so the strategy profile where each agent reports M is a Nash equilibrium.

Furthermore, the max-min welfare is $1/n$, which is actually a factor of n worse than the optimal max-min welfare of 1.

3.5.3 Maxmin welfare is Nash-implementable

Our Nash implementation of max-min welfare will use some of the same intuition from our DSE implementation. However, we will have to be careful to avoid bad equilibria like the one described in Section 3.5.2. Let H denote Mechanism 1 (which DSE-implements max-min welfare), and let $H_X(\mathbf{u})$ denote the allocation produced by H when the reported utility profile is $\mathbf{u}$.

Mechanism 2 operates as follows. First, it asks each agent to report not only her own utility function, but the utility function of *every* agent. Recall that the utility function is specified by the set of desired goods; hence, agent i reports $R_1(i) \dots R_n(i)$, where $R_k(i)$ is the set of goods that agent i says agent k desires.

Step 2 defines η_i and $\bar{N}$, which will be used later to penalize agents in a way that aligns incentives. The scalar η_i denotes the number of agents k where what agent i says that agent k wants ($R_k(i)$) conflicts what agent k says that she wants ($R_k(k)$). We will penalize agents for having large values of η_i to incentivize them to come to a consensus. The set $\bar{N}$ contains the set of agents i where for some other agent k, agent i is saying that agent k wants more goods than agent k is saying that she actually wants (i.e., $R_k(k) \subsetneq R_k(i)$). We will penalize this specific type of disagreement more strongly; the reason will become clear in the proof of Theorem 3.5.3.

Step 3 defines $\alpha_i \in [0, 1]$: $\alpha_i = 1$ represents no penalty, and $\alpha_i = 0$ represents an absolute penalty (i.e., that agent will end up with no utility). Specifically, $\alpha_i = 0$ for each agent in $\bar{N}$, and for those agents not in $\bar{N}$, a higher η_i leads to a higher penalty. Next, we use the mechanism H to compute a max-min-optimal allocation $\mathbf{x}'$ *ignoring agents in* $\bar{N}$. It is crucial that we ignore those agents when computing this allocation. Finally, we return an allocation $\mathbf{x}$ which is just $\mathbf{x}'$ with the α_i penalties applied. As usual, x_i and y_i are vectors in $\mathbb{R}^m_{\geq 0}$.

The first thing to notice is that we have solved the problem from Section 3.5.2. We claim that each agent i reporting $R_k(i) = M$ for all $k \in N$ is not longer a Nash equilibrium when agent i's true desired set is a strict subset of M. This is because if agent i shrinks $R_i(i)$ to her true subset, but reports $R_k(i) = M$ for all $k \neq i$, $\bar{N}$ will contain every agent except for her. This means that she would get all of the resources in Step 3, which clearly increases her utility. The other agents would then respond by setting $R_i(k) = R_i(i)$ so that they are no longer in $\bar{N}$, but this at least shows that $R_k(i) = M$ for all $i, k \in N$ is not a Nash equilibrium.

For the rest of this section, we will use $\tilde{\mathbf{R}} = \tilde{R}_1 \dots \tilde{R}_n$ to denote the true desired sets of goods, and $\tilde{\mathbf{u}} = \tilde{u_1} \dots \tilde{u_n}$ to denote the corresponding utility profile.

Lemma 3.5.6. *When each agent reports* $\tilde{\mathbf{R}}$*, Mechanism 2 returns a max-min-optimal allocation.*

Proof. In this case, we have $\bar{N} = \emptyset$ and $\alpha_i = 1$ for all $i \in N$. Thus the $\mathbf{u}$ used in Step 4 is the true utility profile, so H computes a max-min-optimal $\mathbf{y}$ allocation with respect to the true preferences. Since $\alpha_i = 1$ for all $i \in N$, we have $\mathbf{x} = \mathbf{y}$, so Mechanism 2 does indeed return a max-min-optimal allocation. □

Mechanism 2 A mechanism which Nash-implements max-min welfare.

1. Ask each agent i to report $R_1(i) \dots R_n(i)$, where $R_k(i) \subseteq M$ for each $k \in N$.
2. For each $i \in N$, let $\eta_i = |\{k \in N : R_k(k) \neq R_k(i)\}|$. Define the set $\bar{N}$ by $\bar{N} = \{i \in N : \exists k \in N \text{ s.t. } R_k(k) \subsetneq R_k(i)\}$.
3. For each $i \in N$, define $\alpha_i \in [0, 1]$ as follows. If $i \in \bar{N}$, then $\alpha_i = 0$. Otherwise, let $\alpha_i = 1 - \eta_i/n$.
4. Let $\mathbf{u}$ be the utility profile where the set of goods desired by agent i is $R_i(i)$ if $i \notin \bar{N}$, and is $\emptyset$ if $i \in \bar{N}$. Let $\mathbf{y} = H_X(\mathbf{u})$.
5. Return the allocation $\mathbf{x}$ where for each $i \in N$, $x_i = \alpha_i y_i$.

Lemma 3.5.7. *The strategy profile where each agent reports $\tilde{\mathbf{R}}$ is a Nash equilibrium.*

Proof. Suppose the opposite: then there exists an agent i who can report $R'_1(i) \dots R'_n(i)$ and increase her utility. When all agents report $\tilde{\mathbf{R}}$, let $\tilde{\alpha}_i$ be agent i's value of α_i, let $\tilde{\mathbf{y}}$ be the intermediate allocation produced in Step 4, and let $\tilde{\mathbf{x}}$ be the final resulting allocation. When agent i reports $R'_1(i) \dots R'_n(i)$ and all other agents report $\tilde{\mathbf{R}}$, we use α'_i, $\mathbf{y}'$, and $\mathbf{x}'$ analogously.

Thus we have assumed that $\tilde{u}_i(x'_i) > \tilde{u}_i(\tilde{x}_i)$, i.e., she is strictly happier when she deviates and reports $R'_1(i) \dots R'_n(i)$. Since $x'_i = \alpha'_i y'_i$ and $\tilde{x}_i = \tilde{\alpha}_i \tilde{y}_i$, we have $\tilde{u}_i(\alpha'_i y'_i) > \tilde{u}_i(\tilde{\alpha}_i \tilde{y}_i)$. When all agents report $\tilde{\mathbf{R}}$, all agents are in agreement, so $\tilde{\alpha}_i = 1$. Since $\alpha'_i \leq 1$, we have $\alpha'_i \leq \tilde{\alpha}_i$. Therefore we must have $\tilde{u}_i(y'_i) > \tilde{u}_i(\tilde{y}_i)$.

We claim that after this deviation, $\bar{N} = \emptyset$. If $i \in \bar{N}$, she receives zero utility, so such a deviation could not help her. Since all agents $k \neq i$ report the same thing, the only way for agent $k \neq i$ to be in $\bar{N}$ is if what agent k is reporting that agent i wants (which in this case is $\tilde{R}_i$) is a superset of what agent i is reporting that she wants (which in this case is $R'_i(i)$). If $R'_i(i) \subsetneq \tilde{R}_i$, then there is a good $j \in \tilde{R}_i$ that agent i will receive none of, i.e., $y'_{ij} = 0$. This is because the convex program in H will only allocate agent i a portion of good j if good j is in her reported set. Thus $R'_i(i) \subsetneq \tilde{R}_i$ implies that $\tilde{u}_i(y'_i) = 0$. Therefore agent i's utility cannot have improved, which is a contradiction.

Therefore after the deviation, $\bar{N} = \emptyset$. This means that $\mathbf{y}'$ is just the max-min-optimal allocation computed by Mechanism 1 for utility profile $\tilde{R}_1, \tilde{R}_2 \dots R'_i(i) \dots \tilde{R}_n$. But this implies that agent i is improving her utility for the allocation produced by Mechanism 1 by reporting $R'_i(i)$ instead of $\tilde{R}_i$. This contradicts the strategyproofness of Mechanism 1 (by Lemma 3.5.5). Thus $\tilde{u}_i(y'_i) > \tilde{u}_i(\tilde{y}_i)$ is impossible, which implies that each agent reporting $\tilde{\mathbf{R}}$ is in fact a Nash equilibrium. □

Theorem 3.5.3. *Mechanism 2 Nash-implements max-min welfare.*

Proof. We need to show that for each problem instance, Mechanism 2 has at least one Nash equilibrium, and every Nash equilibrium is optimal. Lemma 3.5.7 implies that at least one Nash equilibrium exists, so it remains to show that every Nash equilibrium results in a max-min-optimal allocation.

Consider an arbitrary Nash equilibrium where agent i reports $R_1(i) \dots R_n(i)$. First, note that each agent i can always achieve $i \notin \bar{N}$ and $\alpha_i = 1$ by having $R_k(i) = R_k(k)$ for each $k \neq i$. This also does not restrict what she reports for $R_i(i)$, which is what actually affects the allocation $\mathbf{y}$. Thus in any Nash equilibrium, we have $\bar{N} = \emptyset$, $\alpha_i = 1$ for all $i \in N$, and $R_k(i) = R_k(k)$ for all $i, k \in N$.

Since H allocates a portion of good j to agent i only if j is in agent i's reported set, if $R_i(i) \subsetneq \tilde{R}_i$, then agent i will receive zero utility. Reporting $R_i(i) = \tilde{R}_i$ instead (and still reporting $R_k(i) = R_k(k)$ for $k \neq i$, so that she is not in $\bar{N}$) would give her nonzero utility, so $R_i(i) \subsetneq \tilde{R}_i$ is impossible in a Nash equilibrium. Thus we must have $\tilde{R}_i \subseteq R_i(i)$ for all $i \in N$.

Now suppose that $\tilde{R}_i \subsetneq R_i(i)$. Suppose that agent i reports $R'_i(i) = \tilde{R}_i$ instead (and reports the same $R_k(i)$ for each $k \neq i$). Then since every agent $k \neq i$ is reporting $R_i(k) = R_i(i)$, we now have $R'_i(i) \subsetneq R_i(k)$ for each $k \neq i$. This means that $\bar{N}$ contains every agent except for i, so every agent other than i is ignored when computing the allocation $\mathbf{y}$. This means that the allocation $\mathbf{y}$ gives agent i her maximum possible utility (which is $\min_{j \in \tilde{R}_i} s_j$: the minimum supply of any good she desires). Since agent i is still reporting $R_k(i) = R_k(k)$ for all $k \neq i$, she still has $\alpha_i = 1$, which means that in the final allocation $\mathbf{x}$, she receives her maximum possible utility.

Therefore in any Nash equilibrium, for each $i \in N$, either $R_i(i) = \tilde{R}_i$ (i.e., she is reporting her true set), or $\tilde{R}_i \subsetneq R_i(i)$ and agent i is receiving her maximum possible utility. We proceed by case analysis.

Case 1: Every agent agent is reporting $R_i(i) = \tilde{R}_i$. Since we have $R_k(i) = R_k(k)$ for all $i, k \in N$, each agent must be reporting $\tilde{\mathbf{R}}$. Then by Lemma 3.5.6, we get a max-min-optimal allocation in this case.

Case 2: At least one agent i is reporting $\tilde{R}_i \subsetneq R_i(i)$ and thus is receiving her maximum possible utility of $\min_{j \in \tilde{R}_i} s_j$. Let $\gamma = \min_{j \in \tilde{R}_i} s_j$: then for any allocation $\mathbf{x}'$,

$$\min_{k \in N} u_k(x'_k) \leq u_i(x'_i) \leq \gamma$$

i.e., the value of the max-min objective can never be more than γ. Since H gives each agent the same utility (according to her reported preferences), for all $k \in N$ we have

$$\gamma = \min_{j \in R_k(k)} x_{kj} \geq \min_{j \in \tilde{R}_k} = \tilde{u}_k(x_k)$$

where the inequality is because $\tilde{R}_i \subseteq R_i(i)$. Thus we have $\min_{k \in N} \tilde{u}_k(x_k) = \gamma$. Thus $\mathbf{x}$ must be max-min-optimal.

Therefore we have shown that in either case, the Nash equilibrium must result in a max-min-optimal allocation. We conclude that Mechanism 2 Nash-implements max-min welfare. □

3.6 Conclusion and future work

In this chapter, we showed that every CES welfare function except $\rho = 1$ can be Nash-implemented by an augmented trading post mechanism. This strengthened previous results which only handled Nash welfare [30] or assumed agents did not behave strategically (Chapter 2). Next, we showed that DSE implementation for this problem is generally impossible, with the exception of max-min welfare, where a simple revelation mechanism does indeed DSE-implement max-min welfare. Although this revelation mechanism does not Nash-implement max-min welfare, we were able to Nash-implement max-min welfare with a different mechanism.

We were not able to resolve whether utilitarian welfare is Nash-implementable for bandwidth allocation. Our trading post mechanism breaks down in this setting, since $f_j(b) = b^{1-1} = 1$ is not a valid constraint curve. Maskin's monotonicity approach is not viable either, since utilitarian welfare does not satisfy no veto power. We leave this as an open question.

Another interesting direction would be to extend these results to a wider range of utility functions. Our reduction between price curves and trading post means that if price curve equilibria maximizing CES welfare were shown to exist for a wider range of utility functions, it seems likely that our Nash implementation results would carry over as well (depending on the form of the price curves).

It would also be interesting to consider another dimension of strategic behavior by allowing agents to choose which path in the network to use. In this case, we could write each agent's utility function as $u_i(x_i) = \max_{p \in P_i} \min_{j \in p} x_{ij}$, where P_i is the set of paths from agent i's desired source to desired destination. This is reminiscent of routing games, in that agents are strategically choosing their paths, but still distinct, in that each agent may use the same link in different quantities (i.e., receive different amounts of bandwidth). Although this model is less accurate in terms of how the internet actually works (see Section 3.1), it may be an appropriate model for other situations.

More broadly, we feel that trading post is a powerful mechanism that is able to simulate a price-taking market while also handling strategic behavior. We wonder if trading post, or variants thereof, may be useful in designing mechanisms for other resource allocation problems as well.

Chapter 4

Counteracting inequality in markets via convex pricing

In this chapter, we study the connection between convex pricing and CES welfare in the quasilinear market model. For linear pricing, the First Welfare Theorem states that Walrasian equilibria[1] maximize the sum of agent valuations. This ensures efficiency, but can lead to extreme inequality across individuals. Many real-world markets – especially for water – use *convex* pricing instead, often known as increasing block tariffs (IBTs). IBTs are thought to promote equality, but there is a dearth of theoretical support for this claim.

In this chapter, we study a simple convex pricing rule and show that the resulting equilibria are guaranteed to maximize a CES welfare function. Furthermore, a parameter of the pricing rule directly determines which CES welfare function is implemented; by tweaking this parameter, the social planner can precisely control the tradeoff between equality and efficiency. Our result holds for any valuations that are homogeneous, differentiable, and concave. We also give an iterative algorithm for computing these pricing rules, derive a truthful mechanism for the case of a single good, and discuss Sybil attacks.

4.1 Introduction

Recall that in the quasilinear model, an agent's utility is her value for the resources she obtains (her *valuation*), minus the money she spends (her *payment*). The First Welfare Theorem states that in this setting, the linear-pricing Walrasian equilibria are exactly the allocations maximizing utilitarian welfare, i.e., the sum of agent valuations. Thus linear pricing *implements* utilitarian welfare in Walrasian equilibrium (sometimes abbreviated "WE"). The result is powerful, but also limiting. Maximizing utilitarian welfare yields the most efficient outcome, but may also cause maximal inequality (see Figure 4.1).

[1] The terms "market equilbria" and "Walrasian equilibria" are equivalent. In this section, we use "Walrasian equilibrium" because that is the more common term in the quasilinear model literature.

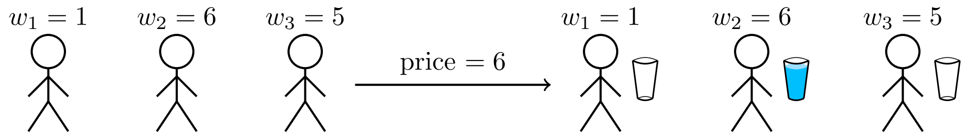

Figure 4.1: An example of how linear pricing can lead to maximal inequality. Consider the three agents above and a single good (say, water), where each agent i's value for x units of the good is $w_i \cdot x$. The unique linear-pricing Walrasian equilibrium sets a price of 6 per unit, which results in agent 2 buying all of the good and the other two agents receiving nothing. More generally, the equilibrium price reflects the maximum anyone is willing to pay, and anyone who is not willing to pay that much is priced out of the market and receives nothing. In contrast, our nonlinear pricing rule always ensures that everyone receives a nonzero amount; see Section 4.2.

One common alternative is *convex* pricing. In this chapter, we study convex pricing rules p of the form

$$p(x_i) = \Big(\sum_j q_j x_{ij}\Big)^{1/\rho}$$

where x_i is *bundle* agent i receives, $x_{ij} \in \mathbb{R}_{\geq 0}$ is the fraction of good j she receives, $q_1, \dots, q_m$ are constants, and $\rho \in (0, 1]$ determines the curvature of the pricing rule. Like linear pricing, p is still *anonymous*, meaning that agents' payments depend only on their purchases (and not on their preferences, for example).

When $\rho = 1$, p reduces to linear pricing. When $\rho < 1$, p is strictly convex, meaning that doubling one's consumption will more than double the price. This will make it easy to buy a small amount, but hard to buy a large amount, which intuitively should lead to a more equal distribution of resources. As the curvature of the pricing rule grows, this effect should be amplified, leading to a different equality/efficiency tradeoff.

Our work seeks to formalize that claim. We will show that the Walrasian equilibria of these convex pricing rules are guaranteed to maximize a *constant elasticity of substitution* (CES) welfare function, where the choice of ρ determines the specific welfare function and thus the precise equality/efficiency tradeoff (Theorem 4.4.1). Our result holds for a wide range of agent valuations.

Convex pricing in the real world. Convex pricing is especially pervasive in the water sector, where such pricing rules are known as *increasing block tariffs* (IBTs) [175], typically implemented with discrete blocks of water (hence the name). IBTs have been implemented and empirically studied in Israel [12], South Africa [38], Spain [86], Jordan [112], and the United States [152], among many other countries.

IBTs are often claimed to promote equality in water access [175], but there has been limited theoretical evidence supporting this (see [127] for one of the only examples). On the other hand, a common concern is that IBTs may lead to poor "economic efficiency" [19, 127]. Our work shows that at least on a theoretical level, convexity of pricing does not necessarily lead to inefficiency: it simply maximizes a different welfare function than the traditional utilitarian one. In particular, it maximizes a CES welfare function.

The Second Welfare Theorem and personalized pricing. The Second Welfare Theorem is perhaps the most famous theoretical result regarding implementation in Walrasian equilibrium. It states any Pareto optimum can be a WE when an arbitrary redistribution of initial wealth is allowed.[2] Another method that achieves the same goal is *personalized pricing*, where different agents can be charged different (linear) prices [82]. In contrast, convex pricing is anonymous: agents purchasing the same bundle always pay the same price.

Each of these approaches has its own pros and cons, and our goal in this chapter is not to claim that convex pricing is "better" than any other approach (or vice versa). Regardless of which approach is "best" in any given situation, convex pricing *is* widely used in practice, and is often claimed to promote equality. Our goal in this chapter is to formally quantify that claim.

More broadly, our work can be thought of as weaving together the previously disjoint threads of CES welfare and convex pricing to provide theoretical support for the oft-cited but rarely quantified claim that IBTs promote equality.

4.2 Results and related work

Main result: convex pricing implements CES welfare maximization in Walrasian equilibrium. Our main result is that for convex pricing of the form $p(x_i) = (\sum_j q_j x_{ij})^{1/\rho}$ for any $\rho \in (0, 1]$[3], a Walrasian equilibrium is guaranteed to exist, and every WE maximizes CES welfare with respect to ρ. This holds for a wide range of agent valuations.

Theorem 4.4.1 (Simplified version)**.** *Assume each valuation is homogeneous of degree r,[4] differentiable, and concave, and fix $\rho \in (0, 1]$. Then an allocation $\mathbf{x} = (x_1, \dots, x_n)$ maximizes CES welfare if and only if there exist constants $q_1, \dots, q_m \in \mathbb{R}_{\geq 0}$ such that for the pricing rule*

$$p(x_i) = \Big(\sum_j q_j x_{ij}\Big)^{1/\rho},$$

$\mathbf{x}$ and p form a WE.

Note that the ρ in $p(x_i)$ is the same ρ for which CES welfare is maximized.

We call the reader's attention to two important aspects of this result. Perhaps most importantly, our result is not simply a reformulation of the First Welfare Theorem: although maximizing CES welfare for valuations $v_1, \dots, v_n$ is equivalent to maximizing utilitarian welfare for valuations $v_1^\rho, \dots, v_n^\rho$, the First Welfare Theorem does not say anything about the agent demands in response

[2] Specifically, for any Pareto optimal allocation, there exists a redistribution of initial wealth which makes that allocation a WE. However, our quasilinear utility model does not have a concept of initial wealth (alternatively, initial wealth is simply an additive constant in agents' utilities which does not affect their behavior), so this result is not as mathematically relevant. See Section 4.2.1 for additional discussion.

[3] The case of $\rho < 0$ is slightly unintuitive, as it can result in agents who care more receiving *less* of the good. Consequently, implementation in WE is impossible; see Theorem 4.8.3.

[4] A valuation is homogeneous of degree r if scaling any bundle by a constant c scales the resulting value by c^r.

to this convex pricing rule. The First Welfare Theorem also does not help with identifying the exact conditions under which Theorem 4.4.1 holds, e.g., homogeneity of valuations.[5]

Secondly, the class of homogeneous, differentiable, and concave valuations is quite large: it generalizes most of the commonly studied valuations, e.g., linear, Cobb-Douglas, and CES (note that here we are referring to CES agent valuations, not CES welfare functions). Although Leontief valuations are not differentiable, we handle them as a special case and show that the same result holds (Theorem 4.10.1).

The following additional properties are of note:

1. For this class of utilities, Theorem 4.4.1 generalizes the First Welfare Theorem:[6] when $\rho = 1$, $p(x_i)$ yields linear pricing and CES welfare yields utilitarian welfare.

2. The constants $q_1, \dots, q_m$ will be the optimal Lagrange multipliers for a convex program maximizing CES welfare. This connection to duality will be very helpful for computing these WE (see Section 4.5).

3. Our pricing rule is strictly convex for $\rho < 1$, with the curvature growing as ρ goes to 0. The smaller ρ gets, the easier it is to buy a small amount, but the harder it is to buy a large amount. Intuitively, this should prevent any single individual from dominating the market and lead to a more equitable outcome. Furthermore, the marginal price at $x_i = \mathbf{0}$ is zero, which ensures that everyone ends up with a nonempty bundle (in contrast to linear pricing: see Figure 4.1). Theorem 4.4.1 provides a tight relationship between the curvature of the pricing rule and the exact equality/efficiency tradeoff.

Towards an implementation. We also prove several supporting results: in particular, regarding implementation. The WE from Theorem 4.4.1 can always be computed by asking each agent for her entire utility function, and then solving a convex program for maximizing CES welfare maximization to obtain the optimal Lagrange multipliers $q_1, \dots, q_m$. However, this is not very practical: people are generally not able to articulate a full cardinal utility function, and even if they are, doing so could require transmitting an enormous amount of information. Section 4.5 presents our first supporting result: an iterative algorithm for computing the WE, where in each step, each agent only needs to report the gradient of her valuation at the current point. Our algorithm is based on the ellipsoid method, and inherits its polynomial-time convergence properties. We recognize that even valuation gradient queries may be difficult for agents to answer, and we leave the possibility of an improved implementation – in particular, a *tâtonnement*[7] – as an open question.

Truthfulness. Our second supporting result considers a different approach to implementation: *truthful* mechanisms. Walrasian equilibria are generally not truthful: agents can lie about their

[5]In fact, not only is homogeneity necessary, but homogeneity of the same degree is necessary: if we allow the degree of homogeneity to differ across agents, the result no longer holds (Theorem 4.8.2).

[6]One direction of the First Welfare Theorem (if $(\mathbf{x}, p)$ is a linear pricing WE, then $\mathbf{x}$ maximizes utilitarian welfare) holds in a much more general setting; see Section 4.11.

[7]A tâtonnement is an iterative algorithm which only asks *demand queries*, i.e., what would each agent purchase given the current prices. Demand queries may be easier to answer than the valuation gradient queries in our algorithm.

preferences to affect the equilibrium prices for their personal gain.[8] For $\rho = 1$, the Vickrey-Clarke-Groves (VCG) mechanism is known to truthfully maximize utilitarian welfare [135]. For the case of a single good and any $\rho \in (0, 1)$, we give a mechanism which truthfully maximizes CES welfare (Theorem 4.6.1). We also show that our mechanism is the unique truthful mechanism up to an additive constant in the payment rule (Theorem 4.6.2). The proof of Theorem 4.6.2 is quite involved, and requires techniques from real analysis such as Kirszbraun's Theorem for Lipschitz extensions and the Fundamental Theorem of Lebesgue Calculus.

Negative results. We prove the following negative results. Most importantly, we show that for any $\rho \neq 1$, linear-pricing WE can have arbitrarily poor CES welfare (Theorem 4.8.1); were this not the case, perhaps it would suffice to simply use linear pricing and accept an approximation of CES welfare. Next, note that Theorem 4.4.1 requires each agent's valuation to be homogeneous with the same degree r. We show that when agents' valuations have different homogeneity degrees, there exist instances where no pricing rule can implement CES welfare maximization in WE (Theorem 4.8.2), and thus our assumption is necessary. We also show that CES welfare maximization cannot be implemented in WE for $\rho < 0$ (Theorem 4.8.3), and discuss the special case of $\rho = 0$ (i.e., Nash welfare).

There is an additional crucial issue which any practical implementation of Theorem 4.4.1 would need to address: *Sybil attacks*. A Sybil attack is when a selfish agent attempts to gain an advantage in a system by creating fake identities [69]. Since the pricing rule from Theorem 4.4.1 is strictly convex for $\rho < 1$, an agent can decrease her payment by masquerading as multiple individuals and splitting her purchase across those identities.[9] In Section 4.7, we propose a model for analyzing Sybil attacks in markets, and show that if these attacks are possible, there exist instances where no pricing rule can implement CES welfare maximization in WE (Theorem 4.7.3).[10]

Additional results. In Section 4.9, we explore connections between our results in the quasilinear utility model, and the Fisher market fixed-budget model. Section 4.10 shows that Theorem 4.4.1 extends to Leontief valuations, which are not differentiable (so the main proof does not apply). Leontief valuations have been a focus of prior work, so we find is worthwhile to handle this as a special case.

4.2.1 Related work

The study of markets has a long history in economics [6, 24, 80, 169, 173]. Recently, this topic has received substantial attention in the computer science community as well (see [170] for an

[8] Another interpretation is that WE assumes agents are *price-taking* (i.e., treat the prices are given and do not lie about their preferences to affect the equilibrium prices) and breaks down when agents are *price-anticipating*.

[9] In contrast, for $\rho = 1$, there is nothing to be gained by creating fake identities.

[10] There are combinations of parameters, however, where our pricing rule is naturally robust to Sybil attacks: in particular, when $v_i(\mathbf{x})(1-\rho) \leq \kappa$ (where $v_i(\mathbf{x})$ is agent i's value for the maximum CES welfare allocation and κ is the identity creation cost). This suggests a natural way for an equality-focused social planner to choose a specific value for ρ: estimate the identity creation cost and scale of valuations in the system of interest, and pick ρ to be as small as possible without incentivizing Sybil attacks.

algorithmic introduction). We first provide some important background on the First and Second Welfare Theorems, and then move on to more recent related work.

The First and Second Welfare Theorems. Conceptually, the First Welfare Theorem establishes an efficiency property that any WE must satisfy, and the Second Welfare Theorem deals with implementing a wide range of allocations as WE. The two welfare theorems originate in the context of *Arrow-Debreu* markets [6], which generalize Fisher markets to allow for (1) agents to enter the market with goods (as opposed to just money)[11] and (2) production of goods. The statements of the First and Second Welfare Theorems in that model are, respectively, "Any (linear pricing) WE is Pareto optimal" and "Any Pareto optimal allocation can be a (linear pricing) WE with *transfers*, i.e., under a suitable redistribution of initial wealth".

In the Fisher market and quasilinear utility models, the First Welfare Theorem can be strengthened to "Any (linear pricing) WE maximizes budget-weighted Nash welfare" [72, 73, 170] and "Any (linear pricing) WE maximizes utilitarian welfare", respectively. The version of the Second Welfare Theorem stated above is appropriate for Fisher markets, since agents' budgets constitute the "initial wealth". However, for quasilinear utilities, there is no notion of initial wealth (alternatively, initial wealth is an additive constant in agents' utilities which does not affect their behavior). Thus for quasilinear utilities, allowing transfers actually does not affect the set of WE. This may seem counterintuitive, since the Second Welfare Theorem (which still holds in this setting) states that any Pareto optimum can be a WE. However, Pareto optimality here is referring to agents' overall quasilinear utilities, *not* the agents' valuations. It can be shown that the only allocations which are Pareto optimal with respect to the quasilinear utilities are allocations maximizing utilitarian welfare, which are already covered by the First Welfare Theorem (without transfers).

Thus on a technical level, the Second Welfare Theorem is not helpful in the world of quasilinear utilities. However, even when the Second Welfare Theorem is mathematically relevant, a centrally mandated redistribution of wealth is often out of the question in practice.

CES welfare and the equality/efficiency tradeoff in healthcare. CES welfare have seen substantial use in healthcare under the name of *isoelastic welfare functions*. This began with [171], largely motivated by concerns abeout purely utilitarian approaches to healthcare (i.e., allocating resources to maximize total health in a community, without concern for equality). Since these decisions can affect who lives and who dies, significant effort has been invested into understanding the equality/efficiency tradeoff, with this class of welfare functions serving as a theoretical tool [67, 137, 171]; see Section 4.2.1 for additional discussion.

There have also been several empirical studies aiming to understand the general population's view of the equality/efficiency tradeoff, with results generally indicating a disapproval of purely utilitarian approaches to healthcare [68, 176]. For example, a survey of 449 Swedish politicians found widespread rejection of purely utilitarian decision-making in healthcare, and under some conditions, the respondents were willing to sacrifice up to 15 of 100 preventable deaths in order to ensure equality across subgroups [118].

[11] These are known as "exchange markets" or "exchange economies".

CES welfare and α-fairness in networking. CES welfare functions have also enjoyed considerable attention from the field of networking, under the name of *α-fairness* (the parameter α corresponds to $1-\rho$ in our definition). The α-fairness notion was proposed by [126], motivated in part as a generalization of the prominent *proportional fairness* objective (which is equivalent to Nash welfare) [110]. See [17] and references therein for further background on α-fairness in networking. To our knowledge, a market-based understanding was developed only for proportional fairness, starting with the seminal work of Kelly et al. [110].

The rest of the chapter is organized as follows. Section 4.3 formally defines the model. In Section 4.4, we present our main result: a simple convex pricing rule implements CES welfare maximization in WE for $\rho \in (0,1]$ (Theorem 4.4.1). Section 4.5 presents an iterative algorithm for computing these WE. In Section 4.6, we consider truthful mechanisms for CES welfare maximization. Section 4.7 discusses Sybil attacks, and Section 4.8 presents our negative results. Section 4.9 discusses connections to Fisher markets, Section 4.10 shows that our main result extends to Leontief valuations, Section 4.11 discusses the First Welfare Theorem in more detail, and Section 4.12 provides some proofs omitted from earlier sections.

4.3 Model

We use the same basic terminology and notation as the previous two chapters (which was defined in Chapter 1). We continue to focus on divisible goods, where x_{ij} can be any real number. In this chapter, we normalize the supply of each good to be 1 without loss of generality.

This chapter assumes quasilinear utility: $u_i(x_i) = v_i(x_i) - p_i$ where p_i is the payment charged to agent i. When each agent's payment only depends on the bundle she receives, i.e., $p_i = p(x_i)$, we call p a *pricing rule*. With the exception of Section 4.6, we will focus on pricing rules. For v_i, we make the following standard assumptions throughout the chapter:

1. Nonzero: There exists a bundle x_i such that $v_i(x_i) > 0$.

2. Montone: If $x_{ij} \geq y_{ij}$ for all $j \in M$, then $v_i(x_i) \geq v_i(y_i)$.

3. Normalized: $v_i(0, \dots, 0) = 0$.

Our positive results require the following three additional properties, which we will mention explicitly whenever used:

4. Concave: For any bundles x_i, y_i and constant $\lambda \in [0,1]$, we have $v_i(\lambda x_i + (1-\lambda) y_i) \geq \lambda v_i(x_i) + (1-\lambda) u_i(y_i)$.

5. Homogeneous of degree r: for any bundle x_i and constant $\lambda \geq 0$, $v_i(\lambda x_i) = \lambda^r v_i(x_i)$. For $0 < r < 1$, this models diminishing returns. Note that homogeneity implies normalization, and for monotone and concave v_i, we must have $r \geq 0$ and $r \leq 1$ respectively.

6. Differentiable: for any bundle x_i and all $j \in M$, $\frac{\partial v_i(x_i)}{\partial x_{ij}}$ is defined.

Weighted CES welfare. We use the same CES welfare objective as in previous chapters, but we know also consider *weighed* CES welfare. For multipliers $\mathbf{a} = (a_1, a_2 \dots a_n) \in \mathbb{R}^n_{\geq 0}$ and $\rho \in (-\infty, 0) \cup (0, 1]$, the (weighted) CES welfare of an allocation $\mathbf{x}$ is $\Phi_{\mathbf{a}}(\rho, \mathbf{x}) = \big(\sum_{i \in N} a_i v_i(x_i)^\rho\big)^{1/\rho}$. We will use $\Psi_{\mathbf{a}}(\rho)$ to denote CES welfare maximization, i.e., $\Psi_{\mathbf{a}}(\rho) = \arg\max_{\mathbf{x} \in \mathbb{R}^m_{\geq 0}:\ \sum_i x_{ij} \leq 1\ \forall j} \Phi_{\mathbf{a}}(\rho, \mathbf{x})$. There may be multiple optimal allocations (for example, if there is a good which no one values), so $\Psi_{\mathbf{a}}(\rho)$ denotes a set. Thus $\mathbf{x} \in \Psi_{\mathbf{a}}(\rho)$ denotes that $\mathbf{x}$ has maximum CES welfare. When each agent has the same multiplier (other than Section 4.9, this will always be the case), we simply write $\Phi(\rho, \mathbf{x})$ and $\Psi(\rho)$.

As an illustrative example, consider a single good and valuations that are homogeneous of degree 1. Utilitarian welfare results in the good being entirely allocated to agents with $w_i = \max_k w_k$, with other agents receiving nothing (see Figure 4.1). In contrast, for $\rho < 1$, the unique allocation maximum CES welfare welfare gives the following bundle $x_i \in \mathbb{R}_{>0}$ to each agent i (Lemma 4.6.2): $x_i = \frac{w_i^{\frac{\rho}{1-\rho}}}{\sum_k w_k^{\frac{\rho}{1-\rho}}}$. One natural case is $\rho = 1/2$, which results in a proportional allocation.

4.4 Main result

We begin with our main result: for a wide range of valuations and any $\rho \in (0, 1]$, a simple convex pricing rule leads to CES welfare maximization in Walrasian equilibrium. Our pricing rule has many additional interesting properties; to avoid redundancy, we refer the reader back to our discussion in Section 4.2. On a high level, the proof relies on the KKT conditions for CES welfare maximization and the KKT conditions for each agent's demand set, and uses Euler's Theorem for homogeneous functions to conjoin the two. This will result in the following theorem:

Theorem 4.4.1. *Assume each v_i is homogeneous of degree r, concave, and differentiable. For any $\rho \in (0, 1]$ and any allocation $\mathbf{x}$, we have $\mathbf{x} \in \Psi(\rho)$ if and only if there exist $q_1, \dots, q_m \in \mathbb{R}_{\geq 0}$ such that for the pricing rule*

$$p(x_i) = \rho r^{\frac{\rho - 1}{\rho}} \Big(\sum_{j \in M} q_j x_{ij}\Big)^{1/\rho},$$

$(\mathbf{x}, p)$ is a WE. Furthermore, $q_1, \dots, q_m$ are optimal Lagrange multipliers for Program 4.1.

4.4.1 Proof setup

We begin by setting up the two relevant convex programs and proving several lemmas. For valuations $v_1 \dots v_n$, nonnegative multipliers $\mathbf{a} = a_1 \dots a_n$, and $\rho \in (-\infty, 0) \cup (0, 1]$, consider the following nonlinear program for maximizing CES welfare:

$$\begin{aligned} \max_{\mathbf{x} \in \mathbb{R}^{n \times m}_{\geq 0}} \quad & \frac{1}{\rho} \sum_{i \in N} a_i v_i(x_i)^\rho \\ s.t. \quad & \sum_{i \in N} x_{ij} \leq 1 \qquad \forall j \in M \end{aligned} \tag{4.1}$$

Since the constraints are linear and the objective function is concave (since $\rho \leq 1$), Program 4.1 is a convex program. Program 4.1 depends on ρ, but we will leave this implicit when clear from context: we will simply say "$\mathbf{x}$ is optimal for Program 4.1" as opposed to "$\mathbf{x}$ is optimal for Program 4.1 with respect to ρ". Note also that we are maximizing $\frac{1}{\rho}\sum_{i\in N} a_i v_i(x_i)^\rho$ instead of the true CES welfare $\Phi_{\mathbf{a}}(\rho, \mathbf{x}) = (\sum_{i\in N} a_i v_i(x_i)^\rho)^{1/\rho}$; this will lead to the same optimal allocation $\mathbf{x}$ and will simplify the analysis.

When $\mathbf{a}$ is not specified, we assume that $\mathbf{a} = \mathbf{1}$. Nonuniform multipliers will only be used in Section 4.9 when we consider connections to Fisher markets, but we include them here for completeness.

Next, consider each agent's demand set given a pricing rule p:

$$D_i(p) = \underset{x_i \in \mathbb{R}^m_{\geq 0}}{\arg\max} \ \big(v_i(x_i) - p(x_i)\big) \tag{4.2}$$

When p is convex (as in Theorem 4.4.1), $-p$ is concave. Since v_i is also concave, $v_i(x_i) - p(x_i)$ is concave, so each agent's demand set defines a convex program (Program 4.2). Program 4.2 depends on i, the agent in question, but again we leave this implicit when it is clear from context.

We will also use the following theorem, due to Euler. We include a short proof in Section 4.12.[12]

Theorem 4.4.2 (Euler's Theorem for homogeneous functions)**.** *Let $f: \mathbb{R}^m_{\geq 0} \to \mathbb{R}$ be differentiable and homogenous of degree r. Then for any $\mathbf{b} = (b_1, \dots b_m) \in \mathbb{R}^m_{\geq 0}$, $\sum_{j=1}^m b_j \dfrac{\partial f(\mathbf{b})}{\partial b_j} = rf(\mathbf{b})$.*

Before we state and prove Theorem 4.4.1, we note one other property: for a pricing rule of the form $p(x_i) = c(\sum_{j\in M} q_j x_{ij})^{1/\rho}$ where $c > 0$, good j has nonzero cost (for the purposes of Walrasian equilibrium) if and only if $q_j = 0$.

4.4.2 Proof of Theorem 4.4.1

The proof of Theorem 4.4.1 is divided into three parts. The first part involves setting up the KKT conditions for Programs 4.1 and 4.2. The second assumes that $\mathbf{x} \in \Psi(\rho)$ and proves that $(\mathbf{x}, p)$ is a WE, and the third assumes that $(\mathbf{x}, p)$ is a WE and proves that $\mathbf{x} \in \Psi(\rho)$.

Proof of Theorem 4.4.1. **Part 1: Setup.** Let $\mathbf{q}$ denote the vector $(q_1, \dots, q_m) \in \mathbb{R}^m_{\geq 0}$; then the Lagrangian of Program 4.1 is $L(\mathbf{x}, \mathbf{q}) = \frac{1}{\rho}\sum_{i\in N} v_i(x_i)^\rho - \sum_{j\in M} q_j(\sum_{i\in N} x_{ij} - 1)$.[13] Since Program 4.1 is convex and satisfies strong duality by Slater's condition, the KKT conditions are both necessary and sufficient for optimality. That is, $\mathbf{x}$ is optimal for Program 4.1 (which is equivalent to $\mathbf{x} \in \Psi(\rho)$) if and only if there exist Lagrange multipliers $\mathbf{q} \in \mathbb{R}^m_{\geq 0}$ such that both of the following hold:[14]

[12] The reason we provide a proof is that this theorem is often stated with the requirement of continuous differentiability, but in fact only requires differentiability; to avoid any confusion, we provide a proof only using differentiability.

[13] The expert reader may notice that we have omitted the $\mathbf{x} \in \mathbb{R}^{m\times n}_{\geq 0}$ constraint from the Lagrangian. We do this to slightly simplify the analysis. The effect on the KKT conditions is that stationarity changes from "For all i, j, $\frac{\partial L(\mathbf{x},\mathbf{q})}{\partial x_{ij}} = 0$" to "For all i, j, $\frac{\partial L(\mathbf{x},\mathbf{q})}{\partial x_{ij}} \leq 0$, and the inequality holds with equality when $x_{ij} > 0$".

[14] The KKT conditions also include primal feasibility and dual feasibility. Since we will only work with valid allocations $\mathbf{x}$ and nonnegative $q_1, \dots, q_m$, these two conditions are trivially satisfied.

1. Stationarity: $\frac{\partial L(\mathbf{x}, \mathbf{q})}{\partial x_{ij}} \le 0$ for all i, j. Furthermore, if $x_{ij} > 0$, the inequality holds with equality.

2. Complementary slackness: for all $j \in M$, either $\sum_{i \in N} x_{ij} = 1$, or $q_j = 0$.

For a given (i, j) pair, $\frac{\partial L(\mathbf{x}, \mathbf{q})}{\partial x_{ij}}$ is equal to $v_i(x_i)^{\rho-1}\frac{\partial v_i(x_i)}{\partial x_{ij}} - q_j$, so stationarity for Program 4.1 is equivalent to: $q_j \ge v_i(x_i)^{\rho-1}\frac{\partial v_i(x_i)}{\partial x_{ij}}$ for all i, j, and when $x_{ij} > 0$, the inequality holds with equality.

Next consider Program 4.2, which defines each agent's demand set. This program has no constraints (other than $x_i \in \mathbb{R}^m_{\ge 0}$), so we can ignore complementary slackness. Thus by the KKT conditions, $x_i \in D_i(p)$ if and only if for every $j \in M$, $\frac{\partial v_i(x_i)}{\partial x_{ij}} \le \frac{\partial p(x_i)}{\partial x_{ij}}$, and if $x_{ij} > 0$, the inequality holds with equality (stationarity). We can explicitly compute the partial derivatives of p: $\frac{\partial p(x_i)}{\partial x_{ij}} = r^{\frac{\rho-1}{\rho}} q_j \big(\sum_{\ell \in M} q_\ell x_{i\ell}\big)^{\frac{1-\rho}{\rho}}$.

Part 2: Optimal CES welfare implies WE. Suppose that $\mathbf{x} \in \Psi(\rho)$. Then there exists $\mathbf{q} \in \mathbb{R}^m_{\ge 0}$ such that $q_j \ge v_i(x_i)^{\rho-1}\frac{\partial v_i(x_i)}{\partial x_{ij}}$ for all j, and $q_j = v_i(x_i)^{\rho-1}\frac{\partial v_i(x_i)}{\partial x_{ij}}$ whenever $x_{ij} > 0$. Using the latter in combination with Euler's Theorem for homogeneous functions, for each (i, j) pair we have

$$\begin{aligned}
\frac{\partial p(x_i)}{\partial x_{ij}} &= r^{\frac{\rho-1}{\rho}} q_j \Big(\sum_{\ell \in M} q_\ell x_{i\ell}\Big)^{\frac{1-\rho}{\rho}} \\
&= q_j \Big(r^{-1} \sum_{\ell : x_{i\ell} > 0} q_\ell x_{i\ell}\Big)^{\frac{1-\rho}{\rho}} \\
&= q_j \Big(r^{-1} \sum_{\ell : x_{i\ell} > 0} v_i(x_i)^{\rho-1} \frac{\partial v_i(x_i)}{\partial x_{i\ell}} x_{i\ell}\Big)^{\frac{1-\rho}{\rho}} && \text{(stationarity for } x_{i\ell} \text{ when } x_{i\ell} > 0) \\
&= q_j \Big(r^{-1} v_i(x_i)^{\rho-1} \sum_{\ell \in M} \frac{\partial v_i(x_i)}{\partial x_{i\ell}} x_{i\ell}\Big)^{\frac{1-\rho}{\rho}} \\
&= q_j \big(r^{-1} v_i(x_i)^{\rho-1} r v_i(x_i)\big)^{\frac{1-\rho}{\rho}} && \text{(Euler's Theorem)} \\
&= q_j v_i(x_i)^{1-\rho}
\end{aligned}$$

Thus $\frac{\partial p(x_i)}{\partial x_{ij}} = q_j v_i(x_i)^{1-\rho}$. Next, we claim that $x_i \in D_i(p)$ for all $i \in N$. Fix an agent i; we show by case analysis that x_i satisfies stationarity (for Program 4.2) for each $j \in M$.

Case 1: $q_j = v_i(x_i)^{\rho-1}\frac{\partial v_i(x_i)}{\partial x_{ij}}$. Then $\frac{\partial p(x_i)}{\partial x_{ij}} = q_j v_i(x_i)^{1-\rho} = v_i(x_i)^{\rho-1}\frac{\partial v_i(x_i)}{\partial x_{ij}} v_i(x_i)^{1-\rho} = \frac{\partial v_i(x_i)}{\partial x_{ij}}$, and we are done.

Case 2: $x_{ij} = 0$ and $q_j \ge v_i(x_i)^{\rho-1}\frac{\partial v_i(x_i)}{\partial x_{ij}}$. Then similarly, $\frac{\partial p(x_i)}{\partial x_{ij}} = q_j v_i(x_i)^{1-\rho} \ge v_i(x_i)^{\rho-1}\frac{\partial v_i(x_i)}{\partial x_{ij}} v_i(x_i)^{1-\rho} = \frac{\partial v_i(x_i)}{\partial x_{ij}}$, and again we are done. Therefore $x_i \in D_i(p)$ for all $i \in N$.

Since $\mathbf{x}$ is a valid allocation, $\sum_{i \in N} x_{ij} \le 1$ for all $j \in M$. This, combined with complementary slackness for Program 4.1, is identical to the market clearing condition for Walrasian equilibrium. Thus we have shown that if $\mathbf{x} \in \Psi(\rho)$, there exist $q_1, \ldots, q_m$ (which are optimal Lagrange multipliers for Program 4.1) such that for pricing rule p as defined, $(\mathbf{x}, p)$ is a WE.

Part 3: WE implies optimal CES welfare. This is similar to Part 2. Suppose there exists $\mathbf{q} \in \mathbb{R}^m_{\geq 0}$ such that for pricing rule $p(x_i) = \rho r^{\frac{\rho-1}{\rho}} (\sum_{j\in M} q_j x_{ij})^{1/\rho}$, $(\mathbf{x}, p)$ is a WE. Recall the partial derivatives of p: $\frac{\partial p(x_i)}{\partial x_{ij}} = q_j \big(r^{-1} \sum_{\ell \in M} q_\ell x_{i\ell}\big)^{\frac{1-\rho}{\rho}}$. We multiply each side by $x_{ij} r^{-1}$, and sum both sides over j:

$$r^{-1} \sum_{j\in M} x_{ij} \frac{\partial p(x_i)}{\partial x_{ij}} = \Big(r^{-1} \sum_{\ell\in M} q_\ell x_{i\ell}\Big)^{\frac{1-\rho}{\rho}} r^{-1} \sum_{j \in M} q_j x_{ij}$$
$$r^{-1} \sum_{j:x_{ij}>0} x_{ij} \frac{\partial p(x_i)}{\partial x_{ij}} = \Big(r^{-1} \sum_{\ell\in M} q_\ell x_{i\ell}\Big)^{\frac{1-\rho}{\rho}+1}$$

Since $(\mathbf{x}, p)$ is a WE, we have $x_i \in D_i(p)$ for all $i \in N$. Thus $\frac{\partial v_i(x_i)}{\partial x_{ij}} = \frac{\partial p(x_i)}{\partial x_{ij}}$ whenever $x_{ij} > 0$, so

$$r^{-1} \sum_{j:x_{ij}>0} x_{ij} \frac{\partial v_i(x_i)}{\partial x_{ij}} = \Big(r^{-1} \sum_{\ell\in M} q_\ell x_{i\ell}\Big)^{1/\rho}$$
$$r^{-1} \sum_{j\in M} x_{ij} \frac{\partial v_i(x_i)}{\partial x_{ij}} = \Big(r^{-1} \sum_{\ell\in M} q_\ell x_{i\ell}\Big)^{1/\rho}$$
$$v_i(x_i) = \Big(r^{-1} \sum_{\ell\in M} q_\ell x_{i\ell}\Big)^{1/\rho} \qquad \text{(Euler's Theorem)}$$
$$v_i(x_i)^{\rho-1} = \Big(r^{-1} \sum_{\ell\in M} q_\ell x_{i\ell}\Big)^{\frac{\rho-1}{\rho}}$$

Using this in combination with $\frac{\partial p(x_i)}{\partial x_{ij}} = q_j \big(r^{-1} \sum_{\ell\in M} q_\ell x_{i\ell}\big)^{\frac{1-\rho}{\rho}}$, we get

$$q_j = \frac{\partial p(x_i)}{\partial x_{ij}} \Big(r^{-1} \sum_{\ell\in M} q_\ell x_{i\ell}\Big)^{\frac{\rho-1}{\rho}} = \frac{\partial p(x_i)}{\partial x_{ij}} v_i(x_i)^{\rho-1}$$

Next, we claim that $(\mathbf{x}, \mathbf{q})$ satisfies stationarity for Program 4.1. We proceed by case analysis for each (i, j) pair. Stationarity for Program 4.2 implies that these are the only two possible cases.

Case 1: $\frac{\partial v_i(x_i)}{\partial x_{ij}} = \frac{\partial p(x_i)}{\partial x_{ij}}$. In this case, $q_j = \frac{\partial v_i(x_i)}{\partial x_{ij}} v_i(x_i)^{\rho-1}$, and we are done.

Case 2: $x_{ij} = 0$ and $\frac{\partial v_i(x_i)}{\partial x_{ij}} \leq \frac{\partial p(x_i)}{\partial x_{ij}}$. In this case we have $q_j \geq \frac{\partial v_i(x_i)}{\partial x_{ij}} v_i(x_i)^{\rho-1}$, and we are again done.

Thus $(\mathbf{x}, \mathbf{q})$ satisfies stationarity for Program 4.1. Furthermore, the second condition of Walrasian equilibrium is again identical to the complementary slackness condition. We conclude that $\mathbf{x} \in \Psi(\rho)$, and that $q_1, \dots, q_m$ are optimal Lagrange multipliers for Program 4.1. This completes the proof. $\square$

The following corollary states that under this pricing rule, each agent's resulting value will be proportional to the her payment. This property will be helpful in future sections, and may also be interesting independently.

Corollary 4.4.1.1. *Assume each v_i is homogeneous of degree r, concave, and differentiable, and let $p(x_i) = (\sum_{j\in M} q_j x_{ij})^{1/\rho}$ for some $\mathbf{q} \in \mathbb{R}^m_{\geq 0}$. Then if $x_i \in D_i(p)$, $p(x_i) = \rho r v_i(x_i)$.*

Proof. As before, stationarity for Program 4.2 gives us $\frac{\partial v_i(x_i)}{\partial x_{ij}} = \frac{\partial p(x_i)}{\partial x_{ij}}$ whenever $x_{ij} > 0$. Also note that by definition, p is homogeneous of degree $1/\rho$. Using these two properties in combination with Euler's Theorem, we get

$$\begin{aligned} \frac{\partial v_i(x_i)}{\partial x_{ij}} &= \frac{\partial p(x_i)}{\partial x_{ij}} \quad \text{for all } j \in M \text{ s.t. } x_{ij} > 0 \\ \sum_{j\in M} x_{ij}\frac{\partial v_i(x_i)}{\partial x_{ij}} &= \sum_{j\in M} x_{ij}\frac{\partial p(x_i)}{\partial x_{ij}} \\ r v_i(x_i) &= \frac{1}{\rho} p(x_i) \end{aligned}$$

Multiplying both sides by $1/\rho$ completes the proof. □

4.5 Towards an implementation

Theorem 4.4.1 guarantees the existence of Walrasian equilibria maximizing CES welfare, but says nothing about how to find these equilibria. As discussed in Section 4.2, we could always explicitly ask each agent for her valuation, and directly solve Program 4.1. However, agents are generally not able to articulate their entire valuations, and even if they are, doing so could be extremely tedious.

In this section, we give an iterative algorithm for computing the WE given by Theorem 4.4.1. The algorithm will just compute the optimal allocation; Lemma 4.5.1 shows how the equilibrium pricing rule can easily be obtained once the optimal allocation is in hand. Our algorithm is computationally equivalent to running the general-purpose ellipsoid method on Program 4.1, i.e., it explores the exact same sequence of allocations. The key is that we are able to implement the ellipsoid method only using valuation gradient queries, i.e., "tell me the gradient of your valuation at this point". We immediately inherit the correctness and polynomial-time convergence properties of the ellipsoid algorithm. Throughout this section, we make the same assumptions as in Theorem 4.4.1: each v_i is concave, homogeneous of degree r, and differentiable.

First, recall that the pricing rule from Theorem 4.4.1 takes the form $p(x_i) = \rho r^{\frac{\rho-1}{\rho}} (\sum_{j\in M} q_j x_{ij})^{1/\rho}$. Since ρ and r are constants, it suffices to compute $\mathbf{q} = q_1, \dots, q_m$. Helpfully, Theorem 4.4.1 tells us that if $\mathbf{q}$ are optimal Lagrange multipliers for Program 4.1, then $(\mathbf{x}, p)$ is a WE for any $\mathbf{x} \in \Psi(\rho)$. The next lemma states if we know an $\mathbf{x} \in \Psi(\rho)$, and have access to the gradients of the agents' valuations at $\mathbf{x}$, we can determine optimal Lagrange multipliers.

Lemma 4.5.1. *Let $\mathbf{x} \in \Psi(\rho)$. Then we can determine optimal Lagrange multipliers $\mathbf{q}$ using only $\nabla v_1(x_1), \dots, \nabla v_n(x_n)$.*

Proof. First, using Euler's Theorem for homogeneous functions (Theorem 4.4.2), we can obtain $v_1(x_1), \dots, v_n(x_n)$ using only $\nabla v_1(x_1), \dots, \nabla v_n(x_n)$. Next, since $\mathbf{x} \in \Psi(\rho)$, the KKT conditions for

Program 4.1 imply that whenever $x_{ij} > 0$, $q_j = v_i(x_i)^{\rho-1}\frac{\partial v_i(x_i)}{\partial x_{ij}}$. Fix a $j \in M$. If $x_{ij} > 0$ for some agent i, then $q_j = v_i(x_i)^{\rho-1}\frac{\partial v_i(x_i)}{\partial x_{ij}}$. We know all the values on the right hand side, so we can compute q_j. if $x_{ij} = 0$ for all $i \in N$, then complementary slackness implies that $q_j = 0$. ☐

Thus it suffices to find an allocation $\mathbf{x} \in \Psi(\rho)$, which is equivalent to finding an $\mathbf{x}$ that is optimal for Program 4.1. There are many iterative algorithms for solving convex programs of this form. Furthermore, many only require (1) oracle access to the objective function and its gradient, and (2) *a separation oracle*[15] for the constraint set (and no additional assumptions of strong convexity or other properties). For the sake of specificity, we focus on the *ellipsoid method* [33], but any algorithm with these properties is sufficient for our purposes.

Lemma 4.5.2 ([33])**.** *Let f be a convex function and let $\mathcal{X}$ be a convex set. Consider the program $\min_{x\in\mathcal{X}} f(x)$. Let $\mathcal{E}$ be a ball containing the minimum of f, and suppose there exists a polynomial-time separation oracle for $\mathcal{X}$. Then the ellipsoid method starting from $\mathcal{E}$ requires only oracle access to f and ∇f, and converges to the minimum of f in polynomial time.*

In our case, we have a trivial polynomial-time separation oracle: simply check each constraint to see if it is violated. For the gradient of our objective function, we have $\frac{\partial}{\partial x_{ij}}\Big(\frac{1}{\rho}\sum_{i\in N} v_i(x_i)^\rho\Big) = \frac{\partial v_i(x_i)}{\partial x_{ij}} v_i(x_i)^{\rho-1}$. By Euler's Theorem for homogeneous functions (Theorem 4.4.2), we have

$$\frac{\partial v_i(x_i)}{\partial x_{ij}} v_i(x_i)^{\rho-1} = \frac{\partial v_i(x_i)}{x_{ij}} \Big(r^{-1}\sum_{\ell\in M} x_{i\ell}\frac{\partial v_i(x_i)}{x_{i\ell}}\Big)^{\rho-1} \tag{4.3}$$

Similarly,

$$\frac{1}{\rho}\sum_{i\in N} v_i(x_i)^\rho = \frac{1}{\rho}\sum_{i\in N}\Big(r^{-1}\sum_{j\in M} x_{ij}\frac{\partial v_i(x_i)}{\partial x_{ij}}\Big)^{\rho} \tag{4.4}$$

Therefore for any allocation $\mathbf{x}$, we can compute both the objective function value and the gradient of the objective function using only the gradients of $v_1, \dots, v_n$. The final ingredient we need is an initial ball guaranteed to contain the optimum; we can simply enclose the entire feasible region in a ball of constant radius.

Thus we get the following iterative algorithm for computing the equilibrium pricing rule:

1. Run the ellipsoid method (or any other suitable convex optimization algorithm) to solve Program 4.1.

2. At the start of each iteration, ask each agent i for the gradient of v_i at the current point $\mathbf{x}$.

3. Whenever the algorithm requires the gradient of the objective function at $\mathbf{x}$, compute it via Equation 4.3.

[15] A separation oracle is an algorithm which, given a point x and a convex set $\mathcal{X}$, determines whether $x \in \mathcal{X}$. If $x \notin \mathcal{X}$, it must return a separating hyperplane (if $\mathcal{X}$ is specified by a set of constraints, returning a violated constraint is sufficient).

4. Whenever the algorithm requires the objective function value at $\mathbf{x}$, compute it via Equation 4.4.

Lemma 4.5.2 immediately implies correctness and polynomial-time convergence.

4.5.1 Eliciting the gradients of valuations

The above algorithm (as well as Lemma 4.5.1) requires us to have access the gradients of agents' valuations. We could simply ask each agent for this information explicitly; depending on the application domain, this may or may not be reasonable. An alternative approach is to relate $\nabla v_i(x_i)$ to agent i's willingness to pay. For example, consider the following query to agent i: "Suppose you have already bought the bundle x_i. What is the smallest marginal price for good j such that you would not buy more of good j?" The KKT conditions for agent i's demand set imply that the answer to this question is exactly $\frac{\partial v_i(x_i)}{\partial x_{ij}}$ (assuming that the agent would not buy more if she is indifferent).

Such a query could be implemented in a variety of ways. One possibility would be gradually increasing the hypothetical marginal price of good j in a continuous fashion, and asking agent i to say "stop" when she would no longer buy more of good j (in a "moving-knife"-like fashion). Also, rather asking agents about hypothetical marginal prices, one could build the necessary marginal prices into an actual pricing rule, e.g., even one as simple as $p(x_i) = \sum_{j \in M} c_j x_{ij}$. The choice of implementation would depend heavily on the specific problem setting; our point here is that there are a variety of ways to elicit $\nabla v_i(x_i)$ via queries about what agent i would purchase in different (hypothetical) situations.

4.6 Truthfulness

An alternative approach to implementation is via *truthful* mechanisms. Walrasian equilibria are generally not truthful: agents can sometimes create more favorable equilibrium prices by lying about their preferences. In this section, we present a truthful mechanism for optimizing CES welfare in the case of a single good (Theorem 4.6.1), and show that it is unique up to additive constants in the payment rule (Theorem 4.6.2). Note that uniqueness beyond additive constants in the payment rule can never be achieved without additional assumptions (e.g., individual rationality), since such constants do not affect the behavior of agents.

Before formally stating and proving these results, we mention an important distinction between this section and Section 4.5. Section 4.5 is an implementation of the WE from Theorem 4.4.1 (which we know maximizes CES welfare). In contrast, the truthful mechanism from this section is an implementation of CES welfare maximization directly, not an implementation of the WE from Theorem 4.4.1. Indeed, we know that the payment rule from Theorem 4.4.1 is not truthful, so we must consider a different payment rule if we desire truthfulness.

To define our truthful mechanism we need the following two lemmas, whose proofs appear in Appendix 4.12. The first states that for a single good, homogeneous and differentiable functions take a very simple form. The second states that for a single good, the maximum CES welfare allocations take a very simple form; also, for $\rho \neq 1$, the optimum is unique.

Lemma 4.6.1. *Let $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be differentiable and homogeneous of degree r. Then there exists $c \in \mathbb{R}_{\geq 0}$ such that $f(x) = cx^r$.*

Lemma 4.6.2. *Let $m = 1$ and $v_i(x_i) = w_i x_i^r$ for all $i \in N$ where $r \in (0, 1]$. Then $\rho \in (0, 1]$ and $r\rho \neq 1$, $\mathbf{x} \in \Psi(\rho)$ if and only if*

$$x_i = \frac{{w_i}^{\frac{\rho}{1-r\rho}}}{\sum_{k \in N} {w_k}^{\frac{\rho}{1-r\rho}}}$$

If $\mathbf{x} \in \Psi(\rho)$ and $r = \rho = 1$, then whenever $x_i > 0$, $w_i = \max_{k \in N} w_k$.

We now define our mechanism. For $\rho = 1$, the VCG mechanism truthfully maximizes utilitarian welfare [135], so assume $\rho \in (0, 1)$. We ask each agent i to report w_i (where $v_i(x_i) = w_i \cdot x_i^r$), assume the w_i's are truthful, and output the (unique) optimal allocation $\mathbf{x} \in \Psi(\rho)$ according to Lemma 4.6.2. Let $\mathbf{b} = b_1, \dots, b_n$ be the vector of reported w_i's. We then charge each agent i the following payment:[16]

$$p_i(\mathbf{b}) = \frac{r\rho}{1 - r\rho} \Big(\sum_{k \neq i} b_k^{\frac{\rho}{1-r\rho}} \Big) \int_{b=0}^{b_i} \frac{b^{\frac{r\rho}{1-r\rho}}}{\big(b^{\frac{\rho}{1-r\rho}} + \sum_{k \neq i} b_k^{\frac{\rho}{1-r\rho}}\big)^{r+1}} \, \mathrm{d}b \tag{4.5}$$

This payment is chosen so that the derivative of agent i's utility at $b_i = w_i$ is 0. In particular, let $x_i(\mathbf{b})$ denote agent i's bundle under reports $\mathbf{b}$. Then we will have $\frac{\partial v_i(x_i(\mathbf{b}))}{\partial b_i} = r w_i \frac{\partial x_i(\mathbf{b})}{\partial b_i} x_i(\mathbf{b})^{r-1}$, and $\frac{\partial p_i(\mathbf{b})}{\partial b_i} = r b_i \frac{\partial x_i(\mathbf{b})}{\partial b_i} x_i(\mathbf{b})^{r-1}$, so the derivative of agent i's overall utility will be $\frac{\partial u_i(\mathbf{b})}{\partial b_i} = r(w_i - b_i) \frac{\partial x_i(\mathbf{b})}{\partial b_i} x_i(\mathbf{b})^{r-1}$. This will imply that it is optimal for agent i to truthfully report $b_i = w_i$.

Theorem 4.6.1. *Assume $m = 1$, and that each v_i is homogenous of degree r (with r publicly known), concave, and differentiable. Then for all $\rho \in (0, 1)$, there is a truthful mechanism which outputs an allocation $\mathbf{x} \in \Psi(\rho)$.*

Proof. Since VCG satisfies the claim for $\rho = 1$, assume $\rho \in (0, 1)$. Let $\mathbf{x}(\mathbf{b})$ denote the allocation outputted given reports $\mathbf{b}$, and let $x_i(\mathbf{b})$ denote agent i's bundle: formally, $x_i(\mathbf{b}) = \frac{{b_i}^{\frac{\rho}{1-r\rho}}}{\sum_{k \in N} {b_k}^{\frac{\rho}{1-r\rho}}}$. Since $m = 1$, Lemma 4.6.1 implies that for all $i \in N$, there exists $w_i \in \mathbb{R}_{\geq 0}$ such that $v_i(x) = w_i \cdot x^r$ for all $x \in \mathbb{R}_{\geq 0}$. Then by Lemma 4.6.2, $x_i(\mathbf{b}) \in \Psi(\rho)$, so it remains to prove truthfulness.

Since we assume that each agent's valuation is not identically zero, we have $w_i > 0$. Also, by concavity and monotonicity of v_i, we have $r \in (0, 1]$. Thus $0 < r\rho < 1$. Since we also have $b_i > 0$, all denominators in $p_i(\mathbf{b})$ are nonzero and thus $p_i(\mathbf{b})$ is well-defined.

Let $v_i(\mathbf{b}) = v_i(x_i(\mathbf{b})) = w_i x_i(\mathbf{b})^r$ for brevity, and let $u_i(\mathbf{b}) = v_i(\mathbf{b}) - p_i(\mathbf{b})$ denote agent i's resulting utility under bids $\mathbf{b}$. Note that $x_i(\mathbf{b})$, $v_i(\mathbf{b})$, $p_i(\mathbf{b})$, and $u_i(\mathbf{b})$ are all differentiable with respect to b_i. Also let $\alpha = \frac{\rho}{1 - r\rho}$ for brevity; then $p_i(\mathbf{b}) = r\alpha(\sum_{k \neq i} b_k^\alpha) \int_{b=0}^{b_i} \frac{b^{r\alpha}}{(b^\alpha + \sum_{k \neq i} b_k^\alpha)^{r+1}} \mathrm{d}b$ and $x_i(\mathbf{b}) = \frac{b_i^\alpha}{\sum_{k \in N} b_k^\alpha}$.

[16] Although this integral does not have a simple closed form, it can be expressed via the hypergeometric function.

To prove truthfulness, we need to show that $w_i \in \arg\max_{b_i \in \mathbb{R}_{>0}} u_i(\mathbf{b})$, i.e., truthfully reporting w_i is an optimal strategy for agent i.[17] Since $u_i(\mathbf{b})$ is differentiable with respect to b_i, we have $\frac{\partial u_i(\mathbf{b})}{\partial b_i} = \frac{\partial v_i(\mathbf{b})}{\partial b_i} - \frac{\partial p_i(\mathbf{b})}{\partial b_i}$. The first term on the right hand side is

$$\frac{\partial v_i(\mathbf{b})}{\partial b_i} = r w_i \frac{\partial x_i(\mathbf{b})}{\partial b_i} x_i(\mathbf{b})^{r-1}$$

The second term is

$$\begin{aligned}
\frac{\partial p_i(\mathbf{b})}{\partial b_i} &= r\alpha\Big(\sum_{k\neq i} b_k^\alpha\Big)\frac{b_i^{r\alpha}}{(b_i^\alpha + \sum_{k\neq i} b_k^\alpha)^{r+1}} \\
&= r\alpha\Big(\sum_{k\neq i} b_k^\alpha\Big)\frac{b_i^{r\alpha}}{(\sum_{k\in N} b_k^\alpha)^{r+1}} \\
&= r\alpha\Big(\sum_{k\neq i} b_k^\alpha\Big)\frac{b_i^{\alpha}}{(\sum_{k\in N} b_k^\alpha)^{2}}\Big(\frac{b_i^\alpha}{\sum_{k\in N} b_k^\alpha}\Big)^{r-1} \\
&= r\alpha\Big(\sum_{k\neq i} b_k^\alpha\Big)\frac{b_i^{\alpha}}{(\sum_{k\in N} b_k^\alpha)^{2}} x_i(\mathbf{b})^{r-1}
\end{aligned}$$

Conveniently, we have $\frac{\partial}{\partial b_i}\big(\frac{b_i^\alpha}{\sum_{k\in N} b_k^\alpha}\big) = \alpha\big(\sum_{k\neq i} b_k^\alpha\big)\frac{b_i^{\alpha-1}}{(\sum_{k\in N} b_k^\alpha)^2}$. Thus

$$\begin{aligned}
\frac{\partial p_i(\mathbf{b})}{\partial b_i} &= r b_i \frac{\partial}{\partial b_i}\Big(\frac{b_i^\alpha}{\sum_{k\in N} b_k^\alpha}\Big) x_i(\mathbf{b})^{r-1} \\
&= r b_i \frac{\partial x_i(\mathbf{b})}{\partial b_i} x_i(\mathbf{b})^{r-1}
\end{aligned}$$

Therefore

$$\frac{\partial u_i(\mathbf{b})}{\partial b_i} = r(w_i - b_i)\frac{\partial x_i(\mathbf{b})}{\partial b_i} x_i(\mathbf{b})^{r-1}$$

Since $\frac{\partial x_i(\mathbf{b})}{\partial b_i} > 0$ and $x_i(\mathbf{b})^{r-1} > 0$ for all b_i, this implies

1. For all $b_i < w_i$, $\frac{\partial u_i(\mathbf{b})}{\partial b_i} > 0$.
2. For all $b_i > w_i$, $\frac{\partial u_i(\mathbf{b})}{\partial b_i} > 0$.
3. For $b_i = w_i$, $\frac{\partial u_i(\mathbf{b})}{\partial b_i} = 0$.

Therefore $w_i \in \arg\max_{b_i \in \mathbb{R}_{>0}}$ (in fact, w_i is the unique maximizer). We conclude that the mechanism is truthful. □

From a technical standpoint, the harder task is proving that this mechanism is unique (up to additive constants in the payment rule). We assume without loss of generality that the mechanism

[17] Note that $u_i(\mathbf{b})$ is not concave in b_i, since $p_i(\mathbf{b})$ is not convex in b_i. Thus the KKT conditions do not apply, so we will have to use a different approach.

asks each agent i to report w_i, and let $\mathbf{b} = b_1, \ldots, b_n$ be the vector of reported w_i's. We use the standard notation of $(\mathbf{b}_{-i}, b'_i)$ to denote the vector where the ith entry is b'_i, and the kth entry for each $k \neq i$ is b_k.

The proof takes a real analysis approach, with Kirszbraun's Theorem for Lipschitz extensions [111] playing a central role. On a high level, the proof proceeds as follows: (1) we establish some basic properties of the payment rule, (2) we show that the payment rule must be Lipschitz continuous not including $b_i = 0$, (3) there exists a Lipschitz extension $\hat{p}_i$ including $b_i = 0$ (Kirszbraun's Theorem), (4) since $\hat{p}_i$ is Lipschitz, it is differentiable almost everywhere and is equal to the integral of its derivative, (5) since it has the same derivative (when defined) as the payment rule from Theorem 4.6.1, the payment rules are equal (up to the constant of integration).

Theorem 4.6.2. *Assume $m = 1$, and that each v_i is homogenous of degree r (with r publicly known), concave, and differentiable. Fix $\rho \in (0,1)$, and let Γ be a truthful mechanism which outputs an allocation $\mathbf{x} \in \Psi(\rho)$. Then the allocation rule is the same as in Theorem 4.6.1, and the payment rule $p_i(\mathbf{b})$ is the same up to an additive constant.*

Proof. **Part 1: Setup and basic properties.** Since there is a unique optimal allocation (Lemma 4.6.2), Γ must take $\mathbf{b} = (b_1, \ldots, b_n)$ as honest and output the same allocation $\mathbf{x}(\mathbf{b})$.

It remains to consider the payment rule. Let $p_i(\mathbf{b})$ denote the payment rule from Theorem 4.4.1, and let $\tilde{p}_i(\mathbf{b})$ denote the payment rule for Γ. Given reports $\mathbf{b}$, define $v_i(\mathbf{b})$ as before, and let $u_i(\mathbf{b}) = v_i(\mathbf{b}) - \tilde{p}_i(\mathbf{b})$ be agent i's resulting utility under Γ. From the point of view of a given agent i, the other agents' reports $\mathbf{b}_{-i}$ can be treated as a constant. Thus for brevity, write $x_i(b) = x_i(\mathbf{b}_{-i}, b)$, $p_i(b) = p_i(\mathbf{b}_{-i}, b)$, and $\tilde{p}_i(b) = \tilde{p}_i(\mathbf{b}_{-i}, b)$ for each $i \in N$.

Fix an $i \in N$. Since Γ is truthful, we must have $w_i \in \arg\max_{b_i \in \mathbb{R}_{>0}} u_i(\mathbf{b})$. Then by definition of $u_i(\mathbf{b})$, we have $w_i \in \arg\max_{b_i \in \mathbb{R}_{>0}} \big(w_i x_i(b_i)^r - \tilde{p}_i(b_i)\big)$. Since w_i could be any element of $\mathbb{R}_{>0}$, and Γ must be agnostic to w_i, we must have $b \in \arg\max_{b_i \in \mathbb{R}_{>0}} \big(b x_i(b_i)^r - \tilde{p}_i(b_i)\big)$ for all $b \in \mathbb{R}_{>0}$.

We first claim that $\tilde{p}_i(b)$ is nondecreasing. Suppose the opposite: then there exists exists $b > b'$ such that $\tilde{p}_i(b) < \tilde{p}_i(b')$. But this means that if $w_i = b'$, reporting $b_i = w_i$ is never an optimal strategy, because the payment can be decreased by reporting $b_i = b$, and $x_i(b) \geq x_i(b')$ (since $x_i(b)$ is nondecreasing). Thus $\tilde{p}_i(b)$ is nondecreasing.

Part 2: $\tilde{p}_i$ is Lipschitz continuous. Fix an arbitrary $b_i > 0$. Since $b_k > 0$ for all $k \neq i$, it can be seen from the definition of $x_i(b)$ that $x_i(b)^r$ is continuously differentiable on $[0, b_i]$. Therefore the maximum of $\frac{\mathrm{d}x_i(b)^r}{\mathrm{d}b}$ is a Lipschitz constant for $x_i(b)^r$, so $x_i(b)^r$ is Lipschitz continuous on $[0, b_i]$. Let κ be this Lipschitz constant: then for all $b, b' \in [0, b_i]$, $|x_i(b)^r - x_i(b')^r| \leq \kappa |b - b'|$.

We claim that $\tilde{p}_i$ is Lipschitz continuous on $(0, b_i]$ with constant $b_i \kappa$. Suppose the opposite: then there exist $b, b' \in (0, b_i]$ such that $|\tilde{p}_i(b) - \tilde{p}_i(b')| > b_i \kappa |b - b'|$. Assume without loss of generality that $b > b'$. Since $\tilde{p}_i$ and x_i are both nondecreasing, we then have $\tilde{p}_i(b) - \tilde{p}_i(b') > b_i \kappa (b - b')$ and $x_i(b)^r - x_i(b')^r \leq \kappa(b - b')$.

Since $b \in \arg\max_{b_i \in \mathbb{R}_{>0}} \big(b x_i(b_i)^r - \tilde{p}_i(b_i)\big)$, we have $b x_i(b)^r - \tilde{p}_i(b) \geq b x_i(b')^r - \tilde{p}_i(b')$ and thus $b(x_i(b)^r - x_i(b')^r) \geq \tilde{p}_i(b) - \tilde{p}_i(b')$. Therefore

$$b\kappa(b - b') \geq b\big(x_i(b)^r - x_i(b')^r\big) \geq \tilde{p}_i(b) - \tilde{p}_i(b') > b_i \kappa (b - b')$$

Therefore $\frac{b\kappa}{b_i\kappa} > 1$, which contradicts $b \le b_i$. Therefore $\tilde{p}_i$ is Lipschitz continuous on $(0, b_i]$.

Part 3: Kirszbraun's Theorem. Thus by Kirszbraun's Theorem [111], $\tilde{p}_i$ has a Lipschitz extension to $[0, b_i]$: that is, there exists $\hat{p}_i : [0, b_i] \to \mathbb{R}_{\ge 0}$ such that $\hat{p}_i$ is Lipschitz continuous on $[0, b_i]$, and $\hat{p}_i(b) = \tilde{p}_i(b)$ for $b \in (0, b_i]$.

Part 4: $\hat{p}_i$ is the integral of its derivative. Lipschitz continuity implies absolute continuity [157], so $\hat{p}_i$ is absolutely continuous on $[0, b_i]$. Thus by the Fundamental Theorem of Lebesgue Calculus [157], $\hat{p}_i$ is differentiable almost everywhere on $[0, b_i]$, its derivative $\frac{\mathrm{d}\hat{p}_i(b)}{\mathrm{d}b}$ is integrable over $[0, b_i]$, and

$$\hat{p}_i(b_i) - \hat{p}_i(0) = \int_{b=0}^{b_i} \frac{\mathrm{d}\hat{p}_i(b)}{\mathrm{d}b}\, \mathrm{d}b$$

Part 5: The derivatives of $\hat{p}_i$ and p_i match, so $\hat{p}_i = p_i + c$. Consider a $b > 0$ at which $\hat{p}_i$ is differentiable. Then $\tilde{p}_i$ is also differentiable, so $b \in \arg\max_{b_i \in \mathbb{R}_{>0}} \big(bx_i(b_i)^r - \tilde{p}_i(b_i)\big)$ implies $\frac{\mathrm{d}\tilde{p}_i(b)}{\mathrm{d}b} = b\frac{\mathrm{d}}{\mathrm{d}b}(x_i(b)^r) = rb\frac{\mathrm{d}x_i(b)}{\mathrm{d}b}x_i(b)^{r-1}$.[18] We showed in the proof of Theorem 4.6.1 that $\frac{\partial p_i(\mathbf{b})}{\partial b_i} = rb_i\frac{\partial x_i(\mathbf{b})}{\partial b_i}x_i(\mathbf{b})^{r-1}$; equivalently, $\frac{\mathrm{d}p_i(b)}{\mathrm{d}b} = rb\frac{\mathrm{d}x_i(b)}{\mathrm{d}b}x_i(b)^{r-1}$. Therefore for all $b > 0$ at which $\hat{p}_i$ is differentiable, we have $\frac{\mathrm{d}\hat{p}_i(b)}{\mathrm{d}b} = \frac{\mathrm{d}p_i(b)}{\mathrm{d}b}$.

Since $\hat{p}_i$ is differentiable almost everywhere, we have $\frac{\mathrm{d}\hat{p}_i(b)}{\mathrm{d}b} = \frac{\mathrm{d}p_i(b)}{\mathrm{d}b}$ almost everywhere. Thus $\frac{\mathrm{d}p_i(b)}{\mathrm{d}b}$ is also integrable over $[0, b_i]$, and $\int_{b=0}^{b_i} \frac{\mathrm{d}p_i(b)}{\mathrm{d}b}\,\mathrm{d}b = \int_{b=0}^{b_i} \frac{\mathrm{d}\hat{p}_i(b)}{\mathrm{d}b}\,\mathrm{d}b$ [157]. Therefore

$$\begin{aligned}\hat{p}_i(b_i) &= \hat{p}_i(0) + \int_{b=0}^{b_i} \frac{\mathrm{d}p_i(b)}{\mathrm{d}b}\,\mathrm{d}b \\ &= \hat{p}_i(0) + p_i(b_i)\end{aligned}$$

where the second equality is from the definition of $p_i(\mathbf{b})$.

Therefore for all $b_i > 0$, $\tilde{p}_i(b_i) = \hat{p}_i(0) + p_i(b_i)$, and so $\tilde{p}_i(\mathbf{b}) = \hat{p}_i(0) + p_i(\mathbf{b})$ for all $\mathbf{b}$. Since this holds for all $i \in N$, $\tilde{p}_i(\mathbf{b})$ is exactly the payment rule from Theorem 4.6.1, up to the additive constant of $\hat{p}_i(0)$. □

It is worth noting that this truthful payment rule is quite complex; in particular, it may be hard to convince agents that it is actually in their best interest to be truthful. In contrast, the Walrasian pricing rule from Theorem 4.4.1 is much simpler and more intuitive. That pricing rule is not truthful, but perhaps formal truthfulness is not crucial if a practical iterative implementation is possible. We do not claim that our algorithm from Section 4.5 is truly practical, but it could be a step in the right direction.

[18] Note that the b in $bx_i(b_i)^r$ is a constant from the point of view of the argmax, so it is treated as a constant by the derivative. To be technically precise, we have $(\frac{\mathrm{d}}{\mathrm{d}b_i}bx_i(b_i)^r)|_{b_i=b} = rb\frac{\mathrm{d}x_i(b)}{\mathrm{d}b}x_i(b)^{r-1}$.

4.7 Sybil attacks

In Sections 4.5 and 4.6, we discussed two alternative approaches to implementation: an iterative query-based algorithm, and a truthful mechanism. However, there is an additional crucial issue which any practical implementation must address: since our pricing rule $p(x_i) = (\sum_{j \in M} q_j x_{ij})^{1/\rho}$ is strictly convex for $\rho < 1$, agents have an incentive to create fake identities. In particular, an agent can decrease her payment while receiving the same bundle by splitting the payment over multiple fake identities.[19] This is known as a *Sybil attack*. The truthful payment rule from Section 4.6 is not strictly convex everywhere, but it is strictly convex on some intervals, and thus has the same vulnerability.

Model of Sybil attacks. We model this as follows. Let κ denote the cost of creating a new identity. The cost could reflect inconvenience, risk of getting caught, or other factors, and would depend on the nature of the system. Let η_i be the *multiplicity* of agent i, i.e., the number of identities agent i controls in the system. This includes both fake identities and the agent's single real identity, so we assume that $\eta_i \in \mathbb{N}_{>0}$. For convex p, multiplicity η_i, and a desired bundle for purchase, it is always optimal for agent i to split the purchase evenly across her identities.[20] Thus we can assume that each identity purchases the same bundle x_i, and we define agent i's utility as

$$u_i(x_i, \eta_i) = v_i(\eta_i x_i) - \eta_i p(x_i) - \eta_i \kappa$$

We do not claim that this perfectly models the reality of Sybil attacks; for example, the identity creation cost is arguably sublinear (one someone has created a single fake identity, creating more might become easier). Our goal here is simply to show formally that at least in some cases, CES welfare maximization cannot be robust to Sybil attacks in general.

Walrasian equilibrium. We focus on a Walrasian model of Sybil attacks; the analogous analysis for truthful mechanisms is left as an open question. We define each agent's *Sybil demand set* by

$$D_i(p) = \underset{x_i \in \mathbb{R}^m_{\geq 0}, \eta_i \in \mathbb{N}_{>0}}{\arg\max} \quad u_i(x_i, \eta_i)$$

Note that we require $\eta_i \in \mathbb{N}_{>0}$. We define a *Sybil Walrasian equilibrium* (SWE) to be an allocation $\mathbf{x}$, payment rule p, and vector of multiplicities $\boldsymbol{\eta} = \eta_1, \dots, \eta_n$ such that

1. Each agent receive a bundle in her demand set: $(x_i, \eta_i) \in D_i(p)$ for all $i \in N$.
2. The market clears: for all $j \in M$, $\sum_{i \in N} x_{ij} \leq 1$. Furthermore, for any $j \in M$ with nonzero cost[21], $\sum_{i \in N} x_{ij} = 1$.

[19] Note that for linear prices there is no such incentive.
[20] This is essentially a multidimensional version of Jensen's inequality; see, e.g., [130].
[21] Recall that good j has "nonzero cost" in our pricing rule if $q_j > 0$.

In this section, we will focus on the case of homogeneity degree $r = 1$. The following lemma states that for any pricing rule, a rational agent either creates no fake identities (i.e, $\eta_i = 1$), or creates an unbounded number (and consequently the demand set is empty).

Lemma 4.7.1. *Assume each v_i is concave, differentiable, and homogeneous of degree 1. Let $\rho \in (0, 1]$, define p as in Theorem 4.4.1, and let $\mathbf{x} \in \Psi(\rho)$. Then we have*

$$D_i(p) = \begin{cases} (x_i, 1) & \text{if } v_i(x_i)(1-\rho) \leq \kappa \\ \emptyset & \text{otherwise} \end{cases}$$

where x_i is agent i's bundle in $\mathbf{x}$.

Proof. When v_i is homogeneous of degree 1, for any bundle y_i, we have $u_i(y_i, \eta_i) = \eta_i v_i(y_i) - \eta_i p(y_i) - \eta_i \kappa = \eta_i\big(v_i(y_i) - p(y_i) - \kappa\big)$. Thus given a choice of η_i, y_i must be chosen to maximize $v_i(y_i) - p(y_i) - \kappa$. Let $\mathbf{x} \in \Psi(\rho)$: then by Theorem 4.4.1 y_i optimizes $v_i(y_i) - p(y_i)$ (and thus $v_i(y_i) - p(y_i) - \kappa$) if and only if $y_i = x_i$. Therefore the demand set is equal to

$$D_i(p) = \Big(x_i, \ \operatorname*{arg\,max}_{\eta_i \in \mathbb{N}_{>0}} \ \eta_i\big(v_i(x_i) - p(x_i) - \kappa\big)\Big)$$

That is, the demanded bundle must always be x_i, and η_i is optimized accordingly.

