Mixed-bundling auctions with reserve prices*

Pingzhong Tang
Computer Science Department
Carnegie Mellon University
kenshin@cs.cmu.edu

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University
sandholm@cs.cmu.edu

ABSTRACT

Revenue maximization in multi-item settings is notoriously elusive. This paper studies a class of two-item auctions which we call a *mixed-bundling auction with reserve prices (MBARP)*. It calls VCG on an enlarged set of agents by adding the seller—who has reserve valuations for each bundle of items—and a fake agent who receives nothing nor has valuations for any item or bundle, but has a valuation for pure bundling allocations, i.e., allocations where the two items are allocated to a single agent. This is a strict subclass of several known classes of auctions, including the affine maximizer auction (AMA), λ-aution, and the virtual valuations combinatorial auction (VVCA). As we show, a striking feature of MBARP is that its revenue can be represented in a simple closed form as a function of the parameters. Thus, we can solve first-order conditions on the parameters and obtain the optimal MBARP. The optimal MBARP yields significantly higher revenue than prior auctions for which the revenue-maximizing parameters could be solved for in closed form: separate Myerson auctions, pure-bundling Myerson auction, VCG, and mixed-bundling auction without reserve prices. Its revenue even exceeds that obtained via simulation within broader classes: VVCA and AMA.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence — Multi agent system; J.4 [**Social and Behavior Science**]: Economics

General Terms

Theory, Economics

Keywords

Auction, optimal auction, combinatorial auction, revenue maximization, bundling, reserve price

1. INTRODUCTION

Perhaps one of the most important open problem in combinatorial auctions (CAs), and mechanism design at large, is to design revenue-maximizing (aka optimal) auctions. Specifically, the problem is, for the seller, to design an auction that maximizes her expected revenue, subjected to the incentive compatibility (IC) and individual rationality (IR) constraints, given the information about bidders' valuation distributions but not the actual values. There has been a significant amount of research on this topic, but even the 2-item case with additive valuations is open. In fact, the problem of designing an optimal CAs (even in a one-bidder setting) is NP-hard [3], so a general concise characterization cannot exist (unless P=NP). This is in contrast to the one-item setting, where the problem was elegantly solved, by Myerson [11]. This was later generalized to multiple identical units of one item [10].

In this paper, we consider a setting with two items, where each bidder's valuation functions are additive, that is, a bidder's valuation for the bundle of the two items is the sum of his valuations for the individual items. Thus, each bidder has a two-dimensional type: he has a valuation for the first item and a valuation for the second item. We consider the revenue optimization problem within a general symmetric class of auction: *mixed-bundling auction with reserve price (MBARP)*. It is a subclass of existing classes of auctions such as *affine maximizer auctions (AMAs)* [14], λ*-autions* [6], and *virtual valuations combinatorial auctions (VVCAs)* [8, 9, 15]. The dominant-strategy IC of MBARPs follows from the fact that each of the above three classes are also dominant-strategy IC.

While MBARPs are a narrower class than the three others, it is general enough to incorporate the ideas of reserve pricing and bundling, which are known to increase revenue in many settings. Specifically, MBARP calls VCG on an enlarged set of agents by adding the seller—who has a reserve valuation for each bundle—and a fake agent who receives nothing and has no valuation for any item or bundle, but has a valuation for pure bundling allocations (i.e., ones where both items are allocated to the same (any) bidder).

As we show, one striking feature of MBARP is that its revenue can be represented in a simple closed form as a function of its parameters. Thus, we can solve first-order conditions on the parameters and obtain the optimal MBARP. We give a system of equations of the optimal parameters in general and solve it for some simple yet common settings: for instance, in the most trivial setting, where there are two agents and each agent's valuation for each item is drawn uniformly on [0,1], our optimal auction yields expected revenue 0.871, which significantly outperforms VCG (0.667), separate Myerson auctions (0.833) (i.e, selling the two items via a sequence of separate optimal single-item auctions), pure-bundling Myerson auction (0.800) and mixed bundling auction without reserve prices (0.786). More surprisingly, it even outperforms the optimal empirical results returned by sampling and approximation on its supersets AMA (0.860) and VVCA (0.838) [15].

We wish to emphasize the generality of our approach to analyz-

*This material is based upon work supported by the National Science Foundation under grants IIS-0905390, IIS-0964579, and CCF-1101668.

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

ing the problem. We start from one bidder's (say i's) perspective, find critical values of the other bidder's (say j's) valuations such that these critical values completely partition j's type spaces while within each partition, the expected revenue of i can be represented with the same function of the auction parameters. The expected revenue from i is then the sum/integration of revenues from all parts of that partition. Another aspect of the generality of our approach is that we reduce the revenue maximization problem of an N-agent auction to that of a 2-agent auction. This can be achieved by, again from one bidder's perspective, using a fake agent to simulate the maximal values of other bidders' valuations. This extends and generalizes Riley and Samuelson's analysis of optimal 1-item auctions [13] to CAs.

2. RELATED RESEARCH

This section reviews in more detail well-known classes of auctions related to this paper.

2.1 The VCG mechanism and reserve prices

The most famous mechanism in CAs is the *Vickrey-Clarke-Groves mechanism (VCG)* [17, 2, 4], where the welfare maximizing allocation is chosen and each bidder i pays the sum of the others' valuations had i not participated minus the sum of the others' actual valuations. It is not hard to see that the VCG can yield revenue arbitrarily far from optimal:

EXAMPLE 2.1. *This can be seen even in one-item one-bidder setting where the bidder's valuation for the item is uniform on [0,1]. The second-price auction (aka Vickrey auction) would give the item to the bidder and obtain 0 revenue since there is no competition. The optimal auction (aka Myerson auction) would offer the item at price of 0.5 and obtain expected revenue of* $\frac{1}{2} \times 0.5 = 0.25$.

