

Collusion-Resistant Mechanisms for Single-Parameter Agents

Andrew V. Goldberg Jason D. Hartline

Microsoft Research – Silicon Valley
1065 La Avenida, Mountain View, CA 94062
`{goldberg,hartline}@microsoft.com`

May 2004

Technical Report
MSR-TR-2004-40

We consider the problem of designing mechanisms with the incentive property that no coalition of agents can engage in a collusive strategy that results in an increase in the combined utility of the coalition. For single parameter agents, we give a characterization that essentially restricts such mechanisms to those that post a “take it or leave it” price to for each agent in advance. We then consider relaxing the incentive property to only hold with high probability. In this relaxed model, we are able to design approximate profit maximizing auctions and approximately efficient auctions. We also give a general framework for designing mechanisms for single parameter agents while maintaining the coalition incentive property with high probability. In addition, we give several results for a weaker incentive property from the literature known as *group strategyproofness*.

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

1 Introduction

A significant recent trend in theoretical computer science is to include economic and game theoretic considerations in the design of algorithms and protocols. This work considers the case where part of the input to the algorithm is provided by *selfish agents*. Not only must the algorithm produce a desired outcome, but it must do so in the presence of potential manipulation of its inputs by the selfish agents attempting to game the system in order to obtain outcomes that favor their own interests. In this body of research, the dominant paradigm is *truthful mechanism design*, and many problems have been considered. These include scheduling, e.g., [24, 2]; shortest paths, e.g., [3, 10]; minimum spanning trees [27]; digital good auctions [16]; online auctions [4, 6]; multicast auctions, e.g., [11, 19, 12]; and combinatorial auctions, e.g., [21, 25, 1]. An assumption made in most works in the area of truthful mechanism design, dating back to and including its origin in the Vickrey-Clarke-Groves (VCG) mechanism [28, 7, 18], is that the selfish agents do not collude with each other.

It is well acknowledged that collusion is a problem in mechanism design. There are many case studies of auctions run in practice where collusion has had significant effects (e.g., Treasury auctions [13, 17] and FCC spectrum auctions [9]). Robinson [26] gives theoretical evidence that demonstrates that the most celebrated truthful mechanism, the Vickrey auction, which sells to the highest bidder at the second highest price, is vulnerable to collusion. Nonetheless, theoretical work on understanding collusion and especially designing mechanisms that are resistant to collusion is limited [8, 20]. With the Internet and advances in cryptography which may make collusion easier and even more of a problem, it is especially important to understand collusion and how to prevent it.

Much of the recent work in designing mechanisms that resist collusion only considers a limited model in which agents might collude: a mechanism is *group strategyproof* (see e.g., [23, 11, 19]) if in any coalition, manipulation, i.e., a non-truth-telling strategy that strictly benefits some member of the coalition will also strictly hurt another member. Implicit in this definition is the assumption that an agent that benefited from the coalition strategy will not payoff an agent that suffered a loss. There are natural instances where collusion in group strategyproof mechanisms results in a negligible loss for one member of the coalition and substantial gain for another. We give one such example in Section 7 along with further motivation for moving to stronger notions of collusion resistance.

In this paper we propose a strong generalization the notion of truthfulness to include the possibility that groups of the agents may collude and exchange side-payments. We say a mechanism is *t-truthful* if truth-telling is an optimal strategy for all agents even when it is possible for agents to form coalitions of size t or less and possibly exchange side-payments with other agents in their coalition. This is equivalent to requiring that non-truth-telling does not result in an increase the total utility of the agents in a coalition. With this strong notion of collusion resistance, it is not necessary to have a model (e.g., as in [22]) for how agents might be able to agree on a collusive strategy because rational agents will have no incentive to form coalitions.

We will consider mechanisms for any *single-parameter agent* problem (see e.g., [21, 2, 1]). In a single-parameter agent problem each agent has a publicly known partitioning of possible outcomes into two sets, the *reject* set and the *accept* set. It is assumed that agent i has valuation zero for any outcome in the reject set and private valuation v_i for any outcome in the accept set. For auction-like problems, agent i 's accept set is simply the set of allocations where agent i is allocated their desired good and the reject set is the set of allocations where i is not allocated their desired good. The truth-telling strategy for agent i would be to report to the mechanism (a.k.a. bid) their true private value, v_i .

As our main impossibility result (Section 3), we will characterize t -truthful mechanisms for $t \geq 2$ as precisely those mechanisms that offer each agent a “take it or leave it” price that is independent of the bids of all agents.¹ We refer to these mechanisms as *posted-price* mechanisms. This result implies that, without any advance knowledge of the input bids, no t -truthful mechanism can optimize any nontrivial objective function. Furthermore, it implies that the standard objectives of profit maximization and economic efficiency are impossible to approximate even in the least restrictive scenarios (Section 4).

¹We view this as an impossibility result because it implies that no “nontrivial” mechanism is t -truthful.

In spite of this impossibility result, we would still like to obtain mechanisms that prevent the agents from effectively colluding. To this end, we follow the approach of [1, 14] and relax the definition of t -truthfulness to allow for mechanisms that are only t -truthful with high probability. Such a mechanism would produce an outcome such that with high probability in coin flips made by the mechanism, no coalition can obtain a higher total utility by non-truthful bidding.

As our main positive result, we use the *consensus estimate* technique from [15] to obtain mechanisms that are t -truthful with high probability and either approximate profit maximization (Section 5) or economic efficiency (Section 6) for single-item multi-unit auctions. These results follow from a general approach which may be useful in obtaining mechanisms that are t -truthful with high probability for other mechanism design problems. The approach is based on using the consensus estimate technique to approximate summary information in a way that with high probability is non-manipulable by a coalition of t agents. This approximate summary information can then be used in a parameterized posted-price mechanism that with high probability results in a feasible near-optimal solution. For this approach to be successful, there must exist summary information that changes very little as a function of the bids of any coalition of t agents. There also must be a posted price mechanism that can use approximate summary information to obtain a near optimal outcome.

We conclude the paper by reconsidering the standard notion of group strategyproofness (Section 7). We study the interplay between group strategyproofness and randomness. We also modify the k -Vickrey auction to obtain a randomized auction that approximates Vickrey’s outcome. The resulting auction approximates Vickrey auction’s efficiency and profit maximization properties and is group strategyproof. Nonetheless, the auction does not seem to be much more resistant to collusion than the Vickrey auction, which is known to be vulnerable [26, 8, 13, 17]. This is additional evidence supporting the theory that group strategyproofness is not a sufficient property to prevent collusion, and motivating notions such as t -truthfulness with high probability.

