

Approximating Optimal Combinatorial Auctions for Complements Using Restricted Welfare Maximization*

Pingzhong Tang^{1,2} and Tuomas Sandholm¹

kenshin@cs.cmu.edu, sandholm@cs.cmu.edu

¹ Computer Science Department, Carnegie Mellon University

² Department of Computer Science, HKUST

Abstract

The VCG mechanism is the gold standard for combinatorial auctions (CAs), and it maximizes social welfare. In contrast, the revenue-maximizing (aka optimal) CA is unknown, and designing one is NP-hard. Therefore, research on optimal CAs has progressed into special settings. Notably, Levin [1997] derived the optimal CA for complements when each agent’s private type is one-dimensional. (This does not fall inside the well-studied “single-parameter environment”.) We introduce a new research avenue for increasing revenue where we poke holes in the allocation space—*based on the bids*—and then use a welfare-maximizing allocation rule within the remaining allocation set. In this paper, the first step down this avenue, we introduce a new form of “reserve pricing” into CAs. We show that Levin’s optimal revenue can be 2-approximated by using “monopoly reserve prices” to curtail the allocation set, followed by welfare-maximizing allocation and Levin’s payment rule. A key lemma of potential independent interest is that the expected revenue from any truthful allocation-monotonic mechanism equals the expected virtual valuation; this generalizes Myerson’s lemma [1981] from the single-parameter environment. Our mechanism is close to the gold standard and thus easier to adopt than Levin’s. It also requires less information about the prior over the bidders’ types, and is always more efficient. Finally, we show that the optimal revenue can be 6-approximated even if the “reserve pricing” is required to be symmetric across bidders.

1 Introduction

Combinatorial auctions (CAs) are a key method for resource (and task) allocation in multiagent systems. One of the most important open problems in 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/mechanism that maximizes her expected revenue, given the information about bidders’ valuation distributions but not the actual values. Despite the problem’s importance and the simplicity of stating it, the problem has eluded researchers for decades. It is open even for the simplest settings such as those with two distinct items for sale and bidders having general two-dimensional private valuation (type) spaces. In fact, designing an optimal CA is NP-hard (even in the one-agent setting) [Conitzer & Sandholm, 2004]. This is in contrast to the one-item setting, where a well-known concise characterization of the optimal auction, called Myerson’s auction, exists [Myerson, 1981]. This was later generalized to multiple copies of one item [Maskin & Riley, 1989]. Over the last 30 years, a few optimal auctions in other special settings have been designed, including Levin’s auction for complements [Levin, 1997], which is also the focus of this paper, multi-item auctions for settings where the bidders’ valuations can take on either of only two values [Armstrong, 2000; Avery & Hendershott, 2000], a 1-item auction for financially constrained bidders [Pai & Vohra, 2008], sponsored search auctions [Iyengar & Kumar, 2006], and a procurement auction where both costs and capacities are private information [Iyengar & Kumar, 2008].

Interestingly, for symmetric settings where valuation distributions are identical across bidders, Myerson’s auction coincides with the second-price auction [Vickrey, 1961] with an appropriately set reserve price. Hartline and Roughgarden [2009] further show that, even if the distributions are asymmetric, a Vickrey auction with appropriately set bidder-specific reserve prices can 2-approximate Myerson’s optimal revenue. The generalization of the Vickrey auction to multiple items is the Vickrey-Clarke-Groves (VCG) mechanism [Groves, 1973; Clarke, 1971]. It is the gold standard for CAs. Likhodedov and Sandholm [2005] 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).

These kinds of results are interesting as they shed light on the tradeoff between revenue and simplicity of the auction. This is because it is generally agreed that optimal auctions like Myerson’s are more complicated and usually impractical. They also tend to require more information about the bidders’

*This material is based upon work supported by the National Science Foundation under IIS grants 0905390 and 0964579. This work is also partially supported by HK RGC GRF 616909.

valuation distributions than the VCG-style mechanisms. We will see that in this paper also.

In this paper, we explore further how far the idea of approximating optimal auctions with variations of the VCG can go, by looking at a different setting, namely Levin’s setting for selling complementary items [Levin, 1997]. Levin gives an auction that generalizes Myerson’s auction to a CA setting where bidders’ valuations exhibit complementarity, but each bidder’s valuations are parameterized by a single private-information variable (aka the bidder’s *type*). This setting captures a variety of real-world domains where the seller knows the bidders’ preference structure (for example, she may know that bidder i always values one item twice as much as another), but not the magnitude of the preferences (which may, in practice, depend on the bidder’s income, budget, and other factors unknown to the seller).

Although the setting of this paper deals with one-dimensional type spaces, it should not be confused with the widely-studied, but simpler, *single-parameter environment* (cf. [Nisan *et al.*, 2007, Chapter 9]). That environment includes a set of bidders and a specified set of *feasible sets* of bidders. Each feasible set denotes a set of bidders that can be served simultaneously. Each bidder has a single number representing his valuation of being served, and his valuation is zero if not served. Examples include 1-item auctions (where each singleton bidder is a feasible set) and auctions for multiple identical units with unit-demand bidders (where each size k set of bidders is a feasible set, where k is the number of units). Thus, a bidder i can only end up with one of two “outcomes”: (1) get served/allocated and enjoy utility v_i , or (2) not get served and obtain utility 0. In contrast, even in a 2-item Levin environment, the bidder can win item 1, item 2, or both, and can have a different valuation for each of these three bundles (and valuation zero for winning nothing).

