

Optimal Clearing of Supply/Demand Curves*

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

Markets are important coordination mechanisms for multi-agent systems, and market clearing has become a key application area of AI algorithms. We study optimal clearing in the ubiquitous setting where there are multiple indistinguishable units for sale. The sellers and buyers express their bids via supply/demand curves. Discriminatory pricing leads to greater profit for the party who runs the market than non-discriminatory pricing. We show that this comes at the cost of computation complexity. For piecewise linear curves we present a fast polynomial-time algorithm for nondiscriminatory clearing, and show that discriminatory clearing is \mathcal{NP} -complete (even in a very special case). We then show that in the more restricted setting of linear curves, even discriminatory markets can be cleared fast in polynomial time. Our derivations also uncover the elegant fact that to obtain the optimal discriminatory solution, each buyer's (seller's) price is incremented (decremented) *equally* from that agent's price in the quantity-unconstrained solution.

Introduction

Commerce is moving online to an increasing extent, and there has been a significant shift to dynamic pricing via auctions (one seller, multiple buyers), reverse auctions (one buyer, multiple sellers), and exchanges (multiple buyers, multiple sellers). These market types have also become key coordination methods in multiagent systems. These trends have led to an increasing need for fast market clearing algorithms. Also, recent electronic commerce server prototypes such as *eMediator* (Sandholm 2002b) and *AuctionBot* (Wurman, Wellman, & Walsh 1998) have demonstrated a wide variety of new market designs, leading to the need for new clearing algorithms.

There has been a recent surge of AI interest (e.g., (Sandholm 2002a; Fujishima, Leyton-Brown, & Shoham 1999; Hoos & Boutilier 2000)) in clearing combinatorial auctions where bids can be submitted on bundles of distinguishable items, potentially multiple units of each (Sandholm 2002b; Leyton-Brown, Tennenholtz, & Shoham 2000; Sandholm *et al.* 2002). There has also been recent work on clearing combinatorial reverse auctions (Sandholm 2002b; Sandholm *et*

al. 2002) and combinatorial exchanges (Sandholm 2002b; Sandholm *et al.* 2002). The clearing problem in a combinatorial market is \mathcal{NP} -complete (Rothkopf, Pekeć, & Harstad 1998), inapproximable (Sandholm 2002a), and in certain variants even finding a feasible solution is \mathcal{NP} -complete (Sandholm *et al.* 2002). On the other hand, markets where there is only one unit of one item for sale are trivial to clear.

In this paper we study a setting which is in between. We study the ubiquitous market setting where there are multiple *indistinguishable* units of an item for sale. This setting is common in markets for stocks, bonds, electricity, bandwidth (e.g., RateXChange), oil, pork bellies, memory chips, CPU time, etc. We study the problem where the bids are known up front. This is the case in many business-to-business markets. The algorithms can also be used in markets where bids arrive over time (and the market is cleared periodically, for example, every 5 minutes, or after some number of bids have arrived since the last clearing) and in multi-stage markets (where a tentative clearing is carried out after each round of bidding).

The naive approach to bidding in a multi-unit market would require the bidders to express their offers as a list of points, for example (\$2 for 1 unit) XOR (\$5 for 2 units) XOR (\$6 for 3 units), etc. The mapping from quantities to prices can be represented more compactly by allowing each bidder to express his offer as a price-quantity curve (supply curve for a seller, demand curve for a buyer). Such curves are natural ways of expressing preferences, are ubiquitous in economics (Mas-Colell, Whinston, & Green 1995), and are becoming common in electronic commerce as well (Sandholm 2002b; Lavi & Nisan 2000; Lupien & Rickard 1997).

In classic economic theory of supply and demand curves (called partial equilibrium theory (Mas-Colell, Whinston, & Green 1995)), the market is cleared as follows. First, the supply curves of the sellers are aggregated, and the demand curves of the buyers are aggregated. Then, the market is cleared at some per-unit price for which supply equals demand (there may be multiple solutions). This way of clearing the market maximizes social welfare.

