Market Clearing Algorithms

Tuomas Sandholm

sandholm@cs.cmu.edu
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213

Subhash Suri

suri@cs.ucsb.edu
Department of Computer Science
University of California
Santa Barbara, CA 93106

Abstract

Commerce is moving online to an increasing extent, and in electronic commerce there has been a significant shift to dynamic pricing via auctions, reverse auctions, and exchanges. We study optimal clearing of markets in the ubiquitous setting where there are multiple indistinguishable units of an item for sale. This setting is common in markets for stocks, bonds, electricity, bandwidth, oil, pork bellies, memory chips, CPU time, etc. The sellers and buyers express their bids via supply/demand curves. We focus on the natural and classical economic setting where each seller's supply increases and each buyer's demand decreases as the price increases (otherwise, absurd outcomes can occur [18]). Discriminatory pricing leads to greater (or equal) economic surplus for the market than non-discriminatory pricing. However, this comes at the cost of computational complexity.

We first study the expressive case where the supply/demand curves are piecewise linear and not necessarily continuous (such curves can approximate any curve arbitrarily closely). We present a polynomial-time algorithm for clearing non-discriminatory markets, and show that clearing discriminatory markets is $\mathcal{NP}$-Complete (even in the setting where there is only one player on one side of the market, that player has a constant supply (demand) curve, and the players on the other side of the market have step function demand (supply) curves).

We then show that in the more restricted setting of linear supply/demand curves, even discriminatory markets can be cleared in polynomial time. Our derivations also uncover the following interesting fact. To obtain the optimal discriminatory solution, each buyer's (seller's) price is incremented (decremented) *equally* from her individual quantity-unconstrained solution. As this price change proceeds, some players' quantities go to zero, at which point those players are excluded from the market. There is no need to consider re-entry of the excluded players' curves in the process.

1 Introduction

Commerce is moving online to an increasing extent, and in electronic commerce there has been a significant shift to dynamic pricing via auctions, reverse auctions, and exchanges. The range of applications includes trading in stock markets [1], procurement in business-to-business commerce, bandwidth allocation in communication networks [2, 3, 7], as well as resource allocation in operating systems [19] and computational grids [14, 20].

These trends have led to an increasing need for fast market clearing algorithms [6, 13]. Also, recent electronic commerce server prototypes such as *eMediator* [16] and *AuctionBot* [21] have demonstrated a wide variety of new market designs, leading to the need for new clearing algorithms. In addition, there has been a recent surge of interest in clearing algorithms for combinatorial auctions [4, 9, 12, 15, 17] where bids can be submitted on bundles of items. Combinatorial auctions allow the bidders to express complementarity and substitutability among multiple distinguishable items (possibly multiple units of each).

In this paper we study the ubiquitous market setting where there are multiple *indistinguishable* units of an item for sale (or there are multiple units of multiple items for sale, but the different items can be treated independently as belonging to different markets). This setting is common in markets for stocks, bonds, electricity, bandwidth (see, e.g., RateXChange), oil, pork bellies, memory chips, CPU time, etc. The naive approach to bidding would require the bidders to express their offers as a sequence of points, for example ($2 for 1 unit) XOR ($5 for 2 units) XOR ($6 for 3 units), etc. However, the mapping from quantities to prices can be represented more compactly by allowing each bidder to express his offer as a price-quantity curve. Price-quantity curves are ubiquitous in economics [11], and are becoming common in electronic commerce as well [8, 10, 16].

In this paper we analyze the inherent complexity of, and design algorithms for, clearing auctions, reverse auctions, and exchanges with supply/demand curves. We view these as some of the most natural problems in the intersection of computer science and economics.

In classic economic theory of supply and demand curves (called partial equilibrium theory [11]), the market is cleared as follows. The supply curves of the sellers are aggregated. Also, the demand curves of the buyers are aggregated. Then, any price for which supply meets demand is an equilibrium (there may be multiple solutions). An important point of partial equilibrium theory is that all buyers and sellers share a common price.

This restriction, however, seems unnecessary, and in most real markets, such as stock markets, different parties are cleared at different prices. The use of multiple prices allows an economic *surplus* to be created that increases welfare in the system. The surplus can be allocated to the market maker, or it can be divided among the market participants (how it is divided can affect the bidders' incentives for revealing their preferences, but we do not address incentives in this paper).

We focus on the computational implications of different pricing schemes on clearing markets with supply/demand curves. We first study *non-discriminatory* markets where there are two prices: one shared by the sellers, and another shared by the buyers. We then study *discriminatory* markets where each participant may be cleared at a different price. We show that markets with piecewise linear supply/demand curves are clearable in polynomial time under non-discriminatory pricing, but $\mathcal{NP}$-Complete to clear under discriminatory pricing. With linear supply/demand curves, even discriminatory markets can be cleared in polynomial time. Our clearing algorithms are all $O(N \log N)$.

2 The Market Model

We consider trading of multiple, indistinguishable units of a single item. The market has two kinds of participants: buyers and sellers. Say there are n sellers and m buyers. Without loss of generality, we assume that no agent is both a buyer and a seller (if an agent is both, we treat his as two separate agents).

Each seller expresses his willingness to sell via a *price-quantity* curve, called his *supply curve*. The supply curve is a function $s : R^+ \to R^+$ from non-negative *unit* prices to non-negative supply quantity. Thus, if the unit price is p, then the seller with supply curve s is willing to supply $s(p)$ units of the good.[1] Similarly, each buyer submits a price-quantity *demand* curve $d : R^+ \to R^+$.