By Corollary 4.4.1.1, $p(x_i) = \rho v_i(x_i)$, so $v_i(x_i) - p(x_i) - \kappa = v_i(x_i)(1-\rho) - \kappa$. Thus if $v_i(x_i)(1-\rho) \leq \kappa$, then 1 is an optimal choice for η_i, so $(x_i, 1) \in D_i(p)$. However, if $v_i(x_i)(1-\rho) > \kappa$, there is no optimal choice for η_i: specifically, η_i goes to infinity. Thus if $v_i(x_i)(1-\rho) > \kappa$, $D_i(p) = \emptyset$. □

This immediately implies that if $\mathbf{x} \in \Psi(\rho)$ satisfies $v_i(x_i)(1-\rho) \leq \kappa$ for all $i \in N$, the convex pricing rule from Theorem 4.4.1 is naturally robust to Sybil attacks.

Theorem 4.7.1. *Assume each v_i is concave, differentiable, and homogeneous of degree 1. Let $\mathbf{x} \in \Psi(\rho)$ for $\rho \in (0, 1]$, and define p as in Theorem 4.4.1. Then if $v_i(x_i)(1-\rho) \leq \kappa$ for all $i \in N$, $(\mathbf{x}, p, \mathbf{1})$ is a SWE.*

Proof. By Lemma 4.7.1, we have $(x_i, 1) \in D_i(p)$ for all $i \in N$ in this case. Theorem 4.4.1 implies that the market clearing condition is met, so $(\mathbf{x}, p, \mathbf{1})$ is a SWE. □

In other words, if the identity creation cost is small, ρ is close to 1, and/or agents valuations are not too large, we need not worry about Sybil attacks. As discussed in Section 4.2, this suggests one possible way for a social planner to choose a value of ρ: estimate κ and the magnitude of valuations, and choose ρ to be as small as possible without incentivizing Sybil attacks.

On the other hand, if $v_i(x_i)(1-\rho) > \kappa$, how bad are the consequences? Theorem 4.7.2 states that an agent's valuation at equilibrium has a hard cap at $\dfrac{\kappa}{1-\rho}$. This provides a hard maximum on the CES welfare in any SWE with p thus defined: in particular, the CES welfare is at most $\big(\sum_{i\in N} \big(\frac{\kappa}{1-\rho}\big)^\rho\big)^{1/\rho} = n^{1/\rho}\frac{\kappa}{1-\rho}$. In general, each $v_i(x_i)$ (and thus the CES welfare) can be arbitrarily large, so Theorem 4.7.2 implies an unbounded ratio between the optimal CES welfare and that of any SWE with this p.

Theorem 4.7.2. *Assume each v_i is concave, differentiable, and homogeneous of degree 1. Let $\rho \in (0, 1]$, and define p as in Theorem 4.4.1. Then for any allocation $\mathbf{x}$ and multiplicities $\boldsymbol{\eta}$ such that $(\mathbf{x}, p, \boldsymbol{\eta})$ is a SWE, we have*

$$v_i(x_i) \leq \frac{\kappa}{1-\rho}$$

Proof. Suppose $(\mathbf{x}, p, \boldsymbol{\eta})$ is a SWE for some allocation $\mathbf{x}$ and multiplicities $\boldsymbol{\eta}$: then each $(x_i, \eta_i) \in D_i(p)$ for all $i \in N$; Thus $D_i(\mathbf{p}) \neq \emptyset$, so Lemma 4.7.1 implies that $v_i(x_i)(1-\rho) \leq \kappa$, and consequently, $v_i(x_i) \leq \frac{\kappa}{1-\rho}$. □

The next natural question is, can we circumvent this by using a different pricing rule? Theorem 4.7.3 answers this in the negative. The counterexample uses an instance with a single good; recall that x_i denotes a scalar in this case.

Theorem 4.7.3. *Let $m = 1$, $v_1(x_1) = wx_1$, and $v_i(x_i) = x_i$ for all $i \neq 1$. Let $(\mathbf{x}, p, \boldsymbol{\eta})$ be any SWE. Then for all $i \neq 1$,*

$$v_i(x_i) \leq \frac{\kappa}{w-1}$$

Proof. Let $(\mathbf{x}, p, \boldsymbol{\eta})$ be any SWE. Fix an arbitrary $i \in N$. As in Lemma 4.7.1, we have $u_i(x_i, \eta_i) = \eta_i(v_i(x_i) - p(x_i) - \kappa)$. Since $(x_i, \eta_i) \in D_i(p)$, we must have $\eta_i \in \arg\max_{\eta_i' \in \mathbb{N}_{>0}} \eta_i(v_i(x_i) - p(x_i) - \kappa)$ (note that we are not assuming anything about the bundle x_i). Since $\arg\max_{\eta_i' \in \mathbb{N}_{>0}} \eta_i'(v_i(x_i) - p(x_i) - \kappa)$ cannot be the empty set, we must have $v_i(x_i) \leq p(x_i) + \kappa$ and $\eta_i = 1$.

Focusing on agent 1, we further claim that $v_1(x_i) \leq p(x_i) + \kappa$ for any $i \neq 1$. Suppose not: then agent 1 could purchase x_i and set $\eta_1 = \infty$ to increase her utility. Thus $v_1(x_i) \leq p(x_i) + \kappa$ for each $i \neq 1$. Now looking at the optimization for $i \neq 1$, we have $v_i(x_i) \geq p(x_i)$. Combining this with $v_1(x_i) \leq p(x_i) + \kappa$, we get $v_1(x_i) \leq v_i(x_i) + \kappa$.

Plugging in our definitions of v_1 and $v_{i \neq 1}$, we get $wx_i \leq x_i + \kappa$, so $x_i(w-1) \leq \kappa$. Substituting back in the definition of v_i, we get $v_i(x_i) \leq \dfrac{\kappa}{w-1}$ for all $i \neq 1$, as required. □

Although the bound in Theorem 4.7.3 is different from that in Theorem 4.7.2, the implication is the same: this is a hard maximum on the value obtained by any agent other than agent 1. As κ goes to zero, the fraction of the good agent 1 receives approaches 1, so the outcome approaches the maximum utilitarian welfare outcome (where agent 1 receives the entirety of the good). Therefore by Theorem 4.8.1, the CES welfare at any Sybil Walrasian equilibrium (for any pricing rule) can be arbitrarily bad in comparison to the optimal CES welfare. Thus in general, when Sybil attacks are possible, it is impossible to implement any bounded approximation of CES welfare maximization in Walrasian equilibrium.

4.8 Negative results

Even when Sybil attacks are not possible, there are limitations to implementation in WE. This section presents several relevant counterexamples.

4.8.1 Linear pricing poorly approximates CES welfare for $\rho \neq 1$

Recall that for an allocation $\mathbf{x}$, $\Phi(\rho, \mathbf{x})$ denotes the CES welfare of $\mathbf{x}$. In contrast, $\Psi(\rho)$ denotes the set of allocations with optimal CES welfare with respect to ρ.

Our first negative result relates to linear pricing. In particular, can linear pricing guarantee a reasonable approximation of CES welfare? We show that the answer is no, justifying the need for nonlinear pricing. In particular, for any $\rho \in (0,1)$, the gap between the CES welfare of any linear pricing equilibrium and the optimal CES welfare can be arbitrarily large.

Note that as ρ goes to zero, $\frac{1}{\rho} - 1$ goes to infinity, so the denominator of the bound (and thus the gap in CES welfare) in the following theorem can indeed be arbitrarily large.

Theorem 4.8.1. *Let $m = 1$, $\rho \in (0,1]$, $v_1(x) = (1+\varepsilon)x$ for some $\varepsilon > 0$, and $v_i(x) = x$ for all $i \neq 1$. Suppose $(\mathbf{x}, p)$ is a WE where p is linear. Then*

$$\frac{\Phi(\rho, \mathbf{x})}{\max_{\mathbf{y}} \Phi(\rho, \mathbf{y})} \leq \frac{1+\varepsilon}{n^{\frac{1}{\rho}-1}}$$

Proof. By the First Welfare Theorem, $\mathbf{x}$ must maximize utilitarian (i.e., $\rho = 1$) welfare. Thus by Lemma 4.6.2, $\mathbf{x}$ must give the entire good to agent 1: $x_1 = 1$ and $x_i = 0$ for $i \neq 1$. Thus the CES welfare of $\mathbf{x}$ with respect to ρ is

$$\Phi(\rho, \mathbf{x}) = \Big(\sum_{i \in N} v_i(x_i)^\rho\Big)^{1/\rho} = \big((1+\varepsilon)^\rho\big)^{1/\rho}$$

In contrast, consider the allocation $\mathbf{y}$ such that $y_i = 1/n$ for all $i \in N$:

$$\Phi(\rho, \mathbf{y}) = \Big(\sum_{i \in N} v_i(1/n)^\rho\Big)^{1/\rho} \geq \Big(\sum_{i \in N} (1/n)^\rho\Big)^{1/\rho} = \Big(n(1/n)^{1/\rho}\Big)^{1/\rho} = n^{\frac{1}{\rho}-1}$$

Thus $\max_{\mathbf{y}} \Phi(\rho, \mathbf{y}) \geq n^{\frac{1}{\rho}-1}$, as required. □

4.8.2 Theorem 4.4.1 does not extend to nonuniform homogeneity degrees

In this section, we show that for all $\rho \in (0,1)$, Theorem 4.4.1 does not extend to the case where different v_i's have different homogeneity degrees. This shows that our result is tight in the sense that it is necessary to require the same homogeneity degree.

We begin with the following lemma, which is a standard property of strictly concave and differentiable functions: it essentially states that any such function is bounded above by any tangent line. This lemma is sometimes called the "Rooftop Theorem".

Lemma 4.8.1. *Let $f : \mathbb{R} \to \mathbb{R}$ be strictly concave and differentiable. Then for all $a, b \in \mathbb{R}$ where $a \neq b$, $f(a) < f(b) + f'(b)(a-b)$, where f' denotes the derivative of f.*

The next lemma is also quite standard; we provide a proof for completeness.

Lemma 4.8.2. *Let $f : \mathbb{R} \to \mathbb{R}$ be strictly concave and differentiable, and let $x, a_1, \ldots, a_k$ be non-negative reals such that $\sum_{i=1}^k a_i = 0$. Then $\sum_{i=1}^k f(x + a_i) < kf(x)$.*

Proof. The lemma follows from Lemma 4.8.1 and arithmetic:

$$\begin{aligned}
\sum_{i=1}^k f(x+a_i) <& \sum_{i=1}^k (f(x) + f'(x)(x + a_i - x)) \\
=& \sum_{i=1}^k f(x) + f'(x)\sum_{i=1}^k a_i \\
=& \sum_{i=1}^k f(x) + f'(x) \cdot 0 \\
=& kf(x)
\end{aligned}$$

□

We are now ready to present our counterexample.

Theorem 4.8.2. *Let $n = 2$ and $m = 1$, and for $x \in \mathbb{R}_{\geq 0}$, let $v_1(x) = x$ and $v_2(x) = \sqrt{2x}$. Then for all $\rho \in (0,1)$, there exists no allocation $\mathbf{x} \in \Psi(\rho)$ and pricing rule $p : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ such that $(\mathbf{x}, p)$ is a WE.*

Proof. Suppose for sake of contradiction that such $\mathbf{x}, p$ do exist. We first claim that $x_1 > x_2$. Suppose the opposite: then $x_2 \geq 1/2 \geq x_1$. Thus $v_2(x_2) \geq 1 > 1/2 \geq v_1(x_1)$. We also have $\frac{\partial v_2(x_2)}{\partial x_2} = \frac{1}{\sqrt{2x_2}} \leq 1 = \frac{\partial v_1(x_1)}{\partial x_1}$. Thus $v_2(x_2) > v_1(x_2)$ and $\frac{\partial v_2(x_2)}{\partial x_2} \leq \frac{\partial v_1(x_1)}{\partial x_1}$. Since $\rho < 1$, $\rho - 1 < 0$, so we have $v_2(x_2)^{\rho-1}\frac{\partial v_2(x_2)}{\partial x_2} < v_1(x_1)^{\rho-1}\frac{\partial v_1(x_1)}{\partial x_2}$. But this contradicts $\mathbf{x} \in \Psi(\rho)$, so we have $x_1 > x_2$ as claimed.[22]

Since $(\mathbf{x}, p)$ is a WE, we must have $x_i \in D_i(p)$ for both agents i. Thus for any $x \neq x_i$, $v_i(x_i) - p(x_i) \geq v_i(x) - p(x)$. Therefore

$$\begin{aligned}
v_1(x_1) - p(x_1) \geq v_1(x_2) - p(x_2) \quad &\text{and} \quad v_2(x_2) - p(x_2) \geq v_2(x_1) - p(x_1) \\
v_1(x_1) + v_2(x_2) - p(x_1) - p(x_2) &\geq v_1(x_2) + v_2(x_1) - p(x_1) - p(x_2) \\
v_1(x_1) + v_2(x_2) &\geq v_1(x_2) + v_2(x_1) \\
v_1(x_1) - v_1(x_2) &\geq v_2(x_1) - v_2(x_2)
\end{aligned}$$

Since $x_1 > 1/2 > x_2$ and $x_1 + x_2 = 1$, let $x_1 = 1/2 + \varepsilon$ and $x_2 = 1/2 - \varepsilon$. Then we have $v_1(x_1) - v_1(x_2) = 2\varepsilon$. For $v_2(x_1) - v_2(x_2)$, we have

$$\begin{aligned}
v_2(x_1) - v_2(x_2) &= \sqrt{1+2\varepsilon} - \sqrt{1-2\varepsilon} \\
&= \frac{(\sqrt{1+2\varepsilon} - \sqrt{1-2\varepsilon})(\sqrt{1+2\varepsilon} - \sqrt{1-2\varepsilon})}{\sqrt{1+2\varepsilon} + \sqrt{1-2\varepsilon}}
\end{aligned}$$

[22] This immediately implies $\frac{\partial v_2(x_2)}{\partial x_2} < 1 = \frac{\partial v_1(x_1)}{\partial x_1}$, which, in combination with $x_1 > x_2$, rules out convex p. However, we still need to rule out non-convex p.

$$= \frac{(1+2\varepsilon) - (1-2\varepsilon)}{\sqrt{1+2\varepsilon} + \sqrt{1-2\varepsilon}}$$
$$= \frac{4\varepsilon}{\sqrt{1+2\varepsilon} + \sqrt{1-2\varepsilon}}$$

Applying Lemma 4.8.2 with $f(x) = \sqrt{x}$, $x = 1$, $k = 2$, and $(a_1, a_2) = (2\varepsilon, -2\varepsilon)$, we get $\sqrt{1+2\varepsilon} + \sqrt{1-2\varepsilon} < 2$. Thus $v_2(x_1) - v_2(x_2) > 4\varepsilon/2 = 2\varepsilon = v_1(x_1) - v_1(x_2)$. However, this contradicts $v_1(x_1) - v_1(x_2) \geq v_2(x_1) - v_2(x_2)$, as we showed above. We conclude that there is no $\mathbf{x} \in \Psi(\rho)$ and pricing rule p such that $(\mathbf{x}, p)$ is a WE. □

4.8.3 CES welfare maximization for $\rho \leq 0$

In this section, we show that there is no pricing rule supporting CES welfare maximization for any $\rho < 0$. For $\rho = 0$ (i.e., Nash welfare), the situation is slightly different. We do show, however, that Nash welfare maximization cannot be supported by a differentiable pricing rule.

Theorem 4.8.3. *Consider the instance with $n = 2$, $m = 1$, $v_1(x) = x$ and $v_2(x) = 2x$. Then for every $\rho < 0$, there is no pricing rule p and allocation $\mathbf{x} \in \Psi(\rho)$ such that $(\mathbf{x}, p)$ is a WE.*

Proof. For any $\rho < 0$ and any $\mathbf{x} \in \Psi(\rho)$, we must have $x_1 > x_2$. Assume $(\mathbf{x}, p)$ is a WE for some pricing rule p: then $x_1 \in D_1(p)$, so $v_1(x_1) - p(x_1) \geq v_1(x_2) - p(x_2)$. Thus $p(x_1) \leq p(x_2) + v_1(x_1) - v_2(x_2) = p(x_2) + x_2 - x_1$. Therefore

$$\begin{aligned}
v_2(x_1) - p(x_1) &\geq 2x_1 - (p(x_2) + x_2 - x_1) \\
&= 3x_1 - x_2 - p(x_2) \\
&> 2x_1 - p(x_2) \\
&> 2x_2 - p(x_2) \\
&= v_2(x_2) - p(x_2)
\end{aligned}$$

Thus agent 2 would rather purchase x_1 than x_2, so $x_2 \notin D_2(p)$. Therefore $(\mathbf{x}, p)$ is not a WE. □

For $\rho = 0$, the situation is different. Recall that Fisher market equilibrium always maximizes Nash welfare, and we can simulate Fisher market budgets by setting

$$p(x_i) = \begin{cases} 0 & \text{if } \sum_{j \in M} q_j x_{ij} \leq 1 \\ \infty & \text{otherwise} \end{cases}$$

where $q_1 \dots q_m$ are the optimal Lagrange multipliers in the convex program for maximizing Nash welfare. Gale and Eisenberg's famous result implies that for such a pricing rule, a WE always exists, and all WE maximize Nash welfare [72, 73]. Note that for $\sum_{j \in M} q_j x_{ij} > 1$, $p(x_i) = \infty$ can be implemented by setting $\frac{\partial p(x_i)}{\partial x_{ij}}$ to be at least $\max_{i \in N} \max_{x_i \in [0,1]^m} \frac{\partial v_i(x_i)}{\partial x_{ij}}$. This ensures that no agent purchases a bundle x_i such that $\sum_{j \in M} q_j x_{ij} > 1$.

The above pricing rule is somewhat artificial, however. One natural question is whether Nash welfare maximization can be implemented with a differentiable pricing rule. We next show that the answer is no.

Theorem 4.8.4. *Consider the instance with $n = 2$, $m = 1$, $v_1(x) = x$ and $v_2(x) = 2x$. Then there is no allocation $\mathbf{x}$ maximizing Nash welfare and differentiable pricing rule p such that $(\mathbf{x}, p)$ is a WE.*

Proof. Suppose the opposite: that such $\mathbf{x}, p$ exist. The unique $\mathbf{x}$ maximizing Nash welfare must have $x_1 = x_2 = 1/2$. Since p, v_1, and v_2 are all differentiable, we have $x_i \in D_i(p)$ if and only if $\frac{\mathrm{d}p(x_i)}{\mathrm{d}x_i} = \frac{\mathrm{d}v_i(x_i)}{\mathrm{d}x_i}$. Since $x_1 = x_2$, we have $\frac{\mathrm{d}p(x_1)}{\mathrm{d}x_1} = \frac{\mathrm{d}p(x_2)}{\mathrm{d}x_2}$. Thus implies $\frac{\mathrm{d}v_1(x_1)}{\mathrm{d}x_1} = \frac{\mathrm{d}v_2(x_2)}{\mathrm{d}x_2}$, which is a contradiction. We conclude that no such $\mathbf{x}, p$ exist. □

4.9 Connections to Fisher markets

The focus of this chapter is on markets for quasilinear utilities, where agents can spend as much money as they want, and the amount spent is incorporated into their resulting utility. The other predominant market model assumes each agent i has a finite budget B_i of money to spend, and has no value for leftover money (in general, this implies that each agent i spends exactly B_i). This is called the *Fisher market* model.[23] In this section, we explore connections between our results and the Fisher market model.

In the Fisher market model, each agent's utility $u_i(x_i)$ is simply $v_i(x_i)$. For pricing rule p, the Fisher market demand set is given by

$$D_i^F(p) = \underset{x_i \in \mathbb{R}^m_{\geq 0}:\ p(x_i) \leq B_i}{\arg\max}\ v_i(x_i)$$

We will reserve the notation $D_i(p)$ for the demand set in the quasilinear case, i.e., $D_i(p) = \arg\max_{x_i \in \mathbb{R}^m_{\geq 0}} (v_i(x_i) - p(x_i))$.

For an allocation $\mathbf{x}$, agent budgets $\mathbf{B} = (B_1, \dots, B_n)$, and a pricing rule p, $(\mathbf{x}, \mathbf{B}, p)$ is a *Fisher market Walrasian equilibrium* if (1) $x_i \in D_i^F(p)$ for all $i \in N$, and (2) $\sum_{i \in N} x_{ij} \leq 1$ for all $j \in M$, and if good j has nonzero cost, $\sum_{i \in N} x_{ij} = 1$.[24] These are the same two conditions for Walrasian equilibrium in quasilinear markets: the only change is the definition of the demand set. To distinguish, we will use the terms "Fisher WE" and "quasilinear WE".

4.9.1 CES welfare maximization in Fisher markets

In the quasilinear model, agents can express not only their relative values between goods, but also the absolute scale of their valuation (i.e., the "intensity" of their preferences) by choosing how much money to spend. In contrast, agents in the Fisher market model spend exactly their budget, and so have no way to express the absolute scale of their valuation. This should make us pessimistic

[23] There are also more general versions of this model that allow each agent's initial endowment to be goods instead of money ("exchange economies") and/or allow production ("Arrow-Debreu markets").

[24] Recall that good j has nonzero cost for j if there is a bundle x_i such that $x_{i\ell} = 0$ for all $\ell \neq j$, but $p(x_i) > 0$.

about the possibility of CES welfare maximization in the Fisher market model in general. Indeed, consider a single good and two agents with valuations $v_1(x) = x$, $v_2(x) = 2x$. For any $\rho > 0$, any optimal allocation $\mathbf{x} \in \Psi(\rho)$ has $x_2 > x_1$. But if $B_1 = B_2$, any Fisher market Walrasian equilibrium will always have $x_1 = x_2$, since both agents simply spend their entire budget on the single good.

However, in general we can convert a quasilinear WE to a Fisher WE if the agents' budgets are sized appropriately. Specifically, we need agent i's budget to be exactly the amount she pays in the quasilinear WE:

Theorem 4.9.1. *Suppose $(\mathbf{x}, p)$ is a quasilinear WE, and let $B_i = p(x_i)$. Then $(\mathbf{x}, \mathbf{B}, p)$ is a Fisher WE.*

Proof. For all $i \in N$, x_i is affordable to agent i under p by definition of B_i. Suppose there were another bundle y_i such that $p(y_i) \leq B_i$ but $v_i(y_i) > v_i(x_i)$. That would contradict $x_i \in D_i(p)$ for the quasilinear case, since $u_i(y_i) = v_i(y_i) - p(y_i) > v_i(x_i) - p(y_i) \geq v_i(x_i) - p(x_i) = u_i(x_i)$. Therefore $x_i \in D_i^F(p)$ for all $i \in N$. Furthermore, the market clearing conditions for Fisher WE and quasilinear WE are identical. We conclude that $(\mathbf{x}, \mathbf{B}, p)$ is a Fisher WE. ☐

Combining the above result with Theorem 4.4.1 gives us the following corollary for CES welfare maximization:

Corollary 4.9.1.1. *Assume each v_i is homogeneous of degree r, concave, and differentiable. Let $\rho \in (0, 1]$, and $p(x_i) = \rho r^{\frac{\rho-1}{\rho}} (\sum_{j \in M} q_j x_{ij})^{1/\rho}$, where $q_1, \dots, q_m$ are optimal Lagrange multipliers for Program 4.1. For $\mathbf{x} \in \Psi(\rho)$, let $B_i = p(x_i)$ for all $i \in N$, Then $(\mathbf{x}, \mathbf{B}, p)$ is a Fisher WE.*

Perhaps the more interesting connection relates to the welfare function being optimized. In the case of linear pricing, the Fisher market Walrasian equilibria are exactly the budget-weighted maximum Nash welfare allocations.[25] One natural question is whether the Fisher market equilibria from Theorem 4.9.1 also optimize a budget-weighted CES welfare function. We answer this in the affirmative. Recall that we define $\Phi_{\mathbf{B}}(\rho, \mathbf{x}) = \big(\sum_{i \in N} B_i v_i(x_i)^\rho\big)^{1/\rho}$, and $\Psi_{\mathbf{B}}(\rho) = \arg\max_{\mathbf{x}} \Phi_{\mathbf{B}}(\rho, \mathbf{x})$.

Lemma 4.9.1. *Assume each v_i is concave and differentiable. Let $\mathbf{x}'$ be any allocation, let $a_i = v_i(x'_i)$ for each $i \in N$, and let $\rho \in (0, 1]$. Then $\mathbf{x}' \in \Psi(\rho)$ if and only if $\mathbf{x}' \in \Psi_{\mathbf{a}}(\rho - 1)$.*

Proof. When $\rho = 1$, $\rho - 1 = 0$, so Program 4.1 does not apply, and we must handle this case separately. We first consider $\rho \neq 1$. The Lagrangian for Program 4.1 for $\Psi_{\mathbf{a}}(\rho - 1)$ is $L(\mathbf{x}, \mathbf{q}) = \frac{1}{\rho-1} \sum_{i \in N} a_i v_i(x_i)^{\rho-1} - \sum_{j \in M} q_j (\sum_{i \in N} x_{ij} - 1)$. The KKT conditions imply that $\mathbf{x} \in \Psi_{\mathbf{a}}(\rho - 1)$ if and only if there exist Lagrange multipliers $q_1, \dots, q_m$ such that:

1. Stationarity: $\dfrac{\partial L(\mathbf{x}, \mathbf{q})}{\partial x_{ij}} = a_i v_i(x_i)^{\rho-2} \dfrac{\partial v_i(x_i)}{\partial x_{ij}} \leq 0$ for all i, j.[26] Furthermore, if $x_{ij} > 0$, the inequality holds with equality.

2. Complementary slackness: for all $j \in M$, either $\sum_{i \in N} x_{ij} = 1$, or $q_j = 0$.

[25] Recall that Nash welfare corresponds to $\rho = 0$, and the budget-weighted Nash welfare of an allocation $\mathbf{x}$ is $\prod_{i \in N} v_i(x_i)^{B_i}$.

[26] Note that since $\mathbf{x}'$ is a fixed allocation, a_i is just some constant, so the differentiation does not affect it.

For Program 4.1 for $\Psi(\rho)$, as before we have $L'(\mathbf{x}, \mathbf{q}) = \frac{1}{\rho}\sum_{i\in N} v_i(x_i)^\rho - \sum_{j\in M} q_j(\sum_{i\in N} x_{ij} - 1)$. Thus the KKT conditions imply that $\mathbf{x} \in \Psi(\rho)$ if and only if there exist $q'_1, \ldots, q'_m \in \mathbb{R}_{\geq 0}$ such that (1) $v_i(x_i)^{\rho-1}\frac{\partial v_i(x_i)}{\partial x_{ij}} \leq q_j$ for all i, j, and when $x_{ij} > 0$, the inequality holds with equality, and (2) for all $j \in M$, either $\sum_{i\in N} x_{ij} = 1$, or $q_j = 0$. Note that if $q_j = q'_j$ for all $j \in M$, the complementary slackness conditions become equivalent.

Next, for $\mathbf{x} = \mathbf{x}'$ we have

$$v_i(x'_i)^{\rho-1}\frac{\partial v_i(x'_i)}{\partial x'_{ij}} = v_i(x'_i)v_i(x'_i)^{\rho-2}\frac{\partial v_i(x'_i)}{\partial x'_{ij}} = a_i v_i(x'_i)^{\rho-2}\frac{\partial v_i(x'_i)}{\partial x'_{ij}}$$

Therefore for given q_j, we have $q_j \geq v_i(x'_i)^{\rho-1}\frac{\partial v_i(x'_i)}{\partial x'_{ij}}$ if and only if $q_j \geq a_i v_i(x'_i)^{\rho-2}\frac{\partial v_i(x'_i)}{\partial x'_{ij}}$, and $q_j = v_i(x'_i)^{\rho-1}\frac{\partial v_i(x'_i)}{\partial x'_{ij}}$ if and only if $q_j \geq a_i v_i(x'_i)^{\rho-2}\frac{\partial v_i(x'_i)}{\partial x'_{ij}}$.

Now suppose $\mathbf{x}' \in \Psi(\rho)$. Then there exist $q_1, \ldots, q_m \in \mathbb{R}_{\geq 0}$ that satisfy both stationarity and complementary slackness. Then as we showed above, $\mathbf{x}'$ and $q_1, \ldots, q_m$ satisfy stationarity for $\Psi_{\mathbf{a}}(\rho - 1)$. Furthermore, the complementary slackness conditions are equivalent, so we have $\mathbf{x}' \in \Psi_{\mathbf{a}}(\rho - 1)$.

Similarly, suppose $\mathbf{x}' \in \Psi_{\mathbf{a}}(\rho - 1)$. Then there exist $q_1, \ldots, q_m$ satisfying stationarity and complementary slackness, so the same $q_1, \ldots, q_m$ along with $\mathbf{x}'$ satisfy the KKT conditions for $\Psi(\rho)$. Therefore $\Psi(\rho)$, and we conclude that $\mathbf{x}' \in \Psi(\rho)$ if and only if $\mathbf{x}' \in \Psi_{\mathbf{a}}(\rho - 1)$ for $\rho \neq 1$.

All of the above was for $\rho \neq 1$; it remains to handle the case of $\rho = 1$. In this case, we can use the same KKT conditions for $\Psi(\rho)$, but must use a different convex program for $\Psi_{\mathbf{a}}(\rho-1)$. Consider the following convex program for maximizing Nash welfare (i.e., CES welfare for $\rho = 0$):

$$\max_{\mathbf{x}\in\mathbb{R}^{n\times m}_{\geq 0}} \sum_{i\in N} a_i \log v_i(x_i) \qquad (4.6)$$
$$s.t. \sum_{i\in N} x_{ij} \leq 1 \qquad \forall j \in M$$

This is known as the Eisenberg-Gale program [72, 73]. In this case, the stationarity condition requires that $\frac{\partial}{\partial x_{ij}} a_i \log v_i(x_i) = a_i v_i(x_i)^{-1}\frac{\partial v_i(x_i)}{\partial x_{ij}} \leq q_j$ for all i, j, and when $x_{ij} > 0$, the inequality holds with equality. Since $\rho = 1$ here, we have $\rho - 2 = -1$. Thus the stationarity condition for $\Psi_{\mathbf{a}}(\rho - 1)$ requires that $av_i(x_i)^{\rho-2}\frac{\partial v_i(x_i)}{\partial x_{ij}} \leq q_j$ for all i, j (and if $x_{ij} > 0$, this holds with equality). This is exactly what we had above, and since we are using the same KKT conditions for $\Psi(\rho)$, this case reduces to the case for $\rho \neq 1$. Therefore for $\rho = 1$, $\mathbf{x}' \in \Psi(\rho)$ if and only if $\mathbf{x}' \in \Psi_{\mathbf{a}}(\rho - 1)$. □

Combining Theorem 4.9.1 and Lemma 4.9.1, we get:

Theorem 4.9.2. *Assume each v_i is homogeneous of degree r, concave, and differentiable, let $\rho \in (0, 1]$, let $q_1, \ldots, q_m \in \mathbb{R}_{\geq 0}$, and let $p(x_i) = \rho\big(\sum_{j\in M} q_j x_{ij}\big)^{1/\rho}$. Given $\mathbf{x} \in \Psi(\rho)$, let $B_i = p(x_i)$. Then all of the following hold:*

1. *$(\mathbf{x}, p)$ is a quasilinear WE.*

2. $(\mathbf{x}, \mathbf{B}, p)$ *is a Fisher WE.*

3. $\mathbf{x} \in \Psi_{\mathbf{B}}(\rho - 1)$

Proof. The first and second conditions hold by Theorems 4.4.1 and 4.9.1, respectively. Then Corollary 4.4.1.1 implies that $p(x_i) = r\rho v_i(x_i)$. Let $a_i = v_i(x_i) = \frac{B_i}{r\rho}$. Thus by Lemma 4.9.1, we have $\mathbf{x} \in \Psi_{\mathbf{a}}(\rho - 1)$. Since scaling all agents' multipliers by the same factor does not affect $\Psi_{\mathbf{a}}(\rho - 1)$, we have $\mathbf{x} \in \Psi_{r\rho\mathbf{a}}(\rho - 1) = \Psi_{\mathbf{B}}(\rho - 1)$, as required. ☐

It is worth noting that the special case of Theorem 4.9.1 for $\rho = 1$ and Leontief utilities with $w_{ij} \in \{0, 1\}$[27] is implied by the work of Kelly et al. [110].

4.10 CES welfare maximization for Leontief valuations

We say that v_i is *Leontief* if there exist weights $w_1, \dots, w_m \in \mathbb{R}_{\geq 0}$ such that

$$v_i(x_i) = \min_{j:\ w_{ij} \neq 0} \frac{x_{ij}}{w_{ij}}$$

Leontief valuations are not differentiable, and so Theorem 4.4.1 does not apply. In this section, we handle Leontief valuations as a special case. Although there are many non-differentiable valuations we could consider, Chapters 2 and 3 focused on Leontief valuations, so we find it worthwhile to show that our result does indeed extend to this case.

Recall Program 4.1:

$$\begin{aligned} \max_{\mathbf{x} \in \mathbb{R}^{n \times m}_{\geq 0}} \ & \frac{1}{\rho} \sum_{i \in N} v_i(x_i)^\rho \\ s.t. \ & \sum_{i \in N} x_{ij} \leq 1 \qquad \forall j \in M \end{aligned}$$

We will work with a specialized version of this for Leontief utilities:

$$\begin{aligned} \max_{\mathbf{x} \in \mathbb{R}^{n \times m}_{\geq 0}, \boldsymbol{\alpha} \in \mathbb{R}^{m}_{\geq 0}} \ & \frac{1}{\rho} \sum_{i \in N} \alpha_i^\rho \\ s.t. \ & w_{ij}\alpha_i \leq x_{ij} \qquad \forall i \in N, j \in M \\ & \sum_{i \in N} x_{ij} \leq 1 \qquad \forall j \in M \end{aligned} \tag{4.7}$$

where we use $\boldsymbol{\alpha}$ to denote the vector $(\alpha_1, \dots, \alpha_n) \in \mathbb{R}^n_{\geq 0}$.

Also recall each agent's demand set $D_i(p) = \arg\max_{x_i \in \mathbb{R}^m_{\geq 0}} \big(v_i(x_i) - p(x_i)\big)$. Similarly to Program 4.7, we consider the following equivalent (specialized) convex program for agent i's demand

[27] This is also known as the bandwidth allocation setting, where each good represents a link in a network, and agent i has $w_{ij} = 1$ for every link j on a fixed path (and $w_{ij} = 0$ otherwise).

set:

$$\max_{x_i \in \mathbb{R}^m_{\geq 0}, \alpha_i \in \mathbb{R}_{\geq 0}} \big(\alpha_i - p(x_i)\big) \tag{4.8}$$
$$s.t.\ w_{ij}\alpha_i \leq x_{ij} \quad \forall j \in M$$

Theorem 4.10.1. *Assume each v_i is Leontief with weights $w_{i1}, \ldots, w_{im}$. Then for any $\rho \in (0, 1]$ and any allocation $\mathbf{x}$, we have $\mathbf{x} \in \Psi(\rho)$ if and only if there exist $q_1, \ldots, q_m \in \mathbb{R}_{\geq 0}$ such that for the pricing rule $p(x_i) = \rho(\sum_{j \in M} q_j x_{ij})^{1/\rho}$, $(\mathbf{x}, p)$ is a WE. Furthermore, $q_1, \ldots, q_m$ are optimal Lagrange multipliers (for the $\sum_{i \in N} x_{ij} \leq 1$ constraints) for Program 4.7.*

Proof. We first claim that in an optimal solution $\mathbf{x}, \boldsymbol{\alpha}$ to either Program 4.7 or Program 4.8, we have $v_i(x_i) = \alpha_i$ for all $i \in N$: that is, that these programs are doing what we want them to. To see this, note that $\alpha_i \leq x_{ij}/w_{ij}$ for all j with $w_{ij} \neq 0$, so $\alpha_i \leq v_i(x_i)$. Furthermore, at least one constraint involving α_i must be tight: otherwise, we could increase α_i and thus the objective value. In particular, we must have $\alpha_i = \min_{j:\ w_{ij} \neq 0} \frac{x_{ij}}{w_{ij}} = v_i(x_i)$. Thus Program 4.8 is indeed maximizing $v_i(x_i) - p(x_i)$, so $x_i \in D_i(p)$ if and only if (x_i, α_i) is optimal for Program 4.8 (for some α_i). Similarly, Program 4.7 is indeed maximizing $\frac{1}{\rho} \sum_{i \in N} v_i(x_i)^\rho$ subject to $\sum_{i \in N} x_{ij} \leq 1$ for all $j \in M$, so $(\mathbf{x}, \boldsymbol{\alpha})$ is optimal for Program 4.7 if and only if $\mathbf{x}$ is optimal for Program 4.1. Therefore $\mathbf{x} \in \Psi(\rho)$ if and only if $(\mathbf{x}, \boldsymbol{\alpha})$ is optimal for Program 4.7 (for some $\boldsymbol{\alpha}$).

Next, we write the Lagrangian of Program 4.7:[28]

$$L(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda}) = \frac{1}{\rho} \sum_{i \in N} \alpha_i^\rho - \sum_{i \in N} \sum_{j \in M} \lambda_{ij}(w_{ij}\alpha_i - x_{ij}) - \sum_{j \in M} q_j \Big(\sum_{i \in N} x_{ij} - 1 \Big)$$

We have strong duality by Slater's condition, so the KKT conditions are both necessary and sufficient for optimality. That is, $(\mathbf{x}, \boldsymbol{\alpha})$ is optimal if and only if there exist $\mathbf{q} \in \mathbb{R}^m_{\geq 0}$, $\boldsymbol{\lambda} \in \mathbb{R}^{m \times n}_{\geq 0}$ such that all of the following hold:[29]

1. Stationarity for $\mathbf{x}$: $\dfrac{\partial L(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})}{\partial x_{ij}} \leq 0$ for all i, j. Furthermore, if $x_{ij} > 0$, the inequality holds with equality.
2. Stationarity for $\boldsymbol{\alpha}$: $\dfrac{\partial L(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})}{\partial \alpha_i} \leq 0$ for all $i \in N$. Furthermore, if $\alpha_i > 0$, the inequality holds with equality.
3. Complementary slackness for $\mathbf{q}$: for all $j \in M$, either $\sum_{i \in N} x_{ij} = 1$, or $q_j = 0$.
4. Complementary slackness for $\boldsymbol{\lambda}$: for all $i \in N$, $j \in M$, either $w_{ij}\alpha_i = x_{ij}$ or $\lambda_{ij} = 0$.

Similarly, let L'_i denote the Lagrangian of Program 4.8 for agent i:

$$L'_i(x_i, \alpha_i, \lambda_i) = \alpha_i - p(x_i) - \sum_{j \in M} \lambda_{ij}(w_{ij}\alpha_i - x_{ij})$$

[28] As in the proof of Theorem 4.4.1, we omit the $\mathbf{x} \in \mathbb{R}^{m \times n}_{\geq 0}$ constraint from the Lagrangian incorporate it into the KKT conditions instead.

[29] As in the proof of Theorem 4.4.1, primal and dual feasibility are trivially satisfied.

where $\lambda_i = (\lambda_{i1}, \dots, \lambda_{im}) \in \mathbb{R}^m_{\geq 0}$. We again have strong duality, so the KKT conditions are again necessary and sufficient. Let $L'_i(x_i, \alpha_i, \lambda_i)$ denote the Lagrangian of this program; then (x_i, α_i) is optimal for Program 4.8 if and only if all of the following hold:

1. Stationarity for x_i: $\frac{\partial L'_i(x_i,\alpha_i,\lambda_i)}{\partial x_{ij}} \leq 0$ for all $j \in M$. If $x_{ij} > 0$, the inequality holds with equality.
2. Stationarity for α_i: $\frac{\partial L'_i(x_i,\alpha_i,\lambda_i)}{\partial \alpha_i} \leq 0$. If $\alpha_i > 0$, the inequality holds with equality.
3. Complementary slackness for λ_i: for all $i \in N$, $j \in M$, either $w_{ij}\alpha_i = x_{ij}$ or $\lambda_{ij} = 0$.

We will claim that $(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})$ is optimal for Program 4.7 if and only if for all $i \in N$, $(x_i, \alpha_i, \alpha_i^{1-\rho}\lambda_i)$ is optimal for Program 4.8. Essentially, we show that if complementary slackness holds (for either program), the stationarity conditions are equivalent. To begin, we can explicitly compute the relevant partial derivatives for given $\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda}$, with $p(x_i) = \rho(\sum_{j\in M} q_j x_{ij})^{1/\rho}$:

$$\frac{\partial L(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})}{\partial x_{ij}} = \lambda_{ij} - q_j$$
$$\frac{\partial L(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})}{\partial \alpha_i} = \alpha_i^{\rho-1} - \sum_{j\in M} \lambda_{ij} w_{ij}$$
$$\frac{\partial L'_i(x_i, \alpha_i, \lambda'_i)}{\partial x_{ij}} = \lambda'_{ij} - q_j\Big(\sum_{\ell\in M} q_\ell x_{i\ell}\Big)^{\frac{1-\rho}{\rho}}$$
$$\frac{\partial L'_i(x_i, \alpha_i, \lambda'_i)}{\partial \alpha_i} = 1 - \sum_{j\in M} \lambda'_{ij} w_{ij}$$

Part 1: ($\Longrightarrow$) Suppose that $\mathbf{x} \in \Psi(\rho)$. Then there exist $\boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda}$ such that the KKT conditions for Program 4.7 are satisfied for $(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})$. We first claim that $\alpha_i > 0$ for all $i \in N$. Suppose not: stationarity implies that $\alpha_i^{\rho-1} \leq \sum_{j\in M} \lambda_{ij} w_{ij}$, but since $\rho - 1 < 0$, the left hand side is not defined for $\alpha_i = 0$. Thus $\alpha_i > 0$.

Therefore by stationarity for α_i, we have $\alpha_i^{\rho-1} = \sum_{j\in M} \lambda_{ij} w_{ij}$. Let $\lambda'_{ij} = \alpha_i^{1-\rho}\lambda_{ij}$ for all i, j[30]. Then $\alpha_i^{\rho-1} = \sum_{j\in M} \lambda_{ij} w_{ij}$ is equivalent to $1 = \sum_{j\in M} \lambda'_{ij} w_{ij}$, and thus $\frac{\partial L'_i(x_i,\alpha_i,\lambda'_i)}{\partial \alpha_i} = 0$ for all $i \in N$. Thus for all $i \in N$, $(x_i, \alpha_i, \lambda'_i)$ satisfies stationarity for α_i for Program 4.8.

We now turn to the x_{ij} variables. Stationarity for x_{ij} in Program 4.7 implies that $\lambda_{ij} = q_j$ whenever $x_{ij} > 0$. Furthermore, complementary slackness for λ_{ij} implies that if $\lambda_{ij} > 0$, $w_{ij}\alpha_i = x_{ij}$. Thus whenever $q_j > 0$ and $x_{ij} > 0$, $w_{ij}\alpha_i = x_{ij}$ and $\lambda_{ij} = q_j$. Therefore for all i, j,

$$\begin{aligned}\frac{\partial L'_i(x_i, \alpha_i, \lambda'_i)}{\partial x_{ij}} &= \lambda'_{ij} - q_j\Big(\sum_{\ell: q_\ell, x_{i\ell} > 0} q_\ell x_{i\ell}\Big)^{\frac{1-\rho}{\rho}} \\ &= \lambda'_{ij} - q_j\Big(\sum_{\ell\in M} \lambda_{i\ell} w_{i\ell}\alpha_i\Big)^{\frac{1-\rho}{\rho}}\end{aligned}$$

[30]Note that this is *not* defining λ'_{ij} to be a function of α_i. This is defining λ'_{ij} based on a fixed value of α_i: in particular, the value from $(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})$, which we assumed to be optimal for Program 4.7. Consequently, the derivatives in the KKT conditions treat λ'_{ij} as a constant.

$$
\begin{aligned}
&= \lambda'_{ij} - q_j\Big(\alpha_i \sum_{\ell \in M} \lambda_{i\ell} w_{i\ell}\Big)^{\frac{1-\rho}{\rho}} \\
&= \lambda'_{ij} - q_j(\alpha_i \alpha_i^{\rho-1})^{\frac{1-\rho}{\rho}} \\
&= \alpha_i^{1-\rho}\lambda_{ij} - q_j\alpha_i^{1-\rho} \\
&= \alpha_i^{1-\rho}\frac{\partial L(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})}{\partial x_{ij}}
\end{aligned}
$$

We have $\frac{\partial L(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})}{\partial x_{ij}} \le 0$ for all i, j by stationarity (and the inequality holds with equality when $x_{ij} > 0$), so $\frac{\partial L'_i(x_i, \alpha_i, \lambda'_i)}{\partial x_{ij}} \le 0$ for all $j \in M$ (and the inequality holds with equality when $x_{ij} > 0$). Thus for each $i \in N$, $(x_i, \alpha_i, \lambda'_i)$ satisfies stationarity for Program 4.8 for x_{ij} for all $j \in M$.

As mentioned above, we have $w_{ij}\alpha_i = x_{ij}$ whenever $\lambda_{ij} > 0$. Since $\lambda'_{ij} > 0$ if and only if $\lambda_{ij} > 0$, we have $w_{ij}\alpha_i = x_{ij}$ whenever $\lambda'_{ij} > 0$. Thus for each $i \in N$, $(x_i, \alpha_i, \lambda'_i)$ satisfies complementary slackness for Program 4.8. Therefore $(x_i, \alpha_i, \lambda'_i)$ satisfies the KKT conditions, and thus is optimal for Program 4.8. Therefore $x_i \in D_i(p)$ for all $i \in N$. The complementary slackness condition for $\mathbf{q}$ is identical to the market clearing condition, so we conclude that $(\mathbf{x}, p)$ is a WE.

Part 2: ($\Longleftarrow$) Suppose that $(\mathbf{x}, p)$ is a WE, where $p(x_i) = \rho(\sum_{j \in M} q_j x_{ij})^{1/\rho}$ for constants $q_1, \dots, q_m \in \mathbb{R}_{\ge 0}$. Then $x_i \in D_i(p)$ for all $i \in N$, so there exists $\boldsymbol{\alpha}, \boldsymbol{\lambda'}$ such that $(x_i, \alpha_i, \lambda'_i)$ is optimal for Program 4.8 for all $i \in N$.

Thus by stationarity, we have $\frac{\partial L'_i(x_i, \alpha_i, \lambda'_i)}{\partial \alpha_i} \le 0$ and $\frac{\partial L'_i(x_i, \alpha_i, \lambda'_i)}{\partial x_{ij}} \le 0$ for all i, j (and if $\alpha_i > 0$ and $x_{ij} > 0$, the inequalities hold with equality). Using the definition of p, we have $\frac{\partial L'_i(x_i, \alpha_i, \lambda'_i)}{\partial x_{ij}} = \lambda'_{ij} - q_j\big(p(x_i)/\rho\big)^{1-\rho}$. Thus $1 \le \sum_{j \in M} \lambda'_{ij} w_{ij}$ and $\lambda'_{ij} \le q_j\big(p(x_i)/\rho\big)^{1-\rho}$.

We first claim that $\alpha_i > 0$ for all $i \in N$. For each agent i, there must exist $j \in M$ such that $\lambda'_{ij} > 0$ and $w_{ij} > 0$: otherwise $1 \le \sum_{j \in M} \lambda'_{ij} w_{ij}$ would be impossible. Consider any such j: then $0 < \lambda'_{ij} \le q_j\big(p(x_i)/\rho\big)^{1-\rho}$, so we must have $p(x_i) > 0$. Suppose $\alpha_i = 0$: then the optimal objective value of Program 4.8 is $\alpha_i - p(x_i) < 0$. But setting $x_{ij} = 0$ for all $j \in M$ achieves an objective value of 0, so $\alpha_i - p(x_i) < 0$ cannot be optimal. This is a contradiction, and so $\alpha_i > 0$ for all $i \in N$.

Returning to the stationarity conditions, we then have $1 = \sum_{j \in M} \lambda'_{ij} w_{ij}$. Complementary slackness implies that $w_{ij}\alpha_i = x_{ij}$ whenever $\lambda'_{ij} > 0$, so we get

$$
\begin{aligned}
\alpha_i &= \sum_{j \in M} \lambda'_{ij} w_{ij} \alpha_i \\
\alpha_i &= \sum_{j : \lambda'_{ij} > 0} \lambda'_{ij} w_{ij} \alpha_i \\
\alpha_i &= \sum_{j \in M} \lambda'_{ij} x_{ij}
\end{aligned}
$$

Combining this with $\lambda'_{ij} = q_j\big(p(x_i)/\rho\big)^{1-\rho}$ whenever $x_{ij} > 0$ gives us

$$
\alpha_i = \sum_{j : x_{ij} > 0} \lambda'_{ij} x_{ij}
$$

$$
\begin{aligned}
&= \sum_{j \in M} q_j x_{ij} \big(p(x_i)/\rho\big)^{1-\rho} \\
&= \big(p(x_i)/\rho\big)^{1-\rho} \sum_{j \in M} q_j x_{ij} \\
&= \big(p(x_i)/\rho\big)^{1-\rho} \big(p(x_i)/\rho\big)^{\rho} \\
&= p(x_i)/\rho
\end{aligned}
$$

Now let $\lambda_{ij} = \alpha_i^{\rho-1} \lambda'_{ij}$ for all i, j. We claim that $(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})$ satisfies the KKT conditions for Program 4.7. For each (i, j) pair, we have

$$
\frac{\partial L(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})}{\partial \alpha_i} = \alpha_i^{\rho-1} - \sum_{j \in M} \lambda_{ij} w_{ij} = \alpha_i^{\rho-1} \Big(1 - \sum_{j \in M} \lambda'_{ij} w_{ij}\Big) = \alpha_i^{\rho-1} \frac{\partial L'_i(x_i, \alpha_i, \lambda'_i)}{\partial \alpha_i}
$$

Since $\alpha_i > 0$, stationarity for Program 4.8 implies that $\frac{\partial L'_i(x_i, \alpha_i, \lambda'_i)}{\partial \alpha_i} = 0$, so we have $\frac{\partial L(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})}{\partial \alpha_i} = 0$. Next, we have

$$
\begin{aligned}
\frac{\partial L(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})}{\partial x_{ij}} &= \lambda_{ij} - q_j \\
&= \alpha_i^{\rho-1} \lambda'_{ij} - q_j \\
&= \alpha_i^{\rho-1} (\lambda'_{ij} - q_j \alpha_i^{1-\rho}) \\
&= \alpha_i^{\rho-1} \Big(\lambda'_{ij} - q_j \big(p(x_i)/\rho\big)^{1-\rho}\Big) \\
&= \alpha_i^{\rho-1} \frac{\partial L'_i(x_i, \alpha_i, \lambda'_i)}{\partial x_{ij}}
\end{aligned}
$$

Stationarity for Program 4.8 implies that $\frac{\partial L'_i(x_i, \alpha_i, \lambda'_i)}{\partial x_{ij}} \le 0$ for all i, j (and when $x_{ij} > 0$, this holds with equality), so we have $\frac{\partial L(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})}{\partial x_{ij}}$ for all i, j (and when $x_{ij} > 0$, this holds with equality). Thus we have shown that $(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})$ satisfies stationarity for Program 4.7. As before, the market clearing condition is equivalent to complementary slackness for $\mathbf{q}$. By complementary slackness for $\boldsymbol{\lambda}'$ (for Program 4.8), we have $w_{ij} \alpha_i = x_{ij}$ whenever $\lambda'_{ij} > 0$. By definition, $\lambda'_{ij} > 0$ if and only if $\lambda_{ij} > 0$, so this implies the required complementary slackness for $\boldsymbol{\lambda}$ (for Program 4.7). Therefore $(\mathbf{x}, \boldsymbol{\alpha}, \mathbf{q}, \boldsymbol{\lambda})$ satisfies the KKT conditions for Program 4.7, and thus is optimal for that program. We conclude that $\mathbf{x} \in \Psi(\rho)$.

□

4.11 The First Welfare Theorem and linear pricing

Recall our main result:

Theorem 4.4.1. *Assume each v_i is homogeneous of degree r, concave, and differentiable. For any $\rho \in (0, 1]$ and any allocation $\mathbf{x}$, we have $\mathbf{x} \in \Psi(\rho)$ if and only if there exist $q_1, \dots, q_m \in \mathbb{R}_{\ge 0}$ such* *that for the pricing rule*

$$p(x_i) = \rho r^{\frac{\rho-1}{\rho}} \Big(\sum_{j \in M} q_j x_{ij}\Big)^{1/\rho},$$

$(\mathbf{x}, p)$ *is a WE. Furthermore,* $q_1, \dots, q_m$ *are optimal Lagrange multipliers for Program 4.1.*

For this class of valuations, Theorem 4.4.1 for $\rho = 1$ implies the First Welfare Theorem: p becomes a linear pricing rule, and CES welfare for $\rho = 1$ is just utilitarian welfare. In particular, Theorem 4.4.1 implies both the existence of a WE, and that every WE maximizes utilitarian welfare.

Typically, the "the First Welfare Theorem" refers to just half of this: that every WE maximizes utilitarian welfare. The reason is that WE are not always guaranteed to exist: for divisible goods, generally at least concavity or quasi-concavity of valuations is necessary. On the other hand, very few assumptions are needed to show that linear pricing equilibria always maximize utilitarian welfare; for example, divisibility of goods is not needed. We provide a proof of this below.

Theorem 4.11.1 (The First Welfare Theorem)**.** *Let* $\mathcal{X}_i \subset \mathbb{R}^m$ *denote the set of feasible bundles for agent* i *(not necessarily convex, and not necessarily the same for all agents). Let* $D_i(p) = \arg\max_{x_i \in \mathcal{X}_i}(v_i(x_i) - p(x_i))$ *and assume* p *is linear. Then if* $(\mathbf{x}, p)$ *is a WE,* $\mathbf{x}$ *maximizes utilitarian welfare.*

Proof. Since p is linear, there exist $q_1, \dots, q_m$ such that $p(y_i) = \sum_{j \in M} q_j y_{ij}$ for any bundle y_i. Consider an arbitrary feasible allocation $\mathbf{y}$. Since $(\mathbf{x}, p)$ is a WE, we have $x_i \in D_i(p)$, so $v_i(x_i) - p(x_i) \geq v_i(y_i) - p(y_i)$. Therefore

$$\begin{aligned}
v_i(x_i) - \sum_{j \in M} q_j x_{ij} &\geq v_i(y_i) - \sum_{j \in M} q_j y_{ij} \\
\sum_{i \in N} v_i(x_i) - \sum_{i \in N}\sum_{j \in M} q_j x_{ij} &\geq \sum_{i \in N} v_i(y_i) - \sum_{i \in N}\sum_{j \in M} q_j x_{ij} \\
\sum_{i \in N} v_i(x_i) - \sum_{j \in M} q_j \sum_{i \in N} x_{ij} &\geq \sum_{i \in N} v_i(y_i) - \sum_{j \in M} q_j \sum_{i \in N} x_{ij}
\end{aligned}$$

Furthermore, $\sum_{i \in N} x_{ij} = 1$ for all $j \in M$ with $q_j > 0$. Also, since $\mathbf{y}$ is a valid allocation, $\sum_{i \in N} y_{ij} \leq 1$ for all $j \in M$. Therefore

$$\begin{aligned}
\sum_{i \in N} v_i(x_i) - \sum_{j \in M} q_j \sum_{i \in N} x_{ij} &\geq \sum_{i \in N} v_i(y_i) - \sum_{j \in M} q_j \sum_{i \in N} x_{ij} \\
\sum_{i \in N} v_i(x_i) - \sum_{j \in M} q_j &\geq \sum_{i \in N} v_i(y_i) - \sum_{j \in M} q_j \\
\sum_{i \in N} v_i(x_i) &\geq \sum_{i \in N} v_i(y_i)
\end{aligned}$$

Thus the utilitarian welfare of $\mathbf{x}$ is at least as high as that of any other allocation. We conclude that $\mathbf{x}$ maximizes utilitarian welfare. ☐

Note that no assumptions at all were made on the nature of the valuations: all we needed was $x_i \in \arg\max_{y_i \in \mathcal{X}_i}(v_i(y_i) - p(y_i))$, and $\sum_{j \in M} x_{ij} = 1$ whenever $q_j > 0$. The most natural cases are $\mathcal{X}_i = \mathbb{R}^m_{\geq 0}$ (divisible goods) and $\mathcal{X}_i = \{0,1\}^m$ (indivisible goods), but the result does hold more broadly.

4.12 Omitted proofs

Theorem 4.4.2 (Euler's Theorem for homogeneous functions)**.** *Let $f : \mathbb{R}^m_{\geq 0} \to \mathbb{R}$ be differentiable and homogenous of degree r. Then for any $\mathbf{b} = (b_1, \dots b_m) \in \mathbb{R}^m_{\geq 0}$, $\sum_{j=1}^m b_j \dfrac{\partial f(\mathbf{b})}{\partial b_j} = rf(\mathbf{b})$.*

Proof. Fix an arbitrary $\mathbf{b} \in \mathbb{R}^m_{\geq 0}$ and let $g(\lambda) = f(\lambda \mathbf{b})$. Since f is differentiable, so is g, and its derivative is given by the multidimensional chain rule: $\frac{\mathrm{d}g(\lambda)}{\mathrm{d}\lambda} = \sum_{j=1}^m b_j \frac{\partial f(\lambda \mathbf{b})}{\partial b_j}$. Since f is homogeneous of degree r, we have $f(\lambda \mathbf{b}) = \lambda^r f(\mathbf{b})$ for all $\lambda \geq 0$. Thus $g(\lambda) = \lambda^r f(\mathbf{b})$ for all $\lambda \geq 0$, so we can differentiable both sides of this equation to get $\sum_{j=1}^m b_j \frac{\partial f(\lambda \mathbf{b})}{\partial b_j} = r\lambda^{r-1} f(\mathbf{b})$. This holds for all $\lambda \geq 0$, so setting $\lambda = 1$ completes the proof. □

Lemma 4.6.1. *Let $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be differentiable and homogeneous of degree r. Then there exists $c \in \mathbb{R}_{\geq 0}$ such that $f(x) = cx^r$.*

Proof. By Euler's Theorem (Theorem 4.4.2), we have $x\dfrac{\mathrm{d}f(x)}{\mathrm{d}x} = rf(x)$ for all $x \in \mathbb{R}_{\geq 0}$. Let $y = f(x)$. We can solve this differential equation explicitly:

$$\begin{aligned} \frac{1}{y} \cdot \frac{\mathrm{d}y}{\mathrm{d}x} &= \frac{r}{x} \\ \int \frac{1}{y} \cdot \frac{\mathrm{d}y}{\mathrm{d}x}\, \mathrm{d}x &= \int \frac{r}{x}\, \mathrm{d}x \\ \int \frac{1}{y}\, \mathrm{d}y &= r \int \frac{1}{x}\, \mathrm{d}x \\ \ln y &= r \ln x + \ln c \end{aligned}$$

where c (and thus $\ln c$) is some constant. Therefore

$$\begin{aligned} e^y &= e^{r \ln x + \ln c} \\ y &= cx^r \end{aligned}$$

Thus $f(x) = cx^r$, as required. □

Lemma 4.6.2. *Let $m = 1$ and $v_i(x_i) = w_i x_i^r$ for all $i \in N$ where $r \in (0,1]$. Then $\rho \in (0,1]$ and $r\rho \neq 1$, $\mathbf{x} \in \Psi(\rho)$ if and only if*

$$x_i = \frac{{w_i}^{\frac{\rho}{1-r\rho}}}{\sum_{k \in N} {w_k}^{\frac{\rho}{1-r\rho}}}$$

If $\mathbf{x} \in \Psi(\rho)$ and $r = \rho = 1$, then whenever $x_i > 0$, $w_i = \max_{k \in N} w_k$.

Proof. As in Section 4.4, strong duality for Program 4.1 implies that any optimal $\mathbf{x}$ must satisfy the KKT conditions. Thus $\mathbf{x} \in \Psi(\rho)$ if and only if there exists $q \in \mathbb{R}_{\geq 0}$ such that (1) stationarity holds: $\frac{\partial v_i(x_i)}{\partial x_i} v_i(x_i)^{\rho-1} \leq q$ for all $i \in N$, and if $x_i > 0$, the inequality holds with equality, and (2) complementary slackness holds: either $\sum_{i \in N} x_i = 1$, or $q = 0$.

Since we assume that $w_i > 0$ for all $i \in N$, any allocation with $\sum_{i \in N} x_i < 1$ is not Pareto optimal, and thus cannot be optimal for Program 4.1. In other words, we must have $q > 0$. Thus complementary slackness simply requires that $\sum_{i \in N} x_i = 1$, and we can focus on stationarity.

We first handle $\rho = 1$. In this case, $\frac{\partial v_i(x_i)}{\partial x_i} v_i(x_i)^{\rho-1} = \frac{\partial v_i(x_i)}{\partial x_i} = w_i$. Thus if $\mathbf{x} \in \Psi(\rho)$ we must have $w_i \leq q$, and if $x_i > 0$, then $w_i = q$. This implies that $q = \max_{k \in N} w_k$. Thus if $x_i > 0$, then $w_i = \max_{k \in N} w_k$, as required.

For the rest of the proof, we assume $r\rho \neq 1$. Since $r, \rho \in (0, 1]$, we have $0 < r\rho < 1$. By the definition of v_i, for an arbitrary allocation $\mathbf{x}$ and $i \in N$ we have

$$\frac{\partial v_i(x_i)}{\partial x_i} v_i(x_i)^{\rho-1} = (w_i r x_i^{r-1})(w_i^{\rho-1} x_i^{r(\rho-1)}) = r w_i^\rho x_i^{r\rho-1}$$

Thus $q \geq r w_i^\rho x_i^{r\rho-1}$. Since $r\rho < 1$, if $x_i = 0$, then $x_i^{r\rho-1}$ is undefined. Therefore stationarity is satisfied if and only if $q = r w_i^\rho x_i^{r\rho-1}$ for all $i \in N$, which is equivalent to

$$x_i = (q/r)^{\frac{1}{r\rho-1}} w_i^{\frac{\rho}{1-r\rho}} \tag{4.9}$$

Furthermore, if $\mathbf{x}$ satisfies Equation 4.9 for all $i \in N$, then $\sum_{i \in N} x_i = 1$ is equivalent to

$$\sum_{i \in N} (q/r)^{\frac{1}{r\rho-1}} w_i^{\frac{\rho}{1-r\rho}} = 1$$

$$(q/r)^{\frac{1}{r\rho-1}} = \Big(\sum_{i \in N} w_i^{\frac{\rho}{1-r\rho}} \Big)^{-1}$$

Therefore $\mathbf{x}$ satisfies stationarity and complementary slackness (and thus satisfies $\mathbf{x} \in \Psi(\rho)$) if and only if

$$x_i = \frac{w_i^{\frac{\rho}{1-r\rho}}}{\sum_{k \in N} w_k^{\frac{\rho}{1-r\rho}}}$$

as required. □

4.13 Conclusion

In this chapter, we studied a simple family of convex pricing rules, motivated by the widespread use of convex pricing in the real world, especially for water. We proved that these pricing rules implement CES welfare maximization in Walrasian equilibrium, providing a formal quantitative interpretation of the frequent informal claim that convex pricing promotes equality. Furthermore, by tweaking the exponent of the pricing rule, the social planner can precisely control the tradeoff between equality and efficiency. This result also shows that convex pricing is not necessarily economically inefficient, as often claimed; it simply maximizes a different welfare function than the traditional utilitarian one.

Improved implementation is perhaps the most important of the future directions we propose. One concrete possibility is a *tâtonnement*: an iterative algorithm where on each step, each agent reports her demand for the current pricing rule, and the pricing rule is adjusted accordingly. Demand queries are arguably easier for agents to answer than valuation gradient queries. Some implementation questions – in particular, how to deal with Sybil attacks – would likely need to be handled on a case-by-case basis.

Aside from the implementation itself, there is the additional challenge of convincing market designers to consider using this type of convex pricing. Equality is generally thought to be desirable, but sellers may be concerned that this will decrease their revenue. In future work, we hope to show that our pricing rule guarantees a good approximation of the optimal revenue for sellers.

Another possible direction would be CES welfare maximization for indivisible goods. The analogous pricing rule would be $p(S) = (\sum_{j \in S} q_j)^{1/\rho}$, where S is a set of indivisible goods. It seems like very different theoretical techniques would be needed in this setting (along with perhaps a gross substitutes assumption), but we suspect that the same intuition of convex pricing improving equality would hold.

Part II

Axiomatic Private Resource Allocation

Chapter 5

A new fairness axiom: envy-freeness up to any good (EFX)

We now move on to axiomatic objectives for private resource allocation. This field is often referred to as “fair division”, since the axioms are typically motivated by fairness in some way (the discussion in Section 1.1 is relevant). Perhaps the most pervasive fairness axiom is envy-freeness, which states that no agent should prefer another agent’s bundle to her own. In this chapter (in fact, in all of Part II), we focus on indivisible goods, where envy-free allocations do not always exist: consider two agents a single good.

This motivates the study of relaxed versions of envy-freeness. We study the *envy-freeness up to any good* (EFX) property, which states that no player prefers the bundle of another player following the removal of any single good. First, we use the leximin solution to show existence of EFX allocations in several contexts, sometimes in conjunction with Pareto optimality. For two players with valuations obeying a mild assumption, one of these results provides stronger guarantees than the currently deployed algorithm on Spliddit, a popular fair division website. Unfortunately, finding the leximin solution can require exponential time. We show that this is necessary by proving an exponential lower bound on the number of value queries needed to identify an EFX allocation, even for two players with identical valuations. We consider both additive and more general valuations, and our work suggests that there is a rich landscape of problems to explore in the fair division of indivisible goods with different classes of player valuations.

These constitute the first formal results regarding EFX, an axiom which has enjoyed substantial attention since (e.g., [3, 39, 45, 46, 92]).

5.1 Introduction

Fair division has a long history, with the earliest known mechanism for solving the problem dating back to the Bible. No, not war; the cut-and-choose protocol. When Abraham and Lot first arrive in the land of Canaan, Abraham suggests that they divide the land between them. Abraham partitions the land into two parts and lets Lot choose which part he would like to keep.

What makes this procedure "fair"? By dividing the land into two pieces he values equally, Abraham can ensure that he will not envy Lot's piece, regardless of which piece Lot takes. Since Lot presumably chooses his favorite piece, he will not envy Abraham. This means that the cut-and-choose protocol guarantees an *envy-free* allocation, meaning that each player likes their allocation at least as much as any other player's allocation.

The cut-and-choose protocol is defined for two players and *divisible* goods, meaning that each good can be divided into arbitrarily small pieces. In this chapter, we consider the setting of *indivisible* goods, meaning that the resource in question is a set of discrete goods, each of which must be wholly allocated to a single player. Unfortunately, envy-freeness cannot be guaranteed in this setting. We see this even with two players and a single good: one player must receive the good, and the other will surely be envious.

Consequently, other notions of fairness are needed. Budish [35] introduced the concept of *envy-freeness up to one good* (EF1). In an EF1 allocation, player i may envy player j, but the envy could be eliminated by removing a single good from player j's allocation. The good is not actually removed; this is a thought experiment used in the definition of envy-freeness up to one good. An EF1 allocation always exists, and can be computed in polynomial time [119].[1]

Caragiannis et al. [40] proposed another fairness criterion, one which is strictly stronger than EF1, but strictly weaker than full envy-freeness. An allocation is *envy-free up to any good* (EFX) if for any i, j where player i envies player j, removing *any* good from j's allocation would eliminate i's envy. Do EFX allocations always exist? This chapter takes the first steps toward answering this question.

5.1.1 Applications

The non-profit website Spliddit (www.spliddit.org) is one of the most promising applications of fair division theory [98]. Spliddit implements mechanisms for several fair division problems: rent division [84], taxi fare division, credit assignment (i.e., for a group project or academic paper) [50], task distribution [140, 37], and distribution of indivisible goods. These mechanisms are available for public use at no cost: users can simply log in, define what is to be divided, and enter their valuations. Since the site's launch in November 2014, there have been over 60,000 users [40]. The company Fair Outcomes, Inc. (http://www.fairoutcomes.com) offers fair division services in a similar vein.

Another compelling fair division application is allocating courses among students. Students have preferences regarding which courses they would like to take, but each course has a limited capacity.

[1] The algorithm of [119] was originally published in 2004 with a different property in mind, as the EF1 property was not proposed until 2011 by [35].

The Wharton School at the University of Pennsylvania now uses a theoretically-grounded mechanism titled Course Match to fairly allocate courses among MBA students, which has led to demonstrably higher satisfaction and perceived fairness among students [35, 36].

A major selling point of these services is that their solutions are *guaranteed* to satisfy certain fairness properties. For example, in the case of distribution of indivisible goods on Spliddit, users *know* that the solution will be envy-free up to one good and Pareto optimal [40]. Our hope is that further work in the area of fair division of indivisible goods will allow user-facing services like Spliddit, Fair Outcomes, Inc., and Course Match to offer users even stronger fairness guarantees.

5.1.2 Prior work

A detailed survey of the fair division literature is outside the scope of this chapter, and we discuss only the works most closely related to ours. See e.g., [27, 128, 28] for further background.

Lipton et al. [119] gave an algorithm whose solution is guaranteed to be EF1 for general valuations. By a *valuation*, we mean a function specifying a player's value for each bundle she might receive. By *general*, we mean that the only assumptions imposed on valuation functions are normalization (the value of the empty set is 0), and monotonicity (adding goods to a bundle cannot make it worse).

Their algorithm allocates the goods in rounds and ensures that the partial allocation at the end of each round is EF1. At the beginning of each round, an unenvied player is identified; if no such player exists, there must be a cycle of envy, and bundles can be swapped along such cycles until no cycles of envy remain. An arbitrary good is then given to this unenvied player. This player may become envied after receiving this good, but the envy could be eliminated by removing the good she just received (since she was unenvied prior to receiving that good). This ensures that whenever player i envies player j, the envy could be eliminated by removing the most recent good given to player j, so the resulting allocation is EF1.

Caragiannis et al. [40] studied the case where valuations are additive, meaning that each player's value for a set of goods is the sum of her values for the individual goods. They showed that the allocation maximizing the product of players' utilities (the maximum Nash welfare solution) is guaranteed to be both EF1 and Pareto optimal, assuming valuations are additive. In contrast, the algorithm of Lipton et al. [119] does not guarantee a Pareto optimal allocation.[2]

Caragiannis et al. [40] also proposed the fairness criterion of envy-freeness up to any good, and left the possible existence of EFX allocations as an open problem. We are not aware of any results regarding EFX allocations prior to this work.

However, there has been substantial follow-up work on this topic. Perhaps the largest breakthrough was that of [45], who showed that EFX allocations are guaranteed to exist for three agents with additive valuations.

[2]Suppose there are two players with additive valuations v_1, v_2 over three goods, a, b, c, where $v_1(\{a\}) = 3, v_1(\{b\}) = 2, v_1(\{c\}) = 4$ and $v_2(\{a\}) = 4, v_2(\{b\}) = 3, v_2(\{c\}) = 2$. The algorithm of Lipton et al. [119] could first allocate a to player 1, then b to player 2, and finally c also to player 2. The resulting allocation is EF1, but giving $\{c\}$ to player 1 and $\{a, b\}$ to player 2 would be better for both players.

	$n = 2$, add	$n = 2$, gen	$n \geq 2$, gen + id	$n > 2$, add	$n > 2$, gen
$\frac{1}{2}$ EFX	✓ (Thm. 5.4.3)	✓ (Thm. 5.4.3)	✓ (Thm. 5.4.2)	✓ (Thm. 5.6.1)	?
EFX	✓ (Thm. 5.4.3)	✓ (Thm. 5.4.3)	✓ (Thm. 5.4.2)	?	?
EFX + PO (nmu)	✓ (Thm. 5.5.5)	✗ (Thm. 5.5.6)	✓ (Thm. 5.5.4)	?	✗ (Thm. 5.5.6)

Table 5.1: A summary of our existence results. Here n is the number of players. "add", "gen", "id", and "nmu" refer to additive valuations, general valuations, identical valuations, and nonzero marginal utility, respectively. "✓" indicates that the type of allocation specified by the row is guaranteed to exist in the setting specified by the column, while "✗" indicates that we give a counterexample, and "?" indicates an open question.

We briefly describe several other models for fair division of indivisible goods. Brams et al. [25] and Aziz et al. [10] assumed that players express only an ordinal ranking over the goods, as opposed to exact values. Certain tasks become easier in this domain, but important information is arguably lost by only considering rankings. Randomized allocations have also been considered (e.g., [18, 37]), but this is not suitable for the applications we are most interested in, where the outcome is only used once. Dickerson et al. [62] took a probabilistic approach, and showed that envy-free allocations are likely to exist when the number of goods is at least a logarithmic factor larger than the number of players. While illuminating, this does not directly bear on our goal: determining when fair allocations are guaranteed to exist, and how they can be computed.

5.1.3 Our contributions

We consider the EFX property in a variety of contexts; our main existence results are given in Table 5.1.

Exponential query complexity lower bound

Section 5.3 presents our most technically involved result: an exponential lower bound on the number of value queries required by a deterministic algorithm to find an EFX allocation. This is done via a reduction from local search on a class of graphs known as the Odd graphs, for which we prove an exponential lower bound. In combination with results due to Dinh and Russell [63] and Valencia-Pabon and Vera [168], this yields an analogous exponential lower bound for randomized algorithms. Dobzinski et al. [64] also use a local search reduction to prove a lower bound on the number of value queries required to find a certain type of equilibrium in a simultaneous second price auction, for bidders with XOS (i.e., fractionally subadditive) valuations. We hope that this lower bound technique will be useful in other contexts as well.

Our lower bound holds even for two players with identical submodular valuations. In stark contrast, the algorithm of Lipton et al. [119] finds an EF1 allocation in polynomial time for general and possibly distinct valuations, and for any number of players. This suggests that EFX is indeed a significantly stronger fairness guarantee than EF1, and deserves further study.

Positive EFX results

Many of our positive results rely on the leximin solution. The leximin (a portmanteau of "lexicographic" and "maximin") solution selects the allocation which maximizes the minimum utility; then, if there are multiple allocations which achieve that minimum utility, it chooses among those the allocation which maximizes the second minimum utility, and so on. The leximin solution was developed as a metric of fairness in and of itself [149, 161, 160], and has been used before in fair division, though typically for randomized allocations (e.g. [18]).

In Section 5.4, we show that when players have general but identical valuations, a modification of the leximin solution is EFX. By *identical valuations*, we mean that all players have the same valuation. This result also yields a cut-and-choose-based protocol for two players with general and possibly distinct valuations that is guaranteed to produce an EFX allocation. This is consistent with our exponential lower bound, as it is well known that finding the leximin solution can require exponential time for general valuations (e.g. [66]).[3]

After a long line of follow-up work, it was recently shown

These positive results contrast with the state-of-the-art for possibly distinct valuations and three or more players, where even for additive valuations, the guaranteed existence of an EFX allocation remains an open question ("despite significant effort," according to [40]).

EFX and Pareto optimality

In Section 5.5, we consider Pareto optimality. In economics, an outcome is Pareto optimal (PO) if there is no way to make one player better off without making another player worse off. We show that even in simple cases, it is possible that no EFX allocation is also PO. However, these cases rely on a player having zero value for a good being added to her bundle.

We propose the assumption that adding a good to a player's bundle strictly improves the player's value for that bundle, and refer to this as "nonzero marginal utility". We view this as quite a weak assumption: especially in real-world settings, one might expect a player to always prefer to have a good than not.

Under this assumption, we show that for two players with additive valuations, the leximin solution is both EFX and PO.[4] We also show that for any number of players with general but identical valuations, the leximin solution is EFX and PO. Finally, we give a counterexample where for two players with distinct general valuations, no EFX allocation is PO (even assuming nonzero marginal utility).

Comparison to Spliddit in the two player case

Perhaps of most practical importance is our result that for two players with additive valuations and nonzero marginal utility, the leximin solution is both EFX and PO. This provides stronger

[3] We mention that Section 5.8 shows that for two players with additive valuations, an EFX allocation can be computed in polynomial time by a different method.

[4] When discussing the leximin solution for players with different valuations, we assume that each player's value for the entire set of goods is the same: were this not true, we could simply rescale the valuations as needed and use the leximin solution over the rescaled valuations.

	a	b	c
player 1	5	3	1
player 2	5	1	3

Figure 5.1: An instance where our algorithm provides stronger guarantees than the algorithm currently deployed on Spliddit. Here two players have additive valuations over three goods, a, b, and c. By symmetry, assume a is given to player 1. Spliddit selects the maximum Nash welfare solution, which gives $\{a, b\}$ to player 1 and $\{c\}$ to player 2. This is EF1 and PO, but not EFX, since player 2 would still envy player 1 after the removal of b. Our algorithm returns the unique (up to symmetry) EFX and PO allocation, which gives $\{a\}$ to player 1 and $\{b, c\}$ to player 2.

guarantees than the currently deployed algorithm on Spliddit, which selects the maximum Nash welfare solution, and only guarantees an allocation which is EF1 and PO.[5]

This manifests even in simple examples, such as the instance given by Figure 5.1. By symmetry, assume that player 1 receives good a. The maximum Nash welfare solution selects the allocation which maximizes the product of utilities: in this case, that would give player 1 a and b, and player 2 only c. This allocation is EF1, because player 2 would not envy player 1 if good a were removed from player 1's bundle. However, the allocation is not EFX, because player 2 *would* envy player 1 even if good b were removed from player 1's bundle.

In contrast, our algorithm returns the unique (up to symmetry) EFX and PO allocation, which gives a to player 1 and b and c to player 2. We suggest that this is also the more intuitively fair allocation. Furthermore, the assumption of nonzero marginal utility seems especially reasonable in the case of two players with additive valuations: if a player is truly indifferent to some good, one could imagine simply giving that good to the other player and excluding it from the fair division process entirely.[6] We do note that Spliddit's current algorithm does not require the assumption of nonzero marginal utility, however. Neither approach has a clear advantage in terms of computational efficiency: both the leximin solution and maximum Nash welfare solution are NP-hard to compute, even for two players with additive valuations.[7]

Finally, in Section 5.6 we propose an approximate version of EFX, and show that a $\frac{1}{2}$-EFX allocation always exists when players have subadditive (possibly distinct) valuations.

More broadly, our results span additive, submodular, subadditive, and general valuations, and identify separations between these classes from a fair division perspective. For example, we show that assuming nonzero marginal utility and two players with additive valuations, an allocation which is both EFX and PO is guaranteed to exist, while there is a counterexample for two players with general valuations. Such valuation classes have already played a central role in the development of algorithmic mechanism design over the past 15 years (e.g. [117]), and they may well prove equally important in the fair division of indivisible goods.

[5]Spliddit only considers additive valuations. This is because each user need only report m values to specify an entire additive valuation; in contrast, an exponential number of values can be required to specify a general valuation.

[6]A vindictive player might object to this: she may be unhappy with the other player receiving a good "for free", even if she has zero value for the good herself. We argue that this constitutes having nonzero value for the good, and that a player has zero value for a good only if she is truly indifferent.

[7]For two players with identical additive valuations, the leximin solution gives each player half the total value if and only if the valuation is a "yes" instance of the partition problem. The reduction is less obvious for maximum Nash welfare; see e.g. [148].

5.2 Model

Although we use the same general resource allocation framework (defined in Section 1.2), e.g., we have the same set of agents $N = \{1, \ldots, n\}$ and set of goods M with $m = |M|$. However, we deviate from Part I in several respects: most crucially, we assume that goods are indivisible, meaning that any feasible allocation $\mathbf{x}$ must have $x_{ij} \in \{0, 1\}$. For simplicity, we denote each agent's bundle as a set $A_i = \{j \in M : x_{ij} = 1\}$. We refer to an allocation as *partial* if only a subset $S \subset M$ of the goods are allocated. When "partial" is omitted, it means that all goods have been allocated.

Each player i has a value for each subset of M, specified as a *valuation function*[8] $v_i : 2^M \to \mathbb{R}_{\geq 0}$. Throughout the chapter, we assume normalization, meaning that $v_i(\emptyset) = 0$, and monotonicity (a.k.a. "free disposal"), meaning that $v_i(S) \leq v_i(T)$ whenever $S \subseteq T$. When we refer to "general valuations," we mean the set of all valuations that satisfy these two properties.

A special type of valuation is an *additive* valuation, where $v_i(S) = \sum_{g \in S} v_i(\{g\})$ for every $S \subseteq M$. Thus m parameters (one for each good) implicitly specify the 2^m values of the valuation. The majority of the literature on computational fair division, with both divisible and indivisible goods, focuses on additive valuations. There are also many interesting subclasses of valuations that generalize additive valuations. For example, our main lower bound result (Theorem 5.3.3) holds for *submodular* valuations, which are valuation functions v that satisfy "diminishing returns":

$$v(S \cup \{x\}) - v(S) \geq v(T \cup \{x\}) - v(T)$$

for every $S \subseteq T$ and $x \notin T$. One of our positive results, Theorem 5.6.1, will hold for *subadditive* valuations. A valuation v is subadditive if

$$v(S) + v(T) \geq v(S \cup T)$$

Every additive valuation is submodular, and every submodular valuation is subadditive.

Our objective is to find a "fair" allocation. Many different notions of fairness have been studied, with envy-freeness being one of the most prominent (see e.g., [27, 128, 28] for further background).

Definition 5.2.1. *An allocation A is* envy-free *if for all i and j,*

$$v_i(A_i) \geq v_i(A_j).$$

We say that i *envies* j if $v_i(A_i) < v_i(A_j)$. Unfortunately, an envy-free allocation does not always exist in the context of indivisible goods. This is clear even with two players and one good: the player who does not receive the good will envy the other, assuming they both have nonzero value for the good. Furthermore, determining whether an envy-free allocation exists is NP-complete [21]: with two players and identical additive valuations, this is the partition problem.

[8] Since there are no monetary payments in this model, the utility function is equal to the valuation function, so we just use "valuation function". This is contrast to the quasilinear utility model, where the utility is the valuation minus the payment.

Consequently, a relaxed version of envy-freeness has been studied, called envy-freeness up to one good [35, 40].

Definition 5.2.2. *An allocation A is* envy-free up to one good *(EF1) if for all i, j where i envies j,*[9]

$$\exists\ g \in A_j \text{ such that } v_i(A_i) \geq v_i(A_j \backslash \{g\}).$$

That is, i may envy j, but there is a good in j's bundle such that if it were removed, i would no longer envy j. An EF1 allocation always exists, and can be computed in polynomial time, even for general valuations [119].

Furthermore, Caragiannis et al. [40] showed that for additive valuations, an allocation which is both EF1 and Pareto optimal always exists; in particular, the maximum Nash welfare solution [129, 108, 148] is guaranteed to satisfy both properties. Caragiannis et al. [40] also proposed a new fairness notion, one which is strictly weaker than envy-freeness, but strictly stronger than EF1.

Definition 5.2.3. *An allocation A is* envy-free up to any good *(EFX) if, for all i, j,*

$$\forall g \in A_j, \quad v_i(A_i) \geq v_i(A_j \backslash \{g\}).$$

In words, EFX demands that removing *any* good from j's bundle would guarantee that i does not envy j. Next, we define the standard notion of Pareto optimality.

Definition 5.2.4. *An allocation A is Pareto optimal if there is no other allocation B where*

$$\forall i \in [n], \quad v_i(B_i) \geq v_i(A_i), \quad \text{and}$$
$$\exists j \in [n], \quad v_j(B_j) > v_j(A_j)$$

Finally, we define an approximate version of EFX. In Section 5.6, we will give an algorithm which produces a $\frac{1}{2}$-EFX allocation for any number of players with subadditive valuations.

Definition 5.2.5. *An allocation A is c-EFX if for all i, j,*

$$\forall g \in A_j, \quad v_i(A_i) \geq c \cdot v_i(A_j \backslash \{g\})$$

where $0 \leq c \leq 1$.

5.3 Query complexity lower bound

We begin with our most technically involved result: an exponential lower bound on the number of value queries required by any deterministic algorithm to compute an EFX allocation. Our lower

[9]The "where i envies j" clause is necessary, or the condition would technically fail when $A_j = \emptyset$. This is not an issue for the definition of EFX, however.

bound will hold even for two players, and even if their valuations are restricted to be identical and submodular.[10]

In Section 5.3.1, we introduce the local search problem that we will reduce from. In Section 5.3.2, we prove that finding an EFX allocation is at least as hard as solving local search on a particular class of graphs. In Section 5.3.3, we show that any deterministic algorithm which finds a local maximum on this class of graphs requires an exponential number of queries. This will imply that the problem of finding an EFX allocation has exponential query complexity as well. Finally, in Section 5.3.4, we extend this lower bound to randomized algorithms.

5.3.1 Local search

The Local Search problem takes as input an undirected graph $G = (V, E)$ and an oracle function $f : V \to \mathbb{R}$. The goal is to find a local maximum $a \in V$, where $f(a) \geq f(b)$ for all $(a, b) \in E$. Since there exists a global maximum, there must exist at least one local maximum. We are interested in the number of queries required to find a local maximum, where a query to $a \in V$ returns $f(a)$. Queries are the only method by which an algorithm can discover information about f (i.e., it is given as a "black box"). All other operations are free in this model—we count only the number of queries. Queries can be adaptive, with an algorithm's choice of which vertex to query next depending on the results of previous queries.

For a graph G, the *deterministic query complexity* of Local Search on G is the minimum number of queries required by any deterministic algorithm to solve Local Search on G (for a worst-case choice of f). Formally, let $D[LS(G)]$ be the deterministic query complexity of Local Search on G. Then $D[LS(G)] = \min_{\Gamma} \max_{f} T_{LS}(G, f, \Gamma)$, where the minimizer ranges over all deterministic algorithms Γ, the maximizer ranges over all functions $f : V \to \mathbb{R}$, and $T_{LS}(G, f, \Gamma)$ is the number of queries used by the algorithm Γ to find a local maximum of f on G.