The VCG can be complemented with *reserve prices* to increase revenue. There are several definitions of reserve prices in CAs. We consider one that is most commonly seen: the seller pretends to have some reserve valuations for the different bundles, and the VCG with reserve prices is simply to apply the VCG on the set of all agents including the seller.

2.2 λ-auction

Another well-known technique to increase revenue in CAs is via bundling of items [12]. One notable example of bundling in CAs is the λ-auction [6]. It is the VCG, but with a fake bidder who does not receive any items, but has valuations towards allocations (instead of bundles).

In this paper, we consider one special subset of λ-auctions where the fake bidder is only interested in the allocations where the entire set of items is allocated to (any) one bidder. The resulting auction works as if it gave a discount for the bidders who are interested in the whole bundle. This auction is called a *mixed-bundling auction (MBA)*. Mixed-bundling auctions can also be complemented with reserve prices, which are defined by further including the seller into consideration. This class is called *mixed-bundling auction with reserve prices (MBARP)*. We will give a formal definition later in Section 4.

Jehiel et. al. used certain local arguments to show that some mixed bundling auction must yield higher revenue than any pure-bundling auction (where the only valid allocations are those that give the whole bundle to one agent) and any welfare-maximization auction [6]. We operate in a different direction, by directly calculating the closed-form expression of the expected revenue and obtaining the optimal revenue within MBARPs. As we shall see, our analysis and evaluations confirm the arguments of Jehiel et. al.

2.3 Myerson's auction, affine maximizer auctions (AMAs), and virtual valuations combinatorial auctions (VVCAs)

The third intuition to improve revenue in CAs is via virtual valuation [11]. The idea is to transform bidder's valuation by some function and then cast the welfare maximization on the virtual valuations; each bidder pays the lowest valuations he could have reported and still won his current bundle. In principle, the function can be anything, as long as it preserves the IC constraint. For example, for Bayes-Nash IC, Myerson's virtual transformation is defined as $\tilde{v}_i = v - \frac{1-F_i(v_i)}{F_i'(v_i)}$, where v_i is agent i's valuation and F_i is the cumulative distribution function of v_i.

However, for dominant-strategy IC, it turns out that affine transformations of the valuation are the only ones that satisfy the IC constraint for an unrestricted valuation space [14]. Affine maximizer auctions (AMAs) apply some affine transformation to get the virtual valuations, and then use the VCG on those virtual valuations (a fake bidder with valuations for the allocations is also included in the bidder set). Formally, the allocation rule in an AMA is to maximize $\sum_{i=0}^{n} \mu_i v_i(a) + \lambda(a)$, where μ's are constants and $\lambda(a)$ are the fake bidder's valuation for allocation a. Clearly, λ-auctions are AMAs where the affine transformation is the identity function $f(v) = v$, i.e., $\mu_i = 1$ for all i.

Virtual valuations combinatorial auctions (VVCAs) [8, 9, 15] are AMAs with the restriction $\lambda(a) = \sum_i \lambda_i(a_i)$, where $\lambda_i(a_i)$ only depends on what bidder i receives.

In summary, MBARPs are a subset of VVCAs and of λ-auctions, which, in turn, are subsets of AMAs.

2.4 Approximation results

An optimal auction may involve reserve pricing, bundling, and virtual valuations. How much revenue can one obtain with auctions that are simpler in form? This question motivated a recent line of research that focuses on designing simple auctions that yield revenue within a factor of optimal. Likhodedov and Sandholm [9, 15] give a logarithmic approximation of optimal multi-item auctions with a variation of the VCG, for two classes of settings: (1) additive valuations (where each bidder's valuation for a bundle is the sum of his valuations for the items in the bundle), and (2) unlimited supply (such as in digital music stores).

Recall that, in symmetric settings (settings where valuation distributions are identical across bidders), Myerson's auction coincides with a Vickrey auction [17] with the so-called monopoly reserve (i.e., a reserve valuation at which Myerson's virtual valuation function equals 0). Hartline and Roughgarden show that in the asymmetric single-parameter environment, the optimal auction, which is Myerson's auction, can be 2-approximated by a Vickrey auction with monopoly reserve prices [5]. Tang and Sandholm [16] study Levin's setting for complements [7] and prove that optimal revenue can be 2-approximated by using monopoly reserve price to curtail the allocation set, followed by welfare-maximizing allocation and Levin's payment rule. They also show that the optimal revenue can be 6-approximated even if the "reserve pricing" is required to be symmetric across bidders. Chawla et. al. prove several constant bounds of approximating optimal auctions in multi-dimensional type space (even though the optimal auction is unknown) using sequential posted prices [1].

3. THE SETTING

We consider a setting with one seller who has two indivisible distinguishable items for sale. There are a set $N = \{1, 2, \ldots, n\}$ of bidders. Each bidder $i \in N$ has a valuation v_1^i for the first item,

v_2^i for the second item, and valuation $v_1^i + v_2^i$ for the bundle of both items. The seller has zero valuation for any bundle. (However, she can pretend she has reserve valuations to obtain a higher revenue.)

An allocation is denoted by a vector $\vec{x}^i = (x_1^i, x_2^i)$ for each bidder i, where $x_j^i \in \{0, 1\}$ is the amount of item j allocated to i. The payment from bidder i to the seller is denoted by p_i. Given a $\vec{x}^i$ and p_i, bidder i's utility function is

$$u_i(\vec{x}^i, p_i) = v_1^i x_1^i + v_2^i x_2^i - p_i$$

This means that the bidders have quasi-linear, additive utility functions and no externalities.

As usual, each v_b^i is treated by all but i as a random variable distributed on $[0, 1]$ according to a cumulative distribution function F_b^i, which admits a non-negative and bounded density function f_b^i. We assume that $F_b^i = F_b^j$ for all $i, j \in N$ and all the $v's$ are independent. Our approach also applies to settings where v_1^i and v_2^i are dependent and $F_b^i \neq F_b^j$. That is, v_1^i and v_2^i are distributed according to some joint cumulative distribution $F^i(v_1^i, v_2^i)$. We make these assumptions only for the ease of presentation. We also use the standard information model where each bidder i knows his own type $v^i = (v_1^i, v_2^i)$ but others do not. The distribution over types is common knowledge.

