

Markets Beyond Nash Welfare for Leontief Utilities

Ashish Goel

Reyna Hulett

Benjamin Plaut

{ashishg, rmhulett, bplaut}@stanford.edu

Abstract

We study the allocation of divisible goods to competing agents via a market mechanism, focusing on agents with Leontief utilities. 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.

1 Introduction

In a market, buyers and sellers exchange goods according to some sort of pricing system. One of the most important concepts in the study of markets is *market equilibrium*, which describes when the supply provided by the sellers and the demands of the buyers exactly match. Market equilibrium theory dates back to Walras's seminal work in 1874 [35]. In 1954, Arrow and Debreu finally showed that under some mild assumptions, a market equilibrium is guaranteed to exist for a wide class of utility functions [1]. This includes Leontief utilities, which will be our focus.

The simplest mathematical model of a market is a *Fisher market*, first proposed in 1891 by Irving Fisher (see [7] for a modern introduction). A Fisher market consists of a set of goods available for sale, and a set of agents, each with a fixed amount of money to spend. It is usually assumed that agents have no value for leftover money. In Fisher markets, each good j has a single real-number price p_j , and the cost of buying some quantity x of good j is $p_j \cdot x$. We refer to such prices as *linear*, meaning that the cost is proportional to the quantity purchased. A market equilibrium assigns a price to each good such that when each agent purchases her favorite bundle of goods that is affordable under these prices, the demand exactly matches the supply.

There are three motivations behind this work. First, in real market economies, prices are often not linear, and depend on the quantity purchased¹. We refer to prices of this form as *price curves*.

¹One 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 paper; see Section 2.2 for additional discussion.

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. Israel’s pricing policy for water is a good example of this, where each additional unit of water costs more than the previous one [3]. 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 paper 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 that will maximize some objective function. The objective function of a social planner is typically referred to as *welfare*. There are many different social welfare functions, the most well studied being utilitarian welfare (the sum of agent utilities), Nash welfare (the product of agent utilities) [22, 18], and max-min welfare (the minimum agent utility [27, 30, 31])². Max-min welfare can be seen as caring only about equality across individuals. The utilitarian welfare measures overall good across the entire population, possibly at the expense of certain individuals. The Nash welfare is something of a compromise between these two extremes.

Eisenberg and Gale famously showed that for linear prices and a large class of agent utilities (including Leontief), the market equilibria correspond exactly to the allocations maximizing Nash welfare [12, 13]. This result is powerful, but also limiting: 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 paper 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.

1.1 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 [19] 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, and the agent has unweighted Leontief utility over the set of goods corresponding to her desired route.

In the bandwidth allocation setting, price curves correspond naturally to a signaling mechanism that provides congestion signals (eg. in the form of a packet mark or drop) and an end-point protocol such as TCP [8] corresponds naturally to agent responses. It has been known that different marking schemes (such as RED and CHOCe [15, 24]) and versions of TCP lead to different objective functions [23], with CES welfare (also known as “ α -fairness”) being one such objective [6, 20]. However, a market-based understanding was developed only for Nash Welfare, starting with the seminal work of Kelly et al. [19].

²The utilitarian welfare is also known as the Benthamite welfare, after Jeremy Bentham. The max-min welfare is also known as the Rawlsian welfare, after John Rawls, or the egalitarian welfare.

1.2 CES welfare functions

For any constant $\rho \in (-\infty, 0) \cup (0, 1]$, the *constant elasticity of substitution* (CES) welfare function is given by

$$\left(\sum_{i \in N} u_i^\rho \right)^{1/\rho}$$

where u_i is agent i 's utility. Setting $\rho = 1$ yields utilitarian welfare, and the limits as $\rho \rightarrow -\infty$ and $\rho \rightarrow 0$ yield max-min welfare and Nash welfare, respectively. The smaller ρ is, the more the social planner cares about individual equality (max-min welfare being the extreme case of this), and the larger ρ is, the more the social planner cares about overall societal good (utilitarian welfare being the extreme case of this). For this reason, ρ is called the *inequality aversion* parameter.

This class of welfare functions was first proposed by Atkinson [2] and further developed by [5]. The CES welfare function (as opposed to the CES agent utility function) has received very little attention in the computational economics community, despite being extremely influential in the traditional economics literature with [2] having over 8000 citations.

