Auctions for Share-Averse Bidders

Mahyar Salek and David Kempe

Department of Computer Science, University of Southern California, CA 90089-0781, USA, {salek, dkempe}@usc.edu

Abstract. We introduce and study *share-averse auctions*, a class of auctions with allocation externalities, in which items can be allocated to arbitrarily many bidders, but the valuation of each individual bidder decreases as the items get allocated to more other bidders. For single-item auctions where players have incomplete information about each others' valuation, we characterize the truthful mechanism that maximizes the auctioneer's revenue, and analyze it for some interesting cases.
We then move beyond single-item auctions, and analyze single-minded combinatorial auctions. We derive sufficient conditions for a truthful allocation in this setting. We also obtain a $\sqrt{m}$-approximation algorithm for maximizing social welfare, which is essentially tight unless P=NP.

1 Introduction

Consider the problem of selling a piece of technological or financial advice. In principle, such information can be sold to all participating bidders at no marginal cost to the seller. However, in reality, the value of the information to each individual bidder decreases the more other bidders receive the information, since the winner(s) of the auction will not obtain as strong a technological or financial advantage over the losers as they would have otherwise. A similar scenario can arise for physical items: for instance, the value a shared network infrastructure, a road, or a park decreases in the number of others who have access to it.

The preceding examples motivate the study of auctions for *share-averse buyers* (*share-averse auctions* for brevity): auctions in which items can in principle be allocated to arbitrarily many bidders, but the valuation of each individual bidder decreases as the items get allocated to more other bidders. Share-averse auctions fall broadly in the scenario of *auctions with allocation externalities* [10, 2, 8]. They differ in that the externalities take on a simpler form: we assume that agents care only about the number of other players sharing items with them, but not about their identity, or the allocation of items which the player does not share. Furthermore, traditional models [10, 2, 8] for externalities use additive terms, whereas in the share-averseness model, multiplicative decreases appear more appropriate.

Our main results in this paper are twofold. First, we extend the seminal work of Myerson [14] and characterize optimal auctions for share-averse bidders if only a single item is auctioned off. We then focus on the case where all bidders have the same share-averseness response function f, and derive a partial characterization of the optimal allocation rule in those cases. As a special case, we recover a result by Maskin and Riley [13] on revenue maximization for multi-unit auctions if bidders are unit-demand, i.e., they need at most one copy of an item.

Second, we consider the case of allocating bundles to share-averse single-minded bidders in a combinatorial auction. In our model, the value ascribed to the bundle by such a bidder depends on the *maximum* number of other bidders she shares any item with. For this problem, we characterize sufficient conditions for a truthful mechanism in the spirit of [12], and provide a (tight) $\sqrt{m}$ approximation mechanism.

2 Related Work

Share-averseness is a negative allocation externality among the winners. Auctions with externalities [9, 8, 10, 2] are often studied in economics both for revenue maximization and efficiency. Many of these scenarios have externalities affecting the *loser* of auctions, whereas our results are based in a reduction in utility for the *winners*. Jehiel et al. [9, 8] study both informational and allocative type-independent externalities. Brocas [2] looks at the extension where externalities depend on the types of both the winner and loser of the good. Recently, Ghosh et al. [4] looked at the computational challenges of allocation with externalities and showed inapproximability results for general case. In their result, the utility depends not only on the number of bidders sharing the item, but also on the identity of the winner set. This makes the problem significantly more complex.

Our work also relates to the problem of allocating public goods or clubs subject to congestion [3]. Public goods are defined as being shared by more than one agent. *Congestion* describes the decrease in utility to the individuals as a result of the sharing. Much of the work on clubs and public goods focuses on the issues of cost sharing and incentive compatibility (see, e.g., [7]). While there has been some work on equilibria in games between different clubs trying to maximize profits (e.g., [15]), these tend to focus on the competition between multiple clubs vying for customers rather than an optimal auction for membership in one club with given size.

Approximation algorithms and truthful mechanisms for combinatorial auctions [16] have recently received a lot of attention. Much of the focus has been on the single-minded case. With m denoting the number of items, Lehman et al. [12] were the first to show that a simple greedy algorithm gives a $\sqrt{m}$-approximation (which is best possible unless P=NP). Gonen et al. [6, 5] use linear programming to extend the results to the more general case of *Packing Integer Programs* (PIPs), where multiple copies of each item are available. Later, Briest et al. [1] improve their result to a truthful $m^{1/b}$ approximation algorithm, where b is the minimum of the multiplicities of all items. We will use both algorithms as a black box in deriving our approximation result for single-minded bidders.

3 Single-Item Auctions with a Prior

In this section, we focus on the special case of selling a single item to risk-neutral bidders. The set of all n bidders is denoted by $N := \{1, \dots, n\}$. Each bidder i has a private valuation v_i, if she is allocated the item exclusively. If she shares the item with k other bidders, her valuation decreases to $v_i \cdot f_i(k)$. We call f_i the *share-averseness function* of bidder i, and require that $f_i(0) = 1$ and f_i is monotonically non-increasing. We assume that share-averseness functions are common knowledge, as opposed to valuations, which are private.