By the revelation principle, it suffices to consider the set of direct-revelation auctions, which begin by soliciting a type from each bidder and then specify an allocation and payment for each bidder. We shall consider the direct-revelation auctions that are dominant-strategy IC and ex-post IR. A mechanism is *(weakly) dominant-strategy IC* if misreporting one's type cannot yield a higher utility for the bidder, no matter what other bidders report. A mechanism is *ex-post IR* if participation yields a non-negative utility, no matter what other bidders report.

4. MIXED-BUNDLING AUCTION WITH RESERVE PRICES

In this section, we describe the *mixed-bundling auction with reserve prices (MBARP)*. As mentioned, it is a subclass of AMA, λ-auction and VVCA. One remarkable feature of MBARP is that its expected revenue can be written in closed form as a function of its parameters, and we can optimize its revenue over the parameters.

Treat the seller as a special bidder, indexed 0, with reserve value a for item 1, b for item 2, and $a + b$ for both. Define another fake bidder, indexed $n+1$, who is never allocated any item, but has value c if some bidder $i \in N$ is allocated both items. MBARP is VCG executed on this extended set of agents $N = \{0, 1, \ldots, n + 1\}$.

DEFINITION 4.1. *Given a type profile $\vec{v} = (\vec{v}^1, \ldots, \vec{v}^n)$, MBARP is defined by its allocation rule $\vec{x} = (\vec{x}^1, \ldots, \vec{x}^n)$ and payment rule $\vec{p} = (p^1, \ldots, p^n)$.*

- *The allocation rule $\vec{x}$ is,*

$$\vec{x}(\vec{v}) = arg\ max_{\vec{x}} = \{\sum_{j=0}^{n}(v_1^j x_1^j + v_2^j x_2^j) + v^{n+1}(\vec{x})\}.$$

By definition, $v^{n+1}(\vec{x}) = c$ if $x_1^i = x_2^i = 1$ for some $1 \leq i \leq n$ and $v^{n+1}(\vec{x}) = 0$ otherwise.

- *The payment rule $\vec{p}$ is,*

$$p^i(\vec{v}) = -(\sum_{j \neq i}^{n}(v_1^j x_1^j + v_2^j x_2^j) + v^{n+1}(\vec{x})) + \sum_{j \neq i}^{n}(v_1^j \tilde{x}_1^j + v_2^j \tilde{x}_2^j) + v^{n+1}(\vec{\tilde{x}}),$$

where $\vec{\tilde{x}}(\vec{v}) = \vec{x}^{-i}(\vec{v}^{-i})$. In other words, $\vec{\tilde{x}}$ is the welfare-maximizing allocation without the participation of i.

LEMMA 4.1.

1. *MBARP is dominant-strategy IC and* ex post *IR.*
2. *A bidder with the lowest type (that is, $v_1^i = v_2^i = 0$) gets 0 utility.*
3. *Given the allocation rule of MBARP, this particular payment rule, $\vec{p}$, yields the highest revenue among all the payment rules that satisfy 1 and 2.*

Proof: The proof of 1 and 2 mirror those of the VCG. The proof of 3 follows from [6, Lemma 1]. ■

5. OPTIMAL PARAMETERS FOR THIS AUCTION CLASS

In this section, we develop a systematic approach to calculating the revenue of MBARP. This approach is general and can be applied to other strategy-proof two-item auctions.

In what follows, we first analyze the case with two bidders and then reduce those with more than two bidders to the two-bidder case.

5.1 Two-bidder setting

Suppose there are two bidder, i and j. Given the values of a, b, and c, as well as j's report (v_1^j, v_2^j), we compute the formula for i's expected payment. For this purpose, let us first look at i's allocation space. The definition of MBARP (the allocation rule and the payment rule) implies that i's allocation space must be of the shape in Figure 1. Formally:

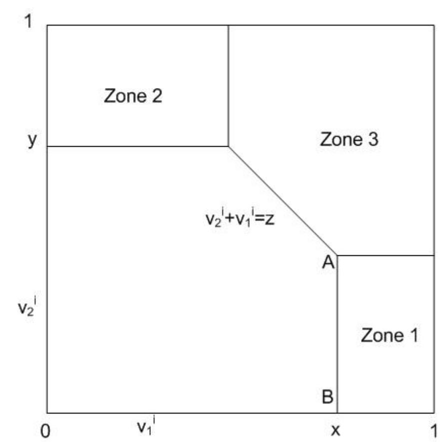

Figure 1: Bidder i's allocation space.

LEMMA 5.1.

- *The allocation space of bidder i is of exactly the shape as in Figure 1, i.e., lines are either vertical, horizontal or of -1 slope, with x, y, z being undecided variables.*
- *Bidder i gets item 1 for the price of x when his type is in Zone 1, gets item 2 for the price of y when his type is in Zone 2 and gets both items for the price of z when his type is in Zone 3.*

- *The expected payment of i is*

$$R(x,y,z) = x \cdot \int_{(v_1^i,v_2^i)\in 1} f_1^i(v_1^i) f_2^i(v_2^i) dv_1^i dv_2^i + y \cdot \int_{(v_1^i,v_2^i)\in 2} f_1^i(v_1^i) f_2^i(v_2^i) dv_1^i dv_2^i + z \cdot \int_{(v_1^i,v_2^i)\in 3} f_1^i(v_1^i) f_2^i(v_2^i) dv_1^i dv_2^i.$$

Proof: The first two claims follow from an exhaustive case analysis immediately following this lemma. Given the first two claims, the third one is straightforward. ■

5.1.1 *First part of the analysis: $a \geq c$ and $b \geq c$*

In this section we consider those parameter values that satisfy $a \geq c$ and $b \geq c$. We will consider the other parameter values in the next section.