2 Definitions

We adopt the general model of *single-parameter agents* (e.g., [21, 2, 1]). For single-parameter agents, the possible outcomes of the mechanisms can be partitioned into two sets *accept* and *reject*. It is assumed that these partitionings are public knowledge. We will let x_i be an indicator for whether the outcome is in agent i ’s accept ($x_i = 1$) or reject ($x_i = 0$) set. Each agent has a private value v_i representing the difference in its value for the reject and accept outcomes. Thus, an agent’s valuation for outcomes with $x_i = 1$ is v_i and zero otherwise. A mechanism for single parameter agents will compute an outcome, which we denote by $\mathbf{x} = (x_1, \dots, x_n)$ and prices $\mathbf{p} = (p_1, \dots, p_n)$. After the mechanism is run, agent i will be required to pay the mechanism p_i . Agent i ’s *utility* is given by the difference between their valuation and their payment, i.e., $u_i = x_i v_i - p_i$. We assume that each agent’s goal is to maximize their utility.

For structured optimization problems, the mechanism may incur a cost for producing a given outcome. We assume that there is a *cost function*, $c(\cdot)$, on outcomes, \mathbf{x} , that specifies this cost. Additionally, some outcomes may be infeasible. We will assume that the cost of infeasible outcomes is infinity.

To make this formulation more concrete, here are a few examples of auction related problems cast into this framework. By taking $c(\mathbf{x}) = 0$ for all \mathbf{x} , we obtain the *unlimited supply auction* problem which models the problem of selling identical units of a single item to bidders that each want at most one unit (e.g., for a digital good) [16]. Similarly, the k -unit *limited supply auction* has $c(\mathbf{x}) = 0$ if $\sum_i x_i \leq k$ and $c(\mathbf{x}) = \infty$ otherwise. The *non-excludable public good* problem (with fixed cost C) has $c(\mathbf{x}) = 0$ when $x_i = 0$ for all i , $c(\mathbf{x}) = C$ when $x_i = 1$ for all i , and $c(\mathbf{x}) = \infty$ otherwise. The *single-parameter combinatorial auction* [1] has $c(\mathbf{x}) = 0$ for all feasible allocations, i.e., ones that do not over-allocate any items, and ∞ otherwise.

The *profit* of a mechanism is simply the difference between the payments made to the mechanism by the agents and the cost incurred by the mechanism for the outcome selected. For prices \mathbf{p} and allocation \mathbf{x} :

$$\text{Profit} = \sum_i p_i - c(\mathbf{x}).$$

The economic *efficiency* of a mechanism is the social welfare of the solution it produces. This is the difference between the sum of the valuations of the agents and the cost. For valuations \mathbf{v} and allocation \mathbf{x} :

$$\text{Efficiency} = \sum_i x_i v_i - c(\mathbf{x}).$$

The outcome that maximizes the efficiency is said to be *efficient*. The two economic objectives we will consider in this paper *profit maximization* and *efficiency maximization*.

The mechanisms we consider are single-round, sealed bid mechanisms. The mechanism takes as input bids, $\mathbf{b} = (b_1, \dots, b_n)$, and a cost function $c(\cdot)$ and computes an outcome \mathbf{x} and prices \mathbf{p} . We will make the standard assumptions of *no positive transfers* and *voluntary participation* which together imply that $p_i = 0$ if $x_i = 0$ and $p_i \leq b_i$ otherwise, see e.g., [23].

We will be considering randomized mechanisms. In a randomized mechanism the allocation, $\mathbf{x} = (x_1, \dots, x_n)$, and prices, $\mathbf{p} = (p_1, \dots, p_n)$, are random variables. Therefore the agent utilities, $u_i = v_i x_i - p_i$, are also random variables. We denote by \bar{x}_i , \bar{p}_i , and \bar{u}_i be the expected value of x_i , p_i , and u_i respectively. Note that because x_i is an indicator variable, \bar{x}_i is simply the probability that agent i is accepted.

The following notation will be convenient. Given a mechanism in general, x_i and thus \bar{x}_i can be a function of the entire input \mathbf{b} . Let \mathbf{b}_{-i} denote the bid vector with bid i replaced by a '?', i.e., $\mathbf{b}_{-i} = (b_1, \dots, b_{i-1}, ?, b_{i+1}, \dots, b_n)$. If we fix \mathbf{b}_{-i} we can view \bar{x}_i as a function of b_i which we will denote by $\bar{x}_i^{(\mathbf{b}_{-i})}(b_i)$ (and similarly for $\bar{p}_i^{(\mathbf{b}_{-i})}(b_i)$ and $\bar{u}_i^{(\mathbf{b}_{-i})}(b_i)$). If \mathbf{b}_{-i} is implicit, we will just write $\bar{x}_i(b_i)$, $\bar{p}_i(b_i)$, etc.

We now discuss the incentive properties that will be the focus of this paper. These incentive properties are useful because analysis of mechanisms that do not have these properties involves making assumptions about agents' strategies and the agents' knowledge of other agent strategies, assumptions that we do not wish to make here. More motivation can be found, for example, in [24, 2, 16].

Definition 2.1 (Truthful in Expectation) *A randomized mechanism is truthful in expectation if for any agent i , regardless of the actions of any other agents, agent i 's expected utility, \bar{u}_i , is at its maximum when agent i bids its true valuation, $b_i = v_i$. Formally, for all \mathbf{b}_{-i} and b_i :*

$$\bar{u}_i^{(\mathbf{b}_{-i})}(v_i) \geq \bar{u}_i^{(\mathbf{b}_{-i})}(b_i).$$

Theorem 2.1 *(See e.g. [2]) For \mathbf{b}_{-i} fixed, a mechanism is truthful in expectation if and only if the probability that an agent wins as a function of its bid, $\bar{x}_i(b_i)$, is monotonically increasing and the payment rule satisfies:²*

$$\bar{p}_i(b_i) = b_i \bar{x}_i(b_i) - \int_0^{b_i} \bar{x}_i(b) db.$$

In the deterministic case the definition of truthfulness reduces to the following:

Definition 2.2 (Deterministic Truthfulness) *A deterministic mechanism is truthful if for any agent i , regardless of the actions of any other agents, agent i 's utility, u_i , is at its maximum when agent i bids its true valuation, $b_i = v_i$.*

Definition 2.3 (Randomized Truthfulness) *A randomized mechanism is truthful if it is a probability distribution over deterministic truthful mechanisms.*

We now describe our extensions of the standard notions of truthfulness to take into account the possibility that agents may collude. This definition captures the case where the agents may exchange side payments to insure that the collusive strategy is beneficial for all participants of the coalition.

Definition 2.4 (t -Truthfulness) *A mechanism is t -truthful in expectation if, for any coalition of size t and any value of the bids of agents not in the coalition, the sum of the expected utilities of the agents in the coalition is maximized when all agents in the coalition bid their true valuations.*

²This implies that give the allocation rule, the payment rule is fixed.

A deterministic mechanism is t -truthful if it is t -truthful in expectation, and a randomized mechanism is t -truthful if it is a distribution over deterministic t -truthful mechanisms.