The difficulty of local search depends on the graph G. The *Kneser graph* $K(n, k)$ is the graph whose vertices are the size k subsets of $[n]$, where two vertices are adjacent if and only if their corresponding subsets are disjoint. The star of our lower bound argument is the *Odd graph*, $K(2k+1, k)$. The most famous Odd graph is the Petersen graph (Figure 5.2).

5.3.2 Local search on $K(2k + 1, k)$ reduces to finding an EFX allocation

The EFX Allocation problem takes as input the set of players $N = [n]$, the set of goods $M = [m]$, and a list of valuations $(v_1, v_2, \ldots, v_n)$. In general, the goal is to find an EFX allocation, or determine that none exists. The only method by which an algorithm can discover information about the v_i's is through *value queries*, where upon querying the valuation v_i at the set S, the algorithm learns $v_i(S)$. Our lower bound applies even for a version of the problem that we will show to be total (Theorem 5.4.2), meaning that an EFX allocation is guaranteed to exist.

Consider the special case of the EFX Allocation problem where all valuations are identical. We will show in Section 5.4 that an EFX allocation is guaranteed to exist in this setting. We

[10] There is of course no hope for an exponential lower bound for additive valuations, since m value queries suffice to reconstruct an entire additive valuation.

can define the deterministic query complexity $D[EFX_{id}(n,m)]$ as the minimum number of queries required to find an EFX allocation for a set of players $N = [n]$ and a set of goods $M = [m]$, given a single valuation v where an EFX allocation is known to exist. Formally, $D[EFX_{id}(n,m)] = \min_{\Gamma} \max_{v} T_{EFX}(N, M, v, \Gamma)$, where $T_{EFX}(N, M, v, \Gamma)$ denotes the number of queries required by the algorithm Γ to find an EFX allocation for players N with valuation v over goods M. Since this is a special case of the general EFX ALLOCATION problem, the deterministic query complexity of the general EFX ALLOCATION problem is at least $D[EFX_{id}(n,m)]$.

We now state and prove our main result of Section 5.3.2. We will use $M = [2k+1]$ for some integer k.

Theorem 5.3.1. *The deterministic query complexity of the* EFX ALLOCATION *problem satisfies*

$$D[EFX_{id}(2, 2k+1)] \geq D[LS(K(2k+1, k))],$$

even for two players with identical submodular valuations.

Proof. Let $T = D[EFX_{id}(2, 2k+1)]$; then there exists an algorithm Γ for finding an EFX allocation which uses at most T queries, regardless of v. We will construct an algorithm Γ' for LOCAL SEARCH which also uses at most T queries, regardless of f. Formally, $\max_{v} T_{EFX}(\{1,2\}, M, v, \Gamma) = T$, and we will construct Γ' such that $\max_{f} T_{LS}(K(2k+1, k), f, \Gamma') \leq T$.

Define the algorithm Γ' on input $(K(2k+1, k), f)$ as follows. For each $S \subseteq [2k+1]$, define $v(S)$ as

$$v(S) = \begin{cases} 2|S| & \text{if } \ |S| < k \\ 2k - \dfrac{1}{1 + e^{(f(S))}} & \text{if } \ |S| = k \\ 2k & \text{if } \ |S| > k. \end{cases}$$

Then run Γ on $(\{1,2\}, [2k+1], v)$ to obtain an EFX allocation (A_1, A_2), and return A_1 if $|A_1| < |A_2|$ and A_2 otherwise. We will show that the returned set corresponds to a local maximum in $K(2k+1, k)$ (see Figure 5.2).

For brevity, define

$$\delta(S) = -\frac{1}{1 + e^{f(S)}}.$$

We note that $-1 < \delta(S) < 0$ for all S, and that $\delta(S)$ is strictly increasing with $f(S)$. Any other function satisfying these properties would work as well.

We first argue that any EFX allocation returned by Γ must give one player exactly k goods. Suppose that this were not the case. Then one player must receive fewer than k goods; without loss of generality, assume $|A_2| < k$, and thus $|A_1| > k+1$. Therefore $v(A_2) \leq 2k-2$ and $v(A_1) = 2k$.

For an arbitrary $g \in A_1$, we have $|A_1 \backslash \{g\}| > k$. Therefore there exists a $g \in A_1$ such that $v(A_1 \backslash \{g\}) = 2k > v(A_2)$, so the allocation cannot be EFX. Thus any EFX allocation must give one player exactly k goods. Therefore Γ will return a set of size k, which corresponds to a valid vertex of $K(2k+1, k)$.

Without loss of generality, assume $|A_1| = k+1$ and $|A_2| = k$. Then $v(A_1) = 2k$ and $v(A_2) =$

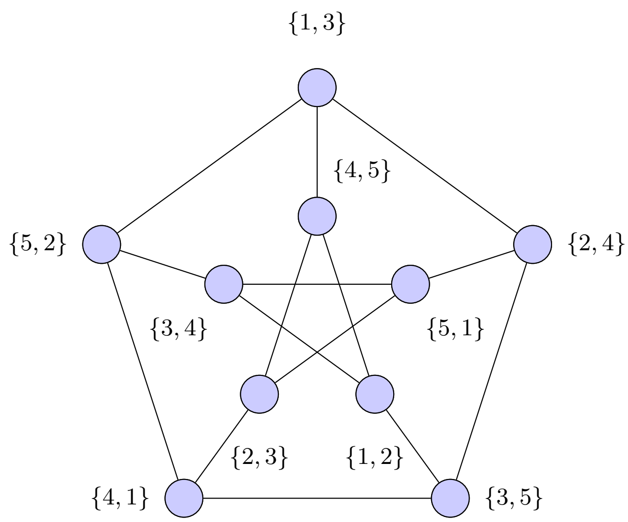

Figure 5.2: The graph shown is $K(2k+1,k)$ for $k = 2$, also known as the Petersen graph. Each vertex corresponds to a size 2 subset of [5]. Suppose the allocation where $A_1 = \{1, 2, 3\}$ and $A_2 = \{4, 5\}$ is EFX. Since $v(S_1) \geq v(S_2)$ if and only if $f(S_1) \geq f(S_2)$, we have $f(\{4, 5\}) \geq f(\{1, 2\})$, $f(\{4, 5\}) \geq f(\{2, 3\})$, and $f(\{4, 5\}) \geq f(\{1, 3\})$. Therefore {4,5} is a local maximum in this graph.

$2k + \delta(A_2) < 2k$, so $v(A_1) > v(A_2)$. Therefore the allocation $A = (A_1, A_2)$ is EFX if and only if $v(A_2) \geq v(A_1 \backslash \{g\})$ for all $g \in A_1$.

We can rewrite this condition as $v(A_2) \geq v(S)$ for all $S \subseteq A_1$ where $|S| = k$. For any $|S| = k$, we have $v(A_2) - v(S) = \delta(A_2) - \delta(S)$. Since δ is strictly increasing with $f(S)$, we have $v(A_2) \geq v(S)$ if and only if $f(A_2) \geq f(S)$. Therefore an allocation (A_1, A_2) is EFX if and only if $f(A_2) \geq f(S)$ for all $S \subseteq A_1$ where $|S| = k$.

Observe that $S \subseteq A_1$ if and only if $S \cap A_2 = \emptyset$. Therefore an allocation (A_1, A_2) is EFX if and only if $f(A_2) \geq f(S)$ for all $S \subseteq M$ where $|S| = k$ and $S \cap A_2 = \emptyset$. This is exactly the definition of A_2 being a local maximum in $K(2k+1,k)$. Therefore an allocation (A_1, A_2) is EFX if and only if A_2 is a local maximum in $K(2k+1,k)$.

Thus Γ' correctly solves Local Search. Furthermore, since Γ' uses no queries outside of running Γ, and Γ uses at most T queries, Γ' also uses at most T queries. Therefore

$$D[EFX_{id}(2, 2k+1)] \geq D[LS(K(2k+1,k))].$$

It remains to show that v is submodular. For any $S \subseteq M$ and $x \in M \backslash S$, we have

$$v(S \cup \{x\}) - v(S) = \begin{cases} 2 & \text{if } |S \cup \{x\}| < k \\ 2 + \delta(S \cup \{x\}) & \text{if } |S \cup \{x\}| = k \\ -\delta(S) & \text{if } |S \cup \{x\}| = k+1 \\ 0 & \text{if } |S \cup \{x\}| > k+1. \end{cases}$$

Therefore $v(S \cup \{x\}) - v(S)$ is non-increasing with $|S|$, since $-1 < \delta(S) < 0$ for all S. Thus $v(X \cup \{x\}) - v(X) \geq v(Y \cup \{x\}) - v(Y)$ whenever $|X| < |Y|$. If $X \subseteq Y$, either $|X| < |Y|$ or $X = Y$. When $X = Y$, we trivially have $v(X \cup \{x\}) - v(X) = v(Y \cup \{x\}) - v(Y)$. Thus we have $v(X \cup \{x\}) - v(X) \geq v(Y \cup \{x\}) - v(Y)$ whenever $X \subseteq Y$, and so v is submodular. □

5.3.3 Query complexity of local search on Odd graphs

In this section, we show that finding a local maximum on $K(2k+1, k)$ has exponential query complexity, completing our lower bound on the number of queries required to find an EFX allocation.[11]

The role of boundaries

For a graph $G = (V, E)$ and a set $S \subseteq V$, define the *boundary* $B(S)$ of S as the set of vertices that are not in S but are adjacent to a vertex in S. Formally, $B(S) = \{a \in V \backslash S : \exists b \in S, (a, b) \in E\}$. The next result, due to [120], implies that local search is hard in graphs that only have large boundaries.

Lemma 5.3.1 ([120])**.** *For any graph $G = (V, E)$ and integers t and c,*

$$D[LS(G)] \geq \min\big(t, \min_S \{|B(S)| : c - t \leq |S| \leq c\}\big).$$

Proof Sketch. We sketch a proof for the benefit of the reader. The proof follows an adversary argument. Let G_u be the subgraph induced by the still-unqueried vertices. While G_u remains connected, suppose the adversary simply returns increasing values for each query. Then the only way for a local maximum to be created is to query a vertex a after querying all of a's neighbors.

Furthermore, while G_u remains connected and contains at least one unqueried vertex, the most recently queried vertex a must have an unqueried neighbor b: if not, G_u must have been disconnected prior to the most recent query. The adversary is free to toggle which of a and b is a local maximum (or possibly neither, if there are more unqueried vertices). Thus while at least one vertex has not been queried and the graph of unqueried vertices remains connected, it cannot be determined where the graph has a local maximum.

Thus the only strategy to counteract the adversary is to perform a sort of binary search. First, we must disconnect the graph of unqueried vertices. At least one of the resulting components must contain a local maximum, and Llewellyn et al. [120] show how we can always identify one such component based on the query results so far. Thus we can recurse on that component, and the process repeats. Llewellyn et al. [120] call this the *separation game*. An example of the separation game being played on a path is given by Figure 5.3.

By this logic, we will have to eventually disconnect a "fairly large" component: if it is too small, the adversary is free to place the local maximum in another larger component. Specifically, Llewellyn et al. [120] show that for any integers t and c, the adversary can force us to either query t vertices, or disconnect a set of vertices S where $c - t \leq |S| \leq c$.

[11] A similar lower bound for local search on $K(2k+1, k)$ was proved (using different arguments) in [64].

In order to disconnect a set of vertices S, every vertex on the boundary of S must be queried. Thus at least $\min\big(t, \min_S\{|B(S)| : c - t \leq |S| \leq c\}\big)$ must be queried, as claimed. □

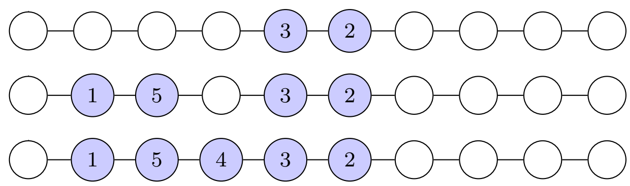

Figure 5.3: An example of the separation game played on a path. After two central vertices are queried, returning values 3 and 2 as shown, we know that there must be a local maximum in the left half. Next, we bisect the left half by querying two more vertices, which return values 1 and 5. At this point, we know that either the vertex with value 5 or the vertex immediately to its right must be a local maximum, and only one more query is required to determine which. In this case, the local maximum is the vertex with value 5.

Boundaries of Kneser graphs

In light of Lemma 5.3.1 and our interest in Kneser graphs, the natural next step it to understand boundary sizes in Kneser graphs. The next lemma is due to Zheng [178].

Lemma 5.3.2 ([178])**.** *Let $\mu_G(r)$ denote $\min_{|S|=r} |B(S)|$. Then for all $1 \leq r \leq \binom{n}{k}$,*

$$\mu_{K(n,k)}(r) \geq \binom{n}{k} - \frac{1}{r}\binom{n-1}{k-1}^2 - r$$

We include a proof for completeness. In it, we make use of the following variant of the Erdős-Ko-Rado theorem. Call the set families $\mathcal{X}$ and $\mathcal{Y}$ *cross-intersecting* if $X \cap Y \neq \emptyset$ for all $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$.

Lemma 5.3.3 ([125])**.** *If $\mathcal{X}$ and $\mathcal{Y}$ are cross-intersecting families of size-k subsets of $[n]$, then*

$$|\mathcal{X}||\mathcal{Y}| \leq \binom{n-1}{k-1}^2$$

Note that the inequality in Lemma 5.3.3 holds with equality (for $k \leq n/2$) when $\mathcal{X}$ and $\mathcal{Y}$ both consist of all subsets of size k that contain the element 1.

Proof. (Of Lemma 5.3.2.) For any S, we can partition V into S, $B(S)$, and $V \backslash (S \cup B(S))$. An example of this is shown in Figure 5.4. Consider an arbitrary $a \in V \backslash (S \cup B(S))$. We know that $a \notin S$ and $a \notin B(S)$, so there is no $b \in S$ where $(a, b) \in E$. Therefore for all $a \in V \backslash (S \cup B(S))$ and $b \in S$, a and b are not adjacent. Recall that a and b are adjacent in $K(n, k)$ if $a \cap b = \emptyset$. Therefore for all $a \in V \backslash (S \cup B(S))$ and $b \in S$, $a \cap b \neq \emptyset$. Thus S and $V \backslash (S \cup B(S))$ are cross-intersecting families.

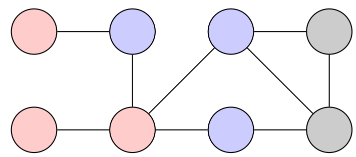

Figure 5.4: The partitioning of an arbitrary graph into S, $B(S)$, and $V\backslash(S \cup B(S))$. In this example, S is the set of pink vertices, $B(S)$ is the set of blue vertices, and $V\backslash(S \cup B(S))$ is the set of gray vertices.

Therefore by Lemma 5.3.3, we have $|S||V\backslash(S \cup B(S))| \leq \binom{n-1}{k-1}^2$. Let $r = |S|$. Then $|V\backslash(S \cup B(S))| \leq \frac{1}{r}\binom{n-1}{k-1}^2$. Therefore for all S,

$$\begin{aligned}|B(S)| &= |V| - |V\backslash(S \cup B(S))| - |S| \\ &= \binom{n}{k} - |V\backslash(S \cup B(S))| - r \\ &\geq \binom{n}{k} - \frac{1}{r}\binom{n-1}{k-1}^2 - r\end{aligned}$$

and so $\mu_{K(n,k)}(r) = \min_{|S|=r} |B(S)| \geq \binom{n}{k} - \frac{1}{r}\binom{n-1}{k-1}^2 - r$. □

We will only be interested in $K(2k+1, k)$, so we will simply write

$$\mu(r) = \mu_{K(2k+1,k)}(r).$$

Similarly, let

$$\beta(r) = \binom{2k+1}{k} - \frac{1}{r}\binom{2k}{k-1}^2 - r.$$

Then $\mu(r) \geq \beta(r)$ for all r.

We next prove a lemma building on Lemma 5.3.2.

Lemma 5.3.4. *Let $r_{max} = \binom{2k}{k-1}$. Then for the graph $K(2k+1, k)$ and any $r^* \leq r_{max}$,*

$$\min_S\{|B(S)| : r^* \leq |S| \leq r_{max}\} \geq \beta(r^*)$$

Proof. We begin by examining the expression $\beta(r) - \beta(r-1)$:

$$\begin{aligned}\beta(r) - \beta(r-1) &= -\frac{1}{r}\binom{2k}{k-1}^2 - r + \frac{1}{r-1}\binom{2k}{k-1}^2 + r - 1 \\ &= \left(\frac{1}{r-1} - \frac{1}{r}\right)\binom{2k}{k-1}^2 - 1 \\ &= \frac{1}{r(r-1)}\binom{2k}{k-1}^2 - 1.\end{aligned}$$

Therefore $\beta(r) - \beta(r-1) \geq 0$ when $r(r-1) \leq \binom{2k}{k-1}^2$. If $r \leq r_{max}$, then $r(r-1) < r^2 \leq r_{max}^2 = \binom{2k}{k-1}^2$. Thus $\beta(r) \geq \beta(r-1)$ when $r \leq r_{max}$. Iterating this inequality yields $\beta(r^*) \leq \beta(r)$

whenever $r^* \leq r \leq r_{max}$.

We can rewrite $\min_S\{|B(S)| : r^* \leq |S| \leq r_{max}\}$ as

$$\begin{aligned}\min_S\{|B(S)| : r^* \leq |S| \leq r_{max}\} &= \min_{r:\ r^* \leq r \leq r_{max}} \min_{|S|=r} |B(S)| \\ &= \min_{r:\ r^* \leq r \leq r_{max}} \mu(r) \\ &\geq \min_{r:\ r^* \leq r \leq r_{max}} \beta(r)\end{aligned}$$

where the last step is due to Lemma 5.3.2. Since $\beta(r^*) \leq \beta(r)$ whenever $r^* \leq r \leq r_{max}$, $\min_{r:\ r^* \leq r \leq r_{max}} \beta(r) = \beta(r^*)$. Therefore $\min_S\{|B(S)| : r^* \leq |S| \leq r_{max}\} \geq \beta(r^*)$, as required. $\square$

Local search on $K(2k+1,k)$

We are now ready to prove our result on $D[LS(K(2k+1,k))]$.

Theorem 5.3.2. *For all k,*

$$D[LS(K(2k+1,k))] \in \Omega\left(\frac{1}{k}\binom{2k+1}{k}\right).$$

Proof. Let $c = r_{max} = \binom{2k}{k-1}$, and let $t = \frac{1}{2k+1} r_{max}$. so $c - t = \frac{2k}{2k+1} r_{max}$. Then by Lemma 5.3.1,

$$D[LS(K(2k+1,k))] \geq$$

$$\min\left(\frac{1}{2k+1} r_{max}, \min_S\left\{|B(S)| : \frac{2k}{2k+1} r_{max} \leq |S| \leq r_{max}\right\}\right)$$

By Lemma 5.3.4,

$$\begin{aligned}\min_S\left\{|B(S)| : \frac{2k}{2k+1} r_{max} \leq |S| \leq r_{max}\right\} &\geq \beta\left(\frac{2k}{2k+1} r_{max}\right) \\ &= \binom{2k+1}{k} - \frac{2k+1}{2k \cdot r_{max}} r_{max}^2 - \frac{2k}{2k+1} r_{max} \\ &\geq \binom{2k+1}{k} - \left(\frac{2k+1}{2k} + 1\right) r_{max}.\end{aligned}$$

Using the identity $\binom{n}{k} = \frac{n}{k}\binom{n-1}{k-1}$ for any n, k, we have $\binom{2k+1}{k} = \frac{2k+1}{k}\binom{2k}{k-1} = \frac{2k+1}{k} r_{max}$. Thus we have

$$\begin{aligned}\min_S\left\{|B(S)| : \frac{2k}{2k+1} r_{max} \leq |S| \leq r_{max}\right\} & \\ &\geq \left(\frac{2k+1}{k} - \frac{2k+1}{2k} - 1\right) r_{max} \\ &= \frac{4k+2-2k-1-2k}{2k} r_{max} \\ &= \frac{1}{2k} r_{max}.\end{aligned}$$

Therefore,

$$\begin{aligned} D[LS(K(2k+1,k))] &\geq \min\left(\frac{1}{2k+1}r_{max}, \frac{1}{2k}r_{max}\right) \\ &= \frac{1}{2k+1}r_{max} \\ &\in \Omega\left(\frac{1}{k}r_{max}\right). \end{aligned}$$

Since $\binom{2k+1}{k} = \frac{2k+1}{k}r_{max}$, we have

$$D[LS(K(2k+1,k))] \in \Omega\left(\frac{1}{k}\binom{2k+1}{k}\right).$$

□

Theorem 5.3.1 and Theorem 5.3.2 together imply our main result of Section 5.3.

Theorem 5.3.3. *The deterministic query complexity of the* EFX Allocation *problem satisfies*

$$D[EFX_{id}(2, 2k+1)] \in \Omega\left(\frac{1}{k}\binom{2k+1}{k}\right),$$

even for two players with identical submodular valuations.

5.3.4 Randomized query complexity

Our reduction from Local Search to EFX Allocation also yields an exponential lower bound for randomized algorithms for free, thanks to results due to Dinh and Russell [63] and Valencia-Pabon and Vera [168]. Let $R[LS(G)]$ be the minimum number of queries required to solve Local Search on G by a randomized algorithm: the algorithm should output a local maximum with probability at least 2/3 (say) over its internal coin flips. Formally, $R[LS(G)] = \min_{\Gamma_R} \max_f T(G, f, \Gamma_R)$, where Γ_R ranges over the set of randomized algorithms.

Similarly, let $R[EFX_{id}(2, 2k+1)]$ be the minimum number of queries required by a randomized algorithm to find an EFX allocation for two players with identical valuations, and $2k+1$ goods (again with correctness probability at least 2/3, say).

Theorem 5.3.4 ([63])**.** *If $G = (V, E)$ is a vertex transitive graph with diameter d, then*

$$R[LS(G)] \in \Omega\left(\frac{\sqrt{|V|}}{d \cdot \log |V|}\right)$$

Since $K(2k+1, k)$ is vertex transitive, the last piece of the puzzle is the following theorem,

Theorem 5.3.5 ([168])**.** *The diameter of $K(2k+1, k)$ is k.*

With these two tools in hand, Theorem 5.3.6 requires only a short proof.

Theorem 5.3.6. *The randomized query complexity of the* EFX Allocation *problem satisfies*

$$R[EFX_{id}(2, 2k+1)] \in \Omega\Big(\sqrt{\binom{2k+1}{k}\frac{1}{k^2}}\Big)$$

even for two players with identical submodular valuations.

Proof. Since $|V| = \binom{2k+1}{k}$ and $\log\Big(\binom{2k+1}{k}\Big) \in O(\log(4^k)) = O(k)$, we have

$$R[LS(K(2k+1, k))] \in \Omega\Big(\sqrt{\binom{2k+1}{k}\frac{1}{dk}}\Big)$$

by Theorem 5.3.4. Thus by Theorem 5.3.5, we have

$$R[LS(K(2k+1, k))] \in \Omega\Big(\sqrt{\binom{2k+1}{k}\frac{1}{k^2}}\Big)$$

The reduction used to prove that $D[EFX_{id}(2, 2k+1)] \geq D[LS(K(2k+1, k))]$ can equivalently be used to show that

$$R[EFX_{id}(2, 2k+1)] \geq R[LS(K(2k+1, k))].$$

Therefore $R[EFX_{id}(2, 2k+1)] \in \Omega\Big(\sqrt{\binom{2k+1}{k}\frac{1}{k^2}}\Big)$. □

While this bound is not as strong as our deterministic lower bound (Theorem 5.3.3), it does establish that even a randomized algorithm requires an exponential number of queries to find an EFX allocation.

5.4 Existence of EFX allocations for general but identical valuations

We mentioned in the previous section that an EFX allocation is guaranteed to exist when all players have the same valuation: this section proves that claim. Specifically, we show that a modified version of the leximin solution is guaranteed to be EFX for general but identical valutions. This also yields a cut-and-choose-based protocol for two players with general and possibly distinct valuations.

5.4.1 The leximin solution

The leximin solution selects the allocation which maximizes the minimum utility of any player. If there are multiple allocations which achieve that minimum utility, it chooses among those the one which maximizes the second minimum utility, and so on. This implicitly specifies a comparison operator $\prec$, which is given by Algorithm 3, and constitutes a total ordering over allocations.

The operator $\prec$ takes as input two allocations A and B, and the list of player valuations $(v_1 ... v_n)$. The players are ordered by utility, and according to some arbitrary but consistent tiebreak for players with the same utility (for example, by player index). The comparison terminates when the ℓth player in A's ordering X^A has different utility from the ℓth player in B's ordering X^B.

The leximin solution is the global maximum under this ordering. The leximin solution is trivially PO, since if it were possible to improve the utility of one player without decreasing the utility of any other player, the new allocation would be strictly larger under $\prec$.

Standard leximin is not EFX

Unfortunately, the standard leximin solution is not always EFX, even for identical valuations. Consider two players with the same (non-additive) valuation v over two goods a and b. Define v by

$$v(S) = \begin{cases} 0 & \text{if } S = \{a\} \\ 1 & \text{if } S = \{b\} \\ 2 & \text{if } S = \{a, b\} \end{cases}$$

By symmetry, suppose without loss of generality that player 1 receives good b. Define the allocation A by $A_1 = \{b\}$ and $A_2 = \{a\}$, and define the allocation B by $B_1 = \{a, b\}$ and $B_2 = \emptyset$.

Since player 2 (the minimum utility player) is indifferent between A and B, leximin selects allocation B because it maximizes the value of player 1 (the second minimum utility player). However, A is EFX, while B is not: player 2 envies player 1 even after the removal of a from B_1.[12]

5.4.2 The leximin++ solution

Our fix is that after maximizing the minimum utility, we maximize the size of the bundle of the player with minimum utility, before maximizing the second minimum utility. Then we maximize the second minimum utility, followed by the size of the second minimum utility bundle, and so on. Thus giving good a to the lower utility player (player 2) is preferable, and so the EFX allocation A is chosen over B.

We call this the *leximin++ solution*. The leximin++ solution induces a comparison operator $\prec_{++}$, also given in Algorithm 3. Similarly to $\prec$, the players are ordered by increasing utility, and then according to an arbitrary but consistent tiebreak among players with the same utility.[13] The comparison terminates when the ℓth player in X^A differs in utility or bundle size from the ℓth player in X^B, with utility being checked before bundle size.

It may not be immediately clear that $\prec_{++}$ specifies a total ordering, but this is in fact the case. The proof of Theorem 5.4.1 appears in Section 5.7.

Theorem 5.4.1. *The comparison operator $\prec_{++}$ specifies a total ordering.*

[12] This example will be relevant again in Section 5.5 as an instance where there is no allocation which is both EFX and PO.

[13] The tiebreak method must be consistent to ensure that $\prec_{++}$ is a total ordering. Consider two players with the same valuation v, and a single good a where $v(\{a\}) = 0$. Suppose $a \in A_1$. Since both players have zero utility, if the tiebreak method were not required to be consistent, both $\{1, 2\}$ and $\{2, 1\}$ would be valid player orderings for A. Consider running $A \prec_{++} A$. If player 2 were considered first in the A on the left, and player 1 were considered first in the A on the right, the operator would see that player 1 has a larger bundle than player 2, and return true.

Algorithm 3 Leximin and Leximin++ comparison operators

function LEXIMINCMP$(A, B, (v_1 ... v_n))$ ▷ Returns true if $A \prec B$ (strictly)
  $X^A \leftarrow$ ordering of players by increasing utility $v_i(A_i)$, then by some arbitrary but consistent tiebreak method for players with the same utility
  $X^B \leftarrow$ corresponding ordering of players under B
  for each $\ell \in [n]$ **do**
    $i \leftarrow X^A_\ell$ ▷ ℓth player in the ordering X_A
    $j \leftarrow X^B_\ell$ ▷ ℓth player in the ordering X_B
    if $v_i(A_i) \neq v_j(B_j)$ **then**
      return $v_i(A_i) < v_j(B_j)$
  return false ▷ In this case, A and B are equal

function LEXIMIN++CMP$(A, B, (v_1 ... v_n))$ ▷ Returns true if $A \prec_{++} B$ (strictly)
  $X^A \leftarrow$ same as in LEXIMINCMP
  $X^B \leftarrow$ same as in LEXIMINCMP
  for each $\ell \in [n]$ **do**
    $i \leftarrow X^A_\ell$
    $j \leftarrow X^B_\ell$
    if $v_i(A_i) \neq v_j(B_j)$ **then**
      return $v_i(A_i) < v_j(B_j)$
    if $|A_i| \neq |B_j|$ **then**
      return $|A_i| < |B_j|$
  return false

We are now ready to prove our main result of this section.

Theorem 5.4.2. *For general but identical valuations, the leximin++ solution is EFX.*

Proof. Let A be an allocation that is not EFX. We will show that A is not the leximin++ solution.

Since A is not EFX, there exist players i, j and $g \in A_j$ where $v(A_i) < v(A_j \backslash \{g\})$. Then any player with utility $\min_k v(A_k)$ must also have utility strictly less than $v(A_j \backslash \{g\})$, so assume with loss of generality that $i = \arg\min_k v(A_k)$. If there are multiple players with minimum utility in A, let i be the one considered last in the ordering X^A, according to the arbitrary but consistent tiebreak method.

Define a new allocation B where $B_i = A_i \cup \{g\}$, $B_j = A_j \backslash \{g\}$, and $B_k = A_k$ for all $k \notin \{i, j\}$. We will show that $A \prec_{++} B$.

Let S be the set of players appearing before i in X^A. We know i is considered last among the players with minimum utility by assumption, so S is exactly the set of players with minimum utility, other than i. Note that neither i nor j are in S.

Since the only bundles that differ between allocations A and B are that of i and j, we have $A_k = B_k$ for all $k \in S$. Thus for all $k \in S$, $v(B_k) = v(A_k) = v(A_i)$. Since $v(B_j) > v(A_i)$, j must occur after every player in S in X^B.

Because $A_i \subset B_i$, we have $v(B_i) \geq v(A_i)$. If $v(B_i) > v(A_i)$, i must occur after every player in S in X^B, since $v(B_i) > v(B_k)$ for all $k \in S$. If $v(B_i) = v(A_i)$, i is still considered after every player in S according to the arbitrary but consistent tiebreak method. Thus i occurs after every player in S

Algorithm 4 Find an EFX allocation for two players with general valuations via cut-and-choose

```
function CutAndChoose(m, v_1, v_2)
    (A_1, A_2) ← Leximin++Solution(2, m, v_1)    ▷ Player 1 uses the leximin++ solution to cut,
    if v_2(A_1) ≥ v_2(A_2) then                  ▷ and player 2 chooses.
        return (A_2, A_1)
    else
        return (A_1, A_2)
```

in X^B in either case, which shows that the first $|S|$ players in X^B are the players in S, in the same order they occur in X^A.

Therefore the leximin++ comparison will not have terminated before reaching position $|S|+1$ in the orderings. Let T be the set of players appearing after i in X^A: note that $j \in T$. By assumption, of the players with minimum utility in A, i appears last in X^A. Therefore all players after i in X^A do not have minimum utility, so $v(A_k) > v(A_i)$ for all $k \in T$. Recall that $v(B_j) > v(A_i)$ and that for all $k \in T \backslash \{j\}$, $v(B_k) = v(A_k)$. Thus $v(B_k) > v(A_i)$ for all $k \in T$.

We know that $X^A_{|S|+1} = i$. If $X^B_{|S|+1} = i$, we have $|A_i| < |B_i|$ (and possibly also $v(A_i) < v(B_i)$), so $A \prec_{++} B$ returns true. If $X^B_{|S|+1} = k$ for some $k \neq i$, then $k \in T$. Therefore $v(A_i) < v(B_k)$, so $A \prec_{++} B$ returns true in this case as well.

Since $A \prec_{++} B$, A cannot be the leximin++ solution. Therefore the leximin++ solution must be EFX. □

We now show how Theorem 5.4.2 can easily be used to find an EFX allocation for two players with general and possibly distinct valuations.[14] Our algorithm for this follows from the observation that any player can partition the goods into k bundles that are mutually EFX from her viewpoint, simply by computing the leximin++ solution with k copies of herself.

Algorithm 4 is a straightforward adaptation of the cut-and-choose protocol. Player 1 partitions the goods into two bundles using the leximin++ solution, and player 2 chooses her favorite bundle.

Theorem 5.4.3. *For two players with general (not necessarily identical) valuations, Algorithm 4 returns an EFX allocation.*

Proof. By Theorem 5.4.2, the allocation is EFX from player 1's viewpoint regardless of which bundle she receives. Player 2 receives her favorite bundle, so the resulting allocation is EFX from her viewpoint as well. □

5.4.3 Limitations of leximin++

Unfortunately, the leximin++ solution may not be EFX when players have different valuations. For example, consider two players with valuations $v_1(S) = |S|$ and $v_2(S) = \epsilon|S|$, for some small $\epsilon > 0$. As long as player 1 receives at least one good, she will have utility at least 1. However, player 2 will always have utility less than 1 for a suitably small ϵ. Thus the leximin++ solution gives a single

[14] The two-player case is not trivial. For example, our lower bound in Theorem 5.3.3 already applies with two players (even with identical valuations).

good to player 1 and the rest to player 2, which will cause player 1 to envy player 2 in violation of EFX.

One might hope that this could be remedied by assuming that all players have the same value for the entire set of goods (or rescaling valuations as necessary if this is not the case). Unfortunately, the set of additive valuations given by Figure 5.5 thwarts this hope.

	a	b	c	d
player 1	14	3	2	1
player 2	7	6	4	3
player 3	20	0	0	0

Figure 5.5: An example where the leximin++ solution fails to be EFX even when all players have the same value for the entire set of goods.

We claim that the allocation $A = (\{b,d\}, \{c\}, \{a\})$ is the only allocation where all players have utility at least 4. To see this, first observe that good a must go to player 3, or player 3 has zero utility. Then the only way to give players 1 and 2 each utility at least 4 is to give $\{b,d\}$ to player 1 and $\{c\}$ to player 2.

Since A is the only allocation which gives all players utility at least 4, A must be the leximin++ solution. However, A is not EFX, because $v_2(\{c\}) < v_2(\{b,d\}\backslash\{d\})$.

We mentioned at the beginning of this section that the leximin solution is trivially PO. The leximin++ solution does not share this guarantee. Indeed, this is necessary in order for the leximin++ solution to be EFX, since it is impossible to simultaneously guarantee EFX and Pareto optimality, even for identical valuations (Theorem 5.5.2). However, that example relies on zero value goods. We will show in the next section that if zero value goods are disallowed, the leximin solution becomes EFX as well as PO in two contexts.

5.5 Pareto optimality

In this section, we examine when EFX and Pareto optimality can be guaranteed simultaneously. We begin by showing that if a player is wholly indifferent to a good being added to her bundle (zero marginal utility), EFX and Pareto optimality can be mutually exclusive even in simple cases.

Theorem 5.5.1. *If zero marginal utility is allowed, there exist additive valuations where no EFX allocation is also PO, even for two players.*

Proof. Consider the following additive valuations:

	a	b	c
player 1	2	1	0
player 2	2	0	1

Since $v_1(\{c\}) = 0$ but $v_2(\{c\}) > 0$, $c \in A_2$ in any PO allocation. Similarly, $b \in A_1$ in any PO allocation.

By symmetry, assume without loss of generality that $a \in A_1$, so $A_1 = \{a, b\}$ and $A_2 = \{c\}$. Then $v_2(\{c\}) = 1$, but $v_2(A_1 \backslash \{b\}) = v_2(\{a\}) = 2$, so the allocation is not EFX.

Therefore no allocation is both EFX and PO. □

A similar example exists for general and identical valuations. This example was also used in Section 5.4 to show that the leximin solution may not be EFX when zero marginal utility is allowed.

Theorem 5.5.2. *If zero marginal utility is allowed, there exist general and identical valuations where no EFX allocation is also PO, even for two players.*

Proof. Consider two players with the same valuation v, and two goods a and b. Define v by

$$v(S) = \begin{cases} 0 & \text{if } S = \{a\} \\ 1 & \text{if } S = \{b\} \\ 2 & \text{if } S = \{a, b\} \end{cases}$$

By symmetry, assume without loss of generality that $b \in A_1$. If $A_1 = \{a, b\}$, then $v(A_2) = v(\emptyset) = 0$, but $v(A_1 \backslash \{a\}) = v(\{b\}) > 0$, so the allocation is not EFX.

Therefore in any EFX allocation, $a \in A_2$. But $v(\{a\}) = v(\emptyset) = 0$ and $v(\{a, b\}) > v(\{b\})$. Thus giving a to player 1 strictly increases player 1's value, without changing player 2's value, so the allocation is not PO.

Therefore no allocation is both EFX and PO. □

On the other hand, if valuations are required to be additive and identical, it is possible to guarantee EFX and Pareto optimality simultaneously, even with zero marginal utility. However, this is an extremely restrictive setting that we mention mostly for completeness; we consider this a very minor result. The proof of Theorem 5.5.3 appears in Section 5.7.

Theorem 5.5.3. *For additive and identical valuations, there exists an allocation that is both EFX and PO (even allowing zero marginal utility).*

5.5.1 Nonzero marginal utility

The negative results of Theorem 5.5.1 and Theorem 5.5.2 both break down if players are assumed to have strictly positive utility for any good being added to their bundle. Formally, we say that a valuation v has nonzero marginal utility if for every set $S \subset [m]$ and $g \notin S$, $v(S \cup \{g\}) - v(S) > 0$.

We feel that this is a reasonable assumption in practice, as $v(S \cup \{g\}) - v(S)$ is allowed to be arbitrarily small, and one might expect players in real world situations to always prefer to have a good than not.

Positive results from leximin

Under the assumption of nonzero marginal utility, the leximin solution is guaranteed to be both EFX and PO for any number of players with general but identical valuations, and for two players with (possibly distinct) additive valuations.

Theorem 5.5.4. *For general but identical valuations with nonzero marginal utility, the leximin solution is EFX and PO.*

Proof. We follow a very similar analysis to the proof of Theorem 5.4.2. Let A be an allocation that is not EFX. Then there exist players i, j and $g \in A_j$ where $v(A_i) < v(A_j \backslash \{g\})$. Again assume without loss of generality that $i = \arg\min_k v(A_k)$, and if there are multiple players with minimum utility in A, let i be the one considered last in the ordering X^A.

Define the same new allocation B where $B_i = A_i \cup \{g\}$, $B_j = A_j \backslash \{g\}$, and $B_k = A_k$ for all $k \notin \{i, j\}$. When zero marginal utility is allowed, the leximin++ modification of considering bundle size is necessary because otherwise if $v_i(B_i) = v_i(A_i)$, it could be the case that $B \prec A$. When zero marginal utility is disallowed, this modification is not necessary because $v_i(B_i) > v_i(A_i)$ always.

The proof of Theorem 5.4.2 can be used nearly verbatim to show that $A \prec B$ (simply omit the sentences handling the case where $v(B_i) = v(A_i)$, since we now have $v(B_i) > v(A_i)$, due to the nonzero marginal utility of v). Thus A is not the leximin solution, so the leximin solution is EFX.

As noted before, the leximin solution is trivially Pareto optimal, since if any player could be made better off without hurting any other player, that new allocation would be strictly larger under $\prec$. ☐

We now show that assuming nonzero marginal utility, the leximin solution is EFX and PO for two players with additive valuations. For this theorem, we will assume that $v_i([m]) = 1$ for all i: were this not the case, we could easily define $v'_i(S) = v_i(S)/v_i([m])$, and find the leximin solution according to v'. Additivity is necessary for Theorem 5.5.5 so that $v_i(A_1) < v_i(A_2)$ implies $v_i(A_1) < 1/2$, and so that $v_i(A_1) \geq v_i(A_2)$ implies $v_i(A_1) \geq 1/2$.

The proof is similar to those of Theorem 5.4.2 and Theorem 5.5.4, in that we consider an arbitrary allocation A that is not EFX, and show that it cannot be the leximin solution by constructing an allocation B such that $A \prec B$. However, the allocation B is constructed differently here.

Theorem 5.5.5. *For two players with additive valuations (not necessarily identical) with nonzero marginal utility, the leximin solution is EFX and PO.*

Proof. Let A be an allocation that is not EFX. Then there exist players i, j and $g \in A_j$ where $v_i(A_i) < v_i(A_j \backslash \{g\})$. Without loss of generality, assume $i = 1$ and $j = 2$.

We know that $v_1(A_1) < v_1(A_2)$, so $v_1(A_1) < 1/2$. If $v_2(A_2) < v_2(A_1)$, the players could swap bundles to increase both of their utilities, so A could not be the leximin solution. Therefore assume $v_2(A_2) \geq v_2(A_1)$, and so $v_2(A_2) \geq 1/2$.

Define two new bundles $S_1 = A_1 \cup \{g\}$ and $S_2 = A_2 \backslash \{g\}$. Then define a new allocation B where $B_1 = \underset{S \in \{S_1, S_2\}}{\arg\min}\ v_2(S)$ and $B_2 = \underset{S \in \{S_1, S_2\}}{\arg\max}\ v_2(S)$.

Since player 2 received her favorite of S_1 and S_2, we still have $v_2(B_2) \geq 1/2$. We have $v_1(S_2) = v_1(A_2 \backslash \{g\}) > v_1(A_1)$ by our original assumption that A is not EFX, and we have $v_1(S_1) = v_1(A_1 \cup \{g\}) > v_1(A_1)$ by the nonzero marginal utility of v_1. Therefore regardless of which bundle player 1 receives, $v_1(B_1) > v_1(A_1)$.

Thus B has a higher minimum utility than A, so A cannot be the leximin solution. Therefore the leximin solution is EFX in this setting, and it remains trivially PO. ☐

Assuming nonzero marginal utility, Theorem 5.5.5 provides stronger guarantees than the currently deployed algorithm on Spliddit, which only guarantees an EF1 and PO allocation. As described in Section 5.1.3, this manifests even in simple cases.

We also argue that the assumption of nonzero marginal utility is particularly reasonable in the case of two players with additive valuations, since if a player is truly indifferent to some good, perhaps that good could simply be given to the other player and excluded from the fair division process entirely.

Counterexample for two players with general valuations

Finally, we show that EFX and Pareto optimality cannot be guaranteed simultaneously for general and distinct valuations, even with the assumption of nonzero marginal utility.

Theorem 5.5.6. *There exist general valuations where no EFX allocation is also PO, even for two players with nonzero marginal utility.*

Proof. We construct a set of valuations for which there is no EFX allocation that is also PO.

Let $n = 2$ and $M = \{a, b, c, d, e\}$. Let $\alpha_1 = \{a\}, \beta_1 = \{b, d\}, \gamma_1 = \{a, c, d\}$ and $\alpha_2 = \{b\}, \beta_2 = \{a, d\}, \gamma_2 = \{b, d, e\}$. The key properties will be $\alpha_1 \subset \beta_2 \subset \gamma_1$ and $\alpha_2 \subset \beta_1 \subset \gamma_2$.

Define each player's valuation v_i by

$$v_i(S) = \begin{cases} 3 + \epsilon(|S| - 3) & \text{if } \gamma_i \subseteq S \\ 2 + \epsilon(|S| - 2) & \text{if } \beta_i \subseteq S \text{ and } \gamma_i \not\subseteq S \\ 1 + \epsilon(|S| - 1) & \text{if } \alpha_i \subseteq S \text{ and } \beta_i, \gamma_i \not\subseteq S \\ \epsilon|S| & \text{otherwise} \end{cases}$$

where ϵ is some small positive value (.1 would suffice). Adding a good to a bundle always increases the value of the bundle by at least ϵ, so v_i satisfies nonzero marginal utility. Also, note that the valuations are symmetric across players, since α_i, β_i, and γ_i are symmetric across players.

We have the following implications:

$$\begin{aligned} \gamma_i \not\subseteq S &\implies v_i(S) < 3 \\ \beta_i, \gamma_i \not\subseteq S &\implies v_i(S) < 2 \\ \alpha_i, \beta_i, \gamma_i \not\subseteq S &\implies v_i(S) < 1 \end{aligned}$$

By Theorem 5.4.3, an EFX allocation $A = (A_1, A_2)$ must exist. Suppose $\gamma_i \subseteq A_i$ for some i: by symmetry, suppose $i = 1$. Since $\beta_1 \cap \beta_2 \cap \gamma_1 \cap \gamma_2 = \{d\} \neq \emptyset$, we have $\beta_2, \gamma_2 \not\subseteq A_2$, so $v_2(A_2) < 2$. Furthermore, β_2 is a strict subset of A_1: specifically, $\beta_2 \subseteq A_1 \backslash \{c\}$. Therefore $v_2(A_1 \backslash \{c\}) \geq v_2(\beta_2) = 2$, which is strictly larger than $v_2(A_2)$. Therefore if $\gamma_i \subseteq A_i$ for either i, A is not EFX.

Now suppose $\beta_i \subseteq A_i$ for some i: again suppose $i = 1$. Similarly, $\beta_2, \gamma_2 \not\subseteq A_2$. In this case, we also have $\alpha_2 \not\subseteq A_2$, since $\alpha_2 \cap \beta_1 \neq \emptyset$. Therefore $v_2(A_2) < 1$. Since $\alpha_2 \subseteq A_1 \backslash \{d\}$, we have $v_2(A_1 \backslash \{d\}) \geq v_2(\alpha_1) = 1$, which is strictly larger than $v_2(A_2)$. Therefore if $\beta_i \subseteq A_i$ for either i, A

is not EFX. Since A is EFX by assumption, we have $\beta_i, \gamma_i \not\subseteq A_i$ for both i, and so $v_i(A_i) < 2$ for both i.

We next claim that $\alpha_i \subseteq A_i$ for both i. Suppose $\alpha_1 \not\subseteq A_1$: then $\alpha_1 \subseteq A_2$. Therefore $v_1(A_1) < 1$, and $v_1(A_2) \geq 1$, so player 1 envies player 2. If there exists $g \in A_2 \backslash \alpha_1$, then g could be removed and player 1 would still envy player 2. Thus if $|A_2| \geq 2$, A is not EFX, so we have $|A_2| = 1$. But then $\alpha_2 \subseteq A_1$ and $|A_1| \geq 2$, so player 1 is envied in violation of EFX. Thus we have $\alpha_1 \subseteq A_1$, and by symmetry, $\alpha_2 \subseteq A_2$.

One of the players has at least three goods; by symmetry, suppose $|A_1| \geq 3$. Since $\alpha_1 \subseteq A_1$ and $\beta_1, \gamma_1, \alpha_2 \not\subseteq A_1$, we have $A_1 = \{a, c, e\}$ and $A_2 = \{b, d\}$.

Consider the allocation $B = (B_1, B_2) = (\{a, c, d\}, \{b, e\})$. Player 2 is indifferent between $\{b, d\}$ and $\{b, e\}$, so $v_2(B_2) = v_2(A_2)$. But $\gamma_1 \subseteq B_1$, so $v_1(B_1) > v(A_1)$. Thus player 1 is strictly better off in B, and no player is worse off. Therefore A is not PO, and so no EFX allocation is PO. □

One last attempt to salvage EFX and PO in this setting might be to require a strict ranking over bundles, i.e., not allow player 2 to be indifferent between $\{b, d\}$ and $\{b, e\}$. However, even that would not work, because we can easily set $v_2(\{b, e\}) > v_2(\{b, d\})$, in which case both players are strictly better off in B.

This counterexample and our query complexity lower bound show that EFX is a very demanding fairness property, even for two players. In the next section, we complement these negative results by showing that an approximate version of EFX is satisfiable for any number of players with subadditive valuations.

5.6 Existence of $\frac{1}{2}$-EFX allocations for subadditive valuations

The possible existence of EFX allocations for possibly distinct valuations and $n \geq 3$ remains an open question, even for additive valuations. However, we are able to achieve an approximate version of EFX, for any number of players with (possibly distinct) subadditive valuations. Recall that an allocation A is c-EFX if for all i, j, and for all $g \in A_j$, $v_i(A_i) \geq c \cdot v_i(A_j \backslash \{g\})$. In words, an allocation is c-EFX if for all i, j, and $g \in A_j$, i's value for her own bundle is at least c times her value for j's bundle after removing g. For example, 1-EFX is equivalent to standard EFX. In this section, we give an algorithm that is guaranteed to return a $\frac{1}{2}$-EFX allocation for any number of players with subadditive valuations.

To describe our algorithm, we must first define the *envy graph*. The envy graph of an allocation A has a vertex for each player, and a directed edge from i to j if player i envies player j. Here we mean full envy (i.e. $v_i(A_i) < v_i(A_j)$), not just envy in violation of EFX. It will be necessary for the envy graph in our algorithm to be acyclic; we now show that we can always ensure this. The following lemma is adapted from Lipton et al. [119].

Lemma 5.6.1. *Let $A = (A_1, A_2 ... A_n)$ be a c-EFX allocation with envy graph $G = (V, E)$, where G contains a cycle. Then there exists another allocation $B = (B_1, B_2 ... B_n)$ with envy graph H where B is also c-EFX, and H has no cycles.*

Proof. We first show that there exists another c-EFX allocation $A' = (A'_1 ... A'_n)$ with envy graph G', where G' has strictly fewer edges than G.

Let $c = (1, 2...|c|)$ be a cycle in G. Thus $v_i(A_i) < v_i(A_{(i \bmod |c|)+1})$ for all $i \in c$. Define a new allocation A' where $A'_i = A_{(i \bmod |c|)+1}$ for all i, and let $G' = (V', E')$ be the envy graph for A'. It is clear that A' is a permutation of A.

Suppose A' is not c-EFX: then there exist $i, j \in N$ and $g \in A'_j$ where $v_i(A') < c \cdot v_i(A'_j \backslash \{g\})$. Since A' is a permutation of A, there exists $k \in N$ where $A_k = A'_j$, so $v_i(A'_i) < c \cdot v_i(A_k \backslash \{g\})$. Observe that $v_i(A'_i) > v_i(A_i)$ if $i \in c$, and $v_i(A'_i) = v_i(A_i)$ otherwise. Thus $v_i(A_i) \leq v_i(A'_i) < c \cdot v_i(A_k \backslash \{g\})$, and so A is also not c-EFX. Therefore if A is c-EFX, then A' is also c-EFX.

Note that the number of edges from $V' \backslash c$ into c is unchanged. Also, the number of edges from c into $V' \backslash c$ has decreased or stayed the same, since the utility of every player in c has strictly increased. Furthermore, for each $i \in c$, the number of players in c whom i envies has decreased by at least one. This shows that G' has strictly fewer edges than G.

If G' still contains a cycle, we can apply this process again to obtain G'', G''', and so on. Since the number of edges strictly decreases each time, we can apply this process at most $|E|$ times before we obtain a envy graph without a cycle. ☐

Algorithm 5 gives pseudocode for our algorithm. Initially all goods are in the pool P, and we proceed in rounds until P is empty, maintaining the invariant that the partial allocation at the end of each round is EFX. The function EliminateEnvyCycles uses Lemma 5.6.1 to ensure that the graph at the beginning of each round is acyclic. Since the envy graph is acyclic, we can always find an unenvied player j, and give an arbitrary good g^* from P to her.

It is possible that this will cause another player i to envy j in violation of $\frac{1}{2}$-EFX. In this case, we return all of i's current bundle to P, and let i's new bundle be just $\{g^*\}$. The key insight is that in order for i to go from not envying j to envying j in violation of $\frac{1}{2}$-EFX, adding g^* to A_j must have caused $v_i(A_j)$ to at least double. We will use that fact, along with the subadditivity of v_i, to show that $v_i(\{g^*\})$ must be larger than i's value for her bundle at the beginning of the round. Thus if i envies any player, it remains consistent with $\frac{1}{2}$-EFX. Any envy directed towards i will be fully EFX, since i will only have one good.

On each round, either P decreases in size (in the case where g^* remains with j), or the sum of utilities increases (in the case where g^* is instead given to i because i envies j in violation of $\frac{1}{2}$-EFX). Thus we can use a potential function argument to show that Algorithm 5 terminates (although it may take a non-polynomial number of rounds).

Theorem 5.6.1. *For subadditive valuations, Algorithm 5 returns a $\frac{1}{2}$-EFX allocation.*

Proof. We refer to each iteration of the while-loop as a round. We first show that the partial allocation at the end of each round is $\frac{1}{2}$-EFX. Then we will show that the algorithm is guaranteed to terminate.

Let A_k^ℓ be the bundle of player k at the beginning of round ℓ, and let B_k^ℓ denote the bundle of player k just before EliminateEnvyCycles is run on round ℓ. Let $A^\ell = (A_1^\ell ... A_n^\ell)$ and $B^\ell = (B_1^\ell ... B_n^\ell)$. In this proof, we use k and k' to denote a generic player; i and j refer exclusively to the variables in the while-loop.

Algorithm 5 Find an $\frac{1}{2}$-EFX allocation for n players with subadditive valuations

1: **function** GetApxEFXAllocation$(n, m, (v_1 ... v_n))$
2: $\quad P \leftarrow [m]$ $\triangleright$ Initially, all goods are in the pool
3: $\quad$ **for each** $i \in [n]$ **do**
4: $\quad\quad A_i \leftarrow \emptyset$
5: $\quad$ **while** $P \neq \emptyset$ **do**
6: $\quad\quad g^* \leftarrow \text{pop}(P)$ $\triangleright$ Remove an arbitrary good from P,
7: $\quad\quad j \leftarrow \text{FindUnenviedPlayer}(A_1, A_2 ... A_n)$ $\triangleright$ and give it to an unenvied player
8: $\quad\quad A_j \leftarrow A_j \cup \{g^*\}$
9: $\quad\quad$ **if** $\exists i \in [n], g \in A_j$ such that $v_i(A_i) < \frac{1}{2} v_i(A_j \backslash \{g\})$ **then**
10: $\quad\quad\quad P \leftarrow P \cup A_i$ $\triangleright$ Return i's old allocation to the pool,
11: $\quad\quad\quad A_j \leftarrow A_j \backslash \{g^*\}$ $\triangleright$ and give i just $\{g^*\}$
12: $\quad\quad\quad A_i \leftarrow \{g^*\}$
13: $\quad\quad (A_1, A_2 ... A_n) \leftarrow \text{EliminateEnvyCycles}(A_1, A_2 ... A_n)$ $\triangleright$ Ensure the envy graph is acyclic
14: $\quad$ **return** $(A_1, A_2 ... A_n)$

We proceed by induction on ℓ. Initially, all players have empty bundles, which trivially satisfies $\frac{1}{2}$-EFX. Thus assume the partial allocation at the beginning of round ℓ is $\frac{1}{2}$-EFX. We will show that the partial allocation at the beginning of round $\ell + 1$ is $\frac{1}{2}$-EFX. The partial allocation at the beginning of round $\ell + 1$ $A^{\ell+1}$ is equal to EliminateEnvyCycles(B^ℓ). Thus by Lemma 5.6.1, it suffices to show that B^ℓ is $\frac{1}{2}$-EFX.

If the body of the if-statement (lines 10-12) is not executed, the allocation B^ℓ is $\frac{1}{2}$-EFX by definition. Thus assume the body of the if-statement is executed. Then $B_j^\ell = A_j^\ell$, because g^* was added and then removed. Thus for all $k \neq i$, $B_k^\ell = A_k^\ell$.

We say that a pair (k, k') is $\frac{1}{2}$-EFX in B^ℓ if $v(B_k^\ell) \geq \frac{1}{2} v(B_{k'} \backslash \{g\})$ for all $g \in B_{k'}^\ell$. We know that A^ℓ is $\frac{1}{2}$-EFX by assumption. Therefore since $B_k^\ell = A_k^\ell$ for all $k \neq i$, all pairs (k, k') where $k \neq i$ and $k' \neq i$ remain $\frac{1}{2}$-EFX in B^ℓ. Furthermore, since $B_i^\ell = \{g^*\}$, the pair (k, i) is $\frac{1}{2}$-EFX for all players k, since $B_i^\ell \backslash \{g\} = \emptyset$ for all $g \in B_i^\ell$.

It remains only to show that the pairs (i, k) are $\frac{1}{2}$-EFX for all players k. We do this by showing that $v_i(B_i^\ell) > v_i(A_i^\ell)$. The fact that this inequality is strict will be important later in showing that the algorithm terminates.

We know that j was unenvied at the beginning of round ℓ, so $v_i(A_i^\ell) \geq v_i(A_j^\ell)$. Since the body of the if-statement executed, we also know that there exists $g \in A_j^\ell \cup \{g^*\}$ such that $v_i(A_i^\ell) < \frac{1}{2} v_i(A_j^\ell \cup \{g^*\} \backslash \{g\})$. Thus $v_i(A_i^\ell) < \frac{1}{2} v_i(A_j^\ell \cup \{g^*\})$, which will be all we need. Therefore,

$$v_i(A_i^\ell) < \frac{1}{2} v_i(A_j^\ell \cup \{g^*\}) \tag{5.1}$$

$$\leq \frac{1}{2}(v_i(A_j^\ell) + v_i(\{g^*\})) \tag{5.2}$$

$$\leq \frac{1}{2}(v_i(A_i^\ell) + v_i(\{g^*\})) \tag{5.3}$$

where 5.2 follows from 5.1 due to v_i being subadditive, and 5.3 follows from 5.2 due to $v_i(A_i^\ell) \geq$

$v_i(A_j^\ell)$. Therefore,

$$\begin{aligned} v_i(A_i^\ell) - \frac{1}{2}v_i(A_i^\ell) &< \frac{1}{2}v_i(\{g^*\}) \\ v_i(A_i^\ell) &< v_i(\{g^*\}) \\ v_i(A_i^\ell) &< v_i(B_i^\ell) \end{aligned}$$

Consider an arbitrary player $k \neq i$. Since A^ℓ is $\frac{1}{2}$-EFX, we have $v_i(A_i^\ell) \geq \frac{1}{2}v_i(A_k^\ell \backslash \{g\})$ for all $g \in A_k^\ell$. Since $v_i(B_i^\ell) > v_i(A_i^\ell)$ and $B_k^\ell = A_k^\ell$ for all $k \neq i$, we have $v_i(B_i^\ell) \geq \frac{1}{2}v_i(B_k^\ell \backslash \{g\})$ for all $g \in B_k^\ell$ as well. Therefore the pair (i,k) is $\frac{1}{2}$-EFX for all players k.

Thus every pair of players is $\frac{1}{2}$-EFX in B^ℓ, so B^ℓ is $\frac{1}{2}$-EFX. This shows that the partial allocation at the end of each round is $\frac{1}{2}$-EFX, and so any allocation returned by the algorithm is $\frac{1}{2}$-EFX.

It remains to show that Algorithm 5 terminates. We use a potential function argument. For round ℓ, define

$$\phi(\ell) = \sum_{k=1}^{n} v(A_k^\ell).$$

We noted above that if round ℓ falls under Case 2, only i's bundle changes, and we have the strict inequality $v_i(B_i^\ell) > v_i(A_i^\ell)$. Therefore $v_i(A_i^{\ell+1}) > v_i(A_i^\ell)$. Thus if round ℓ falls under Case 2, we have $\phi(\ell+1) - \phi(\ell) > 0$.

If round ℓ falls under Case 1, only j's bundle changes, and we have $v_j(A_j^{\ell+1}) \geq v_i(A_i^\ell)$. Therefore if round ℓ falls under Case 1, we have $\phi(\ell+1) - \phi(\ell) \geq 0$.

In any round which falls under Case 1, $|P|$ decreases by one. Therefore if m rounds pass without Case 2 occurring, P becomes empty, and the algorithm terminates. Thus while the algorithm has not terminated, Case 2 must occur at least once every m rounds, and so $\phi(\ell+m) - \phi(\ell) > 0$ for all ℓ.

The number of possible partial allocations is at most $(n+1)^m$: each good can be given to one of the n players, or left in the unallocated pool. Thus the number of distinct values ϕ can take on is at most $(n+1)^m$, and so ϕ can increase at most that many times. Thus after $m(n+1)^m$ rounds, the algorithm must have terminated. □

Finally, we briefly show that $\frac{1}{2}$-EFX and EF1 are incomparable, meaning that neither property implies the other. Recall that an allocation A is EF1 if for all i,j where $A_j \neq \emptyset$, there exists $g \in A_j$ where $v_i(A_i) \geq v_i(A_j \backslash \{g\})$.

Consider the additive valuations on the left, and let $A = (\{a,b\},\{c\})$. A is EF1 because $v_2(A_2) \geq v_2(A_1 \backslash \{a\})$, but A is not $\frac{1}{2}$-EFX because $v_2(A_2) < \frac{1}{2}v_2(A_1 \backslash \{b\})$.

Now consider the valuations on the right, and let $A = (\{a,b,c\},\{d\})$. Then A is not EF1, because $v_2(A_2) < v_2(A_1 \backslash \{g\})$ for all $g \in A_1$, but A is $\frac{1}{2}$-EFX, because $v_2(A_2) \geq \frac{1}{2}v_2(A_1 \backslash \{g\})$ for all $g \in A_1$.

	a	b	c
player 1	3	1	0
player 2	3	0	1

	a	b	c	d
player 1	1	1	1	1
player 2	1	1	1	1

5.7 Additional proofs

Proof of Theorem 5.4.1. To show that $\prec_{++}$, we need to show that for any allocations A, B, and C, $A \prec_{++} A$ is false, and that ($A \prec_{++} B$ and $B \prec_{++} C$) implies $A \prec_{++} C$.

We first show that $A \prec_{++} A$ is false. The key fact is that for a given allocation A, there is only one possible ordering of the players X^A: were this not true, $\prec_{++}$ could fail to produce a total ordering.[15] Therefore on each iteration, the same player is considered from each copy of A. Thus on each iteration, the two bundles compared will be the same, so $A \prec_{++} A$ never terminates until it passes through all $\ell \in [n]$ and returns false at the very end.

It remains to show that ($A \prec_{++} B$ and $B \prec_{++} C$) implies $A \prec_{++} C$. Suppose $A \prec_{++} B$ and $B \prec_{++} C$. Let ℓ_1, ℓ_2, and ℓ_3 be the iterations on which $A \prec_{++} B$, $B \prec_{++} C$ and $A \prec_{++} C$ terminate, respectively. For $x \in \{1, 2, 3\}$, let $i_x = X^A_{\ell_x}$, $j_x = X^B_{\ell_x}$, and $k_x = X^C_{\ell_x}$.

Since $A \prec_{++} B$ terminates on iteration ℓ_1, we have $v(A_{i_1}) < v(B_{j_1})$ or $|A_{i_1}| < |B_{j_1}|$. Similarly, since $B \prec_{++} C$ terminates on iteration ℓ_2, we have $v(B_{i_2}) < v(C_{j_2})$ or $|B_{i_2}| < |C_{j_2}|$.

First we argue that $\ell_3 \geq \min(\ell_1, \ell_2)$. Suppose $\ell < \min(\ell_1, \ell_2)$: then $A \prec_{++} B$ and $B \prec_{++} C$ do not terminate until after iteration ℓ_3. Therefore $v(A_{i_3}) = v(B_{j_3})$, $|A_{i_3}| = |B_{j_3}|$, $v(B_{j_3}) = v(C_{k_3})$, and $|B_{j_3}| = |C_{k_3}|$. Therefore $v(A_{i_3}) = v(C_{k_3})$ and $|A_{i_3}| = |C_{k_3}|$, so $A \prec_{++} C$ could not have terminated on iteration ℓ_3, which is a contradiction. Therefore $\ell_3 \geq \min(\ell_1, \ell_2)$. We proceed by case analysis.

Case 1: $\ell_1 < \ell_2$. Since $B \prec_{++} C$ did not terminate until after iteration ℓ_1, we have $v(B_{j_1}) = v(C_{k_1})$ and $|B_{j_1}| = |C_{k_1}|$. Therefore $v(A_{i_1}) < v(C_{k_1})$ or $|A_{i_1}| < |C_{k_1}|$. We know that $A \prec_{++} C$ cannot have terminated prior to ℓ_1, since $\ell_3 \geq \min(\ell_1, \ell_2) = \ell_1$. Therefore $A \prec_{++} C$ will terminate on iteration ℓ_1 and return true, so $A \prec_{++} C$ holds in Case 1.

Case 2: $\ell_2 < \ell_1$. This case is similar. Since $A \prec_{++} B$ did not terminate until after iteration ℓ_2, we have $v(A_{i_2}) = v(B_{j_2})$ and $|A_{i_2}| = |B_{j_2}|$. Therefore $v(A_{i_2}) < v(C_{k_2})$ or $|A_{i_2}| < |C_{k_2}|$. We know that $A \prec_{++} C$ cannot have terminated prior to ℓ_2, since $\ell_3 \geq \min(\ell_1, \ell_2) = \ell_2$. Therefore $A \prec_{++} C$ will terminate on iteration ℓ_2 and return true, so $A \prec_{++} C$ holds Case 2.

Case 3: $\ell_1 = \ell_2$. In this case we have $i_1 = i_2$, $j_1 = j_2$, and $k_1 = k_2$. Therefore

$$v(A_{i_1}) < v(B_{j_1}) \text{ or } \big(v(A_{i_1}) = v(B_{j_1}) \text{ and } |A_{i_1}| < |B_{j_1}|\big), \text{ and}$$
$$v(B_{j_1}) < v(C_{k_1}) \text{ or } \big(v(B_{j_1}) = v(C_{k_1}) \text{ and } |B_{j_1}| < |C_{k_1}|\big)$$

Note that $v(A_{i_1}) \leq v(B_{j_1})$ and $v(B_{j_1}) \leq v(C_{k_1})$. Therefore if either $v(A_{i_1}) < v(B_{j_1})$ or $v(B_{j_1}) < v(C_{k_1})$, we have $v(A_{i_1}) < v(C_{k_1})$. We know $A \prec_{++} C$ cannot have terminated before $\ell_1 = \ell_2$ since $\ell_3 \geq \min(\ell_1, \ell_2)$, so if $v(A_{i_1}) < v(C_{k_1})$, $A \prec_{++} C$ terminates on iteration ℓ_1 and returns true.

Thus assume $v(A_{i_1}) = v(B_{j_1})$ and $v(B_{j_1}) = v(C_{k_1})$: then $|A_{i_1}| < |B_{j_1}|$ and $|B_{j_1}| < |C_{k_1}|$. Therefore $v(A_{i_1}) = v(C_{k_1})$ and $|A_{i_1}| < |C_{k_1}|$, so $A \prec_{++} C$ terminates on iteration ℓ_1 and returns true. Therefore $A \prec_{++} C$ in Case 3. This shows that ($A \prec_{++} B$ and $B \prec_{++} C$) implies $A \prec_{++} C$,

[15] Consider two players with identical valuations and one good a, where $v(\{a\}) = 0$. Let $A = (\emptyset, \{a\})$. Suppose both (1,2) and (2,1) are valid orderings of the players according to A, and suppose we run $A \prec_{++} A$ with the left hand side A using the ordering (1,2) and the right hand side A using (2,1). Then at $\ell = 1$, $\emptyset$ from the left hand side A will be compared with $\{a\}$ from the right hand side A, and $A \prec_{++} A$ will return true.

and completes the proof. ☐

Proof of Theorem 5.5.3. Let Z be the set of all goods g where $v(\{g\}) = 0$. Therefore for all $g \in M\backslash\{Z\}$, we have $v(\{g\}) > 0$, so v has nonzero marginal utility over the set of goods $M\backslash\{Z\}$.

Let $A = (A_1...A_n)$ be the leximin allocation over $M\backslash\{Z\}$. By Theorem 5.5.4, A is EFX and PO over $M\backslash\{Z\}$.

Let i be the minimum utility player in A. Define a new allocation B over all of M where $B_i = A_i \cup \{Z\}$ and $B_j = A_j$ for all $j \neq i$. Since $v(Z) = 0$, we have $v(B_j) = v(A_j)$ for all j. Therefore since i had minimum utility in A, i also has minimum utility in B.

To see that B is EFX, consider arbitrary players j and k, and any $g \in B_k$. If $i \neq k$, we have $A_k = B_k$. Since A is EFX, we have $v(B_j) = v(A_j) \geq v(A_k\backslash\{g\}) = v(B_k\backslash\{g\})$. If $i = k$, then $v(B_j) \geq v(B_k) \geq v(B_k\backslash\{g\})$, since i has minimum utility in B. This shows that B is EFX.

To see that B is PO, observe that the way the goods in Z are allocated has no effect on the values of the bundles. Therefore the goods in Z have no effect on the Pareto optimality of the allocation, so the Pareto optimality of B follows directly from the Pareto optimality of A. ☐

5.8 A setting where an EFX allocation can be computed quickly

Finally, we describe a setting in which an EFX allocation always exists and can be computed in polynomial time (counting both the value queries and all additional computation done by an algorithm). Our result will hold when players have *additive valuations with identical rankings*, meaning that all players agree on the relative ordering of individual goods. This is, for all players i and j, and for all goods g_1 and g_2, $v_i(g_1) \geq v_i(g_2)$ whenever $v_j(g_1) \geq v_j(g_2)$. This will also yield a polynomial time algorithm for computing an EFX allocation for two players with additive (possibly distinct) valuations.

Requiring identical rankings is not as strong as requiring identical valuations. For example, let $v_1(g_1) = 1, v_1(g_2) = 2, v_1(g_3) = 4$ and $v_2(g_1) = 2, v_2(g_2) = 3, v_2(g_3) = 4$. Then the rankings are identical, but $v_1(\{g_1, g_2\}) < v_1(g_3)$, whereas $v_2(\{g_1, g_2\}) > v_2(g_3)$.

While strong, there are certainly real-world contexts where this assumption makes sense. For example, if the goods are apartments (with differing square footage), airline tickets (with differing numbers of stops and classes of service), or baseball pitchers (with differing statistics), it is plausible that buyers generally agree on which goods are more valuable than others, but disagree on the exact values of these goods.

Our algorithm (Algorithm 6) is reminiscent of our algorithm for finding a $\frac{1}{2}$-EFX allocation for any number of players with subadditive valuations from Section 5.6, in that we allocate the goods in rounds and ensure that the envy graph is acyclic at the beginning of each round. However, here we never return goods to the pool, and allocate the goods in descending order of value.

Recall that Lemma 5.6.1 gives a process that can be used to ensure the envy graph is acyclic: if an envy cycle exists, bundles can be permuted along this cycle such the number of edges in the

Algorithm 6 Find an EFX allocation for additive valuations with identical ranking

```
function GetEFXAllocationSameRanking(n, m, (v_1...v_n))
    P ← Sorted([m])                          ▷ Sort in descending order: P_1 = max(P)
    for each i ∈ [n] do
        A_i ← ∅
    for each i ∈ [m] do
        j ← FindUnenviedPlayer(A_1, A_2...A_n)
        A_j ← A_j ∪ {P_i}
        (A_1, A_2...A_n) ← EliminateEnvyCycles(A_1, A_2...A_n)
    return (A_1, A_2...A_n)
```

envy graph decreases by at least one. The function EliminateEnvyCycles repeatedly performs this process until the envy graph is acyclic.

Theorem 5.8.1. *For additive valuations with identical rankings, Algorithm 6 terminates with an EFX allocation in $O(mn^3)$ time.*

Proof. We first argue that at all times, $v_i(A_j) - v_i(A_i) \leq v_i(g^*)$ where g^* is the good most recently added to what is currently A_j. Since bundles may have been permuted by EliminateEnvyCycles, j may not have been in possession of what is currently A_j at the time g^* was added. This does not affect the proof, however: it is sufficient to interpret A_j as "the bundle that currently belongs to j". Thus instead of saying "i did not envy j at the time", we will say "i did not envy A_j at the time".

Observe that a good is only allocated to a player whom no one envies. Thus directly before g^* was added to A_j, i did not envy A_j: at that point $v_i(A_j) - v_i(A_i) \leq 0$. Therefore directly after g^* was given to j, $v_i(A_j) - v_i(A_i) \leq v_i(g^*)$. Since $v_i(A_i)$ can only have grown since then, we have $v_i(A_j) - v_i(A_i) \leq v_i(g^*)$ until a new good is added to A_j.

Since the goods are allocated in decreasing order of value, the good most recently added to A_j must also be the least valuable good in A_j. Therefore at all times, $v_i(A_j) - v_i(A_i) \leq \min\limits_{g \in A_j} v_i(g)$, and so $v_i(A_i) \geq v_i(A_j) - \min\limits_{g \in A_j} v_i(g)$. For additive valuations, this is equivalent to $v_i(A_i) \geq v_i(A_j \backslash \{g\})$ for all $g \in A_j$. Therefore the allocation at all times is EFX, so the final allocation is EFX.

Finally, we show that Algorithm 6 terminates in $O(mn^3)$ time. Each time a good is allocated, any edges added to the envy graph must point to the recipient. Thus at most n edges are added to the envy graph on each round, and so at most mn edges are added to the graph over the course of the algorithm. Each time a cycle is detected and bundles are permuted along that cycle using Lemma 5.6.1, at least one edge is removed from the graph. Therefore this process is performed at most mn times. Each time this process is performed, we may have to compute a large part of the envy graph, which can take $O(n^2)$ time. Thus the overall running time bound is $O(mn^3)$. □

This algorithm is easily generalizable to general valuations under the condition that all players agree on a single ordering of the marginal values of the goods. Specifically, there must be an ordering of the goods $(g_1, g_2, \ldots, g_m)$ where for any set S, any player i, and all j, we have $v_i(S \cup \{g_j\}) \geq v_i(S \cup \{g_{j+1}\})$. This ordering must be fixed across all sets S. Then instead of allocating goods in descending order of value, we allocate goods in descending order of marginal value, and the analog of Theorem 5.8.1 holds, with essentially the same proof.

Finally, we note that Algorithm 6 can be used to compute an EFX allocation for two players with additive (possibly distinct) valuations in polynomial time. We use a cut-and-choose argument similar to that of Theorem 5.4.3: player 1 runs Algorithm 6 with two copies of herself to find an allocation which will be EFX from her viewpoint, regardless of which bundle she receives. Then player 2 chooses her favorite bundle in the resulting allocation, so the allocation will be fully envy-free from her viewpoint.

5.9 Conclusion and future work

In this chapter, we provided the first general results on the fairness concept of envy-freeness up to any good. Our most technically involved result was an exponential lower bound on the number of queries required by any deterministic algorithm to find an EFX allocation, via a reduction from local search. To complete the lower bound, we proved an exponential lower bound on the number of queries required to find a local maximum on $K(2k+1, k)$. We used results from Dinh and Russell [63] and Valencia-Pabon and Vera [168] to obtain an exponential lower bound for randomized algorithms as well. Our EFX lower bounds hold even for two players with identical submodular valuations.

Next, we showed that for n players with general but identical valuations, a modification of the leximin solution is guaranteed to be EFX. We showed how this result can be adapted into a cut-and-choose protocol for finding an EFX allocation between two players with general and possibly distinct valuations.

We also considered satisfying EFX and Pareto optimality together. We showed that if players are allowed to have zero value for a good being added to their bundle, it is impossible to guarantee EFX and Pareto optimality simultaneously. However, if we assume that a player's value for her bundle is strictly increased by adding any good (even just by some tiny ϵ), the leximin solution is EFX and PO two settings: for n players with general but identical valuations, and for two players with possibly distinct additive valuations. We view the latter result as our result of most practical significance: assuming nonzero marginal utility, it provides stronger guarantees the currently deployed algorithm on Spliddit, even in simple examples. We also gave an algorithm for finding a $\frac{1}{2}$-EFX allocation for any number of players with subadditive valuations. Finally, we discussed a (relatively constrained) set of valuations for which an EFX allocation can be computed in polynomial time.

The ideal next step would be to consider EFX with distinct valuations and more than two players.This problem seems quite challenging, even for the special case of additive valuations. Indeed, Caragiannis et al. [40] were unable to settle the question of whether EFX allocations in that context always exist, "despite significant effort." After substantial follow-up work, it was recently shown that EFX allocations are guaranteed to exist for three players with (possibly distinct) additive valuations [45]. The case of four of more players remains open.

Another direction is to pursue stronger lower bounds for finding an EFX allocation. In particular, communication complexity allows players unlimited computation and queries, and only measures the number of bits transmitted. The cut-and-choose protocol from Section 5.4 constitutes a linear communication protocol for two players with general and possibly distinct valuations to compute an EFX allocation, so any communication complexity lower bound would need to consider more than two players. On the other hand, we know finding an EFX allocation to be hard in the query model even for two players, which suggests an interesting separation.

More generally, communication complexity is one example of a topic that has been studied in algorithmic mechanism design and may be useful in the study of fair division. Another such topic is the hierarchy of complement-free valuations (additive, submodular, subadditive, etc.). Our work already implies separations between these valuation classes from a fair division perspective, and suggests that fair division with different classes of player valuations deserves further study.

Chapter 6

Communication complexity

This chapter continues our study of axiomatic objectives for indivisible private goods. As discussed, envy-freeness allocations may not exist when goods are indivisible. One approach to this difficulty is to consider relaxed versions of envy-freeness: this is what we did in Chapter 5. In this chapter, we take a sort of converse approach: our goal is to efficiently determine whether or not an envy-free allocation exists, from the perspective of communication complexity. We also consider proportionality, another common fairness axiom, and approximations of both proprtionality and envy-freeness. We also study how the complexity varies across valuation classes.

We show that for more than two players (and any combination of other parameters), determining whether a fair allocation exists requires exponential communication (in the number of goods). For two players, tractability depends heavily on the specific combination of parameters, and most of the chapter is focused on the two player setting. Taken together, our results completely resolve whether the communication complexity of computing a fair allocation (or determining that none exist) is polynomial or exponential, for every combination of fairness notion, valuation class, and number of players, for both deterministic and randomized protocols.

6.1 Introduction

An allocationis *envy-free* (EF) if each player's value for her own bundle is at least as much as her value for any other player's bundle. An allocation is *proportional* (Prop) if each player's value for her bundle is at least $1/n$ of her value for the entire set of items, where n is the number of players. We study the problem of finding an envy-free (or proportional) allocation, or showing that none exists.

We also consider approximate versions of these properties: for $c \in [0, 1]$, an allocation is c-EF if each player's value for her own bundle is at least c times her value for any other player's bundle, and an allocation is c-Prop if each player's value for her bundle is at least c/n of her value for the entire set of items. Thus 1-EF and 1-Prop are standard envy-freeness and proportionality, respectively. The same counterexample of two players and a single item shows that these approximate properties also cannot be guaranteed for any $c > 0$.[1]

From a computational complexity viewpoint, this problem is hard even when player valuations are *additive*, meaning that a player's value for a set of items is the sum of her values for the individual items. For two players with identical additive valuations, determining whether a 1-EF or 1-Prop exists is NP-hard, via a simple reduction from the partition problem [21].

It is arguably even more natural to study fair division from a communication complexity perspective, where there is no centralized authority and each player initially knows only her own preferences.

When players have *combinatorial valuations*, their values for a bundle cannot just be decomposed into their values for the individual items.[2] In particular, for m items, a combinatorial valuation may contain 2^m different values. The primary question is to determine whether players need to exchange an exponential amount of information to compute a fair allocation, or whether the problem can be solved using only polynomial communication. This question has not been studied previously, despite the rich literature on communication complexity in combinatorial auctions (e.g. [65, 132, 134]).

Our work can also be thought of as formally studying the difficulty of eliciting different classes of valuations from a fair division standpoint. Additive valuations are typically used in practice (for example on the non-profit website Spliddit [98]) because each player need only report one value for each item to specify the entire valuation. Richer combinatorial valuations allow for more expressiveness, but may be more difficult to elicit. Our work formally studies the tradeoffs between these factors.

6.1.1 Our results

We study the following question: "Given n players and m items, a fairness property $P \in \{\text{EF, Prop}\}$, and a constant $c \in [0, 1]$, how much communication is required to either find a c-P allocation, or show that none exists?"[3] In other words, the problem is to determine whether a c-P allocation exists, and if so, return one. We are primarily interested in whether this can be done with communication polynomial in m. The answer to this question will depend on n, P, and c. We also consider when player valuations are restricted to be submodular or subadditive, as well as deterministic vs. randomized protocols.

All in all, we give a full characterization of the communication complexity for every combination of the following five parameters:

1. Number of players n
2. Valuation class: submodular, subadditive, or general
3. Each $P \in \{\text{EF, Prop}\}$
4. Every constant $c \in [0, 1]$

[1] We generally assume that $c > 0$, since every allocation is both 0-EF and 0-Prop.

[2] An increasing amount of research in fair division considers such combinatorial valuations (e.g., [11, 92]).

[3] We only consider a single c-P property at a time: we do not consider satisfying envy-freeness and proportionality simultaneously. For subadditive valuations, c-EF implies c-Prop, but c-EF and c-Prop are incomparable for general valuations.

5. Deterministic or randomized communication complexity

The importance of the two-player setting

One of our results (Section 6.7) shows that there is no hope for a polynomial communication protocol for more than two players: exponential communication is required for every $n > 2$, for either $P \in \{\text{EF}, \text{Prop}\}$, for any $c > 0$, even for submodular valuations, and even for randomized protocols. The (very important) two-player case is surprisingly rich, however, with multiple phenomena occurring across different valuation classes and constants c. The results for two players in the deterministic setting are summarized in Table 6.1. It is also surprising that there is such a chasm between the two- and three-player cases; for example, there is no analogous chasm for maximizing the social welfare in combinatorial auctions.

Furthermore, in contrast to combinatorial auctions, the two-player setting is fundamental in fair division. Indeed, one of the first known mentions of fair division is in the Bible, when Abraham and Lot use the cut-and-choose method to divide a piece of land. In modern day, one of the primary applications of fair division for indivisible items is divorce settlements, which is fundamentally a two-player setting. Fair Outcomes Inc.[4], a commercial fair division website, only allows for two players. Other applications of fair division, such as dividing an inheritance and international border disputes, are also often two player settings. Unless otherwise mentioned, we assume that $n = 2$ throughout the chapter.

Submodular valuations

We first consider submodular valuations in the deterministic setting (for $n = 2$). We show that full proportionality (1-Prop) requires only polynomial communication (Theorem 6.3.1), whereas full envy-freeness requires exponential communication (Theorem 6.6.1), exhibiting an interesting difference between the two properties.

The hardness result for 1-EF leaves open the intriguing possibility of a polynomial-communication approximation scheme (PAS):[5] for any fixed $c < 1$, is communication cost polynomial in m sufficient? As one of our main results, we prove that this is indeed the case, and we prove it using a reduction to a type of graph we call the "minimal bundle graph" (Theorem 6.4.1). This is our most technically involved argument.

The communication cost of this protocol exponential in $\frac{1}{1-c}$, and so this PAS is not a *fully polynomial-communication approximation scheme* (FPAS), which would require polynomial dependence on $\frac{1}{1-c}$. Our lower bound for 1-EF (Theorem 6.6.1) rules out an FPAS, so our results are still tight.

[4] http://fairoutcomes.com

[5] This is the same idea as a polynomial-time approximation scheme (PTAS), but here we are interested in communication, not time.

Subadditive valuations

The story is different for subadditive valuations, which are treated in Section 6.8. We show that only polynomial communication is required for c-EF when $c \leq 1/2$ (Theorem 6.8.3) and for c-Prop when $c \leq 2/3$ (Theorem 6.8.4). Interestingly, the constants $1/2$ and $2/3$ turn out to be tight: we show that exponential communication is required for c-EF for every constant $c > 1/2$ (Theorem 6.8.5) and for c-Prop for every constant $c > 2/3$ (Theorem 6.8.6). This establishes another interesting difference between the two fairness notions.

General valuations

The story is again different for general valuations, which we consider in Section 6.9. In the deterministic setting, c-EF and c-Prop each require exponential communication for every $c > 0$ (Theorems 6.9.2 and 6.9.1). This resolves the deterministic setting.

It is interesting that hardness (for two players and deterministic protocols) turns out to be monotonic with respect to c, i.e., increasing c cannot make the problem easier. This makes sense intuitively, but we do not have a simple proof of this.

Randomized communication complexity

The c-Prop lower bound for general valuations also holds in the randomized setting for any $c > 0$. However, c-EF admits an efficient randomized protocol for any $c \leq 1$ and general (and hence also subadditive and submodular) valuations. This randomized protocol is based on a reduction to the Equality problem (testing whether two bit strings are identical), which is known to have an efficient randomized protocol. Our randomized protocol for c-EF also carries over to c-Prop for any $c \leq 1$ in the special case of subadditive (and hence also submodular) valuations. This resolves the randomized setting.

One may wonder why we care about deterministic protocols, when randomized protocols do so well. Aside from the technical goal of handling every combination of parameters, in fair division settings with considerable value (e.g., inheritance, divorce settlements), players may be wary of allowing randomization.

Lastly, we briefly consider the maximin share property in Section 6.10, and prove exponential lower bounds in that setting as well.

