

Competitiveness via Consensus

Andrew V. Goldberg*
goldberg@microsoft.com

Jason D. Hartline†
hartline@cs.washington.edu

Abstract

We introduce the following consensus estimate problem. Several processors hold private and possibly different lower bounds on a value. The processors do not communicate with each other, but can observe a shared source of random numbers. The goal is to come up with a consensus lower bound on the value that is as high as possible. We give a solution to the consensus estimate problem and show how it is useful in the context of mechanism design. The consensus problem is natural and may have other applications.

Based on our consensus estimate technique, we introduce Consensus Revenue Estimate (CORE) auctions. This is a class of competitive revenue-maximizing auctions that is interesting for several reasons. One auction from this class achieves a better competitive ratio than any previously known auction. Another one uses only two random bits, whereas the previously known competitive auctions on n bidders use n random bits. Furthermore, a parameterized CORE auction performs better than the previous auctions in the context of mass-market goods, such as digital goods.

1 Introduction

Recent economic and computational trends, such as the negligible cost of duplicating digital goods and the emergence of the Internet, have led to a number of new and interesting dynamic pricing problems. A common approach to such problems, originating in the field of *mechanism design* (see, e.g., [8, 11, 18, 21, 23, 33]), is to develop protocols in which all rational participants are motivated to follow the protocol. Traditionally studied by Economists and Game Theorists, the relevance of mechanism design to Computer Science applications and challenging computational issues in mechanism implementations have recently attracted attention of Computer Scientists. See, e.g., [1, 2, 13, 20, 26, 27, 28, 29].

In this paper we study auctions, which are an im-

portant class of mechanisms. We consider auctions for a set of identical items. We assume that each consumer has a private *utility value*, i.e., the maximum value the consumer is willing to pay for an item. An *auction* takes as input *bids* from each of the consumers and determines which bidders receive an item and at what price. We say the auction is *truthful* (or equivalently, strategy-proof or incentive-compatible) if it is in each consumer's best interest to bid their true utility value.

We are interested in truthful auctions that bring high expected revenue to the seller. Economists have studied such auctions in a Bayesian framework, i.e., assuming a known probability distribution of bidder utilities. A natural approach in this setting is to use the prior distribution knowledge to set reservation prices in the Vickrey-Clarke-Groves mechanism [8, 18, 33]. This approach has been taken in [7, 25]. Auction design problems under weaker assumptions on the prior distribution knowledge have recently been studied in [3, 31]. In contrast to the Bayesian design and analysis framework, in [15, 16] we consider a worst case competitive analysis of auctions. This approach is motivated by on-line algorithms (see, e.g., [6, 32]). We define *competitive auctions* as the truthful auctions that obtain a fraction of the revenue of an optimal auction¹ on every input. An auction is β -*competitive* if on every input its expected revenue is at least the optimal revenue divided by β . We refer to β as the *competitive ratio*.

In [15, 16] we established a framework for competitive auction design and opened the important area of worst-case analysis of profit maximizing mechanisms for further research. Subsequent work on competitive auctions includes extensions to the non-homogeneous item case [17], to the on-line case [4, 5], and to more general mechanisms [14]. However, the area is very rich in open problems. We address some of these problems in this paper and substantially improve upon some of the previous results.

In [14] the optimal competitive ratio is bounded between two and four.² The upper bound is achieved

*Microsoft Research, 1065 La Avenida, Mountain View, CA 94043.

†Computer Science Department, University of Washington.

¹We discuss optimal auctions in Section 2.

²In subsequent work with Michael Saks and Anna Karlin we

by the Sampling Cost Sharing auction. A natural problem we address is that of narrowing the gap between the upper and the lower bounds.

The amount of randomness needed by a competitive auction is another interesting question. An auction is *symmetric* if its outcome depends only on the set of input bids and not on the order the bids are presented in. No symmetric deterministic auction is competitive [15]. Previously known competitive auctions are symmetric and randomized. These auctions are based on random sampling and use one random bit for every bidder. We show that symmetric auctions can be competitive with very little randomization.

A desirable feature of an auction is its ability to achieve competitive ratios close to one in mass-market applications, such as digital goods distribution, where the number of items sold is guaranteed to be large. Results of [15] imply that the parameterized variant of the Dual-price Sampling Optimal Threshold auction, $DSOT_\alpha$, is asymptotically good: its competitive ratio converges to one as α , the minimum number of items guaranteed to be sold, grows. However, the analysis of $DSOT_\alpha$ guarantees competitive ratios that are close to one only for very large values of α . Our new techniques lead to auctions with good competitive ratios for more reasonable values of α .

One of the main contributions of this paper is the definition of a consensus estimate problem and a solution to this problem. Given $\rho > 0$ and $v > 0$, we say that a function g is a ρ -*consensus estimate* of v if for every $w : v/\rho \leq w \leq v$, $g(w)$ is a consensus estimate of v , i.e., $g(w) = g(v)$ and $0 < g(v) \leq v$. This problem has no deterministic solution. We show how to choose g at random so that g is a good consensus estimate on average: For any v , if we define a random variable γ to be $g(v)$ if g is a consensus estimate and zero otherwise, then $\mathbf{E}[\gamma] = \Omega(v)$.