The main impossibility result of this paper is a characterization of t -truthful mechanisms for $t \geq 2$. This characterization is most intuitive in the deterministic case where it simply means that prior to obtaining the bids, the mechanism must post “take it or leave it” prices for each agent. Each agent will then accept or reject its posted price depending on whether it is above or below its valuation. This is generalized by the following definition.

Definition 2.5 (Posted Price) *A mechanism is posted price in expectation if for all i , b_i , \mathbf{b}_{-i} , and \mathbf{b}'_{-i} ,*

$$\bar{x}_i^{(\mathbf{b}_{-i})}(b_i) = \bar{x}_i^{(\mathbf{b}'_{-i})}(b_i).$$

Again, a deterministic mechanism is posted price if it is posted price in expectation, and a randomized mechanism is posted price if it is a randomization over deterministic posted price mechanisms.

The main positive result of the paper is to show that it is possible to construct interesting mechanisms that are collusion resistant for a relaxed definition of t -truthfulness. We follow Archer et al. [1] and make the following definition. Here *high probability* is defined in some parameter of the input, e.g., the number of winners.

Definition 2.6 (t -Truthful with High Probability) *A randomized mechanism is t -truthful with high probability if given any coalition C of size t or fewer, with high probability in the randomness of the mechanism, there is no non-truthtelling strategy for the agents in C that increases their total utility.*

3 The Characterization

In this section we present the characterization of t -truthful mechanisms as posted price. Note that the profit of the mechanism is irrelevant for the incentive properties. Thus we can ignore $c(\cdot)$ and just consider the allocation and the prices that the mechanism produces.

First we show that any posted price in expectation mechanism is t -truthful in expectation (for all t). Then we show that any mechanism that is 2-truthful in expectation is posted price in expectation. This collapses the hierarchy of t -truthful mechanisms showing that any mechanism that is 2-truthful is also t -truthful for all t (the other direction is trivial). Furthermore, it characterizes mechanisms that are t -truthful in expectation as being posted price in expectation.

Lemma 3.1 *Any posted price in expectation mechanism is t -truthful in expectation.*

Proof: Consider agent i in coalition C . Since the mechanism is posted price this agent’s bid does not affect the (expected) utilities of any other agents in the coalition. Thus, to maximize the total (expected) utility of the coalition, this agent should maximize its own (expected) utility. Since any posted price in expectation mechanism is truthful in expectation, this maximum can be achieved by bidding truthfully. The same argument holds for all agents in the coalition and thus any posted price in expectation mechanism is t -truthful in expectation. \square

Lemma 3.2 *Any mechanism that is 2-truthful in expectation is posted price in expectation.*

Proof: Let \mathcal{M} be any mechanism that is 2-truthful in expectation. Consider an agent i and fix the values of all other bids to be $\mathbf{b}_{-i} = \mathbf{v}_{-i}$. Consider the expected utility of agent i , \bar{u}_i , as a function of its bid b_i . From Theorem 2.1 we have:

$$\begin{aligned} \bar{u}_i(b_i) &= v_i \bar{x}_i(b_i) - \bar{p}_i(b_i). \\ &= v_i \bar{x}_i(b_i) - b_i \bar{x}_i(b_i) + \int_0^{b_i} \bar{x}_i(b) db. \end{aligned}$$

Note that for $b_i = v_i$ this simplifies to

$$\bar{u}_i(v_i) = \int_0^{v_i} \bar{x}_i(b) db. \quad (1)$$

If agent i bids $b_i \neq v_i$, its expected loss due to making a suboptimal bid is

$$\begin{aligned} i\text{-loss}_{v_i}(b_i) &= \bar{u}_i(v_i) - \bar{u}_i(b_i) \\ &= \int_0^{v_i} \bar{x}_i(b) db - v_i \bar{x}_i(b_i) + b_i \bar{x}_i(b_i) - \int_0^{b_i} \bar{x}_i(b) db \\ &= (b_i - v_i) \bar{x}_i(b_i) - \int_{v_i}^{b_i} \bar{x}_i(b) db. \end{aligned} \quad (2)$$

Now consider some agent j , $j \neq i$. The probability that j is allocated an item, in general, is a function of all the bids. Keeping $\mathbf{b}_{-\{i,j\}}$ (all bids but bid i and bid j) fixed, we write the probability that j is accepted as a function of b_i and b_j as $\bar{x}_j^{(b_i)}(b_j)$ (and likewise agent j 's utility as $\bar{u}_j^{(b_i)}(b_j)$). In the discussion that follows we will assume that agent j bids $b_j = v_j$. From Equation (1) we have,

$$\bar{u}_j^{(b_i)}(v_j) = \int_0^{v_j} \bar{x}_j^{(b_i)}(b) db.$$

Consider agent j 's expected gain when agent i bids b_i instead of v_i (this ‘‘gain’’ may be negative),

$$\begin{aligned} j\text{-gain}_{v_i}(b_i) &= \bar{u}_j^{(b_i)}(v_j) - \bar{u}_j^{(v_i)}(v_j) \\ &= \int_0^{v_j} \left[\bar{x}_j^{(b_i)}(b) - \bar{x}_j^{(v_i)}(b) \right] db. \end{aligned} \quad (3)$$

Since we have assumed that the mechanism \mathcal{M} is 2-truthful, it must be that the combined expected utility of agent i and j when bidding (v_i, v_j) is at least their utility when bidding (b_i, v_j) . For all v_i , b_i , and v_j ,

$$\begin{aligned} j\text{-gain}_{v_i}(b_i) &\leq i\text{-loss}_{v_i}(b_i) \\ &\leq i\text{-loss}_{v_i}(b_i) + i\text{-loss}_{b_i}(v_i) \\ &= [b_i - v_i] [\bar{x}_i(b_i) - \bar{x}_i(v_i)]. \end{aligned}$$

In the above sequence of equations, $i\text{-loss}_{b_i}(v_i)$, intuitively represents the loss for agent i would incur by non-truthfully bidding v_i when their actual valuation is b_i . This quantity is always positive and is given by equation (2).

Equation 3 implies that for any bid b'_i , $j\text{-gain}_{v_i}(b_i) = j\text{-gain}_{v_i}(b'_i) + j\text{-gain}_{b'_i}(b_i)$. For the remainder of this argument we assume that $b_i > v_i$; the analogous argument can be made for the $v_i > b_i$ case. If we divide the interval from $[v_i, b_i]$ in to k equally sized segments, $b_i^{(0)}, \dots, b_i^{(k)}$ then the total gain for agent j when agent i bids b_i instead of v_i is the sum of the gains for each segment:

$$\begin{aligned} j\text{-gain}_{v_i}(b_i) &= \sum_{\ell=1}^k j\text{-gain}_{b_i^{(\ell-1)}}(b_i^{(\ell)}) \\ &\leq \sum_{\ell=1}^k \left[b_i^{(\ell-1)} - b_i^{(\ell)} \right] \left[\bar{x}_i(b_i^{(\ell-1)}) - \bar{x}_i(b_i^{(\ell)}) \right] \\ &= \frac{1}{k} [b_i - v_i] [\bar{x}_i(b_i) - \bar{x}_i(v_i)]. \end{aligned}$$

This holds for any k , and in the limit we get:

$$j\text{-gain}_{v_i}(b_i) \leq 0.$$

Of course, $j\text{-gain}_{v_i}(b_i) = -j\text{-gain}_{b_i}(v_i)$ and which implies that $j\text{-gain}_{v_i}(b_i)$ is zero. Because this is true for all values of agent j 's valuation, v_j , it must be that $\bar{x}_j^{(b_i)}(v_j) = \bar{x}_j^{(b'_i)}(v_j)$ for all b_i and b'_i .