Following standard convention (e.g., [14]), we assume that each bidder's valuation is drawn independently from some distribution $g_i : [\ell_i, r_i] \to \mathbb{R}^+$ over a finite interval $[\ell_i, r_i]$, and that all bidders share this belief. We denote the cumulative distribution function (CDF) of g_i by $G_i(v) = \int_{\ell_i}^{v} g_i(t)dt$. We let $V = [\ell_1, r_1] \times \cdots \times [\ell_n, r_n]$ denote the set of all possible combinations of bidders' values, and $V_{-i} = \times_{j \in N, j \neq i}[\ell_j, r_j]$ the set of possible values of bidders other than bidder i. The *joint distribution* on valuation vectors $\mathbf{v} = (v_1, \dots, v_n)$ is $g(\mathbf{v}) = \prod_{i \in N} g_i(v_i)$. Likewise, we define $\mathbf{v_{-i}} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$ and $g_{-i}(\mathbf{v_{-i}}) = \prod_{j \in N, j \neq i} g_j(v_j)$.

In this setting, we want to derive a truthful mechanism maximizing the auctioneer's revenue. Such a mechanism can be described by functions $a_i : V \to [0, 1]$ and $p_i : V \to \mathbb{R}^+$ for each agent i. We denote the vector of all of these functions by $\mathbf{a}$ and $\mathbf{p}$, respectively. For each vector $\mathbf{v}$ of valuations, $a_i(\mathbf{v})$ is the fraction of the item assigned to bidder i (corresponding to the *visitation rate* in club good theory [3]), and $p_i(\mathbf{v})$ the expected payment of agent i. Given that we allow fractional assignments of items, we need to define the notion of share-averseness more precisely. We set $w_i(\mathbf{v}) = \sum_{j \in N, j \neq i} a_j(\mathbf{v})$ to be the total fractional sharing of bidder i. One way to interpret $w_i(\mathbf{v})$ is as the expected number of bidders that i is sharing with if each bidder j receives the item with probability $a_j(\mathbf{v})$.

We also need to extend the share-averse function f_i to fractional values now. We define $f'_i(t)$ as the convex combination $(\lceil t \rceil - t) \cdot f_i(\lfloor t \rfloor) + (1 - (\lceil t \rceil - t)) \cdot f_i(\lceil t \rceil)$. This definition ensures monotonicity of f'_i. Henceforth, whenever the distinction is clear from the context, we will use f_i to refer to the extension of the function to fractional values.

Remark 1. It may appear natural to explicitly consider the $a_j(\mathbf{v})$ values as *probabilities*, and assign the item to each bidder *independently* with probability $a_j(\mathbf{v})$. The disadvantage of this approach is that the expected utility of bidder i now depends not only on $w_i(\mathbf{v})$, but also on the *exact fractional assignments of each other agent*, violating our framework of share-averseness.

There is a natural interpretation of the function f'_i defined above. If the item is shared over a period of time, then standard network flow techniques can be used to efficiently find an assignment over time in which *each* bidder i shares the item with $\lceil w \rceil$ other bidders for a $1 - (\lceil w \rceil - w)$ fraction of time, and with $\lfloor w \rfloor$ other bidders for the remaining $\lceil w \rceil - w$ fraction. (Here, $w = w_i(\mathbf{v})$.) From this flow argument, we can also derive an actual distribution letting us interpret the $a_i(\mathbf{v})$ as probabilities. After finding a period of time with corresponding assignments, simply define a distribution over allocations by drawing a uniformly random point in time, and then taking the allocation at that time.

The utility of player i under valuations $\mathbf{v}$ is then $a_i(\mathbf{v}) \cdot v_i \cdot f_i(w_i(\mathbf{v})) - p_i(\mathbf{v})$. Therefore, the *expected* utility of player i with valuation v_i is

$$u_i(v_i) = \int_{V_{-i}} (a_i(\mathbf{v}) \cdot v_i \cdot f_i(w_i(\mathbf{v})) - p_i(\mathbf{v})) \cdot g_{-i}(\mathbf{v_{-i}}) d\mathbf{v_{-i}},$$

where $\mathbf{v} = (v_i, \mathbf{v_{-i}})$, and $d\mathbf{v_{-i}} = dv_1 \cdots dv_{i-1} dv_{i+1} \cdots dv_n$. The expected utility of the seller from this auction is

$$\hat{u} = \int_V \sum_{i \in N} p_i(\mathbf{v}) g(\mathbf{v}) d\mathbf{v}.$$

In order to ensure that the auction mechanism is feasible and truthful, the payments and allocated fractions will have to satisfy voluntary participation and incentive compatibility (truthfulness), as captured by the following two conditions for each bidder i:

$$u_i(v_i) \geq 0 \tag{1}$$
$$u_i(v_i) \geq \int_{V_{-i}} (a_i(\hat{v}, \mathbf{v_{-i}}) \cdot v_i \cdot f_i(w_i(\hat{v}, \mathbf{v_{-i}})) - p_i(\hat{v}, \mathbf{v_{-i}})) \cdot g_{-i}(\mathbf{v_{-i}}) d\mathbf{v_{-i}} \forall \hat{v}. \tag{2}$$

An auction mechanism is specified by the functions determining the (fractional) assignments and the payments of each bidder i, i.e., by the pair $(\mathbf{a}, \mathbf{p})$. To simplify subsequent notation, we define

$$Q_i(v) = \int_{V_{-i}} a_i(v, \mathbf{v_{-i}}) \cdot f_i(w_i(v, \mathbf{v_{-i}})) \cdot g_{-i}(\mathbf{v_{-i}}) d\mathbf{v_{-i}}$$

to be the expected conditional fraction of the original valuation bidder i expects having valuation v. We are now ready to characterize feasible truthful and individually rational mechanisms $(\mathbf{a}, \mathbf{p})$.