6.1.2 Ideas behind our protocols

Since the problem is always hard when $n > 2$, all of our upper bounds are in the two-player setting. All of our positive results require the following condition: for any partition of the items into two bundles A_1 and A_2, each player must be happy with at least one of A_1 and A_2. This is always true for envy-freeness: a player is always happy with whichever of A_1 and A_2 she has maximum value for (she could be happy with both bundles if they have equal value to her). This is not satisfied for proportionality in general, for example if a player has value zero for each of A_1 and A_2, but positive value for $A_1 \cup A_2$. However, it is satisfied for subadditive valuations.

	c-EF (deterministic)		c-Prop (deterministic)	
	easy when	hard when	easy when	hard when
general valuations	never	$c > 0$ (Thm. 6.9.2)	never	$c > 0$ (Thm. 6.9.1)
subadditive valuations	$c \leq 1/2$ (Thm. 6.8.3)	$c > 1/2$ (Thm. 6.8.5)	$c \leq 2/3$ (Thm. 6.8.4)	$c > 2/3$ (Thm. 6.8.6)
submodular valuations	$c < 1$ (Thm. 6.4.1)	$c = 1$ (Thm. 6.6.1)	$c \leq 1$ (Thm. 6.3.1)	never

Table 6.1: A summary of our results for the two-player deterministic setting. For both c-EF and c-Prop, we characterize exactly when the problem is easy (i.e., can be solved with communication polynomial in the number of items) and hard (i.e., requires exponential communication). We note that the protocol for Theorem 6.4.1 has communication cost exponential in $\frac{1}{1-c}$, and the corresponding lower bound (Theorem 6.6.1) rules out a protocol with communication cost polynomial in $\frac{1}{1-c}$. See Section 6.1.1 for additional discussion.

All of our deterministic protocols have the same first step: if there is any allocation where player 1 would be happy to receive either bundle, she specifies that allocation to player 2, and player 2 selects her preferred bundle. Player 2 is guaranteed to be happy with at least one of the bundles by the above condition, and player 1 is happy with either bundle in this allocation, so she is happy as well.

The key to the analysis is what happens when there is no allocation such that player 1 is happy with either bundle. It will turn out that the absence of such an allocation implies certain structure in the valuations. The exact structure, and the way the structure is exploited, depends on the setting (valuation class, property P, and constant c).

For example, consider the case of subadditive valuations and $\frac{1}{2}$-EF. We show that if there is no allocation where player 1 is happy with either bundle, then there must exist a single item that player 1 values more than the rest of the items combined. Then player 1 can simply specify that item to player 2. If player 2 is happy with the rest of the items, we have found a satisfactory allocation. Otherwise, there is no satisfactory allocation, since player 1 and player 2 both care about that particular item more than the rest of the items combined.

Furthermore, this protocol gives an additional guarantee. If a c-P allocation is not returned, the protocol will return the fairest allocation possible, i.e., a c'-P allocation where no allocation is c''-P for any $c'' > c'$. For brevity, we will use c^* to refer to the maximum c' such that a c'-P allocation exists.[6] If player 2 determines that a c-P allocation does not exist, then there is a single item g that both players care about more than all of the other items together. One player will have to not receive item g, and the protocol gives g to the player who will be most unhappy otherwise. This yields a c^*-P allocation. In fact, all of our deterministic protocols give this guarantee, although slightly more work is required to achieve it in other settings.

[6] It is possible that $c^* = 0$ (for example, in the case of two players and one item), but our protocol at least certifies that this is the best possible.

Minimal bundles

The reasoning described above is actually a special case of analyzing what we call *minimal bundles*. We say that a bundle is minimal for some player if that player is happy with the bundle, but is not happy with any strict subset of that bundle.[7] The minimal bundles represent the most a player is willing to compromise. If a player does not receive one of her minimal bundles (or a superset thereof), she cannot be happy, by definition. On the other hand, if a player receives one of her minimal bundles (or a superset thereof), she is guaranteed to be happy.[8] Thus it is both necessary and sufficient for each player to receive one of her minimal bundles (or a superset thereof). By this reasoning, it is sufficient for player 1 to specify all of her minimal bundles to player 2: player 2 can then determine if there is an allocation which satisfies her (player 2), while still giving player 1 one of player 1's minimal bundles.

The general Minimal Bundle Protocol (Protocol 8) is as follows. If there is an allocation where player 1 is happy with either bundle, she specifies that allocation to player 2, and we are done. Otherwise, player 1 specifies all of her minimal bundles to player 2, who searches for a satisfactory allocation. If player 2 fails to find one, she declares that no satisfactory allocation exists. There is a final step that is used to guarantee that a c^*-P allocation is returned if no c-P allocation is found; this will be described later.

The key is proving that the number of minimal bundles is polynomial in m, and this analysis varies based on the context. For example, for subadditive valuations and $\frac{1}{2}$-EF, we discussed above how if there is no allocation where player 1 is happy with either bundle, there must be a single item g that she values more than all of the other items together. This means that player 1 has a single minimal bundle: $\{g\}$.

We also use the protocol to give a PAS for EF in the submodular setting: we show that for every fixed $c < 1$, the number of minimal bundles is at most $2(m+1)^{\frac{8}{1-c}}$, and thus the protocol uses polynomial communication for any fixed c. The analysis for this case is technically involved and involves constructing what we call the "minimal bundle graph" for player 1's valuation. The vertices in this graph are the minimal bundles, and two vertices share an edge if the corresponding bundles overlap by exactly one item (it will be impossible for two minimal bundles to overlap by more than one item). For some of these edges, moving the overlapping item between bundles will cause a large change in value: these special edges will play an important role. We will show that the only way to have a large number of minimal bundles is for there to be a large number of these special edges, but submodularity will imply an upper bound on how many special edges can be incident on a single vertex, and hence an upper bound on the total number of special edges.

The Minimal Bundle Protocol is correct for any valuation class, property P, or constant c. However, in some contexts, the number of minimal bundles may be exponential. Our lower bound constructions all involve valuations with an exponential number of minimal bundles.

[7] A similar notion of "minimal bundles" features prominently in [26].

[8] We assume monotonicity: adding items to a bundle cannot decrease its value.

6.1.3 Related work

In the previous chapter, we discussed general related work in the area of fair division (Section 5.1.2), so in this section, we focus specifically on the relationship with communication complexity.

Communication complexity was first studied by [177]. The paper most relevant to our work is [134], which shows that maximizing social welfare requires exponential communication, even for two players with submodular valuations. Furthermore, they show that for general valuations, any constant factor approximation of the social welfare requires exponential communication to compute. Although they do not mention envy-freeness, proportionality, or fair division, some of their arguments can be adapted to prove exponential lower bounds for some (but not all) of the cases that we study.

A recent and complementary line of work is presented in [29]. They study the communication complexity of fair division with *divisible* goods (also known as "cake cutting"), where each resource can be divided into arbitrarily small pieces. Their paper complements this chapter with no overlap. Together, our works give a comprehensive picture of the communication complexity of fair division in both the indivisible and divisible models.

The organization of the rest of the chapter is as follows. Section 6.2 formally presents the model. Section 6.3 presents our 1-Prop protocol for submodular valuations. In Section 6.4, we discuss the PAS for 1-EF for submodular valuations. Section 6.5 describes our general lower bound approach, and proves a lemma that we will use to prove lower bounds later on in a standardized way. Section 6.6 uses that lemma to prove hardness for 1-EF for submodular valuations, which shows that the PAS from Section 6.4 is optimal. Section 6.7 shows that the problem is always hard for more than two players, even for submodular valuations and even in the randomized setting. The rest of the chapter is focused on resolving the two player case. Section 6.8 presents the upper and lower bounds for subadditive valuations. Section 6.9 considers general valuations, and also handles the randomized two player setting. Table 6.1 will be complete after this section. Finally, we consider the maximin share property (to be defined later) in Section 6.10.

6.2 Model

We assume the same resource allocation model as in Chapter 5: each agent i has a valuation v_i, which may be submodular, subadditive, or unrestricted, and our goal is determine an allocation $A = (A_1, \ldots, A_n)$. In this chapter, our fairness axioms of interest are approximate proportionality and envy-freeness:

Definition 6.2.1. *An allocation $A = (A_1, \ldots, A_n)$ is c-EF for some $c \in [0, 1]$ if for all $i, j \in N$,*

$$v_i(A_i) \geq c \cdot v_i(A_j)$$

Definition 6.2.2. *An allocation $A = (A_1, \ldots, A_n)$ is c-Prop for some $c \in [0, 1]$ if for all $i \in N$,*

$$v_i(A_i) \geq c \cdot \frac{v_i(M)}{n}$$

Thus 1-EF is standard envy-freeness, and 1-Prop is standard proportionality.

We will say that a player is (c, P)-*happy* with an allocation A if property c-P is satisfied from her viewpoint. Specifically, when $P = \text{EF}$, we will say that player i is (c, P)-happy with allocation A if $v_i(A_i) \geq c \cdot v_i(A_j)$ for all j. For $P = \text{Prop}$, we will say a player i is (c, P)-happy if $v_i(A_i) \geq \frac{c}{n} v_i(M)$. We will typically leave P implicit, and just say that player i is c-happy. We sometimes also leave c implicit, and just say that player i is happy.

An instance of Fair Division consists of a set of players N, a set of items M, player valuations $(v_1 ... v_n)$, a fairness property $P \in \{\text{EF}, \text{Prop}\}$, and a constant $c \in [0, 1]$. The goal is to find an allocation satisfying c-P, or show that none exists.

Two players

We use the following additional terminology when $n = 2$. For a player i, we will use $\bar{i}$ to denote the other player. For an allocation $A = (A_1, A_2)$, let $\overline{A}$ be the allocation (A_2, A_1). Also, when $n = 2$, knowing player i's bundle uniquely determines the overall allocation, since player $\bar{i}$ simply has every item not in player i's bundle. Therefore, with slight abuse of notation, we say that player i is c-happy with bundle S if player i is c-happy with the allocation A where $A_i = S$ and $A_{\bar{i}} = M \backslash S$.

6.2.1 Communication complexity

We assume that each player knows only her own valuation v_i, and does not know anything about other players' valuations. In order to solve an instance of Fair Division, players will need to exchange information about their valuations. We assume that all players know N, M, P, and c. Since there 2^m subsets of M, specifying a bundle requires m bits. We will use v^{size} to refer to the number of bits required to represent a value $v_i(S)$. We assume that v^{size} is polynomial in m, otherwise sending even a single value would rule out a polynomial communication protocol.

A (deterministic) protocol Γ specifies which player should speak (and what she should say) as a function of the messages sent so far, and terminates when a player declares that an allocation A satisfies c-P, or when a player declares that no c-P allocation exists. For fixed N, M, P, and c, we define the communication cost of a protocol Γ to be the maximum number of bits Γ sends across all player valuations $v_1 ... v_n$. Formally, let $C_\Gamma(N, M, (v_1 ... v_n), P, c)$ be the number of bits that Γ communicates when run on the fair division instance $(N, M, (v_1 ... v_n), P, c)$. Then the communication cost of Γ is $\max_{(v_1 ... v_n)} C_\Gamma(N, M, (v_1 ... v_n), P, c)$.

We define the deterministic communication complexity $D(n, m, P, c)$ as the minimum communication cost of any protocol Γ which correctly solves Fair Division for n players, m items, property P and constant c. Formally,

$$D(n, m, P, c) = \min_{\Gamma} \max_{(v_1 ... v_n)} C_\Gamma([n], [m], (v_1 ... v_n), P, c)$$

where Γ ranges over all correct deterministic protocols.

In a randomized protocol Γ_R, each player also has access to an infinite stream of random bits. The protocol should correctly solve Fair Division with probability $2/3$ (say) over these random bits.

Like the deterministic setting, the communication cost of Γ_R is the number of bits Γ_R communicates for a worst-case choice of $v_1 ... v_n$. We can similarly define the randomized communication complexity $R(n, m, P, c)$ as the minimum communication cost of any randomized protocol Γ_R which correctly solves Fair Division with probability at least 2/3. Formally,

$$R(n, m, P, c) = \min_{\Gamma_R} \max_{(v_1 ... v_n)} C_{\Gamma_R}([n], [m], (v_1 ... v_n), P, c)$$

where Γ_R ranges over all correct randomized protocols. If valuations are restricted to be subadditive or submodular, the problem may become easier, so $D(n, m, P, c)$ and $R(n, m, P, c)$ may be affected. We use $D_{subadd}(n, m, P, c)$ and $D_{submod}(n, m, P, c)$ to denote the deterministic communication complexity when valuations are restricted to be subadditive and submodular, respectively (and similarly for $R_{subadd}(n, m, P, c)$ and $R_{submod}(n, m, P, c)$). The following relationships are immediate, for all n, m, P, and c:

$$R(n, m, P, c) \leq D(n, m, P, c)$$
$$D_{submod}(n, m, P, c) \leq D_{subadd}(n, m, P, c) \leq D(n, m, P, c)$$
$$R_{submod}(n, m, P, c) \leq R_{subadd}(n, m, P, c) \leq R(n, m, P, c)$$

Another factor that may affect the communication complexity is how the players gain access to random bits. In the *public-coin* model, the players can also see other players' streams of random bits; in the *private-coin* model, each player sees only her own stream. This distinction is not significant in our setting, however, due to the following theorem from [131].

Theorem 6.2.1 ([131])**.** *Suppose there exists a public-coin randomized protocol with communication cost C on ℓ bits of input. Then there exists a private-coin randomized protocol with communication cost $O(C + \log \ell)$.*

Thus we will assume all randomized protocols to be public-coin for the rest of the chapter.

Finally, we mention the multiparty (i.e., $n > 2$) communication complexity model. There is more than one such model: for example, do players communicate in a peer-to-peer fashion, or is each message broadcast for all of the players to see? We discuss in Section 6.7 how this turns out not to matter in our setting.

6.3 An upper bound for 1-Prop with submodular valuations

This section presents our first result: a deterministic protocol for 1-Prop, when there are two players, and when valuations are submodular. The protocol will communicate just $m + 1$ values and a single bundle. Our protocol either finds a 1-Prop allocation or a c^*-Prop allocation. Recall that c^* is the maximum c such that a c-P allocation exists.[9] We prove the following theorem:

[9] Technically, our protocol always returns a c^*-Prop allocation, since we only consider $c \in [0, 1]$. We state the 1-Prop case separately in the theorem because it is handled separately in the protocol.

k	1	2	3
$v_1(g_1, \ldots, g_k)$	2	2	3
δ_k^M	2	0	1

Figure 6.1: An example of a possible valuation v_1 over three goods, and the corresponding values for δ_k^S.

Theorem 6.3.1. *For two players with submodular valuations, Protocol 7 has communication cost at most $(m+1)v^{size} + m$, and either returns a 1-Prop allocation or a c^*-Prop allocation. This also implies that for any $c \in [0, 1]$,*

$$D_{submod}(2, m, Prop, c) \leq (m+1)v^{size} + m$$

To see that the theorem also implies $D_{submod}(2, m, \text{Prop}, c) \leq (m+1)v^{size} + m$ for any c, suppose that the protocol returns a c^*-Prop allocation where $c^* < 1$: then we know that no allocation is c'-Prop for any $c' > c^*$, so a c-Prop allocation exists if and only if $c^* \geq c$. Thus Protocol 7 either finds a c-Prop allocation or shows that none exists, for any $c \in [0, 1]$.

It will be important that the following condition is satisfied in this setting:

Condition 6.3.1. *For every allocation A, each player is happy with at least one of A and $\overline{A}$.*

Recall that for an allocation $A = (A_1, A_2)$, $\overline{A} = (A_2, A_1)$. This condition is satisfied for proportionality with subadditive valuations (and hence also satisfied for submodular valuations):

$$\max\big(v_i(A_1), v_i(A_2)\big) \geq \frac{1}{2}\big(v_i(A_1) + v_i(A_2)\big) \geq \frac{1}{2}v_i(A_1 \cup A_2) = \frac{1}{2}v_i(M) \geq \frac{c}{2}v_i(M)$$

for all $c \in [0, 1]$. Thus player i is always happy if she receives the bundle $\arg\max\big(v_i(A_1), v_i(A_2)\big)$. Also, we assume in this section that $v_1(M) = 1$, without loss of generality: were this not the case, we could simply rescale v_1 as needed.

Let $M = (g_1, g_2 ... g_m)$ be an arbitrary ordering of the items. We assume that this ordering is publicly known. Consider starting from the empty set and adding the items in M one at a time in this order. We define δ_k^M as player 1's marginal value of adding g_k in this process: $\delta_k^M = v_1(g_1, g_2 \ldots g_{k-1}, g_k) - v_1(g_1, g_2 \ldots g_{k-1})$. Note that δ_k^M is not equal to $v_i(\{g_k\})$ in general, because of submodularity.

The protocol is as follows. The first step is common to all of our deterministic protocols: player 1 checks if there is an allocation A where she is happy with both A and $\overline{A}$. If so, player 2 can choose whichever she prefers, and we are done by Condition 6.3.1. If this fails, the following condition is satisfied:

Condition 6.3.2. *There is no allocation A for which player 1 is happy with both A and $\overline{A}$.*

In this case, player 1 sends the values $(\delta_1^M, \delta_2^M ... \delta_m^M)$ to player 2. For every bundle S, player 2 needs to be able to figure out whether player 1 likes S or $M \backslash S$. To do this, player 2 simply pretends that player 1's valuation is additive where δ_k^M is the value of item g_k. Formally, let $\chi(S) = \sum_{g_k \in S} \delta_k^M$: player 2 pretends that $v_1(S) = \chi(S)$. This will not be a perfect estimate of v_1, of course, but player

2 does not need to know the exact value of $v_1(S)$: she only needs to know whether player 1 is happy with S.

Lemma 6.3.1 shows that this actually works: assuming Condition 6.3.2, $v_1(S) \geq 1/2$ if and only if $\chi(S) \geq 1/2$. We informally argue why this is case. Crucially, submodularity implies that $\chi(S)$ will be an underestimate of $v_1(S)$: $v_1(S) \geq \chi(S)$ for all S. Since $\chi(S) + \chi(M \backslash S) = \sum_{k=1}^m \delta_k^M = v_1(M)$, either $\chi(S) \geq 1/2$ or $\chi(M\backslash S) \geq 1/2$. Say $\chi(S) \geq 1/2$: then $v_1(S) \geq \chi(S) \geq 1/2$, so player 1 is happy with S. Then by Condition 6.3.2, we know that player 1 is not happy with $M\backslash S$. Therefore, for any bundle S, player 2 can correctly use χ as a proxy for v_1 to determine which of S and $M\backslash S$ player 1 is happy with. Thus χ is sufficient for her to determine whether or not a 1-Prop allocation exists, and if so, find one. This lemma is the heart of Protocol 7.

Step 4, $S^*(v_i)$, and $\mathbf{c}_i(S^*(v_i))$ are necessary only for finding a c^*-Prop allocation if no 1-Prop allocation is found. For a bundle S and property P, let $\mathbf{c}_i^P(S)$ be the maximum $c' \leq 1$ such that player i is c'-happy with S. For example, $\mathbf{c}_i^{Prop}(S) = \min\left(1, \frac{2v_i(S)}{v_i(M)}\right) = \min(1, 2v_i(S))$, since we assumed $v_i(M) = 1$. Although this section considers only proportionality, we allow for either $P \in \{\text{EF, Prop}\}$ in our definitions, since we will use this terminology again in later sections. We will typically leave P implicit and write $\mathbf{c}_i(S)$.

For each player i, we define a special bundle

$$S^*(v_i) = \operatorname*{arg\,max}_{S \subseteq M:\ \mathbf{c}_i(S) < c} \mathbf{c}_i(S)$$

In words, $S^*(v_i)$ is the bundle that player i is the most happy with, out of all of the bundles she is not fully happy (i.e., c-happy) with.[10]

Protocol 7 Protocol for two players with submodular valuations to either find a 1-Prop allocation or a c^*-Prop allocation.

Private inputs: v_1, v_2
Public inputs: the ordering of $M = \{g_1, g_2 ... g_m\}$

1. If there exists an allocation A where player 1 is happy with both A and $\overline{A}$, player 1 sends that allocation to player 2. If player 2 is happy with A, she declares that A is 1-Prop, otherwise she declares that $\overline{A}$ is 1-Prop.

2. If there is no such allocation A, player 1 sends the values $(\delta_1^M, \delta_2^M ... \delta_m^M)$ to player 2, along with $S^*(v_1)$ and the value $\mathbf{c}_1(S^*(v_1))$.

3. Player 2 first checks if there exists any bundle S where $\chi(S) \geq 1/2$ and $v_2(M\backslash S) \geq 1/2$. If so, she declares that the allocation $(S, M\backslash S)$ is 1-Prop.

4. If not, player 2 computes $S^*(v_2)$, $\mathbf{c}_2(S^*(v_2))$, and $i = \arg\max_{i' \in \{1,2\}} \mathbf{c}_{i'}(S^*(v_{i'}))$. Let A be the allocation where $A_i = S^*(v_i)$ and $A_{\bar{i}} = M\backslash S^*(v_i)$. Player 2 then declares that A is $\mathbf{c}_i(S^*(v_i))$-Prop, and that $c^* = \mathbf{c}_i(S^*(v_i))$.

[10] Although incentives are not the focus of this chapter, we mention that Step 4 makes Protocol 7 easily manipulable. Specifically, it is a dominant strategy for player 1 report $\mathbf{c}_1(S^*(v_1)) = 0$ (i.e., if I am not fully happy, I am not happy at all, so you should make me fully happy). The same reasoning applies to Protocol 8, which has the same Step 4.

It will be useful for the analysis to define δ_i^S for an arbitrary bundle S. First, let

$$S_{\leq k} = \{g_j \in S \mid j \leq k\}$$

For example, $S_{\leq 0} = \emptyset$ and $S_{\leq m} = S$ for all S. Also, whenever $g_k \in S$, we have $S_{\leq k} = S_{\leq k-1} \cup \{g_k\}$. Let $\delta_k^S = v_1(S_{\leq k}) - v_1(S_{\leq k-1})$. Note that for all S, $v_1(S) = \sum_{k=1}^m \delta_k^S$.

Lemma 6.3.1. *Assuming Condition 6.3.2, for any bundle S, $v_1(S) \geq 1/2$ if and only if $\chi(S) \geq 1/2$.*

Proof. We first claim that for any bundle S and any item $g_k \in S$, $\delta_k^S \geq \delta_k^M$. We have

$$\delta_k^S = v_1(S_{\leq k}) - v_1(S_{\leq k-1}) = v_1(S_{\leq k-1} \cup \{g_k\}) - v_1(S_{\leq k-1})$$

and

$$\delta_k^M = v_1(M_{\leq k}) - v_1(M_{\leq k-1}) = v_1(M_{\leq k-1} \cup \{g_k\}) - v_1(M_{\leq k-1})$$

Since $S_{\leq k-1} \subseteq M_{\leq k-1}$, we have $v_1(S_{\leq k-1} \cup \{g_k\}) - v_1(S_{\leq k-1}) \geq v_1(M_{\leq k-1} \cup \{g_k\}) - v_1(M_{\leq k-1})$ by submodularity. Thus $\delta_k^S \geq \delta_k^M$ for all k and S. Therefore for any bundle S,

$$v_1(S) = \sum_{g_k \in S} \delta_k^S \geq \sum_{g_k \in S} \delta_k^M = \chi(S)$$

so $v_1(S) \geq \chi(S)$ for all $S \subseteq M$.

Suppose $\chi(S) \geq 1/2$: then we immediately have $v_1(S) \geq 1/2$ by the above argument. Suppose $v_1(S) \geq 1/2$. Then by Condition 6.3.2, $v_1(M \backslash S) < 1/2$. Therefore $\chi(M \backslash S) < 1/2$. Next, we have

$$\chi(S) + \chi(M \backslash S) = \sum_{g_k \in S} \delta_k^M + \sum_{g_k \in M \backslash S} \delta_k^M = \sum_{k=1}^m \delta_k^M = v_1(M) = 1$$

Since $\chi(M \backslash S) < 1/2$, we have $\chi(S) \geq 1/2$. ☐

Theorem 6.3.1. *For two players with submodular valuations, Protocol 7 has communication cost at most $(m+1)v^{size} + m$, and either returns a 1-Prop allocation or a c^*-Prop allocation. This also implies that for any $c \in [0, 1]$,*

$$D_{submod}(2, m, Prop, c) \leq (m+1)v^{size} + m$$

Proof. If the protocol terminates in step 1, just one bundle is communicated (and zero values), which requires m bits. Thus in this case, the communication cost is $m \leq (m+1)v^{size} + m$. If the protocol does not terminate in step 1, then the m values $(\delta_1^M ... \delta_m^M)$ are sent, plus the bundle $S^*(v_1)$, plus the value $\mathbf{c}_1(S^*(v_1))$. By definition of $\mathbf{c}_1^{Prop}$, $\mathbf{c}_1(S^*(v_1))$ requires a single value to communicate.

Thus in this case, $m+1$ values and one bundle are communicated, so the communication cost is $(m+1)v^{size} + m$. Therefore the communication cost bound is satisfied.

It remains to prove correctness. Suppose the protocol terminates in step 1. By Condition 6.3.1, player 2 is happy with at least one of A and $\overline{A}$. Therefore player 2 is happy with whichever of A and

$\overline{A}$ she declares to be 1-Prop. Player 1 is happy with both A and $\overline{A}$, so she is also happy. Therefore if the protocol terminates in step 1, the declared allocation is in fact 1-Prop.

Suppose the protocol does not terminate in step 1. We assume Condition 6.3.2 for the remainder of the proof. Suppose player 2 declares that $(S, M \backslash S)$ is 1-Prop in step 3: then

$$\chi(S) \geq 1/2 \quad \text{and} \quad v_2(M \backslash S) \geq 1/2$$

Thus by Lemma 6.3.1, $v_1(S) \geq 1/2$, so $(S, M \backslash S)$ is indeed a 1-Prop allocation.

So suppose the protocol does not terminate until step 4. We first claim that no 1-Prop allocation exists. Suppose that a 1-Prop allocation A does exist: then $v_i(A_i) \geq 1/2$ for both i. Since the protocol did not terminate in step 1, we have Condition 6.3.2. Thus by Lemma 6.3.1, $\chi(A_1) \geq 1/2$. Let $S = A_1$: then

$$\chi(S) \geq 1/2 \quad \text{and} \quad v_2(M \backslash S) = v_2(A_2) \geq 1/2$$

so the protocol should have terminated in step 3, which is a contradiction.

Therefore no 1-Prop allocation exists. It remains to show that we return a c^*-Prop allocation in this case. Let $i = \arg\max_{i' \in \{1,2\}} \mathbf{c}_{i'}(S^*(v_{i'}))$ as computed by player 2 in step 4. Let A be the allocation returned by the protocol in this case: $A_i = S^*(v_i)$ and $A_{\bar{i}} = M \backslash S^*(v_i)$.

We first claim that A is $\mathbf{c}_i(S^*(v_i))$-Prop. Player i is $\mathbf{c}_i(S^*(v_i))$-happy with A by definition, and we claim that player $\bar{i}$ is 1-happy with A. If $\bar{i}$ were not 1-happy with A, then she must be 1-happy with $\overline{A}$ by Condition 6.3.1. Furthermore, since player i is not 1-happy with A, she must be 1-happy with $\overline{A}$ also by Condition 6.3.1. But then both players are 1-happy with $\overline{A}$, which is a contradiction.

Thus A is $\mathbf{c}_i(S^*(v_i))$-Prop. Suppose that $c^* \neq \mathbf{c}_i(S^*(v_i))$: then there exists an allocation A' where A' is c-Prop for some $c > \mathbf{c}_i(S^*(v_i))$. We know that player i cannot be happier than $\mathbf{c}_i(S^*(v_i))$-happy without being 1-happy, so player i must be 1-happy with A'. That implies that player 2 is not 1-happy with A', since no allocation makes both players 1-happy in this case. But then the happiest player $\bar{i}$ can be is $\mathbf{c}_{\bar{i}}(S^*(v_{\bar{i}}))$, and $\mathbf{c}_{\bar{i}}(S^*(v_{\bar{i}})) \leq \mathbf{c}_i(S^*(v_i))$ by assumption. Thus for any allocation, there is a player who is at most $\mathbf{c}_i(S^*(v_i))$-happy. Therefore no allocation is c-Prop for any $c > \mathbf{c}_i(S^*(v_i))$. □

6.4 PAS for EF with submodular valuations

In this section, we prove our other positive result for specifically submodular valuations: a deterministic protocol for c-EF when $c < 1$, and when there are two players. This is our most technically involved result. We prove the following theorem:

Theorem 6.4.1. *For two players with submodular valuations and any $c < 1$, Protocol 8 has communication cost at most $2m(m+1)^{\frac{8}{1-c}} + 2v^{size}$, and either returns a c-EF allocation or a c^*-EF allocation. This also implies that*

$$D_{submod}(2, m, EF, c) \leq 2m(m+1)^{\frac{8}{1-c}} + 2v^{size}$$

bundle S	$v_i(S)$	1-Prop?	minimal?
$\{g_1\}$	2	no	N/A
$\{g_2\}$	4	no	N/A
$\{g_3\}$	5	yes	yes
$\{g_1, g_2\}$	6	yes	yes
$\{g_1, g_3\}$	7	yes	no
$\{g_2, g_3\}$	9	yes	no
$\{g_1, g_2, g_3\}$	10	yes	no

Figure 6.2: An example demonstrating the minimal bundle property for $P = \text{Prop}$ and $c = 1$. This instance involves a valuation v_i over three goods. Since $v_i(M) = v_i(\{g_1, g_2, g_3\}) = 10$ in this case, player i is happy with S if and only if $v_i(S) \geq 10/n = 5$. For example, player i is happy with $\{g_1, g_3\}$, but that bundle is not minimal, since player i is also happy with $\{g_3\}$. In contrast, $\{g_1, g_2\}$ is minimal, since player i is happy with neither $\{g_1\}$ nor $\{g_2\}$.

for any $c < 1$.

This constitutes a polynomial-communication approximation scheme (PAS): the communication cost approaches infinity exponentially as c goes to 1, but for any fixed constant $c < 1$, it is polynomial in m.[11]

We use much of the same terminology from Section 6.3: in particular, $\mathbf{c}_i^P(A)$, $S^*(v_i)$, $S_{\leq k}$, and δ_k^S. Also, recall the following condition:

Condition 6.3.1. *For every allocation A, each player is happy with at least one of A and $\overline{A}$.*

This is satisfied for c-EF for any $c \in [0, 1]$, even for general valuations: if $v_i(A_i) \geq v_i(A_{\bar{i}})$, player i is happy with A. Otherwise, $v_i(A_{\bar{i}}) \geq v_i(A_i)$, so she is happy with $\overline{A}$.

Our PAS protocol will use the minimal bundle analysis discussed in Section 6.1.2. For a fixed constant c, we say that a bundle S is *minimal* for a particular player if that player is c-happy with S, but for all $g \in S$, she is not c-happy with $S \backslash \{g\}$. We use $\mathcal{S}$ to denote the set of player 1's minimal bundles: each $S \in \mathcal{S}$ is a minimal bundle for player 1. Also, in this section, we assume that $v_1(M) = 1$.

6.4.1 The protocol

We now describe Protocol 8, also known as the Minimal Bundle Protocol. Although we only consider envy-freeness in this section, we define Protocol 8 for either $P \in \{\text{EF, Prop}\}$. We will use this same protocol in Section 6.8.1 to prove upper bounds for both envy-freeness and proportionality in the subadditive case.

First, if there is an allocation A where player 1 is happy with both A and $\overline{A}$, we are done: player 2 chooses her favorite of A and $\overline{A}$, and she is guaranteed to be happy with at least one them by Condition 6.3.1. If there is no such allocation A, player 1 sends the set $\mathcal{S}$ of all of her minimal bundles to the other player. We will prove that in our setting, the number of minimal bundles sent in step 2 must be polynomial in m. Specifically, we will show that $|\mathcal{S}| < 2(m+1)^{\frac{8}{1-c}}$.

[11] Because the dependence on $\frac{1}{1-c}$ is exponential, this constitutes a PAS but not an FPAS. An FPAS is ruled out in Section 6.6.

Protocol 8 Protocol for two players to either find a c-P allocation or a c^*-P allocation.

Private inputs: v_1, v_2
Public inputs: P, c

1. If there exists an allocation A where player 1 is happy with both A and $\overline{A}$, player 1 sends that allocation to player 2. If player 2 is happy with A, she declares that A is c-P, otherwise she declares that $\overline{A}$ is c-P.
2. If there is no such allocation A, player 1 sends the set $\mathcal{S}$ of her minimal bundles to player 2. She also sends the bundle $S^*(v_1)$ and the value $\mathbf{c}_1(S^*(v_1))$.
3. Player 2 first checks if there exists a bundle $S \in \mathcal{S}$ where player 2 is happy with $M \backslash S$. If so, she declares that $(S, M \backslash S)$ is c-P.
4. If not, player 2 computes $S^*(v_2)$ and $i = \arg\max_{i' \in \{1,2\}} \mathbf{c}_{i'}(S^*(v_{i'}))$. Let A be the allocation where $A_i = S^*(v_i)$ and $A_{\bar{i}} = M \backslash S^*(v_i)$. Player 2 then declares that A is $\mathbf{c}_i(S^*(v_i))$-P, and that $c^* = \mathbf{c}_i(S^*(v_i))$.

The minimal bundles represent the most player 1 is willing to compromise while still being happy: she does not require anything more than a minimal bundle, but she is not happy with any strict subset of any of her minimal bundles. In this way, receiving a minimal bundle is both necessary and sufficient for player 1 to be happy. Using this reasoning, we will show that knowing $\mathcal{S}$ is sufficient for player 2 to find a c-P allocation or show that none exists. Finally, step 4 is identical to that of Protocol 7, and is used to find a c^*-P allocation when no c-P allocation exists.

6.4.2 Correctness

We now formally prove the correctness of Protocol 8. We will prove a few helpful lemmas before proving the main correctness lemma (Lemma 6.4.4).

Lemma 6.4.1. *If Protocol 8 declares an allocation to be c-P, the allocation is in fact c-P.*

Proof. The only two steps that can declare an allocation to be c-P are steps 1 and 3. Suppose the protocol declares an allocation to be c-P in step 1. Then by assumption, there exists an allocation A where player 1 is happy with both A and $\overline{A}$. If player 2 declares A to be c-P, then both players are happy with A, and the claim is satisfied. If player 2 declares $\overline{A}$ to be c-P, then she was not happy with A. By Condition 6.3.1, player 2 is happy with $\overline{A}$. Thus $\overline{A}$ is c-P in this case, so the lemma is satisfied if the protocol terminates in step 1.

Suppose the protocol declares an allocation to be c-P in step 3. Then the allocation declared can be written as $(S, M \backslash S)$ for some $S \in \mathcal{S}$. Since S is minimal, player 1 is happy with S by assumption, and player 2 only declares an allocation to be c-P in this step if she is happy with it. Thus the lemma is satisfied in this case as well. □

Lemma 6.4.2. *Player 1 is happy with a bundle S if and only if there exists a minimal bundle T where $T \subseteq S$.*

Proof. ($\Longrightarrow$) Suppose player 1 is happy with bundle S. If S is minimal, we are done, so assume S is not minimal. Then there exists $g \in S$ where player 1 is happy with $S \backslash \{g\}$. If $S \backslash \{g\}$ is not minimal, there again exists some $g' \in S \backslash \{g\}$ that we can remove, and this process can be repeated until we obtain some minimal subset of S.

($\Longleftarrow$) Suppose there exists a minimal bundle T where $T \subseteq S$. Then by monotonicity, $v_1(S) \geq v_1(T)$. Since T is minimal, player 1 is happy with T. If $P = \text{Prop}$, this is sufficient to show that player 1 is happy with S. If $P = \text{EF}$, it is also necessary to note that $v_1(M \backslash S) \leq v_1(M \backslash T)$, again by monotonicity. Thus the claim holds for both $P \in \{\text{EF}, \text{Prop}\}$. ☐

Lemma 6.4.3. *Protocol 8 declares an allocation to be c-P if and only if a c-P allocation exists.*

Proof. If no c-P allocation exists, the protocol does not declare any allocation to be c-P by Lemma 6.4.1. Thus assume a c-P allocation A exists. Then player 1 is happy with A_1, so by Lemma 6.4.2, there exists $S \in \mathcal{S}$ where $S \subseteq A_1$. Then $A_2 \subseteq M \backslash S$, so by monotonicity, player 2 is happy with $M \backslash S$. Thus if the protocol has not already terminated, player 2 will declare will declare $(S, M \backslash S)$ to be c-P. Then by Lemma 6.4.1, the declared allocation is in fact c-P, so the claim is satisfied in this case.

If the protocol terminated before player 2 considered S in step 3, the protocol declared some other allocation to be c-P, and the declared allocation is again c-P by Lemma 6.4.1 in this case. Thus the claim is satisfied in both cases. ☐

Finally, we show that the protocol correctly returns a c^*-P allocation if no c-P allocation exists. Recall the definitions of $S^*(v_i)$ and $\mathbf{c}(S)$: $\mathbf{c}_i(S)$ is the maximum $c' \leq 1$ where player i is c'-happy with S, and $S^*(v_i) = \arg\max_{S \subseteq M:\ \mathbf{c}_i(S) < c} \mathbf{c}_i(S)$. In words, $S^*(v_i)$ is the bundle that makes player i the most happy, out of all the bundles that do not make her c-happy. For $P = \text{EF}$, $\mathbf{c}_i(S) = \min\left(1, \frac{v_i(S)}{v_i(M \backslash S)}\right)$.

Lemma 6.4.4. *Protocol 8 either returns a c-P allocation or a c^*-P allocation.*

Proof. If a c-P allocation exists, Lemma 6.4.3 implies that the protocol correctly returns one, so the claim is satisfied in this case.

Suppose no c-P allocation exists: then the protocol does not declare an allocation to be c-P, again by Lemma 6.4.3. Thus the protocol does not terminate until step 4. Let $i = \arg\max_{i' \in \{1,2\}} \mathbf{c}_{i'}(S^*(v_{i'}))$ as computed by player 2 in step 4. Let A be the allocation returned by the protocol in this case: $A_i = S^*(v_i)$ and $A_{\bar{i}} = M \backslash S^*(v_i)$.

First observe that A is $\mathbf{c}_i(S^*(v_i))$-P: this is because player i is $\mathbf{c}_i(S^*(v_i))$-happy with A, and player $\bar{i}$ is c-happy with A. Suppose that A is not c^*-P: then there exists an allocation A' where A' is c''-P for some $c'' > \mathbf{c}_i(S^*(v_i))$. We know that player i cannot be happier than $\mathbf{c}_i(S^*(v_i))$-happy without being c-happy, so player i must be c-happy with A'. That implies that player 2 is not c-happy with A', since no allocation makes both players c-happy in this case. But then the happiest player $\bar{i}$ can be is $\mathbf{c}_{\bar{i}}(S^*(v_{\bar{i}}))$, and $\mathbf{c}_{\bar{i}}(S^*(v_{\bar{i}})) \leq \mathbf{c}_i(S^*(v_i))$ by assumption. Thus for any allocation, there is a player who is at most $\mathbf{c}_i(S^*(v_i))$-happy. Therefore $c^* = \mathbf{c}_i(S^*(v_i))$. ☐

6.4.3 Communication cost

It remains to bound the communication cost. This will primarily consist of proving an upper bound on the number of minimal bundles player 1 sends to player 2. We will go through a series of helpful lemmas before proving the final theorem.

The upper bound on the number of minimal bundles will depend on there being no allocation A for which player 1 is happy with both A and $\overline{A}$: recall that if there is such an allocation, then Protocol 8 terminates after step 1 and does not even send the set of minimal bundles $\mathcal{S}$. This condition was defined in Section 6.3.

Condition 6.3.2. *There is no allocation A for which player 1 is happy with both A and $\overline{A}$.*

Let $\Delta(S, g)$ be player 1's marginal value for adding item g to bundle S. Formally, $\Delta(S, g) = v_1(S \cup \{g\}) - v_1(S)$. Also, let $\alpha = \frac{1-c}{2}$.

The idea behind Lemma 6.4.5 is the following. Because of Condition 6.3.2, we have $c \cdot v_1(S) > v_1(M \backslash S)$ whenever player 1 is happy with S. If S is minimal, then moving any $g \in S$ to $M \backslash S$ will invert this inequality: $v_1(S \backslash \{g\}) < c \cdot v_1\big((M \backslash S) \cup \{g\}\big)$. Lemma 6.4.5 uses this to show that at least one of $\Delta(S \backslash \{g\}, g)$ and $\Delta(M \backslash S, g)$ has to be fairly large.

Lemma 6.4.5. *Assuming Condition 6.3.2, for every minimal bundle S and every good $g \in S$,*

$$\max\Big(\Delta(S \backslash \{g\}, g), \Delta(M \backslash S, g)\Big) \geq \alpha$$

Proof. Since S is minimal, for every good $g \in S$, we know that player 1 is not happy with $S \backslash \{g\}$. Specifically,

$$v_1(S \backslash \{g\}) < c \cdot v_1\big((M \backslash S) \cup \{g\}\big)$$

so by definition of Δ, we have

$$v_1(S) - \Delta(S \backslash \{g\}, g) < c \cdot \big(v_1(M \backslash S) + \Delta(M \backslash S, g)\big) = c \cdot v_1(M \backslash S) + c \cdot \Delta(M \backslash S, g)$$

We also know that player 1 is happy with S. Thus by Condition 6.3.2, player 1 is not happy with $M \backslash S$, so $v_1(M \backslash S) < c \cdot v_1(S)$. Adding this to the above equation yields

$$\begin{aligned} v_1(S) - \Delta(S \backslash \{g\}, g) + v_1(M \backslash S) &< c \cdot v_1(M \backslash S) + c \cdot \Delta(M \backslash S, g) + c \cdot v_1(S) \\ (1-c)v_1(S) + (1-c)v_1(M \backslash S) &< \Delta(S \backslash \{g\}, g) + c \cdot \Delta(M \backslash S, g) \\ (1-c)v_1(S) + (1-c)v_1(M \backslash S) &< \Delta(S \backslash \{g\}, g) + \Delta(M \backslash S, g) \\ (1-c)\big(v_1(S) + v_1(M \backslash S)\big) &< \Delta(S \backslash \{g\}, g) + \Delta(M \backslash S, g) \\ (1-c)v_1(M) &< \Delta(S \backslash \{g\}, g) + \Delta(M \backslash S, g) \end{aligned}$$

where the last step follows from submodularity (actually just subadditivity).

Since $v_1(M) = 1$ by assumption, we have

$$\Delta(S\backslash\{g\}, g) + \Delta(M\backslash S, g) \geq 1 - c$$
$$\max\Big(\Delta(S\backslash\{g\}), \Delta(M\backslash S, g)\Big) \geq \frac{1-c}{2} = \alpha$$

□

Next, we define a directed graph $G = (V, E)$ which we call the *minimal bundle graph*. The vertex set V is the set of minimal bundles. With slight abuse of notation, we will use S and T to refer both to minimal bundles and to the corresponding vertices in V. We define the edge set E by

$$E = \big\{(S, T) \mid \exists g \in S \text{ where } T \subseteq (M\backslash S) \cup \{g\}\big\}$$

The next three lemmas establish some useful properties of the minimal bundle graph.

Lemma 6.4.6. *Assuming Condition 6.3.2, let $(S,T) \in E$, and let g be a good in S such that $T \subseteq (M\backslash S) \cup \{g\}$. Then $g \in T$.*

Proof. Suppose $g \notin T$: then $S \subseteq M\backslash T$. Since S is minimal, player 1 is happy with S. Thus by monotonicity, player 1 is also happy with $M\backslash T$. But player 1 is also happy with T, because T is minimal. This contradicts Condition 6.3.2, so we must have $g \in T$. □

Lemma 6.4.7. *Assuming Condition 6.3.2, if $(S,T) \in E$, then there is a unique $g \in S$ where $T \subseteq (M\backslash S) \cup \{g\}$.*

Proof. Suppose there exist $g_1, g_2 \in S$ where $g_1 \neq g_2$, $T \subseteq (M\backslash S) \cup \{g_1\}$, and $T \subseteq (M\backslash S) \cup \{g_2\}$. Then by Lemma 6.4.6, $g_1 \in T$ and $g_2 \in T$. But this contradicts $T \subseteq (M\backslash S) \cup \{g_1\}$, because $g_2 \in S\backslash\{g_1\}$, so $g_2 \notin (M\backslash S) \cup \{g_1\}$. Therefore $g_1 = g_2$. □

Using Lemma 6.4.7 for each edge $(S,T) \in E$, let $g(S,T)$ be the unique good such that $T \subseteq (M\backslash S) \cup \{g(S,T)\}$.

Lemma 6.4.8. *Assuming Condition 6.3.2, if $(S,T) \in E$, then $(T,S) \in E$. Furthermore, $g(T,S) = g(S,T)$.*

Proof. Suppose $(S,T) \in E$: then $T \subseteq (M\backslash S) \cup \{g(S,T)\}$. By Lemma 6.4.6, we have $g(S,T) \in T$. Since $T \subseteq (M\backslash S) \cup \{g(S,T)\}$, we have $S\backslash\{g(S,T)\} \subseteq M\backslash T$. Therefore $S \subseteq (M\backslash T) \cup \{g(S,T)\}$, and so $(T,S) \in E$ and $g(S,T) = g(T,S)$. □

The next lemma is because there are $|S|$ items in S that we could move to $M\backslash S$. The proof uses Lemma 6.4.7 to show that each of them will yield a different minimal bundle T, so this constitutes $|S|$ distinct edges (S,T).

Lemma 6.4.9. *The out-degree of each bundle $S \in V$ is at least $|S|$.*

Proof. Let $S = \{g_1, g_2 ... g_{|S|}\}$. We first claim that for all $g_j \in S$, there exists $T_j \in V$ where $T_j \subseteq (M \backslash S) \cup \{g_j\}$. Consider some $g_j \in S$. Because S is minimal, we know that player 1 is not happy with $S \backslash \{g_j\}$. Therefore player 1 must be happy with $(M \backslash S) \cup \{g_j\}$. Then by Lemma 6.4.2, there exists $T_j \subseteq (M \backslash S) \cup \{g_j\}$ where T_j is minimal. Therefore $(S, T_j) \in E$.

By Lemma 6.4.7, $g_j = g(S, T_j)$ is unique. Thus for all $g \in S$ where $g \neq g(S, T_j)$, we have $g \notin T_j$. This implies that each T_j is distinct. Thus $(S, T_1), (S, T_2) ... (S, T_{|S|})$ are all distinct edges in E, so the out-degree of S is at least $|S|$. ☐

Next, we define a set of edges $E_+ \subseteq E$ by

$$E_+ = \{(S, T) \mid \Delta(S \backslash \{g(S, T)\}, g(S, T)) \geq \alpha\}$$

This is the set of "special edges" alluded to in Section 6.1.2.

The informal argument for the next lemma is as follows. By Lemma 6.4.8, we have $(S, T) \in E$ if and only if $(T, S) \in E$. Then Lemma 6.4.5 (combined with submodularity) implies that at least one of $\Delta(S \backslash \{g(S, T)\}, g(S, T)) \geq \alpha$ and $\Delta(T \backslash \{g(S, T)\}, g(S, T)) \geq \alpha$ is true, so at least one of (S, T) and (T, S) must be in E_+.

Lemma 6.4.10. *Assuming Condition 6.3.2,* $|E_+| \geq |E|/2$.

Proof. Let (S, T) be some edge in E: then by Lemma 6.4.8, $(T, S) \in E$. It suffices to show that for every edge $(S, T) \in E$, at least one of (S, T) and (T, S) are in E_+. Assume $(S, T) \notin E_+$: otherwise we are done. Then

$$\Delta\big(S \backslash \{g(S, T)\}, g(S, T)\big) < \alpha$$

Thus by Lemma 6.4.5,

$$\Delta\big(M \backslash S, g(S, T)\big) \geq \alpha$$

Since (T, S) is an edge in the graph, $S \subseteq (M \backslash T) \cup \{g(S, T)\}$. Therefore $S \backslash \{g(S, T)\} \subseteq M \backslash T$. Thus by submodularity, $\Delta(S \backslash \{g(S, T)\}, g(S, T)) \geq \Delta(M \backslash T, g(S, T)) \geq \alpha$. Therefore $(S, T) \in E_+$. ☐

Lemma 6.4.11 follows from a simple counting argument.

Lemma 6.4.11. *For any integers m and ℓ,* $\sum\limits_{j=0}^{\ell} \dbinom{m}{j} \leq (m+1)^{\ell}$.

Proof. The left-hand-side is number of subsets of $[m]$ of size at most ℓ. The right-hand-side is the number of ways to select ℓ elements from $[m] \cup \{d\}$, where each element can be selected multiple times, and including ordering. We think of d as a dummy element. For each subset $S \subseteq [m]$ counted by $\sum_{j=0}^{\ell} \binom{m}{j}$, we represent it in $(m+1)^{\ell}$ as follows: first select element d $\ell - |S|$ times, and then select the elements in S in any order. Thus each subset of $[m]$ counted by the left-hand-side is represented in a unique way by the right-hand-side, and so $\sum_{j=0}^{\ell} \binom{m}{j} \leq (m+1)^{\ell}$. ☐

We are now ready to prove the final theorem. Recall the following definitions from Section 6.3:

$$S_{\leq k} = \{g_j \in S \mid j \leq k\}$$

$$\delta_k^S = v_1(S_{\le k}) - v_1(S_{\le k-1})$$

Theorem 6.4.1. *For two players with submodular valuations and any $c < 1$, Protocol 8 has communication cost at most $2m(m+1)^{\frac{8}{1-c}} + 2v^{size}$, and either returns a c-EF allocation or a c^*-EF allocation. This also implies that*

$$D_{submod}(2, m, EF, c) \le 2m(m+1)^{\frac{8}{1-c}} + 2v^{size}$$

for any $c < 1$.

Proof. Correctness of Protocol 8 follows from Lemma 6.4.4, so it remains only to bound the communication cost.

We prove that the number of minimal bundles is (strictly) less than $2(m+1)^{\frac{8}{1-c}} = 2(m+1)^{4/\alpha}$, assuming Condition 6.3.2. Let $\beta = 4/\alpha$, and suppose that the number of minimal bundles is at least $2(m+1)^{4/\alpha} = 2(m+1)^{\beta}$. By Lemma 6.4.11, the number of minimal bundles of size at most β is at most $(m+1)^{\beta}$. Thus there are at least $(m+1)^{\beta}$ minimal bundles S where $|S| > \beta$.

So at least half of the minimal bundles have size more than β. Let $G = (V, E)$ be the minimal bundle graph. Then by Lemma 6.4.9, at least half of the minimal bundles in V have out-degree more than β. Therefore $|E| > \beta|V|/2$. Then by Lemma 6.4.10, $|E_+| > \beta|V|/4 = |V|/\alpha$.

For a bundle S, let X_+^S be the set of out-edges from S that are in E_+. Formally,

$$X_+^S = \{(S,T) \in E \mid \Delta(S\backslash\{g(S,T)\}, g(S,T)) \ge \alpha\}$$

and we can define the corresponding goods by $g(X_+^S) = \{g \in S \mid \Delta(S\backslash\{g\}, g) \ge \alpha\}$.

We next show that there must exist a minimal bundle $S \in V$ where $|X_+^S| > 1/\alpha$. Suppose that $|X_+^S| \le 1/\alpha$ for all $S \in V$: then

$$|E_+| \le |V|/\alpha$$

which contradicts $|E_+| > |V|/\alpha$. Therefore there exists some bundle S with $|X_+^S| > 1/\alpha$. By definitions, we have

$$v_1(S) = \sum_{k=1}^{m} \delta_k^S = \sum_{k: g_k \in S} \delta_k^S = \sum_{k: g_k \in S} \Delta(S_{\le k-1}, g_k) \ge \sum_{k: g_k \in g(X_+^S)} \Delta(S_{\le k-1}, g_k)$$

Because $S_{\le k-1} \subseteq S$ and $g_k \notin S_{\le k-1}$, we have $S_{\le k-1} \subseteq S\backslash\{g_k\}$. Therefore by submodularity,[12]

$$\sum_{k: g_k \in g(X_+^S)} \Delta(S_{\le k-1}, g_k) \ge \sum_{k: g_k \in g(X_+^S)} \Delta(S\backslash\{g_k\}, g_k) \ge \sum_{k: g_k \in g(X_+^S)} \alpha = \alpha|X_+^S| > 1$$

But $v_1(M) = 1$, so this is a contradiction. Therefore the number of minimal bundles is less than $2(m+1)^{\frac{8}{1-c}}$.

[12]This is the crucial use of submodularity: that we can add in the items in S one by one, and the value of the set increases by at least $\Delta(S\backslash\{g_k\}, g_k)$ each time. This allows us to pump the value of S over $v_1(M)$.

Thus the number of minimal bundles is at most $2(m+1)^{\frac{8}{1-c}} - 1$. If the protocol terminates in step 1, just one bundle is communicated (and zero values), so the communication cost bound is trivially satisfied. Suppose the protocol does not terminate in step 1: then player 1 sends at most $2(m+1)^{\frac{8}{1-c}} - 1$ minimal bundles, as well as $S^*(v_1)$. Thus at most $2(m+1)^{\frac{8}{1-c}}$ bundles are sent, each of which require m bits to communicate.

Player 1 also sends $\mathbf{c}_i(S^*(v_1))$. By definition of $\mathbf{c}_i^{EF}$, $\mathbf{c}_i(S^*(v_1))$ can be expressed as the ratio of two values, each of which takes v^{size} bits to communicate. Therefore the total communication cost is

$$2m(m+1)^{\frac{8}{1-c}} + 2v^{size}$$

as required. □

We will show formally in Section 6.6 that Theorem 6.4.1 is tight, meaning that exponential communication can be required when $c = 1$. To see why the minimal bundle argument fails for $c = 1$, consider an additive (and hence submodular) valuation over an even number of items, where the value of each item is one. Then a bundle is minimal if and only if it contains exactly half the items, and there are exponential number of such bundles.

6.5 Lower bound approach

In Section 6.6, we will prove a lower bound that matches the PAS from Section 6.4. Before we do that, we describe our general lower bound approach in this section. All of our lower bounds will rely on reductions from two well-known problems in communication complexity: determining whether two bit strings are equal, and determining whether two bit strings are disjoint. Let x_i denote the bit string held by player i, and let x_{ij} denote the jth bit of x_i. An input (x_1, x_2) is a yes-instance of the EQUALITY problem if and only if $x_{1j} \neq x_{2j}$ for all j. An input (x_1, x_2) is a yes-instance of the DISJOINTNESS problem if and only if there exists no j such that $x_{1j} = x_{2j} = 1$. The following lemma states that DISJOINTNESS is hard in the randomized setting (and thus also in the deterministic setting).

Lemma 6.5.1 ([106, 151])**.** *Any randomized protocol which solves* DISJOINTNESS *for bit strings of length ℓ has communication cost $\Omega(\ell)$.*

The following well-known lemma states that EQUALITY is hard in the deterministic setting.

Lemma 6.5.2. *Any deterministic protocol which solves* EQUALITY *for bit strings of length ℓ has communication cost at least ℓ.*

Perhaps surprisingly, EQUALITY admits a constant communication randomized protocol, due to [177].

Lemma 6.5.3 ([177])**.** *There exists a randomized protocol which solves* EQUALITY *and has communication cost $O(1)$.*

The protocol for Lemma 6.5.3 asks each player to compute the inner product mod 2 of her bit string and a random string. The protocol then compares those inner products. The Principle of Deferred Decisions can be used to show that this protocol arrives at the correct answer with probability at least 3/4. Lemma 6.5.3 will be a key element of our randomized upper bound in Section 6.9.3.

All of our lower bounds have the following structure. Given two bit strings x_1 and x_2 of length $\ell = \Omega(\binom{2k}{k})$, we construct a corresponding instance of Fair Division with $O(k)$ items. In the two player case, each index in the bit strings will correspond to a possible allocation that gives each player k items.

Our constructed instance will have that a property that a c-P allocation exists if and only if (x_1, x_2) is a no-instance[13] of Equality (for a deterministic lower bound), or a no-instance of Disjointness (for a randomized lower bound). Thus if there existed a protocol for Fair Division with communication cost less than $\Omega(\binom{2k}{k})$, it could also be used to solve Equality or Disjointness in communication less than $\Omega(\ell)$. This is impossible according to Lemmas 6.5.1 and 6.5.2, so any protocol for Fair Division requires exponential communication.

Using this framework, all that is needed to prove a lower bound for a particular set of parameters (property P, constant c, and a valuation class) is:

1. Given bit strings x_1 and x_2 of length $\Omega(\binom{2k}{k})$, define how to construct a corresponding instance of Fair Division with $O(k)$ items.

2. Show that a c-P allocation exists in the constructed instance if and only if (x_1, x_2) is a no-instance of Equality or Disjointness.

3. Show that the valuations in the constructed instance of Fair Division are of the desired valuation class.

More specifically, our Fair Division instance will have two players and $2k$ items. Valuations will be constructed such that a player will never be happy if she receives fewer than k items, so both players will have to receive exactly k items. There are $\binom{2k}{k}$ allocations which give each player k items, and this gives rise to the exponential communication lower bound.

In fact, we can do this in a very standardized way for the two player deterministic case. Given bit strings of length $\frac{1}{2}\binom{2k}{k}$, we define a list of allocations $\mathcal{T} = (T_1, T_2 ... T_{|\mathcal{T}|})$ where each $T_j = (T_{j1}, T_{j2}) \in \mathcal{T}$ is an allocation giving each player k items: $|T_{j1}| = |T_{j2}| = k$. There two important properties we will need $\mathcal{T}$ to have. First, $\mathcal{T}$ should not contain every such allocation: in particular, for any allocation $A \in \mathcal{T}$, $\overline{A} \notin \mathcal{T}$.[14] Second, the order of allocations in $\mathcal{T}$ cannot depend on the input strings. This order is arbitrary, but publicly known. Note that $|\mathcal{T}| = \frac{1}{2}\binom{2k}{k}$.

Lemma 6.5.4 states that under this approach, all that is necessary to complete the lower bound is to construct valuations satisfying three particular properties. The exact way valuations are constructed will depend on what class we wish them to belong to (general, subadditive, or submodular). We only prove the lemma for the c-EF in the two player deterministic setting. A similar result is

[13]Note that no-instances of Equality or Disjointness become instances where a c-P allocation *does* exist.
[14]Recall that for $A = (A_1, A_2)$, $\overline{A} = (A_2, A_1)$.

possible for other settings, but this is only setting where we prove enough different lower bounds to make it worth having a separate lemma.

For a bit string x_i, let $\overline{x_i}$ denote the string obtained by flipping every bit: $x_{ij} \neq \overline{x_{ij}}$ for all j. We will define two new bit strings, y_1 and y_2, by $y_1 = x_1$ and $y_2 = \overline{x_2}$. Also, recall that for a player i, $\bar{i}$ denotes the other player.

The lemma relies on three conditions. Condition 6.5.1 states that neither player is happy with any bundle containing fewer than k items: then any c-P allocation must either be A or $\overline{A}$ for some $A \in \mathcal{T}$. Condition 6.5.2 states that player i is unhappy receiving $T_{j\bar{i}}$ when $y_{ij} = 1$ (and happy receiving T_{ji}). Condition 6.5.3 states that player i is unhappy receiving T_{ji} when $y_{ij} = 0$ (and happy receiving $T_{j\bar{i}}$). Thus we want to find an index j where either $y_{1j} = y_{2j} = 1$, in which case the allocation (T_{j1}, T_{j2}) is c-P, or where $y_{1j} = y_{2j} = 0$, in which case the allocation (T_{j2}, T_{j1}) is c-P. Therefore we are looking for an index where $y_{1j} = y_{2j}$, which is equivalent to $x_{1j} \neq x_{2j}$. This is exactly the EQUALITY problem.

Lemma 6.5.4. *Given bit strings x_1, x_2, each of length $\frac{1}{2}\binom{2k}{k}$ for some integer k, let $M = [2k]$ and $N = [2]$. Let $y_1 = x_1$ and $y_2 = \overline{x_2}$, and let c be some constant. Let $\mathcal{T} = (T_1, T_2 ... T_{|\mathcal{T}|})$ be a list of allocations as described above. Suppose v_1, v_2 can be constructed such that the following conditions are met:*

Condition 6.5.1. *For all $|S| < k$ and both i, $v_i(S) < c \cdot v_i(M \backslash S)$.*

Condition 6.5.2. *Whenever $y_{ij} = 1$, $v_i(T_{j\bar{i}}) < c \cdot v_i(T_{ji})$.*

Condition 6.5.3. *Whenever $y_{ij} = 0$, $v_i(T_{ji}) < c \cdot v_i(T_{j\bar{i}})$.*

Then any deterministic protocol which finds a c-EF allocation for two players requires exponential communication. Specifically,

$$D(2, 2k, EF, c) \geq \frac{1}{2}\binom{2k}{k}$$

Proof. We reduce from EQUALITY. Given bit strings x_1 and x_2 of length $\frac{1}{2}\binom{2k}{k}$ for some integer k, we construct the following instance of FAIR DIVISION. Let $N, M, (y_1, y_2)$, and $\mathcal{T}$ be as defined in the statement of Lemma 6.5.4. Also assume that v_1 and v_2 satisfy Conditions 6.5.1, 6.5.2, and 6.5.3.

Suppose that (x_1, x_2) is a no-instance of EQUALITY: then there exists j where $x_{1j} \neq x_{2j}$. Therefore $y_{1j} = y_{2j}$. If $y_{1j} = y_{2j} = 1$, then by Condition 6.5.2,

$$v_i(T_{ji}) > \frac{1}{c} v_i(T_{j\bar{i}}) \geq c \cdot v_i(T_{j\bar{i}})$$

for both i. Thus the allocation T_j is c-EF, because each player i receives T_{ji}. If $y_{1j} = y_{2j} = 0$, then by Condition 6.5.3,

$$v_i(T_{j\bar{i}}) > \frac{1}{c} v_i(T_{ji}) \geq c \cdot v_i(T_{ji})$$

for both i. Thus the allocation $\overline{T_j}$ is c-EF, because each player i receives $T_{j\bar{i}}$. Therefore if (x_1, x_2) is a no-instance of EQUALITY, there exists an allocation satisfying c-EF.

Suppose that (x_1, x_2) is a yes-instance of EQUALITY: then for every j, $y_{1j} \neq y_{2j}$. For any allocation A where $|A_i| < k$ for some i, we have $v_i(A_i) < c \cdot v_i(A_{\bar{i}})$ by Condition 6.5.1. Thus A cannot be c-EF whenever $|A_i| < k$ for some i.

Now consider an arbitrary allocation A where $|A_1| = |A_2| = k$. For any such allocation, there must exist j where either $A = T_j$, or $A = \overline{T_j}$. Since $y_{1j} \neq y_{2j}$, there exists a player i where $y_{ij} = 0$, and $y_{\bar{i}j} = 1$. Then by Condition 6.5.3, $v_i(T_{ji}) < c \cdot v_i(T_{j\bar{i}})$. Also, $v_{\bar{i}}(T_{j\bar{\bar{i}}}) < c \cdot v_{\bar{i}}(T_{j\bar{i}})$ by Condition 6.5.2, where $\bar{\bar{i}} = i$ represents the player other than $\bar{i}$. Thus $v_{\bar{i}}(T_{ji}) < c \cdot v_{\bar{i}}(T_{j\bar{i}})$.

Therefore neither player is happy with bundle T_{ji}. But since either $A = T_j$ or $A = \overline{T_j}$, there must be a player who receives T_{ji}, is hence is not happy. Thus no allocation where $|A_1| = |A_2| = k$ can be c-EF, no allocation is c-EF. □

This lemma will be useful in a variety of settings. In the next section, we will use this lemma to prove a lower bound for 1-EF that matches the PAS from Section 6.4.

6.6 1-EF is hard for submodular valuations

In this section, we use the general approach described in Section 6.5 to show that 1-EF requires exponential communication, even for two players with submodular valuations. This shows that the PAS for this setting from Section 6.4 is the best we can hope for.

Formally, Section 6.4 showed that $D_{submod}(2, m, \text{EF}, c)$ is polynomial in m when $c < 1$. We now show that $D_{submod}(2, m, \text{EF}, c)$ is exponential when $c = 1$. Section 6.3 showed that $D_{submod}(2, m, \text{Prop}, c)$ is polynomial for any c, so there is no lower bound necessary there. Thus this section resolves the deterministic submodular case for two players.

Theorem 6.6.1. *For two players with submodular valuations, any deterministic protocol which determines whether a* 1*-EF allocation exists requires an exponential amount of communication. Specifically,*

$$D_{submod}(2, 2k, EF, 1) \geq \frac{1}{2}\binom{2k}{k}$$

Proof. Given bit strings of length $\frac{1}{2}\binom{2k}{k}$ for some integer k, define $M, N, (y_1, y_2)$, and $\mathcal{T}$ as in Lemma 6.5.4. We need only to construct submodular valuations v_1, v_2 such that Conditions 6.5.1, 6.5.2, and 6.5.3 are met. We define each v_i by

$$v_i(S) = \begin{cases} 3|S| & \text{if } \ |S| < k \\ 3k & \text{if } \ |S| > k \\ 3k & \text{if } \ S = T_{ji} \ \text{ and } \ y_{ij} = 1 \\ 3k & \text{if } \ S = T_{j\bar{i}} \ \text{ and } \ y_{ij} = 0 \\ 3k-1 & \text{if } \ S = T_{ji} \ \text{ and } \ y_{ij} = 0 \\ 3k-1 & \text{if } \ S = T_{j\bar{i}} \ \text{ and } \ y_{ij} = 1 \end{cases}$$

Importantly, for every bundle S with $|S| = k$, there exists exactly one pair (i, j) such that $S = T_{ji}$. Thus if $|S| = k$, S falls under exactly one of the last four cases in the definition of v_i.

If $|S| < k$, we have $|M \backslash S| > k$, so $v_i(S) < 3k = v_i(M \backslash S)$. This satisfies Condition 6.5.1. Suppose $y_{ij} = 1$ for some i, j: then $v_i(T_{\bar{j}\bar{i}}) = 3k - 1 < 3k = v_i(T_{ji})$, so Condition 6.5.2 is satisfied. Suppose $y_{ij} = 0$ for some i, j: then similarly, $v_i(T_{ji}) = 3k - 1 < 3k = v_i(T_{\bar{j}\bar{i}})$. Thus Condition 6.5.3 is satisfied as well.

It remains to show that the valuations are submodular. To do this, we examine $v_i(S \cup \{g\}) - v_i(S)$, for any bundle S and item $g \notin S$.

$$v_i(S \cup \{g\}) - v_i(S) = \begin{cases} 3 & \text{if } |S \cup \{g\}| < k \\ 2 \text{ or } 3 & \text{if } |S \cup \{g\}| = k \\ 0 \text{ or } 1 & \text{if } |S \cup \{g\}| = k + 1 \\ 0 & \text{if } |S \cup \{g\}| > k + 1 \end{cases}$$

Therefore $v_i(S \cup \{g\}) - v_i(S)$ is non-increasing with $|S|$. Thus $v_i(X \cup \{g\}) - v_i(X) \geq v_i(Y \cup \{g\}) - v_i(Y)$ whenever $|X| < |Y|$. If $X \subseteq Y$, either $|X| < |Y|$ or $X = Y$. When $X = Y$, we trivially have $v_i(X \cup \{g\}) - v_i(X) = v_i(Y \cup \{g\}) - v_i(Y)$. Thus we have $v_i(X \cup \{g\}) - v_i(X) \geq v_i(Y \cup \{g\}) - v_i(Y)$ whenever $X \subseteq Y$, and so v_i is submodular.

□

Recall that Section 6.4 gave a PAS for this setting, where for any fixed c, communication at most $2(m+1)^{\frac{8}{1-c}}$ is required. In a fully polynomial-communication approximation scheme (FPAS), the dependence in $\frac{1}{1-c}$ is required to be polynomial. The PAS from Section 6.4 is not an FPAS, since the dependence on $\frac{1}{1-c}$ is exponential.

The above proof of Theorem 6.6.1 actually shows that for any $c > \frac{3k-1}{3k} = \frac{3m-2}{3m}$, exponential communication is required. This does not contradict the PAS from Section 6.4, because $\frac{3m-2}{3m}$ is not a fixed constant (it depends on m). However, this does rule out the possibility of an FPAS. To see this, suppose an FPAS existed, and consider some $c > \frac{3m-2}{3m}$. Then the FPAS would have communication cost polynomial in $\frac{1}{1-c}$. We have

$$\frac{1}{1-c} > \frac{1}{1 - \frac{3m-2}{3m}} = \frac{3m}{2}$$

so the communication cost is polynomial of m. But the proof of Theorem 6.6.1 shows communication exponential in m is required, which is a contradiction.

Finally, we note that the proof of Theorem 6.6.1 can easily be adapted to prove exponential lower bounds on the communication complexity of maximizing Nash welfare (the product of player valuations) or egalitarian welfare (the minimum player valuation).

6.7 Everything is hard for more than two players

In this section, we show that FAIR DIVISION requires an exponential amount of communication whenever there are more than two players: even when randomization is allowed, even for submodular valuations, and for any $c > 0$. This will allow us to focus on the two player setting for the rest of the chapter.

Before proving the theorems, we discuss the multiparty (i.e., $n > 2$) communication complexity model. As mentioned in Section 6.2, there is more than one such model. This will turn out not to matter in our setting. The reason is that our lower bounds will hold even when only player 1 and player 2 have private valuations, and the valuations of all other players are public information. One can think of the other players as not really being agents, and just being a (publicly known) part of the input. Thus we never actually consider multiparty communication. In this way, the theorem that we are really proving is that when there are more than two FAIR DIVISION players, the problem is hard in the two-party communication complexity model.

We first prove hardness for envy-freeness, and then reduce envy-freeness to proportionality. Recall that DISJOINTNESS has randomized communication complexity $\Omega(\ell)$, where ℓ is the length of the bit strings (Lemma 6.5.1).

Theorem 6.7.1. *For any $n > 2$ and any $c > 0$, any randomized protocol which determines whether a c-EF allocation exists requires an exponential amount of communication, even for submodular valuations. Specifically,*

$$R_{submod}(n, 2k + n - 2, EF, c) \in \Omega\left(\binom{2k}{k}\right)$$

for any $n > 2$ and $c > 0$.

Proof. We reduce from DISJOINTNESS. Given bit strings x_1 and x_2 of length $\binom{2k}{k}$, we construct a fair division instance as follows. Although there will be more than two players, there are only two bit strings. Let player 1 hold x_1 and player 2 hold x_2, and the other players will have no bit strings.

Let $M_1 = [2k], M_2 = \{g_3...g_n\}$, and $M = M_1 \cup M_2$: note that $|M| = 2k + n - 2$. Let $N = [n]$. We define a similar list of allocations $\mathcal{T} = (T_1, T_2...)$, where $T_j = (T_{j1}, T_{j2})$. Here each T_j is an allocation over only M_1, and for just two players. Any such allocation A where $|A_1| = |A_2| = k$ is in $\mathcal{T}$ (and so is $\overline{A}$). Note that $|\mathcal{T}| = \binom{2k}{k}$. For $i \in \{1, 2\}$, v_i is given by

$$v_i(S) = \begin{cases} k & \text{if } \ g_3 \in S \\ |S|c & \text{if } \ |S| < k \ \text{ and } \ g_3 \notin S \\ kc & \text{if } \ |S| > k \ \text{ and } \ g_3 \notin S \\ (k - \frac{1}{2})c & \text{if } \ |S| = k \ \text{ and } \ g_3 \notin S \ \text{ and } \ S \cap M_2 \neq \emptyset \\ kc & \text{if } \ \exists j \ \ S = T_{ji} \ \text{ where } \ x_{ij} = 1 \ \text{ and } \ g_3 \notin S \\ (k - \frac{1}{2})c & \text{if } \ \exists j \ \ S = T_{ji} \ \text{ where } \ x_{ij} = 0 \ \text{ and } \ g_3 \notin S \end{cases}$$

Every allocation giving each player k items occurs in $\mathcal{T}$ exactly once. Thus when $|S| = k$ and

$S \subset M_1$, exactly one of the last two cases occurs, and any such j must be unique. For $i > 2$, $v_i(S)$ is given by

$$v_i(S) = \begin{cases} 1 & \text{if } g_i \in S \\ 0 & \text{otherwise} \end{cases}$$

Suppose that (x_1, x_2) is a no-instance of DISJOINTNESS: then there exists j where $x_{1j} = x_{2j} = 1$. Consider the allocation A where $A_i = T_{ji}$ for $i \leq 2$, and $A_i = \{g_i\}$ for $i > 2$. For $i > 2$, $v_i(A_i) = 1$ and $v_i(A_{i'}) = 0$ for all $i' \neq i$, so each player $i > 2$ is happy. For $i \leq 2$, we have $v_i(A_i) = v_i(T_{ji}) = kc$, and $v_i(A_{i'}) \leq k$ for all i'. Therefore for all i, i', $v_i(A_i) \geq cv_i(A_{i'})$, so A is c-EF.

Suppose that (x_1, x_2) is a yes-instance of DISJOINTNESS: then for every j, there exists i where $x_{ij} = 0$. Suppose that a c-EF allocation $A = (A_1, A_2)$ exists. We first claim that for every $i > 2$, $g_i \in A_i$: if not, $v_i(A_i) = 0$, so player i will envy whichever player receives g_i.

Thus for $i \leq 2$,

$$v_i(A_i) \geq c \cdot v_i(A_3) \geq c \cdot v_i(\{g_3\}) = kc$$

Suppose a player $i \leq 2$ receives strictly fewer than k items in A_i: then $v_i(A_i) < kc$, since none of those items can be g_3. This is a contradiction, so we have $|A_1| = |A_2| = k$. Since $\mathcal{T}$ contains all of the allocations which give each player k items, there must exist j where $A_i = T_{ji}$ for both i, and $v_i(A_i) \geq kc$. But that implies that $x_{1j} = x_{2j} = 1$, which is a contradiction. Therefore no allocation is c-Prop.

It remains to show that the valuations are submodular. For $i > 2$, v_i is trivially submodular. We now we examine $v_i(S \cup \{g\}) - v_i(S)$ for $i \leq 2$, any bundle S, and any item $g \notin S$ where $g_3 \notin S \cup \{g\}$.

$$v_i(S \cup \{g\}) - v_i(S) = \begin{cases} c & \text{if } |S \cup \{g\}| < k \\ c \text{ or } c/2 & \text{if } |S \cup \{g\}| = k \\ c/2 \text{ or } 0 & \text{if } |S \cup \{g\}| = k+1 \\ 0 & \text{if } |S \cup \{g\}| > k+1 \end{cases}$$

Therefore $v_i(S \cup \{g\}) - v_i(S)$ is non-increasing with $|S|$ when $g_3 \notin S \cup \{g\}$. Thus $v_i(X \cup \{g\}) - v_i(X) \geq v_i(Y \cup \{g\}) - v_i(Y)$ whenever $|X| < |Y|$ and $g_3 \notin S \cup \{g\}$. If $X \subseteq Y$, either $|X| < |Y|$ or $X = Y$. When $X = Y$, we trivially have $v_i(X \cup \{g\}) - v_i(X) = v_i(Y \cup \{g\}) - v_i(Y)$. Thus we have $v_i(X \cup \{g\}) - v_i(X) \geq v_i(Y \cup \{g\}) - v_i(Y)$ whenever $X \subseteq Y$ and $g_3 \notin S \cup \{g\}$. Therefore the submodularity condition is satisfied when $g_3 \notin S \cup \{g\}$.

There are two remaining cases: when $g_3 \in S$, or when $g = g_3$. For $g_3 \in S$, $v_i(S \cup \{g\}) - v_i(S) = 0$ for all S and g, so the condition is satisfied in this case. For $g = g_3$, we have $v_i(X \cup \{g_3\}) - v_i(X) = v_i(M) - v_i(X)$ and $v_i(Y \cup \{g_3\}) - v_i(Y) = v_i(M) - v_i(Y)$. If $X \subseteq Y$, we have $v_i(X) \leq v_i(Y)$, so $v_i(X \cup \{g_3\}) - v_i(X) \geq v_i(Y \cup \{g_3\}) - v_i(Y)$. Therefore v_i is submodular for all i. □

We now prove hardness for proportionality for more than two players, by reducing from envy-freeness.

Theorem 6.7.2. *For any $n > 2$ and any $c > 0$, any randomized protocol which determines whether* *a c-Prop allocation exists requires an exponential amount of communication, even for submodular valuations. Specifically,*

$$R_{submod}(n, 2k+n-2, Prop, c) \in \Omega\left(\binom{2k}{k}\right)$$

for any $c > 0$.

Proof. We reduce from FAIR DIVISION for $P = \text{EF}$. Given an input (x_1, x_2), we define v_i as in the proof of Theorem 6.7.1, except using c/n instead of c. That is, for $i \leq 2$,

$$v_i(S) = \begin{cases} k & \text{if } \; g_3 \in S \\ |S|c/n & \text{if } \; |S| < k \; \text{ and } \; g_3 \notin S \\ kc/n & \text{if } \; |S| > k \; \text{ and } \; g_3 \notin S \\ (k-\frac{1}{2})c/n & \text{if } \; |S| = k \; \text{ and } \; g_3 \notin S \; \text{ and } \; S \cap M_2 \neq \emptyset \\ kc/n & \text{if } \; \exists j \;\; S = T_{ji} \; \text{ where } \; x_{ij} = 1 \; \text{ and } \; g_3 \notin S \\ (k-\frac{1}{2})c/n & \text{if } \; \exists j \;\; S = T_{ji} \; \text{ where } \; x_{ij} = 0 \; \text{ and } \; g_3 \notin S \end{cases}$$

It was shown in the proof of Theorem 6.7.1 that these valuations are submodular.

Theorem 6.7.1 implies that $\Omega(\binom{2k}{k})$ communication is required to determine whether a $\frac{c}{n}$-EF allocation exists under these valuations. We will show that under these valuations, an allocation is c-Prop if and only if it is $\frac{c}{n}$-EF. This will imply that determining whether a c-Prop allocation exists is just as hard as whether a $\frac{c}{n}$-EF allocation exists.

In order for an allocation A to be $\frac{c}{n}$-EF or c-Prop, we must have $v_i(A_i) > 0$ for all i. Thus assume $g_i \in A_i$ for all $i > 2$, and we need only consider $i \leq 2$.

Suppose an allocation A is c-Prop: then for $i \leq 2$, $v_i(A_i) \geq \dfrac{c}{n} v_i(M) = \dfrac{kc}{n}$. Since $v_i(A_{i'}) \leq k$ for all i', we have

$$v_i(A_i) \geq \frac{kc}{n} \geq \frac{c}{n} v_i(A_{i'})$$

for all i'. Therefore A is $\frac{c}{n}$-EF.

Suppose an allocation A is $\frac{c}{n}$-EF: then for all i and i', $v_i(A_i) \geq \frac{c}{n} v_i(A_{i'})$. For $i \leq 2$, we have $v_i(A_3) \geq v_i(\{g_3\}) = k$, so

$$v_i(A_i) \geq \frac{c}{n} v_i(A_3) \geq \frac{kc}{n} = \frac{c}{n} v_i(M)$$

so A is c-Prop. □

This resolves the $n > 2$ case for all combinations of other parameters, so we will assume that $n = 2$ for the remainder of the chapter.

6.8 Subadditive valuations

In this section, we consider the deterministic setting for two players with subadditive valuations. In Section 6.8.1, we use the Minimal Bundle Protocol from Section 6.4 to show that c-EF for $c \leq 1/2$

Protocol 8 Protocol for two players to either find a c-P allocation or a c^*-P allocation.

Private inputs: v_1, v_2
Public inputs: P, c

1. If there exists an allocation A where player 1 is happy with both A and $\overline{A}$, player 1 sends that allocation to player 2. If player 2 is happy with A, she declares that A is c-P, otherwise she declares that $\overline{A}$ is c-P.
2. If there is no such allocation A, player 1 sends the set $\mathcal{S}$ of her minimal bundles to player 2. She also sends the bundle $S^*(v_1)$ and the value $\mathbf{c}_1(S^*(v_1))$.
3. Player 2 first checks if there exists a bundle $S \in \mathcal{S}$ where player 2 is happy with $M \backslash S$. If so, she declares that $(S, M \backslash S)$ is c-P.
4. If not, player 2 computes $S^*(v_2)$ and $i = \arg\max_{i' \in \{1,2\}} \mathbf{c}_{i'}(S^*(v_{i'}))$. Let A be the allocation where $A_i = S^*(v_i)$ and $A_{\bar{i}} = M \backslash S^*(v_i)$. Player 2 then declares that A is $\mathbf{c}_i(S^*(v_i))$-P, and that $c^* = \mathbf{c}_i(S^*(v_i))$.

and c-Prop for $c \leq 2/3$ require only polynomial communication. This is the same protocol that yielded the PAS for EF with submodular valuations, but the communication cost analysis will be different. In Section 6.8.2, we show that this is tight, by giving an exponential lower bound for c-EF and c-Prop when c exceeds $1/2$ and $2/3$, respectively.

6.8.1 Upper bounds

In this section, we prove that when players have subadditive valuations, the Minimal Bundle Protocol (Protocol 8) can be used to solve Fair Division for $\frac{1}{2}$-EF and $\frac{2}{3}$-Prop with polynomial communication. In fact, we will show that if a satisfactory allocation is not found in step 1, there must exist a single item g where $v_1(\{g\}) > v_1(M \backslash \{g\})$. This will imply that the only minimal bundle is $\{g\}$. Protocol 8 is restated here for the convenience of the reader.

In Section 6.4, we proved correctness of this protocol for any setting, so it remains only to prove the communication cost bound for this setting.

Let $\alpha \in (0, 1]$ be some constant. Let $\eta^P(\alpha)$ be the maximum $c \leq 1$ for which any allocation A is guaranteed to be c-P, given $v_i(A_i) \geq \alpha v_i(A_{\bar{i}})$ for both i. For example, $\eta^{EF}(\alpha) = \alpha$. We will write $\eta^P(\alpha) = \eta(\alpha)$ and leave P implicit. Lemma 6.8.1 is strongest for $\alpha = 1/2$, but we find it insightful to prove the theorem for any $\alpha \leq 1/2$.

Also, recall that Condition 6.3.1 is satisfied for c-Prop with subadditive valuations, for any c: for any allocation A, each player must be happy with at least one of A and $\overline{A}$.

Lemma 6.8.1. *For two players with subadditive valuations, $\alpha \in (0, 1/2]$, and $c = \eta(\alpha)$, Protocol 8 has communication cost at most*

$$2(m + v^{size})$$

Proof. If the protocol terminates in step 1, a single allocation is communicated, which requires m bits. Thus the claim is satisfied in this case.

If the protocol does not terminate in step 1, the only communication happens in step 2. For a bundle S, $\mathbf{c}_i(S)$ is defined as the ratio of two values: $\frac{v_i(S)}{v_i(M\backslash S)}$ for EF, and $\frac{2v_i(S)}{v_i(M)}$ for Prop[15]. Thus communicating $\mathbf{c}_1(S^*(v_1))$ requires $2v^{size}$ bits. The only other information transmitted is the bundle $S^*(v_1)$ and $\mathcal{S}$. Communicating $S^*(v_1)$ requires m bits, and $\mathcal{S}$ requires $|\mathcal{S}|m$ bits. Thus the communication cost of the protocol is

$$m(|\mathcal{S}| + 1) + 2v^{size}$$

It remains to show that if the protocol does not terminate in step 1, $|\mathcal{S}| = 1$.

By Condition 6.3.1, for every allocation A, player 1 is happy with at least one of A and $\overline{A}$. Thus $|\mathcal{S}| \geq 1$, so let S be a minimal bundle in $\mathcal{S}$. Since player 1 is happy with S, we know that she is not happy with $M\backslash S$, or the protocol would have terminated in step 1. Suppose $\alpha v_1(S) \leq v_1(M\backslash S)$: then player 1 is $\eta(\alpha)$-happy with $M\backslash S$. Since $c = \eta(\alpha)$, this means player 1 is happy with $M\backslash S$, which is a contradiction. Therefore $\alpha v_1(S) > v_1(M\backslash S)$. This also implies that $v_1(S) > 0$.

Also, since S is minimal, player 1 is not happy with $S\backslash\{g\}$ for all $g \in S$. Therefore player 1 is happy with $(M\backslash S) \cup \{g\}$ for all $g \in S$. By the same argument as above, we have $\alpha v_1(M\backslash S) \cup \{g\}) > v_1(S\backslash\{g\})$.

Since $v_i(S) > 0$, S must be nonempty, so let g be an arbitrary item in S. By subadditivity of v_1, we have

$$v_1(M\backslash S) + v_1(\{g\}) \geq v_1\Big((M\backslash S) \cup \{g\}\Big)$$

Similarly,

$$\begin{aligned} v_1(S\backslash\{g\}) + v_1(\{g\}) &\geq v_1(S) \\ v_1(S\backslash\{g\}) &\geq v_1(S) - v_1(\{g\}) \end{aligned}$$

Therefore

$$\begin{aligned} v_1(M\backslash S) + v_1(\{g\}) &\geq v_1\Big((M\backslash S) \cup \{g\}\Big) \\ &> \frac{1}{\alpha} v_1(S\backslash\{g\}) \\ &\geq \Big(\frac{1}{\alpha} - 1\Big) v_1(S\backslash\{g\}) + v_1(S) - v_1(\{g\}) \end{aligned}$$

Since $v_1(S) > \frac{1}{\alpha} v_1(M\backslash S)$, we have

$$\begin{aligned} v_1(M\backslash S) + v_1(\{g\}) &> \Big(\frac{1}{\alpha} - 1\Big) v_1(S\backslash\{g\}) + \frac{1}{\alpha} v_1(M\backslash S) - v_1(\{g\}) \\ 2v_1(\{g\}) &\geq \Big(\frac{1}{\alpha} - 1\Big) v_1(S\backslash\{g\}) + \Big(\frac{1}{\alpha} - 1\Big) v_1(M\backslash S) \end{aligned}$$

[15]Technically only one value is needed for Prop, since we can assume that $v_1(M) = 1$, so only $v_1(S)$ is needed. However, since two values are needed for EF, we ignore this.

$$v_1(\{g\}) \geq \frac{1}{2}\Big(\frac{1}{\alpha} - 1\Big)\Big(v_1(S\backslash\{g\}) + v_1(M\backslash S)\Big)$$

By subadditivity of v_1, we have

$$\begin{aligned} v_1(S\backslash\{g\}) + v_1(M\backslash S) &\geq v_1\Big((S \cup (M\backslash S))\backslash\{g\}\Big) \\ &= v_1(M\backslash\{g\}) \end{aligned}$$

Therefore,

$$\begin{aligned} v_1(\{g\}) &> \frac{1}{2}\Big(\frac{1}{\alpha} - 1\Big)v_1(M\backslash\{g\}) \\ v_1(\{g\}) &\geq \alpha v_1(M\backslash\{g\}) \end{aligned}$$

where the final step is due to $0 < \alpha \leq 1/2$.

