Divisible goods markets with budget constraints: a unification of revenue and welfare

Simon Finster* Paul W. Goldberg† Edwin Lock‡

February 2022

Abstract

Markets with multiple divisible goods and quasilinear buyers have been studied widely from the perspective of revenue and welfare. In general, it is well-known that envy-free revenue-maximal outcomes can result in lower welfare than competitive equilibrium outcomes. We study the regime in which buyers have quasilinear utilities with budget constraints, and the seller must find prices and an envy-free allocation that maximises revenue or welfare. Our setup mirrors markets such as ad auctions and the arctic product-mix auction proposed for the exchange of financial assets. Our main result is to show that the unique competitive equilibrium prices are also envy-free revenue-maximal. This coincidence of maximal revenue and welfare is surprising, and breaks down even when buyers have piecewise-linear valuations. In our result, we present a novel characterisation of the set of 'feasible' prices at which demand does not exceed supply, show that this set has an elementwise-minimal price vector, and demonstrate that these prices maximise revenue and welfare. The proof also implies an algorithm for finding this unique price vector.

1 Introduction

A central concern in markets for divisible goods is finding procedures for computing competitive equilibrium [2, 13, 15, 25, 26]. While much of the literature has studied settings in which buyers have no financial limitations, some work has begun to explore the implications of imposing buyer budgets on the existence and computation of equilibrium, mainly with quasilinear preferences [8, 10, 11, 14, 24]. Competitive equilibrium (CE) is then a standard, desirable market objective, as the fundamental theorems of welfare economics imply that a CE allocation also maximises social welfare. In many practical settings (e.g. the arctic auction [21]), however, the alternative objective of envy-free revenue maximisation, together with the option for buyers to express budget constraints, is often at least as attractive. In an envy-free allocation that maximises revenue, buyers receive bundles they demand at market prices, but the seller may prefer to allocate only a subset of her supply [19, 21]. Solutions for the two objectives of maximising revenue or welfare do not generally coincide; in fact, in many common economic settings the two objectives diverge significantly in the market outcome.[1]

This paper identifies an important market setting that unifies the two objectives of envy-free revenue maximisation and welfare maximisation. Our market contains multiple buyers and one seller. The seller supplies multiple goods in finite, divisible quantities. Buyers have quasi-linear utilities and budget constraints; that is, every buyer can set an upper limit on the amount they wish to spend. Each buyer has an additive valuation, i.e. a fixed per-unit value for each good.[2] While setups like this are important in practice (we elaborate on ad auctions below), they have received comparatively little attention in theoretical work. Our main contribution is to demonstrate that the two objectives of social

*Nuffield College and Department of Economics, University of Oxford and University of Hagen, simon.finster@fernuni-hagen.de.

†Department of Computer Science, University of Oxford, paul.goldberg@cs.ox.ac.uk

‡Nuffield College and Department of Economics, University of Oxford, edwin.lock@economics.ox.ac.uk

[1] This is the case in divisible goods markets with concave preferences (see Section 1.2). This is also the case in standard indivisible goods markets, for instance, in the strong-substitutes product-mix auction, where competitive equilibrium can be computed efficiently [3] but solving for revenue is APX-hard [19].

[2] Note that budget constraints are especially natural for additive valuations on unbounded domains, as they are necessary to ensure finite quasi-linear demand. Infinite demand would make envy-free allocation impossible.

welfare and envy-free revenue coincide in quasilinear budget-constrained markets: revenue is maximised at competitive equilibrium prices. To prove this, we introduce a novel geometric object, the *feasible region*, which is defined as the set of prices at which, for every good, either the market clears or there is excess supply. This region has a smallest element which, as we show in this paper, is the revenue- and welfare-maximising price.

In applications, buyers with additive valuations are of particular interest. For example, Klemperer [21] introduced the 'arctic auction', originally developed for the government of Iceland, in which buyers have budget-constrained preferences for multiple divisible goods. The market we consider can be interpreted as a special case of this auction. The government planned to use the arctic auction to exchange blocked accounts for other financial assets, e.g. cash or bonds of different quality, available with limited supply.[3] The auction was tailored to this use because buyers could submit a budget, as well as trade-offs between different assets through bid prices. The objective of Klemperer's auction was to maximise revenue subject to envy-freeness, which is attractive for the seller (government) as well as bidders. We show in Section 2.1 that each buyer in the arctic auction has preferences corresponding to the aggregation of multiple budget-constrained quasilinear buyers with additive valuations.

Our market is also a suitable model of many ad auctions as they occur in practice. An identical setting, but with an ex-ante different notion of equilibrium, has recently been addressed by Conitzer et al. [10]. When advertising companies compete for web space (goods) to display their ads, the decision of which publisher to choose and bid for is difficult. It may be intuitive, however, to choose an advertising budget and state demand in terms of 'limit market prices' for multiple, distinct goods, below which the seller, or the market platform, allocates those goods with the highest value for money. This is not only conceptually easier for advertisers, but also practical, feasible, and desirable from the perspective of social welfare *and* revenue, as demonstrated by our results for the budget-constrained quasilinear setting.

The market we study has recently been labelled a 'quasi-Fisher' market (Murray et al. [24]), as it can be considered a generalisation of Fisher markets in which buyers have quasi-linear utilities. In standard Fisher markets, buyers spend their entire budget at any market prices, and so revenue is constant at all prices. In contrast, buyers with quasi-linear utility and budget constraints spend nothing when prices are unacceptably high. Hence, the notion of maximising revenue becomes a viable objective for the seller to pursue. Quasi-Fisher markets have been studied from a (robust) competitive equilibrium perspective [8, 24] as well as from the perspective of Nash social welfare [9]. While competitive equilibrium and Nash social welfare are desirable in terms of stability, designing an auction to maximise revenue while respecting envy-freeness may better reflect the primary interests of the seller without affecting the stability of the allocation.

Our contributions. In this paper, we show that when buyers have budget-constrained utilities with additive valuations, the two objectives of maximising revenue and finding a competitive equilibrium coincide. The unique market-clearing prices are buyer-optimal among all revenue-maximising prices in the sense that they maximise the quantities allocated to each buyer, and thus maximise buyer utilities. Building on the work of Chen et al. [8], our results imply that efficient inner-point methods can be used to find revenue-maximising prices and allocation. Our proof proceeds geometrically and may thus convey insights of independent interest. First we show that the set of 'feasible' prices (in the sense that aggregate demand does not exceed supply) contains an elementwise-minimal price vector. This is not immediately clear, as the set of feasible prices is not convex. We then complete the argument by showing that this elementwise-minimal price vector maximises revenue *and* clears the market. In order to demonstrate the latter, we adapt a price-scaling procedure from Adsul et al. [1]'s simplex-like algorithm for linear Fisher markets to our market setting. Given non-minimal feasible prices, our procedure scales down the prices of a subset of goods while maintaining feasibility.[4] Just as Adsul et al. [1] derive a simplex-like algorithm for solving linear algorithms by repeatedly calling their price-scaling procedure, we see that applying our procedure finitely many times leads to an algorithm for finding the elementwise-minimal prices that may be more efficient and easily implementable in practice than existing interior point methods.

[3] Other related applications may include debt restructuring and the (re-)division of firms between shareholders [23].

[4] The procedure of Adsul et al. [1] scales prices up instead of down. A second difference is that buyers in our setting can choose not to spend (a part of) their budget; recall that buyers in Fisher markets always spend their entire budget. This requires additional pre-processing for buyers that may be indifferent between spending all of the budget or only a fraction thereof at given prices.

Organisation. The remainder of this section discusses related literature (Section 1.1) and highlights the necessity of additive valuations for our result (Section 1.2). In Section 2, we define our model. Our main theorem on the coincidence of revenue and welfare is stated and proved in Section 3. The price-reduction procedure is described in Section 4. In Section 5, we state our theorem on finding market clearing prices, and Section 6 concludes.

1.1 Related literature

Computing equilibrium prices in multi-good markets has raised great interest among economists and computer scientists alike, most notably for the divisible goods market due to Léon Walras [27], for which Arrow and Debreu in their famous result [2] proved the existence of a general market equilibrium. Since then, much algorithmic machinery has been developed for the Arrow-Debreu market (e.g. [15, 20]) and its important special case, the Fisher market [6]. Fisher markets are a simple, economic setting with well-known tractable algorithms for the computation of competitive market equilibria. The famous Eisenberg-Gale convex program [16], originally introduced for linear utilities, has later been studied in the context of various types of preferences such as utilities particular spending constraints [5, 12, 18] or additional transaction costs [7].

Our model of budget-constrained buyers with quasilinear utilities and additive valuations can be understood as a 'quasi-Fisher' market [24]. In the linear Fisher setting, money has no intrinsic value: buyers do not gain any utility from any leftover budget. In contrast to the Arrow-Debreu market, prices are not scalable in neither the linear nor the quasilinear Fisher market setting. Moreover, contrasting with the linear Fisher market, in the budget-constrained buyer market with quasilinear utilities, demand is not invariant to scaling prices, in the sense that, for every buyer, there exists a limit price for each good beyond which the buyer does not demand the good.

While linear Fisher markets have received significant attention in the literature (e.g. Adsul et al. [1], Devanur et al. [13], Orlin [26]), their counterpart with quasilinear utilities has appeared in various guises. The first results on quasilinear Fisher markets were developed by Chen et al. [8]. They show that competitive equilibria can be computed in polynomial time, using interior point algorithms for convex programmes. Cole et al. [9] consider the problem of maximising Nash social welfare in, mainly in linear Fisher markets spending restricted utility models, but also with extensions to Leontief, CES utility, and quasilinear utility.

In contrast with this literature, our paper considers not only competitive equilibrium and social welfare as solution concept and market objective, but also maximising revenue subject to envy-freeness. Our unifying result demonstrates the importance of the quasilinear setting, from a theoretical perspective as well as in applications. The practical relevance is highlighted also by a recent paper by Conitzer et al. [10]. Their setting is particularly inspired by online ad auctions, and the basic properties of the market, buyers, goods, and budget-constrained preferences with additive values are identical to ours. However, unlike in our setting, each divisible good is sold in an independent, single-unit first-price auction, where only the highest bidders can win a positive quantity. Conitzer et al. [10] introduce the solution concept of *first price pacing equilibria* (FPPE), in which the bids submitted correspond to the buyers values scaled (uniformly for each buyer) by a multiplicative factor, the pacing multiplier.[5] Interestingly, this at first sight unrelated auction procedure can also be solved using the modified Eisenberg-Gale convex programme by Chen et al. [8]. Moreover, Conitzer et al. [10] show that the unique FPPE corresponds to a competitive equilibrium in the sense of our setting, that is, in the overarching market for all goods with budget-constrained, quasilinear buyers. While the authors show that the FPPE is revenue-maximal among all budget-feasible pacing multipliers and corresponding allocations, our work implies that the FPPE is indeed revenue-maximising subject to envy-freeness in the entire market.

1.2 Introductory examples

The purpose of the first example is to demonstrate the difficulties one encounters with diminishing marginal values and motivates our interest in additive valuations. We also introduce some of the geo-

[5] An FPPE is defined as a set of pacing multipliers (one for each buyer) and allocations that satisfy the allocation and pricing rule of standard first-price auctions, as well as budget feasibility, supply feasibility, market clearing for demanded goods, and 'no unnecessary pacing', i.e. a buyer's multiplier equals one if she has unspent budget.

metric properties of our framework in a two-good example with additive valuation, which are natural to work with for our objective.

Example 1. A seller (she) can provide a single good available with supply $s \in \mathbb{R}_+$ at zero marginal costs. There is one buyer (he) with a continuous valuation $v : \mathbb{R} \to \mathbb{R}$ for the good. The buyer has quasi-linear utility $u(p, x) = v(x) - px$, where p denotes the unit price of the good, and a budget $\beta \in \mathbb{R}_+ \cup \{\infty\}$. We aim to find an allocation $x \leq s$ and a price p that maximise social welfare (SW) or the seller's revenue (R). The buyer demands quantity x at price p if $x \in \arg\max_x v(x) - px$ such that $px \leq \beta$.