However, it turns out that the *auctioneer* (that is, the party who runs the market—who is neither a buyer nor a seller) will achieve greater (or equal) profit from the same supply/demand curves by reducing the number of units traded,

*This work was funded by, and conducted at, CombineNet, Inc., 311 S. Craig St., Pittsburgh, PA 15213.
Copyright © 2002, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

and charging one per-unit price to the buyers while paying a lower per-unit price to the sellers.¹ We call such pricing *non-discriminatory* because each buyer pays the same amount per unit, and each seller gets paid the same amount per unit. The auctioneer’s profit can be further improved by moving to *discriminatory* pricing where each seller and each buyer can be cleared at a different per-unit price.

Interestingly, the pricing scheme and the shape of the supply/demand curves significantly impact the computational complexity of clearing the market. We show that markets with piecewise linear curves are clearable in polynomial time under non-discriminatory pricing, but \mathcal{NP} -complete to clear under discriminatory pricing. With linear curves, even discriminatory markets can be cleared in polynomial time.²

The Market Model

The market has 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 him as two separate agents).

Each seller expresses his willingness to sell via a *supply curve* $s : R^+ \rightarrow R^+$ from non-negative unit prices to non-negative supply quantity. Thus, if the unit price is p , the seller is willing to supply $s(p)$ units of the good.³ Similarly, each buyer submits a *demand curve* $d : R^+ \rightarrow R^+$.

We assume that supply/demand curves are *piecewise linear*. Such curves can approximate any curve arbitrarily closely. General supply/demand curves, however, can lead to absurd results (infinite values, zero costs, etc. (Sandholm & Suri 2001a)). Therefore, we make the usual assumption 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. We do not assume that the curves are continuous—they could have discrete “jumps”.

In this paper, we study clearing where the objective is to maximize the auctioneer’s profit. We study two pricing schemes: *non-discriminatory* and *discriminatory*.

Non-Discriminatory Pricing

In a non-discriminatory market, there are two clearing prices, one shared by the sellers and one shared by the buyers. 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 supply/demand curves. In particular, suppose s_i is the supply curve of seller i , and d_j is the demand curve

¹The profit could be allocated entirely to the party who runs the market, or it could 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.

²These issues have been settled for auctions (and part of reverse auctions) (Sandholm & Suri 2001a). We settle the general case: multiple buyers and sellers.

³The 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.

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 is to maximize the auctioneer’s profit: $p_{bid}^* \sum_{j=1}^m d_j(p_{bid}^*) - p_{ask}^* \sum_{i=1}^n s_i(p_{ask}^*)$. Thus the computational problem is to determine the clearing prices p_{ask}^* and p_{bid}^* .

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 $\sum_{i=1}^n s_i(p_i^*) = \sum_{j=1}^m d_j(p_j^*)$, and the auctioneer’s profit to be maximized is $\sum_{j=1}^m p_j^* d_j(p_j^*) - \sum_{i=1}^n p_i^* s_i(p_i^*)$. The computational problem is to determine the clearing price for each seller and buyer.

The profit generated under discriminatory pricing is greater (or equal) than that under non-discriminatory pricing.⁴

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 1-seller, 1-buyer market, non-discriminatory and discriminatory pricing are identical.) Figure 1 illustrates the clearing in this setting.




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

⁴However, non-discriminatory markets offer fairness. Discriminatory markets 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.

Lemma 1 (1-Seller, 1-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 profit-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

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

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 profit is $q \left(\frac{-q+b_d}{a_d} - \frac{q+b_s}{a_s} \right)$. Setting the first derivative of the profit (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 profit-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 profit function is negative, implying that the profit is maximized at this quantity.) The seller and buyer clearing prices, namely, p_{ask}^* and p_{bid}^* are obtained by substituting q into the expressions for p_s and p_d . \square

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 Lemma 1. 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 profit achieved when trade occurs at some quantity $q^* + \varepsilon$, where ε can be positive or negative. (We assume $|\varepsilon| \leq q^*$, which ensures that the trade is feasible and produces non-negative profit.)