We now determine the values of x, y, and z case by case.

Case 1. $v_1^j \geq a, v_2^j \geq b$. By the allocation rule of MBARP, bidder i gets the bundle iff

$$v_1^i + v_2^i + c \geq v_1^j + v_2^j + c \text{ and}$$
$$v_1^i + v_2^j \leq v_1^i + v_2^i + c \text{ and}$$
$$v_2^i + v_1^j \leq v_1^i + v_2^i + c.$$

The first equation ensures that bidder j does not get the bundle and the second (third) is to ensure that bidder i does not end up with item $1(2)$ only. These constraints give Zone 3 in Figure 2. By the

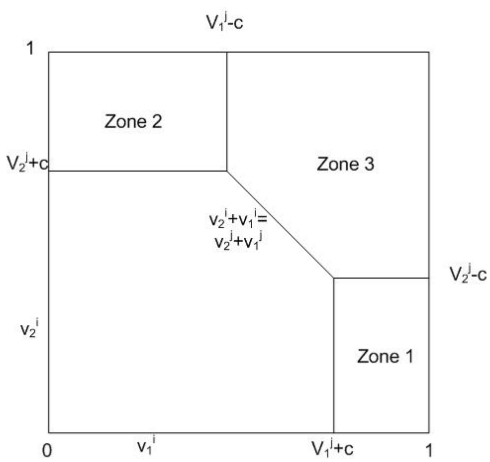

Figure 2: Case 1.

definition of MBARP, bidder i pays $z = v_1^j + v_2^j$ in this region.

Similarly, bidder i gets item 1 only iff

$$v_1^i + v_2^j \geq v_1^j + v_2^j + c \text{ and}$$
$$v_1^i + v_2^j > v_1^i + v_2^i + c.$$

The first equation ensures that bidder j does not get the bundle and the second ensures that i does not get the bundle. These constraints give Zone 1 in Figure 2. Bidder i pays $x = v_1^j + c$.

Similarly, we can derive Zone 2 and $y = v_2^j + c$.

To sum up, the allocation space of bidder i in this case is given by Figure 2. We have $x = v_1^j + c$, $y = v_2^j + c$ and $z = v_1^j + v_2^j$.

Case 2. $v_1^j < a, v_2^j < b$.

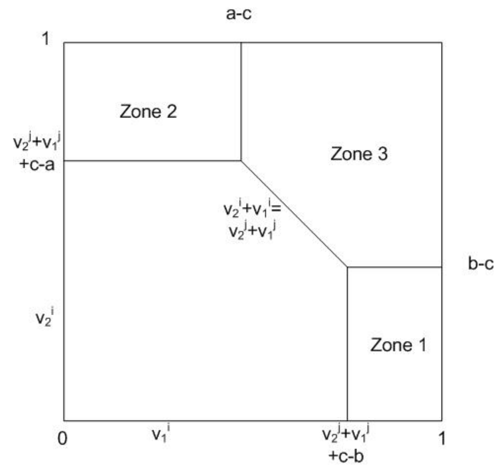

Figure 3: Case 2a.

- **Subcase a**. $v_1^j + v_2^j + c \geq a + b$. By the allocation rule of MBARP, bidder i gets the bundle iff

$$v_1^i + v_2^i + c \geq v_1^j + v_2^j + c \text{ and}$$
$$v_1^i + b \leq v_1^i + v_2^i + c \text{ and}$$
$$v_2^i + a \leq v_1^i + v_2^i + c.$$

  Thus, $z = v_1^j + v_2^j$.

  Bidder i gets item 1 only iff

$$v_1^i + b \geq v_1^j + v_2^j + c \text{ and}$$
$$v_1^i + b > v_1^i + v_2^i + c.$$

  Thus, $x = v_1^j + v_2^j + c - b$. Symmetrically, we have $y = v_1^j + v_2^j + c - a$.

  To sum up, the allocation space is given in Figure 3. We have $x = v_1^j + v_2^j + c - b$, $y = v_1^j + v_2^j + c - a$ and $z = v_1^j + v_2^j$.

- **Subcase b**. $v_1^j + v_2^j + c < a + b$. Bidder i gets the bundle iff

$$v_1^i + v_2^i + c \geq a + b \text{ and}$$
$$v_1^i + b \leq v_1^i + v_2^i + c \text{ and}$$
$$v_2^i + a \leq v_1^i + v_2^i + c.$$

  Thus, $z = a + b - c$.

  Bidder i gets item 1 only iff

$$v_1^i + b \geq a + b \text{ and}$$
$$v_1^i + b > v_1^i + v_2^i + c.$$

  Thus, $x = a$. Symmetrically, $y = b$. The allocation space is given in Figure 4.

  We have $x = a$, $y = b$ and $z = a + b - c$.

Case 3. $v_1^j < a, v_2^j \geq b$.

- **Subcase a**. $v_1^j + v_2^j + c \geq a + v_2^j$. Bidder i gets the bundle iff

$$v_1^i + v_2^i + c \geq v_1^j + v_2^j + c \text{ and}$$
$$v_1^i + v_2^j \leq v_1^i + v_2^i + c \text{ and}$$
$$v_2^i + a \leq v_1^i + v_2^i + c.$$

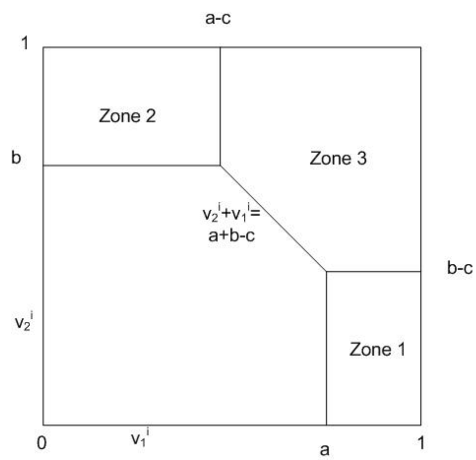

Figure 4: Case 2b.