We complete the proof by showing that for any \mathbf{b} and \mathbf{b}' , $\bar{x}_i^{(\mathbf{b}-i)}(\cdot) = \bar{x}_i^{(\mathbf{b}'-i)}(\cdot)$. The above argument shows that this result holds for \mathbf{b} and \mathbf{b}' that differ in one bid value. The transitivity of equality gives the general result. \square

Combining the two lemmas we have:

Theorem 3.3 *A mechanism is 2-truthful in expectation if and only if it is posted price in expectation.*

Corollary 3.4 *A mechanism is t -truthful in expectation for any $t \geq 2$ if and only if it is posted price in expectation.*

The analogous results follow for deterministic t -truthful mechanisms and randomized t -truthful mechanisms.

4 Lower Bounds

In this section we use the characterizations proved in previous section to show that the natural objectives of profit maximization and efficiency are impossible to approximate with a t -truthful mechanism for the simplest interesting cost functions.

4.1 Profit Maximization

First we consider the objective of profit maximization. Recall that the goal of profit maximization is to maximize $\sum_i p_i - c(\mathbf{x})$. Consider the case of the unlimited supply auction problem, i.e., the cost function that satisfies $c(\mathbf{x}) = 0$ for all \mathbf{x} . Arguably this is the simplest problem we could consider. Following [16], we would like to obtain an auction that approximates the profit of the optimal single price sale, OPT. We prove that no t -truthful mechanism can obtain better than a $\log h$ factor approximation even when all agent valuations are known to be within $[1, h]$.

Theorem 4.1 *The worst case expected profit of any posted price in expectation auction is at most $\text{OPT} / \ln h$ even if the bids are guaranteed to be within the interval $[1, h]$.*

Proof: Let V be chosen randomly from the distribution with $\Pr[V > v] = 1/v$ for $v \in [1, h]$ with a point mass $\Pr[V = h] = 1/h$. Let \mathbf{b} be a random input chosen with all n bids $b_i = V$. The expected optimal profit on this input is $\mathbf{E}[\text{OPT}] = n\mathbf{E}[V] = n(1 + \ln h)$.

Claim: For agent i , the expected payment given any posted price in $[1, h]$ is 1 (This is easy to see for any randomized posted price auction, for a posted price in expectation auction see below for a proof).

The claim implies that any posted price mechanism has an expected profit of n . Thus, there exists an input \mathbf{b} such that the auction profit is a $\ln h$ factor from OPT in expectation. \square

Proof of Claim: Given the input in the range $[1, h]$ we can assume without loss of generality that the auction satisfies $\bar{x}_i(1) = 0$ and $\bar{x}_i(h) = 1$. Any auction that does not satisfy this condition can be modified to satisfy it without decreasing the auction's profit.

Consider agent i with a random valuation V such that $\Pr[V > v] = 1/v$. The probability mass function for the agent's value is $1/v^2$ for $v \in [1, h]$ with a point mass with weight $1/h$ at value h . Using Theorem 2.1,

we get:

$$\begin{aligned}
\mathbf{E}[\bar{p}_i(V)] &= \int_1^h \frac{\bar{p}_i(v)}{v^2} dv + \frac{\bar{p}_i(h)}{h} \\
&= \int_1^h \frac{\bar{x}_i(v)}{v} dv - \int_1^h \int_1^v \frac{\bar{x}_i(b)}{v^2} db dv + \frac{1}{h} \left[\bar{x}_i(h)h - \int_1^h \bar{x}_i(b)db \right] \\
&= \int_1^h \frac{\bar{x}_i(v)}{v} dv - \int_1^h \bar{x}_i(b) \int_v^h v^{-2} dv db + 1 - \frac{1}{h} \\
&= \int_1^h \frac{\bar{x}_i(v)}{v} dv - \int_1^h \bar{x}_i(b) \left[\frac{1}{v} - \frac{1}{h} \right] db + 1 - \frac{1}{h} \\
&= 1.
\end{aligned}$$

4.2 Efficiency

Now we consider the problem of obtaining an approximately efficient outcome via a posted price mechanism. Recall that economic efficiency means that the social welfare, defined as $\sum_i x_i v_i - c(\mathbf{x})$, is maximized. Note that for the unlimited supply auction problem considered above, the trivial mechanism that allocates to all agents at price zero is both posted price and efficient. Instead we consider the next simplest problem, the limited supply auction problem (with k units). For the k -item auction, the cost function satisfies $c(\mathbf{x}) = 0$ for \mathbf{x} with at most k winners and $c(\mathbf{x}) = \infty$ otherwise. We show the following impossibility result.

Theorem 4.2 *No posted price in expectation mechanism for the limited supply auction problem gives a constant approximation to the efficient outcome (in expectation).*

Proof: We give this proof for the case that $k = 1$, for the general case apply the same argument duplicating each bidder k times. Assume for a contradiction that there is an auction \mathcal{A} that is posted price in expectation and gives a c -approximation to the efficient outcome for some integer c . Define the $(c + 1)$ -bid input $\mathbf{b}^{(i)}$ with $b_i = 1$ and all other bids $b_j = 0$. Since \mathcal{A} is c -approximation, on $\mathbf{b}^{(i)}$ we have $\bar{x}_i(1) \geq 1/c$. Now consider the $(c + 1)$ -bid “all ones” input, $\mathbf{b} = (1, \dots, 1)$. Since \mathcal{A} is posted price in expectation, $\bar{x}_i^{(\mathbf{b}^{(i)})}(1) = \bar{x}_i^{(\mathbf{b} - \mathbf{b}^{(i)})}(1)$. Thus, $\sum_i \bar{x}_i^{(\mathbf{b})}(1) \geq (c + 1)/c > 1$. Since expected number of items sold in $\mathcal{A}(\mathbf{b})$ is more than the number of items available there is non-zero probability of oversale which gives an efficiency of negative infinity and contradicts the assumption that \mathcal{A} is a c -approximation. \square

5 Approximate Profit Maximization

We now consider the problem of profit maximization in an unlimited supply auction, i.e., the number of items for sale is at least the number of agents, n . Nonetheless, it is trivial to adapt the results given below for the limited supply case when $k < n$ items are available. We present an auction that approximates profit maximization and is t -truthful with high probability. First a few definitions.