As an application of the consensus estimate problem, we introduce the class of *COnsensus Revenue Estimate* (CORE) auctions. Intuitively, these auctions use consensus estimate to compute a lower bound on the optimal auction revenue that is independent of each individual bid, and use this bound as a target revenue of the cost-sharing auction. Our analysis of CORE auctions leads to the following results:

- We show a CORE auction that is 3.39-competitive. This improves the previous best competitive ratio of four.

¹have improved the lower bound to 2.42.

- We describe a competitive CORE auction that uses only two random bits. This essentially closes the question of how much randomness a symmetric competitive auction needs.
- We introduce a parameterized CORE auction, $CORE_\alpha$, that achieves competitive ratios close to one for more realistic values of α than $DSOT_\alpha$. For example, for $\alpha = 1,000$, we get the competitive ratio of 1.0465, i.e., the auction revenue is less than 5% away from optimal.

The fact that the CORE auctions achieve better competitive ratios and use less randomness demonstrates that techniques behind CORE are interesting and powerful.

2 Definitions and Background

We consider single-round, sealed-bid auctions for a set of identical items available in unlimited supply. As in the previous work on competitive auctions, the results can be extended to the case when the number of available items is limited.

DEFINITION 2.1. A single-round sealed-bid auction mechanism *is one where*:

- *Given the bid vector $\mathbf{b} = (b_1, \dots, b_n)$, i.e., the input, the mechanism computes an outcome allocation, $\mathbf{x} \in \{0, 1\}^n$ and prices, $\mathbf{p} \in \mathbb{R}^n$; i.e., the output. If $x_i = 1$ bidder i wins (i.e. receives the item) and pays price p_i , otherwise we say that bidder i loses.*
- *We assume that $0 \leq p_i \leq b_i$ for all winning bidders and that $p_i = 0$ for all losing bidders (these are the assumptions of no positive transfers and voluntary participation; see, e.g., [24]).*
- *The auctioneer's profit is $\mathcal{A}(\mathbf{b}) = \sum_i p_i$.*

We say the *mechanism is randomized* if the procedure used to compute the allocations and prices is randomized. Otherwise, the mechanism is *deterministic*. Note that if the mechanism is randomized, the profit of the mechanism, the output prices, and the allocation are random variables.

We say that an *auction is symmetric* if the price and allocation vectors are independent of the order of the bids (i.e., they depend only on the set of input bids). All auctions discussed in this paper are symmetric.

We use the following private value model for bidders. Each bidder i has a private utility value u_i , representing the true maximum they are willing to pay for an item. Bidders are *rational*, i.e., each

bidder bids so as to maximize their *profit*, $u_i x_i - p_i$. Bidders have full knowledge of the auction mechanism. Bidders do not collude.

In the rest of this section we review the notions of truthfulness and bid-independence. We describe the competitive framework that we use as a performance metric for analyzing auctions. We review an interesting special case of cost sharing mechanism [24] which our auctions use. Finally, we discuss combining auctions that are competitive on sets of restricted of inputs to get auctions that are competitive on the union of the sets.

2.1 The Bid-Independent Characterization

of Truthful Auctions. We say that a deterministic auction is *truthful* if, for each bidder i and any choice of bid values for all other bidders, bidder i 's profit is maximized by bidding their utility value. We say that a randomized auction is *truthful* if it can be described as a probability distribution over deterministic truthful auctions. Note that with this notion of truthfulness, the probability that a bidder's profit exceeds a value v is simultaneously maximized for every v by bidding truthfully. In the remainder of this paper, when considering truthful auctions, we assume that $b_i = u_i$.

Next we describe a useful characterization of truthful auctions using the notion of *bid-independence*. We define deterministic bid-independent auctions first. Let \mathbf{b}_{-i} denote the vector of bids \mathbf{b} with b_i removed, i.e., $\mathbf{b}_{-i} = (b_1, \dots, b_{i-1}, ?, b_{i+1}, \dots, b_n)$. We call such a vector *masked*. Given a function f on masked vectors, the *deterministic bid-independent auction defined by f* , \mathcal{A}_f , is:

DEFINITION 2.2. (BID-INDEPENDENT AUCTION, \mathcal{A}_f)

On input \mathbf{b} , for each bidder i do the following:

1. $v_i \leftarrow f(\mathbf{b}_{-i})$.
2. If $v_i \leq b_i$, set $x_i \leftarrow 1$ and $p_i \leftarrow v_i$ (*Bidder i wins*).³
3. Otherwise, set $x_i = p_i = 0$ (*Bidder i loses*).

A *randomized bid-independent auction* is a probability distribution over bid-independent auctions. For these auctions, $f(\mathbf{b}_{-i})$ is a non-negative real-valued random variable.

³In fact, bid-independence allows the inequality, $v_i \leq b_i$, to be strict or non-strict at the discretion of $f(\mathbf{b}_{-i})$. Thus f can specify whether to accept b_i at price v_i if b_i is in (v_i, ∞) or $[v_i, \infty)$. For the auctions presented in this paper this subtlety is not important.