Thus player 1 is $\eta(\alpha)$-happy with the bundle $\{g\}$ by definition. Since the protocol did not terminate in step 1, player 1 must not be happy with $M\backslash\{g\}$. Therefore player i is happy with a bundle S if and only if $g \in S$, and so the only minimal bundle is $\{g\}$. □

Theorem 6.8.1 is immediately implied by the combination of Lemma 6.4.4 (correctness) and Lemma 6.8.1 (communication cost).

Theorem 6.8.1. *For two players with subadditive valuations and $c = \eta(1/2)$, Protocol 8 has communication cost at at most $2(m + v^{size})$, and either returns a c-P allocation or a c^*-P allocation.*

Theorem 6.8.1 immediately implies the following result.

Theorem 6.8.2. *For two players with subadditive valuations, a property P, and any constant $c \leq \eta^P(1/2)$, there exists a deterministic protocol with communication cost $2(m + v^{size})$ which solves* Fair Division*. Formally,*

$$D_{subadd}(2, m, P, c) \leq 2(m + v^{size})$$

for any $c \leq \eta^P(1/2)$.

Proof. Run Protocol 8 to either find a $\eta^P(1/2)$-P allocation, or to find a c'-P allocation where c' is the best possible. If a $\eta^P(1/2)$-P allocation exists, then a c-P allocation exists, since $c \leq \eta^P(1/2)$. If a c^*-P allocation is returned where $c^* < \eta^P(1/2)$, then by definition of c^*, a c-P allocation exists if and only if $c^* \geq c$. □

Theorem 6.8.3 is a direct consequence of Theorem 6.8.2 since $\eta^{EF}(\alpha) = \alpha$, and Theorem 6.8.4 requires only a short proof.

Theorem 6.8.3. *For two players with subadditive valuations, $P = EF$, and any constant $c \leq 1/2$, there exists a deterministic protocol with communication cost $2(m + v^{size})$ which solves* Fair Division*. Formally,*

$$D_{subadd}(2, m, EF, c) \leq 2(m + v^{size})$$

for any $c \le 1/2$.

Theorem 6.8.4. *For two players with subadditive valuations,* $P = Prop$*, and any constant* $c \le 2/3$*, there exists a deterministic protocol with communication cost* $2(m + v^{size})$ *which solves* Fair Division*. Formally,*

$$D_{subadd}(2, m, Prop, c) \le 2(m + v^{size})$$

for any $c \le 2/3$.

Proof. By Theorem 6.8.1, we need only show that $\eta^{Prop}(1/2) \ge 2/3$. Suppose $v_1(A_1) \ge \alpha v_1(A_2)$. Then $v_1(A_2) \le \frac{1}{\alpha} v_1(A_1)$, and by subadditivity of v_1, we have

$$\begin{aligned} v_1(M) &= v_1(A_1 \cup A_2) \\ &\le v_1(A_1) + v_1(A_2) \\ &\le v_1(A_1) + \frac{1}{\alpha} v_1(A_1) \\ &= \frac{\alpha + 1}{\alpha} v_1(A_1) \end{aligned}$$

Therefore $v_1(A_1) \ge \frac{2\alpha}{\alpha+1}\left(\frac{1}{2} v_1(M)\right)$, so $\eta^{Prop}(\alpha) \ge \frac{2\alpha}{\alpha+1}$. Therefore $\eta^{Prop}(1/2) \ge 2/3$. □

Since $\eta^{Prop}(1/2) \ge 2/3$, any $\frac{1}{2}$-EF A allocation is also $\frac{2}{3}$-Prop. However, a c'-EF allocation where c' is the maximum possible EF approximation ratio does not necessarily achieve the maximum possible approximation ratio for Prop. Consider the case where $M = \{g_1, g_2\}$ and the players valuations are given by

$$v_1(S) = \begin{cases} 9 & \text{if } S = M \\ 7 & \text{if } S = \{g_1\} \\ 2 & \text{if } S = \{g_2\} \end{cases} \qquad v_2(S) = \begin{cases} 4 & \text{if } S = M \\ 4 & \text{if } S = \{g_1\} \\ 1 & \text{if } S = \{g_2\} \end{cases}$$

and $v_i(\emptyset) = 0$ for both i. There is no $\frac{1}{2}$-EF allocation or $\frac{2}{3}$-Prop allocation in this instance. The allocation achieving the maximum EF approximation ratio is $A = (\{g_2\}, \{g_1\})$ which is $\frac{2}{7}$-EF. On the other hand, the allocation achieving the maximum Prop approximation ratio is $\overline{A}$, which is $\frac{1}{2}$-Prop.

6.8.2 Lower bounds

In this section, we show that $\frac{1}{2}$-EF and $\frac{2}{3}$-Prop are the best we can do deterministically for two players with subadditive valuations. We first prove that c-EF is hard for any $c > 1/2$, and then show that the same construction also proves hardness for c-Prop when $c > 2/3$.

We will use Lemma 6.5.4, which gives a standardized way to prove deterministic lower bounds for EF for two players. Recall that a list of allocations $\mathcal{T} = (T_1, T_2 ...)$ is defined where each $T_j = (T_{j1}, T_{j2})$ is an allocation giving each player k items. Also, for every such allocation A, exactly one of A and $\overline{A}$ appears in $\mathcal{T}$. All that is needed to complete the reduction is to show how to construct valuations v_1, v_2 such that the following conditions are satisfied:

Condition 6.5.1. *For all $|S| < k$ and both i, $v_i(S) < c \cdot v_i(M \backslash S)$.*

Condition 6.5.2. *Whenever $y_{ij} = 1$, $v_i(T_{j\bar{i}}) < c \cdot v_i(T_{ji})$.*

Condition 6.5.3. *Whenever $y_{ij} = 0$, $v_i(T_{ji}) < c \cdot v_i(T_{j\bar{i}})$.*

Theorem 6.8.5. *For two players with subadditive valuations and any $c > 1/2$, any deterministic protocol which determines whether a c-EF allocation exists requires an exponential amount of communication. Formally,*

$$D_{subadd}(2, 2k, EF, c) \geq \frac{1}{2}\binom{2k}{k}$$

for any $c > 1/2$.

Proof. Given bit strings of length $\frac{1}{2}\binom{2k}{k}$ for some integer k, define $M, N, (y_1, y_2)$, and $\mathcal{T}$ as in Lemma 6.5.4. We need only to construct subadditive valuations v_1, v_2 such that Conditions 6.5.1, 6.5.2, and 6.5.3 are met. We define each v_i by

$$v_i(S) = \begin{cases} 0 & \text{if} \quad |S| = 0 \\ 1 & \text{if} \quad 0 < |S| < k \\ 2 & \text{if} \quad k < |S| < 2k \\ 3 & \text{if} \quad |S| = 2k \\ 2 & \text{if} \quad \exists j \quad S = T_{ji} \quad \text{where} \quad y_{ij} = 1 \\ 2 & \text{if} \quad \exists j \quad S = T_{j\bar{i}} \quad \text{where} \quad y_{ij} = 0 \\ 1 & \text{if} \quad \exists j \quad S = T_{ji} \quad \text{where} \quad y_{ij} = 0 \\ 1 & \text{if} \quad \exists j \quad S = T_{j\bar{i}} \quad \text{where} \quad y_{ij} = 1 \end{cases}$$

When $|S| = k$, S falls under exactly one of the last four cases in the definition of v_i.

If $|S| < k$, we have $|M \backslash S| > k$, so $v_i(S) \leq 1$ and $v_i(M \backslash S) \geq 2$. Thus for any $c > 1/2$, $v_i(S) < c \cdot v_i(M \backslash S)$, so Condition 6.5.1 is met. Suppose $y_{ij} = 1$ for some i, j: then $v_i(T_{j\bar{i}}) = 1$ and $v_i(T_{ji}) = 2$, so again $v_i(S) < c \cdot v_i(M \backslash S)$ for any $c > 1/2$. Suppose $y_{ij} = 0$ for some i, j: then similarly, $v_i(T_{ji}) = 1 < c \cdot 2 = c \cdot v_i(T_{j\bar{i}})$ for any $c > 1/2$. Thus Condition 6.5.3 is satisfied as well.

It remains to show that v_i is subadditive for both i. Specifically, we need to show that for any S and T, $v_i(S) + v_i(T) \geq v_i(S \cup T)$. If either $S = \emptyset$ or $T = \emptyset$, this trivially holds, so suppose $|S| > 0$ and $|T| > 0$. We proceed by case analysis.

Case 1: $|S \cup T| < 2k$. Then $v_i(S \cup T) \leq 2$. Since $|S| > 0$ and $|T| > 0$, we have

$$v_i(S) + v_i(T) \geq 1 + 1 \geq 2 \geq v_i(S \cup T)$$

Case 2: $|S \cup T| = 2k$. Then $v_i(S \cup T) = 3$. Since $v_i(S) \geq 1$ and $v_i(T) \geq 1$, it remains to show that at least one of $v_i(S) \geq 2$ and $v_i(T) \geq 2$ is true. Since $S \cup T = M$ in this case, we have $M \backslash S \subseteq T$.

Observe that under these valuations, for any allocation A where $v_i(A_i) \leq 1$, we have $v_i(A_{\bar{i}}) \geq 2$. Thus if $v_i(S) \leq 1$, then $v_i(M \backslash S) \geq 2$, so $v_i(T) \geq 2$. Since v_i only takes on integer values in this proof, if $v_i(S) > 1$, we have $v_i(S) \geq 2$. Thus at least one of $v_i(S) \geq 2$ and $v_i(T) \geq 2$ is true, so the claim is satisfied in this case. Thus v_i is subadditive for both i. □

To prove hardness for proportionality, we reduce from envy-freeness.

Theorem 6.8.6. *For two players with subadditive valuations and any $c > 2/3$, any deterministic protocol which determines whether a c-Prop allocation exists requires an exponential amount of communication. Formally,*

$$D_{subadd}(2, 2k, Prop, c) \geq \frac{1}{2}\binom{2k}{k}$$

for any $c > 2/3$.

Proof. We reduce from Fair Division for $P = \text{EF}$. Given an input (x_1, x_2), we define v_i as in the proof of Theorem 6.8.5. By Theorem 6.8.5, for any $c' > 1/2$, at least $\frac{1}{2}\binom{2k}{k}$ communication is required to determine whether a c'-EF allocation exists under these valuations. We will show that under these valuations, for any $c > 2/3$ and any $c' > 1/2$, an allocation A is c-Prop if and only if is c'-EF: thus the lower bound of Theorem 6.8.5 will apply to c-Prop for $c > 1/2$ as well.[16]

Suppose an allocation A is c'-EF for some $c' > 1/2$: then $v_i(A_i) \geq c' v_i(A_{\bar{i}}) > \frac{1}{2} v_i(A_{\bar{i}})$. Under these valuations, for any allocation A where $v_i(A_i) \leq 1$, we have $v_i(A_{\bar{i}}) \geq 2$. Thus $v_i(A_i)$ must be strictly greater than 1. Since these valuations only take on integer values, this implies that $v_i(A_i) \geq 2$ for both i. Therefore

$$v_i(A_i) \geq 2 \geq \frac{3}{2} = \frac{1}{2} v_i(M) \geq \frac{c}{2} v_i(M)$$

for every $c > 2/3$, so A is c-Prop for every $c > 2/3$.

Now suppose that A is c-Prop for some $c > 2/3$: then $v_i(A_i) \geq \frac{c}{2} v_i(M) = \frac{3c}{2} > 1$ for both i. Thus we again have $v_i(A_i) \geq 2$ for both i, since these valuations only take on integer values. This also implies that $|A_i| > 0$ for both i, which means that $|A_i| < 2k$ for both i. Therefore $v_i(A_{\bar{i}}) \leq 2$ for both i, so we have

$$v_i(A_i) \geq 2 \geq v_i(A_{\bar{i}}) \geq c' v_i(A_{\bar{i}})$$

for any $c' > 1/2$. Therefore A is c'-EF for every $c' > 1/2$.

□

Theorems 6.8.5 and 6.8.6 resolve the deterministic subadditive case. We now move on to general valuations, and give the last few results we need to complete Table 6.1.

[16]It is actually sufficient to show that for any $c > 2/3$, there exists such a $c' > 1/2$, but we prove that this holds for any $c' > 1/2$.

6.9 General valuations

This section covers the remaining settings for envy-freeness and proportionality. In Section 6.9.1, we show that c-Prop is hard for general valuations for any $c > 0$, in both the randomized and deterministic settings. Section 6.9.2 gives a similar lower bound for c-EF for any $c > 0$, but only for deterministic protocols. In Section 6.9.3, we show that there actually exists an efficient randomized protocol for c-EF for any $c \in [0, 1]$. We also show that this protocol works for proportionality in the subadditive case, again for any $c \in [0, 1]$. These results conclude our study of envy-freeness and proportionality.

6.9.1 Proportionality randomized lower bound

Recall that Disjointness on bit strings of length ℓ has randomized communication complexity $\Omega(\ell)$. (Lemma 6.5.1).

Theorem 6.9.1. *For two players with general valuations and any $c > 0$, any randomized protocol which determines whether a c-Prop allocation exists requires an exponential amount of communication. Specifically*

$$R(2, 2k, Prop, c) \in \Omega\left(\binom{2k}{k}\right)$$

for any $c > 0$.

Proof. We reduce from Disjointness. Given bit strings x_1 and x_2 of length $\binom{2k}{k}$, we construct an instance of Fair Division as follows. Let $N = [2]$ be the set of players, and let $M = [2k]$ be the set of items. Let $\mathcal{T} = (T_1, T_2 ... T_{|\mathcal{T}|})$ be an arbitrary ordering of all of the allocations which give each player k items: for any allocation $A = (A_1, A_2)$ where $|A_1| = |A_2| = k$, there exists j where $A_i = T_{ji}$ for both i. Both A and $\overline{A}$ appear in $\mathcal{T}$. Note that $|\mathcal{T}| = \binom{2k}{k}$. Each player i's valuation is defined by

$$v_i(S) = \begin{cases} 0 & \text{if } |S| < k \\ 1 & \text{if } |S| > k \\ 1 & \text{if } \exists j \;\; S = T_{ji} \;\; \text{where} \;\; x_{ij} = 1 \\ 0 & \text{if } \exists j \;\; S = T_{ji} \;\; \text{where} \;\; x_{ij} = 0 \end{cases}$$

Exactly one of the last two cases occur when $|S| = k$, and any such j is unique.

Suppose that (x_1, x_2) is a no-instance of Disjointness: then there exists j where $x_{1j} = x_{2j} = 1$. Consider the allocation $T_j = (T_{j1}, T_{j2})$. Then for both i, $v_i(T_{ji}) = 1 = v_i(M) \geq \frac{c}{2} \cdot v_i(M)$, so the allocation T_j satisfies c-Prop.

Suppose that (x_1, x_2) is a yes-instance of Disjointness: then for every j, there exists i where $x_{ij} = 0$. Suppose that a c-Prop allocation $A = (A_1, A_2)$ exists: then $v_i(A_i) \geq \frac{c}{2} \cdot v_i(M) > 0$ for both i. Suppose a player i receives strictly more than k items in A_i: then the other player receives strictly fewer than k items, and has value zero, which is impossible. Thus $|A_1| = |A_2| = k$. Since $\mathcal{T}$ contains all of the allocations which give each player k items, there must exist j where $A_i = T_{ji}$

for both i. But that implies that $x_{1j} = x_{2j} = 1$, which is a contradiction. Therefore no allocation is c-Prop. □

This lower bound is actually much more general than just c-Prop. It holds for any imaginable fairness property (not just c-EF or c-Prop) where (1) player i is always unhappy if $v_i(A_i) = 0$, even if $v_i(A_{\bar{i}})$ is also 0, and (2) player i is always happy if $v_i(A_i) = v_i(M)$. Both c-EF and c-Prop satisfy the first condition. The second condition is satisfied by c-Prop for any c, but c-EF violates this for every c: player i is always happy if $v_i(A_i) = v_i(A_{\bar{i}}) = 0$. We will see in Section 6.9.3 that this leads to an efficient randomized protocol for c-EF, for any $c \in [0, 1]$.

6.9.2 Envy-freeness deterministic lower bound

In this section we prove that for general valuations, c-EF is hard in the deterministic setting for any $c > 0$. We will use Lemma 6.5.4; recall that we need only show how to construct valuations that satisfy the following conditions:

Condition 6.5.1. *For all $|S| < k$ and both i, $v_i(S) < c \cdot v_i(M \backslash S)$.*

Condition 6.5.2. *Whenever $y_{ij} = 1$, $v_i(T_{j\bar{i}}) < c \cdot v_i(T_{ji})$.*

Condition 6.5.3. *Whenever $y_{ij} = 0$, $v_i(T_{ji}) < c \cdot v_i(T_{j\bar{i}})$.*

Theorem 6.9.2. *For two players with general valuations and any $c > 0$, any deterministic protocol which determines whether a c-EF allocation exists requires an exponential amount of communication. Specifically,*

$$D(2, 2k, EF, c) \geq \frac{1}{2}\binom{2k}{k}$$

for any $c > 0$.

Proof. We use Lemma 6.5.4. Given bit strings of length $\frac{1}{2}\binom{2k}{k}$ for some integer k, define $M, N, (y_1, y_2)$, and $\mathcal{T}$ as in Lemma 6.5.4. We need only to construct valuations v_1, v_2 such that Conditions 6.5.1, 6.5.2, and 6.5.3 are met. We define each v_i by

$$v_i(S) = \begin{cases} 0 & \text{if } |S| < k \\ 1 & \text{if } |S| > k \\ 1 & \text{if } \exists j \;\; S = T_{ji} \;\; \text{where} \;\; y_{ij} = 1 \\ 1 & \text{if } \exists j \;\; S = T_{j\bar{i}} \;\; \text{where} \;\; y_{ij} = 0 \\ 0 & \text{if } \exists j \;\; S = T_{ji} \;\; \text{where} \;\; y_{ij} = 0 \\ 0 & \text{if } \exists j \;\; S = T_{j\bar{i}} \;\; \text{where} \;\; y_{ij} = 1 \end{cases}$$

Recall that for every allocation A which gives each player k items, $\mathcal{T}$ (as defined by Lemma 6.5.4) contains exactly one of A and $\overline{A}$. Thus if $|S| = k$, S falls under exactly one of the last four cases in the definition of v_i.

If $|S| < k$, we have $|M \backslash S| > k$, so $v_i(S) = 0 < c = c \cdot v_i(M \backslash S)$. This satisfies Condition 6.5.1. Suppose $y_{ij} = 1$ for some i, j: then $v_i(T_{\bar{j}\bar{i}}) = 0 < c = c \cdot v_i(T_{ji})$, so Condition 6.5.2 is satisfied. Suppose $y_{ij} = 0$ for some i, j: then similarly, $v_i(T_{ji}) = 0 < c = c \cdot v_i(T_{\bar{j}\bar{i}})$. Thus Condition 6.5.3 is satisfied as well. □

6.9.3 A randomized upper bound

Although c-EF is hard for general valuations in the deterministic setting, it admits an efficient randomized protocol for any $c \leq 1$. Fundamentally, this is because the randomized communication complexity of EQUALITY is constant, while its deterministic complexity is the length of the string. Our deterministic lower bound in Section 6.9.2 was based on a reduction from EQUALITY: in this section, we reduce *to* EQUALITY.

Our protocol will actually be much more general than just c-EF. For example, it will also work for c-Prop for subadditive valuations, for any $c \in [0, 1]$. More generally, it solves FAIR DIVISION with two players when (c, P) satisfies two conditions:

Condition 6.3.1. *For every allocation A, each player is happy with at least one of A and $\overline{A}$.*

Condition 6.9.1. *Whether player i is happy does not depend on any valuation other than v_i.*

All of the fairness properties we consider satisfy Condition 6.9.1. The c-EF property satisfies Condition 6.3.1 for any $c \leq 1$. As mentioned before, c-Prop satisfies this for any $c \leq 1$ for subadditive valuations.

Despite being hard in the deterministic setting, EQUALITY admits an efficient randomized protocol (Lemma 6.5.3), as described in Section 6.5. This protocol (let us call it Γ_{EQ}) enables the FAIR DIVISION randomized protocol that we present in this section.

The standard EQUALITY problem is a decision problem, but FAIR DIVISION is a search problem: we must output a satisfactory allocation if one exists. The search version of EQUALITY is to determine whether two bit strings are equal, and if they are not, to return an index where they differ.

Lemma 6.9.1 ([133])**.** *There exists a randomized protocol which solves the search version of* EQUALITY *for two players and has communication cost $O(\log \ell)$, where ℓ is the length of the bit strings.*

The protocol uses a binary search approach. The players first use Γ_{EQ} to check if their strings are equal. If so, the protocol terminates. If not, the players split their strings into a left half and a right half. They again use Γ_{EQ} to check if their left halves are equal: if they are not, the players recurse on the left half, otherwise they recurse on the right half. This process continues until players isolate a single bit which differs in their bit strings[17].

Since Γ_{EQ} is a randomized protocol, it may return an incorrect answer with probability up to 1/3 (say) each time it is run. If we use Γ_{EQ} many times, as required by the above binary search

[17] The protocol described in [133] is actually slightly stronger: they find the most significant bit where the two strings differ. This is because they have a slightly different goal in that paper, for which finding any bit that differs is not sufficient.

argument, the probability Γ_{EQ} returns the correct answer every time may be less than 2/3, which is unacceptable. This makes the protocol a sort of "noisy binary" search. This can be done with total communication $O(\log \ell \log \log \ell)$ using a standard Chernoff bound argument, but [77] shows how this can be done with total communication just $O(\log \ell)$. We refer to protocol from Lemma 6.9.1 as Γ_{EQS}.

We now present our randomized protocol (Protocol 9). Let $\mathcal{T} = (T_1, T_2 ...)$ be a list of every possible allocation (not just those with bundles of a fixed size) in an arbitrary order. The ordering $\mathcal{T}$ will be publicly agreed upon ahead of time; note that this is not a "cheat" in the sense that our lower bounds still apply even if players publicly agree on an ordering of possible allocations. Condition 6.9.1 is necessary for Protocol 9 to be well-defined (step 1 in particular), but will not appear in the proof of Theorem 6.9.3.

The protocol uses a similar construction to the previous lower bounds in that players have exponentially long bit strings, with each index representing a possible allocation, and where $y_{ij} = 1$ if player i is happy with T_j. Similarly to the Equality lower bounds, an index where $y_{1j} = y_{2j}$ implies the existence of a c-P allocation: if $y_{1j} = y_{2j} = 1$, both players are happy with that allocation, and if $y_{1j} = y_{2j} = 0$, both players are happy with the reverse allocation by Condition 6.3.1. This is made formal by the following theorem:

Protocol 9 Randomized protocol for two players to either find an P allocation or determines that none exists, assuming P satisfies Conditions 6.3.1 and 6.9.1.

Private inputs: v_1, v_2
Public inputs: $P, c, \mathcal{T}$

1. Each player i constructs a bit string y_i as follows: for all j where player i is happy with T_j, player i sets $y_{ij} = 1$. For all j where player i is unhappy with T_j, player i sets $y_{ij} = 0$.
2. Player 1 sets $x_1 = y_1$ and player 2 sets $x_2 = \overline{y_2}$.
3. The players run Γ_{EQS} on (x_1, x_2), which either returns an index j where $x_{1j} \neq x_{2j}$, or determines that the two strings are equal.
4. If the two bit strings are equal, the players declare that no c-P allocation exists.
5. If an index j is returned where $x_{1j} = 1$ and $x_{2j} = 0$, the players declare that T_j is a c-P allocation.
6. If an index j is returned where $x_{1j} = 0$ and $x_{2j} = 1$, the players declare that $\overline{T}_j$ is a c-P allocation.

Theorem 6.9.3. *If (c, P) satisfies Conditions 6.3.1 and 6.9.1, then Procotol 9 either finds a c-P allocation or shows that none exists, and uses communication $O(m)$.*

Proof. Suppose Protocol 9 declares that no c-P allocation exists: then $x_{1j} = x_{2j}$ for all j. This implies that $y_{1j} \neq y_{2j}$ for all j. Therefore whenever player 1 is happy with T_j, player 2 is unhappy with T_j, so no c-P allocation exists.

Suppose Protocol 9 returns an index j where $x_{1j} \neq x_{2j}$. If $x_{1j} = 1$ and $x_{2j} = 0$, then $y_{1j} = y_{2j} = 1$. Thus both players are happy with T_j, so T_j is c-P. If $x_{1j} = 0$ and $x_{2j} = 1$, then $y_{1j} = y_{2j} = 0$, so neither player is happy with T_j. Then by Condition 6.3.1, both players are happy with $\overline{T_j}$, so $\overline{T}_j$ is c-P. Therefore Protocol 9 correctly finds a c-P allocation or determines that none exist.

Since the total number of allocations is $O(2^m)$ when $n = 2$, x_1 and x_2 have length $O(2^m)$. Thus Γ_{EQS} has communication cost $O\big(\log(2^m)\big) = O(m)$. Since all other steps require no communication, Protocol 9 uses communication $O(m)$. □

Theorem 6.9.3 immediately implies the following two theorems:

Theorem 6.9.4. *For any $c \in [0, 1]$, Protocol 9 finds an c-EF allocation or shows that none exists, and has communication cost $O(m)$. Formally,*

$$R(2, m, EF, c) \in O(m)$$

Theorem 6.9.5. *For subadditive valuations and any $c \in [0, 1]$, Protocol 9 finds an c-Prop allocation or shows that none exists, and has communication cost $O(m)$. Formally,*

$$R_{subadd}(2, m, Prop, c) \in O(m)$$

Since $R_{submod}(n, m, P, c) \leq R_{subadd}(n, m, P, c) \leq R(n, m, P, c)$, this settles the randomized communication complexities for all settings with two players. The reader can verify that Table 6.1 is now complete.

6.10 Maximin share

Finally, we consider a different fairness property: maximin share. A player's maximin share (MMS) is the maximum value as she could guarantee herself if she gets to divide the items into n bundles, but chooses last. An allocation A is c-MMS for $c \in [0, 1]$ if each player receives at least a c-fraction of her MMS. We use "MMS" to refer to both each player's max-min share and the fairness property itself. Formally,

Definition 6.10.1. *An allocation A is c-MMS if for every player i,*

$$v_i(A_i) \geq \max_{A'=(A'_1 \ldots A'_n)} \min_{j \in [n]} v_i(A'_j)$$

where A' ranges over all possible allocations.

In this section, we prove exponential lower bounds for MMS in two settings: for general valuations and any $c > 0$, and for submodular valuations when $c = 1$. Both lower bounds hold even for two players, and for randomized protocols. Both lower bounds will rely on reductions from Disjointness.

6.10.1 Lower bound for general valuations and any $c > 0$

In this section we show that for general valuations, c-MMS is hard for any $c > 0$, even for randomized protocols and even if there are only two players. We will reduce from 1-Prop, which we know to be hard in this setting (randomized, $n = 2$, general valuations) from Theorem 6.9.1. We say that an allocation A is over a set of items M to mean that $A_1 \cup A_2 = M$. Also, we say that an allocation A is c-Prop for valuations v_1, v_2 if $v_i(A_i) \geq \frac{c}{2} v_i(M)$ for both i. Since we will be reducing between two different FAIR DIVISION instances, we will be dealing with allocations over different sets of items and different sets of valuations.

Theorem 6.10.1. *For two players with general valuations and any $c > 0$, any randomized protocol which determines whether a c-MMS allocation exists requires an exponential amount of communication. Specifically,*

$$R(2, 2k+4, MMS, c) \in \Omega\left(\binom{2k}{k}\right)$$

for any $c > 0$.

Proof. Consider an arbitrary instance of FAIR DIVISION for two players with general valuations v_1, v_2, any $c > 0$, and some set of items M. We want to know whether there exists an allocation A over M which is 1-Prop for v_1, v_2. Let $\alpha_i = \frac{1}{2c} v_i(M)$: then A is 1-Prop if and only if $v_i(A_i) \geq c\alpha_i$ for both i.

We will create a second instance of FAIR DIVISION as follows. Add four items g_1, g_2, g_3, g_4, let $X = \{g_1, g_2, g_3, g_4\}$, and define $M' = M \cup X$. Let $Y_1 = \{g_1, g_2\}, Y_2 = \{g_3, g_4\}, Z_1 = \{g_1, g_3\}$, and $Z_2 = \{g_2, g_4\}$. The set of players is the same. Define the following valuations v_1' and v_2' over M':

$$v_1'(S) = \begin{cases} \alpha_1 & \text{if } \ Y_1 \subseteq S \ \text{ or } \ Y_2 \subseteq S \\ \min(v_1(S \backslash X), c\alpha_1) & \text{if } \ \{g_1, g_4\} \subseteq S \ \text{ and } \ g_2, g_3 \notin S \\ 0 & \text{otherwise} \end{cases}$$

$$v_2'(S) = \begin{cases} \alpha_2 & \text{if } \ Z_1 \subseteq S \ \text{ or } \ Z_2 \subseteq S \\ \min(v_2(S \backslash X), c\alpha_2) & \text{if } \ \{g_2, g_3\} \subseteq S \ \text{ and } \ g_1, g_4 \notin S \\ 0 & \text{otherwise} \end{cases}$$

We first claim that each player i's MMS is exactly α_i. Since $c \leq 1$, we have $v_i'(A_i') \leq \alpha_i$ for all i and for every allocation A' over M'. Thus each player's MMS is at most α_i. In the partition $(Y_1, Y_2 \cup M)$, player 1 has value α_1 for both bundles, so player 1's MMS is at least α_1. Similarly, player 2 has value α_2 for both bundles in the partition $(Z_1, Z_2 \cup M)$. Thus each player i's MMS is exactly α_i.

Therefore an allocation A' over M' is c-MMS for v_1', v_2' if and only if $v_i'(A_i') \geq c\alpha_i$ for both i. It remains to show that there exists such an allocation A' over M' if and only if an there exists a 1-Prop allocation for v_1, v_2 over M.

Suppose A is 1-Prop allocation over M for v_1, v_2: then $v_i(A_i) \geq c\alpha_i$ for both i. Let $A' =$

$(A_1 \cup X, A_2)$: then $v'_i(A'_i) \geq v_i(A_i) \geq c\alpha_i$, so A' is c-MMS for v'_1, v'_2 over M'.

Now suppose A' is a c-MMS allocation for v'_1, v'_2 over M'. Since $c > 0$, we have $v'_i(A'_i) \geq c\alpha_i > 0$ for all i. For all j and j', we have $Y_j \cap Z_{j'} \neq \emptyset$. Thus if player 1 receives Y_1 or Y_2: then player 2 player 2 cannot receive Z_1 or Z_2. Furthermore, player 2 cannot receives $\{g_2, g_3\}$, so $v'_2(A'_2) = 0$, which is a contradiction. Therefore player 1 cannot receive either Y_1 or Y_2. Similarly, if player 2 receives Z_1 and Z_2, player 1 will have value 0. Thus player 2 does not receive Z_1 or Z_2.

Therefore $v_1(A_1) = 0$ unless $\{g_1, g_4\} \subseteq A_1$, and $v_2(A_2) = 0$ unless $\{g_2, g_3\} \subseteq A_2$. Therefore $\{g_1, g_4\} \subseteq A_1$ and $\{g_2, g_3\} \subseteq A_2$. Thus $v'_i(A'_i) = \min(v_i(A'_i \backslash X), c\alpha_i)$ for both i. Since $v'_i(A'_i) \geq c\alpha_i$ for all i, we have $v_i(A'_i \backslash X) \geq c\alpha_i$ for both i.

Define an allocation A where $A_i = A'_i \backslash X$. It is clear that A is an allocation over M. Then $v_i(A_i) \geq c\alpha_i$ for both i, so A is a 1-Prop allocation for v_1, v_2 over M.

Therefore there exists a c-MMS allocation for v'_1, v'_2 over M' if and only if there exists 1-Prop allocation over for v_1, v_2 over M. This completes the reduction, and shows that for any $c > 0$ and any number of items m,

$$R(2, m+4, \text{MMS}, c) \geq R(2, m, \text{Prop}, 1)$$

Therefore by Theorem 6.9.1, we have $R(2, 2k+4, \text{MMS}, c) \in \Omega\left(\binom{2k}{k}\right)$. □

6.10.2 Lower bound for submodular valuations and $c = 1$

We now show that even for two players with submodular valuations, 1-MMS is hard. This does not hold for c-MMS for any c: in fact, a $\frac{1}{3}$-MMS is guaranteed to exist for submodular valuations [92].

Theorem 6.10.2. *For two players with submodular valuations, any randomized protocol which determines whether a 1-MMS allocation exists requires an exponential amount of communication. Specifically*

$$R(2, 2k, MMS, 1) \in \Omega\left(\binom{2k}{k}\right)$$

Proof. We reduce from Disjointness. Given bit strings x_1 and x_2 of length $\binom{2k}{k}$, we construct an instance of Fair Division as follows. Let $N = [2]$ be the set of players, and let $M = [2k]$ be the set of items. We define $Y_1 = \{1...k\}, Z_1 = \{k+1...2k\}, Y_2 = \{2...k+1\}$, and $Z_2 = \{1\} \cup \{k+2...2k\}$.

Let $\mathcal{T} = (T_1, T_2...T_{|\mathcal{T}|})$ be an arbitrary ordering of all of the allocations which give each player k items: for any allocation $A = (A_1, A_2)$ where $|A_1| = |A_2| = k$, there exists j where $A_i = T_{ji}$ for both i. Note that $|\mathcal{T}| = \binom{2k}{k}$. One exception: none of $(Y_1, Z_1), (Z_1, Y_1), (Y_2, Z_2)$, or (Z_2, Y_2) appear in $\mathcal{T}$.

Player i's valuation is given by:

$$v_i(S) = \begin{cases} 3|S| & \text{if } \ |S| < k \\ 3k & \text{if } \ |S| > k \\ 3k & \text{if } \ S = Y_i \ \text{ or } \ S = Z_i \\ 3k-1 & \text{if } \ S = Y_{\bar{i}} \ \text{ or } \ S = Z_{\bar{i}} \\ 3k & \text{if } \ \exists j \ \ S = T_{ji} \ \text{ where } \ x_{ij} = 1 \\ 3k-1 & \text{if } \ \exists j \ \ S = T_{ji} \ \text{ where } \ x_{ij} = 0 \end{cases}$$

These valuations are submodular by the same argument as in the proof of Theorem 6.6.1. Observe that when $|S| = k$, exactly one of the last four cases occur.

Since $v_i(S) \leq 3k$ for all S, player i's MMS is at most $3k$. For both i, (Y_i, Z_i) is a valid allocation. Furthermore, player i has value $3k$ for both bundles in that allocation. Thus each player i's MMS is at least $3k$, so both players' MMS are exactly $3k$.

Suppose that (x_1, x_2) is a no-instance of Disjointness: then there exists j where $x_{1j} = x_{2j} = 1$. Consider the allocation $T_j = (T_{j1}, T_{j2})$. Then for both i, $v_i(T_{ji}) = 3k$, so the allocation T_j satisfies 1-MMS.

Suppose that (x_1, x_2) is a yes-instance of Disjointness: then for every j, there exists i where $x_{ij} = 0$. Suppose a 1-MMS allocation A exists. We first claim that for both i, $A \neq (Y_i, Z_i)$ and $A \neq (Z_i, Y_i)$ for both i. This is because player $\bar{i}$ will have value $3k - 1$, which is less than her MMS. Suppose there is a player i where $|A_i| < k$: then $v_i(A_i) < 3k$, which is impossible. Thus $|A_1| = |A_2| = k$.

Therefore there exists j where $A = T_j$. But since (x_1, x_2) is a yes-instance of Disjointness, there exists i where $x_{ij} = 0$, so $v_i(T_{ji}) = 3k - 1$. This is a contradiction, so no 1-MMS allocation exists. ☐

6.11 Conclusion

In this chapter, we proposed a simple model for the communication complexity of fair division, and solved it completely, for every combination of five parameters: number of players, valuation class, fairness property P, constant c, and deterministic vs. randomized complexity.

More broadly, communication complexity is an example of topic that has been well-studied in algorithmic game theory but not in fair division, despite having a natural fair division analog. We wonder if there are other such topics.

Part III

Public Resource Allocation

Chapter 7

Markets for public decision-making

We now pivot to public resource allocation, where the group makes a single decision that affects all agents, rather than allocating a separate bundle to each agent. In this chapter, we focus on the case where the resources are a set of issues, each with two possible *alternatives*. Each agent has a preferred alternative for each issue, as well as a utility function describing their relative values between issues. We call this a *public decision-making* problem.

We study adaptations of market economies to this setting: issues have prices, and each agent is endowed with artificial currency that she can use to purchase probability for her preferred alternatives (we allow randomized outcomes). We first show that when each issue has a single price that is common to all agents, market equilibria can be arbitrarily bad.

This negative result motivates a different approach. We present a novel technique called *pairwise issue expansion*, which transforms any public decision-making instance into an equivalent Fisher market, the simplest type of private goods market. This is done by expanding each issue into many goods: one for each pair of agents who disagree on that issue. We show that the equilibrium prices in the constructed Fisher market yield a *pairwise pricing equilibrium* in the original public decision-making problem which maximizes Nash welfare. More broadly, pairwise issue expansion uncovers a powerful connection between the public decision-making and private goods settings; this immediately yields several interesting results about public decisions markets, and furthers the hope that we will be able to find a simple iterative voting protocol that leads to near-optimum decisions.

7.1 Introduction

Fair and transparent public decision-making is a key element of a democratic society, but many public decisions are made by government officials behind closed doors. In this chapter, we investigate mechanisms for large-scale public decision-making where citizens directly vote on a set of issues at the same time, focusing on the case where each issue has exactly two alternatives. In particular, we examine connections to private goods allocation. One can think of each issue that is under consideration as a "good", and public decision-making as "allocating" the good to one of the alternatives. We allow randomized outcomes, where the outcome can put nonzero probability on multiple alternatives: this is analogous to divisible private goods, where a good can be split among multiple agents[1]. The fundamental difference is that in private goods allocation, each agent's utility depends only on the bundle of goods she receives; in public decision-making, the group makes a single decision that affects all agents.

Market economies are one of the longest-studied areas in the distributions of private goods. In our setting, we argue that the fixed-budget model is more appropriate than the quasilinear model. The idea is that rather than using "real money", we will endow each agent with a fixed amount of artificial currency which can only be used within our mechanism. This corresponds to the assumption of a fixed budget, as well as the assumption that agents have no value for leftover money. We will focus on the fixed-budget model with linear prices, often called the Fisher market model [24, 80].

7.1.1 Our contribution

We consider adaptations of markets to the public decision-making setting. Many democracy theorists believe that it is unethical (e.g., see [158]) and many democratic countries stipulate that it is illegal to allow citizens to purchase political influence with actual money. Instead, we think of each agent being endowed with the same amount of "artificial currency" that is useful only for voting on these issues; thus our approach to public decision markets is consistent with the spirit of "one person one vote". Prices are assigned to issues, and agents can use their artificial currency to "purchase" probability for their preferred alternatives on the issues they most value[2].

Markets have the desirable property that each agent can choose how to allocate her money across goods, based on their relative values to her. In the context of large-scale public decision-making, this allows agents to express their relative weights for the different issues in a fine-grained way. This is in contrast to approaches like asking agents to rank the issues by importance, which are more limited in expressiveness. Markets have the additional property that the equilibria are "supported" by prices: prices provide a sort of certificate of fairness, in that each agent can verify that she is spending her budget in the best way possible.

The simplest pricing model assigns a single (linear) price to each issue, and all agents are subject to the same set of prices. We refer to this as "per-issue pricing", or just "issue pricing". In the private goods setting, per-good pricing is sufficient to yield a market equilibrium with optimal *Nash welfare*: the product of agent utilities[3]. Unfortunately, we show in Section 7.3 that issue pricing in the public decisions setting can result in very poor equilibria: the Nash welfare of the equilibrium may be a factor of $\Omega(n)$ worse than optimal, where n is the number of agents. The same instance shows that the utilitarian welfare (the sum of agent utilities) and egalitarian welfare (the minimum agent utility) may both be a factor of $\Omega(n)$ worse than optimal as well.

[1] An alternative interpretation is that the issues themselves are divisible: for example, in the case of a city choosing how much money to a particular project, any amount of money is a valid outcome.

[2] This can also be thought of as a private goods market with externalities: each agent's utility depends not only on her own bundle, but also other agents' bundles.

[3] The concept of Nash welfare is due to [129] and [108].

Pairwise issue expansion

This negative result motivates a more complex market model. Our main contribution is a reduction which transforms any public decision-making instance into a private goods Fisher market instance that is "equivalent" in a strong sense. For each issue, we construct a good for each pair of agents who disagree on that issue. The outcome on that issue can be thought of as the result of pairwise negotiations between each pair of agents who disagree. We refer to this reduction as *pairwise issue expansion*. The equilibrium prices of the constructed Fisher market yield a "pairwise pricing equilibrium" in the original public decisions instance. We show that the resulting pairwise pricing equilibrium maximizes Nash welfare in the public decisions instance.

Furthermore, pairwise issue expansion allows us to directly import results for Fisher markets to the public decisions setting. If the utilities in the public decisions instance are in class $\mathcal{H}$ (say, linear utilities), the utilities in the constructed Fisher market will be nested $\mathcal{H}$-Leontief (for example, nested linear-Leontief)[4]. This means that any result which works for Fisher markets with nested $\mathcal{H}$-Leontief utilities can be imported to public decisions instances with utilities in class $\mathcal{H}$. The main Fisher market results we consider are: (1) a strongly polynomial-time algorithm for finding a Fisher market equilibrium with two agents and any utility functions [43], (2) a strongly polynomial-time algorithm for finding a Fisher market equilibrium for Leontief utilities with weights in $\{0, 1\}$ [87], (3) a polynomial-time algorithm for a Fisher market with Leontief utilities which yields a $O(\log n)$ approximation simultaneously for all canonical welfare functions (i.e. Nash welfare, utilitarian welfare, egalitarian welfare, etc) [95], and (4) a discrete-time tâtonnement process for finding the Fisher market equilibrium for nested CES-Leontief utilities that converges in polynomial-time [9]. Thus pairwise issue expansion yields the following results for the public decision-making setting:

1. A strongly polynomial-time algorithm for finding a public decisions market equilibrium with two agents and any utility functions.

2. A strongly polynomial-time algorithm for finding a public decisions market equilibrium for Leontief utilities with weights in $\{0, 1\}$.

3. A polynomial-time algorithm for a public decisions instance with Leontief utilities which yields a $O(\log n)$ approximation simultaneously for all canonical welfare functions (i.e. Nash welfare, utilitarian welfare, egalitarian welfare, etc).

4. A discrete-time tâtonnement process for finding a public decisions market equilibrium for CES utilities that converges in polynomial-time.

These Fisher market results yield the analogous results for public decisions instances for two agents with any utilities, Leontief utilities with weights in $\{0, 1\}$[5], any Leontief utilities, and CES utilities, respectively. We also discuss public decisions tâtonnement in more depth, and show how

[4] These utility classes will be defined and discussed later.

[5] A nested Leontief-Leontief function is still a Leontief function. Incidentally, this also implies that for a public decision making problem where agent utilities are Leontief, we get a Fisher market which has exactly the same form, i.e. with Leontief utilities.

our reduction can be used to implement a tâtonnement process where agents only interact with the public decisions instance, and never see the constructed Fisher market.

More broadly, our work uncovers a powerful connection between private goods allocation and public decision-making. We hope that pairwise issue expansion will have applications in future work as well. One particularly promising direction is to design an iterative local voting scheme akin to prediction markets (see [47]), where agents (or pairs of agents) arrive sequentially and move the current decision vector in their preferred direction subject to offered issue prices. Our proof of the existence of a simple tâtonnement for public decision markets offers hope that such a scheme may be possible.

7.1.2 Related work

We have already discussed in depth the foundational works regarding private goods markets (Section 1.5). Thus we presently focus only on public resource allocation and public goods markets.

Foley's work on Lindahl Equilibria

The market concept most directly relevant to our public decision markets is that of Lindahl equilibria, developed by Foley [82], who showed that personalized prices (i.e., each agent may be assigned a different price for each good with no restrictions) can support any Pareto optimal solution in the context of public goods[6]. Our work can be thought of as improving upon Foley's work to get much stronger properties for the special case of public decision-making. We obtain these stronger properties using a more sophisticated reduction, one which is in fact weaker in the sense that there is a correspondence between the public decisions market and the private market *only at equilibrium*. Our reduction explicitly relies on the fact that agents are in opposition on each issue in the public decisions setting, which is not the case in the public goods setting.

The stronger properties we obtain are as follows. First, Foley's work [82] allows arbitrary personalized prices, whereas we only require pairwise prices: for each issue, there is a price for each pair of agents who disagree on that issue. Our Fisher market can be thought of as negotiating independently with each person that disagrees with you through a normal market; we are not aware of any such simple interpretation that follows from Foley's very general work. Second, in our private goods reduction, a feasible public goods decision (where each agent shares the same societal decision) emerges naturally: we leverage properties of nested-Leontief utilities and the Nash welfare objective function to implicitly represent the feasibility constraints, which allows us to obtain the correspondence *only at equilibrium*. In contrast, Foley adds cone constraints to a private goods market to explicitly enforce the feasibility constraints of the public decision-making problem; these constraints have no natural real-world analogue. Third, we reduce the public goods setting to a Fisher market, arguably the simplest possible and most-studied private goods market. Because of this, our reduction allows us to lift many Fisher market equilibrium results to the public decision-making setting.

[6] In public goods, all agents have nonnegative utility for every good, and the question is how to allocate their money between the goods. In contrast, in public decision-making, agents have opposing preferred alternatives and are in direct competition on each issue. With a careful modification, Foley's work does carry over to the public decision-making setting.

In particular, our reduction allows us to obtain a tâtonnement for public decision-making, even though intermediate steps in the tâtonnement are in a regime where the public and private markets are not in direct correspondence. It is unclear whether this is possible with Foley's construction. We discuss this in technical detail and elaborate on how our work relates to Lindahl equilibria in Section 7.6.

Tâtonnement

As mentioned above, one of our results is a tâtonnement for public decision-making. A tâtonnement is an iterative process which presents agents with a set of prices, asks what they would buy given those prices, and updates the price of each good based on the aggregate demand of each good. A tâtonnement-like process for computing the maximum Nash welfare outcome in participatory budgeting (see e.g. [97] for more on PB) was recently given by Fain et al. [75]. They showed that the maximum Nash welfare outcome can be computed by using a stochastic gradient descent style algorithm. Their algorithm iteratively elicits agents' demands using a process very similar to quadratic voting [115] and updates the current solution accordingly. While this is similar to a tâtonnement, there is one crucial difference. A true tâtonnement (such as the one we present) allows the agents to directly change the current point: the price of each good is updated by a fixed rule based on the aggregate demand of that good. In contrast, the algorithm of [75] moves to a point that is *different* from the one elicited by the quadratic voting. Also, their result also holds only for linear utilities[7].

A tâtonnement, with a similar elicitation scheme, has been shown to work in practice in the participatory budgeting setting [89, 90]. In those works, a new budget is directly elicited from voters, and the mechanism works for ℓ_p normed cost functions. However, their mechanism finds a total welfare maximizing point as opposed to a Nash welfare maximizing outcome. One direction for future work is to adapt the tâtonnement from this work into such an implementable mechanism with a large number of voters.

Inefficiencies of pricing schemes

Another relevant paper from the economics literature is [57], which shows that per-good pricing can lead to inefficiency for public goods. Their examples do not provide bounds on how much worse per-good pricing can be: in contrast, we show that for linear utilities, issue pricing can be a factor of $O(n)$ worse than optimal. Also, we note that it is easy to adapt the examples in Section 3 to show that two other popular market-based approaches, namely Quadratic Voting [115] and Trading Post Prices [164], also do not result in good equilibria with issue pricing in our public decision market setting.

[7] This discussion is thanks to Kamesh Munagala via private correspondence.

Strategic agents

A key property in mechanism design is *strategy-proofness* (or lack thereof). A mechanism is strategy-proof if even a selfish agent would always honestly report her preferences. Most relevant to us is [159], which shows that even for two agents with linear utilities over divisible goods, any mechanism which is both strategy-proof and Pareto optimal[8] is *dictatorial*, meaning that one agent receives all of the resources[9]. Our binary-issue public decisions setting generalizes the two agent private goods setting, and hence we immediately inherit this impossibility result: any mechanism which is both strategy-proof and Pareto optimal is dictatorial. A dictatorial solution is clearly not desirable, and we would like our outcomes to be Pareto optimal, so we assume throughout this chapter that agents honestly report their preferences and do not address the issue of strategic behavior. Other incentive compatibility results for implementation of general classes of social choice functions are discussed in [58].

We note that several works extending Foley relax the assumption that agents report their preferences truthfully, by building voting games in which the equilibrium is one in which truthful reporting is incentive compatible for each agent [99, 172, 107]. Most notably, Groves and Ledyard [99] construct an allocation-taxation scheme – using message passing – for a market with both private goods and public commodities, such that the equilibrium behavior results in a Pareto optimal solution. As in our work, however, their mechanism is still susceptible to a manipulation in which a consumer considers how future prices and the behavior of others are affected by her current decisions.

Other voting schemes, and one person one vote

Other works also propose alternate voting schemes for multiple issues. Storable Votes [41] allows members of a committee to store votes for future meetings so as to spend their votes on issues that matter most to them; the work proves welfare gains in the case of two voters but does not give a principled way to balance the relative importance or cost of different issues, as we do here. In [52], the authors study adaptations of private goods fairness notions (such as proportionality) to the public decision-making context when randomized outcomes are not allowed. In contrast, we allow fractional solutions (i.e. randomized outcomes) and exactly maximize Nash welfare.

Such works, especially this one, may seem to violate the principle of One Person One Vote [83, 101, 32, 5, 91, 109]. In particular, as we propose individual prices, a given issue may "cost" more for one voter than for another. However, as discussed below, these prices are generated in a principled manner – for each issue, there is a single price for each *pair* of voters who disagree on the issue. Furthermore, we note that, at the onset, each voter is allocated the same "budget" through which to vote on issues.

Other works in market equilibria for public goods

Finally, we note that many more strands of literature, too many to detail here, discuss and extend the work of Foley [82] and more generally the idea of equilibria for the funding of public goods. To

[8]An outcome is Pareto optimal if there is no way to improve the utility of any agent without hurting another agent.
[9]A similar result holds for indivisible goods [113].

our knowledge, our public decision-making setting has not been studied as a special case of such public goods markets.

In [22] and [15], the authors ask what happens when the decision to fund a single public good is simply made through a majority vote; in particular, they study under what conditions of voter preferences for the public good and distribution of tax shares of each voter is the funding of the good Pareto optimal. They find that majority vote can fall short of optimal if income is asymmetrically distributed. In Section 7.3, we show that the case with multiple public decisions is far worse: a generalization of majority vote – where each issue has a price – leads to highly suboptimal outcomes, even when everyone is endowed with the same income.

Another strand aims to study the implications of relationships between individuals. For example, in [150], agents are allowed to form coalitions through binding contracts, resulting in inefficiencies. In [74], there are people who can "produce" a given public good and those who "benefit" from that good. These relationships can be represented by a network in a certain way, and the Lindahl outcomes correspond to a solution characterized by the eigenvector centralities of each node. In this work, voters who agree on a given issue end up on the same side of a bipartite graph, resulting in them purchasing the same probability for that issue.

The assumptions and philosophical underpinnings of equilibria theory are also well-studied, as are applications to other fields. Sen [162] challenges the notion that people have consistent preferences that can be elicited. In particular, he posits that people have "commitments" to a particular social group of other people, whose welfare they care about. We note that the assumption of a utility function is nevertheless common, though it is important to be aware of the limitations of such behavioral abstractions. In [153], competitive equilibria is connected to the idea that in capitalism people are given fruits commensurate to their labor, as part of a discussion of the relationship of notions of justice and capitalism. General equilibrium theories are even connected to Structuralism within the philosophy of science [100]. One prominent application of the economics of public goods has been to study environmental (non-)cooperation [44, 121]. Our work extends such applications by connecting market equilibria ideas to voting on different issues in a fair and efficient way, as discussed above.

The rest of the chapter is organized as follows. Section 7.2 formally defines the models of private goods allocation, public decision-making, and Fisher markets. Section 7.3 shows that issue pricing can result in (very) poor equilibria for public decisions markets. Section 7.4 presents the concept of pairwise issue expansion, and shows how this can be used to obtain optimal equilibria, as well as other properties. Section 7.4.3 gives examples of Fisher market results that we can import to the public decisions setting using our reduction. Section 7.5 focuses on a particular such result: tâtonnement. Section 7.6 discusses the connection to [82] and Lindahl equilibria in more depth. Finally, Section 7.7 provides some additional tâtonnement results. Some proofs are deferred to the end of the chapter to avoid interrupting the narrative flow.

7.2 Model

We use the same general resource allocation framework as the previous chapters, defined in Section 1.2. Our focus here is defining clear terminology and notation which accommodates both public and private goods resource allocation. As much as possible, we intentionally use the same notation for the private and public settings, as one of our primary contributions is to highlight the connections between these.

A private goods instance consists of a set of agents $N = [n]$ and a set of goods $M = [m]$; a public decisions instance consists of a set of agents $N = [n]$ and a set of issues $M = [m]$. We will typically use i and k to denote agents, and j and ℓ to denote goods/issues. We assume that issues are binary, meaning that each issue j has two alternatives: 0 and 1. Each agent $i \in N$ has a preferred alternative for each issue j, denoted by a_{ij}, which they truthfully report.

We assume that goods/issues are divisible, meaning that a single good/issue can split among multiple agents. In a public decision instance, divisibility can be interpreted as randomization over alternatives.[10] An outcome of a private goods instance is an allocation $\mathbf{x} \in [0,1]^{m \times n}$. The main constraint on a private goods allocation is that it cannot allocate more than the available supply[11]: $\mathbf{x}$ is valid only if $\sum_{i \in N} x_{ij} \leq 1$ for all $j \in M$. The outcome of a public decisions instance is denoted by $\mathbf{z} = (z^1, \dots, z^m) \in [0,1]^{m \times 2}$, where $z^j = (z^{j,0}, z^{j,1}) \in [0,1]^2$, and $z^{j,a} \in [0,1]$ is the probability put on alternative a for issue j. An outcome $\mathbf{z}$ is valid only if $\sum_{a \in \{0,1\}} z^{j,a} \leq 1$ for all $j \in M$.

7.2.1 Utility functions

In a private goods instance, we use $u_i(\mathbf{x}) \in \mathbb{R}$ to denote i's utility for allocation $\mathbf{x}$; in a public decisions instance, we use $u_i(\mathbf{z}) \in \mathbb{R}$ to denote i's utility for outcome $\mathbf{z}$. In a private goods instance, it is assumed that an agent's utility depends on only the bundle she receives: $u_i(\mathbf{x}) = u_i(x_i)$. In a public decisions instance, agents do not receive separate bundles: instead, the group makes a single decision that affects all agents. We will assume that agents only have utility for their preferred alternative: this will let us standardize notation as follows. For a public decisions outcome $\mathbf{z}$, let $x_{ij}(\mathbf{z}) = z^{j,a_{ij}}$ for all $i \in N$ and $j \in M$ (we will typically write $x_{ij}(\mathbf{z}) = x_{ij}$ for brevity). Then we can define agent i's *public bundle* as $x_i = (x_{i1}, \dots, x_{im})$. An agent's public bundle represents the fraction of the public decision allocated to her preferred alternative, and so we have $u_i(\mathbf{z}) = u_i(x_i)$ in a public decisions instance as well.

Throughout the chapter, we make the following standard assumptions on each agent's utility function u_i:

1. Continuous: $u_i : [0,1]^m \to \mathbb{R}_{\geq 0}$ is a continuous function.

2. Normalized: $u_i(0, 0, ...0) = 0$.

3. Non-constant: There exists a bundle x_i where $u_i(x_i) > 0$.

[10] In our model, divisibility and randomization are equivalent, but this is not always the case: for example, if there were a some sort of global budget constraint, randomization could lead to a non-viable deterministic outcome.

[11] Although the entire supply is typically allocated, it is standard in the private goods literature to allow for outcomes where this does not occur, i.e. $\sum_{i \in N} x_{ij} < 1$. This will be discussed in Section 7.2.2.

4. Monotone: For any bundles x_i and x'_i where $x_{ij} \geq x'_{ij}$ for all j, $u_i(x_i) \geq u_i(x'_i)$.
5. Concave: For any bundles x_i and x'_i and constant $\lambda \in [0,1]$, we have $u_i(\lambda x_i + (1-\lambda)x'_i) \geq \lambda u_i(x_i) + (1-\lambda)u_i(x'_i)$.
6. Homogeneous of degree 1: For any bundles x_i and x'_i and constant $\lambda \geq 0$ where $x_{ij} = \lambda x'_{ij}$ for all j, $u_i(x_i) = \lambda u_i(x'_i)$.

While many of our results hold for any utility functions satisfying our six assumptions, some hold only for particular subclasses; we will make it clear when this is the case.

The first five of those are standard assumptions in the market literature. The last is less ubiquitous, but still common: in particular, the vast majority of the popular subclasses of utility functions satisfy this assumption. For example, it is often assumed in real-world applications that utility functions are *linear*, meaning that

$$u_i(x_i) = \sum_{j \in M} w_{ij} x_{ij}$$

where $w_{ij} \geq 0$ is the weight agent i has for good j. Another important class is *Leontief* functions, where

$$u_i(x_i) = \min_{j \in M : w_{ij} \neq 0} \frac{x_{ij}}{w_{ij}}$$

Linear utilities imply that goods are independent, whereas Leontief utilities represent perfect complements: goods that only have value in combination. For Leontief utilities, w_{ij} is the relative proportion agent i needs of good j.

Both linear and Leontief utilities are generalized by the class of *constant elasticity of substitution* (CES) utilities, where $u_i(x_i) = \left(\sum_{j \in M} w_{ij}^{\rho} x_{ij}^{\rho}\right)^{1/\rho}$ for some constant $\rho \in (-\infty, 0) \cup (0, 1]$. Linear utilities are obtained by setting $\rho = 1$, and taking the limit as ρ approaches $-\infty$ yields Leontief utilities. Taking the limit as ρ approaches 0 gives Cobb-Douglas utility functions, which have the form $u_i(x_i) = \left(\prod_{j \in M} x_{ij}^{w_{ij}}\right)^{1/\sum_{j \in M} w_{ij}}$[12].

The reader may not that this functional form is identical to the CES welfare functions we studied in Part I. A *welfare* function is what we want to maximize as a social planner; a *utility* function is what an agent wishes to selfishly maximize. In this chapter, we focus on CES utility functions, but primarily use Nash welfare as our welfare function.

7.2.2 Private goods & Fisher markets

We will use the fixed-budget model for our private goods market model. We will focus on linear prices, i.e., $p(x_i) = \sum_{j \in M} p_j x_{ij}$. With abuse of notation, we will also write $p(x_i) = x_i \cdot p$. The

[12]The weights w_{ij} have different interpretations for Leontief utilities vs other CES utilities. For example, if there is only a single good, the CES utility form reduces to $w_{i1}x_{i1}$ and the Leontief utility form reduces to x_{i1}/w_{i1}. When we say that taking the limit as $\rho \to -\infty$ yields Leontief utilities, we mean that we obtain the form of Leontief utilities (i.e., a minimization over all the goods).

demand set is defined in the standard way:

$$D_i(p) = \arg\max_{x_i \in \mathbb{R}^m_{\geq 0}:\ x_i \cdot p \leq B_i} u_i(x_i)$$

As are the equilibrium conditions:

1. Each agent receives a bundle in her demand set: $x_i \in D_i(p)$.
2. The market clears: for all j, $\sum_{i \in N} x_{ij} \leq 1$. Also, if $p_j > 0$, then $\sum_{i \in N} x_{ij} = 1$.

The most natural case is when all agents have the same budget, in which case the ME is also called the *competitive equilibrium from equal incomes* [169].

Condition 2 states that the demand never exceeds the supply, and that any good whose supply is not fully exhausted must have price zero. This implies that agents have no utility for the leftover goods: otherwise they would simply buy more with no additional cost. Note that agents can demand more of a good than the available supply if the cost is less than their budget. It is the role of prices at equilibrium to ensure that demand does not exceed supply.

Under the first five assumptions on utility functions described in Section 7.2.1, a market equilibrium is guaranteed to exist for any Fisher market instance [6]. With the addition of the sixth assumption (homogeneity of degree 1), the equilibrium allocations are exactly the allocations maximizing the Nash welfare[13]:

$$NW(\mathbf{x}) = \Big(\prod_{i \in N} u_i(x_i)^{B_i}\Big)^{1/\mathcal{B}}$$

where $\mathcal{B} = \sum_{i \in N} B_i$. Maximum Nash welfare allocations can be computed in polynomial time by the celebrated Eisenberg-Gale (EG) convex program [72, 73][14]. Nash welfare will be the primary objective function we seek to maximize.

7.2.3 Public decisions

As in the private markets case, the maximum Nash welfare outcome can be found via a convex program:

$$\max_{\mathbf{z} \in [0,1]^{m \times 2}} \Big(\prod_{i \in N} u_i(\mathbf{z})^{B_i}\Big)^{1/\mathcal{B}} \quad s.t. \quad z^{j,0} + z^{j,1} \leq 1 \ \ \forall j \in M \tag{7.1}$$

The solution to this convex program can be found in polynomial time. This program is very different than the EG program; however, we will show via our reduction that these programs become identical under a transformation of utility functions and issue space.

Furthermore, even without any knowledge of utility functions, a 1/2 approximation of this program emerges. Because issues are binary, we can very easily guarantee each agent half of her maximum possible utility simply by putting equal probability on each alternative, i.e., $z^{j,0} = z^{j,1} = 1/2$.

[13] This is technically the "budget-weighted" Nash welfare, but we will omit "budget-weighted" throughout the chapter.

[14] This correspondence still holds under slightly weaker assumptions that our six assumptions [104].

It follows from concavity and $u_i(0, 0, ..., 0) = 0$ that this also achieves 1/2 of the maximum possible Nash welfare.

Proposition 7.2.1. *Let* Γ *be a public decisions instance* (N, M) *with agent budgets* $B = (B_1 ... B_n)$*, and let* $\mathbf{z}$ *be the outcome where* $z^{j,0} = z^{j,1} = 1/2$ *for all* $j \in M$*. Then* $\frac{\max_{\mathbf{z}'} NW(\mathbf{z}')}{NW(\mathbf{z})} \leq 2$.[15]

In light of this, we would expect any reasonable mechanism for public decision-making to do no worse than this (in terms of Nash welfare), and hopefully do substantially better. Unfortunately, we show in the next section that the natural adaption of Fisher markets to the public decisions setting does no better than this for several important classes of utility functions. Even worse, in the case of linear utilities – the most important class of utilities in practice – the Nash welfare can be a factor of $O(n)$ worse than optimal.

7.3 Inefficiency of public decisions markets with issue pricing

In a Fisher market, each good is assigned a single price which is common to all agents: thus all agents are treated the same, which is desirable for fairness. This approach can be thought of as giving each agent the same number of votes and allowing them to trade freely: the price for a given issue expresses the "exchange rate" for that issue. This section shows that in the public decisions setting, setting a single price for each issue (issue pricing) can result in very poor equilibria. Although we primarily consider Nash welfare in this chapter, the same family of instances will show that the utilitarian welfare (sum of agent utilities) and egalitarian welfare (the minimum agent utility) can also be much worse than optimal. We do not consider this a substantial result; it is conceptually similar to the "free-rider problem" (where an agent benefits from resources that they do not pay for), which is known to lead to inefficiency for public goods. Rather, it is important that we establish inefficiency for per-issue pricing in our specific model in order to motivate a more complex pricing scheme.

7.3.1 Setup

A *public decisions market* (PDM) consists of a public decisions instance (N, M) along with agent budgets $B = (B_1 \dots B_n)$. This definition is independent of the pricing scheme: we use the term "PDM" to describe all notions of markets for the public decisions setting. This section uses the following scheme: each issue has a price, each agent uses her budget to purchase probability for her preferred alternatives, and the total probability placed on an alternative is the sum over agents of the probability purchased for that alternative.

Given per-issue prices $p \in \mathbb{R}^m_{\geq 0}$, agent i's *private bundle* $y_i \in \mathbb{R}^m_{\geq 0}$ refers to the vector of probabilities that agent i purchases. We say that y_i is *affordable* if $y_i \cdot p \leq B_i$. Throughout the chapter, we will use y_i to refer to i's private bundle, and x_i to refer to i's public bundle (recall, as defined in

[15] Whenever we maximize over outcomes of a public decisions instance, i.e., $\max_{\mathbf{z}} NW(\mathbf{z})$, we implicitly assume that only valid outcomes are considered, meaning that $z^{j,0} + z^{j,1} \leq 1$ for all $j \in M$. The same is true when we maximize over outcomes of a private goods instance, and we adopt these conventions throughout the chapter.

Section 7.2, that the public bundle x_i represents the fraction of the public decision allocated to i's preferred alternative). This distinction only matters in the public decisions setting: we use y_i and x_i interchangeably in the private goods setting.

In this section, for private bundles $\mathbf{y} = (y_1 ... y_n)$, the corresponding outcome $\mathbf{z} = (z^1 ... z^m) \in [0,1]^{m \times 2}$ is

$$z^{j,a} = \sum_{i \in N:\ a_{ij} = a} y_{ij}$$

The above definition of $\mathbf{z}$ as a function of $\mathbf{y}$ is specific to the issue pricing scheme. The different pricing scheme discussed in Section 7.4 will define $\mathbf{z}$ differently.

In a Fisher market, an agent's demand set contains the bundles which maximize her utility subject to being affordable. In a PDM with issue pricing, an agent's utility depends not only on her own bundle, but also on other agent's bundles. Thus if we want to define an agent's demand set as the bundles which maximize her utility subject to being affordable, the demand set must depend not only on the prices, but also on the private bundles of other agents. With this in mind, we define the demand set by

$$D_i(p, y_{-i}) = \operatorname*{arg\,max}_{y_i \in \mathbb{R}^m_{\geq 0}:\ y_i \cdot p \leq B_i} u_i(y_{-i}, y_i)$$

where y_{-i} is the list of private bundles other than that of agent i, and with slight abuse of types, $u_i(y_{-i}, y_i)$ is agent i's utility for the outcome when i purchases private bundle y_i and the other agents purchase private bundles y_{-i}[16].

An *issue-pricing market equilibrium* (IME) $(\mathbf{y}, p)$ is a list of private bundles $\mathbf{y}$ and issue prices $p \in \mathbb{R}^m_{\geq 0}$ where

1. Each agent receives a private bundle in her demand set: $y_i \in D_i(p, y_{-i})$.
2. The market clears: for all j, $\sum_{i \in N} y_{ij} \leq 1$. Also, if $p_j > 0$, then $\sum_{i \in N} y_{ij} = 1$.

By the same reasoning as in the private setting, whenever an issue is not sold completely, agents have no utility for the unsold fraction of the issue[17].

In general it is not known whether every PDM admits an IME. However, for several important utility classes, we give an instance where an IME does exist, but where every IME has poor Nash welfare.

7.3.2 Linear utilities

We first show that for linear utilities, an IME always exists. Furthermore, the set of IMEs is identical to the set of private goods MEs that would be obtained if the input were instead treated as a Fisher market (i.e., if each agent's utility only depended on her private bundle).

[16] As in the Fisher market setting, agents are allowed to demand more than 1 unit of an issue if the cost is less than their budget. The interpretation of demanding more than unit probability is difficult, but the prices will ensure that this never occurs in equilibrium.

[17] If some issue j is not sold completely and so $z^{j,0} + z^{j,1} < 1$, one can think of the remaining $1 - z^{j,0} - z^{j,1}$ being allocated to some third option that has no value for any agent.

To see this, we can write agent i's utility for private bundles $\mathbf{y}$ as

$$u_i(\mathbf{y}) = \sum_{j \in M} w_{ij} \sum_{\substack{k \in N: \\ a_{kj} = a_{ij}}} y_{kj} = \sum_{j \in M} w_{ij} y_{ij} + \sum_{j \in M} w_{ij} \sum_{\substack{k \in N \setminus \{i\}: \\ a_{kj} = a_{ij}}} y_{kj}$$

Agent i cannot affect the actions of other agents, and so has no control over the second term. Thus agent i maximizes her utility by maximizing the first term, $\sum_{j \in M} w_{ij} y_{ij}$, which is exactly the utility function of an agent in a Fisher market. This is expressed formally by Theorem 7.3.1, whose proof appears in Section 7.10.

Theorem 7.3.1. *For a PDM (N, M, B) with linear utilities given by weights $w_{ij} \geq 0$, for every list of private bundles $\mathbf{y}$ and list of prices p, $(\mathbf{y}, p)$ is an IME if and only $(\mathbf{y}, p)$ is a ME for the Fisher market (N, M, B) with linear utilities given by the same weights.*

We now define the family of instances that exhibit poor equilibria in the issue pricing model. For any integer $n \geq 2$ and real number $w \geq 0$, let $\Phi(n, w)$ be the PDM defined by $n = m$, $w_{ii} = w$ for all i, $w_{ij} = 1$ for all $j \neq i$, $a_{ii} = 0$, $a_{ij} = 1$ for all $j \neq i$, and $B_i = 1$ for all i. In words, on each issue i, agent i is alone on one side of the issue, and the other $n - 1$ agents are on the opposite side. Each agent i has weight w for issue i, and weight 1 for every other issue.

Our next theorem shows that for linear utilities, the Nash welfare of the IME can be a linear factor worse than optimal. This is especially dreadful in light of how easy it is to achieve half of the optimal Nash welfare via Proposition 7.2.1.

Theorem 7.3.2. *For any $\epsilon > 0$, $\Phi(n, 1 + \epsilon)$ with linear utilities has a unique equilibrium $(\mathbf{y}, p)$, where*

$$\frac{\max_{\mathbf{z}'} NW(\mathbf{z}')}{NW(\mathbf{y})} \geq \frac{n - 1}{1 + \epsilon}$$

The proof appears in Section 7.10, but we give some intuition here. We observe that an agent's demand set in a Fisher market always maximizes her "bang-per-buck" ratio: w_{ij}/p_j. To see this, suppose agent i spends some money on a good that does not maximize her bang-per-buck ratio: she could instead spend the same amount of money to get strictly more utility by spending it on a good with maximum bang-per-buck. By Theorem 7.3.1, this property carries over to the IME.

By symmetry, every issue will have the same price. Since $w_{ii} > w_{ij}$ for all i and for all $j \neq i$, agent i's bang-per-buck ratio is maximized only by good i. Thus each agent i spends all of her budget on good i. This leads to the outcome where $y_{ii} = 1$ for all i, and $y_{ij} = 0$ for all $j \neq i$. Thus $z^{j,0} = 1$ for all $j \in M$. The utility of each agent for this outcome $1 + \epsilon$, so the Nash welfare is also $1 + \epsilon$. But in the outcome where $z^{j,1} = 1$, for all j, each agent has utility $n - 1$, so the Nash welfare is $n - 1$. This yields the desired bound of $(n - 1)/(1 + \epsilon)$.

If we used $w_{ii} = w_{ij} = 1$ for all i, j, the outcome where $z^{j,0} = 1$ would still be an IME. However, there would now be many more IMEs, including ones with optimal Nash welfare. By setting $w_{ii} = 1 + \epsilon$ instead of $w_{ii} = 1$, we can make the outcome where $x_{ii} = 1$ for all i the unique equilibrium. This same issue is not present for Cobb-Douglas and CES utilities with $\rho \in (-\infty, 0) \cup (0, 1)$, which we examine in the next section.

7.3.3 Other utilities

We briefly mention two results we have for other classes of utility functions. Using the same Φ construction, Theorems 7.3.3 and 7.3.4 state that the Nash welfare of an IME cannot be much better than 1/2 for Cobb-Douglas utilities and CES utilities, respectively. The formal proofs of Theorems 7.3.3 and 7.3.4 appear in Section 7.10.1.

Theorem 7.3.3. *For any IME* $(\mathbf{y}, p)$ *of* $\Phi(n,1)$ *with Cobb-Douglas utilities,*

$$\frac{\max_{\mathbf{z}'} NW(\mathbf{z}')}{NW(\mathbf{y})} \geq \frac{2 - 2/n}{(n-1)^{1/n}}$$

Theorem 7.3.4. *For any IME* $(\mathbf{y}, p)$ *of* $\Phi(n,1)$ *with CES utilities for parameter* $\rho \in (-\infty, 0) \cup (0,1)$*,*

$$\frac{\max_{\mathbf{z}'} NW(\mathbf{z}')}{NW(\mathbf{y})} \geq 2(1 - 1/n)^{1/\rho}$$

As the number of agents approaches infinity, the bounds in Theorems 7.3.3 and 7.3.4 approach 2. This means that for those classes of utility functions, the issue pricing market model cannot be guaranteed to do better than simply picking the midpoint on every issue (Proposition 7.2.1). The situation is even worse for linear utilities, where the Nash welfare of an IME can be arbitrarily worse than the optimal Nash welfare.

One may wonder why Cobb-Douglas and CES utilities with $\rho \in (-\infty, 0) \cup (0,1)$ do not fail as badly as linear utilities on this family of instances. On a high level, the reason is that both Cobb-Douglas and CES utilities exhibit diminishing returns: the more one buys of a particular good, the less value it adds. This leads to agents splitting their money across multiple goods, regardless of their weights on the individual goods. As a result, small changes in agents' weights end up not affecting their purchases too much. In contrast, for linear utilities, an agent might spend her entire budget on a single good: in fact, if there is a unique good which maximizes her bang-per-buck, she must spend her entire budget on that good. This is exactly the property we use in our inefficiency example, where the fact the $w_{ii} = 1 + \epsilon > w_{ij}$ for $j \neq i$ causes agent i to spend her entire budget on good i.

These negative results motivate a more complex market model, which we present in the next section.

7.4 Pairwise issue expansion and pairwise pricing

In this section, we describe a more complex model of a public decisions market, which relies on *pairwise pricing*: for each issue, there will be a price for each pair of agents who disagree on that issue. We then present our main result: a reduction from any PDM to an equivalent Fisher market. This reduction, which we call *pairwise issue expansion*, can be used to construct a *pairwise pricing equilibrium* that maximizes the Nash welfare. Specifically, for any PDM, we will construct a Fisher market such that market equilibria in the Fisher market correspond to pairwise pricing equilibria in the original PDM which maximize Nash welfare.

The section is organized as follows. Section 7.4.1 introduces pairwise issue expansion and gives an informal argument for correctness. Section 7.4.2 gives some additional notation and setup, and states our theorems. The formal proofs of correctness are somewhat technical and are deferred to the end of the chapter, to improve narrative flow. Finally, Section 7.4.3 discusses some Fisher market results that this reduction allows us to immediately lift to the public decisions setting.

7.4.1 Pairwise issue expansion

For any PDM Γ, we construct a Fisher market $R(\Gamma)$ as follows. The set of agents $N = \{1...n\}$ and their budgets $B_1...B_n$ will be the same. Every issue $j \in M$ will become $O(n^2)$ goods in $R(\Gamma)$. Specifically, for every issue j, there will one good for each pair of agents who disagree on issue j. Let $R^g(M)$ be the set of goods in $R(\Gamma)$: then

$$R^g(M) = \big\{(i, k, j) \mid j \in M,\ i, k \in N,\ a_{ij} \neq a_{kj}\big\}$$

We we will refer to goods (k, k', j) where $i \in \{k, k'\}$ as agent i's "pairwise goods". Note that (i, k, j) and (k, i, j) refer to the same good.

If y_i is a bundle associated with Γ (denoted $y_i \sim \Gamma$), then $y_i \in \mathbb{R}_{\geq 0}^{|M|}$. If y_i is a bundle associated with $R(\Gamma)$ (denoted $y_i \sim R(\Gamma)$), then $y_i \in \mathbb{R}_{\geq 0}^{|R^g(M)|}$.

We will use j to represent issues in M and ℓ to represent goods in $R^g(M)$. We also use $y_{i(ikj)}$ to denote $y_{i\ell}$ when $\ell = (i, k, j)$.

In order to purchase α units of issue j in the PDM, agent i will need to purchase at least α units of all of her pairwise goods for issue j. Formally, agent i's utility for a bundle $y_i \in \mathbb{R}_{\geq 0}^{|R^g(M)|}$ is

$$u_i\left(\min_{\substack{k\in N:\\ a_{i1}\neq a_{k1}}} y_{i(ik1)},\ \min_{\substack{k\in N:\\ a_{i2}\neq a_{k2}}} y_{i(ik2)},\ \dots\ \min_{\substack{k\in N:\\ a_{im}\neq a_{km}}} y_{i(ikm)}\right)$$

Agent i's utility is as if she purchased $\min\limits_{k\in N: a_{ij}\neq a_{kj}} y_{i(ikj)}$ probability of each issue j in the PDM Γ. For example, if agent i's utility in Γ is linear with weights w_{ij}, her utility in $R(\Gamma)$ would be

$$\sum_{j\in M} w_{ij}\Big(\min_{\substack{k\in N:\\ a_{ij}\neq a_{kj}}} y_{i(ikj)}\Big)$$

These utility functions are *nested Leontief*; this will be discussed formally in Section 7.4.2.

Figure 7.1 gives a graphical representation of $R(\Gamma)$ for five agents and a single issue j, where $a_{1j} = a_{2j} = a_{3j} = 0$ and $a_{4j} = a_{5j} = 1$. An edge from an agent to a good indicates that that agent desires that good. One key aspect of pairwise issue expansion is that on each issue j, each agent is in competition with everyone she disagrees with, and not in competition with anyone she agrees with.

We first argue informally for the correctness of the reduction (by "correctness", we mean that

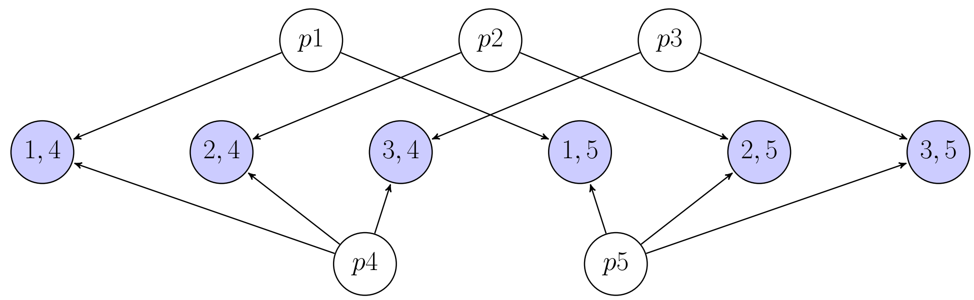

Figure 7.1: A graphical representation of the constructed Fisher market $R(\Gamma)$ for five agents and a single issue j, where $a_{1j} = a_{2j} = a_{3j} = 0$ and $a_{4j} = a_{5j} = 1$.

the equilibria in the Fisher market $R(\Gamma)$ correspond to equilibria in the PDM Γ). Agent i will only ever spend money on her pairwise goods, because other goods do not affect her utility. Because of the nested Leontief structure of the utilities in $R(\Gamma)$, for a fixed issue j, agent i will buy the same amount of each of her pairwise goods: buying a larger amount of one of the goods would not increase her utility (because it would not increase the minimum), so she would be wasting money. Thus for a fixed issue, agent i buys the same amount of each of her pairwise goods (though this can differ across issues).

So suppose that for each issue j, agent i buys α_{ij} of each of her pairwise goods for that issue. If $R(\Gamma)$ is at equilibrium, every agent k who disagrees with agent i on issue j can receive at most $1 - \alpha_{ij}$ of good (i, k, j), since the total supply of each good is 1. As argued above, agent k will never buy more than $1 - \alpha_{ij}$ of any of her pairwise goods on issue j, because of the nested Leontief structure. Thus every agent k who disagrees with agent i on issue j will buy exactly $1 - \alpha_{ij}$ of each of her pairwise goods for issue j. This leaves exactly α_{ij} for everyone who agrees with agent i on issue j. Thus in equilibrium, everyone who agrees with agent i buys α_{ij} of their pairwise goods on issue j, and everyone who disagrees with agent j buys $1 - \alpha_{ij}$ of their pairwise goods on issue j.

This means that when $R(\Gamma)$ is in equilibrium, whenever two agents agree on an issue, they buy the same amount of their pairwise goods for that issue, and whenever they disagree, the amounts they buy sum to 1. Let $\mathbf{z}$ be the outcome where $z^{j,a_{ij}} = \alpha_{ij}$ and $z^{j,1-a_{ij}} = 1 - \alpha_{ij}$ for all $j \in M$. Then $\mathbf{z}$ is a valid outcome of the PDM. Also, because $R(\Gamma)$ is a Fisher market, an equilibrium price vector assigns a single price to each good: this yields a price for each pairwise disagreement on each issue. This leads to the pairwise pricing equilibrium notion, which $\mathbf{z}$ as defined above will satisfy.

Furthermore, we know that any Fisher market equilibrium maximizes Nash welfare. The agents will have the same utilities in both the PDM and the constructed Fisher market at equilibrium, so the Fisher market equilibrium will respond to a pairwise pricing equilibrium which maximizes Nash welfare in the PDM.

Finally, we mention that this reduction can be generalized to d-ary issues under the assumption that each agent has utility for at most one alternative per issue. Instead of one good for each pair of agents who disagree, there would be one good for each set of d agents where each agent has a different preferred alternative, and a similar argument will hold.

7.4.2 Additional setup and formal theorem statements

Some additional notation will be useful. We define relations R and $R_{\leftarrow}$ which will map bundles and prices between Γ and $R(\Gamma)$.

For a bundle $y_i \sim \Gamma$, we define a corresponding bundle $R^b(y_i) \sim R(\Gamma)$ by

$$R^b(y_i)_{(kk'j)} = \begin{cases} y_{ij} \text{ if } i \in \{k, k'\} \\ 0 \text{ if } i \notin \{k, k'\} \end{cases} \qquad \forall (k, k', j) \in R^g(M)$$

where $R^b(y_i)_{(kk'j)}$ denotes the quantity of good (k, k', j) in bundle $R^b(y_i)$. For a bundle $y_i \sim R(\Gamma)$, the corresponding bundle $R^b_{\leftarrow}(y_i) \sim \Gamma$ is defined by

$$R^b_{\leftarrow}(y_i)_j = \min_{\substack{k \in N: \\ a_{ij} \neq a_{kj}}} y_{i(ikj)} \qquad \forall j \in M$$

where $R^b_{\leftarrow}(y_i)_j$ denotes the quantity of issue j in bundle $R^b_{\leftarrow}(y_i)$. Also, for a list of private bundles $\mathbf{y} \sim \Gamma$, we use $R^b(\mathbf{y})$ to refer to the list of private bundles in $R(\Gamma)$ where agent i's bundle is $R^b(y_i)$. Similarly, for any $\mathbf{y} \sim R(\Gamma)$, $R^b_{\leftarrow}(\mathbf{y})$ is a list of private bundles in Γ where agent i's bundle is $R^b_{\leftarrow}(y_i)$.

It is important to note that while the equilibria of Γ and $R(\Gamma)$ coincide, the correspondence is not always meaningful for non-equilibrium outcomes. In particular, not every $y_i \sim R(\Gamma)$ satisfies $y_i = R^b(R^b_{\leftarrow}(y_i))$: for example if $y_{i(kk'j)} > 0$ when $i \notin \{k, k'\}$.

Let u_i be agent i's utility function in Γ. Then agent i's utility function in $R(\Gamma)$ is given by

$$u^R_i(y_i) = u_i(R^b_{\leftarrow}(y_i))$$

This is equivalent to the definition of agent utilities given in Section 7.4.1: simply subtitute the definition of $R^b_{\leftarrow}(y_i)$. Also note that for any $y_i \sim \Gamma$, we have $y_i = R_{\leftarrow}(R^b(y_i))$, and so $u_i(y_i) = u^R_i(R^b(y_i))$.

We would also like to relate prices in Γ and $R(\Gamma)$. Since $R(\Gamma)$ is a Fisher market, any price vector p associated with $R(\Gamma)$ (denoted $p \sim R(\Gamma)$) assigns a single price to each good $\ell \in R^g(M)$: $p \in \mathbb{R}^{|R^g(M)|}_{\geq 0}$. We will be considering per-person per-issue prices for the PDM Γ, so any set of prices p associated with Γ (denoted $p \sim \Gamma$) assigns one price to each person $i \in N$ for each issue $j \in M$: $p \in \mathbb{R}^{m \times n}_{\geq 0}$.

