

Advances in Theoretical Economics

Volume 3, Issue 1

2003

Article 5

Market Research and Market Design

Sandeep Baliga*

Rakesh Vohra†

*Northwestern University, baliga@nwu.edu

†Northwestern University, r-vohra@kellogg.northwestern.edu

Copyright ©2003 by the authors. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher, bepress. *Advances in Theoretical Economics* is one of *The B.E. Journals in Theoretical Economics*, produced by The Berkeley Electronic Press (bepress). <http://www.bepress.com/bejte>.

Market Research and Market Design

Sandeep Baliga and Rakesh Vohra

Abstract

We study trading models when the distribution of signals such as costs or values is not known to traders or the mechanism designer when the profit-maximizing trading procedure is designed. We present adaptive mechanisms that simultaneously elicit this information (market research) while maintaining incentive compatibility and maximizing profits when the set of traders is large (market design). First, we study a monopoly pricing model where neither the seller nor the buyers know the distribution of values. Second, we study a model with a broker intermediating trade between a large number of buyers and sellers with private information about their valuations and costs. We show that when the set of traders becomes large our adaptive mechanisms achieve the same expected profits for the monopolist and the broker as when the distribution of signals is common knowledge.

KEYWORDS: Market Design, Wilson Doctrine

1 Introduction

Suppose second-hand electronic equipment, say a certain brand of CD player, is for sale. Buyers and sellers of the CD player have different values for it and use a broker to intermediate trade. The broker makes a profit by charging a bid-ask spread. In a profit-maximizing mechanism, the spread depends on the distributions of buyers' and sellers' values as these determine the demand and supply curves.¹ Alternatively, the CD player manufacturer may sell new equipment directly to buyers. The firm must know the buyers' demand curve to calculate the profit-maximizing price. The demand curve depends on the distribution of buyers' values.² Therefore, in the standard paradigm, it is assumed that traders have private information or signals about their costs or values but the *distribution* of the signals is common knowledge among the traders and also the mechanism designer. As this distribution is altered, so is the optimal mechanism.

Wilson ([23] and [24])³ has argued that economic institutions should work well in a wide variety of settings and should be independent of the details of the environment since they may not be known when the mechanism is designed. For example, the product being sold might be new and innovative or demand and supply conditions may be changing over time. One way to deal with this difficulty is to conduct market research and use the results in the design of the trading mechanism. There are a number of problems with this approach. First, it may be time-consuming to do the research. Second, to the extent that the subjects in the research are potential participants in the subsequent market, they may have an incentive to misrepresent information. Third, the market research may be based on a small sample and hence may not accurately capture market conditions.

¹Profit-maximizing and/or efficient mechanisms have been studied by Myerson and Satterthwaite [20] and Gresik and Satterthwaite [11].

²Bulow and Roberts [4] building on the work of Myerson [19] study monopoly pricing with incomplete information.

³“Game theory has a great advantage in explicitly analyzing the consequences of trading rules that presumably are really common knowledge; it is deficient to the extent that it assumes other features to be common knowledge, such as one agent's probability assessment about another's preferences or information.

I foresee the progress of game theory as depending on successive reductions in the base of common knowledge required to conduct useful analyses of practical problems. Only by repeated weakening of the common knowledge assumptions will the theory approximate reality,” Wilson [24]

We display *adaptive mechanisms* that simultaneously conduct market research while determining which trades to execute and at what prices. The mechanisms respond to the information elicited from agents, do not depend ex ante on the knowledge of the distribution of signals and maximize profits when the set of traders is large. Both a broker or a firm selling directly to consumers could use the mechanisms we study to conduct market research and market design.

One way to relax the assumptions of the standard model is to require that the distribution of values be common knowledge amongst the agents only. For example, Caillaud and Robert [5] study a revenue-maximizing auction when the seller does not know the distribution of types but the bidders *do* know the details of the environment. But then, as in complete information implementation models (see Moore [18] for a survey), a mechanism can be designed where the bidders equilibrium behavior reveals the distribution at no extra cost to the seller. Such mechanisms rely on the distribution being common knowledge among the agents. An alternative has been to focus on mechanisms that do not require cross-reporting and are distribution free. The goal here is to show that these “sub-optimal” mechanisms are close to optimal as the number of agents increases. In this case one studies the asymptotic properties of a Bayesian equilibrium of the mechanism. For example, in a seminal paper, Wilson [23] shows that a Bayesian equilibrium of a large, sealed-bid, double auction is incentive-efficient in the sense of Holmstrom and Myerson [13]. In later work, Rustichini, Satterthwaite and Williams investigate the k -double auction and show that one of its Bayesian equilibria is asymptotically efficient. While the rules of the double-auction are distribution-free, Bayesian equilibrium relies on players knowing the distribution of values.⁴

Continuing in this vein, one could also require that the distribution is not common knowledge among the traders and replace Bayesian incentive compatibility by ex post incentive compatibility. For example, Maskin [14] and

⁴Gul and Postlewaite [12] take a mechanism design approach and show in a more general environment that a Bayesian equilibrium of their mechanism implements a nearly-efficient, incentive compatible, individually rational allocation when the economy is replicated sufficiently often. However, the rules of their mechanism as well as the solution concept they utilize rely on both the mechanism designer and the players knowing the distribution of states and types. In particular, the mechanism they study relies on determining the state of the world that is most likely given players’ reports of their types and then calculating and implementing a perturbed, competitive equilibrium for the artificial economy that has the same distribution of players’ types as this state.

Dasgupta and Maskin [7] study efficient auctions which are distribution-free in an interdependent-value environment (the research surveyed above deals with private values). In their paper neither the players nor the mechanism designer know the distribution of signals (but agents do know the value functions). They study implementation of ex-post efficient allocations while we study implementation of profit-maximizing ones. This distinction is key as ex-post efficiency is a distribution-free concept (the ex-post efficient allocation in an auction setting requires that the object be given to the player with the highest value) while the profit-maximizing allocation rules in our trading models necessarily depend on the distribution of types. When the seller has beliefs about the types, he can seek to maximize expected revenue with respect to these beliefs. For the resulting mechanism to be immune to the unknown beliefs of the agents, the mechanism has to be ex post incentive compatible. In particular, Myerson [19] shows that the profit-maximizing auction in the independent private values case is, say, a second-price auction with a reserve price, where the reserve price depends on the distribution of valuations. If the mechanism designer also does not know the distribution of types, it is not clear how to approach implementation of this allocation rule. Hence, recent work on distribution-free implementation, such as Bergemann and Morris [2], assumes that the rule being implemented is itself distribution-free.

When the set of traders is large and types are independent draws from the same distribution but the distribution is not known, we resolve this issue. For this case, we show how to elicit the distribution from agents while maintaining incentive compatibility and maximizing profits. Segal [21], in independent and contemporaneous work, has studied a model similar to the first monopoly pricing model we present below. He takes a Bayesian approach assuming that the designer has a prior over the distributions from which the values are drawn. We take a sampling approach and do not assume that the designer has a prior.