Lemma 1. *The following conditions are necessary and sufficient for* $(\mathbf{a}, \mathbf{p})$ *to be feasible, truthful, and individually rational.*

1. *Monotonicity: For each bidder i, if $v \leq v'$, then $Q_i(v) \leq Q_i(v')$.*
2. *Individual Rationality: For each bidder i and valuation v, $u_i(v) \geq 0$.*
3. *(Extended) Incentive Compatibility (EIC): The expected utility function of each bidder i satisfies $u_i(v) = u_i(\ell_i) + \int_{\ell_i}^{v} Q_i(t) dt$.*

Proof. The proof is similar to Lemma 2 from [14]. We only sketch it here due to space constraints. The utility of agent i with true valuation v_i, but reporting a different valuation $\hat{v}$, is

$$\begin{aligned}
&\int_{V_{-i}} (a_i(\hat{v}, \mathbf{v_{-i}}) \cdot v_i \cdot f_i(w_i(\hat{v}, \mathbf{v_{-i}})) - p_i(\hat{v}, \mathbf{v_{-i}})) \cdot g_{-i}(\mathbf{v_{-i}}) d\mathbf{v_{-i}} \\
&= \int_{V_{-i}} (a_i(\hat{v}, \mathbf{v_{-i}}) \cdot (\hat{v} + (v_i - \hat{v})) \cdot f_i(w_i(\hat{v}, \mathbf{v_{-i}})) - p_i(\hat{v}, \mathbf{v_{-i}})) \cdot g_{-i}(\mathbf{v_{-i}}) d\mathbf{v_{-i}} \\
&= u_i(\hat{v}) + (v_i - \hat{v}) Q_i(\hat{v}).
\end{aligned}$$

Thus, incentive compatibility for bidder i is equivalent to requiring that

$$u_i(v_i) \geq u_i(\hat{v}) + (v_i - \hat{v})Q_i(\hat{v}), \tag{3}$$

for all $v_i, \hat{v} \in [\ell_i, r_i]$. The rest of the proof is nearly identical to [14]. ■

The next theorem captures the notion that a mechanism is truthful if and only if the allocation rule for each bidder is monotone, and the prices are defined appropriately. The proof is similar to the proof of Lemma 3 by Myerson [14], and due to space constraints, we defer it to the full version of this paper.

Theorem 1. *Given the allocation functions* $a_1, \ldots, a_n$*, let payment functions* $\hat{p}_i$ *be defined as* $\hat{p}_i(\mathbf{v}) = a_i(\mathbf{v})v_i f_i(w_i(\mathbf{v})) - \int_{\ell_i}^{v_i} a_i(t, \mathbf{v_{-i}}) f_i(w_i(t, \mathbf{v_{-i}}))dt$ *for valuations* $\mathbf{v} = (v_1, \ldots, v_n)$*. Then, a share-averse auction is truthful if and only if the allocation functions* a_i *satisfy the monotonicity condition* $Q_i(v) \leq Q_i(v')$ *when* $v \leq v'$*. Furthermore, the revenue-maximizing auction maximizes*

$$\int_V \sum_{i \in N} \left(v_i - \frac{1-G_i(v_i)}{g_i(v_i)}\right) \cdot a_i(\mathbf{v}) f(w_i(\mathbf{v})) g(\mathbf{v}) d\mathbf{v}. \tag{4}$$

Remark 2. Note that we focus here only on the revenue maximization problem. The problem of maximizing *social welfare* is much simpler in the single-item case. The following mechanism can be easily seen to maximize social welfare and be truthful.

For each k, let S_k be a set of at most k elements maximizing $\sum_{i \in S} v_i f_i(k-1)$, with an arbitrary tie breaking rule consistent over all k. Given k, one can compute S_k with simple sorting. The welfare maximizing mechanism simply picks the set S_k with the largest social welfare. The proof uses a simple exchange argument both to show that the optimum uses an integral allocation and picks S_k. To make the mechanism truthful, one can simply charge each bidder the VCG payments, and standard arguments prove incentive compatibility.

3.1 Regular Auctions

The term $v_i - \frac{1-G_i(v_i)}{g_i(v_i)}$ in Theorem 1 is traditionally called the *virtual valuation* (or *effective bid*) of agent i. The function $c_i(x) = x - \frac{1-G_i(x)}{g_i(x)}$ is called the *virtual valuation function*. If all virtual valuation functions c_i are strictly increasing, the auction is called *regular*. Regularity is a standard assumption in auction theory, and for the rest of this section, we will focus on regular auctions. We also make the further assumption that all bidders have the same share-averseness function, i.e., $f_i = f$ for all i.