Thus, $z = v_1^j + v_2^j$.

Similarly, $x = v_1^j + c$ and $y = v_1^j + v_2^j + c - a$.

The allocation space is given in Figure 5.

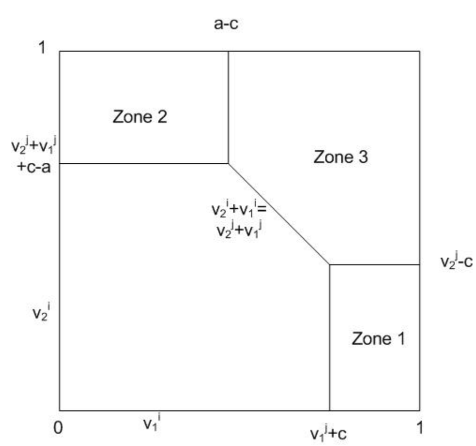

Figure 5: Case 3a.

We have $x = v_1^j + c$, $y = v_1^j + v_2^j + c - a$ and $z = v_1^j + v_2^j$.

- **Subcase b**. $v_1^j + v_2^j + c < a + v_2^j$. Bidder i gets the bundle iff

$$v_1^i + v_2^i + c \geq a + v_2^j \text{ and}$$
$$v_1^i + v_2^i \leq v_1^i + v_2^i + c \text{ and}$$
$$v_2^i + a \leq v_1^i + v_2^i + c.$$

Thus, $z = a + v_2^j - c$.

Similarly, $x = a$ and $y = v_2^j$.

The allocation space is given in Figure 6.

We have $x = a$, $y = v_2^j$ and $z = a + v_2^j - c$.

Case 4. $v_1^j \geq a, v_2^j < b$. This case is symmetric to Case 3.

- **Subcase a**. $v_1^j + v_2^j + c \geq b + v_1^j$. We have $y = v_2^j + c$, $x = v_1^j + v_2^j + c - b$ and $z = v_1^j + v_2^j$.

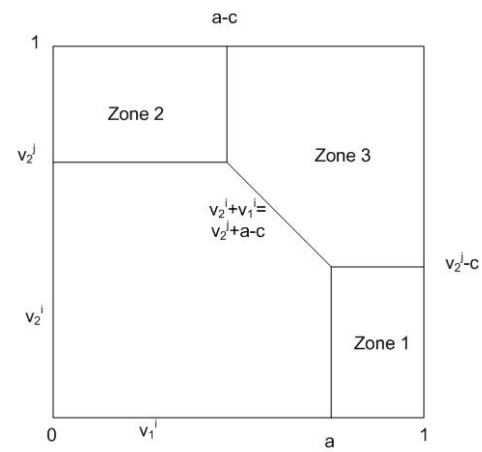

Figure 6: Case 3b.

- **Subcase b**. $v_1^j + v_2^j + c < b + v_1^j$. We have $y = b$, $x = v_1^j$ and $z = b + v_1^j - c$.

The above case-by-case analysis forms a complete partition of bidder j's type space, as shown in Figure 7, with Case 1 corresponding to Zone P_1, Case 2, Subcase a corresponding to Zone P_{2a}, and so on.

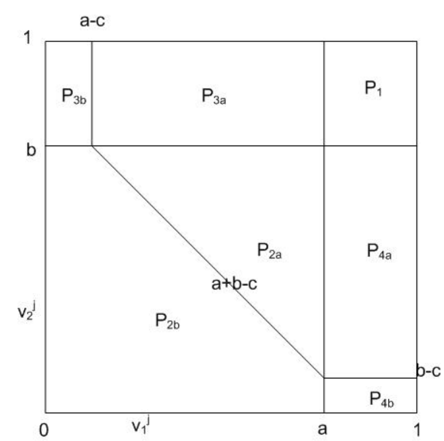

Figure 7: Partition of bidder j's type space.

We are now ready to give a closed form expression for bidder i's expected payment $r_0^i(a, b, c)$.

$$\begin{aligned}
& r_0^i(a,b,c) = \\
& \int_{(v_1^j, v_2^j) \in P_1} f_1^j(v_1^j) f_2^j(v_2^j) R(v_1^j + c, v_2^j + c, v_1^j + v_2^j) dv_1^j dv_2^j \\
+ & \int_{(v_1^j, v_2^j) \in P_{2a}} f_1^j(v_1^j) f_2^j(v_2^j) R(v_1^j + v_2^j + c - b, \\
& \qquad v_1^j + v_2^j + c - a, v_1^j + v_2^j) dv_1^j dv_2^j \\
+ & \ldots \\
+ & \int_{(v_1^j, v_2^j) \in P_{4b}} f_1^j(v_1^j) f_2^j(v_2^j) R(v_1^j, b, b + v_1^j - c) dv_1^j dv_2^j.
\end{aligned}$$

In the symmetric setting (where $f^i = f^j$), bidder j's expected payment also equals $r_0^i(a, b, c)$. This leads to the following theorem.

THEOREM 1. *For $a \geq c, b \geq c$,*

- *The expected revenue of MBARB is $2r_0^i(a,b,c)$;*
- *The optimal parameters a, b, c are given by*[1]

$$\frac{\partial r_0^i(a,b,c)}{\partial a} = \frac{\partial r_0^i(a,b,c)}{\partial b} = \frac{\partial r_0^i(a,b,c)}{\partial c} = 0.$$

In the asymmetric 2-bidder setting, $r_0^j(a,b,c)$ can be obtained by swapping the role of i and j in $r_0^i(a,b,c)$. Then the expected revenue of MBARP is $r_0^i(a,b,c) + r_0^j(a,b,c)$.

5.1.2 The remaining parts of the analysis: $(a < c,\ b \geq c)$ or $(a \geq c,\ b < c)$ or $(a < c,\ b < c)$

Now let us reexamine Figure 7. The existence of P_{3b} relies on the fact that $a - c \geq 0$. The underlying explanation is that when $a < c$, the condition of Case 3, Subcase b, which is $v_1^j + v_2^j + c < a + v_2^j$, never holds. In other words, the condition of Case 3, Subcase a holds trivially. This reasoning leads to the reduction of Figure 7 into Figure 8 when $a < c$ and $b \geq c$.