We assume that supply/demand curves are *piecewise linear*. They need not be continuous—they can can have discrete "jumps". Such piecewise linear curves are quite general; they can approximate any supply or demand curve arbitrarily closely. Arbitrary demand and supply curves, however, can lead to absurd results (infinite values, zero costs, etc. [18]), and thus it is customary to place some natural restrictions on their slopes. We assume that *supply curves* are *upward sloping* and *demand curves* are *downward sloping*. This is economically reasonable in that higher prices increase supply and decrease demand.

The goal of a market is to facilitate trading between buyers and sellers (not necessarily one-to-one) in order to optimize some objective function. Many different objective functions are conceivable, depending on how the market is to be used. In this paper, we use *maximizing surplus* as the objective. The *surplus* is defined as the quantity-weighted sum of cleared demand prices minus the quantity-weighted sum of cleared supply prices. In other words, if the market buys the goods from the sellers and sells them to the buyers using their price-quantity curves, then surplus is the *profit* that the market generates. Assuming that the buyers and sellers reveal their preferences truthfully with their supply/demand curves, this objective leads to the highest possible social welfare among all the participants. So, whether the market designer is selfish who wants to maximize his profit collected, or benevolent who wants to maximize social welfare, he will prefer to use that same objective.

We study two different clearing policies: *non-discriminatory* and *discriminatory* pricing.

2.1 Non-Discriminatory Pricing

In a non-discriminatory market, there are two clearing prices, one shared by the sellers and one shared by the buyers. The quantity sold or bought by each participant is then determined by the participant's price-quantity curve, evaluated at the clearing price.

Specifically, suppose the market clears the sellers at the unit price p^*_{ask}, and the buyers at the unit price p^*_{bid}. These prices uniquely determine the quantity supplied by each seller, and the quantity bought by each buyer, using their price-quantity curves. In particular, suppose s_i is the supply curve of seller i, and d_j is the demand curve of buyer j. Then, for the solution to be feasible, supply must equal demand:

$$\sum_{i=1}^{n} s_i(p^*_{ask}) \;=\; \sum_{j=1}^{m} d_j(p^*_{bid}).$$

Subject to this, the goal of the market is to maximize the surplus:

$$p^*_{bid} \sum_{j=1}^{m} d_j(p^*_{bid}) \;-\; p^*_{ask} \sum_{i=1}^{n} s_i(p^*_{ask})$$

Thus, in a *non-discriminatory* market, the computational problem is to determine the seller and buyer clearing prices p^*_{ask} and p^*_{bid}.

2.2 Discriminatory Pricing

In a discriminatory market, the market can clear each seller and each buyer at a distinct unit price. Specifically, suppose seller i is cleared at unit price p^*_i, and buyer j is cleared at unit price p^*_j. Then, the feasibility condition of supply meeting demand is

[1] The model also covers the possibility that the seller (buyer) is willing to accept a higher (lower) price for the same quantity. However, given our objective, in any optimal solution, each party is cleared exactly on his supply/demand curve.

$$\sum_{i=1}^{n} s_i(p_i^*) \;=\; \sum_{j=1}^{m} d_j(p_j^*),$$

and the surplus to be maximized is

$$p_j^* \sum_{j=1}^{m} d_j(p_j^*) \;-\; p_i^* \sum_{i=1}^{n} s_i(p_i^*).$$

In such a *discriminatory* market, the computational problem is to determine the clearing price for each seller and buyer so as to maximize the surplus.

The surplus generated under discriminatory pricing is greater (or equal) than that under non-discriminatory pricing. However, non-discriminatory markets offer a stronger notion of fairness. Discriminatory markets, however, do offer a weak form of *ex ante* fairness: they are anonymous in the sense that had two players swapped their bids, their allocations would also have been swapped.

3 Clearing Preliminaries

We begin with the simple case of a *one seller, one buyer* market, where the seller has an upward sloping linear supply curve $q = a_s p - b_s$, and the buyer has a downward sloping linear demand curve $q = -a_d p + b_d$, where $a_s, a_d > 0$, and $b_s, b_d \geq 0$. The following elementary lemmata will be useful throughout the paper. (Observe that in a single-seller, single-buyer market, non-discriminatory and discriminatory pricing are identical.) Figure 1 illustrates the clearing in this setting.

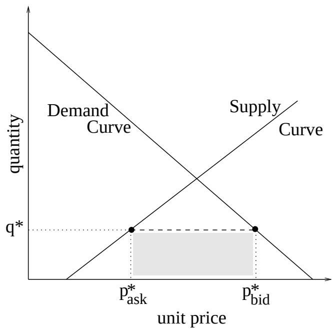

Figure 1: Single-seller, single-buyer market with linear supply/demand curves. The surplus equals the area of the shaded rectangle. Viewed geometrically, the clearing occurs at quantity q^* which is half the height of the triangle formed by the supply and demand lines.