In terms of the virtual valuations, we can state the auctioneer's objective as maximizing $\int_V \sum_{i \in N} c_i(v_i) \cdot a_i(\mathbf{v}) f(w_i(\mathbf{v})) g(\mathbf{v}) d\mathbf{v}$. Without loss of generality, we will assume for the remainder of this section that bidders are sorted by virtual valuations, i.e., $c_1(v_1) \geq c_2(v_2) \geq \cdots \geq c_n(v_n)$.

Theorem 2. *Without loss of generality, the optimal mechanism for general share-averse bidders has the following properties:*

1. *The allocations are monotone non-increasing, i.e.,* $a_i(\mathbf{v}) \geq a_{i+1}(\mathbf{v})$ *for all* i.
2. *If* $c_i(v_i) < 0$*, then* $a_i(\mathbf{v}) = 0$.
3. *For every index* i *with* $a_{i+1}(\mathbf{v}) > 0$*, we have* $a_i(\mathbf{v}) + a_{i+1}(\mathbf{v}) \geq 1$.

Notice that the theorem implies that there can be at most one bidder i with allocation $0 < a_i(\mathbf{v}) < \frac{1}{2}$.

Proof Sketch. An easy calculation using the monotonicity of f shows that swapping the allocations of j and $j+1$ cannot decrease the utility of the auctioneer. Therefore, the optimal allocation is monotone non-increasing using a simple exchange argument. (Details are deferred to the full version.)

If there is a bidder with negative virtual valuation who has a (fractional) allocation, it is easy to see that the auctioneer's revenue strictly increases by taking away that bidder's allocation.

If there is a j such that $a_j(\mathbf{v})+a_{j+1}(\mathbf{v}) \leq 1$, then consider the new assignment giving bidder j an allocation of $a_j(\mathbf{v}) + a_{j+1}(\mathbf{v})$, and bidder $j+1$ an allocation of 0. We obtain that

$$\begin{aligned}
&\hat{u}^{\mathrm{OPT}} - \int_V \sum_{i \neq j, j+1} c_i(v_i) a_i(\mathbf{v}) f(w_i(\mathbf{v})) \\
&= \textstyle\int_V c_j(v_j) a_j(\mathbf{v}) f(w_j(\mathbf{v})) + c_{j+1}(v_{j+1}) a_{j+1}(\mathbf{v}) f(w_{j+1}(\mathbf{v})) g(\mathbf{v}) d\mathbf{v} \\
&\leq \textstyle\int_V + c_j(v_j) \cdot (a_j(\mathbf{v}) + a_{j+1}(\mathbf{v})) \cdot f(w_j(\mathbf{v}) - a_{j+1}(\mathbf{v})) g(\mathbf{v}) d\mathbf{v} \\
&= \hat{u}^{\mathrm{OPT}'} - \textstyle\int_V \sum_{i \neq j, j+1} c_i(v_i) a_i(\mathbf{v}) f(w_i(\mathbf{v})),
\end{aligned}$$

and can therefore repeatedly perform such alterations until the third condition is satisfied. ■

3.2 Convex Share-Averseness Functions

A very natural further restriction on f is that it is convex over its entire support. Intuitively, this corresponds to bidders losing their sensitivity to more and more other bidders sharing the item: the addition of the 100^{th} bidder causes less marginal loss in utility than the addition of the second bidder. If f is convex, we can derive stronger conditions on the allocated fractions than Theorem 2.

Theorem 3. *Under the optimal mechanism for convex share-averse bidders, at most one bidder* j *will obtain a fractional allocation* $0 < a_j(\mathbf{v}) < 1$.

Proof. By Theorem 2, the allocations in OPT are sorted, and no bidder with negative virtual valuation obtains an allocation. Suppose that in OPT, there is a j such that $1 > a_j(\mathbf{v}) \geq a_{j+1}(\mathbf{v}) > 0$. By Theorem 2, we know that $a_j(\mathbf{v}) + a_{j+1}(\mathbf{v}) \geq 1$. We construct an alternate solution, where bidder j's new allocation is 1, and bidder $(j+1)$'s is $a_j(\mathbf{v}) + a_{j+1}(\mathbf{v}) - 1$.

Define $W := \sum_{i \neq j, j+1} a_i(\mathbf{v})$. We can then see that $a_j(\mathbf{v}) \geq a_{j+1}(\mathbf{v})$ implies

$$w_j(\mathbf{v}) = W + a_{j+1}(\mathbf{v}) \leq W + a_j(\mathbf{v}) = w_{j+1}(\mathbf{v}),$$

and because f is monotone non-increasing, $f(w_j(\mathbf{v})) \geq f(w_{j+1}(\mathbf{v}))$. Therefore, the convexity of f implies that $f(w_j(\mathbf{v}) - \delta) - f(w_j(\mathbf{v})) \geq f(w_{j+1}(\mathbf{v})) - f(w_{j+1}(\mathbf{v}) + \delta)$, for any $\delta \geq 0$. In other words, since agent j is currently not sharing as much as agent $j+1$, reducing her load by δ gives a larger increase than the decrease of agent $j+1$ by increasing her load by δ. Setting $\delta = 1 - a_j(\mathbf{v})$, we can now use the above reasoning to derive