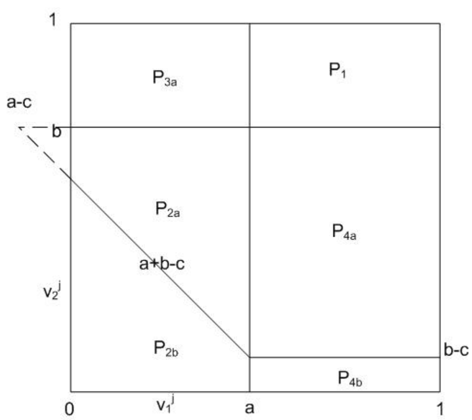

Figure 8: Partition of bidder j's type space when $a < c$ and $b \geq c$.

The corresponding x's y's and z's in each case are still the same as those when $a \geq c$ and $b \geq c$, except that Case 3b does not exist. We denote the expected revenue from i in this case by $r_1^i(a,b,c)$.

For the same reason, we can get revenue $r_2^i(a,b,c)$ when $a \geq c$ and $b < c$.

When $a < c$ and $b < c$, we need to further distinguish two subcases: $a+b-c \geq 0$ and $a+b-c < 0$, since when $a+b-c < 0$, P_{2b} also diminishes.

We denote the revenues for the two subcases above by $r_3^i(a,b,c)$ and $r_4^i(a,b,c)$.

So, the maximal revenue is given by $max_{k=0,\ldots,4}\{r_k^i(a,b,c)\}$. This completes the analysis of the two-bidder case.

5.2 More than two bidders

The allocation space grows rapidly as the number of bidders increases. The case-by-case enumeration becomes unmanageable even with 3 bidders. However, by applying order statistics, we are able to calculate bidder i's expected revenue, by treating all the other bidders simply as one bidder whose valuation for items is distributed according to the $(N-1)$th-order statistic of the valuations from these bidders. In other words, what really matters to bidder i is the greatest valuation of the remaining bidders on each bundle.

In what follows, we shall consider the symmetric setting only. That is, bidders' values are drawn from the same distributions: $F_1(v_1)$ for the first item and $F_2(v_2)$ for the second item. We denote by $F_1^k(v_1^k)$ the (k)th (with $(N-1)$th being the greatest) order statistic of $N-1$ i.i.d variables drawn from $F_1(v_1)$. Analogously, we denote by $F_2^k(v_2^k)$ the (k)th order statistic of $N-1$ i.i.d variables drawn from $F_2(v_2)$. We denote by $F_{sum}^k(v_{sum}^k)$ the (k)th order statistic of $N-1$ i.i.d variables which are sums of $N-1$ pairs of variables, one drawn from $F_1(v_1)$ and the other drawn from $F_2(v_2)$.

Fixing realized valuations v_1^{N-1} and v_1^{N-1}, and treating them as valuations of a fake bidder j for the two items respectively, we have two cases to consider:

- **Case 1.** v_1^{N-1} and v_2^{N-1} are realized by the same actual bidder. This happens with probability $p_1 = \frac{1}{N-1}$ since all the bidders are symmetric. In this case, we can completely reduce the analysis to our previous 2-bidder case, with $v_1^j = v_1^{N-1}$ and $v_2^j = v_2^{N-1}$. We denote the expected revenue from bidder i for this realization by $R_1^i(a,b,c)|(v_1^{N-1}, v_2^{N-1})$.
- **Case 2.** v_1^{N-1} and v_2^{N-1} are realized by different actual bidders. This happens with probability $p_2 = \frac{N-2}{N-1}$. There are two subcases.
  - **Subcase 1.** $v_{sum}^{N-1} + c \leq v_1^{N-1} + v_2^{N-1}$. This happens with probability $p_{21} = F_{sum}^{N-1}(v_1^{N-1} + v_2^{N-1} - c)$. We can reduce the analysis to the 2-bidder case, with a slight modification where the bundling parameter c for bidder j is 0 if he gets both items and bidder i remains unchanged. We can still go though a similar case-by-case analysis of bidder i's expected revenue since Lemma 5.1 still holds. (For example, (x,y,z) in Case 1 should now be changed to $(v_1^{N-1}, v_2^{N-1}, v_1^{N-1} + v_1^{N-1} - c)$. Similarly, one can derive x's, y's, and z's in other cases.) We denote the expected revenue from bidder i for this part by $R_{21}^i(a,b,c)|(v_1^{N-1}, v_2^{N-1})$. This is an expectation over v_{sum}^{N-1} as well as over i's own valuations.
  - **Subcase 2.** $v_{sum}^{N-1} + c > v_1^{N-1} + v_2^{N-1}$ but $v_{sum}^{N-1} < v_1^{N-1} + v_2^{N-1}$. This happens with probability $p_{22} = F_{sum}^{N-1}(v_1^{N-1} + v_2^{N-1}) - F_{sum}^{N-1}(v_1^{N-1} + v_2^{N-1} - c)$. This case means that to get the bundle, i must have a valuation no less than $v_{sum}^{N-1} + c$ (instead of $v_1^{N-1} + v_2^{N-1}$). This reduces the analysis to a 2-bidder auction where the seller's valuations are given by $(a, b, \max\{a+b, v_{sum}^{N-1} + c\})$. The analysis is also similar to that in Section 5. Denote the expected revenue from bidder i for this part by $R_{22}^i(a,b,c)|(v_1^{N-1}, v_2^{N-1})$.