Lemma 3.1 (Single-Seller, Single-Buyer Unconstrained Trading) *Consider a market of one seller, with upward sloping supply $q = a_s p - b_s$, and one buyer, with downward sloping demand $q = -a_d p + b_d$, where $a_s, a_d > 0$, and $b_s, b_d \geq 0$. Then, the surplus-maximizing trade occurs at quantity*

$$q^* \;=\; \frac{1}{2}\left(\frac{a_s b_d - a_d b_s}{a_s + a_d}\right)$$

The clearing prices for the seller and the buyer are

$$p^*_{ask} \;=\; \frac{1}{2}\left(\frac{b_s}{a_s} \;+\; \frac{b_s + b_d}{a_s + a_d}\right), \qquad\qquad p^*_{bid} \;=\; \frac{1}{2}\left(\frac{b_d}{a_d} \;+\; \frac{b_s + b_d}{a_s + a_d}\right)$$

PROOF. The proof is based on simple calculus and is given in the appendix. □

In general, when participants put price or quantity constraints on their curves, the trade will not occur at the "unconstrained" optimal quantity q^* determined by the preceding lemma. The following lemma determines the effect of moving the traded quantity away from q^*. We consider the same one-seller, one-buyer market as above, and compute the surplus achieved when trade occurs at some quantity $q^* + \varepsilon$, where ε can be positive or negative. (We assume that $|\varepsilon| \leq q^*$, which ensures that the trade is feasible, and produces non-negative surplus.)

Lemma 3.2 *In the single-seller, single-buyer market, if the trade occurs at quantity $q^* + \varepsilon$, then the surplus is*

$$q^*(p^*_{bid} - p^*_{ask}) \; - \; \varepsilon^2 \left(\frac{1}{a_s} + \frac{1}{a_d} \right).$$

That is, the surplus shrinks *(quadratically) with $|\varepsilon|$.*

PROOF. Let us consider the case of $\varepsilon > 0$; the other case is analogous. When the traded quantity is increased by ε, the seller's clearing unit price *increases* by ε / a_s, and the buyer's unit clearing price *decreases* by ε / a_d. The new surplus, therefore, equals

$$(q^* + \varepsilon) \left((p^*_{bid} - \frac{\varepsilon}{a_d}) \; - \; (p^*_{ask} + \frac{\varepsilon}{a_s}) \right),$$

which can be written as

$$q^*(p^*_{bid} - p^*_{ask}) \; - \; \varepsilon^2 (\frac{1}{a_s} + \frac{1}{a_d}) \; + \; \varepsilon (p^*_{bid} - p^*_{ask}) \; - \; \varepsilon q^* (\frac{1}{a_s} + \frac{1}{a_d})$$

The proof is completed by showing that the 3rd and the 4th terms of this expression are equal, and therefore cancel each other out. To see this equality, we substitute the expressions for q^*, p^*_{bid}, and p^*_{ask} from Lemma 3.1.

$$\begin{aligned} \varepsilon (p^*_{bid} - p^*_{ask}) &= \varepsilon \left(\frac{b_d}{2a_d} + \frac{b_s + b_d}{2(a_s + a_d)} - \frac{b_s}{2a_s} - \frac{b_s + b_d}{2(a_s + a_d)} \right) = \frac{\varepsilon}{2} \left(\frac{b_d}{a_d} - \frac{b_s}{a_s} \right) \\ \varepsilon q^* (\frac{1}{a_s} + \frac{1}{a_d}) &= \frac{\varepsilon}{2} \left(\frac{a_s b_d - a_d b_s}{a_s + a_d} \right) \left(\frac{1}{a_s} + \frac{1}{a_d} \right) = \frac{\varepsilon}{2} \left(\frac{b_d}{a_d} - \frac{b_s}{a_s} \right) \end{aligned}$$

Thus, the new surplus equals $q^*(p^*_{bid} - p^*_{ask}) \; - \; \varepsilon^2 \left(\frac{1}{a_s} + \frac{1}{a_d} \right)$, which completes the proof. □

Our next lemma states an easy corollary of Lemma 3.2 for the case where the buyer and seller curves are quantity-constrained. That is, suppose the buyer's curve is the downward-sloping linear function $q = -a_d p + b_d$, but restricts quantity to the range $[q'_d, q''_d]$, and the seller's curve is the upward-sloping linear function $q = -a_d p + b_d$, but restricts the quantity to the range $[q'_s, q''_s]$. What trade maximizes the surplus? Of course, the only feasible trades that can occur are those in the quantity range that is common to both. So, assume that $[q', q'']$ is the intersection of the intervals $[q'_d, q''_d]$ and $[q'_s, q''_s]$.

Lemma 3.3 (Single-Seller, Single-Buyer Bounded Trading) *Consider the single-seller, single-buyer market, with linear price-quantity curves, where the buyer and the seller can only trade in the quantity range $[q', q'']$. The surplus-maximizing trade occurs either at q^* (if $q^* \in [q', q'']$), or at that endpoint of the range $[q', q'']$ which is* closer *to q^*.*

PROOF. Of course, if the unconstrained trade quantity q^* is in the feasible range, then the surplus is maximized at q^*. Otherwise, Lemma 3.2 shows that the surplus shrinks quadratically with the deviation ε from q^*. Thus, the optimal trade occurs at the feasible point closest to q^*, which is an endpoint of the range $[q', q'']$. □

4 Non-Discriminatory Markets

In this section, we show how to clear a market with multiple buyers and sellers, each with *piecewise linear* demand/supply curves. The clearing policy of the market is non-discriminatory; that is, all buyers pay the same unit price p^*_{bid}, and all sellers receive the same unit price p^*_{ask}. Recall that then our goal is to maximize the surplus subject to supply meeting demand:

$$\text{maximize} \quad p^*_{bid} \sum_{j=1}^{m} d_j(p^*_{bid}) \; - \; p^*_{ask} \sum_{i=1}^{n} s_i(p^*_{ask}) \qquad \text{s.t.} \qquad \sum_{i=1}^{n} s_i(p^*_{ask}) \; = \; \sum_{j=1}^{m} d_j(p^*_{bid})$$

Our approach is to *aggregate* the demand and the supply curves separately, and then reduce the problem to the single-seller, single-buyer case. The single seller is the aggregate of all sellers, and the single buyer is the aggregate of all buyers. The key idea here is that since *all buyers are cleared at the same price* p^*_{bid}, we can infer quantities sold to each individual buyer by evaluating their curves at p^*_{bid}. The same holds for the sellers.