$$
\begin{aligned}
&\hat{u}^{\mathrm{OPT}} - \int_V \sum_{i \neq j, j+1} c_i(v_i) a_i(\mathbf{v}) f(w_i(\mathbf{v})) \\
&= \textstyle\int_V c_j(v_j) a_j(\mathbf{v}) f(w_j(\mathbf{v})) + c_{j+1}(v_{j+1}) a_{j+1}(\mathbf{v}) f(w_{j+1}(\mathbf{v})) g(\mathbf{v}) d\mathbf{v} \\
&\leq \textstyle\int_V c_j(v_j) a_j(\mathbf{v}) f(w_j(\mathbf{v}) - \delta) + c_{j+1}(v_{j+1}) a_{j+1}(\mathbf{v}) f(w_{j+1}(\mathbf{v}) + \delta) g(\mathbf{v}) d\mathbf{v} \\
&< \textstyle\int_V c_j(v_j) f(w_j(\mathbf{v}) - \delta) + c_{j+1}(v_{j+1}) \cdot (a_{j+1}(\mathbf{v}) - \delta) \cdot f(w_{j+1}(\mathbf{v}) + \delta) g(\mathbf{v}) d\mathbf{v} \\
&= \hat{u}^{\mathrm{OPT}'} - \textstyle\int_V \sum_{i \neq j, j+1} c_i(v_i) a_i(\mathbf{v}) f(w_i(\mathbf{v})).
\end{aligned}
$$

The first inequality used the convexity observation along with the fact that $c_j(v_j) a_j(\mathbf{v}) \geq c_{j+1}(v_{j+1}) a_{j+1}(\mathbf{v})$ by the sorting. The second inequality used monotonicity of f and the sorting $c_j(v_j) > c_{j+1}(v_{j+1})$. By repeating such reallocations, we derive an allocation with at most one fractional $a_j(\mathbf{v})$. ■

Remark 3. A proof similar to Theorem 3 derives the optimal mechanism for multi-unit auctions with regular virtual valuations. It thus recovers Proposition 4 of Maskin and Riley [13]. A multi-unit auction with k items can be modeled by bidders with the share-averseness function $f(x) = 1$ for $x \leq k-1$, and $f(x) = 0$ for $x > k-1$. An easy calculation using Theorem 2 shows that the optimal mechanism assigns the item fully to the first $\min(j, k)$ bidders, and not at all to the remaining ones, where $j \leq n$ is the largest index such that $c_j(v_j) \geq 0$. The payment of a winning agent i can then be easily derived to be the *threshold bid*, the lowest bid with which agent i could have been assigned the item.

4 Single-minded Combinatorial Auctions

In this section, we extend the study of share-averse auctions to the combinatorial setting in which there is more than one item. The set of all items is $M := \{1, \ldots, m\}$. Bidders are *single-minded*. That is, for each bidder i, there exists a set S_i such that $v_i^*(S) = v_i^*(S_i)$ for all $S \supseteq S_i$, and $v_i^*(S) = 0$ otherwise. These are the bidders' valuations if they do not share any items in S.

An assignment $\mathbf{B} = (B_1, B_2, \ldots, B_n)$ of items $B_i \subseteq M$ to bidders need not have disjoint bundles. (However, we restrict our focus to mechanisms that assign items only integrally.) We now use $w_i(\mathbf{B}) = \max_{j \in B_i}(n_{i,j})$ to denote the *maximum* number of other bidders that i shares any of her items with, where $n_{i,j} = |\{i' \neq i \mid j \in B_{i'}\}|$ is the number of users sharing item j with bidder i. The valuation of bidder i is $v_i(\mathbf{B}) = v_i^*(B_i) \cdot f(w_i(\mathbf{B}))$. Notice that we assume in this section that all bidders have the same share-averseness function f.

Remark 4. Naturally, the *maximum* number of other bidders is not the only possible measure of sharing. One could instead consider a (weighted) average, for example. The maximum number appears natural in settings where the items are combined in a physical way, and limited access to any single item causes a bottleneck. An investigation of other aggregations is left for future work.

A bid b_i comprises a pair (S, v). Since both S_i and $v_i^*(S_i)$ are private information, bidders can be strategic both about the set and the valuation they declare. The vector of the bids of all bidders is denoted by $\mathbf{b}$. The bids of all bidders except bidder i are denoted by $\mathbf{b_{-i}}$.

4.1 A sufficient condition for a truthful mechanism

Lehmann et al. [12] proved that an allocation rule for single-minded combinatorial auctions gives rise to a truthful mechanism if the allocation rule is monotone and exact, in the sense that each bidder i is either allocated her desired set S_i or the empty set, and increasing one's bid can never result in moving from receiving S_i to receiving the empty set. We extend these conditions as follows:

1. *Exactness:* For each bidder i, either $B_i(\mathbf{b}) = S_i$ or $B_i(\mathbf{b}) = \emptyset$.
2. *Allocation Monotonicity:* If $S' \subseteq S$ and $v' \geq v$, and $B_i(\mathbf{b}) \neq \emptyset$ with $b_i = (S, v)$, then $w_i(\mathbf{b}) \geq w_i((S', v'), \mathbf{b_{-i}})$. The is, by requesting a smaller set and bidding higher, a bidder can only share with fewer other bidders.