A third relaxation is to eliminate the assumption that the traders and the mechanism designer know that values are independent draws from an unknown distribution. Goldberg et al. [10] investigate such a model where a monopolist has a constant marginal cost of production (in Deshmukh et al. [8], their results are extended to the case of a broker intermediating trade between buyers and sellers with private information). To describe their results, we need some notation. Let D be the set of all (monotone) dominant strategy mechanisms and v a profile of bidder valuations. For each $d \in D$,

let $d(v)$ denote the revenue obtained when using mechanism d on the profile v . Let $b(v)$ denote the revenue achieved by some benchmark mechanism that has full knowledge of v . Their goal is to identify mechanisms $d \in D$ that make $\inf_v \frac{d(v)}{b(v)}$ as large as possible. For some choices of $b(v)$ they show that for every $d \in D$ there is a profile v such that $\frac{d(v)}{b(v)}$ can be made arbitrarily small. This is the case when the benchmark mechanism uses full knowledge of v to extract all the surplus or when it charges a uniform price to all buyers. For other choices of $b(v)$ they identify members of D for which $\inf_v \frac{d(v)}{b(v)}$ is a non-zero constant. This is the case when the benchmark mechanism is forced to sell at least two units for all profiles. As these results suggest, there is no natural benchmark against which to compare their dominant strategy mechanisms.

We present two models. The first is a simple monopoly pricing model in which buyers have private information about their valuations. Our second model is one where there are many buyers and sellers and they both have private information about their valuations and costs. A broker acts as an intermediary between the two sides of the market and maximizes profits. However, he does not know the distributions of values and costs. As buyers and sellers are drawn from different distributions their so-called virtual values⁵ have to be compared to determine whether trade should be allowed. Also, as neither the buyers nor the sellers know the distribution we cannot use Bayesian equilibrium as our solution concept and instead require that agents play dominant strategies. Finally, while the main focus of the paper is on implementing profit-maximizing rules, we will mention how efficient allocations can be implemented in a distribution-free manner in the settings we study.

The main idea we use is easy to explain. As the distribution of signals is unknown, the monopolist and the broker must elicit the distribution of values from the traders themselves. This is what we call *market research*. However, the traders have an incentive to misrepresent their information as it may affect the final price they pay. For example, Priceline executes any trade that has positive surplus given the announcements of a buyer and a seller, charges them the prices they announce and pockets any difference. This gives buyers, say, the incentive to lie about their valuations. Priceline cannot calculate virtual values as it does not know the distribution of signals. Moreover,

⁵Virtual values are related to marginal revenues (see Bulow and Roberts [4]).

as traders lie about their valuations, it is difficult to infer the distribution of valuations from their announcements. We deal with these problems by designing a procedure where no agent's report about his valuation is used to determine the price he pays should he buy or sell an object. This ensures incentive compatibility in dominant strategies and is an extension of the familiar idea used by Vickrey [22]. As all traders announce their information truthfully, we can use their reports to calculate estimated distributions and densities. When there are a large number of traders, the estimates are close to the truth. Finally, our mechanisms always leave zero surplus to buyers with the lowest possible valuation (and sellers with the highest possible cost) and, as the number of traders becomes large, trades are executed if and only if they would be implemented when the distributions are common knowledge. Therefore, our adaptive mechanisms maximize profits as if the distributions and densities are common knowledge when the mechanism is designed.

McLean and Postlewaite [16]⁶ propose an interesting definition of information smallness which is related to our approach: an agent is informationally small if his information does not change the probability assessment of a common value state by very much, given the signals of the other agents. In our model, traders are informationally small with respect to the estimation of the distribution of signals. However, McLean and Postlewaite's [16] analysis concerns the relationship between informational smallness, incentive compatibility and efficiency and is not formally related to our results.

2 Monopoly Pricing

Consider a monopolist with a constant marginal cost selling to a group of $B = \{1, 2, \dots, M\}$ buyers. Each buyer i is interested in at most one unit of the good and his value v_i for the good is an independent draw from a common distribution F with support $[\underline{v}, \bar{v}]$ and a strictly positive density f . The buyer alone knows his value and the seller's marginal cost is normalized to zero. This is the monopoly-pricing version of the canonical *symmetric, independent private values model* introduced by Myerson [19].

In the standard model, it is assumed that F is common knowledge among the buyers and the seller. In that case, a mechanism that maximizes the seller's expected profit per capita is to sell to any buyer who is willing to purchase at a price p^* chosen to maximize $p[1 - F(p)]$. To understand why

⁶See also Gul and Postlewaite [12].

this is the case, it is useful to think of bidder i 's "demand" for the good at price p as

$$x(p) = 1 - F(p)$$

as only types of bidder i with valuations above p purchase the object. Therefore, total revenue expressed as a function of price is

$$R(p) = p(1 - F(p))$$

and p^* maximizes total revenue and hence profits as marginal costs are normalized to zero.⁷

In this model, as there is constant marginal cost, the monopolist's profit-maximization problem is separable across buyers. Moreover, in the optimal mechanism, no buyer is affected by the purchasing decisions of any other buyer. Hence, each buyer has a dominant strategy to purchase if and only if his valuation is above the price offered by the seller. Indeed, the buyers do not need to know the distribution F to implement this strategy. However, it is necessary for the seller to know F to calculate the profit-maximizing price p^* .

We assume buyers' values are conditionally independent draws from some common distribution F but neither the buyers nor the seller know the distribution. It may be useful to imagine an ex ante stage where Nature picks the distribution from which the values are subsequently drawn. Therefore, we allow buyers' values to be correlated. However, we do not explicitly model the ex ante stage or the seller's beliefs over Nature's move. It is in this sense that our mechanism is distribution-free. Despite this, we show that, when the number of buyers is large, the seller can make the same expected profit per capita as when he knows the distribution. We estimate the distribution from the sample of announced values while maintaining incentive compatibility so the announcements are honest. The idea is to have each potential buyer bid to purchase one unit of the good and to set a separate price for each buyer as a function of the bids. This *function* is set before the auction. To describe the *adaptive monopoly pricing* procedure, let b_i be the bid submitted by buyer i . For each $x \in \mathfrak{R}$ and $i \in B$ let:

$$F_M^i(x) = \frac{|\{j \in B \setminus i : b_j \leq x\}|}{M - 1}.$$

⁷If marginal costs were constant at c , total profits are $(p - c)(1 - F(p))$.