A *piecewise linear* price-quantity curve consists of one or more bounded linear curves. We do not require the curve to be continuous (i.e., the quantity can "jump" between pieces). Let us consider the aggregation of buyer curves; the seller curves are handled in the same way. Consider a set of piecewise linear demand curves $d_1, d_2, \ldots, d_n$. Their *aggregate curve* is a piecewise linear function $D : R^+ \to R^+$ such that $D(p)$ is the total demand at unit price p. That is, $D(p) \; = \; d_1(p) + d_2(p) + \ldots + d_n(p)$, where $d_i(p)$ is the demand by curve i at unit price p. The aggregation of *linear* functions leads to a linear function. Thus, if a price interval $[p_1, p_2]$ does not contain the *breakpoints* of any of the demand curves, then the aggregate curve in the interval $[p_1, p_2]$ has the form $q = (\sum_i a_i)p + \sum_i b_i$, where a_i and b_i are the coefficients of the component linear curves.

Given piecewise linear curves, their aggregate is easily computed in $O(K \log K)$ time, by a sweep-line algorithm, where K is the total number of linear pieces in all the curves.

Despite the fact that the input demand curves are downward sloping, the aggregate demand curve need not be downward sloping. Similarly, the aggregate supply curve need not be upward sloping. This can lead to the problem of having multiple prices for a given quantity. We can rectify this by computing the *rightmost* envelope of the aggregate demand curve, and *leftmost* envelope of the aggregate supply curve. In other words, for each quantity q, we just keep the point with the maximum demand price, and the minimum supply price. This does not affect the solution space, since in a surplus-maximizing market, all clearings would occur only at these *envelope* prices.

Our non-discriminatory clearing algorithm can now be described as follows. (Recall that individual demand curves are downward sloping, and individual supply curves are upward sloping.) See Figure 2.

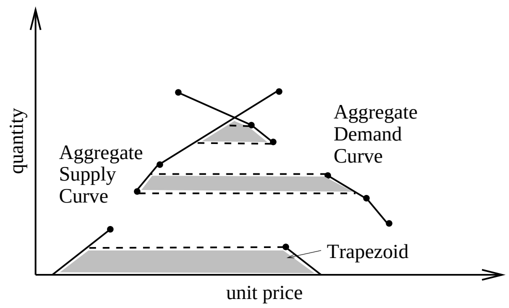

Figure 2: Non-discriminatory market. The figure shows decomposition of the feasible region into four trapezoids. Each trapezoid corresponds to a single-seller, single-buyer market. The aggregate demand curve is intentionally drawn to be discontinuous and not downward sloping (although each piece is downward sloping).

Algorithm ND-Market

1. Compute the piecewise linear aggregate demand curve D, and the aggregate supply curve S.
2. Let the *feasible space* denote the set of points (p, q) for which $D(p)$ and $S(p)$ both exist and $S(p) \leq D(p)$; that is, there is both aggregate demand and aggregate supply for q, and the aggregate demand price is no smaller than the aggregate supply price.
3. Decompose the feasible region into trapezoids, by "drawing" horizontal lines through each breakpoint of D or S.
4. The market clearing problem for each trapezoid corresponds to the single-seller, single-buyer bounded trade (cf. Lemma 3.2), so it can be solved in $O(1)$ time. Since there are $O(K)$ trapezoids, the time complexity of this step is $O(K)$.
5. The maximum-surplus solution over all trapezoids is the optimal solution. Once the clearing prices, p^*_{bid} and p^*_{ask}, are determined, we can evaluate each seller curve at p^*_{ask} and each buyer's curve at p^*_{bid} to determine the quantity sold by each seller, and bought by each buyer.

We summarize this result in the following theorem.

Theorem 4.1 *Consider a single-item, multi-unit market with multiple sellers and buyers, where each seller (buyer) has an upward (downward) sloping piecewise linear curve. Then, the surplus-maximizing clearing of such a market using non-discriminatory pricing can be determined in $O(K \log K)$, where K is the total number of pieces in all of the piecewise linear curves.*

5 Discriminatory Markets

In a discriminatory market, the market can clear each seller and buyer at a distinct unit price. We now study the complexity of clearing discriminatory markets.

5.1 Intractability with Piecewise Linear Curves

In sharp contrast to a non-discriminatory market, we show that clearing a discriminatory market with piecewise linear curves is $\mathcal{NP}$-Complete. In fact, this complexity jump occurs even for the simplest piecewise linear curves: *step functions.* The reduction is from the well-known **knapsack** problem [5], and applies even to the restricted case of **one seller and multiple buyers**.

We define a *step function* demand curve as a tuple (p_j, q_j), indicating a buyer's willingness to buy q_j units at or below the unit price p_j; the buyer is not willing to buy any units at price strictly greater than p_j.