Given arbitrary (but fixed) bids $\mathbf{b_{-i}}$ by all bidders except i, and a fixed set S, we define the j^{th} *critical bid* τ_j^i of bidder i to be the infimum of all v such that $B_i((S, v), \mathbf{b_{-i}}) = S$ and $w_i((S, v)) \leq j$. It then follows immediately from allocation monotonicity that $\tau_1^i \geq \tau_2^i \geq \cdots \geq \tau_{m-1}^i$, and that if bidder i bids less than τ_{m-1}^i, she does not receive any items. Based on the critical values, we define the following payment structure:

$$\pi_j^i = \begin{cases} 0 & \text{if } j = m \\ f(m-1)\tau_{m-1}^i & \text{if } j = m-1 \\ (f(j) - f(j+1))\tau_j^i + \pi_{j+1}^i & \text{if } j < m-1. \end{cases}$$

Expanding the recursive formula gives $\pi_j^i = f(j)\tau_j^i + \sum_{k=j+1}^{m-1} f(k)(\tau_k - \tau_{k-1})$. Given an allocation scheme, we will charge bidder i the amount π_j^i for the unique index j such that $v_i \in (\tau_j^i, \tau_{j-1}^i]$. (If $v_i > \tau_{m-1}^i$, then we define $j = m$.) Note that this payment does not depend on the amount of the agent's bid, but only on the interval which the bid falls into. In the sequel, we assume that the bidder i is fixed, and omit it from the notation where it is clear. The following proposition follows fairly directly from the definition of the payment scheme:

Proposition 1. *If bidder i's bid is denied, her utility is 0. If bidder i bids truthfully, her utility is non-negative.*

The main result of this section is the following theorem:

Theorem 4. *If the allocation rule satisfies Exactness and Monotonicity, then the payment scheme π^i_j yields a truthful implementation.*

Proof. Assume that bidder i desires set S with valuation v, and submits a bid $b' = (S', v')$. By Proposition 1, b' must lead to winning S', and $S' \supseteq S$. By Lemma 3 below, bidding (S, v') gives at least the same utility. In turn, by Lemma 2, the utility of bidding (S, v) is at least that of bidding (S, v'). Hence, it is a dominant strategy to declare (S, v). ∎

Lemma 2. *If bidder i desires set S with valuation v, declaring (S, v) dominates declaring (S, v') for all v'.*

Proof Sketch. Due to space constraints, the proof is deferred to the full version of the paper. The idea is to distinguish several cases. The easy cases are when either (S, v) or (S, v') are losing bids. In those cases, it is easy to show that the utility of the truthful bid dominates the other one. If both bids lead to receiving the set S, then we distinguish whether v or v' leads to more sharing. In both cases, somewhat involved calculations show that the non-truthful declaration cannot lead to higher utility. ∎

Lemma 3. *If bidder i desires set S with valuation v, declaring (S, v) dominates declaring (S', v) for all S'.*

Proof. If $S' \not\supseteq S$, then bidder i can obtain valuation at most 0. Let $\nu = w_i((S, v), \mathbf{b_{-i}})$ and $\nu' = w_i((S', v), \mathbf{b_{-i}})$. Note that $u_i((S, v), \mathbf{b_{-i}}) = f(\nu)(v - \tau_\nu) + \sum_{k=\nu+1}^{m-1} f(k)(\tau_k - \tau_{k-1})$, and $u_i((S', v), \mathbf{b_{-i}}) = f(\nu')(v - \tau_{\nu'}) + \sum_{k=\nu'+1}^{m-1} f(k)(\tau_k - \tau_{k-1})$, where $\nu \le \nu'$ by monotonicity. Define $\phi(S, v) = f(j)$ if $\tau_j < v < \tau_{j-1}$ for $j < m$ and 0 otherwise. Note that by definition $u_i((S, v), \mathbf{b_{-i}}) = \int_0^v \phi(S, v)$. Similarly, $u_i((S', v), \mathbf{b_{-i}}) = \int_0^v \phi(S', v)$.

We show that $\phi(S', x) \le \phi(S, x)$ for all $x \in [0, v]$. This immediately implies $u_i((S', v), \mathbf{b_{-i}}) \le u_i((S, v), \mathbf{b_{-i}})$. If (S, x) is a losing bid, then $u_i((S', v), \mathbf{b_{-i}}) = u_i((S, v), \mathbf{b_{-i}}) = 0$. Otherwise, (S, x) is a winning bid sharing with n_x other winners. By monotonicity, the bid (S', x) would not have been granted with $n'_x < n_x$ bidders, so $x \le \tau_{n_x - 1}$, and $\phi(S', v) \le f(n_x) = \phi(S, v)$. ∎

4.2 A $\sqrt{m}$-approximation algorithm

In this section, we present a mechanism approximating the social welfare of the assignment A, $V(A) = \sum_{i \in N} v_i(B_i(\mathbf{b}))$. We will achieve a $\sqrt{m}$ approximation. Since the problem contains the Set Packing problem as a special case (with share averseness function $f(x) = 1$ for $x = 0$ and $f(x) = 0$ for $x > 0$), it is NP-hard to approximate the problem to within a factor $m^{\frac{1}{2} - \epsilon}$ for any $\epsilon > 0$.