This paper makes the following contributions. We introduce a new research avenue for increasing revenue where we poke holes in the allocation space—*based on the bids*—and then use a welfare-maximizing allocation rule within the remaining allocation set. In this paper, the first step down this avenue, we introduce a new form of “reserve pricing” into CAs. Prior methods have involved inserting an artificial bidder into the CA. Variants of that idea have included ones where the artificial bidder submits bids on items [Likhodedov & Sandholm, 2005], on bundles [Likhodedov & Sandholm, 2004; 2005; Jehiel, Meyer-Ter-Vehn, & Moldovanu, 2007], or on entire allocations [Likhodedov & Sandholm, 2005; Jehiel, Meyer-Ter-Vehn, & Moldovanu, 2007]. This idea has also been analyzed with two artificial bidders [Guruswami *et al.*, 2005]. In contrast, we introduce “reserve pricing” by precluding allocations that do not meet pre-defined criteria. For Levin’s setting we prove that *monopoly reserve prices* followed by welfare-maximizing allocation and Levin’s payment rule 2-approximate optimal revenue. Our mechanism is close to the gold standard (VCG) and thus arguably easier to adopt than Levin’s. It also requires less information about the prior over the bidders’ types, and is more (or equally) efficient (Equation 1). We proceed to prove that a 6-approximation can be achieved using *anonymous reserve prices* followed by welfare-maximizing allocation and Levin’s payment rule. As

an important lemma with potential other applications as well, we show in Levin’s setting that the expected revenue from any truthful allocation-monotonic auction equals its expected virtual valuation. This generalizes Myerson’s lemma from the single-parameter environment.

2 Setting

We study the setting introduced by Levin [1997]. There is one seller with indivisible items for sale. While Levin points out that his results apply to any finite number of items for sale, he derives his results in the 2-item setting. We do the same. Our results, too, apply to any finite number of items since our proofs do not make particular use of the number of items.¹

There are a set of $N = \{1, 2, \dots, n\}$ bidders. Each bidder $i \in N$ has valuation $v_{i1}(\theta_i)$ for the first item, $v_{i2}(\theta_i)$ for the second, and $v_{i1}(\theta_i) + v_{i2}(\theta_i) + v_{i3}(\theta_i)$ for the bundle that includes both items. Here, θ_i is bidder i ’s *type* and $v_{i3}(\theta_i)$ captures complementarity, that is, i ’s additional satisfaction from obtaining the bundle. The seller has zero valuation for any set of items, so her utility equals the revenue, i.e., the sum of payments collected from the bidders.

An allocation is denoted by a vector $\vec{x}_i = (x_{i1}, x_{i2})$ for each bidder i , where $x_{ij} = 1$ if item j is allocated to bidder i , and $x_{ij} = 0$ otherwise. Bidder i ’s utility function is

$$u_i(\theta_i, \vec{x}_i) = v_{i1}(\theta_i)x_{i1} + v_{i2}(\theta_i)x_{i2} + v_{i3}(\theta_i)x_{i1}x_{i2} - p_i,$$

where p_i is the amount bidder i has to pay. We use the standard model of Bayesian mechanism design: (1) for each bidder i , θ_i is distributed on $[\underline{\theta}_i, \bar{\theta}_i]$ according to a cumulative density function F_i , which admits a positive, bounded density function f_i , (2) each bidder i knows his type θ_i , but others, including the seller, do not, (3) the distributions F_i are common knowledge, and (4) the seller and the bidders are risk neutral: each agent tries to maximize his own expected utility.

In order to derive an optimal mechanism for this setting, Levin makes the following three additional standard assumptions, which we also adopt:

- **Assumption 1.** For all $i \in N, j \in \{1, 2, 3\}$, the valuation functions v_{ij} are non-negative and differentiable, and they are increasing in θ_i at a decreasing rate. Formally, $v'_{i1}, v'_{i2} > 0, v'_{i3} \geq 0$ and $v''_{i1}, v''_{i2}, v''_{i3} \leq 0$.
- **Assumption 2.** We also require the usual *monotone hazard rate* condition on the distributions. Formally, defining $h(\theta_i) = \frac{1-F_i(\theta_i)}{f_i(\theta_i)}$, we assume $\frac{dh(\theta_i)}{d\theta_i} \leq 0$ for each bidder i .² Commonly seen distributions, such as the normal, uniform, exponential, and logistic distributions, satisfy this condition.
- **Assumption 3.** For all i and θ_i , $v_{i3}(\theta_i) - h(\theta_i)v'_{i3}(\theta_i) \geq 0$. As we shall see in the beginning of the next section

¹Moving beyond two items, Assumptions 1-2, below, do not change, and Assumption 3 still has to hold for all complementarity parameters, and now there are more of them. For example, in a 3-item auction, there is one such parameter for every pair of items.

²As Levin points out, Assumption 2 is actually stronger than necessary. Together with the other assumptions here, it yields an easy sufficient condition for certain derived distributions to be *regular*, which is really the condition that is needed.

where we define virtual valuations, this means that we assume that the virtual valuation for complementarity is non-negative, regardless of θ_i .

3 Levin’s optimal auction for complements

To describe Levin’s optimal auction in this setting, define *virtual valuation* for $i \in N$, $j \in \{1, 2, 3\}$ as

$$\tilde{v}_{ij}(\theta_i) = v_{ij}(\theta_i) - h(\theta_i)v'_{ij}(\theta_i).$$

The seller is treated as an additional bidder 0 with virtual valuation $\tilde{v}_{01} = \tilde{v}_{02} = \tilde{v}_{03} = 0$.

By the revelation principle, Levin (and we) can focus only on mechanisms where each bidder is asked to report his type. For each type profile, Levin chooses the allocation $\vec{x}_i = (x_{i1}, x_{i2})$ that maximizes the virtual social welfare

$$\sum_{i=1}^N (\tilde{v}_{i1}(\theta_i)x_{i1} + \tilde{v}_{i2}(\theta_i)x_{i2} + \tilde{v}_{i3}(\theta_i)x_{i1}x_{i2}).$$

Levin’s payment rule is as follows.