The revenue maximisation problem is constrained in the sense that the seller cannot enforce a specific allocation to the buyer. She can only set a price, anticipating the buyer's demand. Thus, the seller's revenue is maximised at $(x, p) \in \arg\max_{x,p} px$ such that $x \in \arg\max_x v(x) - px$ and $x \leq s$. Social welfare, on the other hand, is simply maximised at $(x, p) \in \arg\max_{x,p} v(x) - px$ such that $x \leq s$.

To derive optimal (for revenue or social welfare) allocations and prices, we first need to make some assumptions on the buyer's valuation. In the economics literature, diminishing marginal values are a typical assumption on preferences. In that case, the buyer always demands a quantity such that his marginal utility at this quantity is zero. Social welfare maximisation then implies $v'(x) = p$ and $x = s$; therefore, social-welfare maximising prices are always 'market-clearing'.[6] However, if the seller were allowed to adjust the price subject to respecting the buyer's demand, she might be able to not sell the entire supply and extract more revenue. The following proposition shows that this is indeed the case for all strongly concave value functions if the buyer's budget is large.

Proposition 1. *Let v be differentiable and strongly concave with parameter m for some $m > 0$, and let $\beta = \infty$. Then there exists some supply $s \in \mathbb{R}$ so that revenue is not maximised at the market-clearing price.*

When the buyer has a finite budget, the situation is more intricate. Indeed, supply can only lie in the finite interval $X := [0, \max\{x \mid xv'(x) \leq \beta\}$. We define $\tilde{x}$ implicitly by $\tilde{x}v'(\tilde{x}) = \beta$ and we call any bundle $x \in X$ *budget-feasible.*

Proposition 2. *Suppose the buyer has a strongly concave valuation v with parameter m and a finite budget β. If supply s is contained in X with $v'(s) < ms$, or if $m > \frac{v'(\tilde{x})}{\tilde{x}}$, then revenue is not maximised at the market-clearing price.*

This proposition tells us that for any strongly concave valuation, we can find a combination of budget and supply such that the maximisers of the social welfare and the revenue maximisation problem do not coincide. For example, we may require the valuation to be sufficiently concave relative to the budget. Therefore, we consider the class of constant marginal values, i.e. $v(x) = vx$.

Proposition 3. *Let the buyer's valuation be $v(x) = vx$ with budget β. Then the seller's revenue is maximised at market-clearing prices.*

The above proposition is straightforward. Social welfare is maximised at $p = v$ and $x = \frac{\beta}{v}$. Because the buyer spends his entire budget, the seller cannot extract more revenue, and the revenue and welfare maximising allocation and price coincide.

An immediate question to ask is whether this simple logic extends to more general environments and preferences. The answer we present in this paper is affirmative: if any number of buyers have quasi-linear, budget-constrained utility and additive values for any number of goods, then revenue and social welfare are maximised at a unique set of elementwise-minimal prices. The result, however, is not immediate. In the following, we illustrate the difficulty in another simple example with two goods.

Example 2. Two goods A and B are for sale with a supply of $s_A = 3$ and $s_B = 2$. There are three buyers $1, 2,$ and 3 with the following marginal values: $\boldsymbol{v}^1 = (v^1_A, v^1_B) = (2, 3)$, $\boldsymbol{v}^2 = (v^2_A, v^2_B) = (2, 2)$, and $\boldsymbol{v}^3 = (v^3_A, v^3_B) = (4, 2)$. Utilities are quasi-linear, i.e. $u^i(x, p) = \sum_j (v^i_j - p_j)x_j$ for buyer i. Each buyer has a budget of $\beta^1 = \beta^2 = \beta^3 = 1$. We can allocate the 3 units of A and 2 units of B among the three buyers to maximise either revenue or social welfare, but need to respect individual demand. It is not hard to check that at given prices (p_A, p_B) each buyer i will demand a good $j \in \arg\max_{j=1,2} \frac{v^i_j}{p_j}$, if

[6] The welfare theorems hold even in the general version of our model with multiple goods.

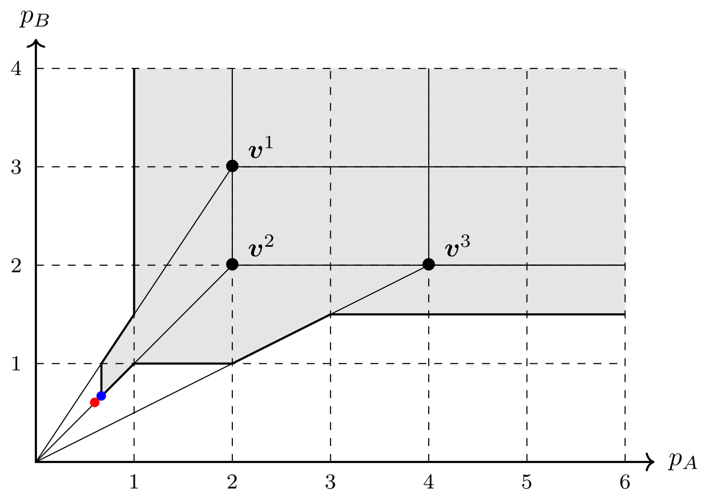

Figure 1: The feasible region in p_A-p_B-space

$v^i_j \geq p_j$. Any good k with $v^i_k < p_k$ will never be demanded by buyer i. This kind of individual demand can be easily represented in price space (more detail in Section 2.1). At some prices, aggregate demand is too large to be satisfied by supply. Prices at which aggregate demand does *not* exceed supply are called *supply-feasible.* The set of supply-feasible prices makes up the *feasible region.* The bids (black dots) and the feasible region (in grey) are illustrated in Fig. 1. Note that the feasible region also includes a short line segment between $p^* := (\frac{3}{5}, \frac{3}{5})$ (red dot) and $p' := (\frac{2}{3}, \frac{2}{3})$ (blue dot).

At prices $(\frac{3}{5}, \frac{3}{5})$, buyer 1 demands $\frac{5}{3}$ of B, buyer 3 demands $\frac{5}{3}$ of A, and buyer 2 demands $x^2_A \in [0, \frac{5}{3}]$ copies of A and $\frac{5}{3} - x^2_A$ of B. With supply $(s_A, s_B) = (3, 2)$, set $x^2_A = \frac{4}{3}$ to clear the market. It is easy to check that indeed any prices on $[p^*, p']$ induce a feasible allocation. All prices $[p^*, p']$ are revenue-maximising. However, only p^* clears the market and constitutes CE prices.

1.3 Preliminaries

For any two n-dimensional vectors $\boldsymbol{v} \in \mathbb{R}^n$ and $\boldsymbol{w} \in \mathbb{R}^n$, we write $\boldsymbol{v} \leq \boldsymbol{w}$ when the inequality holds element-wise. For any $j \in [n]$, $\boldsymbol{e}^j$ denotes the vector whose j-th entry is 1 and other entries are 0. For convenience, we also define $\boldsymbol{e}^0 = \boldsymbol{0}$. The dot product of $\boldsymbol{v}$ and $\boldsymbol{w}$ is denoted $\boldsymbol{v} \cdot \boldsymbol{w}$. For any function $f : A \times B \to \mathbb{R}^n$, we use the implicit summation $f(A', b) = \sum_{a \in A'} f(a, b)$, $f(a, B') = \sum_{b \in B'} f(a, b)$, and $f(A', B') = \sum_{a \in A', b \in B'} f(a, b)$, for any $A' \subseteq A$ and $B' \subseteq B$.

2 Markets, preferences, and objectives

We consider the following market. There are n goods $[n] := \{1, \ldots, n\}$, typically denoted by j and k. We will also work with a notional *null good* 0, and let $[n]_0 = \{0, \ldots, n\}$. A *bundle* of goods, typically denoted by $\boldsymbol{x}$ or $\boldsymbol{y}$ in this paper, is a vector in $\mathbb{R}^n$ whose j-th entry denotes the *quantity* of good j. The seller has a supply bundle $\boldsymbol{s} \in \mathbb{R}^n$ that they wish to sell, partially or completely, by setting uniform market prices $\boldsymbol{p} \in \mathbb{R}^n_+$; a price vector $\boldsymbol{p} \in \mathbb{R}^n$ has a price entry for each of the n goods. When working with the notional null good 0, we implicitly define $p_0 = 0$.

We have m buyers $[m] = \{1, \ldots, m\}$, typically denoted by i, with budgets β^i who compete for the supply of the n distinct goods. Throughout this paper, we assume that every buyer $i \in [m]$ has an *additive valuation* $v^i : \mathbb{R}^n \to \mathbb{R}$ and budget β^i. That is, each buyer i possesses a vector $\boldsymbol{v}^i \in \mathbb{R}^n_+$ so that the valuation is given by $v^i(\boldsymbol{x}) = \boldsymbol{v}^i \cdot \boldsymbol{x}$ for all bundles $\boldsymbol{x} \in \mathbb{R}^n$. Moreover, without loss of generality we can assume that, for every good $j \in [n]$, there exists at least one buyer with a positive valuation v_j for the good. (Otherwise we can simply remove the good from the auction.) Utilities are quasi-linear, i.e. the buyer derives utility $u^i(\boldsymbol{x}, \boldsymbol{p}) = v^i(\boldsymbol{x}) - \boldsymbol{p} \cdot \boldsymbol{x} = (\boldsymbol{v}^i - \boldsymbol{p}) \cdot \boldsymbol{x}$ from receiving bundle $\boldsymbol{x} \in \mathbb{R}^n$ at prices $\boldsymbol{p} \in \mathbb{R}^n$. At any given prices $\boldsymbol{p}$, a buyer i demands the bundles $\boldsymbol{x}$ that maximise their utility $u^i(\cdot, \boldsymbol{p})$,

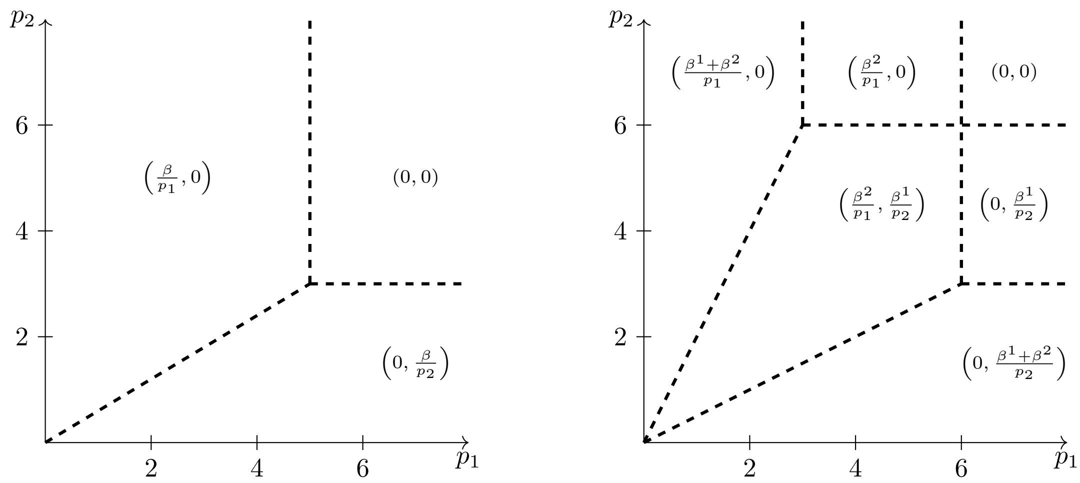

Figure 2: Illustrations of (aggregate) demand correspondences for buyers with quasi-linear utilities, additive valuations and budget constraints. Left: Demand for a single buyer with additive valuation $\boldsymbol{v} = (5,3)$ and budget constraint β divides price-space into three regions. At prices that are element-wise larger than $\boldsymbol{v}$, the buyer demands nothing. At prices $\boldsymbol{p}$ that lie above the line connecting point $\boldsymbol{v}$ to the origin, the buyer demands a quantity β/p_1 of good 1. Conversely, at prices below the line, the buyer demands β/p_2 of good 2. Right: The aggregate demand correspondence of two buyers: the first has valuation $\boldsymbol{v}^1 = (3,6)$ and budget β^1, and the second has valuation $\boldsymbol{v}^2 = (6,3)$ and budget β^2.

subject to not exceeding their budget (expressed by the budget constraint $\boldsymbol{x} \cdot \boldsymbol{p} \leq \beta^i$). This leads to the *demand correspondence* $D^i(\boldsymbol{p}) = \max_{\boldsymbol{x}\in\mathbb{R}^n, \boldsymbol{x}\cdot\boldsymbol{p}\leq\beta^i}(v^i(\boldsymbol{x}) - \boldsymbol{p}\cdot\boldsymbol{x}) = \max_{\boldsymbol{x}\in\mathbb{R}^n, \boldsymbol{x}\cdot\boldsymbol{p}\leq\beta^i}(\boldsymbol{v}^i - \boldsymbol{p})\cdot\boldsymbol{x}$.