These welfare functions also admit an axiomatic characterization:

1. Monotonicity: if one agent's utility increases while all others are unchanged, the welfare function should prefer the new allocation.
2. Symmetry: the welfare function should treat all agents the same.
3. Continuity: the welfare function should be continuous.
4. Independence of common scale: scaling all agent utilities 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 utility in both allocations.
6. The Pigou-Dalton principle: all things being equal, the welfare function should prefer more equitable allocations [10, 26].

Up to monotonic transformations of the welfare function (which of course do not affect which allocations have better welfare than others), the set of welfare functions that satisfy these axioms is exactly the set of CES welfare functions with $\rho \in (-\infty, 0) \cup (0, 1]$, including Nash welfare [21]³. 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.

2 Results and prior work

We assume throughout the paper 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

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

³This actually does not include max-min welfare, which obeys weak monotonicity but not strict monotonicity.

where M is the set of goods, x_{ij} is the amount of good j agent i receives, and w_{ij} is agent i 's (nonnegative) weight for good j . The setting where $w_{ij} \in \{0, 1\}$ for all i and j is also known as bandwidth allocation.

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 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 Results

2.1.1 A necessary and sufficient condition for the existence of price curves.

Section 4 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 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 4. 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 4.3)⁴. This characterization results in a polynomial time algorithm to compute the underlying price curves or show that none exist (Theorem 4.4). Appendix B gives an analogous characterization theorem and polynomial time algorithm for weakly increasing price curves (Theorems B.2 and B.3).

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 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 [25] to complete the proof. The algorithm for computing price curves is a linear program involving the agent-order matrix.

2.1.2 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?

⁴This result extends to the setting of unequal budgets if one instead considers “budget-weighted group-domination-freeness”. We elaborate on this in Section 4.

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

Example 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.

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 welfare⁵. 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 1.1 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 5.1). 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 dual⁶. This result can be thought of as extending the work on price-based congestion control (pioneered by Kelly et al. [19]) 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 5.2). This is analogous to the famous result of Eisenberg and Gale: the linear-pricing equilibrium allocations are exactly the allocations maximizing Nash welfare [12, 13].

One may wonder if Theorem 5.1 could be extended to $\rho = 1$, i.e., maximizing the sum of utilities. Example 1 shows that the answer is no, unfortunately. One may also wonder if Theorem 5.1 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 C.1 in Section C.

2.1.3 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

⁵This 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.

⁶These results extend to agents with unequal budgets if we instead consider the “budget-weighted CES welfare”, i.e., $(\sum_{i \in N} B_i u_i^\rho)^{1/\rho}$, where B_i is agent i ’s budget. We discuss this in Section 5.4. The price curves will take the exact same form.

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. [19], 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 5.2.

2.1.4 Additional results.

We prove two additional results. First, we consider max-min welfare in Appendix A, and show that as long as agents' weights are reasonably normalized, allocations with optimal max-min welfare can be supported. Second, in Appendix B, we give a characterization theorem for weakly increasing price curves (Theorem B.2). 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 B.3).

2.2 Prior work

The study of markets has a long history in the economics literature [35, 33, 1, 7]. Recently, this topic has received significant attention in the computer science community as well (see [34] 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, [16] 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, see Israel's pricing policy for water [3]).

The concept of a social welfare function – a function which encapsulates a societal value system – was first proposed in 1938 [4], and further developed by [28]. For brevity, we will just refer to them as welfare functions. Since then, several different welfare functions have been proposed, the most well-studied being utilitarian welfare, Nash welfare [22, 18], and max-min welfare (the minimum agent utility) [27, 30, 31]. The class of CES welfare functions was first proposed by [2] and further developed by [5], although not under the same name. See [21] for a modern discussion of welfare functions.

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 [29]. This is clearly unacceptable, so we sacrifice strategy-proofness in favor of Pareto optimality. Specifically, we assume throughout the paper that agents always truthfully report their preferences.

The remainder of the paper is organized as follows. Section 3 formally defines the model. Section 4 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 5 gives our second main result: that in the bandwidth allocation setting, every maximum CES welfare allocation can be supported by price curves. At this point we conclude the main paper, and move on to supplementary results. Appendix A shows that allocations with optimal max-min welfare can be supported by price curves in a wide range of settings. In Appendix B, we generalize our characterization theorem from Section 4 to account for weakly increasing price curves. Appendix C 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, Appendix D provides some proofs that are omitted from earlier sections.