THEOREM 2.1. [14] *An auction is truthful if and only if it is equivalent to a bid-independent auction.*

The proof of Theorem 2.1 is constructive. Given a truthful auction \mathcal{A} the proof shows how to obtain the bid-independent function f such that \mathcal{A} is equivalent to \mathcal{A}_f .

2.2 Competitive Auctions. We now review the competitive framework from [14]. In order to evaluate the performance of auctions with respect to the goal of profit maximization, we introduce the optimal single price omniscient auction \mathcal{F} and the related auction $\mathcal{F}^{(m)}$.

The *optimal single price omniscient auction*, \mathcal{F} , is defined as follows: Let \mathbf{b} be a bid vector, and let v_i be the i -th largest bid in \mathbf{b} . Auction \mathcal{F} on input \mathbf{b} determines the value k such that kv_k is maximized. All bidders with $b_i \geq v_k$ win at price v_k ; all remaining bidders lose. The profit of \mathcal{F} on input \mathbf{b} is thus $\mathcal{F}(\mathbf{b}) = \max_{1 \leq k \leq n} kv_k$.

The *optimal single price omniscient auction that sells to at least m bidders*, $\mathcal{F}^{(m)}$, is defined as follows: Let \mathbf{b} and v_i be as in the previous paragraph. Auction $\mathcal{F}^{(m)}$ on input \mathbf{b} determines the value k such that $k \geq m$ and kv_k is maximized. All bidders with $b_i \geq v_k$ win at price v_k ; all remaining bidders lose. The profit of $\mathcal{F}^{(m)}$ on input \mathbf{b} is thus $\mathcal{F}^{(m)}(\mathbf{b}) = \max_{m \leq k \leq n} kv_k$.

We extend the definition of $\mathcal{F}^{(m)}$ to masked vectors by treating the “?” as zero.

We would like an auction to be competitive with \mathcal{F} on every input; however, as shown in [15], if \mathbf{b} is such that a single bidder's utility dominates the total utility of the other bidders, no auction can compete with $\mathcal{F}(\mathbf{b})$. We use two notions of competitiveness for auctions: a general competitiveness for an auction on any set of bids [14] and notion of competitiveness suited for mass-markets where the number of winners of the optimal auction is large [16, 17]. In the latter case, an auction need not be competitive on all inputs.

Let \mathcal{A} be a truthful auction. We say that \mathcal{A} is β -competitive against $\mathcal{F}^{(2)}$ (or just β -competitive) if for all bid vectors \mathbf{b} , the expected profit of \mathcal{A} on \mathbf{b} satisfies

$$\mathbf{E}[\mathcal{A}(\mathbf{b})] \geq \frac{\mathcal{F}^{(2)}(\mathbf{b})}{\beta}.$$

We say that \mathcal{A} is competitive if there exist a constant β such that \mathcal{A} is β -competitive. One can also compare the auction revenue to $\mathcal{F}^{(m)}$ for $m > 2$; however, we do not do so in this paper.

For mass-market auctions (Section 4.2) we use an alternative notion of competitiveness. Given a bid vector \mathbf{b} , let $h(\mathbf{b})$ denote the maximum bid value in

- b. Let \mathcal{A} be a truthful auction. We say that \mathcal{A} is (α, β) -competitive if for all bid vectors \mathbf{b} such that $\mathcal{F}(\mathbf{b}) \geq \alpha h(\mathbf{b})$, the expected profit of \mathcal{A} on \mathbf{b} satisfies

$$\mathbf{E}[\mathcal{A}(\mathbf{b})] \geq \frac{\mathcal{F}(\mathbf{b})}{\beta}.$$

We will use the fact that $\mathcal{F}(\mathbf{b}) \geq \alpha h(\mathbf{b})$ implies that $\mathcal{F}(\mathbf{b})$ has at least α winners.

Note that the first definition of competitiveness is stronger: One can show that the fact that \mathcal{A} is β -competitive against $\mathcal{F}^{(\alpha)}$ implies that \mathcal{A} is (α, β) -competitive. The converse is not necessarily true. Our strongest notion of a competitive auction is that against $\mathcal{F}^{(2)}$.

2.3 Cost Sharing. In this section we review the truthful cost sharing mechanism $\text{CostShare}_{\mathcal{C}}$ [24]. The goal of cost sharing is, given bids \mathbf{b} , to share the cost of a good or service among a subset of the bidders. We restrict our attention to the simple case of sharing a target value C whenever possible.

$\text{CostShare}_{\mathcal{C}}$: Given bids \mathbf{b} , find the largest k such that the highest k bidders can equally share the cost C . Charge each of these bidders C/k . If no subset of bidders can share the cost, the auction has no winners.

Important properties of this auction are as follows:

- $\text{CostShare}_{\mathcal{C}}$ is truthful and symmetric.
- If $C \leq \mathcal{F}(\mathbf{b})$, $\text{CostShare}_{\mathcal{C}}$ has revenue C ; otherwise it has no winners and no revenue.

Since the auction is truthful, it is equivalent to a bid-independent auction for some function f on masked bid vectors. Let $\text{cs}_{\mathcal{C}}$ be this function (Theorem 2.1 guarantees that such a function exists).

2.4 Convex Combination of Auctions. As we shall see, the CORE approach works on bids such that the optimal auction, $\mathcal{F}^{(2)}$, has at least three winners. To get an auction that is competitive on all sets of bids, we use the following result.