The purchase price r_i for bidder i is defined as follows:

$$r_i \in \arg \max \{x(1 - F_M^i(x))\}.$$

Any buyer i whose bid b_i is greater than r_i is sold one unit of the good. Notice that the price a winning bidder pays is not a function of his bid. Therefore, a buyer, in the case he wins, has no incentive to lie as it does not change the amount he pays and can only reduce his payoff to zero if he bids so low that he does not win. A buyer, in the case he loses, also has no incentive to lie as he either does not affect the outcome by underbidding or potentially makes a loss by overbidding to the extent that he wins. Therefore, each buyer has a (weakly) dominant strategy to bid truthfully. This fact, in turn, implies that the buyers do not have to know F for the scheme to work. Moreover, standard arguments via the law of large numbers show that $F_M^i \rightarrow F$ and thus $r_i(1 - F_M^i(r_i)) \rightarrow p^*(1 - F(p^*))$ for all $i \in B$ as $M \rightarrow \infty$ (see Breiman page 283 for example).

Theorem 1 *As $M \rightarrow \infty$, adaptive monopoly pricing achieves the same expected profits per capita as the profit-maximizing selling procedure when the seller knows F .*

Remark 1 *While we have studied profit-maximization here, the efficient allocation can easily be implemented in our context: simply set the price equal to the seller's cost of production.*

3 Markets with a Broker

In the model above, the seller not only produced goods for sale but also designed the mechanism to maximize profits. We now suppose that there is a broker who acts as an intermediary between traders and maximizes his own profits. We first consider the standard model where the distribution over types is common knowledge.

Suppose there are $N = \tau N_0$ sellers each of whom owns an indivisible object that $M = \tau M_0$ buyers want to buy (by a slight abuse of notation we will also refer to the set of sellers and buyers by N and M respectively). Each buyer wants at most one unit of the object and each seller can sell at most the one unit he owns. We denote buyer i 's valuation by v_i and seller j 's valuation by c_j . Buyers' values are independently drawn from a distribution

F and sellers' costs are independently drawn from the distribution H . Both distributions have positive, bounded, differentiable densities, f and h , over the interval $[a, b]$, have bounded second derivatives and satisfy the *monotone hazard rate condition* (we subsequently drop this assumption):

$$\text{virtual values, } v_i - \frac{1 - F(v_i)}{f(v_i)}, \text{ and virtual costs, } c_j + \frac{H(c_j)}{h(c_j)}, \quad (1)$$

are nondecreasing over the interval (a, b) .

Let $v \equiv (v_1, \dots, v_M)$, $c \equiv (c_1, \dots, c_N)$, $v_{-i} \equiv (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_M)$ and $c_{-i} \equiv (c_1, \dots, c_{i-1}, c_{i+1}, \dots, c_N)$. The density $g(v, c) = \prod_{i=1}^M f(v_i) \cdot \prod_{j=1}^N h(c_j)$ is the joint density of all the valuations. We assume agents are risk neutral and have additively separable utility for money and the object. Each agent knows his own valuation but considers the distribution of others' valuations to be distributed as described above. Therefore, $g(v_{-i}, c) = g(v, c)/f(v_i)$ is the joint density of valuations buyer i faces and $g(v, c_{-i}) = g(v, c)/h(c_j)$ is the joint density of valuations seller j faces.

We invoke the Revelation Principle to investigate the properties of incentive compatible and individually rational allocation rules. That is, we will study a direct mechanism where buyers and sellers simultaneously and privately report their valuations into a mechanism which then determines the probabilities with which objects are traded and what payments are executed as a function of the profile of reports. Therefore, let $p_i^\tau(v, c)$ be the probability that buyer i is allocated an object when the profile of buyers' reports is v and the profile of seller reports is c . Similarly, let $q_j^\tau(v, c)$ be the probability that seller j does not make a sale for the profiles of reports v and c . Also, let $r_i^\tau(v, c)$ be the payment made to buyer i when the profile of buyers' reports is v and the profile of seller reports is c (a negative value is a payment made by buyer i). Similarly, let $s_j^\tau(v, c)$ be the payment made to seller j for the profiles of reports v and c . The collection $(\{p_i^\tau\}, \{q_j^\tau\}, \{r_i^\tau\}, \{s_j^\tau\})$ is a direct mechanism.

A mechanism must satisfy *market-clearing* and allocate all N objects to agents for all reports:

$$\sum_{i=1}^M p_i^\tau(v, c) + \sum_{j=1}^N q_j^\tau(v, c) = N$$

for all (v, c) .

To formalize further constraints on a direct mechanism, we define

$$\begin{aligned}\bar{p}_i^\tau(v_i) &= \int \dots \int p_i^\tau(v, c)g(v_{-i}, c)dv_{-i}dc, & \bar{r}_i^\tau(v_i) &= \int \dots \int r_i^\tau(v, c)g(v_{-i}, c)dv_{-i}dc, \\ \bar{q}_j^\tau(c_j) &= \int \dots \int q_j^\tau(v, c)g(v, c_{-j})dvdc_{-j}, & \bar{s}_j^\tau(c_j) &= \int \dots \int s_j^\tau(v, c)g(v, c_{-j})dvdc_{-j}, \\ U_i(v_i) &= \bar{r}_i^\tau(v_i) + v_i\bar{p}_i^\tau(v_i), & V_j(c_j) &= \bar{s}_j^\tau(c_j) - c_j(1 - \bar{q}_j^\tau(c_j)).\end{aligned}$$

Therefore, $U_i(v_i)$ is the expected gains from trade for a buyer with valuation v_i as $\bar{r}_i^\tau(v_i)$ is the expected payment made to him and $\bar{p}_i^\tau(v_i)$ is the probability he acquires an object. Similarly, $V_j(c_j)$ is the expected gains from trade of a seller with valuation c_j as $\bar{s}_j^\tau(c_j)$ is the expected payment made to him and $(1 - \bar{q}_j^\tau(c_j))$ is the probability he sells the object.

The mechanism $(\{p_i^\tau\}, \{q_j^\tau\}, \{r_i^\tau\}, \{s_j^\tau\})$ is *incentive compatible* if and only if, for all buyers, and every v_i and v'_i in $[a, b]$,

$$U_i(v_i) \geq \bar{r}_i^\tau(v'_i) + v_i\bar{p}_i^\tau(v'_i)$$

and all sellers and every c_j and c'_j in $[a, b]$,

$$V_j(c_j) \geq \bar{s}_j^\tau(c'_j) - c_j(1 - \bar{q}_j^\tau(c'_j)).$$

The mechanism $(\{p_i^\tau\}, \{q_j^\tau\}, \{r_i^\tau\}, \{s_j^\tau\})$ is *individually rational* if and only if, for all buyers, and every v_i in $[a, b]$,

$$U_i(v_i) \geq 0$$

and for all sellers, and every c_j in $[a, b]$,

$$V_j(c_j) \geq 0.$$

Finally, as the broker obtains the difference between expected payments to buyers and sellers, his expected profit is

$$U_0 \equiv - \sum_{i=1}^M \int \dots \int r_i^\tau(v, c)dvdc - \sum_{j=1}^N \int \dots \int s_j^\tau(v, c)dvdc.$$