Definition 5.1 *The number of bids in \mathbf{b} with value at least x is denoted $\#_x(\mathbf{b})$.*

Definition 5.2 *Given a set $S \subset \mathbb{R}$, and a real number x , the value $\lfloor x \rfloor_S$ is the largest value $x' \in S$ such that $x' \leq x$, i.e., x rounded down to the nearest element of S . Similarly, the value $\lceil x \rceil_S$ is x rounded up to the nearest element of S .*

Definition 5.3 (APM) *Given input \mathbf{b} and parameters c and α , the Approximate Profit Maximization (APM) auction does as follows:*

1. Sample y uniformly from $[0, 1]$.

2. Let $n_i = \lfloor \#_{\alpha^i}(\mathbf{b}) \rfloor_{\{c^{j+y} : j \in \mathbb{Z}\}}$.
3. Output price α^i that maximizes $\alpha^i n_i$.

The following lemma follows immediately from the definition of APM.

Lemma 5.1 *APM has a worst case profit (over input and random coin flips of the mechanism) of at least factor of $c\alpha$ from the optimal single price sale, OPT.*

Definition 5.4 *For a fixed input \mathbf{b} and a fixed choice of y we say APM has a t -consensus at price α^i if*

$$\lfloor \#_{\alpha^i}(\mathbf{b}) - t \rfloor_{\{c^{j+y} : j \in \mathbb{Z}\}} = \lfloor \#_{\alpha^i}(\mathbf{b}) + t \rfloor_{\{c^{j+y} : j \in \mathbb{Z}\}}.$$

To get some intuition for the importance of this definition, note that if APM is a t -consensus at α^i then it is not possible for any set of t agents to change their bids and get \mathbf{b}' such that $\lfloor \#_{\alpha^i}(\mathbf{b}) \rfloor_{\{c^{j+y} : j \in \mathbb{Z}\}} \neq \lfloor \#_{\alpha^i}(\mathbf{b}') \rfloor_{\{c^{j+y} : j \in \mathbb{Z}\}}$.

Lemma 5.2 [15] *The probability that APM is a t -consensus at price α^i is*

$$1 - \log_c \frac{\#_{\alpha^i}(\mathbf{b}) + t}{\#_{\alpha^i}(\mathbf{b}) - t}.$$

Definition 5.5 *A price α^i is relevant if for some random coin flips in APM and for some set C of t colluders,*

1. APM can output price α^i , and
2. price α^i is lower than the sale price had these t colluders bid truthfully.

We will be showing that with high probability APM obtains a t -consensus simultaneously on all relevant prices. Note that the second part of the definition is important because it is always possible for a coalition to raise the price APM outputs (by bidding ∞).

Lemma 5.3 *For input \mathbf{b} , let α^r be the largest relevant price. Let $m = \#_{\alpha^r}(\mathbf{b}) - t$. Then for any relevant α^i ,*

$$\#_{\alpha^i}(\mathbf{b}) \geq \alpha^{r-i} m / c - t.$$

Proof: For any coalition of t agents manipulating \mathbf{b} to be \mathbf{b}' and any randomization in APM, i.e., any choice of y we have:

$$\alpha^r \lfloor \#_{\alpha^r}(\mathbf{b}') \rfloor_{\{c^{j+y} : j \in \mathbb{Z}\}} \geq \alpha^r m / c \quad \alpha^i \lfloor \#_{\alpha^i}(\mathbf{b}') \rfloor_{\{c^{j+y} : j \in \mathbb{Z}\}} \leq \alpha^i (\#_{\alpha^i}(\mathbf{b}) + t).$$

Thus, for α^i to be relevant, it must be that

$$\alpha^r m / c \leq \alpha^i (\#_{\alpha^i}(\mathbf{b}) + t).$$

Thus,

$$\alpha^{r-i} m / c - t \leq \#_{\alpha^i}(\mathbf{b}).$$

□

Definition 5.6 *For a fixed input \mathbf{b} and a fixed choice of y we say APM is a t -consensus if for all relevant prices α^i , APM is a t -consensus for α^i .*

Notice that for \mathbf{b} and y such that APM is a t -consensus, any coalition of t or fewer agents cannot change their bid values and cause the sale price output by APM to lower. We now prove that the probability that APM is a t -consensus approaches one as the number of agents bidding above the highest relevant price increases. In our proof we will make use of the following fact:

Fact 1 For $0 < a_i < 1$ and $1 \leq i \leq N$, $\prod_{i=1}^N (1 - a_i) \geq 1 - \sum_{i=1}^N a_i$.

Lemma 5.4 Let α^r be the highest relevant price and $m = \#\alpha^r(\mathbf{b}) - t$. The probability that APM is a t -consensus is $1 - \Theta(t/m)$.

Proof: Let $R = \{i : \alpha^i \text{ is relevant}\}$. Consider the probability APM is a t -consensus at α^i for $i \in R$ and apply Lemma 5.3.

$$\begin{aligned} \Pr[t\text{-consensus at } \alpha^i] &= 1 - \log_c \frac{\#\alpha^i(\mathbf{b})+t}{\#\alpha^i(\mathbf{b})-t} \\ &\geq 1 - \log_c \frac{\alpha^{r-i}m/c}{\alpha^{r-i}m/c-2t} \\ &= 1 + \log_c \left(1 - \frac{2tc}{\alpha^{r-i}m}\right). \end{aligned}$$

We want to get a bound on the probability that APM is a t -consensus simultaneously at α^i for all $i \in R$. To do this, we look at the probability that APM is not a t -consensus at each relevant α^i and use the union bound:

$$\begin{aligned} \Pr[t\text{-consensus}] &\geq 1 - \sum_{i \in R} \Pr[\text{not } t\text{-consensus at } \alpha^i] \\ &\geq 1 + \sum_{i \in R} \log_c \left(1 - \frac{2ct}{\alpha^{r-i}m}\right) \\ &\geq 1 + \sum_{j=0}^{\infty} \log_c \left(1 - \frac{2ct}{m} \alpha^{-j}\right) \\ &= 1 + \log_c \left[\prod_{j=0}^{\infty} \left(1 - \frac{2ct}{m} \alpha^{-j}\right) \right]. \end{aligned}$$

Applying Fact 1,

$$\begin{aligned} \Pr[t\text{-consensus}] &\geq 1 + \log_c \left[1 - \sum_{j=0}^{\infty} \frac{2ct}{m} \alpha^{-j} \right] \\ &= 1 + \log_c \left[1 - \frac{2ct}{m} \left(\frac{\alpha}{\alpha-1} \right) \right] \\ &= 1 - \Theta(t/m). \end{aligned}$$

This gives the theorem. □

Lemma 5.5 If APM is a t -consensus then for any coalition of t agents, truthful bidding maximizes the total utility of the coalition.