As with prices, we define $v_0^i = 1$ for the value of the null good. Moreover, for any buyer i and prices $\boldsymbol{p}$, we define the *demanded good set* $J^i(\boldsymbol{p}) = \arg\max_{j\in[n]_0} \frac{v_j^i}{p_j}$. This set of goods is also referred to in the literature (on Fisher markets) as maximising *bang-per-buck*. (Note that $J^i(\boldsymbol{p})$ contains the null good if $\max_{j\in[n]_0} \frac{v_j^i}{p_j} = 1$, as we have $v_0^i = 1$ and $p_0 = 1$ by definition.) This allows us to provide an alternative characterisation of the demand correspondence that we rely on when proving our results below. In particular, Lemma 4 makes the initial observation that bundles only contain quantities of goods that maximise bang-per-buck; in other words, if $\boldsymbol{x}$ is a bundle demanded by buyer i at $\boldsymbol{p}$, then $x_j > 0$ implies $j \in J^i(\boldsymbol{p})$. Moreover, any demanded bundle is the convex combination of the 'extremal' bundles that arise when the entire budget is spent on a single demanded good in $J^i(\boldsymbol{p})$. In Section 2.1, we will see that the arctic product-mix auction introduces a bidding language which starts from this definition of demand to characterise a more general demand type.

Lemma 4. *Fix $\boldsymbol{p} \in \mathbb{R}^n_+$. Let $\boldsymbol{x} \in D(\boldsymbol{p})$ for some buyer with additive valuation $\boldsymbol{v}$, budget β and demanded good set $J(\boldsymbol{p})$. If $x_j > 0$, then $j \in J(\boldsymbol{p})$. Moreover $D(\boldsymbol{p}) = \operatorname{conv}\left\{\frac{\beta}{p_j}\boldsymbol{e}^j \mid j \in J(\boldsymbol{p})\right\}$.*

Proof. Suppose, for contradiction, we have a non-null good j with $x_j > 0$ and $j \notin J(\boldsymbol{p})$. Then there exists another good $k \in J(\boldsymbol{p})$ with $\frac{v_j}{p_j} < \frac{v_k}{p_k}$. If k is the null good, then we define a new bundle $\boldsymbol{x}'$ from $\boldsymbol{x}$ by not spending on good j, so $\boldsymbol{x}' = \boldsymbol{x} - x_j\boldsymbol{e}^j$. It is immediate that $\boldsymbol{x}'$ is budget-feasible in the sense that it does not exceed the buyer's budget β at $\boldsymbol{p}$, as $\boldsymbol{x}'\cdot\boldsymbol{p} < \boldsymbol{x}\cdot\boldsymbol{p} \leq \beta$. Moreover, its utility is $(\boldsymbol{v}-\boldsymbol{p})\cdot\boldsymbol{x}' = (\boldsymbol{v}-\boldsymbol{p})\cdot\boldsymbol{x} + x_j(p_j - v_j) > (\boldsymbol{v}-\boldsymbol{p})\cdot\boldsymbol{x}$. The last inequality follows from our assumption that $\frac{v_j}{p_j} < \frac{v_0}{p_0} = 1$, and it contradicts our assumption that $\boldsymbol{x}$ maximises utility. Now suppose k is not the null good. Then we let $\boldsymbol{x}'$ be the bundle obtained from $\boldsymbol{x}$ by spending nothing on good j and instead using the same amount of money to buy (additional) quantities of good k, so $\boldsymbol{x}' := \boldsymbol{x} + x_j\left(\frac{p_j}{p_k}\boldsymbol{e}^k - \boldsymbol{e}^j\right) \geq 0$. Then $\boldsymbol{x}'$ does not exceed the buyer's budget, as we can verify $\boldsymbol{p}\cdot\boldsymbol{x}' = \boldsymbol{p}\cdot\boldsymbol{x} \leq \beta$. Moreover, its utility is $(\boldsymbol{v}-\boldsymbol{p})\cdot\boldsymbol{x}' = (\boldsymbol{v}-\boldsymbol{p})\cdot\boldsymbol{x} + x_j\left(\frac{p_j}{p_k}v_k - v_j\right) = (\boldsymbol{v}-\boldsymbol{p})\cdot\boldsymbol{x} + x_jp_j\left(\frac{v_k}{p_k} - \frac{v_j}{p_j}\right) > (\boldsymbol{v}-\boldsymbol{p})\cdot\boldsymbol{x}$.

Next we consider the second claim $D(\boldsymbol{p}) = \operatorname{conv}\left\{\frac{\beta}{p_j}\boldsymbol{e}^j \mid j \in J(\boldsymbol{p})\right\}$. Fix $\boldsymbol{x} \in D(\boldsymbol{p})$, and let β_j denote, for each good $j \in [n]_0$, the amount that the buyer spends on j to obtain x_j. If $0 \in J(\boldsymbol{p})$, then we let

$\beta_0 = \beta - \beta_{[n]}$ be the remainder of the budget that the buyer doesn't spend. Hence we have $\beta_{J(\boldsymbol{p})} = \beta$ and $\beta_j = 0$ for $j \notin J(\boldsymbol{p})$. Moreover, $x_j = \frac{\beta_j}{p_j}$ for all $j \in [n]$, and thus $\boldsymbol{x} = \sum_{j \in J(\boldsymbol{p})} \frac{\beta_j}{\beta} \frac{\beta}{p_j} \boldsymbol{e}^j$, so $\boldsymbol{x}$ lies in the convex hull. Now we show that all bundles in the convex hull lie in $D(\boldsymbol{p})$. Note that reallocating spending β_j on good j to another good $k \in J(\boldsymbol{p})$ leads to another bundle, the bundle $\boldsymbol{x}'$ as defined above, that does not exceed the buyer's budget and maximises utility. The latter follows as $\frac{v_k}{p_k} = \frac{v_j}{p_j}$ for $j, k \in J(\boldsymbol{p})$. Repeatedly reallocating spending until the bundle only contains positive quantities of a single good $j \in J(\boldsymbol{p})$ shows that the vertices of the convex hull are in $D(\boldsymbol{p})$. Finally, note that all bundles in the convex hull are convex combinations of the vertices, and taking convex combinations retains both budget-feasibility and maximal utility. $\square$

Expenditure and allocation. When solving the auction, the seller sets market prices and determines an allocation to each buyer. Note that the allocation and prices directly control how much each buyer spends on each good. We capture this expenditure formally, as follows. An expenditure function $e : [m] \times [n]_0 \to \mathbb{R}_+$ expresses, for every buyer $i \in [m]$ and every good $j \in [n]_0$, the amount $e(i, j)$ that buyer i spends on good j. For each expenditure, we can define the corresponding *allocation* $\pi : [m] \times [n]_0 \to \mathbb{R}_+$ denoting the quantity $\pi(i, j) = \frac{e(i,j)}{p_j}$ of good j that buyer i receives. Moreover, for notational convenience we can define vectors $e(i)$ and $\pi(i)$ for each buyer i as $e(i)_j = e(i, j)$ and $\pi(i)_j = e(i, j)$ for all goods $j \in [n]$. The vector $e(i)$ captures how much buyer i spends on each good, while $\pi(i)$ denotes the bundle of goods that i receives. In particular, implicit summation allows us to express aggregate spending and demand as $e([m])$ and $\pi([m])$.

Given an expenditure e, we say that good $j \in [n]$ is *sated* if $e([m], j) = p_j s_j$, and *unsated* if $e([m], j) < p_j s_j$. If the seller wishes not to exceed supply and to allocate each buyer a bundle they demand at the chosen market prices, she needs to satisfy three criteria that are expressed in Definition 1.

Definition 1 (Feasibility and validity)**.** Let $e : [m] \times [n]_0 \to \mathbb{R}_+$ be an expenditure. Then e is:

(i) *Supply-feasible* at prices $\boldsymbol{p}$ if supply weakly exceeds aggregate demand for all goods j, so $e([m], j)/p_j \leq s_j$ for all $j \in [n]$.

(ii) *Budget-feasible* at prices $\boldsymbol{p}$ if every buyer i spends weakly less than its budget on non-null goods. The remaining (non-negative) amount is always spent on the null good, so $e(i, [n]_0) = \beta^i$ for all buyers $i \in [m]$.

(iii) *Demand-valid* at prices $\boldsymbol{p}$ if each buyer only spends on goods that maximise her bang-per-buck, so $e(i, j) > 0$ implies $j \in I^i(\boldsymbol{p})$.

An expenditure is *feasible* if (i) and (ii) hold, and it is *valid* if (iii) also holds. We extend the same terminology to the allocations π corresponding to feasible or valid expenditures defined by $\pi(i, j) = \frac{e(i,j)}{p_j}$. Note that for any valid allocation π at prices $\boldsymbol{p}$, the bundle of goods that buyer i is allocated under π lies in $D^i(\boldsymbol{p})$.

Definition 2. Prices $\boldsymbol{p}$ are *feasible* if there exists an expenditure e valid at $\boldsymbol{p}$.

In Section 3.1, we show that the set of feasible prices forms a lower semi-lattice, and in particular has an *elementwise-minimal price vector* that is dominated by all other feasible prices.

Market objectives. We consider two separate objectives for the seller. She may set prices with the objective of maximising her revenue, or she may set prices to maximise social welfare. Intuitively (and we prove this below), at the social optimum, the seller's entire supply is allocated. To maximise revenue, the seller may prefer to raise prices and retain some of her supply.

1. **Envy-free revenue maximisation.** Given a valid allocation π at prices $\boldsymbol{p}$, the revenue for the seller is given by $\boldsymbol{p} \cdot \pi([m]) = e([m])$. Hence, we define the indirect revenue function as $R(\boldsymbol{p}) := \max_\pi \boldsymbol{p} \cdot \pi([m])$ subject to π being valid. A solution to the *envy-free revenue maximisation problem* consists of a price vector $\boldsymbol{p}$ and allocation π that maximises $R(\boldsymbol{p})$.

2. **Welfare maximisation.** Recall that the buyers are price takers. Thus, social welfare is defined with respect to the submitted values. A solution to the *social welfare maximisation problem* is a price vector and allocation $(\boldsymbol{p}, \pi)$ that maximises social welfare $\sum_{i \in [m]} \boldsymbol{v}^i \cdot \pi(i)$ subject to π being a valid allocation.

3. **Competitive equilibrium.** Given a list of buyers $[m]$ and a supply vector $\boldsymbol{s}$, prices $\boldsymbol{p}$ and allocation π (at $\boldsymbol{p}$) form a competitive equilibrium if, and only if, π is valid at $\boldsymbol{p}$ and $\pi([m], j) = s_j$ for all goods j with $p_j > 0$.

It is well-known that the welfare theorems hold: competitive equilibrium maximises welfare, and any efficient allocation can be supported by prices to constitute a competitive equilibrium. We provide direct proofs of this in the appendix. For this reason, we focus on the two objectives of envy-free revenue maximisation and competitive equilibrium.