Our first result in this section is familiar from the work of Myerson and

Satterthwaite [20] and Gresik and Satterthwaite [11] and we omit its proof:

Theorem 2 For any incentive compatible, individually rational mechanism $(\{p_i^\tau\}, \{q_j^\tau\}, \{r_i^\tau\}, \{s_j^\tau\})$ with a broker,

$$\begin{aligned} \bar{p}_i^\tau & \text{ is weakly increasing for all buyers,} \\ \bar{q}_j^\tau & \text{ is weakly increasing for all sellers} \end{aligned} \quad (2)$$

and

$$\begin{aligned} & U_0 + \sum_{i=1}^M U_i(a) + \sum_{j=1}^N V_j(b) \\ = & U_0 + \sum_{i=1}^M \min_{v \in [a,b]} U_i(v) + \sum_{j=1}^N \min_{c \in [a,b]} V_j(c) \\ = & \sum_{i=1}^M \int \dots \int \left(v_i - \frac{1 - F(v_i)}{f(v_i)} \right) p_i^\tau(v, c) g(v, c) dv dc \\ & - \sum_{j=1}^N \int \dots \int \left(c_j + \frac{H(c_j)}{h(c_j)} \right) (1 - q_j^\tau(v, c)) g(v, c) dv dc. \end{aligned}$$

From the Theorem, we see that the expected profit of the broker is given by

$$U_0 = \sum_{i=1}^M \int \dots \int \left(v_i - \frac{1 - F(v_i)}{f(v_i)} \right) p_i^\tau(v, c) g(v, c) dv dc \quad (3)$$

$$- \sum_{j=1}^N \int \dots \int \left(c_j + \frac{H(c_j)}{h(c_j)} \right) (1 - q_j^\tau(v, c)) g(v, c) dv dc \quad (4)$$

$$- \sum_{i=1}^M U_i(a) - \sum_{j=1}^N V_j(b).$$

Recall that

$$C_B(v) \equiv v - \frac{1 - F(v)}{f(v)} \quad (5)$$

is the virtual value of a buyer with valuation v and

$$C_S(c) \equiv c + \frac{H(c)}{h(c)} \quad (6)$$

is the virtual cost of a seller with valuation c . Since the monotone hazard rate condition is satisfied by F and H , virtual values and costs, (5) and (6) are increasing in values and costs respectively.

The procedure that maximizes (3) subject to individual rationality is the following: First, set $U_i(a) = V_j(b) = 0$ for all buyers and sellers. Execute trades with positive *virtual* surplus (i.e. when a virtual value is higher than a virtual cost), starting with those with the highest virtual surplus first. This solution satisfies the monotonicity condition (2) only if the virtual value function $C_B(v)$ and the virtual cost function $C_S(c)$ are nondecreasing in v and c respectively. The monotone hazard rate condition implies that this is indeed the case.

There are two difficulties posed by the broker's problem. First, there are externalities between traders at the optimal mechanism even when the distributions are common knowledge. For example, suppose a buyer does not acquire an object at the current profile of reports. If he increases his report, he may succeed in acquiring an object while eliminating another buyer from the set of successful bidders. These externalities are not present in the analogous monopoly problem with constant marginal cost. The externalities are only exasperated when the distributions of values and costs have to be estimated from the reports (after the proof of Theorem 3, we remark in more detail about the additional externalities that arise as a result of the estimation problem). Second, as buyers' and sellers' values are drawn from different distributions, virtual values and costs must be compared to determine who trades. From (5) and (6), notice that the determination of virtual values and costs requires knowledge of density and distribution functions. Therefore, both must be estimated in the trading procedure when the distributions are not known.

We assume buyers' and sellers' types are conditionally independent draws from distributions F and H respectively. The buyers, sellers and the broker do not know the distributions. We have the following result for this case:

Theorem 3 *Suppose the monotone hazard rate condition holds. As $\tau \rightarrow \infty$, the broker's expected profits per capita converge to his expected profits per*

capita when he knows F and H .

Proof: Again, we focus on estimating F and f from buyers reports. Split the buyers into two sets M_A and M_B of equal size (if there are an odd number of buyers, allocate an extra buyer to M_A). By an abuse of notation, we will also denote the cardinality of these two sets by M_A and M_B respectively. Let v_i be the report of buyer i . For each $v \in \mathfrak{R}$ let:

$$F_k(v) = \frac{|\{j \in M_l : v_j \leq v\}|}{M_l}.$$

where $k, l \in \{A, B\}$ with $k \neq l$. Notice that the report of buyer $i \in M_k$ does not affect the calculation of the estimated distribution function buyers in M_k . Now, we turn to the estimation of the density function. Let $M_l = n$ and suppose v_1, v_2, \dots, v_n are n i.i.d. draws from F over the interval $[a, b]$. That is, these are the reported valuations of all buyers in M_l . For any $[x, y] \subset [a, b]$ let $\mu_k(x, y) = |\{j : v_j \in [x, y], j \in M_l\}| = (F_k(y) - F_k(x))n$.

If $y - x$ is sufficiently small, one can approximate $f(v)$ for $v \in [x, y]$ by $\frac{F(y) - F(x)}{y - x}$. Since F is unknown, we can approximate F by F_k which suggests that we estimate the value of $f(v)$ by $\frac{\mu_k(x, y)}{n(y - x)}$. This estimate of $f(v)$ involves two approximations. The first, is to approximate $f(v)$ by $\frac{\int_x^y f(t) dt}{y - x}$ and the second to estimate $\int_x^y f(t) dt$ by $\frac{\mu_k(x, y)}{n}$. The first is a good approximation when $y - x$ is small. The second is a good approximation only when $[x, y]$ contains a large number of points, i.e., $y - x$ is large. The main difficulty is to find a trade-off between these two errors so as to produce a good estimate of f . This is the subject of a large literature on density estimation. We refer the reader to Luc Devroye [9] for an introduction to the literature.