- Bidders who receive nothing pay nothing.
- For a bidder i who wins only one item, say item j , let θ_i^0 be the lowest type he could have reported in order to win that item. He pays $v_{ij}(\theta_i^0)$.
- For a bidder i who wins both items, let θ_i^1 be the lowest type he could have reported in order to win both items and θ_i^0 be the lowest type with which he would have won one item. He pays $v_{i1}(\theta_i^0) + v_{i2}(\theta_i^1) + v_{i3}(\theta_i^1)$ if θ_i^0 corresponds to item 1 and $v_{i1}(\theta_i^1) + v_{i2}(\theta_i^0) + v_{i3}(\theta_i^1)$ if θ_i^0 corresponds to item 2.

A mechanism is (*weakly*) *dominant strategy truthful* if misreporting one’s type cannot yield a higher utility for the bidder, no matter what other bidders report. A mechanism is *ex-post individually rational* if participation yields a non-negative utility, no matter what other bidders report.

Theorem 1 ([Levin, 1997]) *For the Levin setting, the auction above is optimal, among all mechanisms that are ex-post individually rational and weakly dominant strategy truthful.*³

4 A new form of “reserve pricing”: Poking holes in the allocation space based on revealed types

Welfare-maximizing mechanisms choose an allocation that maximizes the sum of the agents’ valuations (as opposed to Levin’s virtual welfare-maximizing allocation). The most famous such mechanism is the *Vickrey-Clarke-Groves mechanism (VCG)* [Vickrey, 1961; Clarke, 1971; Groves, 1973], where each bidder i has to pay the sum of the others’ valuations had i not participated minus the sum of the others’ actual valuations. Any welfare-maximizing mechanism can yield expected revenue arbitrarily far from optimal:

³Myerson’s auction for a single item is simply a Levin’s auction where $v_{i2} = 0$ and $v_{i3} = 0$ for all i , as if item 2 did not exist.

Example 4.1 *This is easy to see even in a 1-item 2-bidder setting. Say that bidder 1’s valuations is $\delta, 0 < \delta < 1$, and bidder 2’s valuation is 1 with probability $1 - \epsilon$ and with probability ϵ , it is uniformly drawn between 0 and 1. In this setting, no mechanism that always maximizes social welfare can yield revenue more than δ . On the other hand, a mechanism that uses a reserve price just below 1 will generate expected revenue $(1 - \epsilon) \cdot 1$. Letting δ and ϵ approach zero, we see that the welfare-maximizing mechanism can yield an arbitrarily small fraction of the available expected revenue.*

As discussed in the introduction, prior “reserve pricing” approaches to increasing revenue in the VCG and its generalizations have been based on inserting one or more artificial bidders into the auction. In contrast, we introduce a different “reserve pricing”, where allocations are precluded if they do not meet some criteria.

When applied to a welfare-maximizing mechanism, we call this template *social welfare with holes* because it “pokes holes” in the allocation space by precluding some allocations, and then maximizes welfare over the remaining allocation space. Poking holes is in contrast to the prior more restricted approach of using artificial bidder(s), which removes allocations from consideration in a more uniform, rigid manner.

In the rest of this paper, we restrict attention to mechanisms where the criteria are such that they can be applied to each bundle-agent pair (including 1-item bundles) $\langle b, i \rangle$ separately, and if the criterion for that $\langle b, i \rangle$ is not met, then all allocations that give exactly bundle b to bidder i are precluded. One such criterion is whether the virtual valuation for $\langle b, i \rangle$ exceeds some constant. One can also include further preclusion criteria. For example, in the next section we preclude $\langle b, i \rangle$ pairs where a subset of b would give higher virtual value for i than b does.

We have to be careful, though, because some instantiations of the new reserve pricing make the mechanism non-truthful:

Example 4.2 *This can be shown already in the 2-item setting with just two bidders, i and j , even if both bidders have additive valuations (i.e., $v_{i3} = v_{j3} = 0$ for all types). Let bidder j have only one type and his valuations for the two items be 1 and 4, respectively. We consider three of bidder i ’s types: $(v_{i1}(\theta_i^1) = 1, v_{i2}(\theta_i^1) = 4)$, $(v_{i1}(\theta_i^2) = 5, v_{i2}(\theta_i^2) = 5)$, and $(v_{i1}(\theta_i^3) = 6, v_{i2}(\theta_i^3) = 6)$. Let θ_i^1 be i ’s lowest type. (One can construct continuous valuation functions v_{i1} and v_{i2} (i.e., continuous functions of θ_i) that include these three points and satisfy all the conditions of the Levin model.) Say the mechanism precludes allocating the bundle of both items to bidder i in θ_i^2 and precludes allocating the bundle or item 1 to bidder i in θ_i^3 . We now calculate bidder i ’s Levin payment. At θ_i^2 , he gets item 1 and pays 1 because θ_i^1 is the lowest type with which he gets item 1. At θ_i^3 , he gets item 2 because of the preclusions and pays 6 because θ_i^3 is the lowest type with which he gets item 2. Thus, his payoff is $6 - 6 = 0$. So, if his true type is θ_i^3 , he has an incentive to lie and report θ_i^2 instead because he would obtain utility $6 - 1 = 5$ instead of 0.*

Nisan and Ronen [2007] show that if one removes *ex ante* some allocations, and then maximizes welfare over the remaining space, and prices using the VCG, the resulting mechanism is truthful. Our approach is much more general because

we allow the removal of allocations to depend on the bidders' reported types. As the above example shows, this increased generality comes at the risk of losing truthfulness.

Of course, used to the extreme, our generality allows any allocation rule, A , to be exactly simulated by one in our new class by having, for each type vector, all but the allocation prescribed by A precluded. However, we will not be using the generality to any extreme extent.