Let \mathcal{A}' and \mathcal{A}'' be auctions such that \mathcal{A}' is β' -competitive on $\mathbf{b} \in \mathcal{B}'$ and \mathcal{A}'' is β'' -competitive on $\mathbf{b} \in \mathcal{B}''$. Consider the auction \mathcal{A} that is a “convex combination” of \mathcal{A}' and \mathcal{A}'' : With probability p , \mathcal{A} runs \mathcal{A}' and otherwise \mathcal{A} runs \mathcal{A}'' . The following result is straight-forward.

LEMMA 2.1. \mathcal{A} is $\max(\beta'/p, \beta''/(1-p))$ -competitive on $\mathbf{b} \in \mathcal{B}' \cup \mathcal{B}''$. For the optimal choice of p , \mathcal{A} is $(\beta' + \beta'')$ -competitive on $\mathbf{b} \in \mathcal{B}' \cup \mathcal{B}''$.

Note that the competitive ratio of \mathcal{A} may be better than the lemma guarantees.

Recall that, by definition, $\mathcal{F}^{(2)}(\mathbf{b})$ has at least two winners. Let \mathcal{B}_{3+} denote the set of all \mathbf{b} such that $\mathcal{F}^{(2)}(\mathbf{b})$ has at least three winners and let \mathcal{B}_2 denote the set of bids such that $\mathcal{F}^{(2)}(\mathbf{b})$ has exactly two winners. If a CORE auction is competitive on \mathcal{B}_{3+} , we obtain an auction competitive for all bids by combining the CORE auction with the Vickrey auction [33]. The Vickrey auction is the bid-independent auction defined by the \max function. Since the Vickrey auction sells to the highest bidder at the price equal to the second highest bid, we have the following result.

LEMMA 2.2. *The Vickrey auction is 2-competitive on $\mathbf{b} \in \mathcal{B}_2$.*

3 The Consensus Estimate Problem

In this section we study the following consensus problem that is the key to our CORE auctions. For a given $\rho > 1$ and $v > 0$, we say that a function g is a ρ -consensus estimate of v if

1. g is a *consensus*: for any w such that $v/\rho \leq w \leq v$, we have $g(w) = g(v)$.
2. $g(v)$ is a nontrivial lower bound on v , i.e., $0 < g(v) \leq v$.

We call $g(v)$ the *consensus value*. Intuitively, we would like to find a good consensus estimate for v , where the higher the consensus value, the higher the quality of the estimate. The *payoff*, γ_g , for a function g is $\gamma_g(v) = g(v)$ if g is a ρ -consensus estimate on v and $\gamma_g(v) = 0$ otherwise.

It is easy to see that no deterministic function g is a ρ -consensus estimate for all v . Consider a deterministic g , fix $v > 0$ and $\rho > 1$, and assume that g is a ρ -consensus estimate of v . An inductive argument shows that for any $i \geq 1$, $g(v) = g(v/\rho^i)$, and therefore $g(v) \leq 0$.

DEFINITION 3.1. *The consensus estimate problem is, for any ρ , to give a distribution \mathcal{G} on functions g such that for any v the expected payoff is large relative to v . That is, the worst case value of $\mathbf{E}[\gamma_g(v)]/v$ is large over choices of v .*

3.1 Consensus Estimate Algorithm. We describe a distribution \mathcal{G} that works well in the following sense: for any v , with g from \mathcal{G} we have $\mathbf{E}[g(v)] = \Omega(v)$, where the constant hidden by the Ω notation is a function of ρ . Our solution uses an additional parameter $c > \rho$. The value of c is chosen as a function of ρ to maximize the quality of the

estimate. Consider the following function g_u^c :

$$g_u^c(v) = v \text{ rounded down to nearest } c^{j+u} \text{ for integer } j.$$

Remark. The definition of g_u^c implies that for any v , $\frac{v}{c} \leq g_u^c(v) \leq v$. Thus if g_u^c is a consensus for v then it is a consensus estimate with value within a factor of c from v .

We define \mathcal{G} as a distribution of functions of the form g_U^c with U chosen uniformly on $[0, 1]$. We repeatedly make use of the following result.

LEMMA 3.1. *For g from \mathcal{G} , $g(v)$ is distributed identically to $c^{U'} v/c$ for U' uniform on $[0, 1]$.*

Proof. Consider a random variable $Y = \log_c g(v)$ and let $t = \log_c v - 1$. Then $\Pr[Y \leq t + x] = \Pr[U' \leq x]$ and therefore Y is uniformly distributed between t and $t + 1$. Thus, $g(v)$ is identical to $c^{U'} v/c$. ■

For U uniform $[0, 1]$, the random variable c^U satisfies $\Pr[c^U \leq z] = \Pr[U \leq \log_c z] = \log_c z = \frac{\ln z}{\ln c}$. The probability density function for c^U is $\pi(x) = 1/(x \ln c)$ for $1 \leq x < c$. To see this, note that $\Pr[c^U \leq z]$ is $\int_1^z \frac{1}{x \ln c} dx = \frac{\ln z}{\ln c}$.

Next we bound the probability that g is a consensus.