Here we use what is called the “histogram” estimate of the density function. Choose m , growing with n , points x_1, x_2, \dots, x_m in $[a, b]$ and a number h_m such that

- $x_1 = a$
- $x_n = b$
- $x_j < x_{j+1}$
- $x_{j+1} - x_j = h_m$

The sequence h_m will go to zero as m grows at a rate to be chosen later. Set

$$f_k(v) = \frac{\mu_k(x_j, x_{j+1})}{nh_m}$$

if $v \in [x_j, x_{j+1}]$. First, observe that

$$\frac{\mu_k(x_j, x_{j+1})}{n} = F_k(v + x_{j+1} - v) - F_k(v - (v - x_j)).$$

For any t we have that $|F_k(t) - F(t)| \leq r_n$ almost surely where $r_n \rightarrow 0$ as $n \rightarrow \infty$ and r_n is independent of t . In fact $r_n = O(\sqrt{\frac{\lg \lg n}{n}})$.⁸ Thus we can bound $\frac{\mu_k(x_j, x_{j+1})}{n}$ above by

$$F(v + x_{j+1} - v) - F(v - (v - x_j)) + 2r_n \quad (7)$$

and below by

$$F(v + x_{j+1} - v) - F(v - (v - x_j)) - 2r_n.$$

Approximating (7) by F 's Taylor expansion yields:

$$\begin{aligned} & [F(v) + (x_{j+1} - v)f(v) + O(|v - x_{j+1}|^2)] - [F(v) - (v - x_j)f(v) + O(|v - x_j|^2)] + 2r_n \\ & \leq (x_{j+1} - x_j)f(v) + O(|v - x_{j+1}|^2) + 2r_n \end{aligned}$$

and a similar argument applies to the lower bound. Since f' is bounded, the constant factors in the remainder term are independent of v . Hence,

$$\frac{(x_{j+1} - x_j)f(v) - O(|v - x_j|^2) - 2r_n}{h_m} \leq f_k(v) \leq \frac{(x_{j+1} - x_j)f(v) + O(|v - x_{j+1}|^2) + 2r_n}{h_m}$$

Simplifying:

$$f(v) - O(h_m) - 2r_n/h_m \leq f_k(v) \leq f(v) + O(h_m) + 2r_n/h_m.$$

Choose $h_m \rightarrow 0$ so that $r_n/h_m \rightarrow 0$. Since r_n is $O(\sqrt{\frac{\lg \lg n}{n}})$ this can always be done. Since $h_m \rightarrow 0$ as n and m go to infinity it follows that $f_k(v) \rightarrow f(v)$ almost surely. In fact $\sup_{x \in [0,1]} |f_k(x) - f(x)| \rightarrow 0$. Furthermore for n and m

⁸See Chung [6].

sufficiently large, $f_k(v) > 0$ almost surely as $f(v) > 0$ for all v .

We note that neither the distribution function F_k nor the density function f_k estimated for buyers in M_k depend on their own reports. We estimate distribution functions H_k and density functions h_k , $k \in \{A, B\}$, for sellers in a similar fashion.

Let

$$C_B^k(v) \equiv v - \frac{1 - F_k(v)}{f_k(v)} \quad (8)$$

be the *estimated virtual value* of buyer $i \in M_k$ ($k \in \{A, B\}$) with valuation v and

$$C_S^k(c) \equiv c + \frac{H_k(c)}{h_k(c)} \quad (9)$$

be the *estimated virtual cost* of seller $j \in N_k$ ($k \in \{A, B\}$) with cost c . Both are well-defined since, for large M and N , f_k and h_k are strictly positive.

As the estimated virtual functions are not necessarily monotonic, we “flatten” them. Let the *estimated flattened virtual value* be

$$\bar{C}_B^k(v) = \max_{v' \in [a, v]} \{C_B^k(v')\} \quad k \in \{A, B\}.$$

The function $\bar{C}_B^k(\cdot)$ is non-decreasing by construction.

A similar procedure for the seller defines a nondecreasing *estimated flattened virtual cost* of sellers in set N_k :

$$\bar{C}_S^k(c).$$

In the trading procedure, two markets, market A and market B , are set up, with buyers in M_k allowed to trade with sellers in N_k ($k \in \{A, B\}$). We utilize reported values and costs to rank the buyers and sellers and the estimated flattened virtual functions to determine which trades to execute. First, in market k , label the buyers and sellers so buyer i has a higher reported value than buyer $i+1$ and seller j has a lower reported cost than seller $j+1$ (we ignore ties as they are a measure zero event). Second, the trading decisions are as follows:

$$\begin{aligned} p_1^{\bar{}}((v_1, \cdot), (c_1, \cdot)) &= 1 - q_1^{\bar{}}((v_1, \cdot), (c_1, \cdot)) = 1 \text{ if } \bar{C}_B^k(v_1) \geq \bar{C}_S^k(c_1) \\ &= 0 \text{ otherwise.} \end{aligned}$$

Continue in this fashion till either all the N_k units are traded or till we reach $q \in M_k$ such that

$$\bar{C}_B^k(v_q) \geq \bar{C}_S^k(c_q) \text{ and } \bar{C}_B^k(v_{q+1}) < \bar{C}_S^k(c_{q+1}).$$

Finally, it remains to construct the payment functions so the individual rationality constraints of the lowest type of buyers and the highest types of sellers bind and incentive compatibility holds. Suppose buyer $q \in M_k$ and seller $q \in N_k$ are the lowest buyer and highest seller to trade (recall the procedure by which we have re-labelled the traders). Choose v^* so that

$$v^* = \min\{v \mid \bar{C}_B^k(v) = \bar{C}_S^k(c_q)\}. \quad (10)$$

Notice that v^* does not depend on buyer i 's report as \bar{C}_B^k is estimated using the reports of buyers in market l . Moreover, as $v^* \geq \bar{C}_B^k(v^*) = \bar{C}_S^k(c_q) \geq c_q \geq a$, we must have $v^* \geq a$.

If buyer i wins an object and $\bar{C}_B^k(v_{q+1}) < \bar{C}_S^k(c_q)$, he pays

$$v^*;$$

if buyer i wins an object and $\bar{C}_B^k(v_{q+1}) \geq \bar{C}_S^k(c_q)$, he pays v_{q+1} . A buyer who does not trade does not make any payment.

We show that buyers have a (weakly) dominant strategy to report their valuation truthfully.

First, the price at which winning buyers execute their trade is independent of their own reports. Also, as the estimated flattened virtual value function is determined by buyers' reports from market l , a buyer in market k cannot affect the estimated flattened virtual values of other buyers in market k . Hence, it is a dominant strategy to tell the truth as by lying either a buyer does not change the outcome or loses an object he prefers to win.

Second, consider a buyer i in market k who does not win an object with positive probability as he truthfully reports a valuation $v_i < v_q$. Recall that estimated flattened virtual valuation function in market k is increasing and does not depend on the reports of buyers in market k . Therefore, by underbidding, buyer i does not win an object. If he overbids to the extent that he wins an object with positive probability, he makes a loss as $v_i < v_q$ and $\bar{C}_B^k(v_i) < \bar{C}_S^k(c_{q+1})$. Therefore, buyer i has a (weakly) dominant strategy to report his valuation truthfully.

Also, $U_i(a) = 0$ as the lowest price a buyer ever pays is a because v^* is

always greater than a and as v_{q+1} must be greater than a .

A similar procedure is used to construct the payment function to sellers and similar arguments shows that all sellers have a dominant strategy to report their costs truthfully and that $U_j(b) = 0$.