Therefore, the expected revenue from bidder i is

$$\int_{(v_1^{N-1}, v_2^{N-1})} f_1^{N-1}(v_1^{N-1}) f_2^{N-1}(v_2^{N-1})$$
$$(p_1 R_1^i(a,b,c)|(v_1^{N-1}, v_2^{N-1}) +$$
$$p_2 R_{21}^i(a,b,c)|(v_1^{N-1}, v_2^{N-1}) +$$
$$p_2 R_{22}^i(a,b,c)|(v_1^{N-1}, v_2^{N-1})) dv_1^{N-1} dv_2^{N-1}.$$

[1]In general, the revenue expression is not necessarily convex and the equations have multiple roots. Therefore, it is necessary to evaluate at each root as well as boundaries to determine the optimal set of parameters. These can be easily done in Mathematica 7, as shown in our simulation in Section 6.

To evaluate the integration above (like in the two-bidder case), we partition j's valuation space and integrate each part separately, since $R_1^i(a,b,c)$, $R_{21}^i(a,b,c)$ and $R_{22}^i(a,b,c)$ take different form for different (v_1^{N-1}, v_2^{N-1}). However, each case above (Case 1, Case 2a, and Case 2b) will give us a different partition. Thus, we actually partition j's valuation space into the *coarsest common refinement* of all the partitions generated by the cases above.

The optimal auction parameters a, b, and c then follow from the same first-order conditions as given in Theorem 1.

6. COMPARISON OF AUCTIONS

This section compares the revenues of various auction classes in three different settings, all of which have two bidders and symmetric valuations:

Setting 1: $f_1(v_1) = f_2(v_2) = 1$ on [0,1].

Setting 2: $f_1(v_1) = 2v_1, f_2(v_2) = 1$ on [0,1].

Setting 3: $f_1(v_1) = 2 - 2v_1, f_2(v_2) = 1$ on [0,1].

While our analysis carries over for all (a, b, c), we restricted ourselves to $0 \leq a \leq 1$ and $0 \leq b \leq 1$ in this section. This is without loss of generality as shown in Section 7.

Auction	Optimal a, b, c	Revenue
r_0	0.577, 0.577, 0.265	0.871
r_1	0.561, 0.791, 0.561	0.848
r_2	0.791, 0.561, 0.561	0.848
r_3	0.770, 0.770, 0.770	0.843
r_4	0.257, 0.257, 1.063	0.833
VCG	0.000, 0.000, 0.000	0.667
Separate Myerson auctions	0.500, 0.500, 0.000	0.833
Pure-bundling Myerson	$r_{Bundle} = 0.816$	0.839
Mixed-bundling auction	0.000, 0.000, 0.333	0.786
AMA*	N/A	0.860
VVCA*	N/A	0.838

Table 1: Expected revenues in Setting 1.

For Setting 1, the expected revenues are listed in Table 1. r_0 to r_4 have the same meaning as in Section 5: they are the revenues of the different cases of MBARP, and the first case (r_o, which corresponds to $a \geq c, b \geq c$) yields the highest revenue in this setting. The VCG mechanism is an instance of MBARP with $a = b = c = 0$. The 8th row, separate Myerson auctions, denotes the auction where two items are sold sequentially through optimal single-item auctions (Myerson auctions). The 9th row, pure-bundling Myerson auction, denotes the optimal auction that sells both items to one agent. The 10th and 11th row, AMA* and VVCA*, denote approximated results returned by experiments [15], because there is no known analytical way to optimize the parameters of AMA and VVCA. All other revenue numbers in this section are calculated analytically (by Mathematica 7.0) and are thus (rounded) exact solutions.

We have the following conclusions:

- MBARP yields the highest revenue, even compared to experimentally "optimized" parameter settings for its supersets: AMA and VVCA.

- The reserve prices a and b play more important roles than the bundling parameter c: separate Myerson auctions yield revenue 0.833 while the mixed-bundling auction without reserve prices yields only 0.786.

- Separate Myerson auctions are inferior to the pure-bundling Myerson auction in this setting. (In fact, this is the case for all the three settings.) This is analogous to a milestone result by Palfrey [12] from the bundling literature stating that the pure-bundling Vickrey auction yields more revenue than separate Vickrey auctions in the two-bidder, N-object setting. This result also suggests an interesting general question: Does Palfrey's theorem hold for Myerson auctions as well?

For Settings 2 and 3, the expected revenues are listed in Tables 2 and 3, respectively. Again, MBARP outperforms the other auctions. Optimal expected revenues for VVCA and AMA are unknown.

Auction	a,b,c	Revenue
MBARP(r_0)	0.641, 0.581, 0.225	1.037
VCG	0.000, 0.000, 0.000	0.867
Separate Myerson auctions	0.577, 0.500, 0.000	1.001
Pure-bundling Myerson	$r_{Bundle} = 0.908$	1.002

Table 2: Expected revenues in Setting 2.

Auction	a,b,c	Revenue
MBARP(r_0)	0.415, 0.575, 0.266	0.709
VCG	0.000, 0.000, 0.000	0.533
Separate Myerson auctions	0.333, 0.500, 0.000	0.672
Pure-bundling Myerson	$r_{Bundle} = 0.694$	0.688

Table 3: Expected revenues in Setting 3.

7. DISCUSSION

In this section we discuss three additional issues.

7.1 Better starting point for automated mechanism design

Sandholm et al. show that, by starting from the VCG and then using hill-climbing algorithms to search through the parameter spaces of AMA or VVCA, the auction ends up with high-revenue (locally optimal) parameters [15, 8, 9]. Using that approach, as shown in Section 6, they obtained expected revenue 0.860 for AMA and 0.838 for VVCA in Setting 1. An easy way to improve their results is to start from the optimal MBARP instead of VCG, since optimal MBARP is also an instance of VVCA and AMA, and yields higher revenue than VCG. In fact, we implement one of the algorithms named BLAMA from [15] and improve the optimal revenue in Setting 1 to 0.872.