2.1 Arctic auctions

The market we consider can be interpreted as an important special case of the arctic product-mix auction market first proposed by Paul Klemperer [23] for the government of Iceland.[7] In this market, the seller has a given supply $\boldsymbol{s} \in \mathbb{R}^n_+$ of multiple divisible goods and wishes to find a valid allocation of some subset of this supply among a finite set of buyers with the goal of maximising revenue. The buyers express their demand preferences by submitting a collection of 'arctic bids', each of which consisting of an n-dimensional vector $\boldsymbol{b} \in \mathbb{R}^n_+$ and a monetary budget $\beta(\boldsymbol{b})$. At a market price *above the stated bid price*, an arctic bid *rejects the good*. At market prices below the stated bid prices, the bid spends its budget on goods that yield the highest 'bang-per-buck', i.e. the highest ratio of value to price $\frac{b_j}{p_j}$. When multiple goods maximise 'bang-per-buck', the bid can arbitrarily divide its expenditure between these goods. Moreover, if the maximal bang-per-buck is 1, then the bid may choose not to spend part of its budget (which we interpret as spending on the null good). Hence, Lemma 4 implies that each arctic bid $\boldsymbol{b}$, interpreted in isolation, induces a quasi-linear demand with additive valuation $v(\boldsymbol{x}) = \boldsymbol{b} \cdot \boldsymbol{x}$ and budget $\beta(\boldsymbol{b})$ as defined in Section 2. Moreover, if we denote the demand correspondence of bid $\boldsymbol{b}$ by $D_{\boldsymbol{b}}$, then the demand correspondence of a collection $[m]$ of bids is defined by the Minkowski sum of demands $D_{\mathcal{B}}(\boldsymbol{p}) = \left\{ \sum_{\boldsymbol{b} \in [m]} \boldsymbol{x}^{\boldsymbol{b}} \mid \boldsymbol{x}^{\boldsymbol{b}} \in D_{\boldsymbol{b}}(\boldsymbol{p}) \right\}$. Equivalently, $D_{\mathcal{B}}$ can be understood as the aggregate demand of $|\mathcal{B}|$ buyers with quasi-linear demand, additive valuations $\boldsymbol{b} \in \mathcal{B}$ and budgets $\beta(\boldsymbol{b})$. By submitting multiple arctic bids, buyers can express richer preferences. Nevertheless, it is straightforward to see from the definition of $D_{\mathcal{B}}$ that for the purposes of solving the auction (for welfare or for revenue), the seller can treat each bid independently, and we can assume without loss of generality that each buyer submits a single bid. Hence our main result on the coincidence of welfare and envy-free revenue holds also for the special case of the arctic auction we describe.

3 Main result

Theorem 5. *If buyers have quasi-linear utility and additive valuations, revenue and welfare are maximised at the unique elementwise-minimal feasible price vector $\boldsymbol{p}^*$.*

To prove this in our general setting, we proceed as follows:

1. A key property of the feasible region is that it forms a lower semi-lattice: that is, there exists an elementwise-minimal price vector that is dominated by all other feasible prices. We show this by first proving that, for any two feasible prices, the element-wise minimum of these prices also belongs to the region.

2. We prove that revenue is maximised at elementwise-minimal feasible prices. Note that revenue may also be maximised at other prices.

[7] The arctic product-mix auction is a variant of the original product-mix auction developed for the Bank of England by Paul Klemperer [21, 22]. See also [4]. Klemperer also chose the name 'arctic' to contrast with the fact that the original (but not the arctic) product-mix auction has close connections to tropical geometry. In the general arctic product-mix auction, the seller can additionally choose cost functions to fine-tune its preferences, while in our model, we assume the seller's costs to be zero. See [17] for some discussion of the general case.

3. It is known that market-clearing prices exist and are unique in our market [8]. We also provide a direct proof of this in the appendix.

4. We provide a procedure that, given prices and a valid allocation with unsated goods, reduces prices and finds a new valid allocation. Our procedure is an adaptation of Adsul et al. [1]'s price-scaling procedure designed for Fisher markets. Our procedure implies that markets must clear at elementwise-minimal feasible prices.

3.1 Elementwise-minimal prices

Recall that prices $\boldsymbol{p}$ are feasible if there exists a valid allocation π at $\boldsymbol{p}$. We show that set of feasible prices form a lower semi-lattice. In particular, there exists a special price vector $\boldsymbol{p}^*$ that is element-wise smaller than all other feasible prices (so that $\boldsymbol{p}^* \leq \boldsymbol{p}$ for all feasible $\boldsymbol{p}$). For any two price vectors $\boldsymbol{p}$ and $\boldsymbol{q}$, we let $\boldsymbol{p} \wedge \boldsymbol{q}$ denote their *element-wise minimum* defined as $(\boldsymbol{p} \wedge \boldsymbol{q})_i = \min\{p_i, q_i\}$. The following lemma is central to our proof.

Lemma 6. *If $\boldsymbol{p}$ and $\boldsymbol{p}'$ are feasible, then so is their element-wise minimum $\boldsymbol{p} \wedge \boldsymbol{q}$.*

Fix feasible prices $\boldsymbol{p}$ and $\boldsymbol{p}'$ with element-wise minimum $\boldsymbol{r} = \boldsymbol{p} \wedge \boldsymbol{p}'$, and let π and π' respectively denote valid allocations at $\boldsymbol{p}$ and $\boldsymbol{p}'$. In order to prove Lemma 6, we construct an allocation τ at $\boldsymbol{r}$ and show that it is valid. We first define the set of goods A in which $\boldsymbol{p}$ is strictly dominated by $\boldsymbol{p}'$, and its complement B, so $A = \{j \in [n] \mid p_j < p'_j\}$ and $B = \{j \in [n] \mid p_j \geq p'_j\}$. Then our allocation τ at $\boldsymbol{r}$ is given by

$$\tau(i, \cdot) = \begin{cases} \pi'(i, \cdot) & \text{if buyer } i \text{ demands some good } j \in [n] \text{ at } \boldsymbol{r}, \\ \pi(i, \cdot) & \text{otherwise.} \end{cases} \tag{1}$$

In order to prove the validity of τ, we first state a technical lemma that establishes the connection between a buyer's demand at $\boldsymbol{p}$, $\boldsymbol{p}'$ and $\boldsymbol{r}$.

Lemma 7. *Suppose buyer i demands good $j \in A$ at $\boldsymbol{r}$. Then they also demands j at $\boldsymbol{p}$ and $J^i(\boldsymbol{p}) \subseteq J^i(\boldsymbol{r})$. Similarly, suppose buyer i demands $j \in B$ at $\boldsymbol{r}$. Then they also demand j at $\boldsymbol{p}'$, and $J^i(\boldsymbol{p}') \subseteq J^i(\boldsymbol{r})$. Moreover, we have $J^i(\boldsymbol{p}') \subseteq B$.*

Proof. Fix a buyer i who demands good $j \in A$ at $\boldsymbol{r}$. As $p_j = r_j$, this implies $\frac{v^i_j}{p_j} = \frac{v^i_j}{r_j} \geq \frac{v^i_j}{r_j} \geq \frac{v^i_j}{p_j}$ for all goods $j \in [n]_0$. The first inequality holds due to the definition of demand, and the second inequality follows from $r_j \leq p_j, \forall j \in [n]_0$. Hence, the buyer demands good j at $\boldsymbol{p}$. For the second claim that $J^i(\boldsymbol{p}) \subseteq J^i(\boldsymbol{r})$, fix a good $k \in J^i(\boldsymbol{p})$. Then we have $\frac{v^i_k}{r_k} \geq \frac{v^i_k}{p_k} \geq \frac{v^i_j}{p_j} = \frac{v^i_j}{r_j} \geq \frac{v^i_l}{r_l}$ for all goods $l \in [n]_0$. The first inequality holds due to $r_k \leq p_k$, and the second and third inequalities follow from the fact that i demands k at $\boldsymbol{p}$ and j at $\boldsymbol{r}$. Hence, if the buyer demands good k at $\boldsymbol{p}$, then they demand k at $\boldsymbol{r}$.

Now suppose that the buyer demands $j \in B$ at $\boldsymbol{r}$. The proof of the first claim is identical to the case $j \in A$. We prove the last claim that $J^i(\boldsymbol{p}') \subseteq B$. Suppose, for contradiction, that i demands a good $k \in A$ at $\boldsymbol{p}'$, and good $j \in B$ at $\boldsymbol{r}$. This implies $\frac{v^i_k}{p'_k} < \frac{v^i_k}{p_k} = \frac{v^i_k}{r_k} \leq \frac{v^i_j}{r_j} = \frac{v^i_j}{p'_j}$, in contradiction to the fact that k is demanded at $\boldsymbol{p}'$. □

We can now prove Lemma 6.

Proof of Lemma 6. Let $\boldsymbol{p}$ and $\boldsymbol{p}'$ be two feasible prices and $\boldsymbol{r} = \boldsymbol{p} \wedge \boldsymbol{p}'$ denote their element-wise minimum. As above, π and π' are valid allocations at $\boldsymbol{p}$ and $\boldsymbol{p}'$, and τ is defined as in (1). First we see that τ is indeed an allocation, as $\tau(i, j) \geq 0$ for all buyers i and goods j. It remains to prove that τ satisfies the three criteria of Definition 1 that define validity. We can partition buyers into two sets: the set $\mathcal{B} \subseteq [m]$ of buyers that demand a good in B at $\boldsymbol{r}$, and the set $\mathcal{A} = [m] \setminus \mathcal{B}$ of buyers that do not. Note that, by Lemma 7, the buyers in $\mathcal{B}$ demand only goods in B at $\boldsymbol{p}'$ and, by definition, the buyers in $\mathcal{A}$ demand only goods in A at $\boldsymbol{p}$. It follows that, under π, each buyer $i \in \mathcal{A}$ is only allocated quantities of goods in A (so $\pi(\mathcal{A}, j) > 0$ only if $j \in A$), and, similarly, every buyer $i \in \mathcal{B}$ only receives quantities of goods in B under π' (so $\pi(\mathcal{B}, j) > 0$ only if $j \in B$).

First we show that τ is supply-feasible at $\boldsymbol{r}$. Recall that τ is supply-feasible if $\tau([m], j) \leq s_j$ for all $j \in [n]$. Fix some good $j \in A$. Then by definition of τ, we have $\tau([m], j) = \pi(\mathcal{A}, j) + \pi'(\mathcal{B}, j)$. Recalling that $\pi'(\mathcal{B}, j) = 0$ and making use of the validity of π at $\boldsymbol{p}$, we get $\tau([m], j) \leq \pi([m], j) \leq s_j$. This immediately implies that τ satisfies part (i) of Definition 1 at $\boldsymbol{r}$ for all goods $j \in A$. Analogously, we can show that τ satisfies (i) at $\boldsymbol{r}$ for all goods $j \in B$ by recalling that $\pi(\mathcal{A}, j) = 0$ for any good $j \in B$. As $A \cup B = [n]$, we have shown that τ is supply-feasible.

Next we argue that τ is budget-feasible at $\boldsymbol{r}$. Indeed, note that any buyer in $\mathcal{A}$ only demands goods in A. Moreover, the prices of goods in A are the same at $\boldsymbol{p}$ and $\boldsymbol{r}$, by construction of $\boldsymbol{r}$. It follows that the buyer spends the same under τ at $\boldsymbol{r}$ as they do under π at $\boldsymbol{p}$. As π is budget-feasible at $\boldsymbol{p}$, we see that τ satisfies part (ii) of Definition 1 for all buyers in $\mathcal{A}$. Similarly, as the buyers in $\mathcal{B}$ only demand goods in B, we apply the same argument to see that τ satisfies part (ii) for all buyers in $\mathcal{B}$. As $[m] = \mathcal{A} \cup \mathcal{B}$, τ is budget-feasible.

Finally, we show the demand-validity of τ at $\boldsymbol{r}$. Lemma 7 and the validity of π at $\boldsymbol{p}$ together imply that, for any $j \in [n]$, we have $\tau(i, j) = \pi(i, j) > 0$ only if $j \in J^i(\boldsymbol{p}) \subseteq J^i(\boldsymbol{r})$, which implies that τ satisfies part (iii) of Definition 1 at $\boldsymbol{r}$ for all buyers in $\mathcal{A}$. Analogously, as $\tau(i, j) = \pi'(i, j)$ for all buyers $i \in \mathcal{B}$, we see that τ satisfies constraint (iii) at $\boldsymbol{r}$ for all buyers in $\mathcal{B}$. Hence, τ is demand-valid at $\boldsymbol{r}$, and thus valid. $\square$

Corollary 8. *There exists an elementwise-minimal price vector $\boldsymbol{p}^*$.*

Proof. Suppose there exists no such price vector. This means that for all feasible $\boldsymbol{p}$, there exists some feasible $\boldsymbol{q}$ with $q_j < p_j$ for at least one good $j \in [n]$. Fix some feasible prices $\boldsymbol{p}$ with the property that $\boldsymbol{p}$ cannot be reduced any further in any direction without breaking feasibility. Such a point must exist, as the feasible region is closed and restricted to $\mathbb{R}^n_+$. By assumption, there exists a feasible price vector $\boldsymbol{q}$ with $q_j < p_j$ for some $j \in [n]$. Now consider $\boldsymbol{p}' = \boldsymbol{p} \wedge \boldsymbol{q}$. By Lemma 6, $\boldsymbol{p}'$ is feasible. But as $\boldsymbol{p}' \leq \boldsymbol{p}$ with $p'_j < p_j$, this contradicts our assumption that $\boldsymbol{p}$ cannot be reduced further. $\square$

3.2 Maximising revenue

In Section 3.1, we established that the set of feasible prices contain a unique elementwise-minimal price vector $\boldsymbol{p}^*$. We now show that revenue is maximised at these prices. Note that we do not assume $\boldsymbol{p}^*$ to be the *only* prices at which revenue is maximised; indeed, there can be many revenue-maximising prices. However, $\boldsymbol{p}^*$ maximises the quantities of goods allocated.