LEMMA 3.2. *For g from \mathcal{G} , the probability g is a consensus estimate is $1 - \log_c \rho$.*

Proof. g is a consensus estimate for v if $g(v) \leq \frac{v}{\rho}$. Using Lemma 3.1, we get

$$\begin{aligned} \Pr\left[g(v) \leq \frac{v}{\rho}\right] &= \Pr\left[c^U v/c \leq \frac{v}{\rho}\right] = \Pr\left[c^U \leq \frac{v}{\rho}\right] \\ &= \log_c(c/\rho) = 1 - \log_c \rho. \end{aligned}$$

■

In the application to auctions, the value of ρ is fixed and we chose c to maximize the expectation of γ . Lemma 3.2 implies $\mathbf{E}[\gamma_g(v)] \geq \frac{v}{c}(1 - \log_c \rho)$. The following theorem gives a better bound.

THEOREM 3.1. *For g from \mathcal{G} defined above, for all v , $\mathbf{E}[\gamma_g(v)] = \frac{v}{\ln c}(\frac{1}{\rho} - \frac{1}{c})$.*

Proof. By Lemma 3.1, $g(v)$ is distributed as $c^U v/c$ for U uniform on $[0, 1]$. Therefore,

$$\begin{aligned} \mathbf{E}[\gamma(v)] &= \frac{v}{c} \int_1^{c/\rho} x \cdot \pi(x) dx + \int_{c/\rho}^c 0 \cdot \pi(x) dx \\ &= \frac{v}{c} \int_1^{c/\rho} \frac{1}{\ln c} dx = \frac{v}{\ln c} \left(\frac{1}{\rho} - \frac{1}{c}\right). \end{aligned}$$

■

Note that for a fixed ρ , one can chose the value of c that maximizes $\mathbf{E}[\gamma_g(v)]/v$.

3.2 Consensus Estimates with One Random Bit.

Given the above consensus estimate solution that uses a random real number chosen uniformly from $[0, 1]$ and the fact that no purely deterministic consensus estimate exists, it is natural to ask how much randomness is necessary. We show how to give a consensus estimate with only one random bit. Choose a constant $c' > \rho$ and let $c = c'^2$. Pick the value of u uniformly from $\{0, 1/2\}$ and use function g_u^c as defined above. Note that for these values of u , we round the revenue estimates to even and odd powers of c' , respectively. Since $c' > \rho$, for any value of v , at most one of these values can be in the interval $[v/\rho, v]$ and therefore the revenue estimates agree with probability of at least $1/2$ and the expected payoff is at least $v/(2c'\rho)$. This gives the following lemma:

LEMMA 3.3. *The consensus estimate solution with g uniform from $\{g_0^{c'}, g_{\frac{1}{2}}^{c'}\}$ and $c' = \rho + \frac{\epsilon}{2\rho}$ gives $\mathbf{E}[\gamma_g(v)] = v/(2\rho^2 + \epsilon)$.*

4 CORE Auctions

In this sections we introduce the class of CORE auctions. Recall from Section 2.3 the cost sharing mechanism CostShare_C and the bid-independent function cs_C that implements it. To explain the intuition behind the CORE auctions, consider the generic bid-independent mechanism \mathcal{A}_f : For each bidder i we use the price $f(\mathbf{b}_{-i})$. If it was possible to compute $\mathcal{F}^{(2)}(\mathbf{b})$ from \mathbf{b}_{-i} then f could use this value and output the price, $\text{cs}_{\mathcal{F}^{(2)}(\mathbf{b})}(\mathbf{b}_{-i})$. This mechanism would achieve revenue $\mathcal{F}^{(2)}(\mathbf{b})$. Unfortunately, we cannot compute $\mathcal{F}^{(2)}(\mathbf{b})$ from \mathbf{b}_{-i} . However, if we can show that removing a bidder does not change $\mathcal{F}^{(2)}$ much, i.e., $\mathcal{F}^{(2)}(\mathbf{b})/\rho \leq \mathcal{F}^{(2)}(\mathbf{b}_{-i})$, then we can use a ρ -consensus estimate for each bidder i to agree on a value r that is close to but less than $\mathcal{F}^{(2)}(\mathbf{b})$. Since this value r is computed bid-independently, it can be used with the bid-independent cost share function.

First we show that removing a bid does not change $\mathcal{F}^{(2)}$ much.

LEMMA 4.1. *If $\mathcal{F}^{(2)}(\mathbf{b})$ has $k \geq 3$ winners, then for any i ,*

$$\frac{k-1}{k} \mathcal{F}^{(2)}(\mathbf{b}) \leq \mathcal{F}^{(2)}(\mathbf{b}_{-i}) \leq \mathcal{F}^{(2)}(\mathbf{b}).$$

When the the number of winners in $\mathcal{F}^{(2)}(\mathbf{b})$ is $k \geq 3$, i.e., $\mathbf{b} \in \mathcal{B}_{3+}$, Lemma 4.1 allows us to estimate $\mathcal{F}^{(2)}(\mathbf{b})$ from $\mathcal{F}^{(2)}(\mathbf{b}_{-i})$. In this case, we can use $\rho = \frac{3}{2} \geq \frac{k}{k-1}$ in combination with our ρ -consensus estimate solution from Section 3 as follows.