Proof: Note that APM uses a single sale price for all agents. Agents bidding at least this sale price win and agents below the sale price lose. In the case of a t -consensus, no group of t colluders can change their bid values to lower the sale price. This follows from the definition of t -consensus and relevant prices and implies the theorem. □

Theorem 5.6 For \mathbf{b} such that there are $m + t$ agents above the highest relevant price, the probability that APM is t -truthful is $1 - \Theta(t/m)$.

Unfortunately, this does not directly give us an auction that is t -truthful with high probability on all inputs. In particular, on inputs where m is of the same order as t , the theorem is useless, even if there is some $m' \gg t$ such that selling m' items also gives a large profit. This problem is easily addressed by parameterizing APM to only consider solutions with a large number of winners. A parameterized version of APM is defined as follows.

Definition 5.7 (APM $_\ell$) Given input \mathbf{b} and parameters α , ℓ , and c , APM $_\ell$ works as follows:

1. Sample y uniformly from $[0, 1]$.
2. Let $n_i = \lfloor \#_{\alpha^i}(\mathbf{b}) \rfloor_{\{c^{j+y} : j \in \mathbb{Z}\}}$.
3. Output price α^i that maximizes $\alpha^i n_i$ subject to the constraint that $n_i \geq \ell$.

It is easy to see that in APM $_\ell$, $m \in \Omega(\ell)$, which implies the following result.

Theorem 5.7 The probability that APM $_\ell$ is t -truthful is $1 - O(t/\ell)$.

Setting the parameter ℓ higher leads to the higher probability of t -truthfulness at the expense of possibly missing an optimal solution that sells to less than ℓ agents.

6 Approximate Efficiency

We now consider the problem of designing an efficient mechanism for the k -unit limited supply auction problem. Recall that for this problem, the k -Vickrey auction, which sells to the highest k bidders at the $k + 1$ st highest bid value, exactly solves the problem of designing an efficient truthful auction when collusion is disallowed. For the case where agents may collude, we show that it is possible to obtain a constant factor approximation to the efficient allocation while simultaneously being t -truthful with high probability in k , the number of units for sale.

Definition 6.1 The Approximate Efficiency Maximization (AEM) auction does the following:

1. Use random sampling to estimate the k th highest bid value:
 - (a) Let S be a sample with each agent chosen independently with probability $1/k$.
 - (b) Let p be the highest bid value in S .
2. Use the consensus technique to estimate the number of bidders above p :
 - (a) Let y be a uniform random variable from $[0, 1]$.
 - (b) Let $n_p = \lceil \#_p(\mathbf{b}) \rceil_{\{2^{i+y} : i \in \mathbb{Z}\}}$.
3. Use p and n_p in a posted price mechanism in which with high probability the item is not oversold:
 - (a) If $n_p \leq k$ output price p for all bidders.
 - (b) Otherwise, output offer price z_i independently for each bidder i with

$$z_i = \begin{cases} p & \text{with probability } k/(2n_p) \\ \infty & \text{otherwise.} \end{cases}$$

4. Touch up: If the item is oversold in Step 3b, run k -Vickrey.

Lemma 6.1 AEM is t -truthful with probability $1 - \Theta(t/k)$.

Proof: Take any coalition C of size t or fewer. Let \mathcal{E}_s be the event that a bidder from C is in the sample S . Let \mathcal{E}_c be the event, conditioned on p , that n_p is not a t -consensus (see the definition of t -consensus in Section 5) and $n_p \geq k$. Finally, let \mathcal{E}_o be the event that items are oversold in Step 3b. If none of the events \mathcal{E}_s , \mathcal{E}_c , and \mathcal{E}_o occur, then the coalition C cannot manipulate their bid values to obtain a higher total profit. We will bound the probability of each of these events by $O(t/k)$. The union bound then gives the theorem.

1. Clearly, $\Pr[\mathcal{E}_s] \leq t/k$, as any of the t bidders in C is sampled independently with probability $1/k$.
2. Let p be the price selected by the algorithm. If $n_p \leq k$, then AEM is t -truthful precisely when \mathcal{E}_s does not occur, i.e., with probability at least $1 - t/k$.

If $n_p > k$ then we must have $\#_p(\mathbf{b}) + t \geq k/2$. By Lemma 5.2 the probability that n_p is not a t -consensus is,

$$\begin{aligned} \Pr[\mathcal{E}_c] &\leq \log \frac{\#_p(\mathbf{b}) + t}{\#_p(\mathbf{b}) - t} \\ &\leq \log \frac{k/2}{k/2 - 2t} \\ &\leq -\log \left(1 - \frac{2t}{k} \right) \\ &\in \Theta(t/k). \end{aligned}$$

3. The expected number of units sold in Step 3b is $\#_p(\mathbf{b})k/(2n_p)$. Since $\#_p(\mathbf{b}) \leq n_p$ this is at most $k/2$. By the Chernoff bound, the probability that the expectation is exceeded by more than a factor of two (i.e., the item is oversold) is at most $e^{-\frac{k}{8}} \in o(\frac{1}{k}) \subset o(\frac{t}{k})$.

□

Lemma 6.2 AEM is a constant approximation the optimal efficiency (in expectation).

Proof: The probability that there are between $k/2$ and k bids at least p is $1/4$. To see this, let x be the probability that none of $k/2$ independent trials, each with probability $1/k$ of success, are successful. Then the probability that none of the top $k/2$ bidders is in the sample is exactly x . Furthermore, the probability that at least 1 of the bidders between $k/2$ and k from the top is sampled is $1 - x$. Since these events are independent, the probability that they both occur is $x(1 - x)$. This is at least the minimum possible value of $x'(1 - x')$ for any $x' \in [0, 1]$, differentiating, this minimum occurs at $x' = 1/2$ showing that $x(1 - x) \geq 1/4$.

The above occurrence implies that $n_p < 2k$, so the probability that each of the top $k/2$ bidders is offered price p is at least $1/4$. Thus, our expected efficiency is at most 32 from optimal. □

7 Group Strategyproof Mechanisms

Definition 7.1 (Group Strategyproof) A mechanism is group strategyproof if for any coalition, any non-truthful coalition strategy that results in a strict gain in utility for some agent in the coalition also results in a strict loss in utility for some other agent in the coalition.

One of the motivations for considering the strong collusion resistance notion of t -truthfulness is the weakness of the more prevalent notion of *group strategyproofness*. Consider the following compelling example. The *cost sharing* mechanism of Moulin and Shenker [23], in the setting of a digital good with a fixed production cost C , works by finding the largest number of agents, k , such that each agent's bid is at least C/k , their fair share of the cost. It sells each of these k agents a copy of the good at price C/k and rejects all other agents. Suppose we have a hundred agents with valuations $\{0.99, 1, 1, 1, \dots, 1, 1, 100\}$ and $C = 100$.