Theorem 5.1 *Consider a discriminatory market where each participant submits a step function supply or demand curve. The problem of determining a surplus-maximizing clearing of the market is $\mathcal{NP}$-Complete. This holds even if one side of the market has only on one participant who submits a constant curve.*

Proof. The proof is given in the appendix. □

In the next subsection, we study the more restricted setting where price-quantity curves are **linear**. We show that in that setting even discriminatory markets can be cleared in polynomial time.

5.2 Fast Algorithms for Linear Curves

As we show, the discriminatory market clearing problem is a convex quadratic program with linear constraints, which could be solved in polynomial time using general techniques. However, we present fast ($O(N \log N)$) and simple specialized algorithms. In addition, our algorithms lend insight into the structure of the problem, such as simple closed-form expressions for clearing prices.

We begin with the simpler problems of one-sided markets, where either one seller is matched with multiple buyers, or one buyer is matched with multiple sellers. The former is commonly referred to as an *auction*, and the latter as a *reverse auction* (used widely for procurement or request-for-quotes in business-to-business commerce).

5.2.1 Reverse Auction: One Buyer, Multiple Sellers

Assume that n sellers bid in a reverse auction to sell multiple indistinguishable units of an item. There is a single buyer in the market who wants to acquire at least Q units of the object, at minimum total cost. Each of the sellers has an upward sloping supply curve: $q = a_i p - b_i$, where $a_i > 0$ and $b_i \geq 0$. Since the market can clear each seller at a different price, the problem can be written as

$$\min \sum_{i=1}^{n} p_i q_i \qquad \text{s.t.} \qquad q_i \;=\; a_i p_i - b_i \quad \text{and} \quad \sum_{i=1}^{n} q_i \;\leq\; Q$$

Eliminating p_i's from the objective yields the quadratic function $\sum q_i \left(\frac{q_i + b_i}{a_i}\right)$, leaving $\sum q_i \geq Q$ as the only constraint. Since the seller curves are upward sloping, it is clear that the quantity constraint is tight on optimality. Thus, we can use the method of Lagrangian multipliers to optimize

$$\min \quad \left(\frac{q_i^2}{a_i} + \frac{b_i q_i}{a_i}\right) \;+\; \lambda(Q - \sum_{i=1}^{n} q_i).$$

Setting each partial derivative with respect to q_i to zero gives $2q_i = a_i\lambda - b_i$. Since $\sum_{i=1}^{n} q_i = Q$, we get

$$\lambda \;=\; \frac{2Q + \sum b_i}{\sum a_i} \tag{1}$$

Substituting the value of λ into q_i yields the clearing quantities and prices:

$$q_i \;=\; -\frac{b_i}{2} + \frac{a_i}{2}\left(\frac{2Q + \sum b_i}{\sum a_i}\right) \qquad\qquad p_i \;=\; \frac{b_i}{2a_i} + \frac{1}{2}\left(\frac{2Q + \sum b_i}{\sum a_i}\right) \tag{2}$$

There is only one difficulty in this solution: some of the quantities q_i may be negative (because we ignored the non-negativity constraint from the mathematical program). In particular, if the quantity $a_i\lambda$ is smaller than b_i, then $q_i < 0$. Such a solution could easily arise even in the case of two sellers, where it might be advantageous to *buy some extra units from one, and sell to the other the excess over* Q. For a simple example, consider two sellers, with curves $q = 100p - 100$ and $q = p - 100$, and suppose $Q = 50$. In this case, the Lagrangian method gives $q_1 = 98.5$ and $q_2 = -48.5$, which is clearly infeasible. We show below how to control the Lagrangian solution to keep it feasible. Basically, we show that sellers whose clearing quantity in Eq. (2) is negative must sell zero quantity in any optimal solution. Our algorithm for clearing a discriminatory price one-sided market with one buyer and many sellers can be described as follows:

Algorithm D-ReverseAuction

1. Index the sellers by increasing value of their smallest feasible price, namely, b_i/a_i. Let S_i denote the set of sellers $\{1, 2, \ldots, i\}$.
2. For $i = 1, 2, \ldots, n$ **do**
   - Compute clearing prices and quantities for the set S_i given by Eq. (2).
   - If **any** $q_j < 0$, **terminate**, and output the clearing computed for set S_{i-1}.
   - If the **maximum** clearing price among all sellers is less than b_{i+1}/a_{i+1}, or if $i = n$, **terminate** and output the solution.

Of course, if none of the clearings involve negative quantities, then the Lagrangian method ensures the correctness of the solution. If all sellers in S_i clear for less than b_{i+1}/a_{i+1}, which is the minimum feasible price for $i+1$, then we can obviously terminate. What remains to be shown is that as soon as some quantity becomes negative, say, for the set S_j, we can disregard the sellers j through n.