3 Model

Let $N = \{1, 2, \dots, n\}$ be the set of agents, and let M be a set of divisible goods, where $|M| = m$. Throughout the paper, we use i and k to refer to agents and j and ℓ to refer to goods. Let s_j be the available supply of good j .⁷ The social planner needs to determine an *allocation* $\mathbf{x} \in \mathbb{R}_{\geq 0}^{n \times m}$, where $x_i \in \mathbb{R}_{\geq 0}^m$ is the *bundle* of agent i , and $x_{ij} \in [0, s_j]$ is the fraction of good j allocated to agent i . An allocation cannot allocate more than the available supply: $\mathbf{x}$ is a valid allocation if and only if $\sum_i x_{ij} \leq s_j$ for all j .

Agent i 's utility for a bundle x_i is denoted by $u_i(x_i) \in \mathbb{R}_{\geq 0}$. The literature studies many subclasses of utility functions. For example, *Leontief* utilities have the form

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

where w_{ij} represents the 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 paper, and we assume that each agent has nonzero weight for at least one good.

Just as agents have utilities over the bundles they receive, we can imagine a social planner who wishes to design a pricing mechanism to maximize some societal welfare function $\Phi(\mathbf{x})$. One can think of Φ as the social planner's utility function, which takes as input the agent utilities, instead of a bundle of goods. The most well-studied welfare functions are the *max-min welfare* $\Phi(\mathbf{x}) = \min_{i \in N} u_i(x_i)$, the *Nash welfare* $\Phi(\mathbf{x}) = (\prod_{i \in N} u_i(x_i))^{1/n}$, and the *utilitarian welfare* $\Phi(\mathbf{x}) = \sum_{i \in N} u_i(x_i)$.⁸ These three welfare functions can be generalized by a CES welfare function:

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

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

3.1 Fisher markets

The simplest market model is a *Fisher market* [7]. In a Fisher market, each good is available for sale and each agent enters the market with a fixed budget she wishes to spend. It is assumed

⁷It is common to normalize the supply of each good to 1, but this is not possible when weights are restricted to be either 0 or 1.

⁸One could also imagine the (arguably less natural) case of a social planner who cares about some agents' utilities more than others, which would manifest as weights in her utility function. We briefly consider this case in the setting of CES welfare with unequal budgets (Section 5.4).

that agents have no value for leftover money; this will imply that each agent always spends her entire budget. Unless otherwise stated, we will assume that all agents have the same budget, and normalize all budgets to 1 without loss of generality.

Informally, a *Fisher market equilibrium* assigns a price $p_j \in \mathbb{R}_{\geq 0}$ to each good j so that the agents' demand equals supply. Formally, for a price vector $\mathbf{p} \in \mathbb{R}_{\geq 0}^m$, the *cost* of bundle x_i is $C_{\mathbf{p}}(x_i) = \sum_{j \in M} x_{ij} p_j$. 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.,

$$D_i(\mathbf{p}) = \arg \max_{x_i \in \mathbb{R}_{\geq 0}^m: C_{\mathbf{p}}(x_i) \leq 1} u_i(x_i).$$

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 through the paper, 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 paper 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 $\mathbf{p} \in \mathbb{R}_{\geq 0}^m$ 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 $\sum_{i \in N} x_{ij} = s_j$.

When all agents have the same budget, this is also called the *competitive equilibrium from equal incomes* [33].

For a wide class of agent utilities, including Leontief utilities, an equilibrium is guaranteed to exist [1]⁹. Furthermore, the equilibrium allocations are the exactly the allocations which maximize Nash welfare¹⁰. This is made explicit by the celebrated Eisenberg-Gale convex program [12, 13], and combinatorial approaches to computing market equilibria [11, 17].

3.2 Price curves

Our paper 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* $f_j : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$. 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 $\mathbf{f} = (f_1, \dots, f_m)$, 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

⁹Specifically, an equilibrium is guaranteed to exist as long as 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.