Finally, as τ goes to infinity, the estimated flattened virtual functions converge to the true ones. Therefore, expected profits per capita converge to those when the distributions are known. ■

Remark 2 *Our analysis of monopoly pricing suggests an estimation procedure where each agent's virtual valuation function is estimated by dropping only his own report. But then, for example, winning bidders may be able to deviate from truthtelling, alter which buyers get to trade and thereby lower the price at which they purchase a good. Suppose buyers q and $q-1$'s reports lie in the same histogram. If buyer $q-1$ increases his report, he lowers the estimated density used to calculate the buyer q 's virtual valuation and this may in turn lower the latter's virtual valuation to the extent that the mechanism does not allow him to trade. This can in turn imply that the price buyer $q-1$ pays is determined by seller $q-1$ and this is lower than the price he pays if he tells the truth and buyer and seller q trade. Hence, this estimation procedure is not ex post incentive compatible.*

Remark 3 *An alternative procedure we could have used is to prevent a proportion εM of the buyers and εN of the sellers from trading at all and used their reports to estimate the distribution and density for the remaining buyers. As the number of buyers goes to infinity, these estimates converge to the truth. However, the rate of convergence is slower than in our procedure which uses half the agents to estimate the distribution and density for the other half. Moreover, this procedure does not make as much revenue/capita as the one we utilize.*

Remark 4 *As pointed out in the Introduction, there are many mechanisms that guarantee efficiency when the number of buyers and sellers is large when the agents know the distribution of signals. To our knowledge, only McAfee [17] studies an environment where neither the mechanism designer nor the agents know the distribution. In his mechanism, the sellers are ranked $c_1 \leq c_2 \leq c_3 \dots$ given their announcements and the buyers are ranked $v_1 \geq v_2 \geq v_3 \dots$. The efficient trade q is such that $v_q \geq c_q$ and $v_{q+1} < c_{q+1}$. Roughly speaking, the mechanism allows the $q-1$ highest value buyers trade with the $q-1$*

lowest value sellers with buyers paying v_q and sellers receiving c_q and an intermediate broker retaining the difference. This mechanism converges to efficiency as the number of traders becomes large.

We now turn to the general case where the virtual functions are not monotonically increasing. For the moment, suppose that the distributions are common knowledge. Now, in the trading procedure, if trades with positive virtual surplus are executed, starting with those with the highest virtual surplus first, the monotonicity condition (2) is violated. Myerson [19] and Baron and Myerson [1] deal with this difficulty by “ironing” the virtual functions so they are monotonic. Virtual surplus is determined using the ironed virtual functions and trades with positive ironed virtual surplus are executed starting with the highest ironed virtual surplus. As the ironed virtual functions are flat over certain portions, traders may be tied and hence the trading procedure must randomize the allocation of objects over them.

We briefly describe the ironing procedure for buyers. Given $C_B(\cdot)$ as in (5), let

$$\kappa(\phi) \equiv C_B((F)^{-1}(\phi))$$

for any ϕ between 0 and 1. Let

$$K(\phi) \equiv \int_0^\phi \kappa(\tilde{\phi}) d\tilde{\phi}.$$

Let

$$\bar{K}(\phi) \equiv \text{conv}K(\phi)$$

be the convex hull of K . In other words, $\bar{K}(\phi)$ is the highest convex function on the interval $[0, 1]$, satisfying $\bar{K}(\phi) \leq K(\phi)$ for all ϕ in $[0, 1]$. Since, \bar{K} is convex, it is differentiable almost everywhere so let

$$\bar{\kappa}(\phi) \equiv (\bar{K})'(\phi)$$

whenever this is defined and extend $\bar{\kappa}(\phi)$ by right-continuity to all $0 \leq \phi \leq 1$. Finally, we define the *ironed virtual valuation function* for buyers as

$$\bar{C}_B(v) \equiv \bar{\kappa}(F(v)).$$

Notice that \bar{C}_B is non-decreasing as $\bar{\kappa}$ is the derivative of a convex function \bar{K} . An ironed virtual cost function can be constructed similarly. The ironed

virtual functions are used to determine the rankings of buyers and sellers and which of them trade. We refer the reader to Myerson [19] and Baron and Myerson [1] for the details.

When the distributions are not common knowledge, the trading mechanism must be used to estimate the densities and distributions to iron the virtual functions while ensuring incentive compatibility and approximating maximum profit for the broker.

Again, two markets, market A and market B , are set up, with buyers in M_k allowed to trade with sellers in N_k ($k \in \{A, B\}$). For buyers in market M_k , define the estimated distribution F_k and density f_k as in the proof of Theorem 3. As F_k is not strictly increasing, let

$$(F_k)^{-1}(\phi) \equiv \min\{v : F_k(v) = \phi\}$$

for any ϕ between 0 and 1. Given $C_B^k(\cdot)$ as in (8), let

$$\kappa_k(\phi) \equiv C_B^k((F_k)^{-1}(\phi))$$

for any ϕ between 0 and 1. Let

$$K_k(\phi) \equiv \int_0^\phi \kappa_k(\tilde{\phi}) d\tilde{\phi}.$$

Let

$$\bar{K}_k(\phi) \equiv \text{conv}K_k(\phi)$$

be the convex hull of K_k . Since, \bar{K}_k is convex, it is differentiable almost everywhere so let

$$\bar{\kappa}_k(\phi) \equiv (\bar{K}_k)'(\phi)$$

whenever this is defined and extend $\bar{\kappa}_k(\phi)$ by right-continuity to all $0 \leq \phi \leq 1$. Finally, we define the *estimated ironed virtual valuation function* for buyers M_k as

$$\bar{C}_B^k(v) \equiv \bar{\kappa}_k(F_k(v)).$$

Notice that \bar{C}_B^k is non-decreasing as $\bar{\kappa}_k$ is the derivative of a convex function \bar{K}_k .

A similar procedure for the seller defines the *estimated ironed virtual cost function* for sellers N_k

$$\bar{C}_S^j(c)$$

which is non-decreasing.

