Competitiveness via Consensus

Andrew V. Goldberg[1]
goldberg@microsoft.com

Jason D. Hartline[2]
hartline@cs.washington.edu

June 2002

Technical Report
MSR-TR-2002-73

We introduce Consensus Revenue Estimate (CORE) auctions. This is a class of competitive 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. A parameterized CORE auction performs better than the previous auctions in the context of mass-market goods, such as digital goods. The improved performance is due to the consensus estimate technique that allows more information to be extracted from the input. This technique is very natural and may be useful in other contexts.

Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
http://www.research.microsoft.com

[1]Microsoft Research, 1065 La Avenida, Mountain View, CA 94043.
[2]Computer Science Department, University of Washington.

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., [7, 10, 17, 20, 22, 32]), 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, 12, 19, 25, 26, 27, 28].

In this paper we study auctions, which are an important 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 [7, 17, 32]. This approach has been taken in [6, 24]. Auction design problems under weaker assumptions on the prior distribution knowledge has been studied in [3, 30]. In contrast to the Bayesian design and analysis framework, in [15, 16] we introduce the competitive analysis of auctions. This approach is motivated by on-line algorithms (see, e.g.,[5, 31]). *Competitive auctions* are truthful and for every input obtain a fraction of the revenue of an optimal auction.[1] 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 establish a framework for competitive auction design and open 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 [14], to the on-line case [4], and to more general mechanisms [13]. 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.

The best previous upper bound on the competitive ratio, achieved by the Sampling Cost Sharing auction [15], is four. A lower bound of two on the competitive ratio is given in [13].[2] 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

[1] We discuss optimal auctions in Section 2.

[2] With Michael Saks and Anna Karlin we have improved this lower bound to 2.42.

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_α, 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_α 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 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.

In this paper we introduce the class of *COnsensus Revenue Estimate auctions (CORE auctions)*. The auctions, based on a new consensus technique, work as follows. For each bid b_i, we delete the bid and use the other bids to compute a threshold price p_i for i with the following property. With high probability, the threshold prices for all auction winners i have the same value, p, and thus the auction outcome is simple. Furthermore, the resulting revenue is close to that of the optimal revenue.

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.
- 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_α, that achieves competitive ratios close to one for more realistic values of α than DSOT_α. 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.

This paper is organized as follows. Section 2 gives definitions and background. Section 3 is devoted to the consensus problem, our main tool for design of the CORE auctions. We discuss the CORE auctions in Section 4. Section 5 contains concluding remarks.

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* $\mathbf{b} = (b_1, \ldots, b_n)$*, 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 standard assumptions of* no positive transfers *and* voluntary participation*; see, e.g., [23]).*
- *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 has a private utility value, representing the true maximum they are willing to pay for an item. We denote by u_i bidder i's utility value.
- Bidders are *rational*, i.e., each bidder bids so as to maximize their *profit*, $u_i x_i - p_i$.
- Bidders bid with 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 [23] 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. For randomized auctions, we use the following notion of truthfulness. 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 $b_i = u_i$ unless mentioned otherwise.

Next we describe a useful characterization of truthful auctions using the notion of *bid-independent auctions*. We define deterministic bid-independent auctions first. Let $\mathbf{b}_{-i}$ denote the vector of bids $\mathbf{b}$ with b_i removed, i.e., $\mathbf{b}_{-i} = (b_1, \ldots, b_{i-1}, ?, b_{i+1}, \ldots, b_n)$. 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.

Theorem 2.3 [13] *An auction is truthful if and only if it is equivalent to a bid-independent auction.*

The proof of Theorem 2.3 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

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

$$\mathcal{F}(\mathbf{b}) = \max_{1 \leq k \leq n} kv_k.$$

An important result from [13] is that for all $\mathbf{b}$ no "reasonable" auction, even one that uses multiple prices, can obtain an expected revenue greater that $\mathcal{F}(\mathbf{b})$.

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.

[3] In 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.

Since $\mathcal{F}$ is an upper bound on auction revenue, we would like an auction to be competitive with $\mathcal{F}$ on every input. As shown in [15], however, if $\mathbf{b}$ is such that a single bidder's utility dominates the total utility of the other bidders, no auction can compete with $\mathcal{F}$ on $\mathbf{b}$. We use two notions of competitiveness for auctions: a general competitiveness for an auction on any set of bids [13] and notion of competitiveness suited for mass-markets where the number of winners of the optimal auction is large [14, 16]. In the latter case, the auctions 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 $\mathbf{b}$. 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 standard truthful cost sharing mechanism CostShare_C [23] and describe its implementation as a bid-independent auction. Theorem 2.3 guarantees that such an implementation exists.