¹⁰The 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 [12] and generalized slightly by [17].

across goods. Each agent’s demand set is defined identically to the Fisher market setting: $D_i(\mathbf{f}) = \arg \max_{x_i \in \mathbb{R}_{\geq 0}^m: C_{\mathbf{f}}(x_i) \leq 1} u_i(x_i)$.

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 $j \in M$ ¹¹.

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

4 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 Appendix B. However, this condition (“locked-agent-freeness”) is somewhat unwieldy. Although weakly increasing price curves are sometimes necessary¹², for now we will consider only *strictly* increasing price curves. The corresponding necessary and sufficient condition is the cleaner notion of group-domination-freeness.

4.1 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 weight¹³ 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

¹¹For 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.

¹²Consider 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 is only possible if the price of good 2 is zero: otherwise, 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.

¹³Note that this is *not* the same weight as the w_{ij} representing an agent’s weight for a good.

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 $\mathbf{b}$ ¹⁴.

The formal definition of this condition is below.

Definition 4.1 (Group-domination-free (GDF)). Let $\mathbf{a} = (a_1, \dots, a_n)$ and $\mathbf{b} = (b_1, \dots, b_n)$ be vectors in $\mathbb{R}_{\geq 0}^n$ 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 $\tau \in \mathbb{R}_{\geq 0}$,

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

and there exists a (j, τ) pair where the inequality is strict. $\mathbf{x}$ is group-domination-free if there do not exist vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}_{\geq 0}^n$ such that $\mathbf{a} \succ \mathbf{b}$ in $\mathbf{x}$.¹⁵

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

4.2 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 3.1). Then the condition that $x_i \in D_i(\mathbf{f})$ for all i can be captured by Lemma 4.1, whose proof appears in Appendix D. 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.

¹⁴See [14] for an introduction to stochastic dominance.

¹⁵The “-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 to the existence of a market equilibrium in this setting.




Figure 1: An illustrative example allocation and the construction of the corresponding agent-order matrix.

Lemma 4.1. *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 4.3 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 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.

First we address observation (1). Consider the possible price curves shown in Figure 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).

$$\begin{bmatrix} 1 & 1 & 1 & 0 & 1 & 0 & -1 \\ 1 & 1 & 0 & 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & 1 & 1 & 1 & -1 \\ 0 & 0 & 0 & 0 & 1 & 1 & -1 \end{bmatrix} \times \begin{bmatrix} 0 \\ 1 \\ -0 \\ -0 \\ -0 \\ -1 \\ 1 \end{bmatrix} = \mathbf{0}$$

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

$$\begin{bmatrix} 1 & 1 & 1 & 0 \\ 1 & 1 & 0 & 0 \\ -\frac{1}{0} & -\frac{0}{0} & -\frac{0}{1} & -\frac{0}{0} \\ -\frac{1}{0} & -\frac{0}{0} & -\frac{1}{1} & -\frac{0}{1} \\ -\frac{0}{0} & -\frac{0}{0} & -\frac{1}{1} & -\frac{1}{1} \\ -\frac{1}{-1} & -\frac{1}{-1} & -\frac{1}{-1} & -\frac{1}{-1} \end{bmatrix} \times \begin{bmatrix} 1 \\ -1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ -\frac{1}{0} \\ -\frac{1}{1} \\ -\frac{0}{0} \\ -\frac{0}{0} \end{bmatrix}$$

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

Figure 2: Example row and column sums of the agent-order matrix from Figure 1b.

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 1b. Within a sub-block, each column will correspond to a non-zero agent allocation (i.e., the non-zero points shown in Figure 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 2a exhibits such a vector $\mathbf{y}$ for the matrix A shown in Figure 1b. $\mathbf{y}$ will represent prices, so we require all the entries to be non-negative, denoted $\mathbf{y} \geq 0$; for strictly increasing price curves, we require $\mathbf{y}$ to be strongly positive¹⁶, 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 4.1 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 positive¹⁷, denoted $A^T\mathbf{z} > \mathbf{0}$. For instance, Figure 2b exhibits such a vector $\mathbf{z}$ for the matrix A shown in Figure 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 [25], 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

¹⁶Recall that a strongly positive vector has every entry greater than 0.

¹⁷Recall that a strictly positive vector has entries in $\mathbb{R}_{\geq 0}$ with at least one entry non-zero.