Lemma 9. *For any two distinct feasible price vectors $\boldsymbol{p}$ and $\boldsymbol{p}'$ with $\boldsymbol{p} \leq \boldsymbol{p}'$ we have $R(\boldsymbol{p}) \geq R(\boldsymbol{p}')$. In other words, the maximum obtainable revenue at $\boldsymbol{p}$ is weakly greater than the revenue obtainable at $\boldsymbol{p}'$.*

Proof. Let π and π' be valid allocations that respectively maximise revenue at $\boldsymbol{p}$ and $\boldsymbol{p}'$. Our goal is to determine a valid allocation τ that achieves a weakly greater revenue at $\boldsymbol{p}$ than π' does at $\boldsymbol{p}'$. As $R(\boldsymbol{p}) \geq \sum_{j \in [n]} p_j \tau([m], j)$ and $R(\boldsymbol{p}') = \sum_{j \in [n]} p_j \pi'([m], j)$, this immediately implies the result.

If $\boldsymbol{p} = \boldsymbol{p}'$, there is nothing to prove. Hence we assume that $S := \{j \in [n] \mid p_j < p'_j\}$, the set of goods which are priced strictly lower at $\boldsymbol{p}$ than at $\boldsymbol{p}'$, is non-empty. Fix a buyer $i \in [m]$. In order to define the new allocation $\tau(i, \cdot)$ to i at $\boldsymbol{p}$, we distinguish between the two cases that $J^i(\boldsymbol{p})$ is, and is not, a subset of S.

Case 1: Suppose buyer i demands a subset of S at $\boldsymbol{p}$, so $J^i(\boldsymbol{p}) \subseteq S$. In this case, we set $\tau(i, \cdot) = \pi(i, \cdot)$. As π is valid, and τ and π are both allocations at the same prices, the buyer spends its entire budget under τ. Moreover, they are only allocated goods that they demand.

Case 2: Suppose $J^i(\boldsymbol{p}) \nsubseteq S$. We note that $J^i(\boldsymbol{p}') \cap S = \emptyset$ and buyer i still demands all goods in $J^i(\boldsymbol{p}')$, so $J^i(\boldsymbol{p}') \subseteq J^i(\boldsymbol{p})$. In this case, we set $\tau(i, \cdot) = \pi'(i, \cdot)$. As the buyer is only allocated goods not in S, and $p_j = p'_j$ for all goods $j \in [n]_0 \setminus S$, it follows that the buyer spends the same at both prices.

Note that, in both cases, the buyer is only allocated goods that they demand. To summarise, we define τ as

$$\tau(i, \cdot) = \begin{cases} \pi(i, \cdot) & \text{if buyer } i \text{ demands a subset of } S \text{ at } \boldsymbol{p}, \\ \pi'(i, \cdot) & \text{otherwise.} \end{cases}$$

We now prove that τ is valid. We have already argued above that all buyers satisfy the demand and budget conditions of Definition 1. It remains to show that aggregate demand does not exceed supply

s_j for any goods $j \in [n]_0$ under τ. Note that a buyer is allocated a good j in S under τ if, and only if, they satisfy Case 1 above. Indeed, in this case we set $\tau(i,j) = \pi(i,j)$. Hence, for any $j \in S$, we have $\tau([m],j) \leq \pi([m],j) \leq s_j$. Similarly, for any $j \notin S$ the buyer will satisfy Case 2, and we get $\tau([m],j) \leq \pi'([m],j) \leq s_j$.

Finally, we see that τ achieves weakly greater revenue at $\boldsymbol{p}$ than π' does at $\boldsymbol{p}'$. To see this, note that each buyer satisfying Case 1 spends its entire budget on non-reject goods and thus contributes a weakly greater amount to revenue at $\boldsymbol{p}$ than at $\boldsymbol{p}'$, while a buyer satisfying Case 2 contributes the same amount to revenue. □

Corollary 10. *Revenue is maximised at elementwise-minimal feasible prices* $\boldsymbol{p}^*$*.*

Proof. Suppose $\boldsymbol{p} \neq \boldsymbol{p}^*$ is a revenue-maximising price vector. As $\boldsymbol{p}^*$ is elementwise-minimal, we have $\boldsymbol{p}^* \leq \boldsymbol{p}$. Then by Lemma 9, we can obtain weakly greater revenue at $\boldsymbol{p}^*$ than at $\boldsymbol{p}$, or $R(\boldsymbol{p}^*) \geq R(\boldsymbol{p})$. But as $\boldsymbol{p}$ maximises R, so does $\boldsymbol{p}^*$. □

4 Reducing prices

Suppose $\boldsymbol{p}$ is a feasible price vector at which no allocation exhausts supply $\boldsymbol{s}$ in all goods. In this case, we present a method to scale a subset of goods uniformly in price by a factor $0 < c < 1$ while retaining feasibility and increasing aggregate demand. As stated in Corollary 11, this implies that the elementwise-minimal prices $\boldsymbol{p}^*$ clear the market. We do this by modifying a price-scaling routine that forms part of a 'simplex-like' algorithm by Adsul et al. [1] for Fisher markets to our market setting; this is not immediate, as buyers can choose not to spend their entire budget if prices are too high. In our approach, we scale prices down instead of up, and apply additional pre-processing for bids that may indifferent between spending on the null good and 'real' goods.

4.1 The procedure for reducing prices

In the following, we describe our subroutine that, given prices $\boldsymbol{p}$ that are not market-clearing and an expenditure e, either modifies e in order to ensure that an unsated good becomes sated, or reduces prices of some goods. Moreover, we claim that after $O(n)$ repetitions of this subroutine, prices of some goods are reduced. Indeed, we start with at most n unsated goods. Each application of the subroutine that fails to reduce prices strictly reduces the number of unsated goods, as no sated goods can revert to being unsated. In particular, this implies that the elementwise-minimal prices $\boldsymbol{p}^*$ clear the market, as we could otherwise reduce $\boldsymbol{p}^*$ in price while maintaining feasibility.

Corollary 11. *Competitive equilibrium is uniquely achieved at elementwise-minimal prices* $\boldsymbol{p}^*$*.*

Proof. It is well-known that market-clearing prices are unique in our market. (Theorem 14 in the appendix provides a direct proof of this.) Suppose that $\boldsymbol{p}^*$ is not market-clearing. Then by Theorem 12, and our discussion above, we can reduce prices of some goods while maintaining feasibility. But this contradicts our assumption that $\boldsymbol{p}^*$ is the elementwise-minimal feasible price vector. □

In order to describe our method of scaling prices, we first introduce the demand and expenditure graphs. Then we present a subroutine that – given prices $\boldsymbol{p}$, an expenditure e and an unsated good j under e at $\boldsymbol{p}$ – either returns scaled-down prices or an expenditure at $\boldsymbol{p}$ under which j is sated. After invoking this subroutine $O(n)$ times, prices will be scaled down or all goods are sated. We note that the subroutine defined here is also called in our algorithm (cf. Section 5) to find the unique market-clearing prices $\boldsymbol{p}^*$.

4.2 The demand and expenditure graphs

In order to describe demand and spending relationships between bids and goods, we introduce two bipartite graphs, the *demand graph* and the *expenditure graph*. Modulating expenditures in appropriately defined sub-trees of these graphs is the central component in our price reduction procedure.

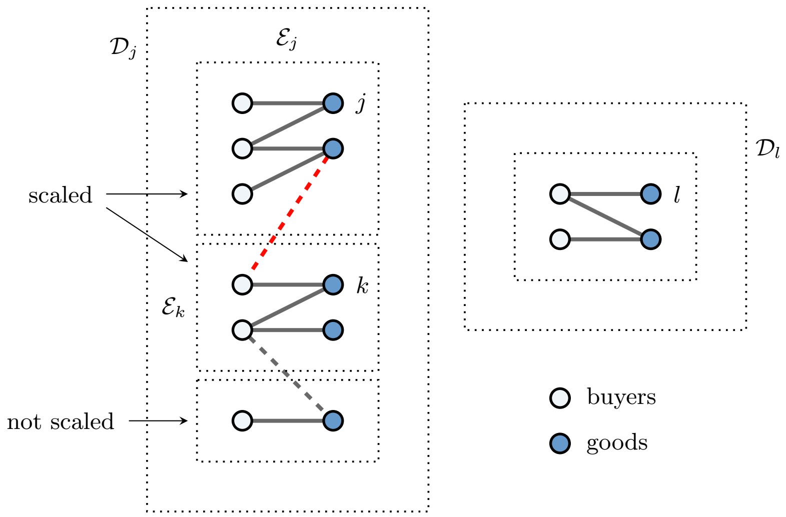

Figure 3: An illustration of the demand and expenditure graphs. Solid and dashed lines respectively indicate edges in $\mathcal{E}$ and $\mathcal{D} \setminus \mathcal{E}$. The component $\mathcal{D}_i$ contains three components of $\mathcal{E}$. If we root $\mathcal{D}_j$ at j, we note that there is an arc (dashed red) from a good in $\mathcal{E}_j$ to a buyer in $\mathcal{E}_k$, and an arc (dashed grey) from a buyer in $\mathcal{E}_k$ to the component below. Hence the tree constructed in Section 4.3, which contains all goods to be scaled down in price, consists of $\mathcal{E}_j$ and $\mathcal{E}_k$ but not the third component.

Definition 3 (Demand graph and expenditure graph)**.** Let e be a valid expenditure at feasible prices $\boldsymbol{p}$. We denote by $\mathcal{D}$ the bipartite *demand graph* on vertex sets $[m]$ and $[n]_0$. There is an edge $(i, j) \in D$ if, and only if, buyer i demands j at $\boldsymbol{p}$. Moreover, we define the *expenditure graph* $\mathcal{E}$ of $\mathcal{D}$ induced by the edges (i, j) with positive expenditure $e(i, j) > 0$. Finally, $\mathcal{D}_j$ and $\mathcal{E}_j$ respectively denote the connected component of $\mathcal{D}$ and $\mathcal{E}$ that contains good j.

Note that $\mathcal{E}_j$ is a subgraph of $\mathcal{D}_j$ for all goods $j \in [n]_0$, and each $\mathcal{D}_j$ can contain multiple connected components of $\mathcal{E}$. Both graphs are illustrated in Fig. 3.

Breaking cycles. Our price reduction method assumes that the demand graph is acyclic. For this reason, we follow Orlin [26] in perturbing buyer valuations by an infinitesimal amount in order to break cycles. Induce an ordering $[m] = \{i^1, i^2, \dots\}$ on the buyers, and perturb each entry $v_j^{i^k}$ of buyer i^k by adding $\varepsilon^{kn} + \varepsilon^j$ for some infinitesimally small $\varepsilon > 0$. In our algorithms, this perturbation can be simulated without a running-time penalty by implementing lexicographic tie-breaking when constructing the demand and expenditure graphs. For details about the perturbation, we refer to [26].

4.3 Reducing prices along a tree

We now describe the subroutine referred to above. The full procedure is stated in Algorithm 1. Correctness is proved in Theorem 12.

Identifying goods to reduce in price. Suppose j is an unsated good under expenditure e at $\boldsymbol{p}$. We describe a procedure which either identifies a set $J \subseteq [n]$ of goods to reduce in price, or finds a new expenditure e' at $\boldsymbol{p}$ under which good j is sated. We first construct a directed tree T by rooting the (undirected) tree $\mathcal{D}_j$ at good j and then remove all subtrees rooted at endpoint goods j of arcs $(i, j) \in \mathcal{D}_j \setminus \mathcal{E}$. This is illustrated in Fig. 3. For notational convenience, we also let T_v denote the subtree of T rooted at v for any vertex $v \in T$.