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 cost sharing when the cost is a fixed value C if there are any winners and zero if there are no winners.

> CostShare_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.

Important properties of this auction are as follows:

- CostShare_C is truthful and symmetric.
- If $C \leq \mathcal{F}(\mathbf{b})$, CostShare_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 cs_C be this function.

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 2.4 *$\mathcal{A}$ 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 [32]. 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 2.5 *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, for a function g is $\gamma(v) = g(v)$ if g is a ρ-consensus estimate on v and $\gamma(v) = 0$ otherwise.

It is easy to see that there is no deterministic function g that is a ρ-consensus estimate for all v.

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

3.1 Consensus Estimate Algorithm

We describe a distribution $\mathcal{G}_\rho$ that works well in the following sense: for any v, $\mathbf{E}[g(v)] = \Omega(v)$, 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 :

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

Note that if g_u^c is a consensus, then this definition implies that g_u^c is a lower bound on v. Thus to show that g_u^c is a consensus estimate, it is sufficient to show that it is a consensus.

Remark. The definition of g_u^c implies that for any v, $\frac{v}{c} \leq g_u^c(v)$. This if g is a consensus, the consensus value is 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 from $[0, 1]$. We repeatedly make use of the following result.

Lemma 3.2 *For a fixed w such that $v/c \leq w \leq v$, $g(w)$ is distributed identically to $vc^{U'}/c$ for U' uniform on* $[0, 1]$.

Proof. Consider a random variable $Y = \log_c g(w)$ and let $t = \log_c v - 1$. Then $\mathbf{Pr}[Y \leq t + x] = \mathbf{Pr}[U \leq x]$ and therefore Y is uniformly distributed between t and $t + 1$. Thus, $g(w)$ is identical to vc^U/c. ■

For U uniform $[0, 1]$, the random variable c^U satisfies $\mathbf{Pr}\left[c^U \leq z\right] = \mathbf{Pr}[U \leq \log_c z] = \log_c z = \frac{\ln z}{\ln c}$. The probability density function for c^U is $\pi(x) = 1/(x \ln c)$ for $1 \leq x < c$. To see this, note that $\mathbf{Pr}\left[c^U \leq z\right]$ 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 3.3 *For u uniform on $[0, 1]$, the probability that g is a consensus estimate is* $1 - \log_c \rho$.

Proof. g is a consensus estimate of $g(v) \leq \frac{v}{\rho}$. Using Lemma 3.2, we get

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

■

In the application to auctions, the value of ρ is fixed and we chose c to maximize the expectation of γ. Lemma 3.3 implies $\mathbf{E}[\gamma(v)] \geq \frac{v}{c}(1 - \log_c \rho)$. The following theorem gives a better bound.

Theorem 3.4 *For $\mathcal{G}$ defined above, for all r,* $\mathbf{E}[\gamma(r)] = \frac{r}{\ln c}(\frac{1}{\rho} - \frac{1}{c})$.

Proof. By Lemma 3.2, $g(v)$ is distributed as vc^U/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 $\mathbf{E}[\gamma(v)]/v$.

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 values 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 $\gamma(v) \geq v/(2c'\rho)$. This gives the following lemma:

Lemma 3.5 *The consensus estimate solution $\mathcal{G}_\rho = \{g_0^{c'}, g_{\frac{1}{2}}^{c'}\}$ with $c' = \rho + \frac{\epsilon}{2\rho}$ gives $\mathbf{E}[\gamma(v)] = v/(2\rho^2 + \epsilon)$.*

4 CORE Auctions

In this sections we introduce the class of CORE auctions. Intuitively, CORE auctions compute consensus estimate r on $\mathcal{F}^{(2)}(\mathbf{b})$ from $\mathbf{b}_{-i}$'s and then run $\text{CostShare}_r(\mathbf{b})$. A formal description is somewhat more involved because we need to deal with the case when consensus is not achieved.

The following lemma is straight-forward but important.