Lemma 5.2 *Suppose in the algorithm* **D-ReverseAuction**, *the first occurrence of a negative clearing quantity is for S_j, then all the sellers $j, j+1, \ldots, n$ sell zero quantity in an optimal solution of the one-sided single-buyer, multi-seller market.*

PROOF. Let j be the smallest index for which some quantity in the Lagrangian solution is negative. It is in fact easy to show that the negative quantity is q_j, the quantity of seller j. Let λ_j and λ_{j-1} denote the values of the Lagrangian multipliers for S_j and S_{j-1}, respectively. Since $q_j = \frac{1}{2}(a_j\lambda_j - b_j) < 0$, we obtain $\lambda_j < \frac{b_j}{a_j}$. This implies that the *minimum feasible price* of seller j is strictly larger than λ_j. Since $\lambda_j = \frac{2Q+\sum_{i=1}^{j} b_i}{\sum_{i=1}^{j} a_i}$, and $\lambda_{j-1} = \frac{2Q+\sum_{i=1}^{j-1} b_i}{\sum_{i=1}^{j-1} a_i}$, simple algebra shows that

$$\lambda_{j-1} < \frac{b_j}{a_j} \tag{3}$$

We now consider the *aggregate* supply curve of the sellers in S_{j-1}.[2] The aggregate linear function has the equation $q = \sum_{i=1}^{j-1} a_i\, p - \sum_{i=1}^{j-1} b_i$. The clearing price for $q = Q$ units is

$$\frac{Q + \sum_{i=1}^{j-1} b_i}{\sum_{i=1}^{j-1} a_i} < \frac{2Q + \sum_{i=1}^{j-1} b_i}{\sum_{i=1}^{j-1} a_i} < \frac{b_j}{a_j}.$$

That is, the minimum unit price for seller j is strictly larger than the *non-discriminatory* price at which Q units can be bought from the first $j-1$ sellers. Thus, using discriminatory price clearing also, it never helps to buy from the seller j. This completes the proof. □

Finally, the algorithm D-ReverseAuction can be implemented to run in $O(n \log n)$ time, for n sellers, as follows. The key is to maintain the Lagrangian multiplier λ_i at each iteration of the algorithm, which takes $O(1)$ time to update. Since the sellers are sorted in increasing order of b_i/a_i, it is clear that the highest clearing price belongs to the most recently added seller i. Similarly, since we only need to check if i's quantity is negative, we only need to compute q_i in round i of the algorithm, which takes $O(1)$ time. Thus, the running time of the algorithm is dominated by the initial sorting.

Theorem 5.3 *Consider a single-sided market with one buyer and n sellers, where each seller has an upward sloping linear supply curve. We can determine the minimum cost discriminatory price clearing for buying Q units of the object in total time $O(n \log n)$.*

[2]Since the maximum clearing price for S_{j-2} is at least b_{j-1}/a_{j-1}, it follows that buying Q units in the aggregate involves all $j-1$ sellers.

5.2.2 Auction: One Seller, Multiple Buyers

A similar solution holds when the market has one seller, with Q units to sell, and many buyers.[3] In this case, the goal is to maximize the revenue for the seller. Let the (downward sloping) demand curve of the jth buyer be $q = -a_j p + b_j$, for $j = 1, 2, \ldots, n$, where $a_j > 0, b_j \geq 0$. The *unconstrained* solution for the market is to sell exactly $\frac{1}{2} b_j$ units to buyer j. However, if the total number of units available is insufficient, that is, $Q < \frac{1}{2} \sum_j b_j$, then we solve the problem using the Lagrangian multiplier method, and obtain the following clearing quantities and prices:

$$q_i = \frac{b_i}{2} - \frac{a_i}{2} \left(\frac{\sum b_i - 2Q}{\sum a_i} \right) \qquad p_i = \frac{b_i}{2a_i} + \frac{1}{2} \left(\frac{\sum b_i - 2Q}{\sum a_i} \right) \tag{4}$$

In other words, if the unconstrained solution is quantity-infeasible, we increase the price for each buyer **by the same amount** until the demand reduces to Q. In increasing the price, if some buyer's curve reaches a point of infeasibility (its demand goes to zero), then we remove that buyer from the set, and recalculate the Lagrangian multiplier.

Theorem 5.4 *Consider a singled-sided market with one seller and n buyers, where each buyer has a downward sloping linear demand curve. We can determine the maximum-revenue discriminatory-price clearing for selling (at most) Q units of the object in total time $O(n \log n)$.*

5.2.3 Exchange: Multiple Sellers, Multiple Buyers

We now describe how to clear discriminatory-price many-to-many exchanges using single-sided markets (auctions and reverse auctions) as building blocks. Due to lack of space, we only sketch the main ideas here; details will appear in the full paper. We plot the "aggregate quantity vs. aggregate revenue" curve for both the sellers and the buyers. On the demand side, let $D(q)$ denote the maximum revenue achieved by selling exactly q units to the buyers and, on the supply side, let $S(p)$ denote the minimum cost of procuring q units from the sellers. Let us assume that both the sellers and the buyers are sorted in increasing order of $\frac{b_i}{a_i}$ — these are the roots of the supply and demand curves, that is, the *maximum* feasible prices for buyers, and the *minimum* feasible prices for sellers.

For the D curve, the starting point is the quantity $q_{\max} = \sum_{j=1}^{m} b_j/2$, which is the optimal quantity sold when quantity is *unconstrained*, as was discussed in Section 5.2.2. The corresponding revenue is $D(q_{\max}) = \sum \frac{b_j^2}{4a_j}$. To determine the aggregate revenue as we decrease the aggregate demand, we use the clearing expressions in Eq. (4). So, the clearing price of each buyer is uniformly increased by $\frac{1}{2}\lambda_0$, where $\lambda_0 = \frac{1}{2} \left(\frac{\sum_{j=1}^{m} b_j - 2q}{\sum_{j=1}^{m} a_j} \right)$ is a function of the quantity q being sold. Consequently, the revenue $D(q)$ is a *quadratic* function of q.