If T does not contain the null good 0, we let J consist of the goods of T and are done. Otherwise, we consider the directed path $(j = g^1, b^1, \dots, g^k, b^k, g^{k+1} = 0)$ from good j to null good 0 in T and modify the expenditure e to redirect spending from 0 to j as follows. Let $m > 0$ be the minimum of $p_j s_j - e([m], j)$ and $\min_{i \in [k]} e(b^i, g^i)$. We let $e'(b^i, g^i) = e(b^i, g^i) + m$ and $e'(b^i, g^{i+1}) = e(b^j, g^i) - m$ for all $i \in [k]$. Note that the resulting expenditure is valid and redirecting spending in this way does not add

Algorithm 1 Tree-based price reduction

Input: Feasible prices $\boldsymbol{p}$, valid expenditure e and unsated good j.

Output: Prices $\boldsymbol{p}' \leq \boldsymbol{p}$ and expenditure e', where either $\boldsymbol{p}' = \boldsymbol{p}$ and j is sated, or $p'_k < p_k$ is for some goods $k \in [n]$.

1: Construct tree T by rooting $\mathcal{D}_j$ at j and removing subtrees rooted at goods k of arcs $(i, k) \in \mathcal{D}_j \setminus \mathcal{E}$.
2: **if** $0 \notin T$ **then** go to step 9. **end if**
3: Compute the path $j = g^1, b^1, \ldots, g^k, b^k, g^{k+1} = 0$ from j to null good 0 in T and the minimum m of $p_j s_j - e([m], j)$ and $\min_{i \in [k]} e(g^i, \boldsymbol{b}^i)$. Update e by $e(\boldsymbol{b}^i, g^i) \mathrel{+}= m$ and $e(\boldsymbol{b}^i, g^{i+1}) \mathrel{-}= m$.
4: **if** good j is now sated under e **then**
5:     **return** $\boldsymbol{p}$ and e.
6: **else**
7:     Let i be the smallest such value $i \in [k]$ and remove from T the subtree rooted at g^{i+1}.
8: **end if**
9: Traverse T depth-first to recursively compute $\sum_{k \in T_v} \alpha_k p_k s_k$ for each vertex $v \in T$.
10: Compute $c^{(1)}, c^{(2)}$ and $c^{(3)}$ as specified in Section 4.3, and scaling factor $c^* = \max\left\{c^{(1)}, c^{(2)}, c^{(3)}\right\}$.
11: Traverse T depth-first to recursively compute $\Delta(i, k)$ for each arc in T. Compute e' as described in (2).
12: Let $p'_k = c^* p_k$ for all $k \in T$ and $p'_k = p_k$ otherwise.
13: **return** $\boldsymbol{p}'$ and e'.

edges to the expenditure graph, or cause the expenditure on any goods in $[n]$ to reduce. In particular, no sated good becomes unsated.

If $m = p_j s_j - e([m], j)$, then good j is sated under the new expenditure e', and we are done. Otherwise, if $m < p_j s_j - e([m], j)$, we obtain the tree corresponding to the new expenditure e' by identifying the smallest $i \in [k]$ such that $e'(b^i, g^{i+1}) = 0$ and removing the subtree rooted at g^{i+1} from T. As the tree no longer contains the null good 0, we let J denote the goods in T and are done.

Scaling prices. Now we suppose that we have identified the set of goods J of the tree T not containing 0, and good j is unsated. In order to fully specify our scaling method, it remains to determine the factor $0 < c < 1$ by which to uniformly scale the prices of goods in J so that the resulting prices remain feasible. Let $\boldsymbol{p}'$ denote the prices after scaling, so $p'_k = cp_k$ for $k \in J$ and $p'_k = p_k$ otherwise. We proceed by describing an expenditure e' at $\boldsymbol{p}'$ that is derived from e by adding, or removing, a difference term parameterised by c to, or from, each $e(i, k)$. The scaling factor c is then chosen to ensure that e' is valid. Recall that e' needs to satisfy the conditions in Definition 1 to be a valid expenditure at prices $\boldsymbol{p}'$. We recall, for instance, that the supply constraint imposes $e'([m], k) \leq p'_k s_k$ for all goods $k \in [n]$. In the following, we will impose a stronger supply constraint for all goods $k \in [n] \setminus \{j\}$. Suppose that the proportion of good k's supply aggregately demanded under e at $\boldsymbol{p}$ is given by $\alpha_k \in [0, 1]$, so that $e([m], k) = \alpha_k p_k s_k$ for good k. Then we require that this proportion is maintained under e' at the new prices $\boldsymbol{p}'$, so $e'([m], k) = \alpha_k p'_k s_k$ for all $k \in [n] \setminus \{j\}$. (Hence $e'([m], k) = c\alpha_k p_k s_k$ for $k \in J$ and $e'([m], k) = \alpha_k p_k s_k$ for $k \notin J$.) Note that this trivially implies the supply constraint, as $\alpha_k \leq 1$ and $c < 1$.

The stronger supply condition uniquely determines an expenditure e' for any given c, which can be expressed by adding, or subtracting, the differential terms $\Delta(i, k)$ to, or from, $e(i, k)$. For every arc in T between some buyer i and good k with endpoint $v \in \{i, k\}$, we define $\Delta(i, k) := (1 - c) \sum_{l \in T_v} \alpha_l p_l s_l$. For all other buyers i and k, we let $\Delta(i, k) = 0$. Finally, we define the new expenditure e' as

$$e'(i, k) = \begin{cases} e(i, k) - \Delta(i, k), & \text{if } i \text{ precedes } k \text{ in } T, \\ e(i, k) + \Delta(i, k), & \text{otherwise.} \end{cases} \tag{2}$$

It remains to find the smallest $0 < c < 1$ such that e' is valid. Conceptually, we choose c by reducing from 1 until one of three scenarios occurs. For each scenario $i\{1, 2, 3\}$, we determine the scaling factor $c^{(i)}$ at which this scenario occurs. Our scaling factor is then chosen by computing $c^* = \max\{c^{(1)}, c^{(2)}, c^{(3)}\}$. The scenarios are as follows.

1. Spending $e([m], j)$ on good j increases until it becomes sated. This strictly reduces the number of unsated goods in $\mathcal{D}$.

2. Some $e'(i, k)$ decreases to 0. This breaks up a component in $\mathcal{E}$.

3. Some bid not in T starts demanding a good in T (along with goods outside T). This connects two components in $\mathcal{D}$.

Scenario 1: We note that c must be chosen such that $e'([m], j) = e([m], j) + \Delta([m], j) \leq cp_j s_j$. Substituting $\Delta([m], j) = \sum_{k \in J \setminus \{j\}} (1-c)\alpha_k p_k s_k$ and solving for c, we see that this holds for scaling factor

$$c^{(1)} = \frac{e([m], j) + \sum_{k \in J \setminus \{j\}} \alpha_k p_k s_k}{p_j s_j + \sum_{k \in J \setminus \{j\}} \alpha_k p_k s_k}. \tag{3}$$

As $0 \leq e([m], j) \leq p_j s_j$, factor $c^{(1)}$ is well-defined and $0 < c^{(1)} < 1$.

Scenario 2: Note that $\Delta(i, k) > 0$ for any scaling factor c that is strictly less than 1. Hence, $e'(i, k)$ is non-negative for all arcs from a good k to a buyer i, for any $c < 1$. Now fix an arc from buyer i to good k. In this case, we need to ensure that $\Delta(i, k) \leq e(i, k)$ to guarantee $e'(i, k) \geq 0$. Substituting $\Delta(i, k) = (1-c) \sum_{l \in T_k} \alpha_l p_l s_l$, and solving for c, we get

$$c \geq 1 - \frac{e(i, j)}{\sum_{k \in T_j} \alpha_k p_k s_k}. \tag{4}$$

Hence $c^{(2)}$, as chosen below, guarantees $e'(i, j) \geq 0$ for every buyer i and good j, with equality for at least one pair (i, j).

$$c^{(2)} = \max_{(i,k) \in T} \left(1 - \frac{e(i, k)}{\sum_{l \in T_k} \alpha_l p_l s_l} \right). \tag{5}$$

Note that $0 < e(i, k) \leq \alpha_k p_k s_k \leq \sum_{l \in T_k} \alpha_l p_l s_l$ for all $(i, k) \in T$, so $c^{(2)}$ is well-defined and satisfies $0 < c^{(2)} < 1$.

Scenario 3: Fix a buyer i not in T. Then the buyer continues to demand goods outside T only if $\frac{1}{c} \max_{k \in T} \frac{v_k^i}{p_k} \leq \max_{l \notin T} \frac{v_l^i}{p_l}$. Hence, by solving for c and taking the maximum over all bids not in T, we get

$$c^{(3)} = \max_{i \notin T} \left[\max_{k \in T} \frac{v_k^i}{p_k} \min_{l \notin T} \frac{p_l}{v_l^i} \right]. \tag{6}$$

Note that $1 \leq \max_{k \in T} \frac{v_k^i}{p_k} < \max_{l \in [n]_0} \frac{v_l^i}{p_l}$ for all $i \notin T$, as these bids do not demand any goods in T at $\boldsymbol{p}$. Hence, $c^{(3)}$ is well-defined and $0 < c^{(3)} < 1$.

Theorem 12. *Suppose $\boldsymbol{p}$ is not market-clearing, and e is a valid expenditure at $\boldsymbol{p}$ with an unsated good j. The price-reduction procedure above (Algorithm 1), when given $\boldsymbol{p}$, e and j, either returns a valid expenditure e' at $\boldsymbol{p}$ that sates good j, or it returns reduced prices $\boldsymbol{p}'$ and a valid expenditure e' at $\boldsymbol{p}'$, in polynomial time.*

Proof. Suppose first that the subroutine terminates in line 5. Then the resulting expenditure e' that it returns is trivially valid, and we are done.

Now suppose that the subroutine does not terminate in line 5. Then we note that $\boldsymbol{p}'$ returned by the procedure is strictly smaller than $\boldsymbol{p}$ for all goods $k \in J$, as the scaling factor is strictly less than 1. Hence it suffices to show that the expenditure e' returned by Algorithm 1 after scaling prices of goods J by factor c^* is valid in the sense of Definition 1. By construction of c^*, $e'(i, k)$ is non-negative for all buyers $[m]$ and goods k. We also note that all arcs in T coincide with edges in the demand graph by construction, and if $e(i, k)$ changes, then there is an edge between i and k in the demand graph. Meanwhile, the choice of $c^{(3)}$ guarantees that all buyers $i \notin T$ demand, at $\boldsymbol{p}'$, a superset of the goods they demand at $\boldsymbol{p}$. It follows e is demand-valid.

Next, we verify that e is budget-feasible. Suppose first that i is a buyer not in T. Then, by construction of T, none of the prices of goods that i demands are scaled, and the spending of i is unchanged. Next suppose that buyer i is a leaf in T, and good k is their parent. Then $\Delta(i, k) = \sum_{l \in T_i} \alpha_l p_l s_l = 0$, as the

Algorithm 2 Finding elementwise-minimal feasible prices $\boldsymbol{p}^*$

Input: List of buyers and their valuations in general position.
Output: Element-wise smallest feasible prices $\boldsymbol{p}'$.

1: Let $\boldsymbol{p}$ be the element-wise maximum over the buyers' valuations in $[m]$ and e be the expenditure at which all buyers spend their entire budget on the null good 0.
2: **while** there are unsated goods **do**
3:     Identify an unsated good j under e at $\boldsymbol{p}$.
4:     **while** good j is unsated **do**
5:         Apply tree-based subroutine stated in Algorithm 1 with good j to update $(\boldsymbol{p}, e)$.
6:     **end while**
7: **end while**
8: **return** prices $\boldsymbol{p}$ and expenditure e.

tree T_i contains no goods. Hence, $e'(i,k) = e(i,k) = \beta^i$ and the condition for budget feasibility holds. Finally, suppose buyer i is not a leaf in T. Let g^1 be their parent and $g^2, \ldots, g^k$ be the children. We need to verify that $e'(i,[n]_0) = \beta^i$. As $e(i,[n]_0) = \beta^i$, it suffices to see that $-\Delta(i,g^1) = \sum_{j=2}^{k} \Delta(i,g^j)$. But this follows from the definition of Δ.