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

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

For $k \geq 3$, the lemma allows us to estimate $\mathcal{F}^{(2)}(\mathbf{b})$ by $\mathcal{F}^{(2)}(\mathbf{b}_{-i})$. In combination with the results of Section 3, we obtain a consensus estimate on $\mathcal{F}^{(2)}(\mathbf{b})$ and use it in a bid-independent version of the cost-sharing auction.

Set $\rho = \frac{3}{2}$ and pick $c > a$. Let g be a function picked from a distribution described in Section 3. Let R be a function from masked bid vectors to reals defined by $R(\mathbf{b}_{-i}) = g(\mathcal{F}^{(2)}(\mathbf{b}_{-i}))$. Consider 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.

Remark. Note that it is not enough to find an R with some smoothness property such that $R(\mathbf{b}_{-i}) \approx R(\mathbf{b}_{-j})$ because CostShare_r and $\text{CostShare}_{r+\epsilon}$ can give very different outcomes. To see this, consider 99 bids at value one and and one bid at value 200. CostShare_{100} has 100 winners at price one while $\text{CostShare}_{100.1}$ has one winner at price 100.1. Thus, the we require that the values of $R(\mathbf{b}_{-i})$ be equal.

Lemma 3.3 and Lemma 4.1 imply that with probability $1 - \log_c \frac{3}{2}$, R computes a consensus estimate on $\mathcal{F}^{(2)}(\mathbf{b})$. In this case, our revenue is at least $\mathcal{F}^{(2)}(\mathbf{b})/c$, giving a competitive ratio of $\frac{1}{c}\left(1 - \log_c \frac{3}{2}\right)$. A near-optimal choice of $c = 2.4$ results in a bound of 4.5 on the competitive ratio for the auction $\mathcal{A}_{f_R}$ on $\mathbf{b} \in \mathcal{B}_{3+}$. Combining this auction with the Vickrey auction as described in Section 2.4, we obtain an 6.5-competitive auction.

This analysis can be improved by utilizing the following observations:

- The Vickrey auction is k-competitive on $\mathcal{B}_{3+}$. If we run the Vickrey action with probability p, this adds $p\mathcal{F}^{(2)}(\mathbf{b})/k$ to the expected revenue of the $\mathcal{B}_{3+}$ case.
- The analysis of $\mathcal{A}_{f_R}$ on $\mathbf{b}$ with $k \geq 3$ can be improved by using the expected bound given by Theorem 3.4.

To get the improved bound, we need to optimize the choice of the parameter c and the probability p. Recall that we run the Vickrey auction with probability p and the CORE auction with probability $1 - p$.

Theorem 4.2 *For an appropriate choice of c and p, the CORE auction is 3.39-competitive against* $\mathcal{F}^{(2)}(\mathbf{b})$.

Proof. Assume that $\mathbf{b} \in \mathcal{B}_{3+}$. For the CORE auction application, γ, defined in Section 3, is a lower bound on the auction revenue. Then Theorem 3.4 implies that if the 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

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

The expected revenue for $\mathbf{b} \in \mathcal{B}_2$ is at least $\mathcal{F}^{(2)}(\mathbf{b})\frac{p}{2}$ due to the Vickrey auction. We pick p and c to optimize and balance the two cases. 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 auction has a simple outcome.[4] Another property is as follows.

Lemma 4.3 *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).

[4]We conjecture that no competitive auction always has a simple outcome.

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.

Theorem 4.4 *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, the notion of an (α, β)-competitive auction may be appropriate and leads to better competitive ratios. In this section we introduce a parameterized variant of CORE, CORE$_\alpha$. We assume that $\alpha \geq 2$ is an integer.

Theorem 4.5 *For any $\alpha \geq 2$, there is β such that CORE$_\alpha$ 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$_\alpha$ 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 a and the optimal value of c depends on α. For $\mathbf{b}$

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.

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

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

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 (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.

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 bid-independent extractor auction that generates the revenue that is equal to (or a constant fraction of) the estimated value. Note that an 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 an extractor. This general estimator-extractor approach applies beyond the basic auctions. For example, in an upcoming paper we show how to use this approach in the context of the double-auction problem [9].

By improving the upper bound from 4 to 3.39, we have substantially narrowed the gap between the lower and the upper bounds on the auction competitive ratio. An interesting open question is to further narrow or completely close the gap. A related open question is if our consensus estimator is optimal or if there is a better one. The latter would lead to a better competitive ratio.