Consider the modified problem (which we denote by $P_{a,b}$), in which we modify the share-averseness function as follows: $f_{a,b}(x) = f(a)$ for $x \le a$, $f_{a,b}(x) = f(x)$ for $a \le x \le b$, and $f_{a,b}(x) = 0$ for $x > b$. Then, in the optimum solution, w.l.o.g., each item is shared between a and b times. This modified problem $P_{a,b}$ can be reasonably well approximated, so long as a and b are chosen "close enough".

Lemma 4. *$P_{a,b}$ can be approximated within a factor* $\min(\sqrt{m}, m^{\frac{1}{b}})\frac{f(b)}{f(a)}$.

Proof. Consider the following set packing problem. The sets are exactly the desired sets S_i, with valuations v_i. We impose a hard constraint that each item can be shared at most b times. Thus, we obtain a packing problem with a uniform upper bound of b on the number of times each item can be included.

Briest et al. [1] show that such a packing problem can be approximated to within $\min(\sqrt{m}, m^{\frac{1}{b}})$. Let OPT be the optimum solution to $P_{a,b}$. For each allocated set S_i, the optimum, by definition, obtains valuation at most $f(a)v_i$. Therefore, the optimum solution to the packing problem has valuation at least $\frac{1}{f(a)}V(\text{OPT})$. Thus, the approximate solution returned by the algorithm of [1] has valuation at least $\frac{1}{f(a)\cdot\min(\sqrt{m}, m^{\frac{1}{b}})}$ times that of OPT. For each allocated set S_i, our solution obtains valuation at least $f(b)v_i$, completing the proof. ■

Next, we show that so long as we are willing to incur a constant factor loss in the approximation guarantee, we can restrict our attention to solutions in which the share-averseness function does not take on too small values.

Lemma 5. *Let $\bar{n}$ be the largest j such that $f(j) \geq \frac{1}{n}$. Then, there is a solution OPT' whose value is within a factor 2 of that of OPT, such that no item is shared more than $\bar{n}$ times in OPT'.*

Proof. Let I be the set of all bidders sharing items with more than $\bar{n}$ other bidders. If the total valuation obtained from bidders in I is at most $\frac{1}{2}V(\text{OPT})$, then we allocate each bidder in I the empty bundle. This ensures the condition, while reducing the total value of the solution by at most a factor of 2.

On the other hand, if the total valuation of bidders in I is at least half the optimum valuation, then let $i^* \in I$ be the bidder with highest valuation for her set S_i. If we allocate i^* her bundle and everyone else the empty bundle, we again ensure that no one shares with more than $\bar{n}$ other bidders. Furthermore, the choice of i^* gives us that

$$v_{i^*} \geq \frac{1}{|I|} \cdot \sum_{i \in I} v_i \geq \frac{1}{n} \cdot \sum_{i \in I} (nf(\bar{n}+1))v_i \geq \frac{1}{2}V(\text{OPT}).$$

Here, we used that $|I| \leq n$, and $nf(\bar{n}+1) \leq 1$ by definition of $\bar{n}$. ■

The idea of our approximation algorithm is to solve several problems of the form $P_{a,b}$, and keep the best of the solutions. In order not to lose too large a factor $f(a)/f(b)$, we ensure that each interval has $f(a)$ and $f(b)$ "reasonably close" together. Formally, we define a sequence $d_1, \ldots, d_{k+1}$ by $d_1 = 0$, and $d_{i+1} = \max\{j \mid f(j-1) \geq \frac{1}{2}f(d_i)\}$. We stop when $d_i \geq \bar{n}$, and let $k+1$ be the total length. Notice that $k = O(\log n)$ by Lemma 5. The algorithm solves each of the problems $P_{d_i, d_{i+1}-1}$ using the algorithm of Briest et al. [1], and returns the best of the solutions found. This clearly takes polynomial time.

Theorem 5. *This algorithm gives a solution within a factor $\Omega(\sqrt{m})$ of OPT.*

Proof. Let OPT be the optimum solution. For each $i = 1, \ldots, k$, let O_j denote the set of bidders j who were assigned their set S_j sharing with d other bidders, for $d_i \leq d \leq d_{i+1} - 1$. Because in each solution O_j, there is potentially less sharing than in OPT, we obtain that $\sum_{i=1}^{k} V(O_i) \geq V(\mathrm{OPT})$.