Another difference between our work and that of Nisan and Ronen is that we use Levin pricing while they use VCG pricing. Those two pricing rules are not the same even in the 1-item setting because Levin makes sure that if a bidder wins the item, he pays at least his valuation corresponding to his lowest winning type. The VCG does not have this "reserve price". The difference between the two pricing rules seems minor, however, because they coincide under a mild assumption on the valuation distributions:

Proposition 4.1 *In the Levin setting, under a welfare-maximizing allocation rule, VCG prices and Levin prices coincide if $v_{ij}(\theta_i) = 0$ for all i, j .*

Proof:

- When bidder i receives nothing, the Levin price and the VCG price both equal 0.
- Now consider the situation where bidder i receives exactly one item. Without loss of generality, say this is item 1. Because valuations are nonnegative in the Levin setting, some bidder receives item 2; say, without loss of generality, this is bidder j . There are two cases:
 - **Case 1.** Without i , the bundle would go to j (throughout the proof, j can be a bidder or the auctioneer). One can show that i 's VCG payment is $v_{j1}(\theta_j) + v_{j3}(\theta_j)$. Now define θ_i^0 such that $v_{i1}(\theta_i^0) = v_{j1}(\theta_j) + v_{j3}(\theta_j)$. For bidder i , when his type $\theta_i > \theta_i^0$, we have $v_{i1}(\theta_i) + v_{j2}(\theta_j) \geq v_{i1}(\theta_i^0) + v_{j2}(\theta_j) = v_{j1}(\theta_j) + v_{j2}(\theta_j) + v_{j3}(\theta_j)$. Thus, i gets item 1. Similarly, when his type $\theta_i \leq \theta_i^0$, i loses item 1. Therefore, θ_i^0 is exactly the lowest type with which i gets item 1. Thus the Levin price equals the VCG price.
 - **Case 2.** Without i , item 1 would go to bidder $k \neq j$ and item 2 would still go to bidder j . Define θ_i^0 such that $v_{i1}(\theta_i^0) = v_{k1}(\theta_k)$. Following a similar argument as Case 1, we still have that θ_i^0 is the lowest type with which i gets item 1.
- Now consider the situation where bidder i receives both items. Let θ_i^1 be the lowest type with which he wins both items, and θ_i^0 be the lowest type with which he wins one item, say item 1. There are two cases:
 - **Case 1.** Without i , item 1 would go to bidder j and item 2 would go to bidder k . The VCG payment is $v_{j1}(\theta_j) + v_{k2}(\theta_k)$. According to the welfare-maximizing allocation rule, we have $v_{i1}(\theta_i^0) = v_{j1}(\theta_j)$ and $v_{i1}(\theta_i^1) + v_{i2}(\theta_i^1) + v_{i3}(\theta_i^1) = v_{i1}(\theta_i^1) + v_{k2}(\theta_k)$. Thus $v_{k2}(\theta_k) = v_{i2}(\theta_i^1) + v_{i3}(\theta_i^1)$. Thus, $v_{i1}(\theta_i^0) + v_{j1}(\theta_j) = v_{i1}(\theta_i^0) + v_{i2}(\theta_i^1) + v_{i3}(\theta_i^1)$. In other words, the Levin price equals the VCG price.

– **Case 2.** Without i , the bundle of both items would go to some bidder j . The VCG payment is $v_{j1}(\theta_j) + v_{j2}(\theta_j) + v_{j3}(\theta_j)$. There are two sub-cases:

- * **Subcase 1.** The welfare of the second-best allocation is $v_{j1}(\theta_j) + v_{j2}(\theta_j) + v_{j3}(\theta_j)$. Then we have $v_{i1}(\theta_i^1) + v_{i2}(\theta_i^1) + v_{i3}(\theta_i^1) = v_{j1}(\theta_j) + v_{j2}(\theta_j) + v_{j3}(\theta_j)$. We now prove that $\theta_i^0 = \theta_i^1$ in this case. Suppose, for contradiction, $\theta_i^0 < \theta_i^1$ (of course, we cannot have $\theta_i^0 > \theta_i^1$). Then there exist a k such that $v_{i1}(\theta_i^0) + v_{k2}(\theta_k) = v_{j1}(\theta_j) + v_{j2}(\theta_j) + v_{j3}(\theta_j)$. Since v_{i1} is increasing in θ_i , it must be the case that $v_{i1}(\theta_i^1) + v_{k2}(\theta_k) > v_{j1}(\theta_j) + v_{j2}(\theta_j) + v_{j3}(\theta_j)$. Then the second-best allocation would not be $v_{j1}(\theta_j) + v_{j2}(\theta_j) + v_{j3}(\theta_j)$. Contradiction.
- * **Subcase 2.** The welfare of the second-best allocation is $v_{i1}(\theta_i) + v_{k2}(\theta_k)$. We can use $v_{i1}(\theta_i^1) + v_{i2}(\theta_i^1) + v_{i3}(\theta_i^1) = v_{i1}(\theta_i^1) + v_{k2}(\theta_k)$ to derive $v_{k2}(\theta_k)$. Substituting that $v_{k2}(\theta_k)$ into the equation $v_{i1}(\theta_i^0) + v_{k2}(\theta_k) = v_{j1}(\theta_j) + v_{j2}(\theta_j) + v_{j3}(\theta_j)$ gives us an equality where the left hand side is the Levin price and the right hand side is the VCG price. ■

While we showed in Example 4.2 that some instantiations of our new hole poking method cause loss of truthfulness, we now show that a certain natural set of criteria for precluding allocations yields a mechanism that is truthful. Here the preclusions depend on the types, so this is quite different from the Nisan-Ronen approach. We also show that this instantiation of our new "reserve pricing" scheme yields a 2-approximation of optimal revenue.