We know that a symmetric auction competitive against $\mathcal{F}^{(2)}$ must use at least one random bit, and that there is such an auction that uses two random bits. The question of existence of such an auction that uses a single fair coin flip remains open.

The consensus problem is related to several problems in Computer Science. Our original motivation came from coding theory (see, e.g., [8]), in particular error-correcting coding, where one of the approaches is for the decoder to find a valid codeword that is closest to the received message, which may contain an error. Another related problem is Byzantine agreement (see, e.g., [21, 29]), where several processors, some of which may be faulty, must agree on a common value. The closest relationship is to the private approximation problem, introduced in [11] (see also [18]), where a player i has a private value x_i, and players want to approximate the value of a given function, $f(x_i, \ldots, x_n)$, while revealing as little information about individual values x_i as possible. Although none of these other problems exactly matches our application, we wonder if there is a closer relationship between these or other problems in Computer Science and problems related to the design of competitive mechanisms.

Acknowledgments

We would like to thank Cynthia Dwork, Amos Fiat, Joan Feigenbaum, Anna Karlin, and Joe Kilian for helpful discussions, suggestions, and pointers to the literature.

References

[1] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42nd IEEE Symposium 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. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis.* Cambridge University Press, 1998.

[6] J. Bulow and J. Roberts. The Simple Economics of Optimal Auctions. *The Journal of Political Economy*, 97:1060–90, 1989.

[7] E. H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 11:17–33, 1971.

[8] T.M. Cover and J. A. Thomas. *Elements of Information Theory.* Wiley, New York, 1991.

[9] K. Deshmukh, A. V. Goldberg, J. D. Hartline, and A. R. Karlin. Truthful and Competitive Double Auctions. In *Proc. 10th European Symposium on Algorithms.* Springer–Verlag, 2002.

[10] S. DeVries and R. Vohra. Combinatorial Auctions: A survey. Unpublished manuscript, 2000.

[11] 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.

[12] 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.

[13] A. Fiat, A. V. Goldberg, J. D. Hartline, and A. Karlin. Comeptitive generalized auctions. In *Proc. 34rd ACM Symposium on the Theory of Computing*. ACM Press, 2002. To appear.

[14] 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.

[15] A. V. Goldberg, J. D. Hartline, A. Karlin, and A. Wright. Competitive Auctions. Submitted to *Games and Economic Behavior.*, 2001.

[16] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive Auctions and Digital Goods. In *Proc. 12th Symp. on Discrete Alg.*, pages 735–744. ACM/SIAM, 2001. Also available as InterTrust Technical Report STAR-TR-99.09.01, 1999, http://www.star-lab.com/tr/tr-99-01.html.

[17] T. Groves. Incentives in Teams. *Econometrica*, 41:617–631, 1973.

[18] 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.

[19] 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.

[20] P. Klemperer. Auction theory: A guide to the literature. In *Journal of Economic Surveys*, pages 227–286. 13(3), 1999.

[21] L. Lamport, R. Shostak, and M. Pease. The byzantine generals problem. *ACM Trans. on Prog. Lang. and Sys.*, 4:382–401, 1982.

[22] A. Mas-Collel, M. Whinston, and J. Green. *Microeconomic Theory.* Oxford University Press, 1995.

[23] H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs. *Economic Theory*, to appear.

[24] R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.

[25] N. Nisan. Algorithms for Selfish Agents: Mechanism Design for Distributed Computation. 1999.

[26] N. Nisan and A. Ronen. Algorithmic Mechanism Design. In *Proc. of 31st Symposium on Theory of Computing*, pages 129–140. ACM Press, New York, 1999.

[27] N. Nisan and A. Ronen. Computationally feasible VCG Mechanisms. In *ACM Conference on Electronic Commerce*, pages 242–252, 2000.

[28] C. Papadimitriou. Algorithms, Games and the Internet. In *Proc. of 33rd Symposium Theory of Computing*, pages 749–752. ACM Press, 2001.

[29] M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *JACM*, 27:228–234, 1980.

[30] I. R. Segal. Optimal Pricing Mechanisms with Unknown Demand. *American Economic Review*, to appear.

[31] D. D. Sleator and R. E. Tarjan. Amortized Efficiency of List Update and Paging Rules. *Communications of the ACM*, 28:202–208, 1985.

[32] W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *J. of Finance*, 16:8–37, 1961.