Stiemke to prove Theorem 4.3.

Lemma 4.2 (1.6.4 in [32]). *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 4.3. *Let $\mathbf{x}$ be any allocation that obeys the supply constraints and gives at least one agent a nonempty bundle. Then $\mathbf{x}$ can be supported by strictly increasing price curves if and only if $\mathbf{x}$ 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 4.1, 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 $\tau_j^1, \tau_j^2, \dots, \tau_j^{|X_j|}$ such that $\tau_j^1 < \tau_j^2 < \dots < \tau_j^{|X_j|}$. 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 $f'_j : X_j \rightarrow \mathbb{R}_{\geq 0}$ such that $0 < f'_j(\tau_j^1) < f'_j(\tau_j^2) < \dots < f'_j(\tau_j^{|X_j|}) \forall j$ 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 4.2. 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 4.2 that $\exists \mathbf{y} \gg \mathbf{0}$ such that $A\mathbf{y} = \mathbf{0}$ 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 $0 < f'_j(\tau_j^1) < \dots < f'_j(\tau_j^{|X_j|}) \forall j$. 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 $\mathbf{a} \succ \mathbf{b}$. $\square$

Remark. Since the matrix A constructed in the proof of Theorem 4.3 is over the rationals, we can also assume that the $\mathbf{y}$ or $\mathbf{z}$ obtained via Lemma 4.2 are over the rationals. In particular, we can scale $\mathbf{z}$ to obtain $\mathbf{z}' \in \mathbb{Z}^n$ with $A^T \mathbf{z}' > \mathbf{0}$. 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 4.4. *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.*

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

Proof. As per the proof of Theorem 4.3, there exist strictly increasing price curves supporting $\mathbf{x}$ if and only if there is a solution to the system $A\mathbf{y} = \mathbf{0}, \mathbf{y} \gg \mathbf{0}$. 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 $y_{-1} = 1$, 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. $\square$

One may wonder if Theorem 4.3 generalizes to other classes of utility functions. Unfortunately, the answer in general is no. Example 3 gives an instance with linear utilities that is GDF, but cannot be supported by price curves.

In Section A, 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.

4.3 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 4.2 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.

5 CES welfare

In this section, we consider CES welfare functions:

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

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 5.1). We present this result this in Section 5.1. 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. [19] (Section 5.2). We also give a converse of sorts to Theorem 5.1 (Section 5.3), and briefly discuss the case of unequal budgets (Section 5.4). Throughout this section, we let $R_i = \{j \in M : w_{ij} = 1\}$ for brevity.

5.1 Main CES welfare result

We now state and prove Theorem 5.1. 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 5.1. *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:

$$\begin{aligned} \max_{\substack{\mathbf{x} \in \mathbb{R}_{\geq 0}^{n \times m}, \\ \mathbf{u} = (u_1 \dots u_n) \in \mathbb{R}_{\geq 0}^n}} & \left(\sum_{i \in N} u_i^\rho \right)^{1/\rho} \\ \text{s.t.} & u_i \leq x_{ij} \quad \forall i \in N, j \in R_i \\ & \sum_{i \in N} x_{ij} \leq s_j \quad \forall j \in M \end{aligned}$$

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 $\left(\sum_{i \in N} u_i^\rho\right)^{1/\rho}$ 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.¹⁸ Thus consider a new convex program with objective function

$$\max_{\mathbf{x} \in \mathbb{R}_{\geq 0}^{n \times m}, \mathbf{u} \in \mathbb{R}_{\geq 0}^n} \frac{1}{\rho} \sum_{i \in N} u_i^\rho, \text{ and the same constraints.}$$

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

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

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 $u_i^* = (\sum_{j \in R_i} q_j^*)^{\frac{1}{\rho-1}}$. 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 $x_{ij}^* > 0$.