5 Monopoly reserves in welfare maximization

In this section we use a welfare-maximizing allocation rule, but we preprocess the allocation space to preclude $\langle b, i \rangle$ pairs that have negative virtual value, and we also preclude $\langle b, i \rangle$ pairs where a subset of b would give higher virtual value for i than b does. The motivation is that this will increase revenue over the vanilla welfare maximization because we know from Theorem 1 that revenue maximization is accomplished by maximizing the sum of virtual valuations.

We observe that another way to see this preprocessing is to realize that it is the optimal mechanism for the setting where each bidder, in turn, would be the only bidder. Optimal pricing schemes for one-bidder settings have been called "monopoly reserves" [Hartline & Roughgarden, 2009].

In this section, we describe the *social welfare maximization with monopoly reserve prices* (SW_{MR}):

- **Step 1. "Poke holes".** For each bidder i , the seller defines 5 *monopoly reserve types* $\tilde{v}_{i1}^{-1}(0)$, $\tilde{v}_{i2}^{-1}(0)$ and $(\sum_{j=1}^3 \tilde{v}_{ij})^{-1}(0)$, $(\tilde{v}_{i1} + \tilde{v}_{i3})^{-1}(0)$, $(\tilde{v}_{i2} + \tilde{v}_{i3})^{-1}(0)$. She asks bidders to report their types. For each bidder i , if $v_{i1}(\theta_i) < \tilde{v}_{i1}^{-1}(0)$, or equivalently $\tilde{v}_{i1}(\theta_i) < 0$, we delete all allocations in which i is awarded item 1 only.

Similarly, whenever $\tilde{v}_{i2}(\theta_i) < 0$, we delete all allocation in which i is awarded item 2 only.

Finally, we delete all allocation in which i is awarded both items if $\tilde{v}_{i1}(\theta_i) + \tilde{v}_{i2}(\theta_i) + \tilde{v}_{i3}(\theta_i) < 0$, or $\tilde{v}_{i1}(\theta_i) + \tilde{v}_{i3}(\theta_i) < 0$, or $\tilde{v}_{i2}(\theta_i) + \tilde{v}_{i3}(\theta_i) < 0$.

- **Step 2.** Choose the welfare-maximizing allocation among the remaining allocations (ties can be broken arbitrarily). Use the Levin payment rule.

In this mechanism, for someone (e.g., a bidder) to verify that a correct allocation and payments were reached, all he needs to know about the prior are the five “reserve” numbers from Step 1. (Of course, he also needs to know the bids.) For example, the auctioneer can send those five numbers and the bids to each bidder. In contrast, in Levin’s auction, one needs to know everyone’s exact distribution functions in order to calculate the virtual social welfare, which determines the allocation and payments.

5.1 Truthfulness and individual rationality

Theorem 2 SW_{MR} is (weakly) dominant strategy truthful and ex-post individually rational.

Proof: Ex-post individual rationality follows from the facts that (1) a bidder who wins nothing pays nothing, and (2) a bidder who wins something has a type that is no lower than the types used to determine what he has to pay.

We prove truthfulness via the following lemmas. A mechanism is *allocation monotonic* if fixing other bidders’ reports, if a bidder i reporting θ_i^0 receives some item, he receives that item (and possibly others) for any report greater than θ_i^0 .

Lemma 5.1 SW_{MR} is allocation monotonic.

Proof: Case 1. i receives exactly one item, say item 1, by reporting θ_i^0 . Let j be the agent that wins item 2. According to the allocation rule of SW_{MR} , we know i values item 1 the most among all the bidders who meet their reserve conditions for item 1 (i.e., those whose virtual valuations for item 1 are nonnegative): $v_{i1}(\theta_i^0) \geq v_{k1}(\theta_k)$. We also have

$$\begin{aligned} v_{i1}(\theta_i^0) + v_{j2}(\theta_j) &= \text{maximum social welfare} \\ &\geq \text{maximum social welfare given that } i \text{ receives nothing.} \end{aligned}$$

Now we change θ_i^0 to any $\theta_i > \theta_i^0$. Since $v'_{i1} \geq 0$, the above equations continue to hold except for one case, where the maximum-welfare allocation changes to the one where i receives the bundle. This leads to the conclusion that i cannot end up with nothing by reporting a higher type.

We also need to show that i cannot end up with item 2 instead of item 1. Suppose he did. Then there is some $\theta_i^1 > \theta_i^0$ such that θ_i^1 meets i ’s reserve type for item 2 and $k \neq i$ such that $v_{k1}(\theta_k) + v_{i2}(\theta_i^1) \geq v_{i1}(\theta_i^1) + v_{i2}(\theta_i^1) + v_{i3}(\theta_i^1)$.⁴ The inequality above is equivalent to $v_{k1}(\theta_k) \geq v_{i1}(\theta_i^1) + v_{i3}(\theta_i^1)$. Since we know $v_{i3}(\theta_i^1) \geq 0$ by complementarity,

⁴ θ_i^1 exceeds i ’s reserve type for item 1 because $\theta_i^1 > \theta_i^0$ and θ_i^0 has already met or exceeded the reserve type. It must be the case that θ_i^1 meets or exceeds i ’s reserve types for the bundle because $\tilde{v}_{i3}(\theta) \geq 0$ by Assumption 3 in Section 2. The reasoning above ensures that $v_{i1}(\theta_i^1) + v_{i2}(\theta_i^1) + v_{i3}(\theta_i^1)$ is a valid social welfare.

we have $v_{k1}(\theta_k) \geq v_{i1}(\theta_i^1) > v_{i1}(\theta_i^0)$. This contradicts the assumption that i values item 1 the most.