The parameter λ changes when the first buyer's price reaches the upper bound (his quantity goes to zero); note that the first buyer is the one with the smallest b_j/a_j term. We then update λ to the new value $\lambda_1 = \frac{1}{2} \left(\frac{\sum_{j=2}^{m} b_j - 2q}{\sum_{j=2}^{m} a_j} \right)$, and recompute the quadratic revenue function, and so on. Thus, $D(q)$ consists of m quadratic pieces, starts at quantity $q_{\max}$ and ends at the origin. See Figure 3.

Similarly, we determine the aggregate supply curve $S(q)$, which starts at the origin, and ends at $q_{\max}$. We determine the maximum quantity that can be purchased from the first seller subject to the price being at most equal to the root of the second seller's supply curve which is b_2/a_2 (below this root only the first seller's curve is active). This quantity is determined using Equations (2), and the revenue is obtained by multiplying the quantity by the price. The first piece of $S(q)$ is a quadratic curve that starts from the origin and ends at this point. The curve is obtained from Equations (2). The second piece of $S(q)$ extends from the end point of the first to the quantity that can be purchased using the

[3]We already studied this case in a recent paper [18]. We summarize the key results here because we will use them next in this paper as a component for deriving the algorithm for clearing discriminatory multi-seller, multi-buyer markets.

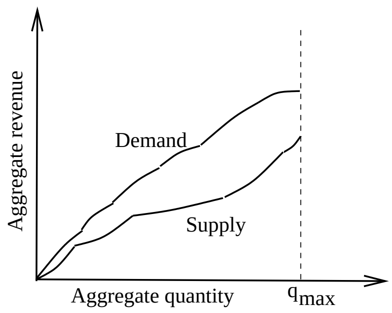

Figure 3: Discriminatory exchange clearing.

first *two* sellers at maximum price b_3/a_3 (where the third seller's curve enters), and so on. We stop when the quantity purchased reaches $q_{\max}$.

Once we have $D(q)$ and $S(q)$, the former with m pieces and the latter with n pieces, the maximum surplus is the maximum *vertical* distance between them, which can be computed easily in $O(n+m)$ time. Computing the curves $D(q)$ and $S(q)$ takes $O(m \log m)$ and $O(n \log n)$ time, respectively. The following theorem summarizes our result.

Theorem 5.5 *A multi-buyer, multi-seller discriminatory-price exchange with linear supply/demand curves can be cleared in $O(N \log N)$ time, where N is the number of participants.*

6 Conclusions and Future Research

We studied optimal clearing of markets in the ubiquitous setting where there are multiple indistinguishable units of an item for sale. The sellers and buyers express their bids via supply/demand curves. We focused on the natural and classical economic setting where each seller's supply increases and each buyer's demand decreases as the price increases (otherwise, absurd outcomes can occur [18]). Discriminatory pricing leads to greater (or equal) economic surplus for the market than non-discriminatory pricing. However, this comes at the cost of computational complexity.

We first studied the expressive case where the supply/demand curves are piecewise linear and not necessarily continuous (such curves can approximate any curve arbitrarily closely). We presented a polynomial-time algorithm for clearing non-discriminatory markets, and showed that clearing discriminatory markets is $\mathcal{NP}$-Complete (even in the setting where there is only one player on one side of the market, that player has a constant supply (demand) curve, and the players on the other side of the market have step function demand (supply) curves).

We then showed that in the more restricted setting of linear supply/demand curves, even discriminatory markets can be cleared in polynomial time. Our derivations also uncovered the following interesting fact. To obtain the optimal discriminatory solution, each buyer's (seller's) price is incremented (decremented) *equally* from her individual quantity-unconstrained solution. As this price change proceeds, some players' quantities go to zero, at which point those players are excluded from the market. There is no need to consider re-entry of the excluded players' curves in the process.

Future research includes studying other restrictions besides linearity that may allow polynomial-time clearing of discriminatory markets, such as continuity assumptions and conditions on second (or higher) derivatives of the supply/demand curves. Also, the question of polynomial-time clearability under other types of supply/demand curves besides piecewise linear ones remains open for both discriminatory and non-discriminatory markets. Finally, we plan to design incremental clearing algorithms for these settings.

References

[1] Ian Domowitz. Automating the continuous double auction in practice: Automated trade execution systems in financial markets. In Daniel Friedman and John Rust, editors, *The Double Auction Market*, volume 14 of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 27–60. Addison-Wesley, 1993.

[2] Joan Feigenbaum, Christos Papadimitriou, and Scott Shenker. Sharing the cost of muliticast transmissions (preliminary version). In *ACM Symposium on Theory of Computing (STOC)*, pages 218–227, 2000.

[3] D Ferguson, C Nikolaou, and Y Yemini. An economy for flow control in computer networks. In *IEEE INFOCOM*, pages 110–118, 1989.

[4] Yuzo Fujishima, Kevin Leyton-Brown, and Yoav Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 548–553, Stockholm, Sweden, August 1999.

[5] Michael R Garey and David S Johnson. *Computers and Intractability.* W. H. Freeman and Company, 1979.

[6] Andrew Goldberg, J Hartline, and A Wright. Competitive auctions and digital goods. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Washington, DC, 2001.

[7] R Karp, E Koutsoupias, C Papadimitriou, and S Shenker. Optimization problems in congestion control. In *FOCS*, 2000.

[8] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 233–241, Minneapolis, MN, 2000.