Let g be a function picked from \mathcal{G} as defined in Section 3. Let R be a function from masked bid

vectors to reals defined by $R(\mathbf{b}_{-i}) = g(\mathcal{F}^{(2)}(\mathbf{b}_{-i}))$. We define the *Basic CORE* auction as the bid-independent auction defined by the following function f_R :

$$f_R(\mathbf{b}_{-i}) = \text{cs}_{R(\mathbf{b}_{-i})}(\mathbf{b}_{-i}).$$

Since g , and therefore R , is chosen from a probability distribution, the auction is randomized. Intuitively, R estimates the optimal revenue.

Note that Lemma 4.1 applies for $\mathbf{b} \in \mathcal{B}_{3+}$, i.e., \mathbf{b} such that $\mathcal{F}^{(2)}(\mathbf{b})$ has $k \geq 3$ winners. We want an auction that is competitive on any \mathbf{b} . Motivated by Section 2.4, we take a convex combination of the above auction with the one-item Vickrey auction (that sells to the highest bidder at the second highest bid value). The Vickrey auction is competitive on $\mathbf{b} \in \mathcal{B}_2$, i.e., \mathbf{b} such that $\mathcal{F}^{(2)}(\mathbf{b})$ has $k = 2$ winners. This convex combination results in an auction that is competitive on all inputs.

DEFINITION 4.1. *The CORE auction is a convex combination of Vickrey (with probability p) and the Basic CORE auction (with probability $1 - p$).*

Theorem 3.1 and Lemma 2.1 can be used to show that the resulting auction is constant-competitive. We give a tighter analysis that utilizes the following observation: The Vickrey auction is k -competitive on \mathcal{B}_{3+} . Therefore if we run the Vickrey auction with probability p , this adds $p\mathcal{F}^{(2)}(\mathbf{b})/k$ to the expected revenue of the \mathcal{B}_{3+} case.

THEOREM 4.1. *For an appropriate choice of c and p , the CORE auction is 3.39-competitive against $\mathcal{F}^{(2)}(\mathbf{b})$.*

Proof. Case 1 ($\mathbf{b} \in \mathcal{B}_{3+}$): For the Basic CORE auction, γ , defined in Section 3, is a lower bound on the auction revenue. Then Theorem 3.1 implies that if the Basic CORE auction is selected, the expected revenue is at least $\frac{\mathcal{F}^{(2)}(\mathbf{b})}{\ln c} \left(\frac{k-1}{k} - \frac{1}{c} \right)$. If the Vickrey auction is selected, the expected revenue is at least $\frac{\mathcal{F}^{(2)}(\mathbf{b})}{k}$. Thus the total revenue is at least

$$\mathcal{F}^{(2)}(\mathbf{b}) \left(\frac{p}{k} + \frac{1-p}{\ln c} \left(1 - \frac{1}{k} - \frac{1}{c} \right) \right).$$

Case 2 ($\mathbf{b} \in \mathcal{B}_2$): The expected revenue for \mathbf{b} is at least $\mathcal{F}^{(2)}(\mathbf{b}) \frac{p}{2}$ due to the Vickrey auction.

We pick p and c to balance the Case 1 and Case 2 revenue. Numeric simulation shows that $c = 2.0$ and $p = 0.59$ is a near-optimal choice. This choice gives a competitive ratio of 3.39. ■

This auction has several interesting properties. One property is that in the “normal” case, i.e., when

the revenue estimates agree or when the Vickrey auction is used, the outcome of the auction is a single sale price with the property that every bidder that bid above the price wins. In general, we have the following result.

LEMMA 4.2. *The CORE auction is at most dual-priced.*

Proof. If we are using Vickrey, the sale price is unique. Otherwise, if we have a consensus estimate, the sale price is also unique. Suppose the estimates disagree. Consider two cases, $k \geq 3$ and $k = 2$.

In the first case, $R(\mathbf{b})$ is between $\frac{k-1}{k}\mathcal{F}^{(2)}(\mathbf{b})$ and $\mathcal{F}^{(2)}(\mathbf{b})$. Thus, some bids will use the sale price from $\text{CostShare}_{R(\mathbf{b})}$ and some will use the price from $\text{CostShare}_{R(\mathbf{b})/c}$. Lower bids will use the former (a higher price) and higher bids will use the latter (a lower price).

In the second case, $\mathcal{F}^{(2)}(\mathbf{b})$ is determined by the two highest bids. If b_i is not one of these bids, then $\mathcal{F}^{(2)}(\mathbf{b}_{-i}) = \mathcal{F}^{(2)}(\mathbf{b})$. The values of $\mathcal{F}^{(2)}$ when one or another of the two bids is removed are the same. It follows that there are only two possible values of $\mathcal{F}^{(2)}(\mathbf{b}_{-j})$, and therefore at most two distinct sale prices. ■

4.1 Random Reals vs. Random Bits. The CORE auction of the previous section needs to select a real-valued random u . To combine this auction with the Vickrey auction, we use another random real, p . Consider a more realistic model of computation that does not allow infinite-precision reals. In this case p and u must be rational. The CORE approach easily adapts to such a model.