Next, we show that e is supply-feasible. For the originally unsated good j, this follows immediately from the choice of our scaling factor $c^* \geq c^{(1)}$ in (3). For all other goods, we verify the stronger supply condition introduced above. Suppose first that k is a good not in T. Then by construction of T, the amount that each bid spends on k remains unchanged. Next suppose good k is a leaf in T, and bid i is its parent. Then the stronger supply constraint requires that $e'(i,k) = ce(i,k) = c\alpha_k p_k s_k$, which agrees with our definition of $\Delta(i,k) = (1-c)\alpha_k p_k s_k$. Finally, suppose good k is not a leaf in T, and let i be its parent and $[m]'$ be its children. Then the stronger supply constraint is satisfied, as

$$e'([m],k) = e([m],k) - \Delta(i,k) + \Delta([m]',k) = \alpha_k p_k s_k - (1-c)\alpha_k p_k s_k = c\alpha_k p_k s_k.$$

To see that the subroutine runs in polynomial time, note that the running time is dominated by traversing the tree with $n+m$ nodes depth-first, which is polynomial in n and m. □

5 Finding market-clearing prices

The simplex-like algorithm of Adsul et al. [1] for linear Fisher markets proceeds by repeatedly applying a price-scaling subroutine until it returns the unique market-clearing prices. Similarly, we can repeatedly apply our subroutine from Section 4 in order to find the unique market-clearing prices in our market in finite time. Note that each call to Algorithm 1 makes progress in the sense that prices decrease or a good is sated. Moreover, the running time analysis we provide in Theorem 13 is analogous to that given in [1], and shows that the algorithm terminates in at most exponential time.

Let $\boldsymbol{p} = \bigvee_{i\in[m]} \boldsymbol{v}^i$ be the element-wise maximum over all valuation vectors. Then it follows that all buyers demand the null good at $\boldsymbol{p}$, while possibly being indifferent between null and other goods. We define the expenditure e at $\boldsymbol{p}$ under which all buyers spend their entire budget on the null good, so for each buyer $i \in [m]$ we let $e(i,0) = \beta^i$ and $e(i,j) = 0$ for all $j \in [n]$. This expenditure is trivially valid.

Of course, no market clears at this initial allocation. We pick any unsated good j and repeatedly call the subroutine of Algorithm 1 with good j, prices $\boldsymbol{p}$ and expenditure e to reduce prices until the market for good j clears (good j becomes sated). Subsequently, we select another unsated good, reduce prices by repeatedly calling Algorithm 1, and repeat this step until all goods are sated.

Theorem 13. *Algorithm 2 finds the market-clearing prices $\boldsymbol{p}^*$ in exponential time.*

Proof. We argue that the number of calls to Algorithm 1 is bounded. Note that the outer loop in Algorithm 2 iterates $O(n)$ times. Indeed, invoking Algorithm 1 does not cause any sated goods to become unsated. Every iteration of the outer loop reduces the number of unsated goods by one, and we start with at most n unsated goods.

We now show that the inner loop of Algorithm 2 applies Algorithm 1 an exponential number of times (in m and n) before good j is sated. Suppose, for contradiction, that the subroutine is called more often.

In this case, we note that on each iteration, the scaling factor is chosen such that scenario 2, scenario 3 or both scenarios occur. Let $\mathcal{D}_j^{(k)}$ denote the connected component of $\mathcal{D}$ containing good j when scenario 2 occurs for the k-th time. In particular, we note that any two consecutive occurrences of scenario 2 can be separated by at most m iterations of scenario 3's. To see this, realise that every time only scenario 3 occurs, T grows through the addition of a buyer, and we have m buyers in total.

We now argue that all $\mathcal{D}_j^{(l)}$ are distinct. As there are exponentially many configurations of $\mathcal{D}_j$, this implies the result. Note that whenever the k-th scenario 2 occurs, we have $e'(i,k) = 0$ for some arc (i,k), and good $k \in T$ is not in T in the next iteration. Hence, the price of good k is not scaled down in the next iteration. Meanwhile, the price of good j is scaled down on every iteration of Algorithm 1. Now suppose, for any k, that $\mathcal{D}_i^{(k)}$ reoccurs at a later stage. This implies that p_j and p_k have been scaled down the same; indeed, this follows by considering the path from j to k in $\mathcal{D}_j^{(k)}$ and applying the conditions of demand. But we have just argued that p_j and p_k have not been scaled down the same, a contradiction.

Finally, we see that the algorithm finds market-clearing prices. To see this, note that it terminates only once all goods are sated. This concludes the proof. □

6 Conclusion

We analyse a market for multiple, divisible goods, in which a unique set of prices exists that induce a socially optimal and revenue-optimal allocation. This co-existence of revenue-optimality and efficiency makes our market compelling in theory as well as highly attractive for sellers, buyers, and market platforms in practice. We provide algorithmic results to derive the efficient and optimal price vector, and expect to improve on these results in future work.

References

[1] B. Adsul, C. S. Babu, J. Garg, R. Mehta, and M. Sohoni. A simplex-like algorithm for linear Fisher markets. *Current Science*, 103(9):1033–1042, 2012.

[2] K. J. Arrow and G. Debreu. Existence of an Equilibrium for a Competitive Economy. *Econometrica*, 22(3):265–290, 1954.

[3] E. Baldwin and P. Klemperer. Understanding Preferences: "Demand Types", and The Existence of Equilibrium with Indivisibilities. *Econometrica*, 87(3):867–932, 2019.

[4] E. Baldwin, P. W. Goldberg, P. Klemperer, and E. Lock. Solving Strong-Substitutes Product-Mix Auctions. arXiv:1909.07313, 2019.

[5] B. Birnbaum, N. R. Devanur, and L. Xiao. Distributed algorithms via gradient descent for fisher markets. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, EC '11, page 127–136, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450302616. doi: 10.1145/1993574.1993594. URL https://doi.org/10.1145/1993574.1993594.

[6] W. C. Brainard and H. E. Scarf. How to Compute Equilibrium Prices in 1891. *American Journal of Economics and Sociology*, 64(1):57–83, 2005.

[7] S. Chakraborty, N. R. Devanur, and C. Karande. Market equilibrium with transaction costs. In A. Saberi, editor, *Internet and Network Economics*, pages 496–504, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. ISBN 978-3-642-17572-5.

[8] L. Chen, Y. Ye, and J. Zhang. A note on equilibrium pricing as convex optimization. In X. Deng and F. C. Graham, editors, *Internet and Network Economics*, pages 7–16, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg. ISBN 978-3-540-77105-0.

[9] R. Cole, N. R. Devanur, V. Gkatzelis, K. Jain, T. Mai, V. V. Vazirani, and S. Yazdanbod. Convex Program Duality, Fisher Markets, and Nash Social Welfare. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC '17)*, EC '17, pages 459–460, Cambridge, Massachusetts, USA, 2017. Association for Computing Machinery.

[10] V. Conitzer, C. Kroer, D. Panigrahi, O. Schrijvers, E. Sodomka, N. E. Stier-Moses, and C. Wilkens. Pacing equilibrium in first-price auction markets. *arXiv preprint arXiv:1811.07166*, 2021.

[11] V. Conitzer, C. Kroer, E. Sodomka, and N. E. Stier-Moses. Multiplicative pacing equilibria in auction markets. *Operations Research*, 2021.

[12] N. R. Devanur and V. V. Vazirani. The spending constraint model for market equilibrium: Algorithmic, existence and uniqueness results. In *Conference Proceedings of the Annual ACM Symposium on Theory of Computing*, STOC '04, pages 519–528, Chicago, Illinois, USA, 2004.

[13] N. R. Devanur, C. H. Papadimitriou, A. Saberi, and V. V. Vazirani. Market Equilibrium via a Primal-Dual Algorithm for a Convex Program. *Journal of the ACM*, 55(5):1–18, 2008.

[14] S. Dobzinski, R. Lavi, and N. Nisan. Multi-unit auctions with budget limits. *Games and Economic Behavior*, 74(2):486–503, 2012.

[15] B. C. Eaves. A finite algorithm for the linear exchange model. *Journal of Mathematical Economics*, 3(2):197–203, 1976.

[16] E. Eisenberg and D. Gale. Consensus of Subjective Probabilities: The Pari-Mutuel Method. *The Annals of Mathematical Statistics*, 30(1):165 – 168, 1959. doi: 10.1214/aoms/1177706369. URL https://doi.org/10.1214/aoms/1177706369.

[17] M. Fichtl. Computing Candidate Prices in Budget-Constrained Product-Mix Auctions. Working paper. https://dss.in.tum.de/files/staff-files/fichtl/BudgetConstrainedPMA.pdf, accessed 12/07/2021.

[18] J. Garg, R. Mehta, M. Sohoni, and N. K. Vishnoi. Towards Polynomial Simplex-Like Algorithms for Market Equilibria. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '13, pages 1226–1242, 2013.

[19] V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *SODA*, volume 5, pages 1164–1173, 2005.

[20] K. Jain. A polynomial time algorithm for computing an Arrow-Debreu market equilibrium for linear utilities. *SIAM Journal on Computing*, 37(1):303–318, 2007.

[21] P. Klemperer. A New Auction for Substitutes: Central Bank Liquidity Auctions, the U.S. TARP, and Variable Product-Mix Auctions. Working Paper, 2008.

[22] P. Klemperer. The Product-Mix Auction: a New Auction Design for Differentiated Goods. *Journal of the European Economic Association*, (8):526–36, 2010.

[23] P. Klemperer. Product-Mix Auctions. Nuffield College Working Paper 2018-W07, http://www.nuffield.ox.ac.uk/users/klemperer/productmix.pdf, 2018.

[24] R. Murray, C. Kroer, A. Peysakhovich, and P. Shah. Robust market equilibria with uncertain preferences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2192–2199, 2020.

[25] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. Algorithmic game theory, 2007. *Book available for free online*, 2007.

[26] J. B. Orlin. Improved Algorithms for Computing Fisher's Market Clearing Prices. In *Proceedings of the Annual ACM Symposium on Theory of Computing*, STOC '10, pages 291–300, Cambridge, Massachusetts, USA, 2010.

[27] L. Walras. Éléments d'économie politique pure ou théorie de la richesse sociale (elements of pure economics, or the theory of social wealth). Lausanne, Paris, 1874. (1899, 4th ed.; 1926, rev ed., 1954, Engl. transl.), 1874.

A On concave functions

The following example can be used to show revenue and welfare do not generally coincide when we move away from additive valuations, even when we have only one good. The valuation function is chosen to be fairly 'benign'. Note that the budget constraints never apply, so this example also applies to general quasi-linear utilities (without budget constraints).

Example 3 (Concave valuation)**.** Consider an auction with a single good available in $s = 3$ units that has a single buyer. The buyer has valuation $v : \mathbb{R} \to \mathbb{R}$ given by $v(x) = \frac{4}{\log 2}(1 - 2^{-x})$ and budget 2. Then revenue is not maximised at market-clearing prices.

Indeed, note that the utility of the buyer for quantity q at price p is $u(x, p) = v(x) - px$, so the buyer's demand $D(p)$ at p is found by solving $v'(x) = p$, which yields $x = -\log_2(\frac{p}{4})$. At $p = 0.5$, we have demand $x = 3$, so p clears the market. Revenue at p is $px = 1.5$. At price $q = 1$, we have demand $y = 2$, so q does not clear the market, but revenue is $qy = 2$, which is greater. Revenue is maximised at $p = \frac{4}{e}$ with a demanded quantity of $\frac{1}{\log(2)}$ and a revenue of $\frac{4}{e \log(2)}$.

Proof of Proposition 1. Recall that at any price p, the buyer demands the bundle x that maximises $v(x) - p$. Given valuation $v(x)$, revenue is maximised at $(x, p) \in \arg\max_{x,p} px$ such that $v'(x) = p$ and $x \leq s$. Thus, maximal revenue given v and s is $v'(x)x$ for some $x \leq s$. Social welfare is maximised at $(x, p) \in \arg\max_{x,p} v(x) - px$ such that $x \leq s$, i.e. revenue at the social optimum is $v'(s)s$. We show that there exists x and s with $x < s$ and $v'(x)x > v'(s)s$. Recall that a differentiable function $f : \mathbb{R} \to \mathbb{R}$ is strongly concave if it satisfies $|f'(x) - f'(y)| \geq m\|x - y\|$ for all distinct $x, y \in \mathbb{R}$ (v does not need to be twice differentiable).