Lemma 2 In the 1-seller, 1-buyer market, if the trade occurs at quantity $q^* + \varepsilon$, then the profit is

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

That is, the profit 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 profit, therefore, equals $(q^* + \varepsilon) \left((p_{bid}^* - \frac{\varepsilon}{a_d}) - (p_{ask}^* + \frac{\varepsilon}{a_s}) \right)$, which can be written as $q^*(p_{bid}^* - p_{ask}^*) - \varepsilon^2 \left(\frac{1}{a_s} + \frac{1}{a_d} \right) + \varepsilon(p_{bid}^* - p_{ask}^*) - \varepsilon q^* \left(\frac{1}{a_s} + \frac{1}{a_d} \right)$. 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 1.

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

Thus, the new profit is $q^*(p_{bid}^* - p_{ask}^*) - \varepsilon^2 \left(\frac{1}{a_s} + \frac{1}{a_d} \right)$. \square

Our next lemma states a corollary of Lemma 2 for the case where the buyer and seller curves are quantity-constrained. 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 $q = a_s p - b_s$, but restricts the quantity to the range $[q'_s, q''_s]$. What trade maximizes the profit? 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 (1-Seller, 1-Buyer Bounded Trading)

Consider the 1-seller, 1-buyer market, with linear supply/demand curves, where the buyer and the seller can only trade in the quantity range $[q', q'']$. The profit-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. If the unconstrained trade quantity q^* is in the feasible range, profit is maximized at q^* . Otherwise, Lemma 2 shows that the profit 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 $[q', q'']$. \square

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, under non-discriminatory pricing. Our approach is to *aggregate* the demand and the supply curves separately, and then reduce the problem to the 1-seller, 1-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.

Let us consider the aggregation of buyer curves; seller curves are handled in the same way. Consider a set of piecewise linear demand curves d_1, d_2, \dots, d_n . Their *aggregate curve* is a piecewise linear function $D : R^+ \rightarrow R^+$ such that $D(p)$ is the total demand at unit price p . That is, $D(p) = d_1(p) + d_2(p) + \dots + d_n(p)$, 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.

The breakpoints of D are the union of the breakpoints of the component curves—the aggregate demand curve changes only when one of the component curves changes. Thus, given a set of n piecewise linear curves each of which has at most k pieces, their aggregate curve D has at most nk breakpoints.

Given n piecewise linear curves, their aggregate is easily computed in $O(nk \log(nk))$ time, by a sweep-line algorithm as follows. Let k be the maximum number of pieces

in any curve. Let p_1, p_2, \dots, p_L , where $L \leq nk$, denote the breakpoints of all the component curves, in sorted order. We scan these breakpoints in right to left order (decreasing order of price), and determine the linear aggregate curve between two consecutive breakpoints. Initially, we compute the linear aggregate function in the range (p_L, ∞) , in $O(n)$ time. Next, as we move to the next breakpoint, at most one linear piece changes—one piece may end and another may begin. (If multiple curves begin or end at the same point, we can enforce an artificial order among those, and consider them one at a time.) We can update the linear aggregate by deleting the coefficients of the leaving curve and adding those of the entering curve, and so each update takes $O(1)$ time. Thus, the complete aggregate curve can be determined in time $O(nk)$, after an initial sorting cost of $O(nk \log(nk))$.

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 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 profit-maximizing market, all clearings occur at these *envelope* prices.

Our non-discriminatory clearing algorithm can now be described as follows (see Figure 2):




Figure 2: *Non-discriminatory market.* The figure shows decomposition of the feasible region into four trapezoids. Each trapezoid corresponds to a 1-seller, 1-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 1-seller, 1-buyer bounded trade (cf. Lemma 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-profit 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 1 *Consider a 1-item, multi-unit market with multiple sellers and buyers, where each seller (buyer) has an upward (downward) sloping piecewise linear curve. Then, a profit-maximizing clearing 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.*

Discriminatory Markets

We now study the complexity of discriminatory markets.