For a price vector $p \sim R(\Gamma)$, we define prices $R^p_{\leftarrow}(p) \sim \Gamma$ by

$$R^p_{\leftarrow}(p)_{ij} = \sum_{\substack{k \in N: \\ a_{ij} \neq a_{kj}}} p_{(ikj)} \qquad \forall i \in N, j \in M$$

where $R^p_{\leftarrow}(p)_{ij}$ is the price of issue j for agent i in price vector $R^p_{\leftarrow}(p)$. In words, $R^p_{\leftarrow}(p)_{ij}$ is the sum of agent i's pairwise prices for issue j. We will also use $R^p_{\leftarrow}(p)_i$ to denote the vector of agent i's prices: $R^p_{\leftarrow}(p)_i = (R^p_{\leftarrow}(p)_{i1} ... R^p_{\leftarrow}(p)_{im})$.

Before we stating our theorems, we should verify that the utilities in $R(\Gamma)$ satisfy the necessary requirements. If the utility functions in Γ are in class $\mathcal{H}$ ($\mathcal{H}$ could be the set of linear utility functions, for example), the utility functions in $R(\Gamma)$ will be *$\mathcal{H}$-nested Leontief.*

Definition 7.4.1. *For some agent i, let $f_{i1}, f_{i2}...f_{iL}$ be Leontief utility functions. Then a utility function u_i is $\mathcal{H}$-nested-Leontief if there exists a utility function $h_i : \mathbb{R}^L_{\geq 0} \to \mathbb{R}_{\geq 0}$ such that $h_i \in \mathcal{H}$, and*

$$u_i(y_i) = h_i\Big(f_{i1}(y_i), f_{i2}(y_i) \dots f_{iL}(y_i)\Big)$$

for any bundle y_i.

In our setting, $L = m$ for all agents, and for each $j \in M$, $f_{ij}(y_i) = R^b_{\leftarrow}(y_i)_j = \min_{k \in N : a_{ij} \neq a_{kj}} y_{i(ikj)}$. Then for each agent i, $u_i^R(y_i) = u_i(R^b_{\leftarrow}(y_i)) = u_i\big(f_{i1}(y_i), f_{i2}(y_i)...f_{iL}(y_i)\big)$.

The next lemma states that as long as h_i and $f_{i1}, f_{i2}...f_{iL}$ satisfy our assumptions on utility functions, their composition does as well.

Lemma 7.4.1. *Suppose that the functions $h_i, f_{i1}, f_{i2} \dots f_{iL}$ are continuous, normalized, concave, homogeneous of degree 1, non-decreasing, and non-constant. Then $u_i = h_i(f_{i1}, f_{i2} \dots f_{iL})$ meets the same conditions.*

We will claim that each market equilibrium in $R(\Gamma)$ corresponds to a *pairwise-pricing market equilibrium* (PME) in Γ. The formal definition of a PME appears in Section 7.9.1. Informally, it is a list of private bundles $\mathbf{y}$ and per-person per-issue prices $p \in \mathbb{R}^{m \times n}_{\geq 0}$ generated by pairwise issue expansion (i.e., $p = R^p_{\leftarrow}(p')$ for some $p' \sim R(\Gamma)$) such that

1. Every agent receives a private bundle in her demand set.
2. Whenever two agents agree on an issue, they purchase the same amount of that issue.
3. Whenever two agents disagree on an issue, they amounts of that issue that they purchase sum to 1.

This is exactly the definition alluded to via the α_{ij} variables in the informal argument given in Section 7.4.1. This leads to the following theorem:

Theorem 7.4.1. *For an allocation $\mathbf{y} \sim R(\Gamma)$ and prices $p \sim R(\Gamma)$, $(\mathbf{y}, p)$ is a ME of the market $R(\Gamma)$ if and only if $(R^b_{\leftarrow}(\mathbf{y}), R^p_{\leftarrow}(p))$ is a PME of the PDM Γ.*

Finally, we wish to claim the maximum Nash welfare outcomes in Γ and $R(\Gamma)$ correspond. We will actually prove this correspondence for all welfare functions, not just the Nash welfare, and even for approximations of welfare functions.

Formally, let $\Psi : \mathbb{R}^n_{\geq 0} \to \mathbb{R}$ be a function. When the n inputs to Ψ are understood to be the n agent utilities for a particular outcome (of a pubic or private instance), we call Ψ a *welfare function*. Because Ψ depends only on the agent utilities, we will use this terminology and notation for both the public and private settings. With slight abuse of types, we will write $\Psi(\mathbf{z}) = \Psi(u_1(\mathbf{z}), u_2(\mathbf{z}), ...u_n(\mathbf{z}))$[18].

[18]Throughout most of the chapter, we use $\mathbf{z}$ to refer to the outcome of a public decisions instance and $\mathbf{x}$ to refer to the outcome of a private goods instance. In this discussion, the welfare functions are the same for both public and private instances, so we will use $\mathbf{z}$ to denote outcomes for both.

Common welfare functions include the utilitarian welfare function, $\Psi(\mathbf{z}) = \sum_{i \in N} u_i(\mathbf{z})$, the egalitarian welfare function, $\Psi(\mathbf{z}) = \min_{i \in N} u_i(\mathbf{z})$, and most importantly for us, the (budget-weighted) Nash welfare function, $\Psi(\mathbf{z}) = \Big(\prod_{i \in N} u_i(\mathbf{z})^{B_i} \Big)^{1/\mathcal{B}}$. We say that an outcome $\mathbf{z}$ is a α-approximation of Ψ if

$$\Psi(\mathbf{z}) \geq \alpha \cdot \max_{\mathbf{z}' \sim \Gamma} \Psi(\mathbf{z}')$$

If $\mathbf{z}$ is an outcome of a public decisions instance, technically $R^b(\mathbf{z})$ does not typecheck, since $\mathbf{z} = (z^1 ... z^m)$ is not a list of bundles. We interpret $R^b(\mathbf{z})$ to mean $R^b(x_1, x_2 ... x_n)$, where x_i is agent i's public bundle as induced by $\mathbf{z}$.

Theorem 7.4.2. *Let Ψ be a welfare function, let Γ be the public decisions instance (N, M) with budgets $B_1 ... B_n$, and let $\alpha \geq 0$. Then $\mathbf{z}$ is an α-approximation of Ψ in Γ if and only if $R(\mathbf{z})$ is an α-approximation of Ψ in $R(\Gamma)$.*

Note that by the same reasoning, $\mathbf{z} \sim R(\Gamma)$ is an α-approximation of Ψ if and only if $R_{\leftarrow}(\mathbf{z}) \sim \Gamma$ is also an α-approximation of Ψ. Thus for any welfare function Ψ and any $\alpha \geq 0$, the α-approximations of Γ and $R(\Gamma)$ correspond exactly.

7.4.3 Lifting Fisher markets results using pairwise issue expansion

In addition to uncovering a surprising conceptual connection, pairwise issue expansion allows us to immediately lift many results from the Fisher market literature to the public decision-making setting. In particular, if a result holds for Fisher market with $\mathcal{H}$-nested Leontief utilities, it holds in the public decisions setting for $\mathcal{H}$ utilities. Any Fisher market result regarding the ME can be lifted using Theorem 7.4.1, and any Fisher market result regarding any approximation of any welfare function can be lifted using Theorem 7.4.2. The following Fisher market results are known:

1. There exists a strongly polynomial-time algorithm for finding a Fisher market equilibrium with two agents and any utility functions [43][19].
2. There exists a strongly polynomial-time algorithm for finding a Fisher market equilibrium for Leontief utilities with weights in $\{0, 1\}$ [87].
3. There exists a polynomial-time algorithm for a Fisher market with Leontief utilities which yields a $O(\log n)$ approximation simultaneously for all canonical welfare functions (i.e. Nash welfare, utilitarian welfare, egalitarian welfare, etc) [95].

The first two can be lifted using Theorem 7.4.1, and the third can be lifted using Theorem 7.4.2. Note that nested Leontief-Leontief functions are just Leontief functions. This yields the following PDM results:

1. There exists a strongly polynomial-time algorithm for finding a PME with two agents and any utility functions.

[19]For completeness, we mention an additional mild condition required for this result: the polytope containing the set of feasible utilities of the two agents must be able to be described via a combinatorial LP.

2. There exists a strongly polynomial-time algorithm for finding a PME for Leontief utilities with weights in $\{0, 1\}$.

3. There exists a polynomial-time algorithm for a PDM with Leontief utilities which yields a $O(\log n)$ approximation simultaneously for all canonical welfare functions (i.e. Nash welfare, utilitarian welfare, egalitarian welfare, etc).

The final result we are interested in lifting is a discrete-time tâtonnement process for finding the Fisher market equilibrium for nested CES-Leontief utilities that converges in polynomial-time [9]. The next section discusses this in more depth.

As a final comment, there are also results of interest that do not immediately fall under the framework of Theorems 7.4.1 and 7.4.2, but which we conjecture could be lifted using our reduction. For example, in Chapter 2 we studied Leontief utilities in the context of price curves. We believe that pairwise issue expansion could be used to generate *pairwise price curves*, where there would be a price curve assigned to each pair of agents who disagree on an issue, instead of a single price. Pairwise issue expansion seems general enough to apply to this type of non-standard market models, but we leave this for future work.

7.5 Public market tâtonnements

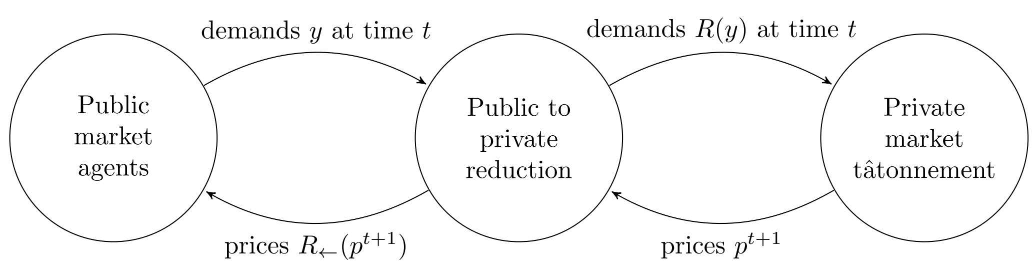

Figure 7.2: $R_{\leftarrow}(\mathcal{T})$ illustration using a hidden private market tâtonnement

In this section, we describe how the reduction immediately leads to existence of several public market tâtonnements. In particular, we show that any Fisher market tâtonnement that works for $\mathcal{H}$-nested utilities yields a PDM tâtonnement for $\mathcal{H}$ utilities.

This does not immediately follow from pairwise issue expansion for several reasons. The first is that tâtonnement deals with approximate equilibria, and Theorem 7.4.1 only considers exact equilibria. Because this correspondence holds for approximate equilibria as well (see proof of Theorem 5.4), any Fisher market tâtonnement that works for $\mathcal{H}$-nested utilities immediately yields an *algorithm* for computing PDM equilibria for $\mathcal{H}$ utilities, but not a tâtonnement. A true PDM tâtonnement would only have access to agents' demands in the PDM, but the resulting algorithm would need to elicit agents' demands in the constructed Fisher market. We handle this by running the Fisher market tâtonnement as a hidden subroutine within the PDM tâtonnement, as demonstrated by Figure 7.2.

Let a Fisher market tâtonnement $\mathcal{T}$ be an iterative algorithm that starts at an initial price vector p^0, and then at each time t,

1. Receives demand set $D_i(p^t)$[20] from each agent i.

2. Updates prices as some function g of the demands, $p^{t+1} = g_{\mathcal{T}}(p^t, D(p^t))$.

As time increases, prices and associated demands approach an approximate equilibrium, for some notion of approximate. Figure 7.2 illustrates the meta-algorithm for public market tâtonnements. From a Fisher market tâtonnement $\mathcal{T}$, let $R_{\leftarrow}(\mathcal{T})$ be the induced public market tâtonnement that initializes an initial price vector p^0 in the hidden Fisher market and then at each time t,

1. Converts prices p^t to public market prices $R_{\leftarrow}(p^t)$ and shows them to agents.

2. Receives demand set $D_i(R_{\leftarrow}(p^t))$ from each agent i.

3. Converts the agent demands to the associated private market demand set $\mathbf{y}^R = \{R(y_i)\}_{y_i \in D_i(R_{\leftarrow}(p^t))}$.

4. Updates prices through the Fisher market tâtonnement function, $p^{t+1} = g_{\mathcal{T}}(p^t, \mathbf{y}^R)$.

We begin with the definitions of approximate equilibria and convergence.

Definition 7.5.1. *A δ-equilibrium $(\mathbf{x}, p)$ in a Fisher market is an allocation $\mathbf{x}$ and prices p where*

1. *Each agent receives a bundle in her demand set: $x_i \in D_i(p)$.*

2. $p_j > \delta \implies \sum_{i \in N} x_{ij} > 1 - \delta$

3. $\forall j, \sum_{i \in N} x_{ij} \leq 1 + \delta$

Note that this definition is introduced in [9].

Definition 7.5.2. *A δ-PME $(\mathbf{y}, p)$ is a list of private bundles $\mathbf{y}$ and per-person per-issue prices $p \in \mathbb{R}^{m \times n}_{\geq 0}$ where*

1. *Each agent receives a private bundle in her demand set: $y_i \in D_i(p_i)$.*

2. *The market approximately clears: there exists a outcome $\mathbf{z} = (z^1 ... z^m) \in [0,1]^{m \times 2}$ where for every issue $j \in M$, all of the following hold:*

   (a) $z^{j,0} + z^{j,1} \leq 1 + \delta$

   (b) *For all $i \in N$, $y_{ij} \leq z^{j,a_{ij}} + \delta$. If $p_{ij} > n\delta$, then $y_{ij} > z^{j,a_{ij}} - \delta$.*

[20] With strictly concave utility functions, each agent's demand at a given price is unique. With linear utilities, the demand set can be expressed as the set of goods that are equally desirable at the given prices. Also, we assume that agents are *price-taking*, meaning that they honestly report their demand given a set of prices, and do not anticipate how prices will change as a result of their actions.

Definition 7.5.3. *A tâtonnement $\mathcal{T}$ has converged to a δ-equilibrium at time T if $\exists \mathbf{y}$ where $(\mathbf{y}, p^T)$ is a δ-equilibrium. Similarly, $R_{\leftarrow}(\mathcal{T})$ has converged to a δ-PME at time T if $\exists \mathbf{y}$ where $(\mathbf{y}, R_{\leftarrow}(p^T))$ is a δ-PME.*

The definition does not imply that *all* demands at the equilibrium prices form an approximate equilibrium, only that there exists an allocation consistent with demands at the equilibrium prices such that the supply constraints are met. However, note that when utility functions are strictly concave, demands are unique.

Our first theorem shows that any Fisher market tâtonnement for $\mathcal{H}$-nested Leontief utility functions yields a PDM tâtonnement for $\mathcal{H}$ utility functions. Theorem 7.5.1, whose proof appears in Section 7.10.1, allows one to lift both convergence and convergence rates from Fisher market tâtonnements.

Theorem 7.5.1. *Consider a Fisher market tâtonnement $\mathcal{T}$. Suppose $\mathcal{T}$ converges to a δ-equilibrium for $\mathcal{H}$-nested leontief utilities in $O(\kappa(m, n, \delta))$ time steps, where n is the number of agents and m the number of goods. Then $R_{\leftarrow}(\mathcal{T})$ converges to a 3δ-PME for the PDM with $\mathcal{H}$ utilities in $O(\kappa(n^2 m, n, \delta))$.*

One Fisher market tâtonnement that we can lift using Theorem 7.5.1 comes from [9], which gives a polynomial-time tâtonnement that converges to a δ-equilibrium for CES-Leontief utilities with $\rho \in (-\infty, 0) \cup (0, 1)$ in polynomial time. By Theorem 7.5.1, this yields a PDM tâtonnement for CES utilities.

We would also like a PDM tâtonnement that works for a wider range of utility functions, especially linear utilities. Section 7.7 presents a stochastic gradient descent style tâtonnement for Fisher markets which converges asymptotically to an equilibrium for all EG utility functions[21], following the framework of [48]. Combined with Theorem 7.5.1, this tâtonnement implies existence of a PDM tâtonnement with asymptotic convergence to an equilibrium for all EG utility functions.

7.6 Lindahl Equilibria

In this section, we show how our public decisions setting corresponds to a natural public goods market in the setting of Lindahl equilibria, and how our reduction can also be used to compute Lindahl prices for this public goods market. The Lindahl Equilibrium has a long history and, at times, the term has been used to mean slightly different things [136]. Foley [82] gave general conditions for the existence of the Lindahl Equilibrium and its correspondence to the core. We first introduce a simplified definition of Lindahl Equilibria.

Definition 7.6.1 ([82, 136])**.** *A Lindahl Equilibrium with m public goods, 1 private good, n agents, and entry private good amounts $\{w_i\}_{i=1}^n$ is an (public goods allocation, private goods allocation, per-person per-issue price) vector $(z^* \in \mathbb{R}_+^m, \{y_i^* \in \mathbb{R}_+\}_{i=1}^n, \{p_i^* \in \mathbb{R}_+^m\}_{i=1}^n)$ such that*

1. *z^* is a solution to $\max_z \left(\sum_{j=1}^n p_j^*\right) z - c(z)$*

[21] functions that meet the 5 conditions from Section 7.2.

2. For each $i \in \{1, \ldots, n\}$, (z^*, y_i^*) *is a solution to* $\max_{(z,y_i)} u_i(z, y_i)$ *subject to* $p_i^* z + y_i \leq w_i$

where $c(z)$ *is the cost to produce the public good vector* z *in terms of the private good and* u_i *is the participant utility function in terms of the public and private goods.*

In a Public Decision Market, the private good is the "influence" of each agent, for which agents have no utility, i.e., influence not spent is lost. Furthermore, there are 2 public goods per issue, 1 for each alternative, and each with the same price for each agent. Similarly,

$$c(z) = \begin{cases} 0 & z^{j,0} + z^{j,1} \leq 1 \ \forall j \\ \infty & \text{else} \end{cases}$$

i.e., if in the case in which each alternative on each issue is implemented with some probability, then there is no cost of using the entire probability, and no possibility of creating more probability.

Lemma 7.6.1. *An equilibrium* $(z^* \in \mathbb{R}_+^{2\times m}, \{y_i^* \in \mathbb{R}_+\}_{i=1}^n, \{p_i^* \in \mathbb{R}_+^{2\times m}\}_{i=1}^n)$ *of the PDM with* m *issues and agent budgets* B_i*, where* z^* *is the decision vector,* p^* *are the per-person per-issue prices from the Fisher market reduction, and final influence vectors are* $y_i^* = 0$*, is a Lindahl Equilibrium with entry private good amounts* $w_i = B_i$*.*

Proof. Condition 2 follows directly from the equilibrium condition that optimal allocation is in the demand set of each agent at the equilibrium prices. Condition 1 requires a bit more work. Note that $\forall j$, for each copy of the good, the sub-price $p_{ikj} = p_{kij}$, where $a_{ij} \neq a_{kj}$. Then, $\forall j, p_j^0 = \sum_{i:a_{ij}=a} \sum_{k:a_{kj}\neq a} p_{ikj} \implies p_j^0 = p_j^1$. Thus, all feasible x are in the solution set in the first condition, and an equilibrium of the PDM is feasible. □

Lemma 7.6.1, alongside Theorem 7.4.1 and the existence of Fisher market equilibria [6] establishes the existence of a Lindahl equilibrium in our setting.

We note that Lemma 7.6.1 further establishes that the solution is in the *core*, as Lindahl Equilibria are in the core [82].

7.6.1 Economies with public goods

The existence of Lindahl Equilibria in Public Good economies (of which Public Decision Markets are a special case, as we will show) was established by Foley [82]. The chief technique is a reduction to Private Goods economies. The reduction yields a non-constructive existence proof, and operates as follows: create a copy of each good for each participant, with equality of the amount of each good enforced through conic constraints. Then, the proof is finished by invoking the existence of equilibria satisfying certain conditions in Private Goods economies, after showing that the additional constraints restricting the cone of production do not violate any assumptions [59]. We note that existence of Lindahl equilibria of the PDM can also be established non-constructively through the same technique, by showing that the resulting market satisfies the assumptions in [82].

One natural question is how Foley's reduction to private goods compares to the reduction in this work. Equation (7.2) contains the convex program to find MNW through Foley's reduction. We use $s \in \{0,1\}$ to denote each side of the issue. Equation (7.3) contains our reduction, which has a nested utility function structure.

$$\begin{aligned}
\max_{\mathbf{x}\in[0,1]^{(2\times m)\times n}} & \Big(\prod_{i\in N} u_i(x_i)^{B_i}\Big)^{1/\mathcal{B}} \\
s.t.\ \ & x_{ij}^0 + x_{ij}^1 \le 1 && \forall j \in M,\ i \in N \\
& x_{ij}^s = x_{kj}^s && \forall j \in M,\ i,k \in N, s \in \{0,1\}
\end{aligned} \tag{7.2}$$

$$\begin{aligned}
\max_{v\in[0,1]^{m\times n},\mathbf{x}\in[0,1]^{\tilde{n}}} & \Big(\prod_{i\in N} u_i(v_i)^{B_i}\Big)^{1/\mathcal{B}} \\
s.t.\ \ & v_{ij} \le x_{i(ikj)} && \forall j \in M,\ i,k \in N, a_{ij} \neq a_{kj} \\
& x_{i(ikj)} + x_{k(ikj)} \le 1 && \forall j \in M,\ i,k \in N, a_{ij} \neq a_{kj}
\end{aligned} \tag{7.3}$$

Where $\tilde{n}$ is 2(# of disagreement pairs across issues). Both programs can be solved in polynomial time. However, Foley's reduction does not obviously resemble a Fisher market due to the extra equality constraints.

Comment 7.6.1. Program (7.2) does not transform into a Fisher market.

Proof. We write the Lagrangian of the Program (7.2), with the objective function written in log form.

$$\begin{aligned}
L(x,p,q) = \sum_i B_i \log(u_i(x_i)) - \sum_{i,j} p_{ij}(x_{ij}^0 + x_{ij}^1) + \sum_{i,j} p_{ij} - \sum_{i\neq k,j,s} q_{ikj}^s (x_{ij}^s - x_{kj}^s) \\
s.t.\ \ p \ge 0, x \ge 0, q_{ikj}^s \ge 0
\end{aligned}$$

This Lagrangian has per-person per-issue prices q_{ikj}^s, p_{ij} for each side of each issue that cannot be trivially turned into per-good prices with separate goods not having joint constraints. If one considers each (i,j,s) tuple a good in the Fisher market, the two goods associated with the two sides of each issue, x_{ij}^0, x_{ij}^1 are coupled through p_{ij}. Similarly, if pair (i,j) corresponds to a good, goods across individuals x_{ij}^s, x_{kj}^s are coupled through q_{ikj}^s in a way that does not resemble supply constraints. Merging the goods for each side, such that the good represents the probability on one of the alternatives for each issue, would violate the non-decreasing constraint for utility functions. Eliminating or combining either of these variables would amount to eliminating the corresponding constraints. This coupling prevents a non-trivial mapping to the Fisher market case, which does not have cross-good constraints but has per-good prices. □

Our reduction fills this gap for Public Decision markets, showing that these prices can emerge from a pure Fisher market with a modification of buyer utilities. In Program (7.3), there is a good (i,k,j) for each (i,k) pair that disagrees on issue j, with i's utility function only dependent on the amount i buys, x_{ikj}. This program thus yields goods with per-good pricing (on the Fisher market goods) and no cross-goods constraints.

Proposition 7.6.1. *Programs* (7.2) *and* (7.3) *are equivalent.*

Proof. Both are equivalent to Program (7.1), which we repeat below:

$$\max_{\mathbf{z}\in[0,1]^{2\times m}} \Big(\prod_{i\in N} u_i(z)^{B_i}\Big)^{1/\mathcal{B}}$$
$$s.t.\quad z^{j,0}+z^{j,1}\le 1 \qquad \forall j\in M$$

Program (7.2) and (7.1) are immediately equivalent by combining variables. The equivalence of the reduction (Theorem 7.4.1) establishes that Program (7.3) and PDM Program (7.1) have the same solution. □

7.7 General tâtonnement with asymptotic convergence in Fisher markets

The prior work discussed in Section 7.5 established deterministic tâtonnements with polynomial time convergence rates only for certain classes of utility functions, or for those that converge to weak equilibria. However, we would like a general PDM converges for a wider class of utilities (in particular, linear utilities). To achieve this, we sacrifice convergence in polynomial time (or any characterization of convergence rates), which has been the primary focus on prior work such as [9, 48][22]. In this section, we present a discrete stochastic gradient descent style tâtonnement for all Fisher markets that result from PDMs with EG utility functions through the reduction. This can then be lifted through Theorem 7.5.1 to yield a general PDM tâtonnement.

This section broadly follows the gradient descent framework for Fisher market tâtonnements from [48], and the tâtonnement can be seen as an asymptotic discretization of their continuous time tâtonnement. The tâtonnement operates on the dual of the EG convex program, which is a function of the prices, and whose gradient is the excess demand (Lemma 7.7.1). We first establish that there exists a bounded, convex region Π in which demands are bounded, with $p^*\in\Pi$[23]. We finish the proof with a standard SGD convergence technique, Lemma 7.7.3.

Let $\phi(p)$ be the objective of the dual of EG convex program. We use D^R to denote demands as the demands are in a Fisher market $R(\Gamma)$ that is constructed from PDM Γ. The following lemma from [48] establishes that $\phi(p)$ is itself a convex function whose gradient is the excess demand.

Lemma 7.7.1 ([48])**.** $\nabla\phi(p)=1-\sum_i D_i^R(p_i)$

[22] Those works take great care to show conditions analogous to strong convexity or Lipschitz continuity of the gradient in the cases of interest.

[23] It is well known that the primal objective function is strictly concave, and so p^*, the optimal dual solution, is unique.

Note that $\phi(p)$ refers to the set of sub-gradients, and $D_i^R(p_i)$ to the set of demands. Even when demands at a given price are not unique (such as with linear utilities), each combination of demands yields a sub-gradient of the dual objective function.

Before being able to apply a canonical gradient descent convergence theorem, we need to establish that there exists a bounded, convex set which contains the optimal price p^* in its interior. We construct such a set next.

Lemma 7.7.2. *$\exists \Pi \subset \mathbb{R}_+^m$ bounded and convex s.t. $p^* \in \arg\max \phi(p) \subset \Pi$, and that $\forall p \in \Pi, \forall i$, $D_i^R(p) < \infty$.*

Proof. We claim $p^* \in [0,2]^m$ in our setting. Let $y_i \in D_i^R(p_i^*)$, and thus $p_j^* y_{ij} \leq B_i = 1 \, \forall j$. By Fisher market equilibrium conditions, $p_j^* > 0 \implies \sum_i y_{ij} = 1$. In our setting, $y_{ij} > 0$ for at most two distinct i. Thus, $\exists i$ s.t. $y_{ij} \geq \frac{1}{2} \implies p_j^* \leq 2$.

Let $p_{\min}$ be any value such that $0 < p_{\min} < \min_{\{j : p_j^* > 0\}} p_j^*$. Then, let $\Pi = [p_{\min}^1, 2] \times \cdots \times [p_{\min}^m, 2]$, where $p_{\min}^j = \begin{cases} 0 & p_j^* = 0 \\ p_{\min} & \text{else} \end{cases}$. Π as defined has the desired properties. □

Throughout, we use $[\cdot]_{\mathcal{X}}$ denote the projection onto a set $\mathcal{X}$. We will also use the following stochastic gradient descent convergence lemma.

Lemma 7.7.3 ([105])**.** *Consider a convex function f on a non-empty bounded closed convex set $\mathcal{X} \subset \mathbb{R}^m$, and use $[\cdot]_{\mathcal{X}}$ to designate the projection operator. Starting with some $x_0 \in \mathcal{X}$, consider the SGD update rule $x_t = [x_{t-1} - \eta_t(\nabla f(x_t) + z_t + e_t)]_{\mathcal{X}}$, where z_t is a zero-mean random variable and e_t is a constant. Let $\mathrm{E}_t[\cdot]$ be the conditional expectation given $\mathcal{F}_t$, the σ-field generated by $x_0, x_1, \ldots, x_t$. If*

$$
\begin{aligned}
& f(\cdot) \textit{ has a unique minimizer } x^* \in \mathcal{X} \\
& \eta_t > 0, \sum_t \eta_t = \infty, \sum_t \eta_t^2 < \infty \\
& \exists C_1 \in \mathbb{R} < \infty \textit{ s.t. } \|\nabla f(x)\|_2 \leq C_1, \forall x \in \mathcal{X} \\
& \exists C_2 \in \mathbb{R} < \infty \textit{ s.t. } \mathrm{E}_t[\|z_t\|^2] \leq C_2, \forall t \\
& \exists C_3 \in \mathbb{R} < \infty \textit{ s.t. } \|e_t\|_2 \leq C_3, \forall t \\
& \sum_t \eta_t \|e_t\| < \infty \textit{ w.p. } 1
\end{aligned}
$$

Then $x_t \to x^*$ w.p. 1 as $t \to \infty$.

We can now construct a stochastic tâtonnement for a Fisher Market $R(\Gamma)$ that is constructed from PDM Γ when agents have any EG utility function in the PDM, as $\mathcal{H}$-nested leontief utility functions remain EG utility functions.

Lemma 7.7.4. *Suppose $p^* \in \arg\min \phi(p)$. Then, $\exists y = (y_1 \ldots y_n)$ s.t. $y_i \in D_i^R(p^*)$, (y, p^*) is a ME.*

Proof. Follows directly from Part 2 of Lemma 5 in [48], that $\arg\max_{x\geq 0} L(x,p) \subseteq D^R(p)$. The ME (x^*, p^*) is such that $x^* \in \arg\max_{x\geq 0} L(x,p^*) \subseteq D^R(p^*)$. ☐

Theorem 7.7.1. *Let agents i in the PDM Γ have utilities $u_i \in \mathcal{H}$ that are concave, continuous, non-decreasing, non-constant, and homogeneous of degree 1. Then there exists a stochastic gradient descent-style tâtonnement for which, as $t \to \infty$, $p^t \to p^*$, where $p^* \in \arg\min \phi(p)$, and $\exists x$ s.t. $x_i \in D_i(p^*)$ and $(x, R_{\leftarrow}(p^*))$ is a PME.*

Proof. Construct a Fisher market $R(\Gamma)$ through the reduction.

Let $\mathcal{T}$ be the following descent in the constructed market. Start with prices $p^0 \in \Pi$, where Π as defined in Lemma 7.7.2. Update prices using the rule $p^{t+1} = [p^t - \eta_t \left(1 - \sum_i \tilde{D}_i^R(p_i^t)\right)]_\Pi$, for $\eta_t = \frac{1}{t}$, and $\tilde{D}_i^R(p) = y_{i,t} + b_{i,t} + z_{i,t}$, for some $y_{i,t} \in D_i^R(p^t)$. Assume $z_{i,t}$ a zero-mean random variable and $b_{i,t}$ a constant that follow the conditions of Lemma 7.7.3.

By Lemma 7.4.1, the implied utility functions still yield Eisenberg-Gale markets and by Lemma 7.7.1, $\nabla\phi(p) = 1 - \sum_i D_i^R(p)$. By Lemma 7.7.2, $\exists C < \infty$ *s.t.* $\|\nabla\phi(p)\| < C, \forall p \in \Pi$. Convergence to prices $p^* \in \arg\min \phi(p)$ follows from Lemma 7.7.3. By Lemma 7.7.4, $\exists \mathbf{y}$, $y_i \in D_i^R(p^*)$ s.t. $(\mathbf{y}, p^*)$ is a ME in the Fisher market. Thus, $\mathcal{T}$ converges asymptotically to a ME.

By Theorem 7.5.1, $\mathcal{T}$ can be lifted to create a tâtonnement $R_{\leftarrow}(\mathcal{T})$ in the PDM that converges asymptotically to a PME.

☐

Note that Π is not known a priori. However, it can be approximated during the gradient descent without affecting convergence: for example, if at any point demand goes to infinity, backtrack and impose a minimum price. Then, if demand goes to 0 with a positive price, lower this minimum price.

Theorem 7.7.1 and Theorem 7.5.1 together create a tâtonnement with asymptotic convergence for Public Decision Markets for general concave, continuous, non-constant, non-decreasing, and homogeneous of degree 1 utility functions.

7.8 Conclusion

In this chapter, we studied adaptations of markets to the public decision-making setting. In Section 7.3, we showed that issue pricing in the public decisions setting can yield very poor equilibria: for linear utilities, the Nash welfare can be a factor of $O(n)$ worse than optimal. This is in contrast to private goods, where per-good pricing is the accepted standard, and yields optimal equilibria. We showed in Section 7.4 that pairwise issue expansion reduces any public decisions market to an equivalent Fisher market, how optimal equilibria can be constructed using this reduction. We used pairwise issue expansion to lift various Fisher market results to the public decision-making context, including tâtonnement, which we discussed in Section 7.5.

Most importantly, our reduction uncovers a powerful connection between the private goods and public decision-making settings that we believe has many possible applications. For example, suppose we had a mechanism for private goods which computes some desirable outcome other than maximum Nash welfare (maybe it computes the allocation which maximizes the minimum utility, for example).

If that algorithm works for nested $\mathcal{H}$-Leontief utilities for private goods, we imagine that it could be immediately lifted to work for $\mathcal{H}$ utilities in for public-decisions. More generally, it seems like more or less any result that applies to $\mathcal{H}$-Leontief utilities for private goods would apply for $\mathcal{H}$ utilities for public decisions. We believe this merits more study.

7.9 Omitted definitions and proofs from Section 4

Section 7.9.1 contains the formal definitions of the pairwise pricing model. Section 7.9.2 contains the formal analysis of R and $R_{\leftarrow}$, leading to Theorems 7.4.1 and 7.4.2.

7.9.1 Pairwise pricing

In the issue pricing model of Section 7.3, the price for an issue was the same for all agents, and the amount of probability put on alternative a on issue j in the outcome (denoted $z^{j,a}$) was the sum of the agents' purchases. Formally, Section 7.3 defined $z^{j,a} = \sum_{k \in N:\ a_{ij}=a} y_{ij}$ where y_{ij} is the probability that agent i purchased on issue j (y_i is agent i's private bundle).

The definition of $z^{j,a}$ will be different here. This section describes a model where agents may have different prices for the same issue. This will allow us to enforce that in equilibrium, all agents who agree on issue j will purchase the same amount of issue j. For every $i \in N$ and $j \in M$ where agent i's price for issue j is nonzero, at equilibrium[24] we will have

$$z^{j,a_{ij}} = y_{ij}$$

The key consequence is that each agent's private and public bundles will be the same in equilibrium, and so $u_i(y_i) = u_i(x_i(\mathbf{z}))$. Thus each agent's utility can be written as a function of only her private bundle: this will enable the reduction to private goods.

We now formally describe the personalized pricing model. For prices $p \in \mathbb{R}_{\geq 0}^{m \times n}$, let p_{ij} be the price for agent i for issue j, and let $p_i = (p_{i1} ... p_{im})$. Formally, a private bundle y_i is affordable if $y_i \cdot p_i = \sum_{j \in M} y_{ij} p_{ij} \leq B_i$. Because we will have $u_i(y_i) = u_i(x_i)$ in equilibrium, we can define agent i's demand to be independent of other agents' private bundles[25]:

$$D_i(p) = \underset{y_i \in \mathbb{R}_{\geq 0}^m:\ y_i \cdot p_i \leq B_i}{\arg\max}\ u_i(y_i)$$

A *personalized-pricing market equilibrium* (PME) $(\mathbf{y}, p)$ is a list of private bundles $\mathbf{y}$ and personalized prices $p \in \mathbb{R}_{\geq 0}^{m \times n}$ where

1. Each agent receives a private bundle in her demand set: $y_i \in D_i(p_i)$.

[24] The outcome will not be well-defined for a list of private bundles not at equilibrium, since agents may have incompatible demands. This will not be important; we mention it only for completeness.

[25] We assume that agents truthfully report their demands according to this definition: recall that we do not consider strategic behavior.

2. The market clears: there exists an outcome $\mathbf{z} = (z^1 ... z^m) \in [0,1]^m$ where for every issue $j \in M$, all of the following hold:

   (a) $z^{j,0} + z^{j,1} = 1$

   (b) For all $i \in N$, $y_{ij} \leq z^{j,a_{ij}}$. If $p_{ij} > 0$, then $y_{ij} = z^{j,a_{ij}}$.

The market clearing condition (Condition 2) is different than in traditional private goods markets. Instead of the sum of the agent's demands being equal to the supply, the condition here is that there is a single outcome that is consistent with every agent's demand. Roughly speaking, this means that whenever two agents agree on an issue, they demand the same quantity of that issue, and whenever two agents disagree, the sum of their demands equals the supply. This can be thought of as all agents buying the "same" private bundle, modulo their preferred alternatives.

At equilibrium, $\mathbf{z}$ is treated as the outcome of the public decisions instance. However, $\mathbf{z}$ may not be unique: if $y_{ij} < z^{j,a_{ij}}$ for some i, j, there may be multiple outcomes compatible with the list of agent demands. The following proposition shows that all outcomes compatible with $\mathbf{y}$ are more or less the same.

Proposition 7.9.1. *Let $(\mathbf{y}, p)$ be a PME. Then for any outcome $\mathbf{z}$ satisfying the market clearing condition, $u_i(\mathbf{z}) = u_i(y_i)$ for all $i \in N$.*

Proof. Fix some agent $i \in N$, and let y'_i be the private bundle where $y'_{ij} = z^{j,a_{ij}}$ for all $j \in M$. For every issue j where $y_{ij} \neq z^{j,a_{ij}}$, we have $p_{ij} = 0$. Thus y_i and y'_i have the same cost. Since y_i is in agent i's demand set, y_i is affordable. Thus y'_i is also affordable. Suppose $u_i(\mathbf{z}) = u_i(y'_i) > u_i(y_i)$: then y_i would not be in agent i's demand set, which is a contradiction. □

Since each agent's private and public bundles are the same at equilibrium in this model, we mostly omit "private" and "public" and just use the term "bundle". We reserve $\mathbf{z}$ for denoting the overall outcome of the PDM, and just use y_i to denote agent i's bundle.

7.9.2 Formal analysis of pairwise issue expansion

We begin with Lemma 7.9.1, which states that as long as agents only buy their pairwise goods, the cost of a bundle $y_i \sim R(\Gamma)$ at prices p is the same as the cost of bundle $R_{\leftarrow}(y_i) \sim \Gamma$ at prices $R_{\leftarrow}(p)$. The proof primary consists of arithmetic and substituting definitions.

Lemma 7.9.1. *Given prices $p \sim R(\Gamma)$ and a bundle $y_i \sim R(\Gamma)$,*

1. $R_{\leftarrow}(y_i) \cdot R_{\leftarrow}(p)_i \leq y_i \cdot p$

2. *Suppose that (1) for any $j \in M$ and $k, k' \in N \backslash \{i\}$ where $y_{i(kk'j)} \neq 0$, we have $p_{(kk'j)} = 0$, and (2) for any $j \in M$ and $k \in N$ where $y_{i(ikj)} \neq \min_{\substack{k \in N: \\ a_{ij} \neq a_{kj}}} y_{i(ikj)}$, we have $p_{(ikj)} = 0$. Then $R_{\leftarrow}(y_i) \cdot R_{\leftarrow}(p)_i = y_i \cdot p$.*

Proof. Suppose y_i is a bundle in $R(\Gamma)$. The cost of y_i at prices p is

$$\begin{aligned}
y_i \cdot p &= \sum_{\ell \in R(M)} y_{i\ell} p_\ell && \text{(by definition)} \\
&= \sum_{j \in M} \sum_{\substack{k,k' \in N: \\ a_{kj} \neq a_{k'j}}} y_{i(kk'j)} p_{(kk'j)} && \text{(rewriting each good } \ell \in R(M) \text{ as a triple } (i,k,j)) \\
&\geq \sum_{j \in M} \sum_{\substack{k \in N: \\ a_{ij} \neq a_{kj}}} y_{i(ikj)} p_{(ikj)} && \text{(only including agent } i\text{'s pairwise goods in the sum)} \\
&\geq \sum_{j \in M} \sum_{\substack{k \in N: \\ a_{ij} \neq a_{kj}}} p_{(ikj)} \min_{\substack{k' \in N: \\ a_{ij} \neq a_{k'j}}} y_{i(ik'j)} && \text{(replacing each } y_{i(ikj)} \text{ with } \min_{\substack{k' \in N: \\ a_{ij} \neq a_{k'j}}} y_{i(ik'j)}) \\
&= \sum_{j \in M} \sum_{\substack{k \in N: \\ a_{ij} \neq a_{kj}}} p_{(ikj)} R_{\leftarrow}(y_i)_j && \text{(by definition)} \\
&= \sum_{j \in M} R_{\leftarrow}(y_i)_j \sum_{\substack{k \in N: \\ a_{ij} \neq a_{kj}}} p_{(ikj)} && \text{(moving } R_{\leftarrow}(y_i)_j \text{ out of the inner sum)} \\
&= \sum_{j \in M} R_{\leftarrow}(y_i)_j R_{\leftarrow}(p)_{ij} && \text{(by definition)} \\
&= R_{\leftarrow}(y_i) \cdot R_{\leftarrow}(p)_i && \text{(by definition)}
\end{aligned}$$

Furthermore, the first inequality holds with equality if for any $j \in M$ and $k, k' \in N \backslash \{i\}$ where $y_{i(kk'j)} \neq 0$, $p_{(kk'j)} = 0$. Similarly, the second inequality holds with equality if for any $j \in M$ and $k \in N$ where $y_{i(ikj)} \neq \min_{\substack{k \in N: \\ a_{ij} \neq a_{kj}}} y_{i(ikj)}$, $p_{(ikj)} = 0$. Therefore under those two assumptions, $R_{\leftarrow}(y_i) \cdot R_{\leftarrow}(p)_i = y_i \cdot p$. □

Lemma 7.9.2 is a simple application of Lemma 7.9.1.

Lemma 7.9.2. *For prices $p \sim R(\Gamma)$ and a bundle $y_i \sim \Gamma$, we have $y_i \cdot R_{\leftarrow}(p)_i = R(y_i) \cdot p$.*

Proof. By definition of $R(y_i)$, we have (1) $y_{i(kk'j)} = 0$ for all $j \in M$ and $k, k' \in N \backslash \{i\}$, and (2) $y_{i(ikj)} = \min_{\substack{k \in N: \\ a_{ij} \neq a_{kj}}} y_{i(ikj)}$ for all $j \in M$ and $k \in N$. Then by Lemma 7.9.1, $y_i \cdot R_{\leftarrow}(p)_i = R(y_i) \cdot p$. □

Lemma 7.9.3 states that if a bundle $y_i \sim R(\Gamma)$ is agent i's demand set $D_i^R(p)$, then (1) y_i contains only agent i's pairwise goods, and (2) for a fixed issue j, y_i contains the same amount of each of her pairwise goods. The proof is based on the informal argument given before: violating either (1) or (2) wastes money that could be spend to increase her utility.

Lemma 7.9.3. *Given prices $p \sim R(\Gamma)$, suppose a bundle $y_i \sim R(\Gamma)$ is in $D_i^R(p)$. Then (1) for any $j \in M$ and $k, k' \in N \backslash \{i\}$ where $y_{i(kk'j)} \neq 0$, we have $p_{(kk'j)} = 0$, and (2) for any $j \in M$ and $k \in N$ where $y_{i(ikj)} \neq \min_{\substack{k \in N: \\ a_{ij} \neq a_{kj}}} y_{i(ikj)}$, we have $p_{(ikj)} = 0$.*

Proof. First, suppose for sake of contradiction that there exists $j \in M$ and $k, k' \in N \backslash \{i\}$ where $p_{(kk'j)} > 0$ and $y_{i(kk'j)} > 0$. Consider the bundle $y'_i \sim R(\Gamma)$ which is identical to y_i, except that $y'_{i(kk'j)} = 0$. Since $R_{\leftarrow}(y_i) = R_{\leftarrow}(y'_i)$, we have $u_i^R(y_i) = u_i^R(y'_i)$. But since $p_j > 0$, the $y_i \cdot p - y'_i \cdot p = y_{i(kk'j)}p_{(kk'j)}$. Consider the bundle $y''_i \sim R(\Gamma)$ where for all $\ell \in R(M)$,

$$y''_{i\ell} = y'_{i\ell} + \frac{y_{i(kk'j)}p_{(kk'j)}}{\sum_{\ell' \in R(M)} p_{\ell'}}$$

Then we have $y''_i \cdot p = y'_i \cdot p + y_{i(kk'j)}p_{(kk'j)} = y_i \cdot p$. Since $y_i \in D_i^R(p)$, y_i is affordable at prices p. Thus y''_i is affordable at prices p.

Finally, we show that $u_i^R(y''_i) > u_i^R(y_i)$. We have $y''_{i\ell} > y'_{i\ell}$ for all $\ell \in R(M)$. Thus for all ℓ, there exists a constant $\alpha_\ell > 1$ where $y''_{i\ell} = \alpha_\ell y'_{i\ell}$. Let $\alpha = \min_{\ell \in R(M)} \alpha_\ell$. Then $y''_{i\ell} \geq \alpha y'_{i\ell}$ for all $\ell \in R(M)$. Because u_i^R is homogenous of degree 1 and monotone, we have $u_i^R(y''_{i\ell}) \geq \alpha \cdot u_i^R(y'_{i\ell}) > u_i^R(y'_{i\ell}) = u_i^R(y_i)$.

Thus we have $u_i^R(y''_{i\ell}) > u_i^R(y_{i\ell})$ and $y''_i \cdot p = y'_i \cdot p$. But then y_i cannot be in agent i's demand set, which is a contradiction.

The second case is similar. Suppose that there exists $j \in M$ and $k \in N$ where $p_{(ikj)} > 0$ and $y_{i(ikj)} > \min_{\substack{k' \in N: \\ a_{ij} \neq a_{k'j}}} y_{i(ik'j)}$. Define the bundle $y'_i \sim R(\Gamma)$ to be identical to y_i, except that $y_{i(ikj)} = \min_{\substack{k' \in N: \\ a_{ij} \neq a_{k'j}}} y_{i(ik'j)}$. Define the bundle $y''_i \sim R(\Gamma)$ by

$$y''_{i\ell} = y'_{i\ell} + \frac{\left(y_{i(ikj)} - \min_{\substack{k' \in N: \\ a_{ij} \neq a_{k'j}}} y_{i(ik'j)}\right)p_{(ikj)}}{\sum_{\ell' \in R(M)} p_{\ell'}}$$

Then $u_i^R(y''_i) > u_i^R(y'_i) = u_i^R(y_i)$, and $y''_i \cdot p = y_i \cdot p$. Thus y_i cannot be in agent i's demand set, which is a contradiction. ☐

Lemma 7.9.4 is a straightforward combination of the previous two lemmas.

Lemma 7.9.4. *Given prices $p \sim R(\Gamma)$, suppose a bundle $y_i \sim R(\Gamma)$ is in $D_i^R(p)$. Then $R_{\leftarrow}(y_i) \cdot R_{\leftarrow}(p)_i = y_i \cdot p$.*

Proof. By Lemma 7.9.3, we have (1) for any $j \in M$ and $k, k' \in N \backslash \{i\}$ where $y_{i(kk'j)} \neq 0$, we have $p_{(kk'j)} = 0$, and (2) for any $j \in M$ and $k \in N$ where $y_{i(ikj)} \neq \min_{\substack{k \in N: \\ a_{ij} \neq a_{kj}}} y_{i(ikj)}$, we have $p_{(ikj)} = 0$. Therefore by Lemma 7.9.1, we have $R_{\leftarrow}(y_i) \cdot R_{\leftarrow}(p)_i = y_i \cdot p$. ☐

Lemma 7.9.5 states that y_i is in agent i's demand set in $R(\Gamma)$ if and only if $R_{\leftarrow}(y_i)$ is in agent i's demand set in Γ. This will not only play an important role in the proof of Theorem 7.4.1, but also later on in tâtonnement.

The majority of the proof of Lemma 7.9.5 is devoted to proving that

$$\max_{\substack{y'_i \sim \Gamma: \\ y'_i \cdot R_{\leftarrow}(p)_i \leq B_i}} u_i(y'_i) = \max_{\substack{y'_i \sim R(\Gamma): \\ y'_i \cdot p \leq B_i}} u_i^R(y'_i)$$

The intuitive argument for the above equality is that the utilities and costs of bundles are the same in both Γ and $R(\Gamma)$. Slightly more formally, for any bundle $y_i \sim \Gamma$, $R(y_i) \sim R(\Gamma)$ has the same utility (by definition) and the same cost (by Lemma 7.9.2). For any bundle $y_i \sim R(\Gamma)$, $u_i^R(y_i) = u_i(R_{\leftarrow}(y_i))$ is also true by definition, but y_i and $R_{\leftarrow}(y_i)$ do not necessarily have the same cost. That is where Lemma 7.9.4 will be important.

Lemma 7.9.5. *Given prices $p \sim R(\Gamma)$ and a bundle $y_i \sim R(\Gamma)$, $y_i \in D_i^R(p)$ if and only if $R_{\leftarrow}(y_i) \in D_i(R_{\leftarrow}(p))$.*

Proof. Lemma 7.9.2 states that $y'_i \cdot R_{\leftarrow}(p)_i = R(y'_i) \cdot p$ for any bundle $y'_i \sim \Gamma$. This implies the following set equivalence:

$$\{y'_i \mid y'_i \sim \Gamma \text{ and } y'_i \cdot R_{\leftarrow}(p)_i \leq B_i\} = \{y'_i \mid y'_i \sim \Gamma \text{ and } R(y'_i) \cdot p \leq B_i\}$$

Next, recall that for any bundle $y'_i \sim \Gamma$, $R_{\leftarrow}(R(y'_i)) = y'_i$, so

$$\{y'_i \mid y'_i \sim \Gamma \text{ and } y'_i \cdot R_{\leftarrow}(p)_i \leq B_i\} = \{R_{\leftarrow}(R(y'_i)) \mid y'_i \sim \Gamma \text{ and } R(y'_i) \cdot p \leq B_i\}$$

For every bundle $y'_i \sim \Gamma$, $R(y'_i)$ is a bundle in $R(\Gamma)$. Therefore we can replace $R(y'_i)$ with y'_i and get

$$\{R_{\leftarrow}(R(y'_i)) \mid y'_i \sim \Gamma \text{ and } R(y'_i) \cdot p \leq B_i\} \subseteq \{R_{\leftarrow}(y'_i) \mid y'_i \sim R(\Gamma) \text{ and } y'_i \cdot p \leq B_i\}$$

Note that the relationship is now subset instead of equality. This is because there may be some $y'_i \sim R(\Gamma)$ that does not equal $R(y''_i)$ for any $y''_i \sim \Gamma$. Combining this subset relationship with the previous equality gives us

$$\{y'_i \mid y'_i \sim \Gamma \text{ and } y'_i \cdot R_{\leftarrow}(p)_i \leq B_i\} \subseteq \{R_{\leftarrow}(y'_i) \mid y'_i \sim R(\Gamma) \text{ and } y'_i \cdot p \leq B_i\}$$

Since $u_i(R_{\leftarrow}(y'_i)) = u_i^R(y'_i)$ by definition, we have

$$\{u_i(y'_i) \mid y'_i \sim \Gamma \text{ and } y'_i \cdot R_{\leftarrow}(p)_i \leq B_i\} \subseteq \{u_i^R(y'_i) \mid y'_i \sim R(\Gamma) \text{ and } y'_i \cdot p \leq B_i\}$$

Taking the max gives us

$$\max\Big(\{u_i(y'_i) \mid y'_i \sim \Gamma \text{ and } y'_i \cdot R_{\leftarrow}(p)_i \leq B_i\}\Big) \leq \max\Big(\{u_i^R(y'_i) \mid y'_i \sim R(\Gamma) \text{ and } y'_i \cdot p \leq B_i\}\Big)$$

which we can rewrite as

$$\max_{\substack{y'_i \sim \Gamma: \\ y'_i \cdot R_{\leftarrow}(p)_i \leq B_i}} u_i(y'_i) \leq \max_{\substack{y'_i \sim R(\Gamma): \\ y'_i \cdot p \leq B_i}} u_i^R(y'_i)$$

Consider an arbitrary $y''_i \in D_i^R(p)$: then

$$u_i^R(y''_i) = \max_{\substack{y'_i \sim R(\Gamma): \\ y'_i \cdot p \leq B_i}} u_i^R(y'_i)$$

and $y''_i \cdot p \leq B_i$. Then by Lemma 7.9.4, $R_{\leftarrow}(y''_i) \cdot R_{\leftarrow}(p)_i = y''_i \cdot p \leq B_i$. Since $R_{\leftarrow}(y''_i) \sim \Gamma$ and $R_{\leftarrow}(y''_i) \cdot R_{\leftarrow}(p)_i \leq B_i$, we have

$$\max_{\substack{y'_i \sim \Gamma: \\ y'_i \cdot R_{\leftarrow}(p)_i \leq B_i}} u_i(y'_i) \geq u_i(R_{\leftarrow}(y''_i))$$

By definition, $u_i^R(y''_i) = u_i(R_{\leftarrow}(y''_i))$, so

$$\max_{\substack{y'_i \sim R(\Gamma): \\ y'_i \cdot p \leq B_i}} u_i^R(y'_i) = u_i^R(y''_i) = u_i(R_{\leftarrow}(y''_i)) \leq \max_{\substack{y'_i \sim \Gamma: \\ y'_i \cdot R_{\leftarrow}(p)_i \leq B_i}} u_i(y'_i) \leq \max_{\substack{y'_i \sim R(\Gamma): \\ y'_i \cdot p \leq B_i}} u_i^R(y'_i)$$

Therefore,

$$\max_{\substack{y'_i \sim \Gamma: \\ y'_i \cdot R_{\leftarrow}(p)_i \leq B_i}} u_i(y'_i) = \max_{\substack{y'_i \sim R(\Gamma): \\ y'_i \cdot p \leq B_i}} u_i^R(y'_i)$$

Finally, suppose $y_i \in D_i^R(p)$: then $y_i \cdot p \leq B_i$, and by Lemma 7.9.4 we have $R_{\leftarrow}(y_i) \cdot R_{\leftarrow}(p)_i = y_i \cdot p \leq B_i$. Also,

$$u_i(R_{\leftarrow}(y_i)) = u_i^R(y_i) = \max_{\substack{y'_i \sim R(\Gamma): \\ y'_i \cdot p \leq B_i}} u_i^R(y'_i) = \max_{\substack{y'_i \sim \Gamma: \\ y'_i \cdot R_{\leftarrow}(p)_i \leq B_i}} u_i(y'_i)$$

so $R_{\leftarrow}(y_i) \in D_i(R_{\leftarrow}(p))$. Suppose $R_{\leftarrow}(y_i) \in D_i(R_{\leftarrow}(p))$: then $R_{\leftarrow}(y_i) \cdot R_{\leftarrow}(p)_i \leq B_i$. Since $y_i = R(R_{\leftarrow}(y_i))$, we have $y_i \cdot p = R_{\leftarrow}(y_i) \cdot R_{\leftarrow}(p)_i \leq B_i$. Also,

$$u_i^R(y_i) = u_i(R_{\leftarrow}(y_i)) = \max_{\substack{y'_i \sim \Gamma: \\ y'_i \cdot R_{\leftarrow}(p)_i \leq B_i}} u_i(y'_i) = \max_{\substack{y'_i \sim R(\Gamma): \\ y'_i \cdot p \leq B_i}} u_i^R(y'_i)$$

Therefore $y_i \in D_i^R(p)$.

□

Recall that for any bundle $y_i \sim \Gamma$, $R_{\leftarrow}(R(y_i)) = y_i$. Thus by Lemma 7.9.5, $R(y_i) \in D_i^R(p)$ if and only if $y_i = R_{\leftarrow}(R(y_i)) \in D_i(p)$. This is expressed by Corollary 7.9.5.1, which will be useful in Section 7.5 when considering tâtonnement processes in the PDM.

Corollary 7.9.5.1 (of Lemma 7.9.5). Given prices $p \sim R(\Gamma)$ and a bundle $y_i \sim \Gamma$, $y_i \in D_i(R_{\leftarrow}(p))$ if and only if $R(y_i) \in D_i^R(p)$.

Theorem 7.4.1. *For an allocation* $\mathbf{y} \sim R(\Gamma)$ *and prices* $p \sim R(\Gamma)$*,* $(\mathbf{y}, p)$ *is a ME of the market* $R(\Gamma)$ *if and only if* $(R^b_{\leftarrow}(\mathbf{y}), R^p_{\leftarrow}(p))$ *is a PME of the PDM* Γ*.*

Proof. ($\Longrightarrow$) Suppose $(\mathbf{y}, p)$ is a ME of the Fisher market $R(\Gamma)$: then $y_i \in D_i^R(p)$ for all $i \in N$, and $\sum_{i\in N} y_{i\ell} = 1$ or $p_\ell = 0$ for all $\ell \in R(M)$. By Lemma 7.9.5, we have $R_\leftarrow(y_i) \in D_i(R_\leftarrow(p))$.

We define $\mathbf{x} = (x^1 ... x^m) \in [0,1]^{m\times 2}$ as follows:

$$x^{j,0} = \max_{i\in N: a_{ij}=0} R_\leftarrow(y_i)_j \tag{7.4}$$
$$x^{j,1} = 1 - x^{j,0}$$

for all $j \in M$. We claim that for all $i \in N$ and $j \in M$, $R_\leftarrow(y_i)_j \le x^{j,a_{ij}}$, and that $R_\leftarrow(y_i)_j = x^{j,a_{ij}}$ if $R_\leftarrow(p)_{ij} > 0$.

We first show that $R_\leftarrow(y_i)_j \le x^{j,a_{ij}}$ for all i, j. When $a_{ij} = 0$, this is true by definition, so assume $a_{ij} = 1$. Since $(\mathbf{y}, p)$ is a ME of $R(\Gamma)$, for any $\ell \in R(M)$, we have $\sum_{k'\in N} y_{k'\ell} \le 1$. Thus for any $k \in N$ where $a_{ij} \ne a_{kj}$ (i.e. $a_{kj} = 0$), we have $\sum_{k'\in N} y_{k'(ikj)} \le 1$.

Also, recall that by definition, $R_\leftarrow(y_i)_j = \min_{k\in N: a_{ij}\ne a_{kj}} y_{i(ikj)}$. Thus $R_\leftarrow(y_i)_j \le y_{i(ikj)}$ for all k. Similarly, $R_\leftarrow(y_k)_j \le y_{k(ikj)}$. Therefore

$$R_\leftarrow(y_i)_j \ + R_\leftarrow(y_k)_j \le y_{i(ikj)} + y_{k(ikj)} \le \sum_{k'\in N} y_{k'(ikj)} \le 1 \qquad \forall k \in N: \ a_{kj} = 0 \tag{7.5}$$

$$R_\leftarrow(y_i)_j \ + \max_{k\in N: a_{kj}=0} R_\leftarrow(y_k)_j \le 1 \tag{7.6}$$

$$R_\leftarrow(y_i)_j \ \le 1 - \max_{k\in N: a_{kj}=0} R_\leftarrow(y_k)_j = x^{j,1} \tag{7.7}$$

Thus $R_\leftarrow(y_i)_j \le x^{j,a_{ij}}$ for all $i \in N$. It remains to show that $R_\leftarrow(y_i)_j = x^{j,a_{ij}}$ whenever $R_\leftarrow(p)_{ij} > 0$.

Suppose for sake of contradiction there exists $i \in N$ and $j \in M$ where $R_\leftarrow(y_i)_j < x^{j,a_{ij}}$ and $R_\leftarrow(p)_{ij} > 0$. Since $R_\leftarrow(y_i)_j = \min\limits_{k\in N: a_{ij}\ne a_{kj}} y_{i(ikj)}$, there must exist $k \in N$ where $a_{ij} \ne a_{kj}$ and $y_{i(ikj)} < x^{j,a_{ij}}$. Since $R_\leftarrow(p)_{ij} = \sum\limits_{k'\in N: a_{ij}\ne a_{k'j}} p_{(ik'j)}$, there must exist k' where $p_{(ik'j)} > 0$.

If $y_{i(ik'j)} > y_{i(ikj)}$, we have $p_{(ik'j)} = 0$ by Lemma 7.9.3. Thus assume $y_{i(ik'j)} \le y_{i(ikj)} < x^{j,a_{ij}}$. We showed above that $R_\leftarrow(y_{k'})_j \le x^{j,a_{k'j}} = 1 - x^{j,a_{ij}}$: thus $y_{k'(ik'j)} \le 1 - x^{j,a_{ij}}$. Therefore $y_{i(ik'j)} + y_{k'(ik'j)}) < x^{j,a_{ij}} + 1 - x^{j,a_{ij}} = 1$. If there exists $i' \notin i, k'$ where $y_{i'(ik'j)} > 0$, then $p_{(ik'j)} = 0$, which is a contradiction. Therefore $\sum_{\ell\in R(M)} y_{i\ell} = y_{i(ik'j)} + y_{k'(ik'j)}) < 1$.

But then by the definition of a ME, we have $p_{(ik'j)} = 0$, which is again a contradiction. Therefore $R_\leftarrow(y_i)_j = x^{j,a_{ij}}$ whenever $R_\leftarrow(p)_{ij} > 0$. This shows that $(R_\leftarrow(y), R_\leftarrow(p))$ is a PME of Γ.

($\Longleftarrow$) Suppose $(R_\leftarrow(\mathbf{y}), R_\leftarrow(p))$ is a PME of the PDM Γ. Then $R_\leftarrow(y_i) \in D_i(R_\leftarrow(p))$ for all $i \in N$, so $y_i \in D_i^R(p)$ by Lemma 7.9.5. Also, there exists $\mathbf{x} = (x^1 ... x^m) \in [0,1]^{m\times 2}$ where for all $i \in N$ and $j \in M$,

1. $x^{j,0} + x^{j,1} = 1$
2. $R_\leftarrow(y_i)_j \le x^{j,a_{ij}}$
3. $R_\leftarrow(y_i)_j = x^{j,a_{ij}}$ whenever $R_\leftarrow(p)_{ij} = 0$.

It remains to show that for all $\ell \in R(M)$, either $\sum_{i\in N} y_{i\ell} = 1$ or $p_\ell = 0$. Suppose for sake of contradiction that there exists $\ell = (i,k,j) \in R(M)$ where $\sum_{k'\in N} y_{k'(ikj)} < 1$ and $p_{(ikj)} > 0$. If there exists $k' \notin \{i,k\}$ where $y_{k'(ikj)} > 0$, then $p_{(ikj)} = 0$ by Lemma 7.9.3. Thus

$$\sum_{k'\in N} y_{k'(ikj)} = y_{i(ikj)} + y_{k(ikj)} < 1$$

Furthermore, if either $y_{i(ikj)} \neq R_{\leftarrow}(y_i)_j$ or $y_{k(ikj)} \neq R_{\leftarrow}(y_k)_j$, we have $p_{(ikj)} = 0$ again by Lemma 7.9.3. Thus $y_{i(ikj)} = R_{\leftarrow}(y_i)_j$ and $y_{k(ikj)} = R_{\leftarrow}(y_k)_j$, so

$$R_{\leftarrow}(y_i)_j + R_{\leftarrow}(y_k)_j < 1$$

Recall that $R_{\leftarrow}(y_i)_j \leq x^{j,a_{ij}}$ and $R_{\leftarrow}(y_k)_j \leq x^{j,a_{kj}}$, and that $x^{j,a_{ij}} + x^{j,a_{kj}} = 1$ since $a_{ij} \neq a_{kj}$. Thus in order for $R_{\leftarrow}(y_i)_j + R_{\leftarrow}(y_k)_j < 1$ to be true, either $R_{\leftarrow}(y_i)_j < x^{j,a_{ij}}$ or $R_{\leftarrow}(y_k)_j < x^{j,a_{kj}}$. By symmetry, suppose $R_{\leftarrow}(y_i)_j < x^{j,a_{ij}}$ without loss of generality. Then because $(R_{\leftarrow}(\mathbf{y}), R_{\leftarrow}(p))$ is a PME, we have $R_{\leftarrow}(p)_{ij} = 0$.

By definition, $R_{\leftarrow}(p)_{ij} = \sum\limits_{k'\in N: a_{ij}\neq a_{k'j}} p_{(ik'j)}$. Since $p_{(ik'j)} \geq 0$ for all i, k', j, we have $p_{(ik'j)}$ for all $k' \in N$ where $a_{ij} \neq a_{k'j}$. But then $p_{(ikj)} = 0$, which is a contradiction. Thus for all $\ell \in R(M)$, either $\sum_{i\in N} y_{i\ell} = 1$ or $p_\ell = 0$. Therefore $(\mathbf{y}, p)$ is a ME of $R(\Gamma)$. ☐

Theorem 7.4.2. *Let Ψ be a welfare function, let Γ be the public decisions instance (N, M) with budgets $B_1 ... B_n$, and let $\alpha \geq 0$. Then $\mathbf{z}$ is an α-approximation of Ψ in Γ if and only if $R(\mathbf{z})$ is an α-approximation of Ψ in $R(\Gamma)$.*

Proof. We first claim that $R(\mathbf{x})$ is a valid allocation in $R(\Gamma)$, meaning that $\sum_{i\in N} R(x_i)_\ell \leq 1$ for all $\ell \in R(M)$. By definition of $R(x_i)$, $R(x_i)_{(kk'j)} = 0$ whenever $i \notin \{k, k'\}$, and $R(x_i)_{(kk'j)} = x_{ij}$ whenever $i \in \{k, k'\}$. Therefore, for all $\ell \in R(M)$,

$$\sum_{i\in N} R(x_i)_\ell = \sum_{i,k,k'\in N} R(x_i)_{(kk'j)} = R(x_i)_{(ikj)} + R(x_k)_{(ikj)} = x_{ij} + x_{kj}$$

By definition of $R(M)$, the fact that good (i,k,j) exists implies that $a_{ij} \neq a_{kj}$. Since $\mathbf{x}$ is a valid outcome of Γ, for all $j \in M$ we must have $x_{ij} + x_{kj} \leq 1$ whenever $a_{ij} \neq a_{kj}$. Therefore $\sum_{i\in N} R(x_i)_\ell \leq 1$. Since this holds for all $\ell \in R(M)$, $R(\mathbf{x})$ is a valid allocation in $R(\Gamma)$.

By definition of $R_{\leftarrow}$ and u_i^R, we have $u_i^R(x_i) = u_i(R_{\leftarrow}(x_i))$. Since Ψ depends only on the agents' utilities, we have $\Psi(\mathbf{x}') = \Psi(R(\mathbf{x}'))$ for any outcome $\mathbf{x}' \sim \Gamma$. Similarly, recall that $u_i(x_i') = u_i^R(R(x_i'))$ for any bundle $x_i' \sim \Gamma$, so $\Psi(\mathbf{x}') = \Psi(R_{\leftarrow}(\mathbf{x}'))$ for any outcome $\mathbf{x}' \sim R(\Gamma)$.

Thus for every possible outcome of Γ, there is an outcome of $R(\Gamma)$ which has the same value of Ψ, and for every possible outcome of $R(\Gamma)$, there is an outcome of Γ which has the same value of Ψ. Therefore we have the numeral equality

$$\max_{\mathbf{x}'\sim\Gamma} \Psi(\mathbf{x}') = \max_{\mathbf{x}'\sim R(\Gamma)} \Psi(\mathbf{x}')$$

Finally, because $\Psi(\mathbf{x}) = \Psi(R(\mathbf{x}))$, we have $\Psi(\mathbf{x}) \geq \alpha \cdot \max_{\mathbf{x}' \sim \Gamma} \Psi(\mathbf{x}')$ if and only if $\Psi(R(\mathbf{x})) \geq \alpha \cdot \max_{\mathbf{x}' \sim R(\Gamma)} \Psi(\mathbf{x}')$. Therefore $\mathbf{x}$ is an α-approximation of Ψ in Γ if and only if $R(\mathbf{x})$ is an α-approximation of Ψ in $R(\Gamma)$. □

7.10 Other omitted proofs

7.10.1 Omitted proofs from Section 7.3

Theorem 7.3.1. *For a PDM (N, M, B) with linear utilities given by weights $w_{ij} \geq 0$, for every list of private bundles $\mathbf{y}$ and list of prices p, $(\mathbf{y}, p)$ is an IME if and only $(\mathbf{y}, p)$ is a ME for the Fisher market (N, M, B) with linear utilities given by the same weights.*

Proof. Let Γ be the PDM (N, M, B) with linear utilities u_i given by weights w_{ij}, and $\tilde{\Gamma}$ be the Fisher market (N, M, B) with linear utilities $\tilde{u}_i$ given by the same weights.

Let $(\mathbf{y}, p)$ be an IME of Γ: then $y_i \in D_i(p, y_{-i})$ for all i. Let x_i be agent i's public bundle in $\mathbf{y}$, let $\mathbf{y}' = (y_{-i}, y'_i)$ for an arbitrary private bundle y'_i, and let x'_i be agent i's public bundle for private bundles $\mathbf{y}'$. Then we have

$$
\begin{aligned}
u_i(y_{-i}, y_i) &= \max_{y'_i:\ y'_i \cdot p \leq B_i} u_i(y_{-i}, y'_i) \\
\sum_{j \in M} w_{ij} x_{ij} &= \max_{y'_i:\ y'_i \cdot p \leq B_i} \sum_{j \in M} w_{ij} x'_{ij} \\
\sum_{j \in M} w_{ij} \sum_{\substack{k \in N: \\ a_{kj} = a_{ij}}} y_{kj} &= \max_{y'_i:\ y'_i \cdot p \leq B_i} \left(\sum_{j \in M} w_{ij} y'_{ij} + \sum_{j \in M} w_{ij} \sum_{\substack{k \in N \setminus \{i\}: \\ a_{kj} = a_{ij}}} y_{kj} \right) \\
\sum_{j \in M} w_{ij} y_{ij} + \sum_{j \in M} w_{ij} \sum_{\substack{k \in N \setminus \{i\}: \\ a_{kj} = a_{ij}}} y_{kj} &= \max_{y'_i:\ y'_i \cdot p \leq B_i} \left(\sum_{j \in M} w_{ij} y'_{ij} + \sum_{j \in M} w_{ij} \sum_{\substack{k \in N \setminus \{i\}: \\ a_{kj} = a_{ij}}} y_{kj} \right) \\
\sum_{j \in M} w_{ij} y_{ij} + \sum_{j \in M} w_{ij} \sum_{\substack{k \in N \setminus \{i\}: \\ a_{kj} = a_{ij}}} y_{kj} &= \max_{y'_i:\ y'_i \cdot p \leq B_i} \left(\sum_{j \in M} w_{ij} y'_{ij} \right) + \sum_{j \in M} w_{ij} \sum_{\substack{k \in N \setminus \{i\}: \\ a_{kj} = a_{ij}}} y_{kj} \\
\sum_{j \in M} w_{ij} y_{ij} &= \max_{y'_i:\ y'_i \cdot p \leq B_i} \sum_{j \in M} w_{ij} y'_{ij} \\
\tilde{u}_i(y_i) &= \max_{y'_i:\ y'_i \cdot p \leq B_i} \tilde{u}_i(y'_i)
\end{aligned}
$$

Also, the total price of y_i is $y_i \cdot p$ in both Γ and $\tilde{\Gamma}$. Let $\tilde{D}_i(p)$ be agent i's demand set for prices p in $\tilde{\Gamma}$: then by the above chain of equations, if $y_i \in D_i(p, y_{-i})$ for any y_{-i}, we have $y_i \in \tilde{D}_i(p)$. Furthermore, the exact same chain of equations in reverse order shows that if $y_i \in \tilde{D}_i(p)$, then $y_i \in D_i(p, y_{-i})$ for all y_{-i}.

Since $y_i \in \tilde{D}_i(p)$, the allocation $\mathbf{y}$ in $\tilde{\Gamma}$ gives each agent a bundle in her demand set given prices p. Also, because $(\mathbf{y}, p)$ is an IME of Γ, we have that $\sum_{i \in N} y_{ij} \leq 1$, and $\sum_{i \in N} y_{ij} = 1$ whenever $p_j > 0$. Therefore $(\mathbf{y}, p)$ is a ME of $\tilde{\Gamma}$.

Now let $(\mathbf{y}, p)$ be a ME of $\tilde{\Gamma}$. Since $y_i \in \tilde{D}_i(p)$ implies $y_i \in D_i(p, y_{-i})$, we have that $\mathbf{y}$ in Γ gives each agent a bundle in her demand set. By the definition of ME, we have $\sum_{i \in N} y_{ij} \leq 1$, and $\sum_{i \in N} y_{ij} = 1$ whenever $p_j > 0$. Therefore $(\mathbf{y}, p)$ is an IME of Γ. □

Theorem 7.3.2. *For any $\epsilon > 0$, $\Phi(n, 1+\epsilon)$ with linear utilities has a unique equilibrium $(\mathbf{y}, p)$, where*

$$\frac{\max_{\mathbf{z}'} NW(\mathbf{z}')}{NW(\mathbf{y})} \geq \frac{n-1}{1+\epsilon}$$

Proof. Let $\Phi'(n, 1+\epsilon)$ be the Fisher market with the same agents, goods, and weights as $\Phi(n, 1+\epsilon)$, also with linear utilities. Let $(\mathbf{y}, p)$ be a ME of $\Phi'(n, 1+\epsilon)$. Then $y_i \cdot p = B_i = 1$ for all i, and so $\sum_{j \in M} p_j = \mathcal{B} = n$.

We next observe for a Fisher market with linear utilities, any (private) bundle in an agent's demand set maximizes her "bang-per-buck" ratio: w_{ij}/p_j. To see this, consider agent i moving δ of her budget to a good that does not maximize her bang-per-buck: this would decrease her utility, and so that bundle cannot be in her demand set.

Suppose for sake of contradiction that there exists ℓ where $p_\ell \neq 1$. Since $\sum_{j \in M} p_j = n$, there must exist ℓ where $p_\ell < 1$. Let $\ell = \min_{j \in M} p_j$. Since $w_{\ell\ell} > w_{\ell j}$ for all $j \neq \ell$, only issue ℓ maximizes agent ℓ's bang-per-buck. Thus there is a single bundle y_i in her demand set, and it consists of her spending her entire budget on issue ℓ. But since $p_\ell < 1 = B_\ell$, agent ℓ purchases more of good ℓ than exists in the supply, and so the market cannot clear. Thus any ME of $\Phi'(n, 1+\epsilon)$ must have $p_j = 1$ for all j.

Now assume that $p_j = 1$ for all j. Since for each agent i, $w_{ii} = w > 1 = w_{ij}$ for all $j \neq i$, the only bundle that maximizes agent i's bang-per-buck consists of her spending her entire budget on issue i. Thus the unique ME is $(\mathbf{y}, p)$ where $y_{ii} = 1$ for all i, and $y_{ij} = 0$ whenever $i \neq j$. Furthermore, by Theorem 7.3.1, $(\mathbf{y}, p)$ is the unique IME of $\Phi(n, 1+\epsilon)$.

The Nash welfare of $\mathbf{y}$ in $\Phi(n, 1+\epsilon)$ is

$$NW(\mathbf{y}) = \Big(\prod_{i \in N} \sum_{j \in M} w_{ij} y_{ij}\Big)^{1/n} = \Big(\prod_{i \in N} 1+\epsilon\Big)^{1/n} = 1+\epsilon$$

Now consider the outcome $\mathbf{z}$ where $z^{j,0} = 0$ and $z^{j,1} = 1$ for all $j \in M$. Let x_{ij} be agent i's public bundle, as usual. Then

$$NW(\mathbf{z}) = \Big(\prod_{i \in N} \sum_{j \in M} w_{ij} x_{ij}\Big)^{1/n} = \Big(\prod_{i \in N} \sum_{j \in M \setminus \{i\}} 1 \cdot x_{ij}\Big)^{1/n} = \Big(\prod_{i \in N} (n-1)\Big)^{1/n} = n-1$$

and therefore

$$\frac{\max_{\mathbf{z}'} NW(\mathbf{z}')}{NW(\mathbf{y})} \geq \frac{n-1}{1+\epsilon}$$

□

We now present the proofs of Theorems 7.3.3 and 7.3.4, which state that issue-pricing equilibria can be inefficient for Cobb-Douglas and CES utilities, respectively.

We first prove a lemma motivated by the following concept. In Section 7.3.2, we described how for linear utilities, the bundles in an agent's demand set maximize her bang-per-buck, in both the public and private settings. This is not true in general for other utilities, since goods are not independent. However, the same concept still applies: agent i will not spend any money on issue j if there is another issue ℓ where she has a higher marginal utility per dollar spent on issue ℓ. This concept will be made formal by examining $\frac{\partial(u_i(x_i))}{\partial x_{ij}}$, which is the partial derivative of agent i's utility with respect to x_{ij}. Although these derivatives may be complicated in general, they are well-behaved for Cobb-Douglas and CES utilities with $\rho \in (-\infty, 0) \cup (0, 1)$.

We will use this concept to show that for Cobb-Douglas utilities, $x_{ij}p_j = \min_{\ell \in M} x_{i\ell}p_\ell$ for any issue j that agent i is spending any money on. For CES utilities with $\rho \in (-\infty, 0) \cup (0, 1)$, we will show that $x_{ij}^{1-\rho}p_j = \min_{\ell \in M} x_{i\ell}^{1-\rho}p_\ell$ (note that $1 - \rho > 0$ since $\rho \in (-\infty, 0) \cup (0, 1)$). Using these two properties, the following lemma will imply that $x_{ij} = 1/2$ for all j, which allows us to compute the Nash welfare.

Lemma 7.10.1. *Let* $(\mathbf{y}, p)$ *be an IME of* $\Phi(n, 1)$ *and let* x_i *be agent* i*'s public bundle as induced by* $\mathbf{y}$*. Suppose that there exists* $c > 0$ *such that for every issue* j *that agent* i *spends any money on,* $x_{ij}^c p_j = \min\limits_{\ell \in M} x_{i\ell}^c p_\ell$*. Then* $x_{ij} = 1/2$ *for all* i *and* j*.*

Proof. The majority of the proof will be dedicated to showing that for every agent i, there must exist an issue j where $x_{ij}^c p_j \leq 1/2^c$. Suppose for sake of contradiction that there exists an agent i where $x_{ij}^c p_j > 1/2^c$ for every issue j.

We first show that there must exist an agent k and issue j where $x_{kj}^c p_j < 1/2$. Because $(\mathbf{y}, p)$ is an IME, all agents exhaust their budgets, so $\sum_{j \in M} p_j = \sum_{k \in N} B_k = n$. Because $|M| = n$ here, there must exist $j \in M$ where $p_j \leq 1$. Since $x_{ij}^c p_j > 1/2^c$, we have $x_{ij} > 1/2$. Let k be any agent where $a_{kj} \neq a_{ij}$: then $x_{kj} < 1/2$, and so $x_{kj}^c p_j < 1/2^c$.

We know by definition of $\Phi(n, 1)$, agents i and k agree on all issues other than i and k: $a_{i\ell} = a_{k\ell}$ whenever $\ell \notin \{i, k\}$. Thus for all issues $\ell \notin \{i, k\}$, $x_{k\ell}^c p_\ell > 1/2^c > x_{kj}^c p_j$. By assumption, agent k only spends money on issues ℓ which minimize $x_{k\ell}^c p_\ell$. Thus agent k does not spend money on any issues besides i and k (note that either $j = i$ or $j = k$).

Therefore amount of money agent k spends in total is $\sum_{\ell \in M} y_{k\ell}p_\ell = y_{kk}p_k + y_{ki}p_i$. Since agent k exhausts her budget, we have $y_{kk}p_k + y_{ki}p_i = B_k = 1$. Thus there must exist $\ell \in \{k, i\}$ where $y_{k\ell}p_\ell \geq 1/2$. Therefore $x_{k\ell}p_\ell \geq 1/2$.

Since $a_{kk} \neq a_{ik}$ and $a_{ki} \neq a_{ii}$, we have $a_{k\ell} \neq a_{i\ell}$. Because $(\mathbf{y}, p)$ is an IME, we have $x_{i\ell} = 1 - x_{k\ell}$. Also, since agent k spends money on issue ℓ, we have $x_{k\ell}^c p_\ell \leq x_{kj}^c p_j < 1/2^c$ by assumption. Therefore

$$\begin{aligned}
x_{k\ell}^c p_\ell &< 1/2^c < x_{i\ell}^c p_\ell \\
x_{k\ell}^c p_\ell &< (1 - x_{k\ell})^c p_\ell \\
x_{k\ell}^c &< (1 - x_{k\ell})^c \\
x_{k\ell} &< (1 - x_{k\ell}) \\
x_{k\ell} &< 1/2
\end{aligned}$$

Since agent k exhausts her budget, we have $y_{kk}p_k + y_{ki}p_i = B_k = 1$, and so $x_{kk}p_k + x_{ki}p_i = 1$. Thus there must exist $\ell \in \{i, k\}$ where $y_{k\ell}p_\ell \geq 1/2$. Because $x_{k\ell}^c p_\ell < 1/2^c$, we have $x_{k\ell}^{1-c}/2^c > x_{k\ell}p_\ell \geq 1/2$. Therefore

$$\begin{aligned} \frac{x_{k\ell}^{1-c}}{2^c} &> \frac{1}{2} \\ \frac{(1/2)^{1-c}}{2^c} &> \frac{1}{2} \\ 1 &> 1 \end{aligned}$$

which is a contradiction. Therefore for every agent i, there exists an issue j where $x_{ij}^c p_j \leq 1/2^c$.

By assumption, if $x_{ij}^c p_j > \min_{\ell \in M} x_{i\ell}^c p_\ell$, then agent i spends no money on issue j. Since there exists an issue j where $x_{ij}^c p_j \leq 1/2^c$, we have that agent i spends no money on any issue j where $x_{ij}^c p_j > 1/2^c$.

Suppose for sake of contradiction that an agent i and issue j exist where $x_{ij}^c p_j > 1/2^c$: then some agent k where $a_{kj} = a_{ij}$ is spending money on issue j. But since $a_{kj} = a_{ij}$, we have $x_{kj}^c p_k > 1/2^c$, so agent k cannot be spending any money on issue j. Therefore for every agent i and every issue j, $x_{ij}^c p_j \leq 1/2^c$.

Suppose for sake of contradiction that there exists an issue j where $p_j \neq 1$. Since $\sum_{\ell \in M} p_\ell = n$, there must exist an issue ℓ where $p_\ell > 1$. Let k be any other agent other than ℓ: then $a_{k\ell} \neq a_{\ell\ell}$. Since $x_{k\ell} + x_{\ell\ell} = 1$, we have $\max(x_{k\ell}, x_{\ell\ell}) \geq 1/2$. Therefore $(\max(x_{k\ell}, x_{\ell\ell}))^c p_\ell > 1/2^c$, which is a contradiction. Therefore $p_j = 1$ for all j.

Finally, suppose there exists an agent i and issue j where $x_{ij} \neq 1/2$, there must exist an agent k where $x_{kj} > 1/2$. Then $x_{kj}^c p_j > 1/2^c$, which is again a contradiction. Therefore for every agent i and issue j, $x_{ij} = 1/2$. □

We are now ready to prove Theorems 7.3.3 and 7.3.4.

Theorem 7.3.3. *For any IME* $(\mathbf{y}, p)$ *of* $\Phi(n, 1)$ *with Cobb-Douglas utilities,*

$$\frac{\max_{\mathbf{z}'} NW(\mathbf{z}')}{NW(\mathbf{y})} \geq \frac{2 - 2/n}{(n-1)^{1/n}}$$

Proof. Let x_i be agent i's public bundle as induced by $\mathbf{y}$. Recall that a Cobb-Douglas utility is given by

$$u_i(\mathbf{y}) = u_i(x_i) = \Big(\prod_{j \in M} x_{ij}^{w_{ij}}\Big)^{1/\sum_{j \in M} w_{ij}}$$

which for $\Phi(n, 1)$, simplifies to

$$u_i(x_i) = \Big(\prod_{j \in M} x_{ij}\Big)^{1/n}$$

Thus for all j, we have

$$\frac{1}{p_j}\frac{\partial(u_i(\mathbf{y}))}{\partial x_{ij}} = \frac{x_{ij}^{\frac{1}{n}-1}}{p_j n}\Big(\prod_{\ell\in M\setminus\{j\}} x_{i\ell}\Big)^{1/n} = \frac{1}{x_{ij}p_j n}\Big(\prod_{\ell\in M} x_{i\ell}\Big)^{1/n} = \frac{1}{x_{ij}p_j n}u_i(x_i)$$

We are going to invoke Lemma 7.10.1 with $c = 1$. Suppose that there exists an agent i and issues j, ℓ such that agent i is spending on issue j, but $x_{ij}p_j > x_{i\ell}p_\ell$. Then

$$\frac{1}{p_j}\frac{\partial(u_i(\mathbf{y}))}{\partial x_{ij}} < \frac{1}{p_\ell}\frac{\partial(u_i(\mathbf{y}))}{\partial x_{i\ell}}$$

Thus there exists some $\delta > 0$ such that if agent i spent δ less on issue j and δ more on issue ℓ, agent i's utility would increase. But x_i is in agent i's demand set, so it cannot be possible for her to increase her utility while staying within her budget. This is a contradiction, so therefore $x_{ij}p_j = \min_{\ell\in M} x_{i\ell}p_\ell$ for all i, j.