Using Lemma 3.3, one can show the following result.

THEOREM 4.2. *For any $\epsilon > 0$, and with appropriate parameter settings, a CORE auction uses two random bits and is $(6 + \epsilon)$ -competitive against $\mathcal{F}^{(2)}$.*

Proof. If $\mathbf{b} \in \mathcal{B}_2$, we chose the Vickrey auction with probability $1/2$. In this case, the competitive ratio is four. For the rest of the proof we assume $\mathbf{b} \in \mathcal{B}_{3+}$; this case determines the competitive ratio.

For $\mathbf{b} \in \mathcal{B}_{3+}$, the revenue in the case when the Vickrey auction is selected is $\mathcal{F}^{(2)}(\mathbf{b})/k$. In the other case, our one random bit consensus estimate algorithm with $c' \approx \frac{3}{2}$ and $\rho = k/(k-1)$ gets an expected consensus value of $v/(2c'\rho) = \frac{v(k-1)}{3k}$. Thus the expected profit is approximately

$$\frac{\mathcal{F}^{(2)}(\mathbf{b})}{2} \left(\frac{1}{k} + \frac{k-1}{3k} \right) \geq \frac{\mathcal{F}^{(2)}(\mathbf{b})}{6}.$$

■

By using more random bits, we get better discrete approximation of the continuous distribution of u and of the optimal value of p of the previous section. With sufficiently many bits, we can get arbitrary close to the 3.39 competitive ratio.

4.2 CORE for Mass-Markets. In the context of mass-market goods, one may be able to assume that the number of items sold is large (e.g., guaranteed to be at least a thousand). In this context, it is possible to design (α, β) -competitive auctions that have relevant performance guarantees that are better than those of worst case auctions. In this section we introduce a parameterized variant of CORE, CORE_α . We assume that $\alpha \geq 2$ is an integer.

THEOREM 4.3. *For any $\alpha \geq 2$, there is β such that CORE_α is (α, β) -competitive and $\beta = 1 - O\left(\frac{1}{\sqrt{\alpha}}\right)$.*

Proof. Note that if $\mathcal{F}(\mathbf{b}) \geq \alpha h(\mathbf{b})$, then for any i we have $(\alpha - 1)\mathcal{F}(\mathbf{b})/\alpha \leq \mathcal{F}(\mathbf{b}_{-i}) \leq \mathcal{F}(\mathbf{b})$. The auction CORE_α computes, for every i , a consensus estimate on $\mathcal{F}(\mathbf{b})$ and the corresponding threshold price for i . Since $k \geq \alpha$, the value of ρ and the optimal value of c depends on α . For \mathbf{b} such that $\mathcal{F}(\mathbf{b}) \geq \alpha h(\mathbf{b})$, our consensus estimate solution for $\rho = \alpha/(\alpha - 1)$ gives expected revenue at least

$$(4.1) \quad \frac{\mathcal{F}(\mathbf{b})}{\ln c} \left(1 - \frac{1}{\alpha} - \frac{1}{c}\right).$$

Let $c = 1 + \frac{1}{\sqrt{\alpha}}$. Using the inequality $\ln(1 + x) \leq x$, we get the following bound on the expected revenue of CORE_α .

$$\begin{aligned} & \frac{\mathcal{F}(\mathbf{b})}{\ln c} \left(1 - \frac{1}{\alpha} - \frac{1}{c}\right) \\ & \geq \mathcal{F}(\mathbf{b}) \sqrt{\alpha} \left(1 - \frac{1}{\alpha} - \frac{1}{1 + 1/\sqrt{\alpha}}\right) \\ & \geq \mathcal{F}(\mathbf{b}) \left(-\frac{1}{\sqrt{\alpha}} + \frac{1}{1 + 1/\sqrt{\alpha}}\right) \\ & \geq \mathcal{F}(\mathbf{b}) \left(1 - \frac{1}{\sqrt{\alpha}} \left(1 + \frac{1}{1 + 1/\sqrt{\alpha}}\right)\right) \\ & \geq \mathcal{F}(\mathbf{b}) \left(1 - \frac{3}{2\sqrt{\alpha}}\right). \end{aligned}$$

■

Using equation (4.1) in a numeric simulation, we can compute near-optimal values of c and β for a given value of α . Table 1 gives values of β for several values of α . The data suggests that for mass-market goods, CORE_α performance may be acceptable.

α	β
2	5.3567
3	3.2833
10	1.7021
100	1.1602
1000	1.0465
10000	1.0144
100000	1.0045

Table 1: Competitive ratios of CORE_α for some values of α .

5 Concluding Remarks

The two key ingredients of a CORE auction are an estimator that computes a consensus lower bound on the optimal solution value in a bid-independent way, and a profit extractor that generates the revenue that is equal to the estimated value. Note that a profit extractor may be easier than a competitive auction to design, as the former has additional information, namely a good lower bound on the auction revenue. For the basic auction problem we address in this paper, the well-known cost sharing auction is a profit extractor. This general estimator-extractor approach applies beyond the basic auctions. As an example it has already been successfully applied to the double auction problem [10].

Our solution to the consensus problem is a natural one. We would like to know if this solution is optimal or if there is a better solution.