Each assignment O_i is a feasible solution to $P_{d_i, d_{i+1}-1}$. Therefore, by Lemma 4, the solution A_i found by the algorithm in the i^{th} iteration satisfies

$$V(A_i) \geq \frac{f(d_i)}{\min(\sqrt{m}, m^{1/(d_{i+1}-1)}) \cdot f(d_{i+1}-1)} \cdot V(O_i) \geq \frac{1}{2(\min(\sqrt{m}, m^{1/(d_{i+1}-1)}))} V(O_i),$$

by the definition of the d_i. Now, consider 2 cases:

1. If $V(O_1) + V(O_2) \geq \frac{1}{2} \cdot V(\mathrm{OPT})$ (i.e., sharing very little can give within a constant factor of the optimum total welfare), then

$$V(A_1) + V(A_2) \geq \frac{1}{2\sqrt{m}} \cdot (V(O_1) + V(O_2)) = \Omega(\frac{1}{2\sqrt{m}}) \cdot V(\mathrm{OPT}).$$

   Therefore, at least one of A_1, A_2 gives an $\Omega(\frac{1}{2\sqrt{m}})$ approximation.
2. If, on the other hand, $V(O_1) + V(O_2) < \frac{1}{2} V(\mathrm{OPT})$, then $\sum_{i=3}^{k} V(O_i) \geq \frac{1}{2} V(\mathrm{OPT})$. Because $d_4 \geq 4$, and thus each item can be allocated at least three times in $P_{d_i, d_{i+1}-1}$ for $i \geq 3$, we know that the algorithm of Briest et al. [1] gives an $\Omega(m^{1/3})$ approximation for each such subproblem. The best of the solutions A_i for $i \geq 3$ is at least as good as the average, i.e., at least

$$\frac{1}{2(k-2)m^{1/3}} V(\mathrm{OPT}) \geq \frac{1}{2m^{1/3}\log(n)} V(\mathrm{OPT}) \geq \frac{1}{2\sqrt{m}} V(\mathrm{OPT}),$$

   so long as m and n are polynomially related. ■

5 Further Directions

In the context of single-item share-averse auctions, a promising direction for future work is to characterize the optimum mechanism more specifically when bidders have different share-averseness functions. Perhaps, stronger assumptions on the distributions could help here. It would also be interesting to draw further connections to the literature on club goods, and consider the effects of multiple competing auctioneers.

In the context of share-averse combinatorial auctions, many directions remain open. It would be desirable to obtain approximation guarantees (nearly) matching those for regular combinatorial auctions, e.g., with submodular valuations. For the single-minded case, our algorithm gives an essentially best-possible approximation guarantee. However, it does not satisfy the monotonicity condition in Theorem 4. A simple randomized variation gives monotonicity in the declared values. However, the more difficult problem is that the algorithm of Briest et al. [1] is not monotone in terms of the number of sharing agents. Modifying the algorithms of [1] to achieve monotonicity in the amount of sharing is an interesting direction for future work.

Another challenge is to obtain exact or tight approximate solutions when bidders have different share-averseness functions. Our algorithms can be generalized to this case, but current results on approximations of PIPs [11, 1] are not quite strong enough to give the tight $\sqrt{m}$ approximation. Finally, we have not yet covered the case where the share-averseness functions are not public knowledge. Designing mechanisms that are also incentive compatible with regard to revealing share-averseness, or mechanisms that learn share-averseness from past bids, is an interesting direction for future work.

Acknowledgments We thank Isabelle Brocas for useful discussions.

References

1. P. Briest, P. Krysta, and B. Vöcking. Approximation techniques for utilitarian mechanism design. In *The 37th ACM Symposium on Theory of Computing*, pages 39–48, 2005.
2. I. Brocas. Auctions with type dependent and negative externalities: the optimal mechanism, November 2007. IEPR working paper.
3. R. Cornes and T. Sandler. *The Theory of Externalities, Public Goods, and Club Goods*. Cambridge University Press, 1996.
4. A. Ghosh and M. Mahdian. Externalities in online advertising. In *17th international conference on World Wide Web*, 2008.
5. R. Gonen and D. Lehmann. Linear programming helps solving large multi-unit combinatorial auctions. In *Electronic Market Design Workshop*, 2001.
6. Rica Gonen and Daniel J. Lehmann. Optimal solutions for multi-unit combinatorial auctions: branch and bound heuristics. In *ACM Conference on Electronic Commerce*, pages 13–20, 2000.
7. M. Jackson and A. Nicoló. The strategy-proof provision of public goods under congestion and crowding preferences. *Journal of Economic Theory*, 115(2):278–308, 2004.
8. P. Jehiel and B. Moldovanu. Efficient design with interdependent valuations. *Econometrica*, 69(5):1237–59, 2001.
9. P. Jehiel, B. Moldovanu, and E. Stacchetti. Multidimensional mechanism design for auctions with externalities. *Journal of economic theory*, 85(2):258–294, 1999.
10. M. Kamien, S. Oren, and Y. Tauman. Optimal licensing of cost reducing innovation. *Journal of Mathematical Economics 21*, pages 483–508, 1992.
11. P. Krysta. greedy approximation via duality for packing, combinatorial auctions and routing. In *30th International Symposium on Mathematical Foundations of Computer Science*, 2005.
12. D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. In *ACM Conference on Electronic Commerce*, pages 96–102. ACM, 1999.
13. E. Maskin and J. Riley. Optimal multi-unit auction. *The economics of missing markets, information and games*, pages 312–335, 1989.
14. R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
15. S. Scotchmer. Profit maximizing clubs. *J. of Public Economics*, 27:25–45, 1985.
16. V. Smith, P. Crampton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006.