7.2 What if $a < 0$ or $a > 1$?

Although our theoretical analysis applies for all tuples of (a, b, c), our Mathematica program (Section 6) was based on the assumption that a and b are both within [0,1]. We now show that our assumption is without loss of generality. First, consider $a < 0$. This case is equivalent to the one where $a = 0$. (Going through the case-by-case analysis, we have only Case 1 and Case 4 left, and both of those cases yield x's, y's and z's that do not depend on a.) Intuitively, this means there is no reserve price on item 1.

Now consider $a > 1$. This case is equivalent to the one where we have revised auction parameters a' and c' as follows: $a' = 1$ and $c' = c - (a - 1)$. (Going through the case-by-case analysis, we have only Case 2 and Case 3, and both cases yield x's, y's and z's that either do not depend on a, or depend only on $c - a$, or have $x > 1$. Each of these three cases will remain the same after we switch to a' and c'.) This means that no bidder can win exactly one item.

7.3 Simple versus optimal mechanisms

A direction that might further improve revenue is to introduce asymmetry into the allocation rule, as Myerson did in the optimal one-item auction [11]. However, there are two concerns. First, we depart from Myerson's framework in the sense that we stick to the paradigm of dominant-strategy IC, while Myerson relaxed it to Bayes-Nash IC.

The second concern is related to the simplicity of the auction. This is part of the reason that Riley and Samuelson [13] restrict themselves on symmetric one-item auctions. (Fortunately, they ended up with Myerson's auction in symmetric settings. By analogy, this suggests that the optimal two-item additive-valuations auctions might lie within the MBARP family.) This is also the major motivation why revenue lower bounds from VCG-like mechanisms have been studied [5, 16].

Because we have a closed-form expression for revenue, we can explicitly trade off revenue for simplicity. For instance, for Setting 1 we may like the simplicity of the MBARP with $(a, b, c) = (0.6, 0.6, 0.3)$. We can calculate that the revenue for this configuration is 0.8696, which is very close to the optimal MBARP. That set of parameters are also noted by Jehiel et. al. [6]. Another nice set of parameters is $(a, b, c) = (\frac{1}{2}, \frac{1}{2}, \frac{1}{4})$, and it yields revenue 0.8609.

8. CONCLUSIONS AND FUTURE WORK

Revenue-maximizing multi-item auctions are perhaps the most important open topic in auction design and mechanism design at large. It is open even in the 2-item additive case. We studied a class of two-item auctions called mixed-bundling auctions with reserve prices (MBARP). The idea is that the allocation rule is biased towards the bids for the whole bundle to increase the probability of selling the bundle together. It also includes reserve pricing to further increase revenue. In fact, it is general enough to include auctions such as the VCG and separate Myerson auctions. A remarkable feature of MBARPs is that the expected revenue can be represented in a simple closed-form expression, and can be optimized easily. We gave a system of equations of the optimal parameters in general and solved it for some canonical settings. The optimal MBARP yields significantly higher revenue than the known auction classes with closed-form revenue expressions. Furthermore, its revenue even exceeds the optimal empirical results returned by sampling and approximation on much broader classes, where the truly optimal expected revenues are difficult to obtain.

There are several directions for future research. We considered the case where the bidders' valuations are additive. It is not hard to see that, with appropriate rotations of the lines in Lemma 5.1, our current approach carries over to calculating the expected revenue when bidders' valuations for the bundle are linear combination of their valuations for the items in that bundle (formally, $\lambda_1 v_1^i + \lambda_2 v_1^i + \lambda_3$, where $\lambda_1, \lambda_2, \lambda_3$ are constants). Perhaps our approach can be generalized even further.

Second, can the bundling parameter, c, be optimized separate from the reserve prices a and b?

Third, can we extend this approach to more than two items? For example, when there are 3 items, we need at least 4 bundling parameters: 3 for each pair of items and 1 for the whole bundle. Does there exist a simple derivation from the optimal 2-item parameters to the optimal 3-item parameter? More generally, can we reduce the analysis of three items to that of two items?

Future work also includes introducing asymmetry into the allocations rule. This might increase revenue even when bidders' valuations are symmetrical. Or, does the revenue-maximizing mechanism lie within the MBARP family?

9. REFERENCES

[1] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. STOC '10, pages 311–320. ACM, 2010.

[2] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 2:19–33, 1971.

[3] V. Conitzer and T. Sandholm. Self-interested automated mechanism design and implications for optimal combinatorial auctions. In *ACM EC*, pages 132–141, 2004.

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

[5] J. D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *ACM EC*, 2009.

[6] P. Jehiel, M. Meyer-Ter-Vehn, and B. Moldovanu. Mixed bundling auctions. *Journal of Economic Theory*, 127(1):494–512, 2007.

[7] J. Levin. An optimal auction for complements. *Games and Economic Behavior*, 18(2):176–192, February 1997.

[8] A. Likhodedov and T. Sandholm. Methods for boosting revenue in combinatorial auctions. In *AAAI*, pages 232–237, 2004.

[9] A. Likhodedov and T. Sandholm. Approximating revenue-maximizing combinatorial auctions. In *AAAI*, pages 267–273, 2005.

[10] E. Maskin and J. Riley. Optimal multi-unit auctions. In *The Economics of Missing Markets, Information, and Games*, chapter 14, pages 312–335. 1989.

[11] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[12] T. R. Palfrey. Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica*, 51(2):463–83, March 1983.

[13] J. G. Riley and W. F. Samuelson. Optimal auctions. *American Economic Review*, 71(3):381–92, June 1981.

[14] K. Roberts. The characterization of implementable social choice rules. In J.-J. Laffont, editor, *Aggregation and Revelation of Preferences*. North-Holland Publishing Company, 1979.

[15] T. Sandholm, A. Likhodedov, and A. Gilpin. Automated design of revenue-maximizing combinatorial auctions. *Operations Research*. Special issue on Computational Economics. To appear.

[16] P. Tang and T. Sandholm. Approximating optimal combinatorial auctions for complements using restricted welfare maximization. In *IJCAI*, pages 379–385, 2011.

[17] W. Vickrey. Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance*, pages 8–37, 1961.