Case 2. i receives both items by reporting θ_i^0 . According to the allocation rule, $v_{i1}(\theta_i^0) + v_{i2}(\theta_i^0) + v_{i3}(\theta_i^0)$ is the maximum of social welfare. Since all three terms in the equation above are increasing in θ_i , $v_{i1}(\theta_i) + v_{i2}(\theta_i) + v_{i3}(\theta_i)$ will maximize welfare for all $\theta_i \geq \theta_i^0$. Thus, the two items will still be allocated to bidder i . ■

Given Lemma 5.1, truthfulness of SW_{MR} follows from the following three lemmas.

Lemma 5.2 A bidder who wins nothing when telling the truth cannot benefit from lying.

Proof: By reporting truthfully, he gets zero utility. If he under-reports, he still wins nothing, pays nothing, and gets zero utility. Now consider over-reporting. If he still wins nothing, then he still gets zero utility. If he wins one item, say item 1, he pays $v_{i1}(\theta_i^0)$, where θ_i^0 is greater than his true type θ_i , so he gets negative utility $v_{i1}(\theta_i) - v_{i1}(\theta_i^0)$. The case where he wins two items is similar. ■

Lemma 5.3 A bidder who wins exactly one item when telling the truth cannot benefit from lying.

Proof: Say, without loss of generality, that the item he wins is item 1. By reporting truthfully he gets utility $v_{i1}(\theta_i) - v_{i1}(\theta_i^0) \geq 0$, with $\theta_i \geq \theta_i^0$. By under-reporting, he wins nothing with zero utility or still gets item 1 with unchanged utility. By over-reporting, he still wins item 1 with unchanged utility or wins both with reduced utility $v_{i1}(\theta_i) + v_{i2}(\theta_i) + v_{i3}(\theta_i) - v_{i1}(\theta_i^0) - v_{i2}(\theta_i^0) - v_{i3}(\theta_i^0)$, with $\theta_i^0 \geq \theta_i \geq \theta_i^0$. ■

Lemma 5.4 A bidder who wins two items when telling the truth cannot benefit from lying.

Proof: Say, without loss of generality, that if he wins exactly one item, he wins item 1. His utility when telling the truth and winning both items is $v_{i1}(\theta_i) + v_{i2}(\theta_i) + v_{i3}(\theta_i) - v_{i1}(\theta_i^0) - v_{i2}(\theta_i^0) - v_{i3}(\theta_i^0)$, with $\theta_i \geq \theta_i^0 \geq \theta_i^0$. By under-reporting, if he still gets two items, the utility is unchanged. If he gets item 1 only, he get reduced utility $v_{i1}(\theta_i) - v_{i1}(\theta_i^0)$. If he wins nothing, he gets zero utility, which is even worse. By over-reporting, he still wins both items with unchanged utility. ■

This completes of the proof of truthfulness of SW_{MR} . ■

5.2 SW_{MR} 2-approximates optimal revenue

Theorem 3 SW_{MR} 2-approximates the optimal revenue.

We prove this starting with a sequence of lemmas.

The famous Myerson’s lemma [1981] shows that in any truthful single-item auction, the expected payment from any bidder i equals i ’s expected virtual valuation. We now give a similar result in our more general setting.

Lemma 5.5 For any auction that satisfies (a) truthfulness, (b) allocation monotonicity and (c) bidders with lowest type always get zero utility, we have that

$$\mathbf{E}_\theta \left(\sum_{i=1}^n p_i \right) = \mathbf{E}_\theta \sum_{i=1}^n (\tilde{v}_{i1}(\theta_i)x_{i1} + \tilde{v}_{i2}(\theta_i)x_{i2} + \tilde{v}_{i3}(\theta_i)x_{i1}x_{i2})$$

Proof: Fixing other bidders' reports, allocation monotonicity tells us there are two cutoff types $\theta_i^0 \leq \theta_i^1$ for bidder i . When his type is below θ_i^0 , he receives nothing; when his type is between θ_i^0 and θ_i^1 , he receives one item; and when his type is above θ_i^1 , he receives both items.

Truthfulness tells us that the cutoffs θ_i^0 and θ_i^1 cannot depend on i 's own report. Truthfulness also tells us that i 's payment for each possible allocation to him is fixed (does not depend on his report) and his utility increases continuously as his allocation changes as his type increases.

Condition (c) then guarantees that bidder i pays exactly $v_{i1}(\theta_i^0)$ when he receives one item and $v_{i2}(\theta_i^1) + v_{i3}(\theta_i^1)$ more for receiving the second item. So, conditions (a), (b), and (c) uniquely define Levin's payment rule.

Suppose bidder i receives item 1 first. His expected payment (again, conditional on fixing the others' bids) is

$$\begin{aligned} \mathbf{E}_{\theta_i}(p_i) &= v_{i1}(\theta_i^0)(1 - F_i(\theta_i^0)) \\ &\quad + (v_{i2}(\theta_i^1) + v_{i3}(\theta_i^1))(1 - F_i(\theta_i^1)). \end{aligned}$$

His expected virtual valuation is

$$\begin{aligned} &\mathbf{E}_{\theta_i}(\tilde{v}_{i1}(\theta_i)x_{i1} + \tilde{v}_{i2}(\theta_i)x_{i2} + \tilde{v}_{i3}(\theta_i)x_{i1}x_{i2}) \\ &= \int_{\theta_i^0}^{\infty} (v_{i1}(\theta_i) - h(\theta_i)v'_{i1}(\theta_i))d\theta_i \\ &+ \int_{\theta_i^1}^{\infty} (v_{i2}(\theta_i) - h(\theta_i)v'_{i2}(\theta_i)) + (v_{i3}(\theta_i) - h(\theta_i)v'_{i3}(\theta_i))d\theta_i \\ &= v_{i1}(\theta_i^0)(1 - F_i(\theta_i^0)) + (v_{i2}(\theta_i^1) + v_{i3}(\theta_i^1))(1 - F_i(\theta_i^1)) \\ &= \mathbf{E}_{\theta_i}(p_i). \end{aligned}$$

We complete the proof of Lemma 5.5 by taking expectations $\mathbf{E}_{\theta_{-i}}$ of the first line and the last line of this equation sequence (which assumed the types of the other bidders, θ_{-i} , were given), and then we sum both sides over all bidders. ■

It is easy to see that both welfare maximization with Levin's payment rule and SW_{MR} satisfy (a), (b), and (c).