When the monotone hazard rate condition holds, buyers and sellers can be ranked according to their values and costs in the optimal mechanism and ties are a zero probability event. When the condition does not hold, agents have to be ranked according to the estimated virtual functions and ties have to be broken with positive probability. To ensure incentive compatibility, values that correspond to the same estimated ironed virtual value must face the same price should they transact. Also, a number of buyers, say, might be tied for the position of losing bidder. The price that is charged in this situation to a buyer who wins with probability one must ensure both that he has no incentive to understate his value and tie with these bidders and that the losing bidders have no incentive to exaggerate and win with probability one. These are the new features of the mechanism used to prove the following theorem:

Theorem 4 *As $\tau \rightarrow \infty$, the broker's expected profits per capita converge to his expected profits when he knows F and H .*

Proof: We utilize the estimated ironed virtual functions both to rank the buyers and sellers and to determine which trades to execute. First, in market k , label the buyers and sellers so buyer i has a higher reported estimated ironed virtual value than buyer $i + 1$ and seller j has a lower reported estimated ironed virtual cost than seller $j + 1$. If there are k buyers, say, with the same estimated ironed virtual valuation, label them randomly so each buyer has the same probability of being allocated a particular label. Second, the trading decisions are as follows:

$$\begin{aligned} p_1^\tau((v_1, \cdot), (c_1, \cdot)) &= 1 - q_1^\tau((v_1, \cdot), (c_1, \cdot)) = 1 \text{ if } \bar{C}_B^k(v_1) \geq \bar{C}_S^k(c_1) \\ &= 0 \text{ otherwise.} \end{aligned}$$

Continue in this fashion till either all the N_k units are traded or till we reach $q \in M_k$ such that

$$\bar{C}_B^k(v_q) \geq \bar{C}_S^k(c_q) \text{ and } \bar{C}_B^k(v_{q+1}) < \bar{C}_S^k(c_{q+1}).$$

This trading procedure uses the estimated virtual functions to mimic the profit-maximizing allocation when the distributions are known. When the set of agents is large, the estimates are close to the truth and our procedure achieves the same expected profit per capita as when the distributions are

known.

It remains to construct the payment functions so the individual rationality constraints of the lowest type of buyers and the highest types of sellers bind and incentive compatibility holds. Suppose buyer $q \in M_k$ and $q \in N_k$ are the lowest buyer and highest seller to trade (recall the procedure by which we have re-labelled the traders). Given v_{q+1} , suppose there are $S = \{1, \dots, s\}$ buyers such that $\bar{C}_B^k(v_i) > \bar{C}_B^k(v_{q+1})$ for $i \in S$ and $R = \{1, \dots, r\}$ buyers such that $\bar{C}_B^k(v_i) = \bar{C}_B^k(v_{q+1})$ for $i \in R$. Notice that, by definition, we must have $q \geq s$ and $s + r \geq q$. Also, given v_{q+1} , let $L(v_{q+1})$ be the highest valuation such that $\bar{C}_B^k(L(v_{q+1})) = \bar{C}_B^k(v_{q+1})$ and let $l(v_{q+1})$ be the lowest valuation such that $\bar{C}_B^k(l(v_{q+1})) = \bar{C}_B^k(v_{q+1})$. Choose v^* so that

$$v^* = \min\{v \mid \bar{C}_B^k(v) = \bar{C}_S^k(c_q)\}.$$

Notice that the variables $L(v_{q+1})$, $l(v_{q+1})$, and v^* all do not depend on buyer i 's report for $i < q + 1$. Moreover, as $v^* \geq \bar{C}_B^k(v^*) = \bar{C}_S^k(c_q) \geq c_q \geq a$, we must have $v^* \geq a$.

If buyer i wins an object by reporting a valuation above $L(v_{q+1})$ and $\bar{C}_B^k(v_{q+1}) < \bar{C}_S^k(c_q)$, he pays

$$v^*;$$

if buyer i wins an object by reporting a valuation above $L(v_{q+1})$ and $\bar{C}_B^k(v_{q+1}) \geq \bar{C}_S^k(c_q)$, he pays

$$\frac{r + s - q}{r + 1}L(v_{q+1}) + \frac{q - (s - 1)}{r + 1}l(v_{q+1});$$

if he wins by reporting a valuation $v_i \in [l(v_{q+1}), L^i(v_{q+1})]$, he pays

$$l(v_{q+1}).$$

A buyer who does not trade does not make any payment.

We show that buyers have a (weakly) dominant strategy to report their valuation truthfully.

First, consider the case where $\bar{C}_B^k(v_{q+1}) < \bar{C}_S^k(c_q)$. In this case, all winning buyers pay v^* . As the price at which winning buyers execute their trade is independent of their own reports, it is a dominant strategy to tell the truth by the standard argument.

Second, consider the case where $\bar{C}_B^k(v_{q+1}) \geq \bar{C}_S^k(c_q)$. Buyer i 's payoff is

$$\frac{r+s-q}{r+1}(v_i - L(v_{q+1})) + \frac{q-(s-1)}{r+1}(v_i - l(v_{q+1})) \quad (11)$$

if he is bidding more than $L(v_{q+1})$;

$$\frac{q-s}{r}(v_i - l(v_{q+1})) \quad (12)$$

if he is bidding between $L(v_{q+1})$ and $l(v_{q+1})$ and 0 otherwise.

Consider a buyer i who wins by truthfully bidding a valuation v_i greater than $L(v_{q+1})$ so $v_i > L(v_{q+1}) \geq l(v_{q+1})$. Therefore, $v_i - l(v_{q+1}) \geq v_i - L(v_{q+1}) > 0$ and buyer i 's expected payoff is positive. By bidding more than v_i , buyer i does not change the probability of trade or the expected price he pays so he has no incentive to overbid. If he underbids to the extent that he does not win an object with positive probability, his expected payoff goes down to zero. If he underbids so that he ties with the r bidders whose estimated ironed virtual valuations are exactly equal to $\bar{C}_B^k(v_{q+1})$, his probability of winning an object is $\frac{q-(s-1)}{r+1}$ as there are now $s-1$ bidders who have higher estimated ironed virtual valuations and $r+1$ bidders who have estimated ironed virtual valuations exactly equal to $\bar{C}_B^k(v_{q+1})$. His expected payoff is then $\frac{q-(s-1)}{r+1}(v_i - l(v_{q+1}))$ which, as $v_i - L(v_{q+1}) > 0$, is lower than his expected payoff (11) from bidding truthfully. Therefore, buyer i has no incentive to underbid in this case.

Next, consider a buyer i who wins with positive probability by truthfully reporting a valuation v_i such that $\bar{C}_B^k(v_i) = \bar{C}_B^k(v_{q+1})$. By definition of $L(v_{q+1})$ and $l(v_{q+1})$, we must have $v_i \in [l(v_{q+1}), L(v_{q+1})]$ and $v_i - l(v_{q+1}) \geq 0 \geq v_i - L(v_{q+1})$ so his expected payoff is non-negative. By underbidding, buyer i either does not change the outcome or does not win an object. Therefore, he has no incentive to underbid as $v_i - l(v_{q+1}) \geq 0$. By overbidding, either buyer i does not change the outcome or he wins for certain if he bids more than $L(v_{q+1})$. In the latter case, there are now $s+1$ bidders whose reported estimated ironed virtual valuations are strictly greater than $\bar{C}_B^k(v_{q+1})$ and $r-1$ bidders whose reported estimated ironed virtual valuations are exactly equal to $\bar{C}_B^k(v_{q+1})$. Therefore, buyer i 's expected payoff is

$$\frac{r+s-q}{r}(v_i - L(v_{q+1})) + \frac{q-s}{r}(v_i - l(v_{q+1}))$$

if he bids more than $L(v_{q+1})$. But, as $v_i - L(v_{q+1}) \leq 0$ this is less than the expected payoff (12) from telling the truth. Therefore, buyer i has no incentive to overbid in this case.