If agents bid truthfully, the mechanism will sell an item to the last agent at price \$100 and reject all other agents. In this case, all agent utilities are zero. If the first agent were to bid \$1, each agent gets the good at \$1. Except for the first and the last agent, all agent utilities are zero. The first agent has a loss of \$0.01 and the last one has a gain of \$99. Even though the cost sharing mechanism is known to be group strategyproof, the first and the last agent can benefit from collusion if side-payments are allowed.

In this section we consider group strategyproofness and the goals of profit maximization and economic efficiency. We show that the definition of group strategyproof permits mechanisms that approximate profit maximization or efficiency and are group strategyproof. Nonetheless, given the weakness of group strategyproofness we are reluctant to believe that these mechanisms offer significantly more protection from collusion than other truthful mechanisms with no guarantees about collusion resistance. In particular, our Approximate Vickrey auction (Section 7.2) which is group strategyproof in expectation does not seem to be much more collusion resistant than the classical Vickrey auction which is well known to be vulnerable to collusion both theoretically, as discussed in [26, 8], and in practice [13, 17].

We begin our discussion by making some general observations about group strategyproofness. In several aspects, the definition of group strategyproofness is more subtle than the definitions of truthfulness or t -truthfulness. The first subtlety is that randomization over deterministic group strategyproof auctions need not be group strategyproof in expectation (and vice versa). The second, which we defer to the end of this section is that with respect to mechanisms that are group strategyproof (in expectation) randomized coalition strategies offer the coalition strictly more power than if they are restricted to use deterministic strategies. Until Section 7.4 we assume that the coalitions follow deterministic strategies.

7.1 Group Strategyproofness and Randomized Mechanisms

Some basic properties of group strategyproofness are very different from those for truthfulness and t -truthfulness. In particular, a randomization over deterministic group strategyproof mechanisms is not necessarily group strategyproof (in expectation). Consider the following three auctions with two bidders.

Auction 1: Set $p_1 = 10$. Set p_2 to 10 if $b_1 \geq 10$ and to 1 otherwise.

Auction 2: (the mirror of Auction 1) Set $p_2 = 10$. Set p_1 to 10 if $b_2 \geq 10$ and to 1 otherwise.

Auction 3: Run Auction 1 or Auction 2 with probability 1/2 each.

First observe that Auction 1 (and Auction 2, by symmetry) is group strategyproof. Bidder 2 cannot affect bidder 1's price. Bidder 1 can only lower bidder 2's price by rejecting price 10. Note that if bidder 1's utility is above 10, this bidder would have performed the suboptimal action of rejecting the item at price 10. Thus, we can conclude that Auction 3 is a randomization over deterministic group strategyproof auctions.

Now we show that auction 3 is not group strategyproof in expectation. Assume that bidder 1 and 2 both have utility values of 11 for the item. Their profits for obtaining an item at price 1 and 10 are 10 and 1 respectively. If they follow the truth-telling strategy, their expected profit will be 1 because they will always obtain the item at price 10. If the two bidders collude and bid 2 each, one will lose and one will obtain the item at price 1. This winner is determined based on the coin flip, so the expected profit of each bidder is 5. Thus, in expectation, both bidders gain from this strategy.

As we show in the next section, a mechanism that is group strategyproof in expectation is not necessarily a randomization over deterministic group strategyproof mechanisms.

7.2 Group Strategyproof and Vickrey

We now give a mechanism that is an $(1 - \epsilon)$ -approximation in expectation to the k -unit Vickrey auction outcome (efficiency and profit) and also group strategyproof in expectation. However, it is not a randomization over deterministic group strategyproof mechanisms

Approximate Vickrey Auction:

1. With probability $1 - \epsilon$, run k -Vickrey.
2. Otherwise, with probability ϵ pick k bidders at random and offer each of them price p chosen from a continuous distribution with support $[0, \infty)$.

Clearly, because the k -Vickrey auction is run with probability $1 - \epsilon$, we have a $(1 - \epsilon)$ -approximation to efficiency and profit in expectation. Next we show that the auction is group strategyproof in expectation. Consider any coalition. First, if the Vickrey price does not decrease by the coalition strategy, then no bidder benefits. Second, if the Vickrey price decreases, this is because the value of the $k + 1$ st bid has decreased. This can only happen if some bidder with value equal to or above the $k + 1$ st bid value lowers their bid below that value. This bidder, however, will not be a winner in k -Vickrey and, because of Step 2, is strictly worse off by not bidding their true value.

7.3 Posted Price Mechanisms

We define a *sequential posted price mechanism* as follows. Given a predetermined ordering of the bidders, π , in round i bidder π_i is offered price z_i . The bidder may accept or reject the offered price. Thus during step i , the mechanism only learns whether the bidder have accepted or rejected the offered price. The price offered in round i can be based on the response of bidders in all previous rounds.

Lemma 7.1 *A posted price mechanism with deterministically chosen ordering π is group strategyproof.*

Proof: Since π is deterministically chosen, in any coalition there is a “first bidder”. This bidder is offered a price z that is not a function of the actions of any of the other bidders in the coalition. Further, in a posted price mechanism, the only way a bidder can affect the coalition with a non-truthful strategy is by rejecting an offered price below their utility or by accepting a price above their utility. In either case such a strategy makes this first bidder in the coalition strictly worst off and thus, a sequential posted price mechanism is group strategyproof. \square

Blum et al. in [6] apply the approach of expert based learning to solve the *online posted price auction problem*. In this problem bidders arrive one at a time and the auctioneer must offer each bidder a price. The bidder either rejects or accepts the price. Clearly, a online posted price auction is a sequential posted price mechanism. The online posted price mechanism given in [6] assumes all bids are between 1 and h and obtains profit of at least $(1 - \epsilon) \text{OPT} - O(h \log h \log \log h/\epsilon)$.

We note very briefly that using the same techniques as for the Approximate Vickrey auction presented above, we can also look at sequential auctions where the auctioneer learns the bidders’ bids after they arrive (instead of just whether or not they accepted the offered price). This allows us to make use of the online auction results of [6] which gives an improved bound of $(1 - \epsilon) \text{OPT} - O(h \log h/\epsilon)$ (this result has been recently improved by Blum and Hartline to give $(1 - \epsilon) \text{OPT} - O(h/\epsilon)$ [5]).

7.4 Group Strategyproofness and Randomized Coalition Strategies

In this section we briefly show that group strategyproofness when randomized coalition strategies are allowed is strictly more restrictive than group strategyproofness without randomized coalition strategies. In particular we show that the group strategyproof approximation of the k -Vickrey auction, discussed in Section 7.2, is not *group strategyproof for randomized strategies*. Take $k = 1$ and consider three bids with values $v_1 = 10$, $v_2 = 9$, and $v_3 = 1$. Assume ϵ is negligibly small. The following randomized coalition strategy benefits both bidder 1 and bidder 2 (assuming bidder 3 follows its optimal strategy of bidding $b_3 = 1$). Bidders 1 and 2 pick their bid independently and uniformly from $[1, 2]$. Each wins with probability $1/2$ and their profit on winning is at least 7 for bidder 2 and 8 for bidder 1. This improves the expected profit of both bidders. Thus a mechanism that is group strategyproof for deterministic strategies is not necessarily group strategyproof for randomized (a.k.a., mixed) strategies.