Lemma 5.6 For any bidder i and allocation (x_{i1}, x_{i2}) , i 's payment plus virtual valuation is no less than i 's valuation:

$$\begin{aligned} p_i + \tilde{v}_{i1}(\theta_i)x_{i1} + \tilde{v}_{i2}(\theta_i)x_{i2} + \tilde{v}_{i3}(\theta_i)x_{i1}x_{i2} \\ \geq v_{i1}(\theta_i)x_{i1} + v_{i2}(\theta_i)x_{i2} + v_{i3}(\theta_i)x_{i1}x_{i2}. \end{aligned}$$

Proof: If i receives nothing, the above holds with equality.

If i receives one item, say item 1, let θ_i^0 be the lowest type with which i would have won item 1 and θ_i^1 be the type when $\tilde{v}_{i1}(\theta_i^1) = 0$. Clearly, $\theta_i \geq \theta_i^0 \geq \theta_i^1$. Thus,

$$\begin{aligned} p_i + \tilde{v}_{i1}(\theta_i) &= v_{i1}(\theta_i^0) + \tilde{v}_{i1}(\theta_i) \\ &\geq v_{i1}(\theta_i^1) + v_{i1}(\theta_i) - h(\theta_i)v'_{i1}(\theta_i) \\ &= h(\theta_i^1)v'_{i1}(\theta_i^1) + v_{i1}(\theta_i) - h(\theta_i)v'_{i1}(\theta_i) \geq v_{i1}(\theta_i). \end{aligned}$$

The last inequality follows from $h(\theta_i)v'_{i1}(\theta_i)$ being weakly decreasing in θ_i .

If i receives two items, let θ_i^1 be the lowest type with which he would have won both and θ_i^0 be the lowest type with which he would have won one item, say item 1. We must have $\theta_i \geq \theta_i^1 \geq \theta_i^0$. Let θ_i^2 be the type such that $\tilde{v}_{i1}(\theta_i^2) = 0$ and θ_i^3 be

the type such that $\tilde{v}_{i2}(\theta_i^3) + \tilde{v}_{i3}(\theta_i^3) = 0$. According to the Step 1 of SW_{MR} we have $\theta_i^0 \geq \theta_i^2$ and $\theta_i^1 \geq \theta_i^3$. Thus,

$$\begin{aligned} &p_i + \tilde{v}_{i1}(\theta_i) + \tilde{v}_{i2}(\theta_i) + \tilde{v}_{i3}(\theta_i) \\ &= v_{i1}(\theta_i^0) + v_{i2}(\theta_i^1) + v_{i3}(\theta_i^1) + v_{i1}(\theta_i) + v_{i2}(\theta_i) + v_{i3}(\theta_i) \\ &\quad - h(\theta_i)v'_{i1}(\theta_i) - h(\theta_i)v'_{i2}(\theta_i) - h(\theta_i)v'_{i3}(\theta_i) \\ &\geq v_{i1}(\theta_i^2) + v_{i2}(\theta_i^3) + v_{i3}(\theta_i^3) + v_{i1}(\theta_i) + v_{i2}(\theta_i) + v_{i3}(\theta_i) \\ &\quad - h(\theta_i)v'_{i1}(\theta_i) - h(\theta_i)v'_{i2}(\theta_i) - h(\theta_i)v'_{i3}(\theta_i) \\ &= h(\theta_i^2)v'_{i1}(\theta_i^2) + h(\theta_i^3)v'_{i2}(\theta_i^3) + h(\theta_i^3)v'_{i3}(\theta_i^3) \\ &\quad - h(\theta_i)v'_{i1}(\theta_i) - h(\theta_i)v'_{i2}(\theta_i) - h(\theta_i)v'_{i3}(\theta_i) \\ &\quad + v_{i1}(\theta_i) + v_{i2}(\theta_i) + v_{i3}(\theta_i) \\ &\geq v_{i1}(\theta_i) + v_{i2}(\theta_i) + v_{i3}(\theta_i). \end{aligned}$$

Again, the last inequality follows from $h(\theta_i)v'_{i1}(\theta_i)$ being weakly decreasing in θ_i . ■

We are now ready to prove Theorem 3.

Proof: Let \bar{x} be the allocation rule of SW_{MR} and \bar{x}^L be that of Levin's optimal auction. Note that Levin's auction has all the allocation restrictions that ours has (and possibly more) and SW_{MR} maximizes welfare in this restricted set. Therefore, our efficiency is greater (or equal) to that of Levin's. Formally, for any θ ,

$$\begin{aligned} &\sum_{i=1}^n v_{i1}(\theta_i)x_{i1} + v_{i2}(\theta_i)x_{i2} + v_{i3}(\theta_i)x_{i1}x_{i2} \\ &\geq \sum_{i=1}^n v_{i1}(\theta_i)x_{i1}^L + v_{i2}(\theta_i)x_{i2}^L + v_{i3}(\theta_i)x_{i1}^Lx_{i2}^L, \quad (1) \end{aligned}$$

According to Lemma 5.5, we have

$$\mathbf{E}_{\theta} \left(\sum_{i=1}^n p_i \right) = \mathbf{E}_{\theta} \sum_{i=1}^n (\tilde{v}_{i1}(\theta_i)x_{i1} + \tilde{v}_{i2}(\theta_i)x_{i2} + \tilde{v}_{i3}(\theta_i)x_{i1}x_{i2}) \quad (2)$$