Therefore by Lemma 7.10.1, we have $x_{ij} = 1/2$ for all i and j. So the Nash welfare of $\mathbf{y}$ is

$$NW(\mathbf{y}) = \Big(\prod_{i\in N}\Big(\prod_{j\in M} x_{ij}\Big)^{1/n}\Big)^{1/n} = \Big(\prod_{i\in N}\Big(\prod_{j\in M} 1/2\Big)^{1/n}\Big)^{1/n} = \Big(\prod_{i\in N} 1/2\Big)^{1/n} = 1/2$$

Consider the outcome $\mathbf{z}'$ where $x'_{ii}(\mathbf{z}') = 1/n$ for all i, and $x'_{ij}(\mathbf{z}') = \dfrac{n-1}{n}$ whenever $j \neq i$. Then

$$NW(\mathbf{z}') = \Big(\prod_{i\in N}\Big(\prod_{j\in M} x'_{ij}\Big)^{1/n}\Big)^{1/n} = \Big(\prod_{i\in N}\Big(\frac{1}{n}\Big(\frac{n-1}{n}\Big)^{n-1}\Big)^{1/n}\Big)^{1/n} = \Big(\frac{1}{n}\Big(\frac{n-1}{n}\Big)^{n-1}\Big)^{1/n}$$

$$= \Big(\frac{1}{n-1}\Big(\frac{n-1}{n}\Big)^{n}\Big)^{1/n} = \frac{1}{(n-1)^{1/n}}\frac{n-1}{n} = \frac{1-1/n}{(n-1)^{1/n}}$$

Therefore

$$\frac{\max_{\mathbf{z}'} NW(\mathbf{z}')}{NW(\mathbf{y})} \geq \frac{2-2/n}{(n-1)^{1/n}}$$

□

Theorem 7.3.4. *For any IME* $(\mathbf{y}, p)$ *of* $\Phi(n, 1)$ *with CES utilities for parameter* $\rho \in (-\infty, 0)\cup(0, 1)$*,*

$$\frac{\max_{\mathbf{z}'} NW(\mathbf{z}')}{NW(\mathbf{y})} \geq 2(1-1/n)^{1/\rho}$$

Proof. Let x_i be agent i's public bundle as induced by $\mathbf{y}$. Recall that a CES utility is given by

$$u_i(\mathbf{y}) = u_i(x_i) = \Big(\sum_{j\in M} w_{ij}^\rho x_{ij}^\rho\Big)^{1/\rho} = \Big(\sum_{j\in M} x_{ij}^\rho\Big)^{1/\rho}$$

and so we have

$$\frac{1}{p_j}\frac{\partial(u_i(\mathbf{y}))}{\partial x_{ij}} = \frac{1}{p_j}\frac{1}{\rho}\rho x_{ij}^{\rho-1}\Big(\sum_{j\in M} x_{ij}^{\rho}\Big)^{\frac{1}{\rho}-1} = \frac{1}{x_{ij}^{1-\rho}p_j}u_i(x_i)^{\rho(\frac{1}{\rho}-1)}$$

This time, we are going to invoke Lemma 7.10.1 with $c = 1-\rho$ (note that since $\rho \in (-\infty, 0)\cup(0,1)$, we have $1-\rho > 0$). Suppose that there exists an agent i and issues j, ℓ such that agent i is spending on issue j, but $x_{ij}^{1-\rho}p_j > x_{i\ell}^{1-\rho}p_\ell$. Then

$$\frac{1}{p_j}\frac{\partial(u_i(\mathbf{y}))}{\partial x_{ij}} < \frac{1}{p_\ell}\frac{\partial(u_i(\mathbf{y}))}{\partial x_{i\ell}}$$

So again there exists some $\delta > 0$ such that if agent i spent δ less on issue j and δ more on issue ℓ, agent i's utility would increase. But x_i is in agent i's demand set, so this is a contradiction for the same reason as in the previous proof. Thus $x_{ij}p_j = \min_{\ell\in M} x_{i\ell}p_\ell$ for all i, j.

Therefore, by Lemma 7.10.1, we have $x_{ij} = 1/2$ for all i and j, so the Nash welfare of $\mathbf{y}$ is

$$NW(\mathbf{y}) = \Big(\prod_{i\in N}\Big(\sum_{j\in M} x_{ij}^{\rho}\Big)^{1/\rho}\Big)^{1/n} = \Big(\prod_{i\in N}\Big(\sum_{j\in M}(1/2)^{\rho}\Big)^{1/\rho}\Big)^{1/n} = \Big(\prod_{i\in N}\frac{n^{1/\rho}}{2}\Big)^{1/n} = \frac{n^{1/\rho}}{2}$$

Consider the outcome $\mathbf{z}'$ where $x'_{ii}(\mathbf{z}') = 0$ for all i and $x'_{ij}(\mathbf{z}') = 1$ whenever $j \neq i$. Then

$$NW(\mathbf{z}') = \Big(\prod_{i\in N}\Big(\sum_{j\in M} x_{ij}^{\rho}\Big)^{1/\rho}\Big)^{1/n} = \Big(\prod_{i\in N}\Big(\sum_{j\in M\setminus\{i\}} 1\Big)^{1/\rho}\Big)^{1/n} = \Big(\prod_{i\in N}(n-1)^{1/\rho}\Big)^{1/n} = (n-1)^{1/\rho}$$

Therefore

$$\frac{\max_{\mathbf{z}'} NW(\mathbf{z}')}{NW(\mathbf{y})} \geq \frac{(n-1)^{1/\rho}}{n^{1/\rho}/2} = 2\Big(\frac{n-1}{n}\Big)^{1/\rho} = 2(1-1/n)^{1/\rho}$$

□

7.10.2 Omitted proofs from Section 7.5

Theorem 7.5.1. *Consider a Fisher market tâtonnement $\mathcal{T}$. Suppose $\mathcal{T}$ converges to a δ-equilibrium for $\mathcal{H}$-nested leontief utilities in $O(\kappa(m, n, \delta))$ time steps, where n is the number of agents and m the number of goods. Then $R_{\leftarrow}(\mathcal{T})$ converges to a 3δ-PME for the PDM with $\mathcal{H}$ utilities in $O(\kappa(n^2m, n, \delta))$.*

Proof. By the reduction defined in Section 7.4, the hidden private market has $O(n^2m)$ goods (1 copy of each good for each pair of agents who disagree on the issue) and n agents.

We next show the tâtonnement $\mathcal{T}$ is being run correctly, i.e., the sequence of prices $(p^0, p^1, p^2 ...)$, alongside some demands $(y_i^t \in D_i^R(p^t))$ converges to a δ-equilibrium. This is not trivial since $\mathcal{T}$ is run with a detour through the PDM. By Corollary 7.9.5.1, given prices $p \sim R(\Gamma)$ and a bundle $y_i \sim \Gamma$, $y_i \in D_i(R_{\leftarrow}(p)) \iff R(y_i) \in D_i^R(p)$. Thus at each step, p^{t+1} is being calculated by $g_{\mathcal{T}}$ based on valid demands in $D_i^R(p^t)$, so the sequence of prices $(p^0, p^1, p^2 ...)$ converges to a δ-equilibrium. Then, by supposition, there exist demands $y \sim R(\Gamma)$ at time T such that $(\mathbf{y}, p^T) \sim R(\Gamma)$ form a δ-equilibrium in the hidden Fisher market, for $T = O(\kappa(n^2 m, n, \delta))$.

Recall Theorem 7.4.1: $(\mathbf{y}, p)$ is a Fisher market equilibrium if and only if $(R_{\leftarrow}(\mathbf{y}), R_{\leftarrow}(p))$ is a PME. The rest of the proof involves showing that Theorem 7.4.1 holds for approximate equilibria as well, as defined. Recall that $R_{\leftarrow}(y_i)_j = \min_{\substack{k \in N: \\ a_{ij} \neq a_{kj}}} y_{i(ikj)}$ $\forall j \in M$, and $R_{\leftarrow}(p)_{ij} = \sum_{\substack{k \in N: \\ a_{ij} \neq a_{kj}}} p_{(ikj)}$ $\forall i \in N, j \in M$. We claim $(R_{\leftarrow}(y), R_{\leftarrow}(p^t))$ forms a 3δ-PME:

1. $y_i \in D_i^R(p^t) \implies R_{\leftarrow}(y_i) \in D_i(R_{\leftarrow}(p^t))$. (Lemma 7.9.5)

2. We define $\mathbf{z} = (z^1 ... z^m) \in [0,1]^{m \times 2}$ as follows, analogously to Equation (7.4) in the proof of Theorem 7.4.1[26]:

$$\begin{aligned} z^{j,0} &= \max_{i \in N: a_{ij} = 0} R_{\leftarrow}(y_i)_j \\ z^{j,1} &= \max(1 - z^{j,0}, 0) \end{aligned}$$

   Then, $\forall j \in M$,

   (a) $z^{j,0} + z^{j,1} \leq 1 + \delta$ follows from the definition and from y part of a δ-equilibrium of a Fisher market.

   (b) For all $i \in N$, $R_{\leftarrow}(y_i)_j \leq z^{j,a_{ij}} + \delta$ follows from Equations (7.5)-(7.7), with 1 replaced with $1 + \delta$ and the $=$ in line (7.7) replaced with $\leq$. Finally,

$$\begin{aligned} R_{\leftarrow}(p)_{ij} > n\delta &\implies \exists \tilde{k} \text{ s.t. } p_{(ikj)} > \delta \\ &\implies y_{i(i\tilde{k}j)} + y_{\tilde{k}(i\tilde{k}j)} > 1 - \delta \quad \text{(Condition 2 of Fisher } \delta\text{-equilibrium)} \end{aligned}$$

   By Lemma 7.9.3, $p_{(i\tilde{k}j)} > 0 \implies y_{i(i\tilde{k}j)} = \min_{\substack{k \in N: \\ a_{ij} \neq a_{\tilde{k}j}}} y_{i(i\tilde{k}j)} = R_{\leftarrow}(y_i)_j$, and $y_{\tilde{k}(i\tilde{k}j)} = R_{\leftarrow}(y_{\tilde{k}})_j$. Then

$$\begin{aligned} R_{\leftarrow}(y_i)_j &> 1 - R_{\leftarrow}(y_{\tilde{k}})_j - \delta \\ &> 1 - z^{j,a_{\tilde{k}j}} - 2\delta && \text{(First part of Condition (b)} \\ &> z^{j,a_{ij}} - 3\delta && \text{definition of } z^{j,a_{ij}}) \end{aligned}$$

   Thus, $R_{\leftarrow}(p)_{ij} > n\delta \implies R_{\leftarrow}(y_i)_j > z^{j,a_{ij}} - 3\delta$

□

[26] In the proof of Theorem 7.4.1, the definition of $z^{j,1}$ was simply $1 - z^{j,0}$. It is necessary to use $\max(1 - z^{j,0}, 0)$ here instead: because this is an approximate equilibrium, it is possible that $z^{j,0} > 1$, which would make $1 - z^{j,0}$ negative.

Chapter 8

A new fairness axiom for public decision-making: equality of power

Ronald Dworkin's *equality of resources* [71], and the closely related concept of envy-freeness, are two of the fundamental axiomatic ideas behind fair allocation of private goods. The appropriate analog to these concepts in a public decision-making environment is unclear, since all agents consume the same "bundle" of resources (though they may have different utilities for this bundle). Drawing inspiration from equality of resources and the Dworkin quote below, we propose that equality in public decision-making should allow each agent to cause equal cost to the rest of society, which we model as equal externality. We term this *equality of power*. The first challenge here is that the cost to the rest of society must be measured somehow, and it is generally impossible to elicit the scale of individual utilities (in the absence of monetary payments). Again drawing inspiration from foundational literature for private goods economies, we normalize each agent's utility so that every agent's marginal utility for additional power is the same. We show that for quadratic utilities, in the large market limit, there always exists an outcome that simultaneously satisfies equal power, equal marginal utility for additional power, and utilitarian welfare maximization with respect to the normalized utilities.

> "*Equality of resources supposes that the resources devoted to each person's life should be equal. That goal needs a metric. The auction proposes what the envy test in fact assumes, that the true measure of the social resources devoted to the life of one person is fixed by asking how important, in fact, that resource is for others. It insists that the cost, measured in that way, figure in each person's sense of what is rightly his and in each person's judgment of what life he should lead, given that command of justice.*"
>
> Ronald Dworkin, *What is Equality? Part II: Equality of Resources*, 1981

8.1 Introduction

In settings where monetary payments are not allowed, it is generally impossible to elicit the absolute scale of agents' utilities. This makes most welfare objectives [1] difficult to maximize. Instead, it is common to focus on some notion of equality or fairness. In the context of pure private goods economies, this is commonly represented, in both analytic philosophy and economics, by the closely related ideas of the envy-freeness [85], competitive equilibrium from equal incomes [169] and equality of resources [71]. There has been a recent surge of interest in these topics – and more generally, axiomatic fair division of resources – in the computational economics community as well.

It is not immediately clear how to adapt these concepts to the public decision-making setting. For example, envy-freeness is not meaningful in such an economy, since all agents "consume" the same outcome; they simply receive different utility from that outcome. In this chapter we propose and analyze a potential solution in a continuous public decision-making environment (i.e., an outcome is a point in $\mathbb{R}^m$, where each of the m dimensions represents a public issue) that we call *equality of power*.

The idea at the foundation of equality of power is that each individual's opinion should be given equal weight. This is widely considered by political theorists to be the defining feature of democracy [55] if not of justice more generally [2, 4, 163]. As Dahl puts it, "The moral judgement that all human beings are of intrinsically equal worth...(requires) that the good or interests of each person must be given equal consideration." Despite this progress on the political theory front, no version of the equal power concept has been formalized technically. How do we formally define "equal weight" of opinions? In this chapter, we propose a formal definition of equality of power, and show that for quadratic utility functions[2] and a large number of agents[3], there always exists an outcome satisfying this definition.

8.1.1 Our contribution

Ronald Dworkin's seminal work in private goods economies suggests that each agent should be allowed to impose equal cost on the rest of society [71]. We model this as *externality*. The externality of an agent is the decrease in welfare (we focus on utilitarian welfare in this chapter) for everyone else caused by the existence of that agent. That is, consider the outcome that would be chosen in the absence of that agent, versus the outcome chosen when that agent is included: the externality is the difference in utilitarian welfare for the rest of the agents between those two outcomes. We define an agent's power to be her externality, and so equal power requires all agents to have the same externality.

However, we cannot define utilitarian welfare in the standard sense, because we do not know the scale of the individual utilities. To define a common scale, we follow the spirit of Dworkin [71] and measure utilities by a metric where the marginal value of additional power for every agent is equal.

[1] Indeed, this applies to any CES welfare function other than Nash welfare: see Section 1.5.
[2] See Section 8.1.1 for a definition.
[3] Specifically, we show that for any finite number of agents, we achieve an approximate version of this equal power outcome, and the approximation error goes to 0 as the number of agents goes to infinity.

This is tightly analogous to the definition of equality of resources in terms of equal units of an artificial auction currency, which is exactly the concept of competitive equilibrium from incomes concept from Varian [169]. For additional intuition, imagine that the social planner has a finite amount of power to allocate. In order to maximize utilitarian welfare, the marginal value of additional power should be the same for each agent: otherwise, moving power from agents with lower marginal value to agents with higher marginal value would increase the utilitarian welfare. We emphasize that the above discussion is not a technical statement, nor is it representative of our actual mathematical model; we include it solely for intuition behind choosing a common scale which equalizes the marginal value of additional power. That said, our choice of common scale is not without drawback, and we think it would be interesting for future work to consider other ways of defining a common utility scale.

Informal statement of results

Our full mathematical model is given in Section 8.2, but we give an informal description here. We assume agents have quadratic utilities: each agent i has an ideal point $y_i \in \mathbb{R}^m$, along with a weight w_{ij} for each issue j. Agent i's utility for an outcome $x \in \mathbb{R}$ is defined as

$$u_i(x) = -\sum_{j=1}^{m} w_{ij}(y_{ij} - x_j)^2$$

where y_{ij} and x_j are the jth coordinates of y_i and x, respectively. Note that $u_i(x)$ is maximized at $x = y_i$.

To define "marginal value for additional power", we use the following elicitation scheme. Consider an outcome x for the public decision-making problem. We ask each agent to move the outcome towards her ideal point, under the constraint that the externality she imposes on the rest of society is at most some constant γ. When γ is uniform across all agents, this satisfies equality of power.

Our goal, then, is the following. We desire a scaling of utilities $\mathbf{c}$ and a public decision-making outcome x[4] such that all of the following hold:

1. Each agent has equal power. This is achieved by having γ be uniform across agents.
2. Each agent has equal marginal utility for additional power with respect to the elicitation scheme described above (allowing each agent to move the outcome towards their ideal point).
3. The net movement in each direction in the above elicitation scheme is 0.
4. The outcome x maximizes utilitarian welfare with respect to $\mathbf{c}$.

For quadratic utilities, we are able to prove existence of such an x and $\mathbf{c}$ in the limit as the number of agents approaches infinity. This leads to Theorem 8.4.2, whose formal statement comes later. Here $\delta_{ij}(x) \in \mathbb{R}$ represents the amount agent i chooses to move from the current point on issue j. We will often simply denote this by δ_{ij}, and denote the desired shift vector of agent i by δ_i.

[4]This scaling, which we denote $\mathbf{c}$, will be a vector assigning a scaling factor to each agent. The outcome x will be a vector in $\mathbb{R}^m$, where m is the number of issues.

Note that we do not need to explicitly require that all agents have equal power, as this is ensured by the elicitation scheme (as long as γ is uniform).

Theorem 8.4.2 (Informal)**.** *When agent utilities are quadratic, there exists an outcome x and a scaling of agent utilities* $\mathbf{c}$ *such that as the number of agents goes to infinity, all of the following hold:*

1. *The net movement along each issue (i.e., $\sum_i \delta_{ij}(x)$) is 0.*

2. *Every agent has the same marginal utility for additional power.*

3. *The outcome x maximizes utilitarian welfare with respect to* $\mathbf{c}$.

The technical statement of the theorem can be found in Section 8.4. Our proof is quite technically involved. Along with some standard (though involved) Lagrangian duality techniques, we use a fixed point argument whereby we show that a particular infinite-dimensional function admits "almost-fixed" points, i.e., points z where z and $f(z)$ are arbitrarily close (we will end up choosing our scaling $\mathbf{c}$ to be an almost-fixed point of this particular function). Our primary technical contribution is a novel technique for proving existence of approximate fixed points; see Section 8.4.1 for a more in-depth discussion.

Dependence of marginal utility and utilitarian welfare on the utility scaling

Both the marginal utilities and the utilitarian welfare are computed with respect to the utility scale $\mathbf{c}$. The reader may be worried that this makes Theorem 8.4.2 circular, but it is important to recognize three things. First, the scaling $\mathbf{c}$ is not a free parameter: it is tied down by our requirement that the marginal value for additional power be equal. Second, as mentioned above, this is strongly inspired by the definition of equality of resources in terms of equal amounts of an artificial currency (see Varian [169]).

Third, and most importantly, we argue that it is not meaningful to ask for equal marginal utilities or welfare maximization with respect to the "true" utilities. This is because, in our model, "true" utilities do not really exist: it is not meaningful for ask for the absolute scale of an agent's utility (since there are no monetary payments). The model is not that we are given agents' true utilities and we are scaling them, the model is that we are *defining* a scale of agents' utilities, since some scale is needed in order to maximize utilitarian welfare. Inspired by [71], we are choosing a scale that equalizes the marginal utilities.

For some intuition, in the one dimensional case, the outcome specified by Theorem 8.4.2 turns out to be the median of the agents' ideal points (see Section 8.3). Furthermore, we argue in Section 8.10 that our solution concept is not trivial, by showing that an "obvious" choice for $\mathbf{c}$ (specifically, giving each agent the same scaling factor c_i) does not work.

Finally, we briefly discuss incentives and computation. Our query to agents – to provide a desired shift from the current point, under the equal power constraint – is an elicitation method, not a mechanism. Consequently, our result should be thought of only as an existence result. We do not consider mechanism design in a formal sense in this chapter, and leave that for future work. We are optimistic about the possibility of an iterative procedure for computing x and $\mathbf{c}$, where on each step, each agent provides a $\delta_i(x)$, and we use $\delta_1 \ldots \delta_n$ to compute the next iterate.

8.1.2 Connections to quadratic voting

It will turn out that our equal power constraint will reduce to a simple quadratic constraint of the form $\sum_j q_j \delta_{ij}^2 \le \gamma$, where each q_j is a positive constant and each j is an issue. Quadratic voting is an increasingly promising voting scheme, both in theory [14, 13, 49, 102, 115, 116, 146, 174] and practice [89, 147, 155]. The fact that our equal power outcome can be implemented with (weighted) quadratic voting leads to a host of promising directions for future work.

In particular, we are optimistic about the possibility of an iterative protocol for computing our desired outcome $(x^*, \mathbf{c})$. As suggested above, consider an iterative algorithm where on each step, we ask each agent for their desired shift $\delta_i(x)$ from the current point x, and use those shifts to compute the next iterate. This algorithm was first studied by Hylland and Zeckhauser in 1979 [102], although instead of the externality constraint (which reduces to $\sum_j q_j \delta_{ij}^2 \le \gamma$), they subjected each agent to the dimension-symmetric quadratic voting constraint of $\sum_j \delta_{ij}^2 \le \gamma$[5]. They show that their procedure converges to a Pareto optimum.

However, we desire something stronger than just a Pareto optimum. Intuitively, by using a dimension-symmetric constraint, their algorithm ignores the fact that some issues are more important to the population than others. The more the rest of society cares about an issue, the more difficult it should be for an individual to affect the outcome on that issue. This is what Dworkin's quote from the beginning of our work captures, and what inspires our equal externality constraint. As discussed above, our equal externality constraint will reduce to a constraint of the form $\sum_j q_j \delta_{ij}^2 \le \gamma$. Each q_j should be interpreted as the aggregate weight society places on issue j.

The distinction between these two constraints is not simply technical. Since each issue is unitless in this model, it not clear what the "right" description of the issue space is, i.e., the right scale for each issue[6]. Our equal externality constraint will be invariant to such rescaling, as intuitively should be the case: if some issue j is rescaled, q_j will simply rescale accordingly. This means that regardless of the representation of the issue space, the outcome described by Theorem 8.4.2 will be the same. However, Hylland and Zeckhauser's algorithm dimension-symmetric algorithm is extremely vulnerable to this: their outcome will depend dramatically on the precise description of the decision space.

For future work, we are interested in the variant of their algorithm where their dimension-symmetric constraint is replaced with our equal power constraint. This leads to another complication: the right scale for each issue (i.e., q_j) is not known a priori. However, we believe that the right scaling can be discovered as the algorithm progresses based on agents' desired shifts. This is similar to how iterative algorithms for computing private goods market equilibria[7] slowly discover the right prices based on agent demands. All in all, we conjecture that this will lead to an iterative algorithm for public decision-making that both maximizes utilitarian welfare, and is consistent with

[5] The dimension-symmetric version of this algorithm has also been studied in [14, 13, 49, 89].
[6] Note that rescaling of the issue space is independent of our scaling $\mathbf{c}$ of the agent utility functions.
[7] Such algorithms are often known as *tâtonnements*.

the spirit of equality of resources and envy-freeness studied by economics and philosophy giants such as Ronald Dworkin, Hal Varian, and many more.

Further connections to quadratic voting and second order methods

For the expert reader, we include a brief discussion of some more technical aspects of these connections. Going back to Dworkin [71], he suggests the use of an auction based on equal initial endowments; while he is not explicit about the auction theory involved, he seems to appeal to the idea of a Walrasian auction to which many auction designs converge in large replications of private goods economies with a fixed number of goods [54, 154]. However, the structure of power and quadratic voting is fundamentally different than the linear pricing of a Walrasian auction. For a large population, each agent is only able to suggest a very small shift δ_i from the current point. In particular, the second and higher derivatives of her utility function with respect to δ_i vanish as the number of agents goes to infinity. In order to capture the remaining first derivative, the "pricing" of the δ_i (i.e., the externality constraint on δ_i) should therefore be a quadratic form rather than a linear function, so that the first derivatives of the constraint are linear.

8.1.3 Other related work

There has been significant recent progress on the theory of public decision-making, some of which with close ties to our work, and some of which using very different approaches. An iterative algorithm which elicits a desired shift from each agent on each step has been studied in [14, 13, 49, 89] and shown to converge under certain assumptions. Furthermore, most of this work does focus specifically on quadratic constraints on the desired shifts. However, none of this work addresses the "weighting" or "rescaling" of dimensions that is crucial to our work (and handled by the q_j constants, as discussed above). For example, [89] focuses on the case where each agent cares about all of the dimensions the same amount.

One can think of the equal power constraint as a pricing mechanism, in the sense that the amount of externality caused (which is equal to $\sum_j q_j \delta_{ij}^2$) is the "price", and each agent has γ units of power to spend. One famous result regarding pricing for public decision-making is that when arbitrary personalized prices are allowed (i.e., the central authority can give agents different prices for the same issue with no restrictions), any Pareto optimal point can be a market equilibrium [82]. (This was also discussed in Chapter 7). Instead of using personalized prices, we subject each agent to exactly the same equal power constraint. In this way, our work is arguably more consistent with the spirit of equality of resources.

We briefly mention several non-market approaches. Storable Votes [41] considers a repeated voting context, and permits agents to store their votes for future meetings. In [52], the authors adapt traditional private goods fairness axioms (such as a proportionality) to the public decision-making context for the case where only a discrete set of outcomes are allowed for each issue. The discrete public decision-making problem is also studied by [76], which considers approximate versions of the core, since the (exact) core is not guaranteed to exist in the discrete version of the problem.

The chapter proceeds as follows. Section 8.2 presents the formal model. Section 8.3 considers the one-dimensional case[8]; this serves as a "warm-up" for the main proof. Since the proof of our main result (Theorem 8.4.2) is quite involved, we use Section 8.4 to set up the main result and provide a detailed roadmap of the proof. We then move on to the formal proof. Section 8.5 contains the fixed point argument that we use to identify our desired outcome x and scaling $\mathbf{c}$. Section 8.6 provides some additional setup before embarking on the rest of the proof. Section 8.7 proves several properties that will be important throughout the proof, such a technical version of the statement "each agent is a small fraction of a large population". Section 8.8 characterizes each agent's desired shift δ_i, and show that under the choice of x and $\mathbf{c}$ from Section 8.5, (almost all) the agents have (almost) the same marginal value for additional power. Section 8.9 handles the last requirement of Theorem 8.4.2, which is that $\sum_j \delta_{ij}(x)$ is (almost) 0 for each issue j. Finally, Section 8.10 shows that an "obvious" choice of $\mathbf{c}$ (specifically, giving every agent the same scaling factor) is not sufficient for our purposes; this section is solely for intuition.

8.2 Model

As in all chapters of the thesis, we use the same basic resource allocation model: we have a set of n agents with preferences over m resources. In this case, each resource represents a public issue. We assume that an outcome for a particular issue is a scalar in $\mathbb{R}$, so an outcome for the overall problem is a vector in $\mathbb{R}^m$. A group of agents need to choose an outcome for this decision-making problem. We assume that each agent is drawn i.i.d. from an integrable probability distribution $p : N \to \mathbb{R}_{\geq 0}$ over possible agent types, where N is the set of agent types. We assume that the distribution is not concentrated too strongly anywhere: specifically, we assume that there exists $p_{max} > 0$ such that $p(i) \leq p_{max}$ for all i. We will use "agents" and "agent types" interchangeably. Since we are holding m fixed and taking n to infinity, we will think of m as a constant (in that we suppress it in asymptotic notation).

Each agent type is specified by an ideal outcome $y_i \in \mathbb{R}^m$ and a weight vector $w_i \in \mathbb{R}^m$. The weight vector represents how much the agent cares about different issues. Let $y_{ij} \in \mathbb{R}$ and $w_{ij} \in \mathbb{R}_{\geq 0}$ denote agent i's ideal outcome and weight for issue j, respectively. Then agent i's utility for an arbitrary point $x \in \mathbb{R}^m$ is

$$u_i(x) = -c_i \sum_{j \in M} w_{ij}(x_j - y_{ij})^2$$

where $c_i \in \mathbb{R}_{>0}$ is the scaling of agent i's utility that we choose. Note that agent i's utility is maximized at $x = y_i$.

Recall that χ denotes the set of feasible outcomes, so in this model, $\chi \subset \mathbb{R}^m$. We assume in this chapter that χ is bounded and convex. Define d_{max} by

$$d_{max} = \sup_{a,b \in \chi} ||a - b||_2$$

[8] Recall that our desired outcome turns out to be the median of the agents' ideal points in this case.

where $||a - b||_2 = \sqrt{\sum_{j \in M} (a_j - b_j)^2}$ is the L_2 norm.

We also assume all agents' weights are bounded above and below. Specifically, we assume there exists $w_{min}, w_{max} > 0$ such that $w_{min} \leq w_{ij} \leq w_{max}$ for all i, j[9]. The assumptions of boundedness of χ and boundedness of w_{ij} for all i, j together imply that the set N is bounded.

Equality of power

Let n be the number of agents sampled from p, and let N_s be the random variable representing this set of sampled agents. Rather than focusing on the welfare of the actual set of agents sampled, we use the *expected* utilitarian welfare with respect to the distribution p. However, since the agents are drawn i.i.d., the law of large numbers implies that the two coincide in the limit as $n \to \infty$ anyway. Specifically, for $x \in \chi$, the expected utilitarian welfare $U(x)$ (which depends on the chosen c_i's) is defined by

$$U(x) = \mathbb{E}_{N_s \sim p}\left[\sum_{i \in N_s} u_i(x)\right] = n \int_{i \in N} p(i) u_i(x)\, \mathrm{d}i$$

For our equal power (i.e., equal externality) constraint, we define externality with respect to the expected utilitarian welfare U, not with respect to the welfare of the actual sampled agents N_s. Formally, given a current point $x \in \chi$, we define the externality of a desired shift $\delta \in \mathbb{R}^m$ to be $U(x) - U(x + \delta)$[10]. Thus given a current point x and a small power constant $\gamma > 0$, we present each agent i with the following convex program:

$$\max_{\delta \in \mathbb{R}^m} u_i(x + \delta) \qquad (8.1)$$
$$s.t. \quad U(x) - U(x + \delta) \leq \gamma$$

Let $\delta_i(x) \in \mathbb{R}^m$ be the optimal solution to Program 8.1 starting from point $x \in \mathbb{R}^m$, and let $\lambda_i(x) \in \mathbb{R}$ be the value of the Lagrange multiplier in the optimal solution. These variables refer only to the optimal solution of agent i's copy of Program 8.1, not any sort of global optimal solution. Also note that this program implicitly depends on the scaling factors $\mathbf{c}$ (through U).

For a convex program with a differential objective function (such as Program 8.1), the Lagrange multiplier represents how much we could improve the objective value if the constraint were relaxed[11]. In our case, the objective function here is agent i's utility for outcome $x + \delta$, and the constraint is enforcing that the power used by agent i is at most γ. Thus the Lagrange multiplier $\lambda_i(x)$ is exactly agent i's marginal value for additional power, and this is what we wish to equalize across agents.

[9] This is really only one assumption, actually: c_i will be invariant to w_i in the sense that if agent i doubles w_i, c_i will halve. This means that only the relative weights matter anyway, so we are essentially assuming that the ratio of each agent's maximum weight dividing by minimum weight is bounded above, i.e., w_i is "well-conditioned".

[10] The careful reader may notice that externality is usually defined as the impact on the welfare of everyone *else*, excluding the agent in question. However, since we assume agents to be drawn i.i.d., this distinction is not important.

[11] See Chapter 5 of [23] for an introduction to this type of perturbation analysis.

Our solution concept

Our solution concept – an *equal-power equal-λ ε-equilibrium* – asks us to chose an outcome $x \in \mathbb{R}^m$ and agent scaling factors $\mathbf{c} \in \mathbb{R}^N_{>0}$ (along with a power constant γ and a particular Lagrange multiplier λ) that satisfies three requirements. First, the expected net movement (the sum of $\delta_i(x)$'s) from the current point is smaller than ε. We use the L_2 norm to express the size of $\sum_i \delta_i(x)$. The second requirement is that all agents except an ε fraction have the same value of $\lambda_i(x)$, up to an ε error. Finally, the selected outcome $x \in \mathbb{R}^m$ must maximize welfare with respect to the chosen agent scaling factors $\mathbf{c} \in \mathbb{R}^N_{>0}$. Since N is a (continuous) distribution over agent types, $\mathbf{c}$ will be an infinite-dimensional vector.

Definition 8.2.1. *An equal-power equal-λ ε-equilibrium is a outcome $x \in \mathbb{R}^m$, agent scaling factors $\mathbf{c} \in \mathbb{R}^N_{>0}$, power constant γ, and marginal utility $\lambda > 0$ such that*

1. $\mathbb{E}_{N_s \sim p}\left[|| \sum_{i \in N_s} \delta_i(x)||_2\right] < \varepsilon$.
2. *The expected number of agents i with $(1-\varepsilon)\lambda \leq \lambda_i(x) \leq \lambda$ is at least $(1-\varepsilon)n$.*
3. *The outcome x maximizes welfare with respect to $\mathbf{c}$, i.e., $x \in \arg\max_{x'} U(x')$.*

We will usually leave x implicit and just write δ_i (which is a vector), δ_{ij} (which is a scalar), and λ_i (which is a scalar). Note that we are only asking for $|| \sum_{i \in N_s} \delta_i||_2$ to be small in expectation, but the law of large numbers ensures that the realized value will converge to the expectation with probability 1 as $|N_s| \to \infty$.

Our goal will be to show that for any $\varepsilon > 0$, there exists a large enough n (number of agents) such that an equal-power equal-λ ε-equilibrium exists (for some choice of λ). Specifically, we will choose a fixed x and $\mathbf{c}$ based on the underlying distribution p, agnostic to the set of agents that are actually sampled. We then show that the approximation error goes to 0 as $n \to \infty$.

In the next section, we show that in the one-dimensional case, our desired outcome is the median of the agents' ideal points.

8.3 Warm-up: one dimension

We view the one-dimensional case as a warm-up in the sense that the result of this section (Theorem 8.3.1) will be subsumed by our result for the m-dimensional case (Theorem 8.4.2). Although the proof of Theorem 8.4.2 is much more technically involved, the general flow of the proof for the one-dimensional case is similar, so we find it instructive to present first. The main difference is that for the m-dimensional case, the equilibrium point is identified as an approximate fixed point of a particular (somewhat complicated) function. In contrast, for the one-dimensional case, we are able to "guess" that the equilibrium point should be the median. There are several additional small differences, such as the specific bound on λ_i. If the reader is confident and wishes to skip this warm-up, we encourage them to proceed directly to Section 8.4.

Since we are working with a single dimension, we have $w_i, y_i \in \mathbb{R}$, and agent i's utility function is $u_i(x) = -c_i w_i (x - y_i)^2$. Define x to be the median of the agents' ideal points. Specifically, choose

$x \in \mathbb{R}$ such that

$$\int_{i \in N} p(i)\ \mathrm{sgn}(x - y_i)\,\mathrm{d}i = 0$$

That is, the probability of sampling an agent i with $y_i \leq x$ is equal to the probability of sampling an agent i with $y_i \geq x$. Since p is continuous, such an x must exist (if there are multiple, choose one arbitrarily).

For each $i \in N$ with $y_i \neq x$, we define c_i to be inversely proportional to her weight w_i and the distance between y_i and x. Agents with $y_i = x$ will turn out to not matter (because this set has measure 0), so we set $c_i = c$ for those agents, where c can be any constant.

$$c_i = \begin{cases} \dfrac{1}{w_i |x - y_i|} & \text{if } y_i \neq x \\ c & \text{if } y_i = x \end{cases}$$

This definition will imply that the outcome is scale-invariant: doubling w_i results in halving c_i, which leads to the same final utility function of $u_i(x) = -\frac{(x-y_i)^2}{|x-y_i|} = -|x - y_i|$. Also, let $q = n \int_{k \in N} p(k)|x - y_k|^{-1}\,\mathrm{d}k$.[12] Since $\int_{k \in N} p(k)|x - y_k|^{-1}\,\mathrm{d}k$ is just some constant (i.e., independent of n), q is $\Theta(n)$.

Theorem 8.3.1. *For $x, \mathbf{c}$ as defined above, there exists a power constant γ such that the following all hold:*

1. $|n \int_{i \in N} p(i)\delta_i(x)\,\mathrm{d}i|$ *goes to 0 as $n \to \infty$.*
2. *For each agent i except a vanishing fraction*[13]*, λ_i goes to $1/\sqrt{q\gamma}$ as $n \to \infty$.*
3. *The outcome x maximizes welfare with respect to $\mathbf{c}$, i.e., $x \in \arg\max_{x'} U(x')$.*

Note that rather than converging to a specific value, λ_i is approaching $1/\sqrt{q\gamma}$, and q is $\Theta(n)$. However, since we are interested in multiplicative differences in λ_i, this is not a problem. For the m-dimensional case, one product of our more complicated setup will be that λ_i converges to a specific value: specifically, $1/\sqrt{\gamma}$.

Since the point of this section is to give intuition for the main proof (and not to actually prove an interesting result), we are less formal and rigorous than we will be in the proof of the main result. There are also a few (uninformative) parts of the proof that we defer entirely until the proof of the main result.

To start, we consider welfare maximization.

Lemma 8.3.1. *The outcome x as defined above maximizes welfare with respect to $\mathbf{c}$.*

Proof. Since U is concave and differentiable, and we are maximizing over an unrestricted domain, x maximizes U if and only if derivative of U at x is 0:

$$\frac{d}{dx}U(x) = \frac{d}{dx} n \int_{i \in N} p(i) u_i(x_i)\,\mathrm{d}i = -2\int_{i \in N} p(i) c_i w_i (x - y_i)\,\mathrm{d}i = -2\int_{i \in N} p(i)\frac{w_i(x - y_i)}{w_i|x - y_i|}$$

[12]Note that although $|y_i - x|^{-1}$ is undefined at $x = y_i$, its integral is indeed well-defined.

[13]That is, the fraction of agents for whom this does not hold should go to 0 as $n \to \infty$.

By definition of x, we have $\int_{i\in N} p(i)\ \mathrm{sgn}(x-y_i)\,\mathrm{d}i = \int_{i\in N} p(i)\frac{x-y_i}{|x-y_i|}\,\mathrm{d}i = 0$. Thus $\frac{d}{dx}U(x) = 0$, so $x \in \arg\max_{x'} U(x)$. □

Next, we obtain an expression for δ_i in terms of λ_i.

Lemma 8.3.2. *For each $i \in N$, we have* $\delta_i = \dfrac{(y_i - x)}{|x-y_i|(|x-y_i|^{-1} + \lambda_i q)}$.

Proof. We begin by writing the Lagrangian of Program 8.1 for an arbitrary agent i:

$$L(\delta_i, \lambda_i) = u_i(x+\delta_i) - \lambda_i\big(U(x) - U(x+\delta_i) - \gamma\big)$$

The KKT conditions imply that the derivative of L with respect to δ_i should be zero for the optimal δ_i:

$$\begin{aligned}
\frac{d}{d\delta_i}L(\delta_i,\lambda_i) &= \frac{d}{d\delta_i}u_i(x+\delta_i) + \lambda_i\frac{d}{d\delta_i}U(x+\delta_i)\\
&= -\frac{2(x+\delta_i-y_i)}{|x-y_i|} - \lambda_i n\int_{k\in N} p(k)\frac{2(x+\delta_i-y_k)}{|x-y_k|}\,\mathrm{d}k\\
&= -\frac{2\delta_i}{|x-y_i|} - \frac{2(x-y_i)}{|x-y_i|} - n\int_{k\in N} p(k)\frac{2\lambda_i\delta_i}{|x-y_k|}\,\mathrm{d}k - 2n\lambda_i\int_{k\in N} p(k)\frac{x-y_k}{|x-y_k|}\,\mathrm{d}k
\end{aligned}$$

By the definition of x, we have $\int_{k\in N} p(k)\ \mathrm{sgn}(x-y_k)\,\mathrm{d}k = \int_{k\in N} p(k)\frac{x-y_k}{|x-y_k|}\,\mathrm{d}k = 0$, so

$$\begin{aligned}
\frac{d}{d\delta_i}L(\delta_i,\lambda_i) &= -\frac{2\delta_i}{|x-y_i|} - \frac{2(x-y_i)}{|x-y_i|} - n\int_{k\in N} p(k)\frac{2\lambda_i\delta_i}{|x-y_k|}\,\mathrm{d}k\\
&= -2\delta_i\Big(|x-y_i|^{-1} + \lambda_i n\int_{k\in N} p(k)|x-y_k|^{-1}\,\mathrm{d}k\Big) - \frac{2(x-y_i)}{|x-y_i|}
\end{aligned}$$

Since $\frac{d}{d\delta_i}L(\delta_i,\lambda_i) = 0$, we get

$$\delta_i = \frac{(y_i-x)}{|x-y_i|\Big(|x-y_i|^{-1} + \lambda_i n\int_{k\in N} p(k)|x-y_k|^{-1}\,\mathrm{d}k\Big)} = \frac{(y_i-x)}{|x-y_i|(|x-y_i|^{-1}+\lambda_i q)}$$

□

We will now use Lemma 8.3.2 to derive explicit bounds on λ_i. Let $\hat{N} = \{i \in N : |x-y_i| \geq 1/n^{1/4}\}$. Clearly as n goes to ∞, the fraction of agents not in $\hat{N}$ goes to 0. Furthermore, we can always choose the power constant γ to be small enough such that every agent in $\hat{N}$ exhausts her power. Thus for each $i \in \hat{N}$, $U(x) - U(x+\delta_i) = \gamma$.

Lemma 8.3.3. *For each $i \in \hat{N}$,* $\dfrac{1}{\sqrt{q}}\Big(\dfrac{1}{\sqrt{\gamma}} - \dfrac{1}{\Omega(n^{1/4})}\Big) \leq \lambda_i \leq \dfrac{1}{\sqrt{q\gamma}}$.

Proof. Using arithmetic (we prove this for the more general setting later: see Lemma 8.8.1), $U(x) - U(x+\delta_i) = \delta_i^2 q$. Therefore $\delta_i^2 = \gamma/q$. Also using Lemma 8.3.2, for each $i \in \hat{N}$ we have

$$\begin{aligned}\frac{(y_i - x)^2}{(y_i - x)^2(|x - y_i|^{-1} + \lambda_i q)^2} &= \gamma/q \\ |x - y_i|^{-1} + \lambda_i q &= \sqrt{q/\gamma} \\ \lambda_i &= \frac{1}{\sqrt{q\gamma}} - \frac{1}{q|x - y_i|}\end{aligned}$$

Clearly we have $\lambda_i \leq 1/\sqrt{q\gamma}$. Since $|x - y_i| \geq 1/n^{1/4}$ for $i \in \hat{N}$, we have $\lambda_i \geq \frac{1}{\sqrt{q}}\left(\frac{1}{\sqrt{\gamma}} - \frac{n^{1/4}}{\sqrt{q}}\right) = \frac{1}{\sqrt{q}}\left(\frac{1}{\sqrt{\gamma}} - \frac{1}{\Omega(n^{1/4})}\right)$. Therefore for each $i \in \hat{N}$, $\frac{1}{\sqrt{q}}\left(\frac{1}{\sqrt{\gamma}} - \frac{1}{\Omega(n^{1/4})}\right) \leq \lambda_i \leq \frac{1}{\sqrt{q\gamma}}$. □

Finally, we need the expected net movement to be small: $\int_{i \in N} p(i)\delta_i(x)\,\mathrm{d}i$ goes to 0 as $n \to \infty$.

Lemma 8.3.4. *As $n \to \infty$, $|\int_{i \in N} p(i)\delta_i(x)\,\mathrm{d}i|$ goes to 0.*

The proof for Lemma 8.3.4 is fairly tedious, even for one dimension. Thus we defer this part of the proof until later, when we formally prove this for the general case.

Lemmas 8.3.1, 8.3.3, and 8.3.4 together imply Theorem 8.3.1.

8.4 Main theorem setup

In this section, we state our main theorem (and one variant of the theorem), and provide a roadmap of the proof. Informally, our main result is:

Theorem 8.4.2 (Informal)**.** *As the number of agents goes to infinity, there exists an outcome x and a scaling of agent utilities $\mathbf{c}$ such that all of the following hold:*

1. *The expected net movement – $||n \int_{i \in N} p(i)\delta_i(x)\,\mathrm{d}i||_2$ – is 0.*
2. *Every agent has the same marginal utility for additional power.*
3. *The outcome x maximizes expected welfare with respect to $\mathbf{c}$.*

This implies that for any $\varepsilon > 0$ and large enough n, there exists an equal-power equal-λ ε-equilibrium.

We state two theorems in this section. The theorem statements refer to a function f that will be defined in Section 8.5. The variable α is a parameter of f that is used to ensure continuity of f, and will be chosen as a function of n.

Most of the chapter is devoted to proving Theorem 8.4.1, which assumes an exact fixed point of f, and presents approximation bounds on our quantities of interest (i.e, λ_i and $||n \int_{i \in N} p(i)\delta_i(x)\,\mathrm{d}i||_2$) as a function of α and n. An intermediate lemma (Lemma 8.6.1), which appears in Section 8.6, shows that there is a choice of α (specifically, $\alpha = n^{-7/8}$) such that the approximation error vanishes as n goes to ∞, also assuming an exact fixed point of f.

In reality, we are not able to prove that f has an exact fixed point. Instead, we show in Section 8.5 that f has an ε-fixed point for each $\varepsilon > 0$ (where $\varepsilon = 0$ would denote an exact fixed point). Theorem 8.4.2 states that we can pick ε small enough that using an ε-fixed point of f is good enough.

Theorem 8.4.1. *Suppose f as defined in Section 8.5 has an exact fixed point $\mathbf{c}$ for any choice of α and n. Let $x_j = \left(\int_{i\in N} p(i)c_i w_{ij}\,\mathrm{d}i\right)^{-1} \int_{i\in N} p(i)c_i w_{ij} y_{ij}\,\mathrm{d}i$. Let α be chosen as a function of n so that $\lim_{n\to\infty} n^{3/2}\alpha = \infty$ and $\lim_{n\to\infty} \alpha^{m/2} n^{m/4} = 0$. Then for any ε, there exists α and n such that $(x, \mathbf{c}, \gamma, 1/\sqrt{\gamma})$ is an equal-power equal-λ ε-equilibrium. Specifically:*

1. $||n \int_{i\in N} p(i)\delta_i(x)\,\mathrm{d}i||_2 \leq \dfrac{\gamma + \sqrt{\gamma}}{O(n^2\alpha^2)} + O(\alpha^{m/2} n^{m/4+1})$.
2. *For all i except an expected $O(\alpha^{m/2}n^{m/4})$ fraction,* $O\Big(\sqrt{\dfrac{n^{3/2}\alpha}{n^{3/2}\alpha+1}}\Big)\dfrac{1}{\sqrt{\gamma}} \leq \lambda_i(x) \leq \dfrac{1}{\sqrt{\gamma}}$.
3. *The outcome x maximizes welfare with respect to $\mathbf{c}$.*

The assumptions of $\lim_{n\to\infty} n^{3/2}\alpha = \infty$ and $\lim_{n\to\infty} \alpha^{m/2} n^{m/4} = 0$ in Theorem 8.4.1 are necessary for a few parts of the proof to work.

Using Theorem 8.4.1 and Lemma 8.6.1, we get our final result:

Theorem 8.4.2. *Let $\mathbf{c}$ be an ε-fixed point of f and let $x_j = \left(\int_{i\in N} p(i)c_i w_{ij}\,\mathrm{d}i\right)^{-1} \int_{i\in N} p(i)c_i w_{ij} y_{ij}\,\mathrm{d}i$. Let $\alpha = n^{-7/8}$ and $m \geq 6$. Then there exists a small enough ε such that all of the following hold:*

1. $||n \int_{i\in N} p(i)\delta_i(x)\,\mathrm{d}i||_2 \leq (\gamma + \sqrt{\gamma})O(n^{-1/4}) + O(n^{-1/8})$.
2. *For all i except an expected $O(n^{-3/4})$ fraction,* $O\Big(\sqrt{\dfrac{n^{5/8}}{n^{5/8}+1}}\Big)\dfrac{1}{\sqrt{\gamma}} \leq \lambda_i(x) \leq \dfrac{1}{\sqrt{\gamma}}$.
3. *The outcome x maximizes welfare with respect to $\mathbf{c}$.*

8.4.1 Proof roadmap

In this section, we state and describe the key lemmas in our proof. The proofs of some of these lemmas involve additional lemmas, but we only include the most important lemmas.

Section 8.5 is devoted to showing that the function f satisfies the *approximate fixed-point property* (AFPP) [16]: for every $\varepsilon > 0$, f admits an ε-fixed point[14]. In our opinion, this is the most interesting part of the overall proof, and constitutes a novel approach for proving existence of approximate fixed points.

The function f will actually be infinite-dimensional, and it turns out (perhaps unsurprisingly) to be challenging to prove fixed point existence for infinite-dimensional functions. Instead, we will analyze a finite-dimensional variant denoted by g_A (A will be any finite set of agents).

We first show that g_A has an exact fixed point for any finite set A (Lemma 8.5.1). In the following lemma statement, α and $s_1 \dots s_{|A|}$ are parameters of g_A.

[14] See Section 8.5 for a formal definition.

Lemma 8.5.1. *For any* $\alpha \in (0, \frac{\sqrt{n}w_{max}^2 d_{max}^2}{w_{min}}]$*, any finite set* $A \subset N$*, and any nonnegative* $s_1 \dots s_{|A|}$ *that sum to 1, the function* g_A *has a fixed point* $\mathbf{c}^* \in \left[\dfrac{w_{min}}{w_{max}^2 d_{max}^2}, \dfrac{\sqrt{n}}{\alpha}\right]^{|A|}$.

We then show that for any $\varepsilon > 0$, there exists an A large enough that any exact fixed point of g can be transformed into an ε-fixed point of f. This part of the argument is quite involved, and uses the following steps: (1) Partition the agent space into arbitrarily small hypercubes. (2) Choose a "representative" from each hypercube in a careful way. (3) Let A be the set of those representatives, let s_ℓ be the measure of p in the ℓth hypercube, and let $\mathbf{c}$ be an exact fixed point of g_A with parameters $s_1 \dots s_{|A|}$. (4) Assign each agent not in A to have the same scaling factor as her representative. (5) Show that for small enough hypercubes, this is an ε-fixed point of f. This results in the following lemma:

Lemma 8.5.3. *The function* f *satisfies AFPP.*

Throughout the rest of the chapter after Section 8.5, we use x and $\mathbf{c}$ to specifically refer to the outcome and agent scaling factors defined in Theorem 8.4.1, not arbitrary outcomes/agent scaling factors. This is primarily for simplicity and brevity.

In Section 8.7, we establish some important properties we will use along the way. First, we show that x as defined in the statement of Theorem 8.4.1 maximizes welfare with respect to the agent scaling factors $\mathbf{c}$ defined in that theorem statement.

Lemma 8.7.1. *The outcome* x *maximizes the utilitarian welfare* $U(x)$ *with respect to* $\mathbf{c}$.

Next, we define $\hat{N}$ as the set of "normal" agents ("normal" will be defined later), and show that the measure of $N \setminus \hat{N}$ is small (i.e., almost all the agents are "normal"). Since we treat m as a constant, the right hand side is $O(\alpha^{m/2} n^{m/4})$; everything else is a constant.

Lemma 8.7.2. *We have* $\displaystyle\int_{i \notin \hat{N}} p(i)\, \mathrm{d}i \le \alpha^{m/2} n^{m/4} \frac{\pi^{m/2} p_{max}}{\Gamma(\frac{m}{2}+1)} \left(\frac{\sqrt{w_{max}}}{w_{min}}\right)^m$.

Lemma 8.7.4 states that each agent is a small fraction of the overall population, in terms of weight $c_i w_{ij}$ on any issue. Note that the left hand side is an integral over the whole population, and the right hand side is a multiple of $c_i w_{ij}$.

Lemma 8.7.4. *For any agent* $i \in N$, $\displaystyle\int_{k \in N} p(k) c_k w_{kj}\, \mathrm{d}k \ge \frac{w_{min}^2}{4 w_{max}^3 d_{max}^2} \sqrt{n} \alpha c_i w_{ij}$.

Section 8.8 to devoted to analyzing properties of δ_i and λ_i. First, we show that δ_i obeys a particular expression in terms of λ_i.

Lemma 8.8.2. *For every agent* i *and issue* j, $\delta_{ij} = \dfrac{c_i w_{ij}(y_{ij} - x_j)}{c_i w_{ij} + \lambda_i q_j}$.

We will define an approximation variable τ_i such that $\lambda_i = \dfrac{1}{\sqrt{\tau_i + \gamma}}$, and show that τ_i is small. This holds for every "normal" agent, i.e., the agents in $\hat{N}$.

Lemma 8.8.5. *For each agent* $i \in \hat{N}$, $\tau_i \le \dfrac{2\gamma^{3/2} m w_{max}}{n^{3/2}\alpha\beta w_{min} - 2\sqrt{\gamma} m w_{max}}$.

We will also show that $\tau_i > 0$ (this part will be easy). This will allow us to upper- and lower-bound the value of λ_i for each $i \in \hat{N}$. The variable η will be defined later, but we will have $\eta = \Theta(n^{3/2}\alpha)$; by assumption of Theorem 8.4.1, $\lim_{n\to\infty} n^{3/2}\alpha = \infty$. This means that $\lim_{n\to\infty} \eta = \infty$ as well, and this range of allowable λ_i shrinks to a single point (specifically, $1/\gamma$).

Lemma 8.8.6. *For all $i \in \hat{N}$, $\sqrt{\frac{\eta}{\eta+1}} \cdot \frac{1}{\sqrt{\gamma}} \leq \lambda_i \leq \frac{1}{\sqrt{\gamma}}$.*

Finally, Section 8.9 proves an upper bound on $\left|\left|n \int_{i\in N} p(i)\delta_i(x)\,\mathrm{d}i\right|\right|_2$: the expected net movement with respect to the current point. Note that $\mathbb{E}_{N_s\sim p}[||\sum_{i\in N_s}\delta_i(x)||_2] = n||\int_{i\in N} p(i)\delta_i(x)\,\mathrm{d}i||_2$.

Lemma 8.9.5. *We have* $\left|\left|n\int_{i\in N} p(i)\delta_i(x)\,\mathrm{d}i\right|\right|_2 \leq \frac{\gamma+\sqrt{\gamma}}{\Omega(n^2\alpha^2)} + O(\alpha^{m/2}n^{m/4+1})$.

Lemmas 8.7.1, 8.8.6, and 8.9.5 will together imply Theorem 8.4.1.

8.5 The fixed point argument

We will choose our agent scaling factors $\mathbf{c}$ to be an (approximate) fixed point of a particular function f (defined below). This section is devoted to constructing this function and showing that it satisfies the approximate fixed point property (defined in Definition 8.5.1).

Defining the function f.

Let $\mathbb{R}^N_{>0}$ be the set of functions $\mathbf{c}: N \to \mathbb{R}_{>0}$. We can think of each function $\mathbf{c}$ as assigning a scaling factor $c(i) > 0$ to each agent type i. For this section of the chapter, we will use the function notation $c(i)$ instead of c_i.

For brevity, for each $j \in M$ define $x_j(\mathbf{c})$ by

$$x_j(\mathbf{c}) = \Big(\int_{i\in N} p(i)c(i)w_{ij}\,\mathrm{d}i\Big)^{-1}\int_{i\in N} p(i)c(i)w_{ij}y_{ij}\,\mathrm{d}i$$

This is a continuous average of all agents' ideal points y_{ij} weighted by $c(i)$ and w_{ij}. We show later that this choice of x maximizes welfare with respect to $\mathbf{c}$ (Lemma 8.7.1).

The function $f: \mathbb{R}^N_{>0} \to \mathbb{R}^N_{>0}$ will take as input a function $\mathbf{c}: N \to \mathbb{R}_{>0}$, and returns a another function $f(\mathbf{c}): N \to \mathbb{R}_{>0}$. Just as $c(i) \in \mathbb{R}_{>0}$ is the scaling factor that $\mathbf{c}$ assigns to agent i, we use $[f(\mathbf{c})](i) \in \mathbb{R}_{>0}$ to denote the scaling factor that $f(\mathbf{c})$ assigns to agent i. For a small $\alpha > 0$ to be chosen later, we define $[f(\mathbf{c})](i)$ by

$$[f(\mathbf{c})](i) = \frac{\sqrt{n}}{\max\Big(\alpha, \sqrt{\sum_{j\in M} w_{ij}^2(y_{ij}-x_j(\mathbf{c}))^2\big(\int_{k\in N} p(k)c(k)w_{kj}\,\mathrm{d}k\big)^{-1}}\Big)}$$

Intuition behind f

Since we will be choosing $\mathbf{c}$ to be an (approximate) fixed point of f, each c_i will end up approximately equal to $[f(\mathbf{c})](i)$. Consequently, the structure of f gives us a qualitative interpretation of the agent scaling factors $\mathbf{c}$.

Ignore α and $\sqrt{n}$ for now. First, notice that f is invariant to the scale of individual utilities. Specifically, if an agent scales up her weights by a constant factor κ, $[f(\mathbf{c})](i) \approx c(i)$ decreases by a factor of κ: w_{ij}^2 becomes $\kappa^2 w_{ij}^2$, then pull κ out of the square root (still ignoring α). The consequence is that our "common utility scale" defined by $\mathbf{c}$ is invariant to individuals scaling up or down their utility functions, as it should be.

Next, think of $\int_{k \in N} p(k)c(k)w_{kj} \, \mathrm{d}k$ is the aggregate weight placed by society on issue j. Thus each term $w_{ij}^2(y_{ij} - x_j(\mathbf{c}))^2 \left(\int_{k \in N} p(k)c(k)w_{kj} \, \mathrm{d}k \right)^{-1}$ is equal to agent i's weight for issue j (i.e., w_{ij}) times her utility loss on issue j (i.e., $w_{ij}(y_{ij} - x_j(\mathbf{c}))$, as a fraction of society's total weight on that issue. Thus the summation in the denominator can be thought of as expressing how much agent i "disagrees" with the rest of society, with respect to the current point $x(\mathbf{c})$. More intense disagreement leads to a larger denominator, and smaller overall value for $c(i)$.

That said, the real reason for this choice of f is technical: in order for (almost) all agents to have (almost) the same value of λ_i, we will need $c(i)^2$ to be proportional to the expression under the square root for (almost) all agents. This is exactly what a fixed point of f gives us (ignoring α).

Finally, the purposes of $\sqrt{n}$ and α are purely technical. The maximization with α is to ensure that there is no discontinuity in f when the expression under the square root is zero (continuity is required for our fixed point analysis). We will show that α can be chosen so that the properties we desire are not affected. The $\sqrt{n}$ is simply to help certain aspects of the math later on work out smoothly.[15]

Approximate fixed points

Ideally, we would like to show that f has a fixed point. As the reader might expect, showing existence of fixed points in infinite-dimensional spaces can be challenging. Instead, we will show that f admits approximate fixed points:[16]

Definition 8.5.1. *Let X be a set. We say that a function $f : \mathbb{R}^X_{>0} \to \mathbb{R}^X_{>0}$ satisfies the approximate fixed point property (AFPP) if for every $\varepsilon > 0$, there exists $\mathbf{c}$ such that $|[f(\mathbf{c})](i) - c(i)| < \varepsilon$ for all $i \in X$. We call such a $\mathbf{c}$ an ε-*fixed point.

The rest of this section is devoted to showing that f satisfies AFPP. To do this, we define a function g_A for any finite set of agents A which serves a finite-dimensional approximation of f. We will show that g_A admits an exact fixed point for any set A. To complete the proof, we will show that picking A to be arbitrarily large allows us to approximate f arbitrarily well.

8.5.1 Showing that f satisfies AFPP

Let $A = \{i_1, i_2 \dots i_{|A|}\}$ be a finite set of agents with nonnegative coefficients $s_1 \dots s_{|A|}$ that sum to 1.With slight abuse of notation, we will use s_{i_k} and s_k interchangeably. We define a function

[15] It is worth noting that f does depend on n (both explicitly in the numerator, and implicitly through α, which will end up depending on n); this is not a problem, however. Since $x(\mathbf{c})$ is a weighted average of the agents' ideal points, scaling all $c(i)$ by the same amount (which is what $\sqrt{n}$ does) will not affect $x(\mathbf{c})$, the equilibrium outcome. With regards to α, we will need α to go to zero $n \to \infty$ so that the error introduced goes to 0. One can think of this as the impact of α going to zero as $n \to \infty$, so that we achieve an exact equal-power equal-λ equilibrium in the limit.

[16] In general, the approximate fixed point property can be defined for any metric space.

$g_A : \mathbb{R}_{>0}^{|A|} \to \mathbb{R}_{>0}^{|A|}$ by

$$[g_A(\mathbf{c})](i) = \frac{\sqrt{n}}{\max\left(\alpha, \sqrt{\sum_{j\in M} w_{ij}^2 (y_{ij} - x_j(\mathbf{c}))^2 \left(\sum_{k\in A} s_k c(k) w_{kj}\right)^{-1}}\right)}$$

for all $i \in A$. That is, g_A takes as input a function $\mathbf{c} : A \to \mathbb{R}_{>0}$ that assigns $c(i)$ to each $i \in A$, and it returns a vector $g_A(\mathbf{c}) \in \mathbb{R}_{>0}^a$ that assigns $[g_A(\mathbf{c})](i)$ each $i \in A$. When $\mathbf{c}$ has a finite domain (such as in the definition of g_A), $x_j(\mathbf{c})$ is defined to be the discrete average of y_i for $i \in A$, weighted by $c(i)$, w_{ij}, and s_i. Formally, $x_j(\mathbf{c}) = (\sum_{i\in A} s_i w_{ij})^{-1} \sum_{i\in A} s_i w_{ij} y_{ij}$. When $\mathbf{c}$ has a continuous domain as in the definition of f, $x_j(\mathbf{c})$ is defined to be the continuous weighted average defined previously.

Lemma 8.5.1 states that for any set A and any small enough α, g_A has a fixed point. The proof uses Brouwer's fixed point theorem:

Theorem 8.5.1 (Brouwer's fixed point theorem)**.** *Let ℓ be a positive integer, let $S \subset \mathbb{R}^\ell$ be convex and compact, and let $f : S \to S$ be continuous. Then there exists $\mathbf{c}^* \in S$ such that $f(\mathbf{c}^*) = \mathbf{c}^*$.*

Lemma 8.5.1. *For any $\alpha \in (0, \frac{\sqrt{n} w_{max}^2 d_{max}^2}{w_{min}}]$, any finite set $A \subset N$, and any nonnegative $s_1 \dots s_{|A|}$ that sum to 1, the function g_A has a fixed point $\mathbf{c}^* \in \left[\frac{w_{min}}{w_{max}^2 d_{max}^2}, \frac{\sqrt{n}}{\alpha}\right]^{|A|}$.*

Proof. Let $S = \left[\frac{w_{min}}{w_{max}^2 d_{max}^2}, \frac{\sqrt{n}}{\alpha}\right]^{|A|}$. Since $c(i) > 0$ for all i for all $\mathbf{c} \in S$, $\sum_{k\in A} s_k c(k) w_{kj}$ is strictly positive. Thus we have $\sum_{j\in M} w_{ij}^2 (y_{ij} - x_j(\mathbf{c}))^2 \left(\sum_{k\in A} s_k c(k) w_{kj}\right)^{-1} \geq 0$, so the denominator of $[g_A(\mathbf{c})](i)$ is always real. Furthermore, since $\alpha > 0$, the denominator is always positive, so the function is well-defined and continuous on S. It is also clear that S is convex and compact (and nonempty as long as $\alpha \leq \frac{\sqrt{n} w_{max}^2 d_{max}^2}{w_{min}}$).

It remains to show that $g_A(\mathbf{c}) \in S$ for all $\mathbf{c} \in S$. First, since the denominator is always at least α, $[g_A(\mathbf{c})](i) \leq \sqrt{n}/\alpha$ for all $\mathbf{c}$ and all i. Next, since $c(i) \geq \frac{w_{min}}{w_{max}^2 d_{max}^2}$ for all $i \in N$ (because $\mathbf{c} \in S$),

$$\begin{aligned}
\sqrt{\sum_{j\in M} w_{ij}^2 (y_{ij} - x_j(\mathbf{c}))^2 \left(\sum_{k\in A} s_k c(k) w_{kj}\right)^{-1}} &\leq \sqrt{\sum_{j\in M} w_{ij}^2 (y_{ij} - x_j(\mathbf{c}))^2 \left(\sum_{k\in A} s_k \frac{w_{min}}{w_{max}^2 d_{max}^2} w_{kj}\right)^{-1}} \\
&= \frac{w_{max} d_{max}}{\sqrt{w_{min}}} \sqrt{\sum_{j\in M} w_{ij}^2 (y_{ij} - x_j(\mathbf{c}))^2 \left(\sum_{k\in A} s_k w_{kj}\right)^{-1}} \\
&= \frac{w_{max} d_{max}}{\sqrt{w_{min}}} \sqrt{\sum_{j\in M} w_{max}^2 (y_{ij} - x_j(\mathbf{c}))^2 \left(\sum_{k\in A} s_k w_{min}\right)^{-1}} \\
&\leq \frac{w_{max}^2 d_{max}}{w_{min}} \sqrt{\sum_{j\in M} (y_{ij} - x_j(\mathbf{c}))^2} \\
&\leq \frac{w_{max}^2 d_{max}^2}{w_{min}}
\end{aligned}$$

Therefore the denominator of $[g_A(\mathbf{c})](i)$ is at most $\frac{w_{max}^2 d_{max}^2}{w_{min}}$, which implies that $[g_A(\mathbf{c})](i) \geq \frac{\sqrt{n} w_{min}}{w_{max}^2 d_{max}^2}$ for all $\mathbf{c} \in S$. Since $n \geq 1$, we have $[g_A(\mathbf{c})](i) \geq \frac{\sqrt{n} w_{min}}{w_{max}^2 d_{max}^2} \geq \frac{w_{min}}{w_{max}^2 d_{max}^2}$, as required.

Therefore $g_A(\mathbf{c}) \in S$ for every $\mathbf{c} \in S$. Thus by Brouwer's fixed point theorem, there exists an $\mathbf{c}^* \in S$ such that $g_A(\mathbf{c}^*) = \mathbf{c}^*$. □

8.5.2 Using g to approximate f

We next show that for certain choices of A, fixed points of g_A are approximate fixed points of f. The proof approach is as follows:

1. Partition the space of agent types into arbitrarily small hypercubes (Lemma 8.5.2 shows that this is possible). Thus all agents within a given hypercube will have arbitrarily similar values of y_{ij} and w_{ij}.

2. Choose a "representative" from each hypercube. The representative for the ℓth hypercube will be chosen such that for each $j \in M$, her ideal point y_{ij} and w_{ij} are equal to the weighted average (within the hypercube) of ideal points and weight vectors, respectively. Such an agent is guaranteed to exist within the same hypercube.

3. Let A be the set of representatives, let s_ℓ be the measure of p in the ℓth hypercube, and let $\mathbf{c}_A$ be a fixed point of g_A for coefficients $s_1 \dots s_{|A|}$.

4. Define $\mathbf{c} : N \to \mathbb{R}_{>0}$ so that each agent's scaling factor $c(i)$ is equal to the scaling factor of her representative under $c_A(i)$. Since every agent is in some hypercube, this fully specifies $\mathbf{c}$.

5. Show that $|[f(\mathbf{c})](i) - c(i)|$ is small.

Lemma 8.5.2. *For some $q > 0$, let $S \subset \mathbb{R}^q$ be a hybercube. Then for any ε, there exists a partition of S into hypercubes $S_1 \dots S_L$ such that for any $\ell \in \{1 \dots L\}$, for all $z, z' \in S_\ell$, $||z - z'||_\infty < \varepsilon$*[17].

Proof. If S is a hypercube, then we can bisect it along every coordinate to create many hypercubes, each with side length half of the original. Starting with N, continue halving the side length in this way until the side length of every hypercube is less than ε. That implies that for any vectors z, z' in the same hypercube, $||z - z'||_\infty < \varepsilon$. □

By assumption, $\chi \subset \mathbb{R}^m$ is bounded, so let $\bar{\chi}$ represent the smallest hypercube that contains χ. Without loss of generality, we can use $\bar{\chi}$ instead, and simply set $p(i) = 0$ for all $y_i \notin \chi$. The set of weight vectors $w_i = (w_{i1} \dots w_{im})$ with $w_{min} \leq w_{ij} \leq w_{max}$ for all i, j is also a hypercube in $\mathbb{R}^m$. Since each agent i is a pair (y_i, w_i), we can write $i \in N \subset \mathbb{R}^{m^2}$, and N too is a hypercube.

Lemma 8.5.3. *The function f satisfies AFPP.*

Proof. Recall the definition of AFPP: we need to show that for any $\varepsilon' > 0$, there exists an ε'-fixed point $\mathbf{c}$ of f. Specifically, we need $|[f(\mathbf{c})](i) - c(i)| < \varepsilon'$ for all $i \in N$.

Part 1: Defining the approximate fixed point c. Fix an $\varepsilon > 0$; later on we will choose ε as a function of ε'. Using Lemma 8.5.2, let $S_1 \dots S_L$ be a partition of N into hypercubes such that for any $\ell \in \{1 \dots L\}$, for all $i, k \in S_\ell$, $||i - k||_\infty < \varepsilon$. This means that for any i, k in the same hypercube and and any $j \in M$, we have

$$|y_{ij} - y_{kj}| < \varepsilon \quad \text{and} \quad |w_{ij} - w_{kj}| < \varepsilon \tag{8.2}$$

[17] We use $||\cdot||_\infty$ to denote the L_∞ norm, which is defined to be the maximum coordinate.

For each $\ell \in \{1 \dots L\}$, let $s_\ell = \int_{i \in S_\ell} p(i)\,\mathrm{d}i$ be the measure of S_ℓ. For each S_ℓ, we will carefully pick a *representative* i_ℓ. For each ℓ, and each j, define $w_{\ell j}^{avg}$ and $y_{\ell j}^{avg}$ by

$$y_{\ell j}^{avg} = \big(\int_{i \in S_\ell} p(i) w_{ij}\,\mathrm{d}i\big)^{-1} \int_{i \in S_\ell} p(i) w_{ij} y_{ij}\,\mathrm{d}i \quad \text{and} \quad w_{\ell j}^{avg} = \frac{1}{s_\ell} \int_{i \in S_\ell} p(i) w_{ij}\,\mathrm{d}i$$

Thus for each $j \in M$, $y_{\ell j}^{avg}$ is a weighted average of y_{ij} for $i \in S_\ell$, and $w_{\ell j}^{avg}$ is a weighted average of w_{ij} for $i \in S_\ell$. In particular, $\min_{k \in S_\ell} y_{kj} \le y_{\ell j}^{avg} \le \max_{k \in S_\ell} y_{kj}$, and $\min_{k \in S_\ell} w_{kj} \le w_{\ell j}^{avg} \le \max_{k \in S_\ell} w_{kj}$. Thus since each S_ℓ is a hypercube, each S_ℓ contains an agent i_ℓ with $w_{i_\ell, j} = w_{\ell j}^{avg}$ and $y_{i_\ell j} = y_{\ell j}^{avg}$ for all $j \in M$.

Define $A = \{i_1, i_2 \dots i_L\}$, and let $\mathbf{c}_A$ be a fixed point of g_A with coefficients $s_1 \dots s_\ell$ (which is guaranteed to exist, by Lemma 8.5.1). Define $\mathbf{c} : N \to \mathbb{R}_{>0}$ so that for each $i \in S_\ell$, $c(i) = c_A(i_\ell)$. In words, we define each agent's scaling factor $c(i)$ to be the same as that of her representative.

Part 2: Properties of c. For any $j \in M$,

$$\begin{aligned}
x_j(\mathbf{c}) &= \big(\int_{k \in N} p(k) c(k) w_{kj}\,\mathrm{d}k\big)^{-1} \int_{k \in N} p(k) c(k) w_{kj} y_{kj}\,\mathrm{d}k && \text{(definition of } x_j(\mathbf{c}) \text{ for continuous } \mathbf{c}) \\
&= \big(\textstyle\sum_{\ell=1}^{L} \int_{k \in S_\ell} p(k) c(k) w_{kj}\,\mathrm{d}k\big)^{-1} \sum_{\ell=1}^{L} \int_{k \in S_\ell} p(k) c(k) w_{kj} y_{kj}\,\mathrm{d}k && \text{(summing integral over hypercubes)} \\
&= \big(\textstyle\sum_{\ell=1}^{L} c_A(i_\ell) \int_{k \in S_\ell} p(k) w_{kj}\,\mathrm{d}k\big)^{-1} \sum_{\ell=1}^{L} c_A(i_\ell) \int_{k \in S_\ell} p(k) w_{kj} y_{kj}\,\mathrm{d}k && \text{(definition of } c(k) \text{ for } k \in S_\ell) \\
&= \big(\textstyle\sum_{\ell=1}^{L} c_A(i_\ell) \int_{k \in S_\ell} p(k) w_{kj}\,\mathrm{d}k\big)^{-1} \sum_{\ell=1}^{L} c_A(i_\ell) y_{\ell j}^{avg} \int_{k \in S_\ell} p(k) w_{kj}\,\mathrm{d}k && \text{(definition of } y_{\ell j}^{avg}) \\
&= \big(\textstyle\sum_{\ell=1}^{L} c_A(i_\ell) s_\ell w_{\ell j}^{avg}\big)^{-1} \sum_{\ell=1}^{L} c_A(i_\ell) y_{\ell j}^{avg} s_\ell w_{\ell j}^{avg} && \text{(definition of } w_{\ell j}^{avg}) \\
&= \big(\textstyle\sum_{\ell=1}^{L} c_A(i_\ell) s_\ell w_{i_\ell, j}\big)^{-1} \sum_{\ell=1}^{L} c_A(i_\ell) s_\ell w_{i_\ell, j} y_{i_\ell, j} && \text{(definitions of } w_{i_\ell, j} \text{ and } y_{i_\ell, j}) \\
&= \big(\textstyle\sum_{k \in A} s_k c_A(k) w_{kj}\big)^{-1} \sum_{k \in A}^{L} s_k c_A(k) w_{kj} y_{kj} && \text{(definition of } A = \{i_1 \dots i_L\}) \\
&= x_j(\mathbf{c}_A) && \text{(definition of } x_j(\mathbf{c}) \text{ for discrete } \mathbf{c})
\end{aligned}$$

Therefore for each $j \in M$, we have $x_j(\mathbf{c}) = x_j(\mathbf{c}_A)$, where the left hand side and right hand side are continuous and discrete weighted averages, respectively.

In the process of the above sequence of equations, we also showed that $\int_{k \in N} p(k) c(k) w_{kj}\,\mathrm{d}k = \sum_{k \in A} s_k c_A(k) w_{kj}$. Using this, and the fact that $\mathbf{c}_A$ is a fixed point of g_A, for all $i \in A$ we have

$$\begin{aligned}
c_A(i) &= \sqrt{n}\Big(\max\big(\alpha, \sqrt{\textstyle\sum_{j \in M} w_{ij}^2 (y_{ij} - x_j(\mathbf{c}_A))^2 \big(\sum_{k \in A} s_k c_A(k) w_{kj}\big)^{-1}}\big)\Big)^{-1} \\
&= \sqrt{n}\Big(\max\big(\alpha, \sqrt{\textstyle\sum_{j \in M} w_{ij}^2 (y_{ij} - x_j(\mathbf{c}))^2 \big(\int_{k \in N} p(k) c(k) w_{kj}\,\mathrm{d}k\big)^{-1}}\big)\Big)^{-1} \\
&= [f(\mathbf{c})](i)
\end{aligned}$$

Since $c(i) = c_A(i)$ for $i \in A$, we therefore have $c(i) = [f(\mathbf{c})](i)$ exactly when $i \in A$.

Part 3: Showing that $|[f(c)](i) - c(i)|$ is small for every $i \in N$. For $i \in A$, we are done. For $i \notin A$, recall that for each ℓ and all $i \in S_\ell$, $c(i) = c(i_\ell)$ by definition, so $c(i) = c(i_\ell) = [f(\mathbf{c})](i_\ell)$. Therefore it suffices to show that $|[f(\mathbf{c})](i) - [f(\mathbf{c})](i_\ell)|$ is small.

For brevity, let $r(i) = \sum_{j \in M} w_{ij}^2 (y_{ij} - x_j(\mathbf{c}))^2 \big(\int_{k \in N} p(k) c(k) w_{kj} \, \mathrm{d}k \big)^{-1}$. Then for all $i \in S_\ell$,

$$
\begin{aligned}
|r(i) - r(i_\ell)| &= \Big(\int_{k \in N} p(k) c(k) w_{kj} \, \mathrm{d}k \Big)^{-1} \Big| \sum_{j \in M} \big(w_{ij}^2 (y_{ij} - x_j(\mathbf{c}))^2 - w_{i_\ell, j}^2 (y_{i_\ell j} - x_j(\mathbf{c}))^2 \big) \Big| && \text{(defn of } r(i)\text{)} \\
&\leq \Big(\int_{k \in N} p(k) c(k) w_{kj} \, \mathrm{d}k \Big)^{-1} \sum_{j \in M} \big| w_{ij}^2 (y_{ij} - x_j(\mathbf{c}))^2 - w_{i_\ell, j}^2 (y_{i_\ell j} - x_j(\mathbf{c}))^2 \big| && \text{(triangle ineq.)} \\
&\leq \Big(\int_{k \in N} p(k) c(k) w_{kj} \, \mathrm{d}k \Big)^{-1} \sum_{j \in M} \big| w_{ij}^2 (y_{ij} - x_j(\mathbf{c}))^2 - (w_{ij} + \varepsilon)^2 (y_{ij} - x_j(\mathbf{c}) + \varepsilon)^2 \big| && \text{(Eq. 8.2)} \\
&= \Big(\int_{k \in N} p(k) c(k) w_{kj} \, \mathrm{d}k \Big)^{-1} \sum_{j \in M} \big| \varepsilon^4 + \varepsilon^2 \big(w_{ij}^2 (y_{ij} - x_j(\mathbf{c}))^2 \big) \big| && \text{(cancel terms)} \\
&\leq \Big(\int_{k \in N} p(k) c(k) w_{kj} \, \mathrm{d}k \Big)^{-1} \sum_{j \in M} (\varepsilon^4 + \varepsilon^2 w_{max}^2 d_{max}^2) && \text{(defn's of } d_{max}, w_{max}\text{)} \\
&\leq \Big(\int_{k \in N} p(k) c(k) w_{min} \, \mathrm{d}k \Big)^{-1} m (\varepsilon^4 + \varepsilon^2 w_{max}^2 d_{max}^2) && \text{(defn of } w_{min}\text{)} \\
&\leq \Big(\int_{k \in N} p(k) \frac{w_{min}}{w_{max}^2 d_{max}^2} w_{min} \, \mathrm{d}k \Big)^{-1} m (\varepsilon^4 + \varepsilon^2 w_{max}^2 d_{max}^2) && \text{(Lemma 8.5.1)} \\
&\leq O(1) \cdot m \cdot (\varepsilon^4 + \varepsilon^2 \cdot O(1)) \\
&= O(\varepsilon^2)
\end{aligned}
$$

Our citation of Lemma 8.5.1 is because Lemma 8.5.1 guarantees that every $c(k) \geq \frac{w_{min}}{w_{max}^2 d_{max}^2}$. Thus for each $i \in S_\ell$, $|r(i) - r(i_\ell)| \leq O(\varepsilon^2)$.

If $[f(\mathbf{c})](i) = [f(\mathbf{c})](i_\ell) = 1/\alpha$, then $c(i) = [f(\mathbf{c})](i_\ell) = [f(\mathbf{c})](i)$, and we are done. Thus assume at least one does not equal $1/\alpha$. Suppose $f[(\mathbf{c})](i) \neq 1/\alpha$ (the argument is symmetric for the other case), and basic algebra gives us

$$
\begin{aligned}
|[f(\mathbf{c})](i) - [f(\mathbf{c})](i_\ell)| &= \Big| \frac{\sqrt{n}}{\sqrt{r(i)}} - \frac{\sqrt{n}}{\max(\alpha, \sqrt{r(i_\ell)})} \Big| \\
&\leq \sqrt{n} \Big| \frac{1}{\sqrt{r(i)}} - \frac{1}{\sqrt{r(i_\ell)}} \Big| \\
&= \frac{\sqrt{n}}{\sqrt{r(i) r(i_\ell)}} \Big| \sqrt{r(i_\ell))} - \sqrt{(r(i)} \Big| \\
&= \frac{\sqrt{n}}{\sqrt{r(i) r(i_\ell)} (\sqrt{r(i)} + \sqrt{r(i_\ell)})} |r(i_\ell) - r(i)| \\
&\leq \frac{\sqrt{n} \cdot O(\varepsilon^2)}{\sqrt{r(i) r(i_\ell)} (\sqrt{r(i)} + \sqrt{r(i_\ell)})}
\end{aligned}
$$

Since $r(i), r(i_\ell) \leq \alpha$, we have

$$
|[f(\mathbf{c})](i) - [f(\mathbf{c})](i_\ell)| \leq \frac{O(\sqrt{n} \varepsilon^2)}{\alpha^3}
$$

and thus $|[f(\mathbf{c})](i) - c(i)| \leq O(\sqrt{n} \varepsilon^2)/\alpha^3$. Now, for a fixed ε, taking $n \to \infty$ does make this bound go to infinity. The key here is that ε can be chosen independently of α and n. That is, for any instantiation of f (i.e., for a fixed n and α), we can pick ε to be as small as we want. In particular, since this bound holds uniformly for all $i \in N$ (i.e., with the same ε), for any $\varepsilon' > 0$, there exists ε such that

$$
|[f(\mathbf{c})](i) - c(i)| < \varepsilon'
$$

for all $i \in N$. We conclude that f satisfies AFPP. $\square$

8.6 Two proofs omitted earlier

With Lemma 8.5.3 in hand, the bulk of the remaining work is to prove Theorem 8.4.1. Before embarking on that task, we provide two quick proofs that were omitted from Section 8.4.

Lemma 8.6.1. *Suppose f as defined in Section 8.5 has an exact fixed point $\mathbf{c}$ for any choice of α and n. Let $x_j = \left(\int_{i\in N} p(i)c_i w_{ij}\,\mathrm{d}i\right)^{-1} \int_{i\in N} p(i)c_i w_{ij} y_{ij}\,\mathrm{d}i$. Let $\alpha = n^{-7/8}$ and $m \geq 6$.*

1. *$||n \int_{i\in N} p(i)\delta_i(x)\,\mathrm{d}i||_2 \leq (\gamma + \sqrt{\gamma})O(n^{-1/4}) + O(n^{-1/8})$.*
2. *For all i except an expected $O(n^{-3/4})$ fraction, $O\Big(\sqrt{\frac{n^{5/8}}{n^{5/8}+1}}\Big)\frac{1}{\sqrt{\gamma}} \leq \lambda_i(x) \leq \frac{1}{\sqrt{\gamma}}$.*
3. *The outcome x maximizes welfare with respect to $\mathbf{c}$.*

Proof. Assume Theorem 8.4.1 holds. We need to show the following:

1. $\lim_{n\to\infty} \alpha^{m/2} n^{m/4} = 0$
2. $\lim_{n\to\infty} n^{3/2}\alpha = \infty$
3. $\lim_{n\to\infty} \alpha^{m/2} n^{m/4+1} = 0$
4. $\lim_{n\to\infty} \alpha^2 n^2 = \infty$

The first two points are necessary as conditions of Theorem 8.4.1, in addition to being necessary for vanishing approximation error. Note that the third point implies the first 1.

Since $m \geq 6$, $O(\alpha^{m/2} n^{m/4+1})$ becomes $O(n^{-3m/16+1}) = O(n^{-18/16+1}) = O(n^{-1/8})$, which does indeed go to zero as $n \to \infty$. This satisfies points 1 and 3.

For the second point, we have $O(n^{3/2}\alpha) = O(n^{3/2-7/8}) = O(n^{5/8})$. Thus $O(n^{3/2}\alpha)$ does go to infinity as $n \to \infty$.

Finally, $\alpha^2 n^2 = n^{2/8} = n^{1/4}$, which goes to infinity as $n \to \infty$. $\square$

Theorem 8.4.2. *Let $\mathbf{c}$ be an ε-fixed point of f and let $x_j = \left(\int_{i\in N} p(i)c_i w_{ij}\,\mathrm{d}i\right)^{-1} \int_{i\in N} p(i)c_i w_{ij} y_{ij}\,\mathrm{d}i$. Let $\alpha = n^{-7/8}$ and $m \geq 6$. Then there exists a small enough ε such that all of the following hold:*

1. *$||n \int_{i\in N} p(i)\delta_i(x)\,\mathrm{d}i||_2 \leq (\gamma + \sqrt{\gamma})O(n^{-1/4}) + O(n^{-1/8})$.*
2. *For all i except an expected $O(n^{-3/4})$ fraction, $O\Big(\sqrt{\frac{n^{5/8}}{n^{5/8}+1}}\Big)\frac{1}{\sqrt{\gamma}} \leq \lambda_i(x) \leq \frac{1}{\sqrt{\gamma}}$.*
3. *The outcome x maximizes welfare with respect to $\mathbf{c}$.*

Proof. An ε-fixed point is guaranteed to exist by Lemma 8.5.3. The key here is that ε can be chosen independently of any other parameters (α, n, γ, etc). Furthermore, $\lambda_i(x)$ and $||n \int_{i\in N} p(i)\delta_i(x)\,\mathrm{d}i||_2$ will be continuous functions of $\mathbf{c}$. Thus for any ε', there exists ε such that perturbing $\mathbf{c}$ by ε' changes both $||n \int_{i\in N} p(i)\delta_i(x)\,\mathrm{d}i||_2$ and each λ_i by at most ε. Since the results are only asymptotic anyway, we can pick ε' small enough that all of the approximations still hold. $\square$

8.7 Important properties to be used later

The rest of the chapter is devoted to proving Theorem 8.4.1. Throughout, we assume that $\mathbf{c}$ is a fixed point of the function f from Section 8.5, and that $x_j = \left(\int_{i\in N} p(i)c_i w_{ij}\,\mathrm{d}i\right)^{-1} \int_{i\in N} p(i)c_i w_{ij} y_{ij}\,\mathrm{d}i$. Recall that we used the function notation of $c(i)$ only for Section 8.5; we use the vector notation c_i for the rest of the chapter.