One desirable consumer “fairness” property for an auction is for winners to all pay a single price and for all bidders that bid above that price to win. More formally we say an auction outcome is *simple* if it has the following structure: there is a price p such that all bids above p win at price p , all bids below p lose, and bids tied with p either all win at price p or all lose. Recently with Anna Karlin, we have shown that no competitive auction always has a simple outcome. On the other hand, CORE auction outcomes are simple in the “normal” case when the auction achieves consensus.

The consensus problem has some similarity to coding theory (see, e.g., [9]), Byzantine agreement (see, e.g., [22, 30]), and to the private approximation problem, [12, 19]. We wonder if there is a deeper relationship with one of these areas.

References

- [1] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42nd IEEE Symp. on Foundations of Computer Science*, 2001.

- [2] A. Archer and E. Tardos. Frugal Path Mechanisms. In *Proc. 13th Symp. on Discrete Alg.* ACM/SIAM, 2002.
- [3] S. Baliga and R. Vohra. Market research and market design. Manuscript. Available via <http://www.kellogg.nwu.edu/faculty/vohra/htm/res.htm>, 2002.
- [4] Z. Bar-Yossef, K. Hildrum, and F. Wu. Incentive-compatible online auctions for digital goods. In *Proc. 13th Symp. on Discrete Alg.* ACM/SIAM, 2002.
- [5] A. Blum, V. Kumar, A. Rudra, and F. Wu. Online Learning in Online Auctions. In *Proc. 14th Symposium on Discrete Algorithms*. ACM/SIAM, 2003.
- [6] A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.
- [7] J. Bulow and J. Roberts. The Simple Economics of Optimal Auctions. *The Journal of Political Economy*, 97:1060–90, 1989.
- [8] E. H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 11:17–33, 1971.
- [9] T.M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley, New York, 1991.
- [10] K. Deshmukh, A. Goldberg, J. Hartline, and A. Karlin. Truthful and Competitive Double Auctions. In *Proc. 10th European Symposium on Algorithms*. Springer-Verlag, 2002.
- [11] S. DeVries and R. Vohra. Combinatorial Auctions: A survey. Unpublished manuscript, 2000.
- [12] J. Feigenbaum, Y. Ishai, T. Malkin, and K. Nissim. Secure Multiparty Computation of Approximations. In *Proc. Int. Col. on Automata, Languages and Prog.* Springer-Verlag, 2001.
- [13] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the Cost of Multicast Transmissions. In *Proc. of 32nd Symposium Theory of Computing*, pages 218–226. ACM Press, New York, 2000.
- [14] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive generalized auctions. In *Proc. 34rd ACM Symposium on the Theory of Computing*. ACM Press, 2002.
- [15] A. Goldberg, J. Hartline, A. Karlin, and A. Wright. Competitive Auctions. Submitted to *Games and Economic Behavior*, 2001.
- [16] A. Goldberg, J. Hartline, and A. Wright. Competitive Auctions and Digital Goods. In *Proc. 12th Symp. on Discrete Alg.*, pages 735–744. ACM/SIAM, 2001.
- [17] A. V. Goldberg and J. D. Hartline. Competitive Auctions for Multiple Digital Goods. In *Proc. 9th European Symposium on Algorithms*, pages 416–427. Springer-Verlag, 2001.
- [18] T. Groves. Incentives in Teams. *Econometrica*, 41:617–631, 1973.
- [19] S. Halevi, R. Krauthgamer, and E. Kushilevitz. Private Approximation of NP-Hard Functions. In *Proc. of 33st Symposium on Theory of Computing*. ACM Press, New York, 2001.
- [20] J. Hershberger and S. Suri. Vickrey Pricing in network routing: Fast payment computation. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001.
- [21] P. Klemperer. Auction theory: A guide to the literature. In *Journal of Economic Surveys*, pages 227–286. 13(3), 1999.
- [22] L. Lamport, R. Shostak, and M. Pease. The byzantine generals problem. *ACM Trans. on Prog. Lang. and Sys.*, 4:382–401, 1982.
- [23] A. Mas-Collel, M. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [24] H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs. *Economic Theory*, to appear.
- [25] R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [26] N. Nisan. Algorithms for Selfish Agents: Mechanism Design for Distributed Computation. 1999.
- [27] N. Nisan and A. Ronen. Algorithmic Mechanism Design. In *Proc. of 31st Symp. on Theory of Computing*. ACM Press, New York, 1999.
- [28] N. Nisan and A. Ronen. Computationally feasible VCG Mechanisms. In *ACM Conference on Electronic Commerce*, pages 242–252, 2000.
- [29] C. Papadimitriou. Algorithms, Games and the Internet. In *Proc. of 33rd Symposium Theory of Computing*, pages 749–752. ACM Press, 2001.
- [30] M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *JACM*, 27:228–234, 1980.
- [31] I. R. Segal. Optimal Pricing Mechanisms with Unknown Demand. *American Economic Review*, to appear.
- [32] D. D. Sleator and R. E. Tarjan. Amortized Efficiency of List Update and Paging Rules. *Communications of the ACM*, 28:202–208, 1985.
- [33] W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *J. of Finance*, 16:8–37, 1961.