8 Discussion and Conclusions

We studied t -truthfulness, a natural notion of collusion resistance that allows for the possibility that colluders may exchange side-payments to redistribute their gains from colluding. We have shown that the only mechanisms that are t -truthful are posted price mechanisms, mechanism that must post “take it or leave it”

offers for each agent before it has access to any of the agents' bids. Under this restriction it is impossible to construct interesting mechanisms. We have shown that if we relax t -truthfulness and consider t -truthfulness with high probability, mechanisms for approximating both profit maximization and efficiency maximization exist. These mechanisms are based on the following general design framework.

- A *summary value*, $S(\mathbf{b})$, is a value that does not change much (e.g., a constant factor) when some agents change their bids. In both APM and AEM we used the summary values, $\#_p(\cdot)$, “the number of bids that are at least p ”.
- The *consensus estimate* of a summary value $S(\cdot)$ is a randomized function $f_S(\cdot)$ which approximates $S(\cdot)$ and with high probability satisfies $f_S(\mathbf{b}') = f_S(\mathbf{b})$ when \mathbf{b} and \mathbf{b}' only differ in a small number of bids. In many cases it is useful for the consensus estimate to be a lower bound, e.g. for APM (or an upper bound, e.g., for AEM) on the actual summary value. Consensus estimate functions for summary values have been shown to exist in [15].
- An *approximate summary mechanism* \mathcal{M}_S is a (possibly randomized) posted price mechanism that is parameterized by summary values, $\mathcal{S} = (S_1, S_2, \dots)$, and that (approximately) achieves the desired objectives given approximations to the summary values. It may require the approximate summary values be upper or lower bounds and it may fail to meet the objective with some low probability in random coin flips it makes (e.g., in AEM).

Using this approach, the design problem for t -truthful with high probability mechanisms becomes that of identifying the appropriate summary values and designing the approximate summary mechanism for the desired objective.

For auction settings where the goods being allocated could potentially be reallocated among the agents after the mechanism is run, an even stronger notion of collusion resistance is interesting, one that allows both side-payments and goods to be exchanged and still truth-telling remains optimal for agents in a coalition. We observe in passing that it is rather simple to adapt our characterization to show that this class of mechanisms is precisely the class of mechanisms where the prices posted to each agent are identical, i.e., $\bar{x}_i(\cdot) = \bar{x}_j(\cdot)$. Both APM and AEM satisfy this stronger property with high probability.

Finally we conclude by considering the observation that in standard truthful mechanism design profit maximization is more difficult than efficiency maximization. This disparity is more apparent for worst case profit maximization, where the work of [16, 12] for profit maximization can be contrasted with the k -Vickrey auction for (the worst case) efficiency. Our experience in designing the collusion resistant auctions APM and AEM indicates an interesting reversal of this phenomenon.

References

- [1] A. Archer, C. Papadimitriou, K. Tawar, and E. Tardos. An Approximate Truthful Mechanism for Combinatorial Auctions with Single Parameter Agents. In *Proc. 14th Symp. on Discrete Alg.* ACM/SIAM, 2003.
- [2] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001.
- [3] A. Archer and E. Tardos. Frugal Path Mechanisms. In *Proc. 13th Symp. on Discrete Alg.* ACM/SIAM, 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 and J. Hartline. Unpublished Manuscript, 2004.
- [6] A. Blum, V. Kumar, A. Rudra, and F. Wu. Online Learning in Online Auctions. In *Proc. 14th Symposium on Discrete Algorithms*. ACM/SIAM, 2003.

- [7] E. H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 11:17–33, 1971.
- [8] P. Cramton and J. Schwartz. Collusion in Auctions. *Annales d'Économie et de Statistique*, 15/16:217–230, 1989.
- [9] P. Cramton and J. Schwartz. Collusive Bidding: Lessons from teh FCC Spectrum Auctions. *Journal of RRegulatory Economics*, 17:229–252, May 2000.
- [10] E. Elkind, A. Sahai, and K. Steiglitz. Frugality in Path Auctions. In *Proc. 15th Symp. on Discrete Alg.* ACM/SIAM, 2004.
- [11] 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.
- [12] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive Generalized Auctions. In *Proc. 34th ACM Symposium on the Theory of Computing*. ACM Press, New York, 2002.
- [13] M. Friedman. Comment on ‘Collusion in the auction market for treasury bills’. *J. of Political Economy*, 9:757–785, 1996.
- [14] A. Goldberg and J. Hartline. Envy-Free Auction for Digital Goods. In *Proc. of 4th ACM Conference on Electronic Commerce*. ACM Press, New York, 2003.
- [15] A. V. Goldberg and J. D. Hartline. Competitiveness via Concensus. In *Proc. 14th Symp. on Discrete Alg.* ACM/SIAM, 2003.
- [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.
- [17] G. Goswami, T.H. Noe, and M.J. Rebello. Collusion in Uniform-Price Auctions: Experimental Evidence and Implications for Treasury Auctions. *Review of Financial Studies*, 72:513–514, 1964.
- [18] T. Groves. Incentives in Teams. *Econometrica*, 41:617–631, 1973.
- [19] K. Jain and V. Vazirani. Applications of approximation to cooperative games, 2001.
- [20] P. Klemperer. Auction theory: A guide to the literature. In *Journal of Economic Surveys*, pages 227–286. 13(3), 1999.
- [21] D. Lehmann, L. I. O’Callaghan, and Y. Shoham. Truth Revelation in Approximately Efficient Combinatorial Auctions. In *Proc. of 1st ACM Conf. on E-Commerce*, pages 96–102. ACM Press, New York, 1999.
- [22] R. McAfee and J. McMillan. Bidding Rings. *The American Economic Review*, 82(3):579–599, 1992.
- [23] H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs: Budget Balance Versus Efficiency. *Economic Theory*, 18:511–533, 2001.
- [24] 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.
- [25] N. Nisan and A. Ronen. Computationally feasible VCG Mechanisms. In *ACM Conference on Electronic Commerce*, pages 242–252, 2000.
- [26] M. Robinson. Collusion and the Choice of Auction. *Rand Journal of Economics*, 16:141–145, 1985.
- [27] K. Talwar. The Price of Truth: Frugality in Truthful Mechanisms. In *Proc. of 20th International Symposium on Theoretical Aspects of Computer Science*. Springer, 2003.
- [28] W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *J. of Finance*, 16:8–37, 1961.