8.7.1 Welfare maximization for quadratic utilities

Lemma 8.7.1. *The outcome x maximizes the utilitarian welfare $U(x)$ with respect to* $\mathbf{c}$.

Proof. The welfare of an outcome x' is

$$U(x') = n\int_{i\in N} p(i)u_i(x')\,\mathrm{d}i = -n\int_{i\in N} p(i)c_i \sum_{j\in M} w_{ij}(x'_j - y_{ij})^2\,\mathrm{d}i$$

Since U is concave and differentiable, and we are interested in an unconstrained maximization, it suffices to show that the gradient of U at x is 0. Specifically, the partial derivative with respect to x_j should be zero for each j:

$$\begin{aligned}\frac{\partial}{\partial x_j}U(x) = &\ -2n\int_{i\in N} p(i)c_i w_{ij}(x_j - y_{ij})\,\mathrm{d}i\\ = &\ \int_{i\in N} p(i)c_i w_{ij}x_j\,\mathrm{d}i - \int_{i\in N} p(i)c_i w_{ij}y_{ij}\,\mathrm{d}i\\ = &\ x_j\int_{i\in N} p(i)c_i w_{ij}\,\mathrm{d}i - \int_{i\in N} p(i)c_i w_{ij}y_{ij}\,\mathrm{d}i\end{aligned}$$

Substituting in the definition of x_j as given in the statement of Theorem 8.4.1:

$$\begin{aligned}\frac{\partial}{\partial x_j}U(x) = &\ \left((\textstyle\int_{k\in N} p(k)c_k w_{kj}\,\mathrm{d}k)^{-1}\int_{k\in N} p(k)c_k w_{kj}y_{kj}\,\mathrm{d}k\right)\int_{i\in N} p(i)c_i w_{ij}\,\mathrm{d}i - \int_{i\in N} p(i)c_i w_{ij}y_{ij}\,\mathrm{d}i\\ = &\ \int_{i\in N} p(i)c_i w_{ij}y_{ij}\,\mathrm{d}i - \int_{i\in N} p(i)c_i w_{ij}y_{ij}\,\mathrm{d}i\\ = &\ 0\end{aligned}$$

Thus x is indeed optimal for U. □

8.7.2 The measure of "unusual" agents is small

Let $\hat{N}$ be the set of agents i for whom $c_i = \sqrt{n}/\sqrt{\sum_{j\in M} w_{ij}^2(y_{ij}-x_j)^2\left(\int_{k\in N} p(k)c_k w_{kj}\,\mathrm{d}k\right)^{-1}}$. In ways that will become clear later, $\hat{N}$ is the set of "normal" agents. We need to show that the number of agents not in $\hat{N}$ shrinks to zero as α goes to zero. Here Γ denotes the gamma function.

Lemma 8.7.2. *We have* $\displaystyle\int_{i\notin\hat{N}} p(i)\,\mathrm{d}i \le \alpha^{m/2}n^{m/4}\frac{\pi^{m/2}p_{max}}{\Gamma(\frac{m}{2}+1)}\left(\frac{\sqrt{w_{max}}}{w_{min}}\right)^m$.

Proof. Since $\mathbf{c}$ is a fixed point of f, if $i\notin\hat{N}$, we have $\alpha \ge \sqrt{\sum_{j\in M} w_{ij}^2(y_{ij}-x_j)^2\left(\int_{k\in N} p(k)c_k w_{kj}\,\mathrm{d}k\right)^{-1}}$.

Also, by the definition of f, we have $c_k \leq \sqrt{n}/\alpha$ for all $k \in N$. Thus for an arbitrary $i \in N$ we have

$$\begin{aligned}
\sqrt{\sum_{j\in M} w_{ij}^2 (y_{ij}-x_j)^2 \Big(\int_{k\in N} p(k) c_k w_{kj}\,\mathrm{d}k\Big)^{-1}} &\geq \sqrt{\sum_{j\in M} w_{ij}^2 (y_{ij}-x_j)^2 \Big(\int_{k\in N} p(k) \frac{\sqrt{n}}{\alpha} w_{kj}\,\mathrm{d}k\Big)^{-1}} \\
&= \frac{\sqrt{\alpha}}{n^{1/4}} \sqrt{\sum_{j\in M} w_{ij}^2 (y_{ij}-x_j)^2 \Big(\int_{k\in N} p(k) w_{kj}\,\mathrm{d}k\Big)^{-1}} \\
&\geq \frac{\sqrt{\alpha}}{n^{1/4}} \sqrt{\sum_{j\in M} w_{min}^2 (y_{ij}-x_j)^2 \Big(\int_{k\in N} p(k) w_{max}\,\mathrm{d}k\Big)^{-1}} \\
&= \frac{\sqrt{\alpha} w_{min}}{n^{1/4}\sqrt{w_{max}}} \sqrt{\sum_{j\in M} (y_{ij}-x_j)^2 \Big(\int_{k\in N} p(k)\,\mathrm{d}k\Big)^{-1}} \\
&\geq \frac{\sqrt{\alpha} w_{min}}{n^{1/4}\sqrt{w_{max}}} \sqrt{\sum_{j\in M} (y_{ij}-x_j)^2} \\
&= \frac{\sqrt{\alpha} w_{min}}{n^{1/4}\sqrt{w_{max}}} ||y_i - x||_2
\end{aligned}$$

Thus if $i \notin \hat{N}$,

$$\begin{aligned}
\alpha &\geq \frac{\sqrt{\alpha}}{n^{1/4}} ||y_i - x||_2 \frac{w_{min}}{\sqrt{w_{max}}} \\
||y_i - x||_2 &\leq \frac{\sqrt{\alpha w_{max}} n^{1/4}}{w_{min}}
\end{aligned}$$

Therefore $i \notin \hat{N}$ only if $||y_i - x||_2 \leq \frac{\sqrt{\alpha w_{max}} n^{1/4}}{w_{min}} w_{min}$. Let B denote the ball with radius $\frac{\sqrt{\alpha w_{max}} n^{1/4}}{w_{min}}$ centered at x. Then we have $\int_{i\notin\hat{N}} p(i)\,\mathrm{d}i \leq \int_{i:y_i\in B} p(i)\,\mathrm{d}i$.

Since $p(i) \leq p_{max}$ for all $i \in N$ by assumption, we have

$$\int_{i:y_i\in B} p(i)\,\mathrm{d}i \leq p_{max} \int_{i:y_i\in B} \mathrm{d}i$$

Since $\int_{i:y_i\in B} \mathrm{d}i$ is just the volume of the m-dimensional unit ball with radius $\frac{\sqrt{\alpha w_{max}} n^{1/4}}{w_{min}}$, we have

$$\int\limits_{i:y_i\in B} \mathrm{d}i = \frac{\pi^{m/2}}{\Gamma(\frac{m}{2}+1)} \Big(\frac{\sqrt{\alpha w_{max}} n^{1/4}}{w_{min}}\Big)^m$$

□

Since we treat m, p_{max}, w_{min} and w_{max} as constants, we can simply write

$$\int_{i\notin\hat{N}} p(i)\,\mathrm{d}i = O(\alpha^{m/2} n^{m/4})$$

8.7.3 Each agent is a small fraction of the population

In this section, we show that the weight contribution by any single agent on any issue (i.e., $c_i w_{ij}$) is a small fraction of the weight of the internet population. The main result is Lemma 8.7.4; we first prove Lemma 8.7.3, which lower bounds the aggregate weight of the whole population on issue j.

Lemma 8.7.3. *For all* $j \in M$, $\int_{k\in N} p(k)c_k w_{kj}\,\mathrm{d}k \geq \frac{w_{min}^2}{4w_{max}^2 d_{max}^2} n$.

Proof. Since $\mathbf{c}$ is assumed to be a fixed point of f, for all $j \in M$ we have

$$\begin{aligned}
\int_{k\in N} p(k)c_k w_{kj}\,\mathrm{d}k &\geq w_{min}\int_{k\in N} p(k)c_k\,\mathrm{d}k \\
&= w_{min}\int_{k\in\hat{N}} \frac{\sqrt{n}p(k)\,\mathrm{d}k}{\max\left(\alpha, \sqrt{\sum_{\ell\in M} w_{k\ell}^2 (y_{k\ell}-x_\ell)^2 \left(\int_{i\in N} p(i)c_i w_{i\ell}\,\mathrm{d}i\right)^{-1}}\right)} \\
&\geq w_{min}\sqrt{n}\int_{k\in N} \frac{p(k)\,\mathrm{d}k}{\sqrt{\sum_{\ell\in M} w_{k\ell}^2 (y_{k\ell}-x_\ell)^2 \left(\int_{i\in N} p(i)c_i w_{i\ell}\,\mathrm{d}i\right)^{-1}}} \\
&\geq w_{min}\sqrt{n}\int_{k\in\hat{N}} \frac{p(k)\,\mathrm{d}k}{\sqrt{\sum_{\ell\in M} w_{k\ell}^2 (y_{k\ell}-x_\ell)^2 \left(\int_{i\in N} p(i)c_i w_{min}\,\mathrm{d}i\right)^{-1}}} \\
&\geq w_{min}^{3/2}\sqrt{n}\sqrt{\int_{i\in N} p(i)c_i\,\mathrm{d}i}\int_{k\in\hat{N}} \frac{p(k)\,\mathrm{d}k}{\sqrt{\sum_{\ell\in M} w_{k\ell}^2 (y_{k\ell}-x_\ell)^2}} \\
&\geq \frac{w_{min}^{3/2}}{w_{max}}\sqrt{n}\sqrt{\int_{k\in N} p(k)c_k\,\mathrm{d}k}\int_{k\in\hat{N}} \frac{p(k)\,\mathrm{d}k}{\sqrt{\sum_{\ell\in M} (y_{k\ell}-x_\ell)^2}} \\
&\geq \frac{w_{min}^{3/2}}{w_{max}d_{max}}\sqrt{n}\sqrt{\int_{k\in N} p(k)c_k\,\mathrm{d}k}\int_{k\in\hat{N}} p(k)\,\mathrm{d}k
\end{aligned}$$

By Lemma 8.7.2, $\int_{k\notin\hat{N}} p(k)\,\mathrm{d}k = O(\alpha^{m/2}n^{m/4})$, and we know that $\lim_{n\to\infty}\alpha^{m/2}n^{m/4} = 0$ by assumption of Theorem 8.4.1. Thus for large enough n, $\int_{k\in\hat{N}} p(k)\,\mathrm{d}k \geq 1/2$, so

$$\begin{aligned}
w_{min}\int_{k\in N} p(k)c_k\,\mathrm{d}k &\geq \frac{w_{min}^{3/2}}{w_{max}d_{max}}\sqrt{n}\sqrt{\int_{k\in N} p(k)c_k\,\mathrm{d}k}\int_{k\in\hat{N}} p(k)\,\mathrm{d}k \\
&\geq \frac{w_{min}^{3/2}}{2w_{max}d_{max}}\sqrt{n}\sqrt{\int_{k\in N} p(k)c_k\,\mathrm{d}k}
\end{aligned}$$

Therefore $w_{min}\int_{k\in N} p(k)c_k\,\mathrm{d}k \geq \frac{w_{min}^{3/2}}{2w_{max}d_{max}}\sqrt{n}\sqrt{\int_{k\in N} p(k)c_k\,\mathrm{d}k}$, so

$$\begin{aligned}
\sqrt{\int_{k\in N} p(k)c_k\,\mathrm{d}k} &\geq \frac{w_{min}^{1/2}}{2w_{max}d_{max}}\sqrt{n} \\
\int_{k\in N} p(k)c_k\,\mathrm{d}k &\geq \frac{w_{min}}{4w_{max}^2 d_{max}^2} n
\end{aligned}$$

Therefore

$$\int_{k\in N} p(k)c_k w_{kj}\,\mathrm{d}k \geq \frac{w_{min}^2}{4w_{max}^2 d_{max}^2} n$$

as required.

□

Lemma 8.7.4. *For any agent $i \in N$,* $\displaystyle\int_{k \in N} p(k) c_k w_{kj} \, \mathrm{d}k \geq \frac{w_{min}^2}{4 w_{max}^3 d_{max}^2} \sqrt{n} \alpha c_i w_{ij}$.

Proof. Since $c_i \leq \sqrt{n}/\alpha$ for all $i \in N$ and $w_{ij} \leq w_{max}$, the right hand side is at most $\frac{w_{min}^2}{4 w_{max}^2 d_{max}^2} n$. Applying Lemma 8.7.3 completes the proof. □

For brevity, let $\beta = \frac{w_{min}^2}{4 w_{max}^3 d_{max}^2}$. Thus $\int_{k \in N} p(k) c_k w_{kj} \, \mathrm{d}k \geq \beta \sqrt{n} \alpha c_i w_{ij}$.

8.8 Characterizing δ_i and λ_i

In this section, we derive an expression for δ_i in terms of λ_i, then provide almost tight upper and lower bounds on λ_i. Lemma 8.8.2 gives the expression for δ_i in terms of λ_i, and Lemma 8.8.6 gives the upper and lower bounds for λ_i.

8.8.1 Deriving an expression for δ_i in terms of λ_i

First, we show that the equal power constraint of Program 8.1 can be reduced to a simpler form.

Lemma 8.8.1. *Then the equal power constraint is equivalent to*

$$\sum_{j \in M} \delta_{ij}^2 \big(n \textstyle\int_{k \in N} p(k) c_k w_{kj} \, \mathrm{d}k \big) \leq \gamma$$

Proof. For quadratic utilities, we can rewrite $U(x + \delta_i)$ as

$$U(x + \delta_i) = U(x) + \delta_i^T (\nabla_x U)(x) + \frac{1}{2} \delta_i^T (\nabla_x^2 U)(x) \delta_i$$

where the gradient is just with respect to x. For brevity, we will omit the parentheses and just write $\nabla_x U(x)$ etc. By Lemma 8.7.1, x maximizes U with respect to $\mathbf{c}$. Therefore $\nabla_x U(x) = 0$, so

$$\begin{aligned} U(x) - U(x + \delta_i) &= U(x) - \big(U(x) + \delta_i^T \nabla_x U(x) + \frac{1}{2} \delta_i^T \nabla_x^2 U(x) \delta_i \big) \\ &= -\frac{1}{2} \delta_i^T \nabla_x^2 U(x) \delta_i \end{aligned}$$

This makes the constraint of equal power reduce to

$$-\frac{1}{2} \delta_i^T \nabla_x^2 U(x) \delta_i \leq \gamma$$

Next, recall that $U(x) = n \int_{i \in N} p(i) u_i(x) \, \mathrm{d}i$ by definition. Therefore

$$\frac{\partial}{\partial x_j} U(x) = \frac{\partial}{\partial x_j} \Big(-n \textstyle\int_{k \in N} p(k) \big(c_k \sum_{j \in M} w_{kj} (x_j - y_{kj})^2 \big) \, \mathrm{d}k \Big) = -2n \int_{k \in N} p(k) c_k w_{kj} (x_j - y_{kj}) \, \mathrm{d}k$$

which means that $\frac{\partial^2}{\partial x_j \partial x_\ell} = 0$ whenever $j \neq \ell$, for $j = \ell$ we have

$$\frac{\partial^2}{\partial x_j^2} U(x) = -2n \int_{k \in N} p(k) c_k w_{kj} \, \mathrm{d}k$$

Thus $\nabla_x^2 U(x)$ is a diagonal matrix with entry $n \int_{k \in N} p(k) c_k w_{kj} \, \mathrm{d}k$ in the jth row, so the equal power constraint simplifies to $\sum_{j \in M} \delta_{ij}^2 \big(n \int_{k \in N} p(k) c_k w_{kj} \, \mathrm{d}k \big) \leq \gamma$, as required. □

For brevity, define q_j by

$$q_j = n \int_{k \in N} p(k) c_k w_{kj} \, \mathrm{d}k$$

Thus the equal power constraint is equivalent to $\sum_{j \in M} q_j \delta_{ij}^2 \leq \gamma$. We can think of q_j as how much the population in aggregate cares about issue j. The expression $\sum_{j \in M} q_j \delta_{ij}^2$ indicates that the more the population cares about issue j, the more power it required to move on that issue.

Lemma 8.8.2. *For every agent i and issue j, $\delta_{ij} = \frac{c_i w_{ij}(y_{ij} - x_j)}{c_i w_{ij} + \lambda_i q_j}$.*

Proof. Define the Lagrangian by

$$\begin{aligned} L(\delta_i, \lambda_i) &= u_i(x + \delta_i) - \lambda_i \big(U(x) - U(x + \delta_i) - \gamma \big) \\ &= \; - c_i \textstyle\sum_{j \in M} w_{ij} (x + \delta_i - y_{ij})^2 - \lambda_i \big(\sum_{j \in M} q_j \delta_{ij}^2 - \gamma \big) \end{aligned}$$

where the second line used Lemma 8.8.1.

For $\delta_i = \mathbf{0}$, we have $U(x) - U(x + \delta_i) = 0 < \gamma$, so we have strong duality by Slater's condition. Since the convex is convex, the optimal solution must satisfy the KKT conditions; in particular, stationarity:

$$\frac{\partial}{\partial \delta_{ij}} L(\delta_i, \lambda_i) = 0$$

for all $j \in M$. Therefore for all j,

$$\begin{aligned} -2c_i w_{ij}(x_j + \delta_{ij} - y_{ij}) - 2\lambda_i q_j \delta_{ij} &= 0 \\ (c_i w_{ij} + \lambda_i q_j)\delta_{ij} + c_i w_{ij}(x_j - y_{ij}) &= 0 \\ \delta_{ij} &= \frac{c_i w_{ij}(y_{ij} - x_j)}{c_i w_{ij} + \lambda_i q_j} \end{aligned}$$

as required. □

8.8.2 Bounding λ_i

For each agent i, we have $m + 1$ unknowns: $\delta_{i1} \dots \delta_{im}$, and λ_i. The previous section gave us m equations: one for each δ_{ij}. For our last equation, we show that we can pick γ small enough such that for most of the agents, the equal power constraint holds with equality. Specifically, the power constraint will hold with equality for every agent in $\hat{N}$ (and we know $N \setminus \hat{N}$ to have small measure by Lemma 8.7.2).

Lemma 8.8.3. *There exists a γ such that the power constraint is tight for all $i \in \hat{N}$, i.e., $\sum_{j\in M} q_j \delta_{ij}^2 = \gamma$.*

Proof. For all agents $i \in \hat{N}$, we know that

$$\begin{aligned}
\alpha &\leq \sqrt{\textstyle\sum_{j\in M} w_{ij}^2 (y_{ij} - x_j)^2 \big(\int_{k\in N} p(k) c_k w_{kj}\, \mathrm{d}k\big)^{-1}} \\
&\leq \sqrt{\textstyle\sum_{j\in M} w_{max}^2 (y_{ij} - x_j)^2 \big(\int_{k\in N} p(k) \frac{w_{min}}{w_{max}^2 d_{max}^2} w_{min}\, \mathrm{d}k\big)^{-1}} \\
&\leq \frac{w_{max}^2 d_{max}}{w_{min}} \sqrt{\textstyle\sum_{j\in M} (y_{ij} - x_j)^2 \big(\int_{k\in N} p(k)\, \mathrm{d}k\big)^{-1}} \\
&\leq \frac{w_{max}^2 d_{max}}{w_{min}} ||y_i - x||_2
\end{aligned}$$

Thus $||y_i - x||_2 \geq \frac{\alpha w_{min}}{w_{max}^2 d_{max}}$. Suppose for sake of contradiction that for all $\gamma > 0$, there is an agent $i \in \hat{N}$ whose power constraint is not tight. That would imply that there are agents in $\hat{N}$ arbitrarily close to x, which we just showed is not true. We conclude that there exists a $\gamma > 0$ such that $\sum_{j\in M} q_j \delta_{ij}^2 = \gamma$ for all $i \in \hat{N}$. □

Note that if Lemma 8.8.3 holds for some $\gamma > 0$, it also holds for any $\gamma' \in (0, \gamma]$. In particular, later on we will require that $\gamma \leq 1$.

Recall that $\hat{N}$ is the set of agents i for whom $c_i = \sqrt{n}/\sqrt{\sum_{j\in M} w_{ij}^2 (y_{ij} - x_j)^2 \big(\int_{k\in N} p(k) c_k w_{kj}\, \mathrm{d}k\big)^{-1}}$. This implies $c_i = 1/\sqrt{\sum_{j\in M} w_{ij}^2 (y_{ij} - x_j)^2 \big(n \int_{k\in N} p(k) c_k w_{kj}\, \mathrm{d}k\big)^{-1}} = 1/\sqrt{\sum_{j\in M} w_{ij}^2 (y_{ij} - x_j)^2 q_j^{-1}}$. For each $i \in \hat{N}$, define τ_i by

$$\tau_i = \frac{c_i^2}{\lambda_i^2} \sum_{j\in M} w_{ij}^2 (y_{ij} - x_j)^2 q_j^{-1} - \sum_{j\in M} q_j \delta_{ij}^2$$

By Lemma 8.8.3,

$$\frac{c_i^2}{\lambda_i^2} \sum_{j\in M} w_{ij}^2 (y_{ij} - x_j)^2 q_j^{-1} = \tau_i + \gamma$$

We know that $c_i = 1/\sqrt{\sum_{j\in M} w_{ij}^2 (y_{ij} - x_j)^2 q_j^{-1}}$ for every $i \in \hat{N}$, so $\frac{1}{\lambda_i^2} = \tau_i + \gamma$, and therefore

$$\lambda_i = \frac{1}{\sqrt{\tau_i + \gamma}}$$

8.8.3 Bounding τ_i

Our goal is to show that for all $i \in \hat{N}$, λ_i is close to $1/\sqrt{\gamma}$. To do this, we need to show that τ_i is small and nonnegative. Nonnegativity is (much) easier, so we begin with that.

Lemma 8.8.4. *For all $i \in \hat{N}$, $\tau_i > 0$.*

Proof. Plugging in the expression for δ_{ij} from Lemma 8.8.2, we have

$$\begin{aligned}\sum_{j\in M} q_j\delta_{ij}^2 &= \sum_{j\in M}\frac{q_j c_i^2 w_{ij}^2 (y_{ij}-x_j)^2}{(c_i w_{ij}+\lambda_i q_j)^2}\\ &< \sum_{j\in M}\frac{q_j c_i^2 w_{ij}^2 (y_{ij}-x_j)^2}{(\lambda_i q_j)^2}\\ &= \frac{c_i^2}{\lambda_i^2}\sum_{j\in M} w_{ij}^2 (y_{ij}-x_j)^2 q_j^{-1}\end{aligned}$$

Therefore

$$\tau_i > \frac{c_i^2}{\lambda_i^2}\sum_{j\in M} w_{ij}^2(y_{ij}-x_j)^2 q_j^{-1} - \frac{c_i^2}{\lambda_i^2}\sum_{j\in M} w_{ij}^2(y_{ij}-x_j)^2 q_j^{-1} = 0$$

□

Next, we prove an upper bound on τ_i. The upper bound expression in Lemma 8.8.5 appears quite complicated, but notice that the denominator contains $n^{3/2}\alpha$, and we know that $\lim_{n\to\infty} n^{3/2}\alpha = \infty$.

Lemma 8.8.5. *For each agent $i \in \hat{N}$, $\tau_i \le \frac{2\gamma^{3/2} m w_{max}}{n^{3/2}\alpha\beta w_{min} - 2\sqrt{\gamma} m w_{max}}$.*

Proof. We begin the proof with some algebraic manipulations on the definition of τ_i:

$$\begin{aligned}\tau_i &= \frac{c_i^2}{\lambda_i^2}\sum_{j\in M} w_{ij}^2(y_{ij}-x_j)^2 q_j^{-1} - \sum_{j\in M} q_j\delta_{ij}^2\\ &= \frac{c_i^2}{\lambda_i^2}\sum_{j\in M} w_{ij}^2(y_{ij}-x_j)^2 q_j^{-1} - \sum_{j\in M} q_j\Big(\frac{c_i w_{ij}(y_{ij}-x_j)}{c_i w_{ij}+\lambda_i q_j}\Big)^2\\ &= c_i^2\sum_{j\in M} w_{ij}^2(y_{ij}-x_j)^2\Big(\frac{1}{\lambda_i^2 q_j} - \frac{q_j}{(c_i w_{ij}+\lambda_i q_j)^2}\Big)\\ &= c_i^2\sum_{j\in M} w_{ij}^2(y_{ij}-x_j)^2\frac{(c_i w_{ij}+\lambda_i q_j)^2 - \lambda_i^2 q_j^2}{\lambda_i^2 q_j (c_i w_{ij}+\lambda_i q_j)^2}\\ &= c_i^2\sum_{j\in M} w_{ij}^2(y_{ij}-x_j)^2\frac{c_i^2 w_{ij}^2 + \lambda_i^2 q_j^2 + 2c_i w_{ij}\lambda_i q_j - \lambda_i^2 q_j^2}{\lambda_i^2 q_j (c_i w_{ij}+\lambda_i q_j)^2}\\ &= c_i^2\sum_{j\in M} w_{ij}^2(y_{ij}-x_j)^2\frac{c_i^2 w_{ij}^2 + 2c_i w_{ij}\lambda_i q_j}{\lambda_i^2 q_j (c_i w_{ij}+\lambda_i q_j)^2}\\ &= \frac{c_i^3}{\lambda_i^2}\sum_{j\in M} w_{ij}^3(y_{ij}-x_j)^2\frac{c_i w_{ij}+2\lambda_i q_j}{q_j(c_i w_{ij}+\lambda_i q_j)^2}\\ &= \frac{2c_i^3}{\lambda_i^2}\sum_{j\in M} w_{ij}^3(y_{ij}-x_j)^2\frac{c_i w_{ij}+2\lambda_i q_j}{q_j(c_i w_{ij}+\lambda_i q_j)(2c_i w_{ij}+2\lambda_i q_j)}\end{aligned}$$

$$\leq \frac{2c_i^3}{\lambda_i^2} \sum_{j \in M} w_{ij}^3 (y_{ij} - x_j)^2 \frac{1}{q_j(c_i w_{ij} + \lambda_i q_j)}$$

By Lemma 8.8.1 (and the definition of q_j), each agent's power constraint is $\sum_{j \in M} \delta_{ij}^2 q_j \leq \gamma$. This means that for all $i \in N$ and $j \in M$, $\sqrt{\delta_{ij}} \leq \sqrt{\gamma} q_j^{-1/2}$. Also recall that $\delta_{ij} = \frac{c_i w_{ij}(y_{ij} - x_j)}{c_i w_{ij} + \lambda_i q_j}$ (Lemma 8.8.2). Therefore

$$\begin{aligned}
\tau_i &\leq \frac{2c_i^3}{\lambda_i^2} \sum_{j \in M} w_{ij}^3 (y_{ij} - x_j)^2 \frac{1}{q_j(c_i w_{ij} + \lambda_i q_j)} \\
&= \frac{2c_i^2}{\lambda_i^2} \sum_{j \in M} w_{ij}^2 (y_{ij} - x_j) \delta_{ij} q_j^{-1} \\
&\leq \frac{2c_i^2 \sqrt{\gamma}}{\lambda_i^2} \sum_{j \in M} w_{ij}^2 |y_{ij} - x_j| q_j^{-1/2} q_j^{-1}
\end{aligned}$$

Next, let $\Delta_i = \max_{j \in M} |y_{ij} - x_j| q_j^{-1/2}$. Recall that $c_i = 1/\sqrt{\sum_{j \in M} w_{ij}^2 (y_{ij} - x_j)^2 q_j^{-1}}$ for every $i \in \hat{N}$. Since $w_{ij} \geq w_{min}$, we have

$$\sqrt{\sum_{j \in M} w_{ij}^2 (y_{ij} - x_j)^2 q_j^{-1}} \geq w_{min} \sqrt{\sum_{j \in M} (y_{ij} - x_j)^2 q_j^{-1}} \geq w_{min} \Delta_i$$

Thus $c_i \leq \frac{1}{w_{min} \Delta_i}$ for all $i \in \hat{N}$. Therefore

$$\tau_i \leq \frac{2c_i \sqrt{\gamma}}{w_{min} \Delta_i \lambda_i^2} \sum_{j \in M} w_{ij}^2 \Delta_i q_j^{-1} = \frac{2c_i \sqrt{\gamma}}{w_{min} \lambda_i^2} \sum_{j \in M} w_{ij}^2 q_j^{-1}$$

Lemma 8.7.4 implies that for all $j \in M$, $\int_{k \in N} p(k) c_k w_{kj} \, \mathrm{d}k \geq \beta \sqrt{n} \alpha c_i w_{ij}$. Using the definition of q_j, we get $q_j \geq \beta n^{3/2} \alpha c_i w_{ij}$, so $c_i w_{ij} q_j^{-1} \leq \frac{1}{\beta n^{3/2} \alpha}$. Therefore

$$\begin{aligned}
\tau_i &\leq \frac{2\sqrt{\gamma}}{w_{min} \lambda_i^2} \sum_{j \in M} w_{ij} (c_i w_{ij} q_j^{-1}) \\
&\leq \frac{2\sqrt{\gamma}}{\beta w_{min} n^{3/2} \alpha \lambda_i^2} \sum_{j \in M} w_{ij} \\
&\leq \frac{2\sqrt{\gamma} m w_{max}}{\beta w_{min} n^{3/2} \alpha \lambda_i^2} \\
&= (\tau_i + \gamma) \frac{2\sqrt{\gamma} m w_{max}}{\beta w_{min} n^{3/2} \alpha}
\end{aligned}$$

We can now solve for τ_i:

$$\begin{aligned}
\tau_i n^{3/2} \alpha \beta w_{min} &\leq 2(\tau_i + \gamma) \sqrt{\gamma} m w_{max} \\
\tau_i (n^{3/2} \alpha \beta w_{min} - 2\sqrt{\gamma} m w_{max}) &\leq 2\gamma^{3/2} m w_{max}
\end{aligned}$$

$$\tau_i \leq \frac{2\gamma^{3/2} m w_{max}}{n^{3/2}\alpha\beta w_{min} - 2\sqrt{\gamma} m w_{max}}$$

Note that we assumed $\beta w_{min} n^{3/2}\alpha - 2\sqrt{\gamma} m w_{max} > 0$ when dividing both sides by that quantity. This is true as $n \to \infty$, since $\lim_{n\to\infty} n^{3/2}\alpha = \infty$ by assumption. □

For brevity, let $\eta = \frac{n^{3/2}\alpha\beta w_{min} - 2\sqrt{\gamma} m w_{max}}{2 m w_{max}}$. Thus $\tau \leq \gamma^{3/2}/\eta$. Since we can always pick $\gamma \leq 1$, we have $\tau \leq \gamma/\eta$. Recall that β is a constant, so $\eta = \Theta(n^{3/2}\alpha)$.

Lemma 8.8.6. *For all $i \in \hat{N}$, $\sqrt{\frac{\eta}{\eta+1}} \cdot \frac{1}{\sqrt{\gamma}} \leq \lambda_i \leq \frac{1}{\sqrt{\gamma}}$.*

Proof. By Lemma 8.8.5,

$$\begin{aligned}
\lambda_i &= \frac{1}{\sqrt{\tau_i + \gamma}} \\
&\geq \frac{1}{\sqrt{\gamma/\eta + \gamma}} \\
&\geq \frac{\sqrt{\eta}}{\sqrt{\gamma + \gamma\eta}} \\
&= \sqrt{\frac{\eta}{\eta+1}} \cdot \frac{1}{\sqrt{\gamma}}
\end{aligned}$$

This satisfies the first inequality. The second follows immediately from the fact that $\tau_i \geq 0$. □

8.9 Bounding the net movement: $||\int_{i\in N} p(i)\delta_i(x)\,\mathrm{d}i||_2$

Finally, we need to show that $||\int_{i\in N} p(i)\delta_i(x)\,\mathrm{d}i||_2$ is small. Recall that by Lemma 8.8.2,

$$\delta_i = \frac{c_i w_{ij}(y_{ij} - x_j)}{c_i w_{ij} + \lambda_i q_j}$$

for all $i \in N$. We start by defining two approximations to δ_i:

$$\begin{aligned}
\delta_i' &= \frac{c_i w_{ij}(y_{ij} - x_j)}{\lambda_i q_j} \\
\delta_i'' &= c_i w_{ij}(y_{ij} - x_j) q_j^{-1}\sqrt{\gamma}
\end{aligned}$$

We first show that $\left|\left|\int_{i\in N} p(i)\delta_i''(x)\,\mathrm{d}i\right|\right|_2 = 0$ exactly. Next, we show that δ_i'' approximates δ_i' well, and that δ_i' approximates δ_i well, for each $i \in \hat{N}$. Finally, we show that the agents not in $\hat{N}$ do not matter too much, since their combined measure is small (Lemma 8.7.2).

Lemma 8.9.1.

$$\left|\left|\int_{i\in N} p(i)\delta_i''(x)\,\mathrm{d}i\right|\right|_2 = 0$$

Proof.

$$\begin{aligned}\left|\left|\int_{i\in N} p(i)\delta_i''(x)\,\mathrm{d}i\right|\right|_2 &= \left|\left|\int_{i\in N} p(i)c_iw_{ij}(y_{ij}-x_j)q_j^{-1}\sqrt{\gamma}\,\mathrm{d}i\right|\right|_2\\ &= \left|\left|\sqrt{\gamma}q_j^{-1}\int_{i\in N} p(i)c_iw_{ij}y_{ij}\,\mathrm{d}i - \sqrt{\gamma}q_j^{-1}\int_{i\in N} p(i)c_iw_{ij}x_j\,\mathrm{d}i\right|\right|_2\\ &= \left|\left|\sqrt{\gamma}q_j^{-1}\int_{i\in N} p(i)c_iw_{ij}y_{ij}\,\mathrm{d}i - \sqrt{\gamma}q_j^{-1}\Big(\int_{i\in N} p(i)c_iw_{ij}\,\mathrm{d}i\Big)x_j\right|\right|_2\end{aligned}$$

Since $x_j = \left(\int_{i\in N} p(i)c_iw_{ij}\,\mathrm{d}i\right)^{-1}\int_{i\in N} p(i)c_iw_{ij}y_{ij}\,\mathrm{d}i$, we have

$$\begin{aligned}\left|\left|\int_{i\in N} p(i)\delta_i''(x)\,\mathrm{d}i\right|\right|_2 &= \left|\left|\sqrt{\gamma}q_j^{-1}\int_{i\in N} p(i)c_iw_{ij}y_{ij}\,\mathrm{d}i - \sqrt{\gamma}q_j^{-1}\int_{i\in N} p(i)c_iw_{ij}y_{ij}\,\mathrm{d}i\right|\right|_2\\ &= 0\end{aligned}$$

as required. □

Lemma 8.9.2. *For each $i \in \hat{N}$,*

$$||\delta_i' - \delta_i||_2 \le \frac{2\gamma d_{max}}{\beta^2 n^3 \alpha^2}$$

Proof. The expression $\delta_i' - \delta_i$ reduces to:

$$\begin{aligned}\frac{c_iw_{ij}(y_{ij}-x)}{\lambda_i q_j} - \frac{c_iw_{ij}(y_{ij}-x)}{c_iw_{ij}+\lambda_i q_j} &= c_iw_{ij}(y_{ij}-x_j)\frac{c_iw_{ij}+\lambda_i q_j - \lambda_i q_j}{\lambda_i q_j(c_iw_{ij}+\lambda_i q_j)}\\ &= \frac{c_i^2 w_{ij}^2(y_{ij}-x_j)}{\lambda_i q_j(c_iw_{ij}+\lambda_i q_j)}\end{aligned}$$

Then using Lemmas 8.7.4 and 8.8.6, we have

$$\begin{aligned}||\delta_i' - \delta_i||_2^2 &= \frac{c_i^4 w_{ij}^4 (y_{ij}-x_j)^2}{\lambda_i^2 q_j^2 (c_iw_{ij}+\lambda_i q_j)^2}\\ &\le \frac{c_i^4 w_{ij}^4 (y_{ij}-x_j)^2}{\lambda_i^4 q_j^4}\\ &\le \frac{d_{max}^2}{\lambda_i^4}\cdot\left(\frac{c_iw_{ij}}{q_j}\right)^4\\ &\le \frac{d_{max}^2}{\lambda_i^4\beta^4 n^6\alpha^4}\\ &\le \frac{\gamma^2}{\beta^4 n^6 \alpha^4}\left(\frac{\eta+1}{\eta}\right)^2 d_{max}^2\end{aligned}$$

As long as $\eta \geq 1$ (which is of course true as n approaches infinity), we have

$$||\delta_i' - \delta_i||_2^2 \leq 4\frac{\gamma^2}{\beta^4 n^6 \alpha^4} d_{max}^2$$

Altogether, this implies that

$$||\delta_i' - \delta_i||_2 \leq \frac{2\gamma d_{max}}{\beta^2 n^3 \alpha^2}$$

as required. □

Lemma 8.9.3. *For all* $i \in \hat{N}$,

$$||\delta_i' - \delta_i''||_2 \leq \frac{\sqrt{\gamma} d_{max}}{\beta \eta n^{3/2} \alpha}$$

Proof. We have

$$\begin{aligned}
\frac{1}{\lambda_i} - \sqrt{\gamma} &\leq \sqrt{\gamma}\Big(\sqrt{\frac{\eta+1}{\eta}} - 1\Big) \\
&= \sqrt{\gamma}\Big(\sqrt{\frac{\eta+1}{\eta}} - \sqrt{\frac{\eta}{\eta}}\Big) \\
&= \sqrt{\gamma}\Big(\sqrt{\frac{\eta+1}{\eta}} - \sqrt{\frac{\eta}{\eta}}\Big) \\
&= \sqrt{\gamma}\frac{\sqrt{\eta+1} - \sqrt{\eta}}{\sqrt{\eta}} \\
&= \sqrt{\gamma}\frac{\eta+1-\eta}{\sqrt{\eta}(\sqrt{\eta+1} + \sqrt{\eta})} \\
&\leq \frac{\sqrt{\gamma}}{\eta}
\end{aligned}$$

Thus $||\delta_i' - \delta_i''||_2$ is bounded by

$$\begin{aligned}
||\delta_i' - \delta_i''||_2 &\leq \Big|\Big|\frac{c_i w_{ij}(y_{ij} - x)}{\lambda_i q_j} - c_i \sqrt{\gamma} q_j^{-1} w_{ij}(y_{ij} - x)\Big|\Big|_2 \\
&\leq \Big|\Big|\frac{\sqrt{\gamma}}{\eta} c_i q_j^{-1} w_{ij}(y_{ij} - x)\Big|\Big|_2 \\
&\leq \frac{\sqrt{\gamma}}{\eta}||c_i q_j^{-1} w_{ij}(y_{ij} - x)||_2 \\
&= \frac{\sqrt{\gamma}}{\eta}\sqrt{(y_{ij} - x_j)^2 (c_i w_{ij} q_j^{-1})^2} \\
&= \frac{\sqrt{\gamma}}{\eta}||y_i - x||_2 c_i w_{ij} q_j^{-1} \\
&\leq \frac{\sqrt{\gamma} d_{max}}{\beta \eta n^{3/2} \alpha}
\end{aligned}$$

□

Lemma 8.9.4.

$$\left|\left|\int_{i\notin\hat{N}} p(i)(\delta_i - \delta_i'')\,\mathrm{d}i\right|\right|_2 = O(\alpha^{m/2}n^{m/4})$$

Proof. For δ_i, we trivially have $||\delta_i||_2 \le d_{max}$. For δ_i'', by Lemma 8.7.4 (and the definition of q_j) we have $c_i w_{ij} q_j^{-1} \le \frac{1}{\beta n^{3/2}\alpha}$. Therefore

$$||\delta_i''||_2 = \sqrt{\gamma} c_i w_{ij} q_j^{-1} ||y_i - x||_2 \le \frac{\sqrt{\gamma} d_{max}}{\beta n^{3/2}\alpha}$$

Since $\lim_{n\to\infty} n^{3/2}\alpha = \infty$ by assumption, we have $||\delta_i''||_2 = O(1)$ (this is a loose bound of course, but sufficient for our purposes). Therefore

$$\begin{aligned}\left|\left|\int_{i\notin\hat{N}} p(i)(\delta_i - \delta_i'')\,\mathrm{d}i\right|\right|_2 &\le \int_{i\notin\hat{N}} p(i)||\delta_i||_2\,\mathrm{d}i + \int_{i\notin\hat{N}} p(i)||\delta_i''||_2\,\mathrm{d}i \\ &\le (d_{max} + O(1)) \int_{i\notin\hat{N}} p(i)\,\mathrm{d}i\end{aligned}$$

By Lemma 8.7.2, $\int_{i\notin\hat{N}} p(i)\,\mathrm{d}i = O(\alpha^{m/2}n^{m/4})$. Since d_{max} is also a constant, we have $\left|\left|\int_{i\notin\hat{N}} p(i)(\delta_i - \delta_i'')\,\mathrm{d}i\right|\right|_2 = O(\alpha^{m/2}n^{m/4})$. □

Lemma 8.9.5. *We have* $\left|\left|n\int_{i\in N} p(i)\delta_i(x)\,\mathrm{d}i\right|\right|_2 \le \frac{\gamma+\sqrt{\gamma}}{\Omega(n^2\alpha^2)} + O(\alpha^{m/2}n^{m/4+1})$.

Proof. We have

$$\begin{aligned}\left|\left|\int_{i\in N} p(i)\delta_i\,\mathrm{d}i\right|\right|_2 &= \left|\left|\int_{i\in N} p(i)(\delta_i'' - \delta_i'' + \delta_i)\,\mathrm{d}i\right|\right|_2 \\ &= \left|\left|\int_{i\in N} p(i)\delta_i''\,\mathrm{d}i\right|\right|_2 + \left|\left|\int_{i\in N} p(i)(\delta_i - \delta_i'')\,\mathrm{d}i\right|\right|_2\end{aligned}$$

By Lemma 8.9.1, $\left|\left|\int_{i\in N} p(i)\delta_i''\,\mathrm{d}i\right|\right|_2 = 0$. Thus

$$\begin{aligned}\left|\left|\int_{i\in N} p(i)\delta_i\,\mathrm{d}i\right|\right|_2 &= \left|\left|\int_{i\in N} p(i)(\delta_i - \delta_i'')\,\mathrm{d}i\right|\right|_2 \\ &= \left|\left|\int_{i\in\hat{N}} p(i)(\delta_i - \delta_i'')\,\mathrm{d}i + \int_{i\notin\hat{N}} p(i)(\delta_i - \delta_i'')\,\mathrm{d}i\right|\right|_2 \\ &\le \left|\left|\int_{i\in\hat{N}} p(i)(\delta_i - \delta_i'')\,\mathrm{d}i\right|\right|_2 + \left|\left|\int_{i\notin\hat{N}} p(i)(\delta_i - \delta_i'')\,\mathrm{d}i\right|\right|_2\end{aligned}$$

where the inequality follows from the triangle inequality of norms. By Lemma 8.9.4, we have $\left|\left|\int_{i\notin\hat{N}} p(i)(\delta_i - \delta_i'')\,\mathrm{d}i\right|\right|_2 = O(\alpha^{m/2}n^{m/4})$, so $\left|\left|\int_{i\in N} p(i)\delta_i\,\mathrm{d}i\right|\right|_2 \le \left|\left|\int_{i\in\hat{N}} p(i)(\delta_i - \delta_i'')\,\mathrm{d}i\right|\right|_2 + O(\alpha^{m/2}n^{m/4})$. For the next sequence of equations, we will use the triangle inequality, Lemmas 8.9.2 and 8.9.3, and $\int_{i\in\hat{N}} p(i)\,\mathrm{d}i \le \int_{i\in N} p(i)\,\mathrm{d}i = 1$.

$$\begin{aligned}\left|\left|\int_{i\in\hat{N}} p(i)(\delta_i - \delta_i'')\,\mathrm{d}i\right|\right|_2 &= \left|\left|\int_{i\in\hat{N}} p(i)(\delta_i'' - \delta_i' + \delta_i' - \delta_i)\,\mathrm{d}i\right|\right|_2 \\ &\le \left|\left|\int_{i\in\hat{N}} p(i)(\delta_i'' - \delta_i')\,\mathrm{d}i\right|\right|_2 + \left|\left|\int_{i\in\hat{N}} p(i)(\delta_i' - \delta_i)\,\mathrm{d}i\right|\right|_2 \\ &\le \int_{i\in\hat{N}} p(i)||\delta_i'' - \delta_i'||_2\,\mathrm{d}i + \int_{i\in\hat{N}} p(i)||\delta_i' - \delta_i||_2\,\mathrm{d}i \\ &\le \frac{2\gamma d_{max}}{\beta^2 n^3\alpha^2} \int_{i\in\hat{N}} p(i)\,\mathrm{d}i + \frac{\sqrt{\gamma}}{\beta\eta n^{3/2}\alpha} d_{max} \int_{i\in\hat{N}} p(i)\,\mathrm{d}i\end{aligned}$$

$$= \frac{2\gamma d_{max}}{\beta^2 n^3 \alpha^2} + \frac{\sqrt{\gamma}}{\beta \eta n^{3/2} \alpha} d_{max}$$
$$= \frac{\gamma}{\Omega(n^3\alpha^2)} + \frac{\sqrt{\gamma}}{\Omega(n^3\alpha^2)}$$

Therefore
$$\left|\left|\int_{i\in N} p(i)\delta_i \,\mathrm{d}i\right|\right|_2 \leq \frac{\gamma}{\Omega(n^3\alpha^2)} + \frac{\sqrt{\gamma}}{\Omega(n^3\alpha^2)} + O(\alpha^{m/2} n^{m/4})$$

To obtain the required bound for $\left|\left|n\int_{i\in N} p(i)\delta_i(x)\,\mathrm{d}i\right|\right|_2$, we simply multiply the above expression by n.

□

Lemmas 8.7.1, 8.8.6, and 8.9.5 together imply Theorem 8.4.1.

8.10 Non-triviality of our solution concept

As mentioned in Section 8.1.1, our definition of an equal-power equal-λ equilibrium may appear circular, since we are maximizing welfare and evaluating λ_i with respect to a utility scale which we get to choose. In Section 8.1.1, we argued that intuitively, this is analogous to the definition of equality by allocating equal amounts of an artificial currency. In this section, we argue that our solution concept is mathematically nontrivial, by (informally) showing that a particular "obvious" choice for $\mathbf{c}$ is not sufficient. Specifically, we show that a uniform $\mathbf{c}$ (i.e., $c_i = c$ for all i) cannot lead to an equal-power equal-λ equilibrium in general. Making this argument formal would require substantial algebra similar to that in Section 8.8.2, which we feel would obscure the primary intuition.

Recall Lemma 8.8.2, where $q_j = n\int_{k\in N} p(k)c_k w_{kj}\,\mathrm{d}k$:

Lemma 8.8.2. *For every agent i and issue j,* $\delta_{ij} = \frac{c_i w_{ij}(y_{ij} - x_j)}{c_i w_{ij} + \lambda_i q_j}$.

Suppose there exists some $c > 0$ such that $c_i = c$ for all $i \in N$. Then δ_{ij} reduces to

$$\delta_{ij} = \frac{c w_{ij}(y_{ij} - x_j)}{c w_{ij} + \lambda_i c n \int_{k\in N} p(k) w_{kj}\,\mathrm{d}k} = \frac{w_{ij}(y_{ij} - x_j)}{w_{ij} + \lambda_i n \int_{k\in N} p(k) w_{kj}\,\mathrm{d}k}$$

As n goes to infinity, the $\lambda_i n\int_{k\in N} p(k)w_{kj}\,\mathrm{d}k$ term dominates the $c_i w_{ij}$ term. Thus we can approximate δ_{ij} by
$$\delta_{ij} \approx \frac{w_{ij}(y_{ij} - x_j)}{\lambda_i n\int_{k\in N} p(k) w_{kj}\,\mathrm{d}k} = \frac{c w_{ij}(y_{ij} - x_j)}{\lambda_i q_j}$$

Recall that for any agent i whose power constraint is not tight, we have $\lambda_i = 0$; this is a standard property of Lagrange multipliers. Clearly for those agents, we cannot have $(1-\varepsilon')\lambda \leq \lambda_i \leq \lambda$ for any $\lambda > 0$. An equal-power equal-λ ε'-equilibrium requires that the above hold for at least a $1-\varepsilon'$ fraction of the agents, so at least a $1-\varepsilon'$ fraction of the agents must have a tight power constraint. Thus in order to achieve an exact equal-power equal-λ equilibrium as $n \to \infty$, almost all of the agents must have a tight power constraint as $n \to \infty$.

Let i be an arbitrary agent with a tight power constraint. By Lemma 8.8.1, this implies that $\sum_{j\in M} \delta_{ij}^2 q_j = \gamma$. Thus we have

$$\sum_{j\in M} \left(\frac{cw_{ij}(y_{ij}-x_j)}{\lambda_i q_j}\right)^2 q_j \approx \gamma$$
$$\sum_{j\in M} c^2 w_{ij}^2 (y_{ij}-x_j)^2 q_j^{-1} \approx \lambda_i^2$$
$$c^2 \sum_{j\in M} w_{ij}^2 (y_{ij}-x_j)^2 q_j^{-1} \approx \lambda_i^2$$

We need every pair of agents $i, i' \in \tilde{N}$ to satisfy $\lambda_i/\lambda_k \approx 1$. That means we need

$$\frac{\sum_{j\in M} w_{ij}^2 (y_{ij}-x_j)^2 (\int_{k\in N} p(k) c_k w_{kj}\,\mathrm{d}k)^{-1}}{\sum_{j\in M} w_{i'j}^2 (y_{i'j}-x_j)^2 (\int_{k\in N} p(k) c_k w_{kj}\,\mathrm{d}k)^{-1}} \approx 1$$

Note that $(\int_{k\in N} p(k) c_k w_{kj}\,\mathrm{d}k)^{-1}$ is just some constant. Consider a pair of agents i and i' such that $|y_{ij}-x_j| > |y_{i'j}-x_j|$ and $w_{ij} > w_{i'j}$ for each $j \in M$: then this ratio can never approach 1, even as $n \to \infty$. For an appropriate choice of distribution p, this means that there will always be constant measure set of agents whose values of λ_i are not close to the Lagrange multipliers of other agents.

This means that in general, choosing the same scaling factor for each agent will not be sufficient for an equal-power equal-λ equilibrium. Furthermore, the above reasoning intuitively suggests that we really do need $c_i \approx 1/\sqrt{\sum_{j\in M} w_{ij}^2 (y_{ij}-x_j)^2 q_j^{-1}}$ for (almost) every $i \in N$.

8.11 Conclusion

In this chapter, we proposed and analyzed the concept of *equal power* for multidimensional continuous public decision-making. Drawing fundamental literature in political philosophy and economics, we argued that that equality of power is a natural analog of equality of resources and envy-freeness for public decision-making. Our main result is that for any $\varepsilon > 0$ and a large enough number of agents, an equal-power equal-λ ε-equilibrium is guaranteed to exist; in other words, we achieve an exact equal-power equal-λ equilibrium in the limit. In our opinion most interesting part of our proof is the novel fixed point argument presented in Section 8.5.

There are many possible directions for future work. The first is the possibility of an iterative algorithm for converging to an equal-power equal-λ equilibrium. As discussed in Section 8.1.2, there is a good reason to be optimistic about the existence of such an algorithm, especially in conjunction with the already extensive quadratic voting literature.

It could also be interesting to extend our results to other utility functions beyond quadratic utilities. A first step could be a "general quadratic utility" of the form $u_i(x) = -(y_i - x)^T W_i (y_i - x)$ for some positive definite matrix W_i. When W_i is a diagonal matrix, this reduces to the form of utility function we studied in this chapter. For a general quadratic utility, we conjecture that the equal power constraint would be reduce to a constraint of the form $\delta_i^T Q \delta_i \le \gamma$. This would correspond to not just a rescaling of the issues, but also a rotation. Our proof does not immediately carry over to general quadratic utility functions, and we suspect that additional mathematical insights are needed.

Possibly the most exciting direction – but also the most ambitious – is extending our results beyond economics that are purely public goods. In general, economies will involve a much richer mix of public goods at different levels of social organization, which are thus partially private when viewed from another resolution (e.g. goods that accrue at the national or city level, but do not spillover beyond these). Efficient, equal budget mechanisms for such societies might offer powerful insights about economic structures that could outperform existing mixtures of capitalism and states.

Funding

My thesis was supported in part by NSF Graduate Research Fellowship under grant DGE-1656518.

Bibliography

[1] B. Adsul et al. "Nash Equilibria in Fisher Market". In: *Proceedings of the Third International Conference on Algorithmic Game Theory*. SAGT'10. Athens, Greece: Springer-Verlag, 2010, pp. 30–41.

[2] "Difference without Domination: Pursuing Justice in Diverse Democracies". In: ed. by D. Allen and R. Somanathan. University of Chicago Press, 2020. Chap. A New Theory of Justice: Difference without Domination.

[3] G. Amanatidis et al. "Maximum Nash welfare and other stories about EFX". In: *Theoretical Computer Science* 863 (2021), pp. 69–85.

[4] E. S. Anderson. "What Is the Point of Equality?" In: *Ethics* 109 (1999), pp. 287–337.

[5] S. Ansolabehere and J. Snyder. *The End of Inequality: One Person, One Vote and the Transformation of American Politics*. New York: W. W. Norton, 2008.

[6] K. J. Arrow and G. Debreu. "Existence of an Equilibrium for a Competitive Economy". In: *Econometrica* 22.3 (1954), pp. 265–290.

[7] E. R. Arunachaleswaran et al. "Fair and Efficient Cake Division with Connected Pieces". In: *International Conference on Web and Internet Economics*. Springer. 2019, pp. 57–70.

[8] A. B. Atkinson. "On the measurement of inequality". In: *Journal of Economic Theory* 2.3 (1970), pp. 244–263.

[9] N. Avigdor-Elgrabli, Y. Rabani, and G. Yadgar. "Convergence of Tâtonnement in Fisher Markets". In: *arXiv:1401.6637 [cs]* (2014). arXiv: 1401.6637.

[10] H. Aziz et al. "Fair Assignment of Indivisible Objects Under Ordinal Preferences". In: *Artificial Intelligence* 227.C (2015), pp. 71–92.

[11] S. Barman and S. K. Krishna Murthy. "Approximation Algorithms for Maximin Fair Division". In: *Proceedings of the 2017 ACM Conference on Economics and Computation*. EC '17. Cambridge, Massachusetts, USA: ACM, 2017, pp. 647–664.

[12] N. Becker. "Water Pricing in Israel: Various Waters, Various Neighbors". In: *Water Pricing Experiences and Innovations*. Ed. by A. Dinar, V. Pochat, and J. Albiac-Murillo. Cham: Springer International Publishing, 2015, pp. 181–199.

[13] D. Benjamin et al. "Aggregating Local Preferences to Guide Marginal Policy Adjustments." In: *The American economic review* 103 (3 2013), pp. 605–610.

[14] D. Benjamin et al. “The Relationship between the Normalized Gradient Addition Mechanism and Quadratic Voting”. In: *Public Choice* 172.1-2 (2017), p. 233.

[15] T. C. Bergstrom. “When Does Majority Rule Supply Public Goods Efficiently?” en. In: *Measurement in Public Choice*. Ed. by S. Strøm. London: Palgrave Macmillan UK, 1981, pp. 75–85.

[16] M. Berinde. “Approximate Fixed Point Theorems”. In: *Studia. Universitatis Babeş-Bolyai. Mathematica* 51.1 (2006), pp. 11–25.

[17] D. Bertsimas, V. F. Farias, and N. Trichakis. “On the Efficiency-fairness Trade-off”. In: *Management Science* 58.12 (2012), pp. 2234–2250.

[18] A. Bogomolnaia and H. Moulin. “Random Matching Under Dichotomous Preferences”. In: *Econometrica* 72.1 (2004), pp. 257–279.

[19] J. J. Boland and D. Whittington. “Water Tariff Design in Developing Countries: Disadvantages of Increasing Block Tariffs (ibts) and Advantages of Uniform Price with Rebate (upr) Designs”. In: *World Bank Water and Sanitation Program, Washington, DC* 37 (2000).

[20] T. Bonald and L. Massoulié. “Impact of Fairness on Internet Performance”. In: *ACM SIGMETRICS Performance Evaluation Review*. Vol. 29. 1. ACM. 2001, pp. 82–91.

[21] S. Bouveret and J. Lang. “Efficiency and envy-freeness in fair division of indivisible goods: logical representation and complexity”. In: *Journal of Artificial Intelligence Research* 32 (2008), pp. 525–564.

[22] H. R. Bowen. “The Interpretation of Voting in the Allocation of Economic Resources”. In: *The Quarterly Journal of Economics* 58.1 (1943), p. 27.

[23] S. P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge, UK ; New York: Cambridge University Press, 2004.

[24] W. C. Brainard and H. E. Scarf. “How to Compute Equilibrium Prices in 1891”. In: *American Journal of Economics and Sociology* 64.1 (2005), pp. 57–83.

[25] S. J. Brams, D. M. Kilgour, and C. Klamler. “Maximin Envy-Free Division of Indivisible Items”. In: *Group Decision and Negotiation* 26.1 (2017), pp. 115–131.

[26] S. J. Brams, D. M. Kilgour, and C. Klamler. “The undercut procedure: an algorithm for the envy-free division of indivisible items”. In: *Social Choice and Welfare* 39.2-3 (2012), pp. 615–631.

[27] S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.

[28] F. Brandt et al. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

[29] S. Brânzei and N. Nisan. “Communication Complexity of Cake Cutting”. In: *ArXiv e-prints* (2017). arXiv: 1709.09876.

[30] S. Branzei, V. Gkatzelis, and R. Mehta. "Nash Social Welfare Approximation for Strategic Agents". In: *Proceedings of the 2017 ACM Conference on Economics and Computation*. EC '17. Cambridge, Massachusetts, USA: ACM, 2017, pp. 611–628.

[31] S. Brânzei et al. "The Fisher Market Game: Equilibrium and Welfare". In: *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*. AAAI'14. Québec City, Québec, Canada: AAAI Press, 2014, pp. 587–593.

[32] R. Briffault. "Who Rules at Home?: One Person/One Vote and Local Governments". In: *The University of Chicago Law Review* 60.2 (1993), pp. 339–424.

[33] S. Bubeck. "Convex Optimization: Algorithms and Complexity". In: *Foundations and Trends® in Machine Learning* 8.3-4 (2015), pp. 231–357.

[34] J. M. Buchanan, R. D. Tollison, and G. Tullock. *Toward a theory of the rent-seeking society*. 4. Texas A & M Univ Pr, 1980.

[35] E. Budish. "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes". In: *Journal of Political Economy* 119.6 (2011), pp. 1061–1103.

[36] E. Budish et al. "Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation". In: *Operations Research* 65.2 (2016), pp. 314–336.

[37] E. Budish et al. "Designing Random Allocation Mechanisms: Theory and Applications". In: *American Economic Review* 103.2 (2013), pp. 585–623.

[38] C. Burger and A. Jansen. "Increasing Block Tariff Structures As a Water Subsidy Mechanism in South Africa: An Exploratory Analysis". In: *Development Southern Africa* 31.4 (2014), pp. 553–562.

[39] I. Caragiannis, N. Gravin, and X. Huang. "Envy-freeness up to any item with high Nash welfare: The virtue of donating items". In: *Proceedings of the 2019 ACM Conference on Economics and Computation*. 2019, pp. 527–545.

[40] I. Caragiannis et al. "The Unreasonable Fairness of Maximum Nash Welfare". In: *Proceedings of the 2016 ACM Conference on Economics and Computation*. EC '16. Maastricht, The Netherlands, 2016, pp. 305–322.

[41] A. Casella. "Storable votes". In: *Games and Economic Behavior* 51.2 (2005). Special Issue in Honor of Richard D. McKelvey, pp. 391–419.

[42] V. Cerf and R. Kahn. "A Protocol for Packet Network Intercommunication". In: *IEEE Transactions on Communications* 22.5 (1974), pp. 637–648.

[43] D. Chakrabarty, N. Devanur, and V. V. Vazirani. "New Results on Rationality and Strongly Polynomial Time Solvability in Eisenberg-Gale Markets". In: *Internet and Network Economics: Second International Workshop, WINE 2006, Patras, Greece, December 15-17, 2006. Proceedings*. Ed. by P. Spirakis, M. Mavronicolas, and S. Kontogiannis. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 239–250.

[44] P. Chander and H. Tulkens. “A Core-Theoretic Solution for the Design of Cooperative Agreements on Transfrontier Pollution”. en. In: *Public goods, environmental externalities and fiscal competition*. Ed. by P. Chander et al. Boston, MA: Springer US, 2006, pp. 176–193.

[45] B. R. Chaudhury, J. Garg, and K. Mehlhorn. “EFX exists for three agents”. In: *Proceedings of the 21st ACM Conference on Economics and Computation*. 2020, pp. 1–19.

[46] B. R. Chaudhury et al. “A little charity guarantees almost envy-freeness”. In: *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2020, pp. 2658–2672.

[47] Y. Chen and D. M. Pennock. “Designing markets for prediction”. In: *AI Magazine* 31.4 (2010), pp. 42–52.

[48] Y. K. Cheung, R. Cole, and N. Devanur. “Tatonnement beyond gross substitutes?: gradient descent to the rescue”. In: *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. ACM, 2013, pp. 191–200.

[49] H. Chung and J. Duggan. “Directional Equilibria”. In: *Journal of Theoretical Politics* 30 (2018), pp. 272–305.

[50] G. de Clippel, H. Moulin, and N. Tideman. “Impartial division of a dollar”. In: *Journal of Economic Theory* 139.1 (2008), pp. 176–191.

[51] R. Cole, V. Gkatzelis, and G. Goel. “Mechanism Design for Fair Division: Allocating Divisible Items without Payments”. In: *EC 2013*. 2013.

[52] V. Conitzer, R. Freeman, and N. Shah. “Fair Public Decision Making”. In: *Proceedings of the 2017 ACM Conference on Economics and Computation*. EC ’17. Cambridge, Massachusetts, USA: ACM, 2017, pp. 629–646.

[53] F. Cowell. *Measuring Inequality*. Oxford University Press, 2011.

[54] M. W. Cripps and J. M. Swinkels. “Efficiency of Large Double Auctions”. In: *Econometrica* 74.1 (2006), pp. 47–92.

[55] R. A. Dahl. *On Political Equality*. Yale University Press, 2006.

[56] H. Dalton. “The Measurement of the Inequality of Incomes”. In: *The Economic Journal* 30.119 (1920), pp. 348–361.

[57] L. Danziger. “A graphic representation of the Nash and Lindahl equilibria in an economy with a public good”. In: *Journal of Public Economics* 6.3 (1976), pp. 295–307.

[58] P. Dasgupta, P. Hammond, and E. Maskin. “The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility”. In: *The Review of Economic Studies* 46.2 (1979), pp. 185–216.

[59] G. Debreu. *Theory of Value*. New York: John Wiley & Sons, 1959.

[60] G. Debreu. “Topological Methods in Cardinal Utility Theory”. In: *Mathematical methods in the social sciences*. Ed. by S. K. Kenneth J Arrow and P. Suppes. Stanford University Press, 1960.

[61] N. R. Devanur et al. "Market Equilibrium Via a Primal–dual Algorithm for a Convex Program". In: *Journal of the ACM (JACM)* 55.5 (2008), p. 22.

[62] J. P. Dickerson et al. "The Computational Rise and Fall of Fairness". In: *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*. AAAI'14. 2014, pp. 1405–1411.

[63] H. Dinh and A. Russell. "Quantum and Randomized Lower Bounds for Local Search on Vertex-Transitive Graphs". In: *Quantum Information and Computation* 10.7–8 (2010), pp. 636–652.

[64] S. Dobzinski, H. Fu, and R. Kleinberg. "On the Complexity of Computing an Equilibrium in Combinatorial Auctions". In: *Proceedings of the Twenty-sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '15. San Diego, California, 2015, pp. 110–122.

[65] S. Dobzinski, N. Nisan, and M. Schapira. "Approximation Algorithms for Combinatorial Auctions with Complement-Free Bidders". In: *Math. Oper. Res.* 35.1 (2010), pp. 1–13.

[66] S. Dobzinski and J. Vondrák. "Communication Complexity of Combinatorial Auctions with Submodular Valuations". In: *Proceedings of the Twenty-fourth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '13. New Orleans, Louisiana, 2013, pp. 1205–1215.

[67] P. Dolan. "The Measurement of Individual Utility and Social Welfare". In: *Journal of health economics* 17.1 (1998), pp. 39–52.

[68] P. Dolan and R. Cookson. "A Qualitative Study of the Extent to Which Health Gain Matters When Choosing between Groups of Patients". In: *Health Policy* 51.1 (2000), pp. 19–30.

[69] J. Douceur. "The Sybil Attack". In: *Proceedings of the International Workshop on Peer-to-Peer Systems*. 2002.

[70] P. Dubey and M. Shubik. "A theory of money and financial institutions. 28. The noncooperative equilibria of a closed trading economy with market supply and bidding strategies". In: *Journal of Economic Theory* 17.1 (1978), pp. 1–20.

[71] R. Dworkin. "What Is Equality? Part Ii: Equality of Resources". In: *Philosophy & Public Affairs* 10.4 (1981), pp. 283–345.

[72] E. Eisenberg. "Aggregation of Utility Functions". In: *Management Science* 7.4 (1961), pp. 337–350.

[73] E. Eisenberg and D. Gale. "Consensus of Subjective Probabilities: The Pari-mutuel Method". In: *The Annals of Mathematical Statistics* 30.1 (1959), pp. 165–168.

[74] M. Elliott and B. Golub. "A Network Approach to Public Goods". In: *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*. EC '13. New York, NY, USA: ACM, 2013, pp. 377–378.

[75] B. Fain, A. Goel, and K. Munagala. "The Core of the Participatory Budgeting Problem". In: *Web and Internet Economics: 12th International Conference, WINE 2016, Montreal, Canada, December 11-14, 2016, Proceedings*. Ed. by Y. Cai and A. Vetta. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016, pp. 384–399.

[76] B. Fain, K. Munagala, and N. Shah. "Fair Allocation of Indivisible Public Goods". In: *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*. Ed. by É. Tardos, E. Elkind, and R. Vohra. ACM, 2018, pp. 575–592.

[77] U. Feige et al. "Computing with noisy information". In: *SIAM Journal on Computing* 23.5 (1994), pp. 1001–1018.

[78] M. Feldman, K. Lai, and L. Zhang. "The Proportional-Share Allocation Market for Computational Resources". In: *IEEE Transactions on Parallel and Distributed Systems* 20.8 (2009), pp. 1075–1088.

[79] W. Feller. *An Introduction to Probability Theory and Its Applications, Vol. 1, 3rd Edition.* 3rd. Wiley, 1968.

[80] I. Fisher. *Mathematical Investigations in the Theory of Value and Prices.* Connecticut Academy of Arts and Sciences, 1892.

[81] S. Floyd and V. Jacobson. "Random early detection gateways for congestion avoidance". In: *IEEE/ACM Transactions on Networking* 1.4 (1993), pp. 397–413.

[82] D. K. Foley. "Lindahl's Solution and the Core of an Economy with Public Goods". In: *Econometrica* 38.1 (1970), p. 66.

[83] M. A. Gabis. "One Person, One Vote". In: *Northern Kentucky Law Review* 5 (1978), p. 241.

[84] Y. (Gal et al. "Which Is the Fairest (Rent Division) of Them All?" In: *Proceedings of the 2016 ACM Conference on Economics and Computation.* EC '16. Maastricht, The Netherlands, 2016, pp. 67–84.

[85] G. Gamow and M. Stern. "Puzzle-math, Edn". In: *Viking Adult, URL http://www. worldcat. org/isbn/0670583359* (1958).

[86] M. A. García-Rubio, A. Ruiz-Villaverde, and F. González-Gómez. "Urban Water Tariffs in Spain: What Needs to Be Done?" In: *Water* 7.4 (2015), pp. 1456–1479.

[87] D. Garg et al. "A Primal-Dual Algorithm for Computing Fisher Equilibrium in the Absence of Gross Substitutability Property". In: *Internet and Network Economics: First International Workshop, WINE 2005, Hong Kong, China, December 15-17, 2005. Proceedings.* Ed. by X. Deng and Y. Ye. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 24–33.

[88] N. Garg, A. Goel, and B. Plaut. "Markets for Public Decision-making". In: *Proceedings of the 14th Conference on Web and Internet Economics.* WINE '18. 2018.

[89] N. Garg et al. "Collaborative Optimization for Collective Decision-making in Continuous Spaces". In: *Proceedings of the 26th International Conference on World Wide Web.* WWW '17. 2017, pp. 617–626.

[90] N. Garg et al. "Iterative Local Voting for Collective Decision-making in Continuous Spaces". In: *Journal of Artificial Intelligence Research* To Appear (2018).

[91] H. Gersbach. "Why one person one vote?" en. In: *Social Choice and Welfare* 23.3 (2004), pp. 449–464.

[92] M. Ghodsi et al. “Fair allocation of indivisible goods: Improvements and generalizations”. In: *Proceedings of the 2018 ACM Conference on Economics and Computation*. 2018, pp. 539–556.

[93] G. Giraud. “Strategic market games: an introduction”. In: *Journal of Mathematical Economics* 39.5 (2003). Strategic Market Games, pp. 355–375.

[94] A. Goel, R. Hulett, and B. Plaut. “Markets beyond Nash Welfare for Leontief Utilities”. In: *Proceedings of the 15th Conference on Web and Internet Economics (WINE ’19)*. 2019.

[95] A. Goel and H. Nazerzadeh. “Price-based Protocols for Fair Resource Allocation: Convergence Time Analysis and Extension to Leontief Utilities”. In: *ACM Trans. Algorithms* 10.2 (2014), 5:1–5:14.

[96] A. Goel and B. Plaut. “Counteracting inequality in markets via convex pricing”. In: *International Conference on Web and Internet Economics*. Springer. 2020, pp. 89–101.

[97] A. Goel et al. “Knapsack Voting: Voting mechanisms for Participatory Budgeting”. In: *ACM Transactions on Economics and Computation (TEAC)* 7.2 (2019), pp. 1–27.

[98] J. Goldman and A. D. Procaccia. “Spliddit: Unleashing Fair Division Algorithms”. In: *SIGecom Exch.* 13.2 (2015), pp. 41–46.

[99] T. Groves and J. Ledyard. “Optimal Allocation of Public Goods: A Solution to the ”Free Rider” Problem”. en. In: *Econometrica* 45.4 (1977), p. 783.

[100] D. W. Hands. “The Structuralist View of Economic Theories: A Review Essay: The Case of General Equilibrium in Particular”. en. In: *Economics & Philosophy* 1.2 (1985), pp. 303–335.

[101] G. M. Hayden. “The False Promise of One Person, One Vote”. In: *Michigan Law Review* 102.2 (2003), pp. 213–267.

[102] A. Hylland and R. Zeckhauser. “A Mechanism for Selecting Public Goods When Preferences Must Be Elicited”. In: *Kennedy School of Government discussion paper D* 70 (1979).

[103] N. Immorlica, B. Plaut, and E. G. Weyl. “Equality of Power and Fair Public Decision-Making”. In: *International Conference on Web and Internet Economics*. Springer. 2020.

[104] K. Jain and V. V. Vazirani. “Eisenberg–gale Markets: Algorithms and Game-theoretic Properties”. In: *Games and Economic Behavior* 70.1 (2010), pp. 84–106.

[105] L. Jiang and J. Walrand. “Scheduling and Congestion Control for Wireless and Processing Networks”. en. In: *Synthesis Lectures on Communication Networks* 3.1 (2010), pp. 1–156.

[106] B. Kalyanasundaram and G. Schintger. “The Probabilistic Communication Complexity of Set Intersection”. In: *SIAM J. Discret. Math.* 5.4 (1992), pp. 545–557.

[107] M. Kaneko. “The ratio equilibrium and a voting game in a public goods economy”. In: *Journal of Economic Theory* 16.2 (1977), pp. 123–136.

[108] M. Kaneko and K. Nakamura. “The Nash Social Welfare Function.” In: *Econometrica* 47.2 (1979), pp. 423–435.

[109] L. Karabarbounis. “One Dollar, One Vote*”. en. In: *The Economic Journal* 121.553 (2011), pp. 621–651.

[110] F. P. Kelly, A. K. Maulloo, and D. K. Tan. "Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability". In: *Journal of the Operational Research society* 49.3 (1998), pp. 237–252.

[111] M. D. Kirszbraun. "Über Die Zusammenziehende Und Lipschitzsche Transformationen". In: *Fundamenta Mathematicae* 22.1 (1934), pp. 77–108.

[112] C. Klassert et al. "Increasing Block Tariffs in an Arid Developing Country: A Discrete/continuous Choice Model of Residential Water Demand in Jordan". In: *Water* 10.3 (2018), p. 248.

[113] B. Klaus and E. Miyagawa. "Strategy-proofness, solidarity, and consistency for multiple assignment problems". In: *International Journal of Game Theory* 30.3 (2002), pp. 421–435.

[114] E. Koutsoupias and C. Papadimitriou. "Worst-case Equilibria". In: *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science*. STACS'99. Trier, Germany: Springer-Verlag, 1999, pp. 404–413.

[115] S. P. Lalley and E. G. Weyl. "Nash Equilbria for Quadratic Voting". In: *arXiv:1409.0264 [cs, math]* (2014). arXiv: 1409.0264.

[116] S. P. Lalley and E. G. Weyl. "Quadratic Voting: How Mechanism Design Can Radicalize Democracy". In: *AEA Papers and Proceedings*. Vol. 108. 2018, pp. 33–37.

[117] B. Lehmann, D. Lehmann, and N. Nisan. "Combinatorial Auctions with Decreasing Marginal Utilities". In: *Proceedings of the 3rd ACM Conference on Electronic Commerce*. EC '01. Tampa, Florida, USA: ACM, 2001, pp. 18–28.

[118] L. Lindholm, M. Rosén, and M. Emmelin. "How Many Lives Is Equity Worth? a Proposal for Equity Adjusted Years of Life Saved." In: *Journal of Epidemiology & Community Health* 52.12 (1998), pp. 808–811.

[119] R. Lipton et al. "On Approximately Fair Allocations of Indivisible Goods". In: *Proceedings of the 5th ACM Conference on Electronic Commerce*. EC '04. New York, NY, USA, 2004, pp. 125–131.

[120] D. C. Llewellyn, C. Tovey, and M. Trick. "Local Optimization on Graphs". In: *Discrete Appl. Math.* 23.2 (1989), pp. 157–178.

[121] K.-G. Mäler. "Chapter 1 Welfare economics and the environment". In: *Handbook of Natural Resource and Energy Economics*. Vol. 1. Elsevier, 1985, pp. 3–60.

[122] E. Maskin. "Nash Equilibrium and Welfare Optimality". In: *The Review of Economic Studies* 66.1 (1999), pp. 23–38.

[123] E. Maskin and T. Sjöström. "Implementation Theory". In: *Handbook of Social Choice Theory Vol. I*. Ed. by K. Arrow, A. Sen, and K. Suzumura. Amsterdam: North Holland, 2002, pp. 237–288.

[124] A. Matros. “Chinese auctions”. In: *GAME THEORY AND MANAGEMENT. Collected abstracts of papers presented on the Fifth International Conference Game Theory and Management/Editors Leon A. Petrosyan and Nikolay A. Zenkevich.–SPb.: Graduate School of Management SPbU, 2011.–268 p. The collection contains abstracts of papers accepted for the Fifth International*. 2011, p. 153.

[125] M. Matsumoto and N. Tokushige. “The exact bound in the Erdős-Ko-Rado theorem for cross-intersecting families”. In: *Journal of Combinatorial Theory, Series A* 52.1 (1989), pp. 90–97.

[126] J. Mo and J. Walrand. “Fair End-to-end Window-based Congestion Control”. In: *IEEE/ACM Transactions on networking* 8.5 (2000), pp. 556–567.

[127] H. Monteiro and C. Roseta-Palma. “Pricing for Scarcity? an Efficiency Analysis of Increasing Block Tariffs”. In: *Water Resources Research* 47.6 (2011).

[128] H. Moulin. “Fair Division and Collective Welfare”. In: MIT press, 2003. Chap. 3.

[129] J. Nash. “The Bargaining Problem”. In: *Econometrica* 18.2 (1950), pp. 155–162.

[130] E. Neuman. “Inequalities Involving Multivariate Convex Functions. Ii”. In: *Proceedings of the American Mathematical Society* 109.4 (1990), pp. 965–974.

[131] I. Newman. “Private vs. Common Random Bits in Communication Complexity”. In: *Inf. Process. Lett.* 39.2 (1991), pp. 67–71.

[132] N. Nisan. “The Communication Complexity of Approximate Set Packing and Covering”. In: *Automata, Languages and Programming: 29th International Colloquium, ICALP 2002 Málaga, Spain, July 8–13, 2002 Proceedings*. Ed. by P. Widmayer et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 868–875.

[133] N. Nisan. “The communication complexity of threshold gates”. In: *Combinatorics, Paul Erdos is Eighty* 1 (1993), pp. 301–315.

[134] N. Nisan and I. Segal. “The communication requirements of efficient allocations and supporting prices”. In: *Journal of Economic Theory* 129.1 (2006), pp. 192–224.

[135] N. Nisan et al., eds. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[136] A. van den Nouweland. “Lindahl and Equilibrium”. In: *Individual and Collective Choice and Social Welfare*. Springer, 2015, pp. 335–362.

[137] V. Ortún. “Contradictions and Trade-offs between Efficiency and Equity”. In: *Institut Borja de Bioetica, Allocation of Resources in Health Care. Barcelona: Fundación MAPFRE Medicina* (1996), pp. 113–24.

[138] J. Padhye et al. “Modeling TCP Throughput: A Simple Model and Its Empirical Validation”. In: *SIGCOMM Comput. Commun. Rev.* 28.4 (1998), pp. 303–314.

[139] R. Pan, B. Prabhakar, and K. Psounis. “CHOKe - a stateless active queue management scheme for approximating fair bandwidth allocation”. In: *Proceedings IEEE INFOCOM 2000. Conference on Computer Communications. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies (Cat. No.00CH37064)*. Vol. 2. 2000, 942–951 vol.2.

[140] E. A. Pazner and D. Schmeidler. "Egalitarian Equivalent Allocations: A New Concept of Economic Equity". In: *The Quarterly Journal of Economics* 92.4 (1978), p. 671.

[141] C.-t. Perng. "On a Class of Theorems Equivalent to Farkas's Lemma". In: *Applied Mathematical Sciences* 11.44 (2017), pp. 2175–2184.

[142] A. C. Pigou. *Wealth and Welfare*. Macmillan and Company, limited, 1912.

[143] B. Plaut. "Optimal Nash Equilibria for Bandwidth Allocation". In: *International Conference on Web and Internet Economics*. Springer, Cham, 2020, pp. 77–88.

[144] B. Plaut and T. Roughgarden. "Almost envy-freeness with general valuations". In: *Proceedings of the 2018 ACM-SIAM Symposium on Discrete Algorithms*. SODA '18. 2018.

[145] B. Plaut and T. Roughgarden. "Communication complexity of discrete fair division". In: *SIAM Journal on Computing* 49.1 (2020), pp. 206–243.

[146] E. A. Posner and E. G. Weyl. "Quadratic Voting and the Public Good: Introduction". In: *Public Choice* 172.1-2 (2017), p. 1.

[147] D. Quarfoot et al. "Quadratic Voting in the Wild: Real People, Real Votes". In: *Public Choice* 172.1-2 (2017), pp. 283–303.

[148] S. Ramezani and U. Endriss. "Nash Social Welfare in Multiagent Resource Allocation". In: *Proceedings of the Twelfth Annual Conference on Agent-Mediated Electronic Commerce*. AMEC '10. 2010, pp. 117–131.

[149] J. Rawls. *A Theory of Justice*. 1st ed. Cambridge, Massachussets: Belknap Press of Harvard University Press, 1971.

[150] D. Ray and R. Vohra. "Coalitional Power and Public Goods". en. In: *Journal of Political Economy* 109.6 (2001), pp. 1355–1384.

[151] A. A. Razborov. "On the distributional complexity of disjointness". In: *International Colloquium on Automata, Languages, and Programming*. Springer. 1990, pp. 249–253.

[152] M. E. Renwick and R. D. Green. "Do Residential Water Demand Side Management Policies Measure Up? an Analysis of Eight California Water Agencies". In: *Journal of environmental economics and management* 40.1 (2000), pp. 37–55.

[153] J. Riley. "Justice Under Capitalism". In: *Nomos* 31 (1989), pp. 122–162.

[154] D. J. Roberts and A. Postlewaite. "The Incentives for Price-taking Behavior in Large Exchange Economies". In: *Econometrica: Journal of the Econometric Society* (1976), pp. 115–127.

[155] A. Rogers. *Colorado Tried a New Way to Vote: Make People Pay – Quadratically*. 2019.

[156] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. The MIT Press, 2005.

[157] H. L. Royden and P. Fitzpatrick. *Real Analysis*. Vol. 32. Macmillan New York, 1988.

[158] D. Satz. *Why Some Things Should Not Be for Sale: The Moral Limits of Markets*. English. Reprint edition. Oxford New York, NY: Oxford University Press, 2012.

[159] J. Schummer. “Strategy-proofness versus efficiency on restricted domains of exchange economies”. In: *Social Choice and Welfare* 14.1 (1996), pp. 47–56.

[160] A. Sen. “Social Choice Theory: A Re-examination”. In: *Econometrica: journal of the Econometric Society* 45.1 (1977), pp. 53–89.

[161] A. Sen. “Welfare Inequalities and Rawlsian Axiomatics”. In: *Theory and decision* 7.4 (1976), pp. 243–262.

[162] A. K. Sen. “Rational Fools: A Critique of the Behavioral Foundations of Economic Theory”. In: *Philosophy & Public Affairs* 6.4 (1977), pp. 317–344.

[163] A. K. Sen. “Democracy As a Universal Value”. In: *Journal of Democracy* 10.3 (1999), pp. 3–17.

[164] L. Shapley and M. Shubik. “Trade Using One Commodity as a Means of Payment”. In: *Journal of Political Economy* 85.5 (1977), pp. 937–68.

[165] R. J. van der Spek. *Commodity Prices in Babylon 385 - 61 Bc.* http://www.iisg.nl/hpw/babylon.php. Accessed December 14 2019. 2005.

[166] R. J. van der Spek and C. Mandemakers. “Sense and Nonsense in the Statistical Approach of Babylonian Prices”. In: *Bibliotheca Orientalis* 60.5 (2003), pp. 521–537.

[167] J. Stoer and C. Witzgall. *Convexity and Optimization in Finite Dimensions I.* Berlin, Heidelberg: Springer Berlin Heidelberg, 1970, pp. 7, 24.

[168] M. Valencia-Pabon and J.-C. Vera. “On the Diameter of Kneser Graphs”. In: *Discrete Math.* 305.1-3 (2005), pp. 383–385.

[169] H. Varian. “Equity, Envy, and Efficiency”. In: *Journal of Economic Theory* 9.1 (1974), pp. 63–91.

[170] V. V. Vazirani. “Combinatorial Algorithms for Market Equilibria”. In: *Algorithmic Game Theory.* Ed. by N. Nisan et al. Cambridge university press, 2007. Chap. 5, pp. 103–134.

[171] A. Wagstaff. “Qalys and the Equity-efficiency Trade-off”. In: *Journal of health economics* 10.1 (1991), pp. 21–41.

[172] M. Walker. “A Simple Incentive Compatible Scheme for Attaining Lindahl Allocations”. en. In: *Econometrica* 49.1 (1981), p. 65.

[173] L. Walras. *Elements of Pure Economics: Or, the Theory of Social Wealth. Translated by William Jaffé.* Published for the American Economic Association and the Royal Economic Society, 1954.

[174] E. G. Weyl. “The Robustness of Quadratic Voting”. In: *Public Choice* 172.1-2 (2017), p. 75.

[175] D. Whittington. “Possible Adverse Effects of Increasing Block Water Tariffs in Developing Countries”. In: *Economic Development and Cultural Change* 41.1 (1992), pp. 75–87.

[176] K. U. Wittrup-Jensen and K. M. Pedersen. *An Empirical Assessment of the Person Trade-off: Valuation of Health, Framing Effects, and Estimation of Weights for Fairness.* Syddansk Universitet, 2008.

[177] A. C.-C. Yao. “Some Complexity Questions Related to Distributive Computing(Preliminary Report)”. In: *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing.* STOC ’79. Atlanta, Georgia, USA, 1979, pp. 209–213.

[178] S. Zheng. “The Vertex Isoperimetric Problem on Kneser Graphs”. Unpublished research project at Massachusetts Institute of Technology. 2015.