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

$$x_{ij}^* = u_i^* = \left(\sum_{j \in R_i} q_j^* \right)^{\frac{1}{\rho-1}}$$

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^* :

$$C_{\mathbf{f}}(x_i^*) = \sum_{\ell \in M} q_\ell^* (x_{i\ell}^*)^{1-\rho} = \sum_{\ell: q_\ell, x_{i\ell} \neq 0} q_\ell^* (x_{i\ell}^*)^{1-\rho} = \sum_{\ell: q_\ell, x_{i\ell} \neq 0} q_\ell^* \left(\sum_{j \in R_i} q_j^* \right)^{\frac{1-\rho}{\rho-1}} = \sum_{\ell: q_\ell, x_{i\ell} \neq 0} \frac{q_\ell^*}{\sum_{j \in R_i} q_j^*}$$

¹⁸We add a factor of $1/\rho$ instead of ρ because this will slightly simplify the analysis.

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 $u_i^* = \left(\sum_{j \in R_i} q_j^* \right)^{\frac{1}{\rho-1}}$ 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 4.1, 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. $\square$

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 $\rho \rightarrow -\infty$ 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 $\rho \rightarrow 0$ 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 $\rho > 0$.

5.2 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 5.1, 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 ¹⁹.

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. [19].

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

¹⁹In general, different values of ρ will lead to different optimal allocations and Lagrange multipliers.

maximize Nash welfare, and gives a tâtonnement process that is akin to gradient descent on the dual program [9]. 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.

5.3 A converse to Theorem 5.1

In this section, we give a converse of sorts to Theorem 5.1: 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 3.1.

The proof of Theorem 5.2 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 5.1, which is based on the fact that the KKT conditions are necessary for optimality. The formal proof appears in Appendix D.

Theorem 5.2. *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.*

5.4 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* $\Phi_{CES}(\mathbf{x}) = (\sum_{i \in N} B_i u_i(x_i)^\rho)^{1/\rho}$, then the proof of Theorem 5.1 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 $C_{\mathbf{f}}(x_i) = \sum_{\ell \in R_i} \frac{B_i q_\ell^*}{\sum_{j \in R_i} q_j^*} = B_i$, 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.

6 Conclusion

In this paper, 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 5.2. We think that the approach of [9] 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 Appendix C some of the challenges that linear utilities pose for analyzing price curves, but perhaps everything would fall into place with the right framework.

Third, future research could consider *quasilinear* utilities. In this paper, we assumed that agents have fixed budgets, and have no value for leftover money (i.e., “fake money”). In the quasilinear setting (i.e., “real money”), agents also choose how much to spend, and each agent’s utility function is equal to the value she derives from her bundle minus the amount she pays. It would be interesting to see whether our results extend to that setting.

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.

References

- [1] Kenneth J. Arrow and Gerard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, 1954.
- [2] Anthony B Atkinson. On the measurement of inequality. *Journal of Economic Theory*, 2(3):244 – 263, 1970.
- [3] Nir Becker. *Water Pricing in Israel: Various Waters, Various Neighbors*, pages 181–199. Springer International Publishing, Cham, 2015.
- [4] Abram Bergson. A reformulation of certain aspects of welfare economics. *The Quarterly Journal of Economics*, 52(2):310–334, 1938.
- [5] Charles Blackorby and David Donaldson. Measures of relative equality and their meaning in terms of social welfare. *Journal of Economic Theory*, 18(1):59 – 80, 1978.
- [6] Thomas Bonald and Laurent Massoulié. Impact of fairness on internet performance. *SIG-METRICS Perform. Eval. Rev.*, 29(1):82–91, June 2001.
- [7] William C. Brainard and Herbert E. Scarf. How to compute equilibrium prices in 1891. *American Journal of Economics and Sociology*, 64(1):57–83, 2005.
- [8] V. Cerf and R. Kahn. A protocol for packet network intercommunication. *IEEE Transactions on Communications*, 22(5):637–648, May 1974.
- [9] Yun Kuen Cheung, Richard Cole, and Nikhil Devanur. Tatonnement beyond gross substitutes?: gradient descent to the rescue. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 191–200. ACM, 2013.