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 **knapsack** problem (Garey & Johnson 1979), 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 2 *Consider a discriminatory market where each participant submits a step function supply or demand curve. Determining a profit-maximizing clearing of the market is \mathcal{NP} -complete. This holds even if one side of the market has only one participant—who submits a constant curve.*

PROOF. We reduce the **knapsack** to our market clearing problem, where there is one seller and multiple buyers. Let $\{(s_1, v_1), (s_2, v_2), \dots, (s_n, v_n), Z\}$ 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 3.




Figure 3: Reduction from knapsack.

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. \square

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 curves have even one breakpoint per curve, market clearing becomes intractable with discriminatory pricing.

Fast Algorithms for Linear Curves

In this section we show that in the more restricted setting where supply/demand curves are **linear**, even discriminatory markets can be cleared in polynomial time. 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 closed-form expressions for prices.

We begin with the simpler problems: *reverse auctions* where one buyer is matched with multiple sellers, and *auctions* where one seller is matched with multiple buyers.

Reverse Auction: One Buyer, Multiple Sellers Say n sellers bid in a reverse auction to sell multiple indistinguishable units of an item. There is one buyer in the mar-

ket. He wants to acquire at least Q units, at minimum total cost. Each seller has an upward sloping supply curve: $q = a_i p - b_i$, where $a_i > 0$ and $b_i \geq 0$. The clearing problem is

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

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

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

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

$$\lambda = \frac{2Q + \sum b_i}{\sum a_i} \quad (1)$$

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

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

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.

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 $\{1, 2, \dots, i\}$.
2. For $i = 1, 2, \dots, n$ **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 4 *Suppose in the algorithm D-ReverseAuction, the first occurrence of a negative clearing quantity is for S_j , then all the sellers $j, j + 1, \dots, n$ sell zero quantity in an optimal solution of the one-sided 1-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

$$\lambda_{j-1} < \frac{b_j}{a_j} \quad (3)$$

We now consider the *aggregate* supply curve of the sellers in S_{j-1} .⁵ 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 . \square

Finally, 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 , 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 run time of the algorithm is dominated by the initial sorting. Put together, we have:

Theorem 3 *Consider a reverse auction with 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 in total time $O(n \log n)$.*

⁵Since the maximum clearing price for S_{j-2} is at least b_{j-1}/a_{j-1} , buying Q units in the aggregate involves all $j - 1$ sellers.

Auction: One Seller, Multiple Buyers A similar solution holds when the market has one seller, with Q units to sell, and many buyers.⁶ The goal is to maximize the seller's revenue. Let the (downward sloping) demand curve of the j th buyer be $q = -a_j p + b_j$, for $j = 1, 2, \dots, n$, 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:

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

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. To summarize:

Theorem 4 ((Sandholm & Suri 2001a)) *Consider an auction with n buyers who have downward sloping linear demand curves. We can determine the revenue-maximizing discriminatory-price clearing for selling (at most) Q units in time $O(n \log n)$.*

Exchange: Multiple Sellers, Multiple Buyers We now describe how to clear discriminatory-price exchanges using auctions and reverse auctions as building blocks. 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. On the supply side, let $S(p)$ denote the minimum cost of procuring q units from the sellers. Assume that both the sellers and 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 . 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); the first

⁶This case was recently solved (Sandholm & Suri 2001a). We summarize some of the key results here because we will use them as components for deriving the algorithm for clearing discriminatory exchanges.

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 4.




Figure 4: *Discriminatory exchange clearing.*

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 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 first two sellers at maximum price b_3/a_3 (where the 3rd 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 profit 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 *In a multi-buyer, multi-seller discriminatory-price exchange with linear supply/demand curves, a profit-maximizing clearing can be determined in $O(N \log N)$ time, where N is the number of participants.*

Conclusions and Future Research

We studied profit-maximizing clearing of markets in the ubiquitous setting where there are multiple indistinguishable units 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 (Sandholm & Suri 2001a)). Discriminatory pricing leads to greater (or equal) profit for the auctioneer than non-discriminatory pricing. However, we showed that this comes at the cost of computational complexity.