Finally, consider a buyer i who does not win an object with positive probability as he truthfully reports a valuation $v_i < l(v_{q+1})$. By lying, either buyer i does not change the outcome or he makes a loss, as $v_i < l(v_{q+1})$, if he overbids to the extent that he wins an object with positive probability.

Therefore, buyer i has a (weakly) dominant strategy to report his valuation truthfully.

Also, $U_i(a) = 0$ as the lowest price a buyer ever pays is a .

A similar procedure is used to construct the payment function to sellers and similar arguments shows that all sellers have a dominant strategy to report their costs truthfully and that $U_j(b) = 0$.

Finally, as τ goes to infinity, the estimated ironed virtual functions converge to the true ones. Therefore, the expected profits per capita converge to those when the distributions are known. ■

4 Conclusion

We now turn to a number of issues not covered by the analysis above.

First, we note that our analysis can (trivially) be extended to auctions with a fixed number of objects for sale: Suppose there are K goods for sale and each buyer wishes to purchase at most one unit. All buyers' valuations are drawn from the same distribution so the distribution is used only in determining the reserve price in the optimal auction. In this case, as the set of buyers becomes large, the probability that K buyers do not have valuations above the reserve price goes to zero. Hence, the seller can run a $K + 1^{th}$ price auction without a reserve and get close to the optimum as the set of buyers gets large. If buyers are drawn from different distributions, then their virtual valuations have to be compared when allocating objects. Suppose that buyer 1's valuation is drawn from a distribution with a support that contains higher valuations than those of any other buyer. As the set of bidders drawn from each distribution becomes large, all the objects are allocated to buyers drawn from the same distribution as buyer 1. Hence, only valuations not virtual valuations must be compared when allocating objects and, again, the probability that winning buyers do not have valuations above the reserve price that is optimal for buyer 1's distribution goes to zero as the set of

buyers become large. Therefore, a $K + 1^{th}$ price auction maximizes revenue as the set of bidders becomes large.

Second, the analysis above assumes that the number of objects sold is fixed. If the set of objects grows faster than the number of bidders (or the support of the distribution of valuations expands sufficiently rapidly with the number of buyers), the issues we have studied above again arise and can be similarly resolved.

Finally, future work might therefore focus on the extension to interdependent values and the more difficult issue of how to design profit-maximizing mechanisms when the distribution is unknown and there are a small number of traders.

5 Colophon

Market Research and Market Design

Sandeep Baliga and Rakesh Vohra
Kellogg S.O.M (M.E.D.S.),
2001 Sheridan Road,
Evanston, IL 60208-2009.

(baliga@kellogg.northwestern.edu and r-vohra@kellogg.northwestern.edu)
July, 2003.

References

- [1] David Baron and Roger Myerson (1982): “Regulating a Monopolist with Unknown Costs,” *Econometrica*, Vol. 50, pp 911-930.
- [2] Dirk Bergemann and Stephen Morris (2002): “Robust Mechanism Design,” mimeo, Yale University.
- [3] Leo Breiman (1992). *Probability*. Philadelphia: S.I.A.M. Press.
- [4] Jeremy Bulow and John Roberts (1989): “The Simple Economics of Optimal Auctions,” *Journal of Political Economy*, Vol. 97, 1060-1090.
- [5] Bernard Caillaud and Jacques Robert (1998): “Implementing the Optimal Auction,” mimeo.

- [6] K. Chung (1949): “An Estimate Concerning the Kolmogorov Limit Distribution,” *Transactions of the American Mathematical Society*, 67, pp 37-50.
- [7] Partha Dasgupta and Eric Maskin (2000): “Efficient Auctions,” *Quarterly Journal of Economics*, Vol. CXV, 341-388.
- [8] Kastubh Deshmukh, Andrew Goldberg, Jason Hartline and Anna Karlin (2002): “Truthful and Competitive Double Auctions,” (Extended Abstract), *European Symposium on Algorithms*. Heidelberg: Springer Verlag.
- [9] Luc Devroye (1987). *A Course in Density Estimation*. Heidelberg: Springer Verlag.
- [10] Andrew Goldberg, Jason Hartline, Anna Karlin and Andrew Wright (2001): “Competitive Auctions,” *Proceedings of the 12th Symposium on Discrete Algorithms*, pp 735-744. ACM/SIAM.
- [11] Thomas Gresik and Mark Satterthwaite (1989): “The Rate at which a Simple Market Converges to Efficiency as the Number of Traders Increases: An Asymptotic Result for Optimal Trading Mechanisms,” *Journal of Economic Theory*, Vol. 48, pp 304-332.
- [12] Faruk Gul and Andrew Postlewaite (1992): “Asymptotic Efficiency in Large Economies with Asymmetric Information,” *Econometrica*, Vol. 60, 1273-1292.
- [13] Bengt Holmstrom and Roger Myerson (1983): “Efficient and Durable Decision Rules with Incomplete Information,” *Econometrica*, Vol. 51, 1799-1819.
- [14] Eric Maskin (1992): “Auctions and Privatization,” in *Privatization* (ed. by H. Siebert).
- [15] Eric Maskin and John Riley (1984): “Monopoly with Incomplete Information,” *RAND Journal of Economics*, 15, pp 171-196.
- [16] Richard McLean and Andrew Postlewaite (1999): “Informational Size and Incentive Compatibility,” mimeo, University of Pennsylvania.

- [17] Preston McAfee (1992): “A Dominant Strategy Double Auction,” *Journal of Economic Theory*, 56, pp 434-450.
- [18] John Moore (1992): “Implementation, Contracts and Renegotiation in Environments with Complete Information,” in *Advances in Economic Theory: Sixth World Congress, Volume 1*, ed. J.-J. Laffont. Cambridge: Cambridge University Press.
- [19] Roger Myerson (1981) : “Optimal Auction Design,” *Mathematics of Operations Research*, Vol. 6, pp58-73.
- [20] Roger Myerson and Mark Satterthwaite (1983): “Efficient Mechanisms for Bilateral Trade,” *Journal of Economic Theory*, Vol. 29, pp 265-281.
- [21] Ilya Segal (2003): “Optimal Pricing Mechanisms with Unknown Demand,” *American Economic Review*, 93(3), pp. 509-529.
- [22] William Vickrey (1961): “Counterspeculation, Auctions and Competitive Sealed Tenders,” *Journal of Finance*, 16, 8-37.
- [23] Robert Wilson (1985): “Incentive Efficiency of Double Auctions,” *Econometrica*, Vol. 53, pp 1101-1115.
- [24] Robert Wilson (1987): “Game-Theoretic Approaches to Trading-Processes,” in *Advances in Economic Theory: Fifth World Congress*, ed. T. Bewley, Cambridge: Cambridge University Press.