- [10] Hugh Dalton. The measurement of the inequality of incomes. *The Economic Journal*, 30(119):348–361, 1920.
- [11] Nikhil R. Devanur, Christos H. Papadimitriou, Amin Saberi, and Vijay V. Vazirani. Market equilibrium via a primal–dual algorithm for a convex program. *J. ACM*, 55(5):22:1–22:18, November 2008.
- [12] E. Eisenberg. Aggregation of Utility Functions. *Management Science*, 7(4):337–350, July 1961.
- [13] Edmund Eisenberg and David Gale. Consensus of Subjective Probabilities: The Pari-Mutuel Method. *The Annals of Mathematical Statistics*, 30(1):165–168, March 1959.
- [14] William Feller. *An Introduction to Probability Theory and Its Applications, Vol. 1, 3rd Edition*. Wiley, 3rd edition, January 1968.
- [15] S. Floyd and V. Jacobson. Random early detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, Aug 1993.
- [16] Duncan K. Foley. Lindahl’s Solution and the Core of an Economy with Public Goods. *Econometrica*, 38(1):66, January 1970.
- [17] Kamal Jain and Vijay V. Vazirani. Eisenberg–Gale markets: Algorithms and game-theoretic properties. *Games and Economic Behavior*, 70(1):84–106, September 2010.
- [18] Mamoru Kaneko and Kenjiro Nakamura. The nash social welfare function. *Econometrica*, 47(2):423–35, 1979.
- [19] F P Kelly, A K Maulloo, and D K H Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49(3):237–252, Mar 1998.
- [20] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5), 2000.
- [21] Hervé Moulin. *Fair Division and Collective Welfare*, chapter 3. MIT Press, January 2003.
- [22] John Nash. The Bargaining Problem. *Econometrica*, 18(2):155–162, April 1950.
- [23] Jitendra Padhye, Victor Firoiu, Don Towsley, and Jim Kurose. Modeling tcp throughput: A simple model and its empirical validation. *SIGCOMM Comput. Commun. Rev.*, 28(4):303–314, October 1998.
- [24] Rong 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)*, volume 2, pages 942–951 vol.2, 2000.
- [25] Cerng-tiao Perng. On a class of theorems equivalent to farkas’s lemma. *Applied Mathematical Sciences*, 11(44):2175–2184, 2017.
- [26] A.C. Pigou. *Wealth and Welfare*. PCMI collection. Macmillan and Company, limited, 1912.

- [27] John Rawls. *A Theory of Justice*. Belknap Press of Harvard University Press, Cambridge, Massachussets, 1 edition, 1971.
- [28] Paul Anthony Samuelson. Foundations of economic analysis. 1947.
- [29] James Schummer. Strategy-proofness versus efficiency on restricted domains of exchange economies. *Social Choice and Welfare*, 14(1):47–56, Dec 1996.
- [30] Amartya Sen. Welfare inequalities and rawlsian axiomatics. *Theory and Decision*, 7(4):243–262, Oct 1976.
- [31] Amartya Sen. Social choice theory: A re-examination. *Econometrica*, 45(1):53–89, 1977.
- [32] Josef Stoer and Christoph Witzgall. *Convexity and Optimization in Finite Dimensions I*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1970.
- [33] Hal Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63 – 91, 1974.
- [34] Vijay V Vazirani. Combinatorial algorithms for market equilibria. *Algorithmic Game Theory*, pages 103–134, 2007.
- [35] L. Walras. *Éléments d’économie politique pure; ou, Théorie de la richesse sociale*. Number v. 1-2. Corbaz, 1874.

Acknowledgements

This research was supported in part by NSF grant CCF-1637418, ONR grant N00014-15-1-2786, and the NSF Graduate Research Fellowship under grant DGE-1656518.

A 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 A.1 states that under a quite general normalization assumption, we can support a max-min allocation.

Theorem A.1. *Suppose there exist strictly increasing functions $g_1, \dots, g_m$ 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 4.3, so $\mathbf{x}$ is GDF, which completes the proof. $\square$

One natural corollary of Theorem A.1 is the following:

Corollary A.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 A.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.²⁰ Furthermore, the following counterexample falls under the even simpler bandwidth allocation setting: $w_{ij} \in \{0, 1\}$ for all i, j .

Theorem A.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

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. $\square$

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 A.3. *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 $\Phi_{MM}(\mathbf{x}) \geq (1 - \epsilon) \max_{\mathbf{x}'} \Phi_{MM}(\mathbf{x}')$.*

²⁰The L_∞ norm is defined as $\max_{j \in M} w_{ij}$.