[9] Daniel Lehmann, Lidian Ita O'Callaghan, and Yoav Shoham. Truth revelation in rapid, approximately efficient combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 96–102, Denver, CO, November 1999.

[10] William A Lupien and John T Rickard. Crossing network utilizing optimal mutual satisfaction density profile. US Patent 5,689,652, granted Nov. 18 to Optimark Technologies, 1997.

[11] Andreu Mas-Colell, Michael Whinston, and Jerry R. Green. *Microeconomic Theory.* Oxford University Press, 1995.

[12] Noam Nisan. Bidding and allocation in combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 1–12, Minneapolis, MN, 2000.

[13] Noam Nisan and Amir Ronen. Algorithmic mechanism design. In *STOC*, pages 129–140, 1999.

[14] Ori Regev and Noam Nisan. The POPCORN market - an online market for computational resources. In *Proc. 1st International Conference on Information and Computation Economies*, pages 148–157, Charleston, SC, October 1998.

[15] Tuomas Sandholm. An algorithm for optimal winner determination in combinatorial auctions. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 542–547, Stockholm, Sweden, 1999.

[16] Tuomas Sandholm. eMediator: A next generation electronic commerce server. In *Proceedings of the Fourth International Conference on Autonomous Agents (AGENTS)*, pages 73–96, Barcelona, Spain, June 2000.

[17] Tuomas Sandholm and Subhash Suri. Improved algorithms for optimal winner determination in combinatorial auctions and generalizations. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 90–97, Austin, TX, 2000.

[18] Tuomas Sandholm and Subhash Suri. Market clearability. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI)*, Seattle, WA, 2001.

[19] Carl A Waldspurger, Tad Hogg, Bernardo Huberman, Jeffrey O Kephart, and W Scott Stornetta. Spawn: A distributed computational economy. *IEEE Transactions on Software Engineering*, 18(2):103–117, 1992.

[20] R Walsky, J Plank, J Brevick, and T Bryan. G-Commerce: Market formulations controlling resource allocation on the computational grid. In *IPDPS*, 2001.

[21] Peter R Wurman, Michael P Wellman, and William E Walsh. The Michigan Internet AuctionBot: A configurable auction server for human and software agents. In *Proceedings of the Second International Conference on Autonomous Agents (AGENTS)*, pages 301–308, Minneapolis/St. Paul, MN, May 1998.

A Proof of Lemma 3.1

PROOF. Consider a trade between the buyer and seller at quantity q. The clearing prices for the seller and buyer are determined from their curves: $p_s = (q + b_s)/a_s$ and $p_d = (-q + b_d)/a_d$. The total surplus is

$$q\left(\frac{-q+b_d}{a_d} - \frac{q+b_s}{a_s}\right)$$

Setting the first derivative of the surplus (with respect to q) to zero, we get

$$-\frac{2q}{a_d} + \frac{b_d}{a_d} - \frac{2q}{a_s} - \frac{b_s}{a_s} = 0,$$

which gives the surplus-maximizing quantity

$$q = q^* = \frac{1}{2}\left(\frac{a_s b_d - a_d b_s}{a_s + a_d}\right)$$

(The second derivative of the surplus function is negative, implying that the surplus is maximized at this quantity.) The seller and buyer clearing prices, namely, p^*_{ask} and p^*_{bid} are easily obtained by substituting q into the expressions for p_s and p_d. □

B Proof of Theorem 5.1 and Remarks

PROOF. We reduce the **knapsack** problem to our market clearing problem, where there is one seller and multiple buyers. Let $\{(s_1, v_1), (s_2, v_2), \ldots, (s_n, v_n),\ Z\}$ be an instance of the knapsack problem—Z is the knapsack capacity, s_i and v_i, respectively, are the size and value of item i. The goal is to choose a subset of items of maximum value with total size at most Z. We create an instance of the discriminatory price market, using step function demand curves, as follows.

The seller has Z units of the goods, and we can scale the prices so that his step function bid is $(0, Z)$—that is, the seller can sell Z units at price 0 or higher, but no more than Z units at any price. The buyer i places a step function bid $(v_i/s_i, s_i)$, meaning he is willing to buy s_i units at lot price v_i (or maximum unit price v_i/s_i), and no units for a higher price. See Figure 4. Since we are using discriminatory pricing, the goal is to choose a subset of buyer bids maximizing the total revenue subject to the quantity constraint Z, which is easily seen to be equivalent to a solution of the knapsack. □

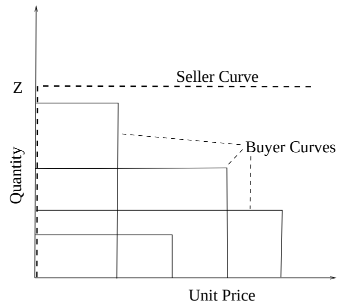

Figure 4: Reduction from knapsack.

Remark: We can modify the construction in the preceding theorem so that it holds even when the demand curves are *continuous and downward concave*, and the supply curves are *continuous and upward convex*. Specifically, if a demand curve is the step function (p, q), then we modify it into a continuous, piecewise linear, downward concave function as follows: the first piece starts at $(0, q + \varepsilon)$ and ends at (p, q); the second piece starts at (p, q) and ends at $(p + \varepsilon', 0)$, for some arbitrarily small $\varepsilon, \varepsilon' > 0$. Finally, we scale all the numbers up so that all numbers become integral. Thus, we conclude that if the price-quantity curves have even one breakpoint per curve, market clearing becomes intractable with discriminatory pricing.