First, note that there exists $\overline{x} < \infty$ such that $v'(x) < mx$ for all $x \geq \overline{x}$, due to strict concavity of v. Fix some supply $s \geq \overline{x}$ and let $\varepsilon = \frac{1}{2}(ms - v'(s))$. As v is strongly concave, we also have $v'(s - \varepsilon) \geq v'(s) + m\varepsilon$ for any ε. Hence,

$$\begin{aligned} v'(s-\epsilon)(s-\epsilon) &\geq (v'(s) + m\epsilon)(s - \epsilon) \\ &= v'(s)s + \epsilon(ms - v'(s) - m\epsilon) \\ &> v'(s)s. \end{aligned}$$

□

Proof of Proposition 2. First, suppose that there exists supply $s \in X$ with $v'(s) < ms$. Then the result follows analogous to the proof of Proposition 1.

We now prove the second part of the statement. The budget constraint is given by $px \leq \beta$. At the demanded quantity, $v'(x) = p$ holds. Thus, for all feasible x, it must hold that $v'(x) \leq \beta/x$, i.e. $x \leq \tilde{x}$. Now we demonstrate that $v'(\tilde{x} - \epsilon) \leq m(\tilde{x} - \epsilon)$ for some small ϵ. Then the result follows from Proposition 2. For some $\delta > 0$, we have

$$\begin{aligned} m(\tilde{x} - \epsilon) &\geq \left(\frac{v'(\tilde{x})}{\tilde{x}} + \delta\right)(\tilde{x} - \epsilon) \\ &= \frac{(v'(\tilde{x}) + \delta\tilde{x})(\tilde{x} - \epsilon)}{\tilde{x}} \\ &\geq v'(\tilde{x} - \epsilon) \end{aligned}$$

The last inequality holds for $\epsilon \to 0$ and some $\delta > 0$. □

B Omitted proofs

It is known [8] that the market we consider has a unique market-clearing price vector (but allocations at these prices may not be unique). Here we provide a direct proof of this.

Theorem 14. *There is a unique price vector* $\boldsymbol{p}^*$ *at which the market is cleared.*

Proof. Suppose we have a valid allocation π that exhausts supply $\boldsymbol{s}$. We want to find the corresponding prices $\boldsymbol{p}$ that support π. First, we formulate the problem of finding $\boldsymbol{p} \geq 0$ as a system of linear inequalities by writing out the budget constraints as equalities and the demand constraints as inequalities. The latter is achieved by adding, for every buyer i who demands a positive amount of good j, the inequalities $\frac{v^i_j}{p_j} \geq \frac{v^i_k}{p_k}, \forall k \in [n]_0$. Our goal is to prove that there is exactly one solution to this system of inequalities. In Section 4, we show that the elementwise-minimal prices $\boldsymbol{p}^*$ clear the market, so the polytope defined by this system has at least one solution.[8]

For contradiction, suppose there is a second feasible price vector $\boldsymbol{p} \geq \boldsymbol{p}^*$ that also clears the market. Then by convexity of the polytope, it follows that $\boldsymbol{p}' = \boldsymbol{p}^* + \varepsilon(\boldsymbol{p} - \boldsymbol{p}^*)$ is also market-clearing for any $0 \leq \varepsilon \leq 1$. From now on, we assume that $\boldsymbol{p}'$ is obtained by letting ε be infinitesimally small. Let $S := \{j \in [n] \mid p^*_j < p_j\}$ denote the prices that change when we move from $\boldsymbol{p}^*$ to $\boldsymbol{p}'$. As $\boldsymbol{p} \neq \boldsymbol{p}^*$, S is non-empty. We first make a technical observation about the demand of each bid at $\boldsymbol{p}^*$ and $\boldsymbol{p}'$.

Observation 1. For any buyer $i \in [m]$, we note: (i) If $J^i(\boldsymbol{p}^*) \not\subseteq S$, then $J^i(\boldsymbol{p}') = J^i(\boldsymbol{p}^*) \setminus S$. (ii) If $J^i(\boldsymbol{p}^*) \subseteq S$, then $J^i(\boldsymbol{p}') \subseteq J^i(\boldsymbol{p}^*)$. This holds because ε was chosen to be infinitesimally small.

Now suppose e' is a valid expenditure at $\boldsymbol{p}'$ that clears the market. Let $\mathcal{S} := \{i \in [m] \mid J^i(\boldsymbol{p}^*) \subseteq S\}$ denote the set of buyers satisfying case (ii) of Observation 1. We take a look at the revenue contributions from each buyer towards goods in S. At $\boldsymbol{p}'$, Observation 1 tells us that only buyers $i \in \mathcal{S}$ spend on goods in S, and they each spend their entire budget on S. Hence, the budget and supply feasibility of e' imply $\sum_{i\in\mathcal{S}} \beta^i = e'(\mathcal{S}, S) = e'([m], S) = \sum_{j\in S} p'_j s_j$. At $\boldsymbol{p}^*$, the buyers of $\mathcal{S}$ also spend their entire budget on S, but it may be the case that other buyers also spend part of their budget on S. Hence, $\sum_{i\in\mathcal{S}} \beta^i = e(\mathcal{S}, S) \leq e([m], S) = \sum_{j\in S} p^*_j s_j$. It follows that $\sum_{j\in S} p'_j s_j \leq \sum_{j\in S} p^*_j s_j$, in contradiction to the fact that $p'_j > p^*_j$ for all $j \in S$. □

B.1 Welfare theorems

We show that the first and second welfare theorem hold in our market. First we define efficiency formally.

Definition 4 (Efficiency)**.** A feasible allocation π is *efficient* if $\sum_{i\in[m],j\in[n]} v_j \pi(i,j) \geq \sum_{i\in[m],j\in[n]} v_j \pi'(i,j)$ for all feasible allocations $\pi' \neq \pi$.

Proposition 15. *Suppose π is a budget-feasible allocation at prices $\boldsymbol{p}$ and fix buyer i. Then $\sum_{j\in[n]_0}(v^i_j - p_j)\pi(i,j) \geq \sum_{j\in[n]_0}(v^i_j - p_j)\pi'(i,j)$ for all budget-feasible allocations π' if, and only if, $\pi(i) \in D^i(\boldsymbol{p})$.*

Proof. This follows immediately from budget-constrained quasi-linearity. □

Theorem 16 (First welfare theorem)**.** *If $(\boldsymbol{p}, \pi)$ is a competitive equilibrium, π is an efficient allocation.*

Proof. Let $(\boldsymbol{p}, \pi)$ be a competitive equilibrium with corresponding expenditure e. Let π' be a feasible allocation at $\boldsymbol{p}$ with expenditure e'. Proposition 15 implies

$$\sum_{i\in[m],j\in[n]} (v^i_j - p_j)\pi(i,j) \geq \sum_{i\in[m],j\in[n]} (v^i_j - p_j)\pi'(i,j).$$

Rearranging and using the fact that $e([m],[n]) = \boldsymbol{p} \cdot \boldsymbol{s} \geq e'([m],[n])$, we have

$$\sum_{i\in[m],j\in[n]} v^i_j \pi(i,j) - \sum_{i\in[m],j\in[n]} v^i_j \pi'(i,j) \geq e([m],[n]) - e'([m],[n]) \geq 0.$$

□

Theorem 17 (Second welfare theorem)**.** *Let $(\boldsymbol{p}, \pi)$ be a competitive equilibrium and π' be another efficient and feasible allocation. Then $(\boldsymbol{p}, \pi')$ is also a competitive equilibrium.*

Proof. By definition of efficiency, we first note that

$$\sum_{i\in[m],j\in[n]} v^i_j \pi(i,j) = \sum_{i\in[m],j\in[n]} v^i_j \pi'(i,j). \tag{7}$$

[8] Indeed, the algorithm in Section 5 finds these market-clearing prices.

Secondly, as π is a competitive equilibrium and π' is feasible, we have $\pi([m]) = \boldsymbol{s}$ and $\pi'([m]) \leq \boldsymbol{s}$. Multiplying both terms with $\boldsymbol{p}$, we get

$$\boldsymbol{p} \cdot \pi([m]) \geq \boldsymbol{p} \cdot \pi'([m]). \tag{8}$$

We now show that equality must hold for (8) by demonstrating that strict inequality contradicts (7). This proves our claim. Indeed, using (8) with strict inequality and Proposition 15, we obtain

$$\begin{aligned}
\sum_{i\in[m],j\in[n]} v_j^i \pi(i,j) &= \sum_{i\in[m],j\in[n]} (v_j^i - p_j)\pi(i,j) + \boldsymbol{p} \cdot \pi([m]) && (9)\\
&> \sum_{i\in[m],j\in[n]} (v_j^i - p_j)\pi'(i,j) + \boldsymbol{p} \cdot \pi'([m]) && (10)\\
&= \sum_{i\in[m],j\in[n]} v_j^i \pi'(i,j). && (11)
\end{aligned}$$

□

C Examples

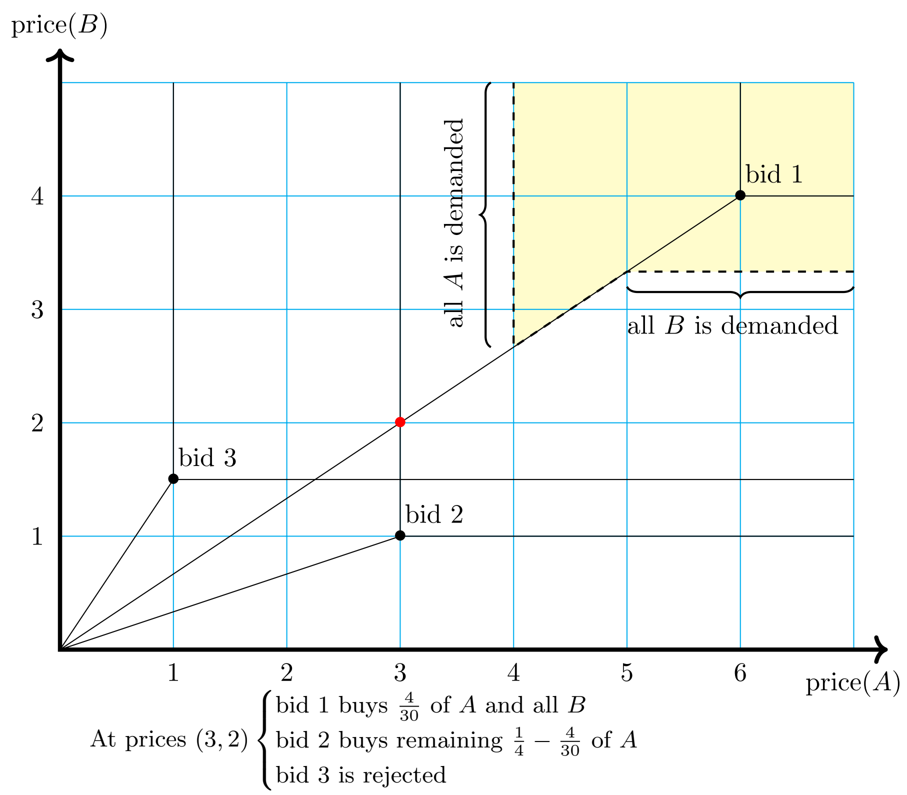

Figure 4: Example with 3 bids each of weight (or, budget) 1. With a supply of $\frac{1}{4}$ units of A and $\frac{3}{10}$ of B, the yellow region is "feasible prices" where at most this supply is demanded. It includes a 1-dimensional line between $(4, \frac{8}{3})$ and $(3, 2)$. It has the key property that for any pair of feasible price vectors, the coordinate-wise minimum of them also belongs to the region. The red dot at $(3, 2)$ is the unique coordinate-wise minimal price where this supply is demanded. Bid 2 is partially fulfilled: we prioritise bid 1 to maximise welfare. Notice that if we increase the weights of bids 2 or 3, that has no effect on the final price vector.