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 A.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 $\mathbf{x}'$,

$$\begin{aligned} \min_{i \in N} \frac{x_{ij}}{w'_{ij}} &\geq \min_{i \in N} \frac{x'_{ij}}{w'_{ij}} \\ \min_{i \in N} \frac{\alpha_i x_{ij}}{w_{ij}} &\geq \min_{i \in N} \frac{\alpha_i x'_{ij}}{w_{ij}} \\ \min_{i \in N} \alpha_i u_i(x_i) &\geq \min_{i \in N} \alpha_i u_i(x'_i) \end{aligned}$$

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 $\lim_{q \rightarrow \infty} \alpha_i = (\sum_{j \in M} w_{ij}^q)^{1/q} = 1$ for all $i \in N$, which completes the proof. $\square$

B Characterization of allocations supported by weakly increasing price curves

In Section 4, 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 4, 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 4.1). This requirement can be expressed by *locked-agent-freeness*.

Definition B.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 4.1. 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 B.2.

The proof of Theorem B.2 is similar to the proof of Theorem 4.3 for strictly increasing price curves. The main difference is that strictly increasing price curves trivially satisfy condition (b) of Lemma 4.1, 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 4.2 (although still equivalent to Farkas’s Lemma [25]), this one due to Motzkin. Recall that $\mathbf{v} > \mathbf{0}$ denotes a strictly positive vector, and $\mathbf{v} \gg \mathbf{0}$ strongly positive.

Lemma B.1 (1.6.1 in [32]). *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 B.2. *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 4.1, 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 $\tau_j^1 < \tau_j^2 < \dots < \tau_j^{|X_j|}$. 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 $f'_j : X_j \rightarrow \mathbb{R}_{\geq 0}$ such that

1. for all $j \in M$, $0 \leq f'_j(\tau_j^1) \leq f'_j(\tau_j^2) \leq \dots \leq f'_j(\tau_j^{|X_j|})$ (weakly increasing)




Figure 3: Example construction of the marginal-cost matrix from an agent-order matrix.

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 B.1. As in the proof of Theorem 4.3, 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 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

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

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 B.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 $\mathbf{y} > \mathbf{0}$, so without loss of generality we can assume the last entry of $\mathbf{y}$ is 1. Thus as before, $A\mathbf{y} = \mathbf{0}$ 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 $\mathbf{x}$. $\square$

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

Theorem B.3. *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.*

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

Proof. As per the proof of Theorem B.2, 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 $y_{-1} = 1$, 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. $\square$

C Counterexamples

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

Example 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 is 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 5 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 C.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 C.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 5.1 given in Section 5, we used duality to show that for a fixed ρ , any maximum CES welfare allocation $\mathbf{x}$ has the form

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

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

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

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

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

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 B.2, 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

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

$$\begin{aligned}
(\rho - 1) \ln(1 - \varepsilon) &\geq \ln 2 - \ln(2 - \varepsilon) \\
\rho &\leq 1 + \frac{\ln 2 - \ln(2 - \varepsilon)}{\ln(1 - \varepsilon)}
\end{aligned}$$

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. $\square$

C.1 Difficulties in analyzing linear utilities

We assumed throughout the paper 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

$$\arg \max_{x_i \in \mathbb{R}_{\geq 0}^m: C_{\mathbf{f}}(x_i) \leq 1} u_i(x_i)$$

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 3).

D Omitted proofs

Lemma 4.1. *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. ($\Leftarrow$) 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, 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

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

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

Example 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.

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

($\implies$) 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 $x_i \in D_i(\mathbf{f})$. $\square$

Theorem 5.2. 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 [12, 13]. 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 5.1. We use the same convex program for maximizing CES welfare, which, as stated in the proof of Theorem 5.1, 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 $\frac{\partial L}{\partial x_{ij}}(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*) = \lambda_{ij}^* - q_j^*$; by definition, this is equal to zero. For $j \notin R_i$, we have $\frac{\partial L}{\partial x_{ij}}(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*) = -q_j^*$. 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 $\frac{\partial L}{\partial u_i}(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \mathbf{q}^*) = u_i^{*\rho-1} - \sum_{j \in R_i} \lambda_{ij}^* = 0$ for 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

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

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

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

$$u_i^{*\rho-1} - \sum_{j \in R_i} q_j^* = 0$$

$$u_i^{*\rho-1} - \sum_{j \in R_i} \lambda_{ij}^* = 0$$

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. $\square$