We first studied the case where the supply/demand curves are piecewise linear and not necessarily continuous (such curves can approximate any curve arbitrarily closely). We presented an $O(K \log K)$ algorithm for clearing non-discriminatory markets, and showed that clearing discriminatory markets is \mathcal{NP} -complete (even if 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 ($O(N \log N)$). Our derivations also uncovered the interesting fact that to obtain the optimal discriminatory solution, each buyer's (seller's) price is incremented (decremented) *equally* from her individual price in the 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 those 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. Yet another interesting question is the impact of side constraints on the complexity of clearing the market (Sandholm & Suri 2001b).

We also plan to design *incremental clearing algorithms* for these settings. Finally, we plan to design *online clearing algorithms* that do nearly as well, without knowing all the supply and demand in advance, as the best clearing algorithm in hindsight. (Online clearing of exchanges with single-unit bids has already been addressed (Blum, Sandholm, & Zinkevitch 2002).)

References

- Blum, A.; Sandholm, T.; and Zinkevitch, M. 2002. Online algorithms for market clearing. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 971–980.
- Fujishima, Y.; Leyton-Brown, K.; and Shoham, Y. 1999. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, 548–553.
- Garey, M. R., and Johnson, D. S. 1979. *Computers and Intractability*. W. H. Freeman and Company.
- Hoos, H., and Boutilier, C. 2000. Solving combinatorial auctions using stochastic local search. In *Proceedings of*

the National Conference on Artificial Intelligence (AAAI), 22–29.

Lavi, R., and Nisan, N. 2000. Competitive analysis of incentive compatible on-line auctions. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, 233–241.

Leyton-Brown, K.; Tennenholtz, M.; and Shoham, Y. 2000. An algorithm for multi-unit combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.

Lupien, W. A., and Rickard, J. T. 1997. Crossing network utilizing optimal mutual satisfaction density profile. US Patent 5,689,652, granted Nov. 18 to Optimark Technologies.

Mas-Colell, A.; Whinston, M.; and Green, J. R. 1995. *Microeconomic Theory*. Oxford University Press.

Rothkopf, M. H.; Pekeč, A.; and Harstad, R. M. 1998. Computationally manageable combinatorial auctions. *Management Science* 44(8):1131–1147.

Sandholm, T., and Suri, S. 2001a. Market clearability. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI)*, 1145–1151.

Sandholm, T., and Suri, S. 2001b. Side constraints and non-price attributes in markets. In *IJCAI-2001 Workshop on Distributed Constraint Reasoning*, 55–61.

Sandholm, T.; Suri, S.; Gilpin, A.; and Levine, D. 2002. Winner determination in combinatorial auction generalizations. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. Early version appeared at the AGENTS-01 Workshop on Agent-Based Approaches to B2B, pp. 35–41, Montreal, Canada, May 2001.

Sandholm, T. 2002a. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence* 135:1–54. First appeared as an invited talk at the First International Conference on Information and Computation Economies, Charleston, SC, Oct. 25–28, 1998. Extended version appeared as Washington Univ., Dept. of Computer Science, tech report WUCS-99-01, January 28th, 1999. Conference version appeared at the International Joint Conference on Artificial Intelligence (IJCAI), pp. 542–547, Stockholm, Sweden, 1999.

Sandholm, T. 2002b. eMediator: A next generation electronic commerce server. *Computational Intelligence*. Special issue on Agent Technology for Electronic Commerce. To appear. Early versions appeared in the Conference on Autonomous Agents (AGENTS-00), pp. 73–96, 2000; AAAI-99 Workshop on AI in Electronic Commerce, Orlando, FL, pp. 46–55, July 1999; and as a Washington University, St. Louis, Dept. of Computer Science technical report WU-CS-99-02, Jan. 1999.

Wurman, P. R.; Wellman, M. P.; and Walsh, W. E. 1998. The Michigan Internet AuctionBot: A configurable auction server for human and software agents. In *Proceedings of the Second International Conference on Autonomous Agents (AGENTS)*, 301–308.