According to Lemma 5.6, we have

$$\begin{aligned} &\mathbf{E}_{\theta} \sum_{i=1}^n (p_i + \tilde{v}_{i1}(\theta_i)x_{i1} + \tilde{v}_{i2}(\theta_i)x_{i2} + \tilde{v}_{i3}(\theta_i)x_{i1}x_{i2}) \\ &\geq \mathbf{E}_{\theta} \sum_{i=1}^n (v_{i1}(\theta_i)x_{i1} + v_{i2}(\theta_i)x_{i2} + v_{i3}(\theta_i)x_{i1}x_{i2}). \quad (3) \end{aligned}$$

Combining (2) and (3), we have

$$2\mathbf{E}_{\theta} \left(\sum_{i=1}^n p_i \right) \geq \mathbf{E}_{\theta} \sum_{i=1}^n (v_{i1}(\theta_i)x_{i1} + v_{i2}(\theta_i)x_{i2} + v_{i3}(\theta_i)x_{i1}x_{i2}). \quad (4)$$

Combining (1) and (4), we have

$$\begin{aligned} &2\mathbf{E}_{\theta} \left(\sum_{i=1}^n p_i \right) \geq \mathbf{E}_{\theta} \sum_{i=1}^n (v_{i1}(\theta_i)x_{i1} + v_{i2}(\theta_i)x_{i2} + v_{i3}(\theta_i)x_{i1}x_{i2}) \\ &\geq \mathbf{E}_{\theta} \sum_{i=1}^n (v_{i1}(\theta_i)x_{i1}^L + v_{i2}(\theta_i)x_{i2}^L + v_{i3}(\theta_i)x_{i1}^Lx_{i2}^L) \\ &\geq \text{revenue of Levin's auction}. \quad (5) \end{aligned}$$

The last inequality follows from the individual rationality of Levin's auction. ■

We conjecture that the approximation bound of 2 is tight.

6 Anonymous “reserve prices” in welfare maximization

In many applications, the seller does not have the freedom to use bidder-specific reserve prices. Instead, she is constrained to treating every bidder symmetrically, i.e., the reserve prices must be anonymous. Recall that in symmetric (that is, with i.i.d bidders) single-item settings, the second-price auction—which maximizes welfare—with a monopoly reserve price is optimal [Myerson, 1981]. Without the i.i.d assumption, a Vickrey auction with some anonymous reserve price 4-approximates optimal revenue [Hartline & Roughgarden, 2009]. How well does a welfare-maximizing mechanism with anonymous reserves do in our more general setting?

Definition 6.1 (SW_{AR} mechanism) *In SW_{AR} , we pretend the seller has valuation a for the first item, b for the second item, and c for the bundle. Then, a welfare-maximizing allocation rule is used, followed by the Levin payment rule.*

Theorem 4 SW_{AR} 6-approximates Levin’s optimal revenue.

We prove the theorem through two lemmas as follows.

Lemma 6.1 *Auction 1 below 3-approximates Auction 2:*

1. For each bidder i (not including the seller), introduce a new bidder whose type is drawn i.i.d from the same distribution as i ’s, and whose valuation function is also the same as i ’s. Then run welfare maximization with the Levin payment rule on this enlarged set of bidders.
2. Levin’s optimal auction on the original set of bidders (not including the seller).

Given Lemma 6.1, Theorem 4 follows Lemma 6.2.

Lemma 6.2 *Let a , b , and c be random variables that simulate $\max_i\{v_{i1}\}$, $\max_i\{v_{i2}\}$ and $\max_i\{v_{i1} + v_{i2} + v_{i3}\}$, respectively, in the original bidder set. SW_{AR} with (a, b, c) is then a 2-approximation of Auction 1 of Lemma 6.1.*

Proof: For expected revenue, the reserve prices (a, b, c) simulate (and thus can replace) all the duplicated bidders. ■

7 Conclusions and future research

We introduced a new research avenue for increasing revenue where we poke holes in the allocation space—based on the bids (revealed types)—and then use a welfare-maximizing allocation rule within the remaining allocation set. In this paper, the first step down this avenue, we introduced a new form of “reserve pricing” into CAs. We showed that the optimal revenue for complements can be 2-approximated by using “monopoly reserve prices” to curtail the allocation set, followed by welfare-maximizing allocation and Levin’s payment rule. A key lemma of potential independent interest is that the expected revenue from any truthful allocation-monotonic mechanism equals the expected virtual valuation; this generalizes Myerson’s lemma [1981] from the single-parameter environment. Like the VCG, our mechanism uses welfare-maximization as the allocation rule. Under a mild assumption on valuation distributions, the payment rules also coincide. So, being close to the gold standard of CAs (VCG),

our mechanisms are likely easier to adopt than Levin’s. Our mechanisms also require much less information about the priors, for example, in order to verify correct execution of the mechanism. Our mechanism is also more efficient than Levin’s. Finally, we showed that the optimal revenue can be 6-approximated even if the “reserve pricing” is required to be symmetric across bidders.

Future work includes studying the tightness of our approximation bounds. Second, are there other interesting instantiations of our new “reserve pricing” framework (beyond our monopoly pricing) that are truthful, individually rational, and have some other desirable properties? Future research also includes extending our “reserve pricing” method—or our hole-poking approach more generally—to other environments, such as those with multi-dimensional type spaces.

References

- [Armstrong, 2000] Armstrong, M. 2000. Optimal multi-object auctions. *Review of Economic Studies* 67(3):455–81.
- [Avery & Hendershott, 2000] Avery, C. N., and Hendershott, T. 2000. Bundling and optimal auctions of multiple products. *Review of Economic Studies* 67(3):483–97.
- [Clarke, 1971] Clarke, E. H. 1971. Multipart pricing of public goods. *Public Choice* 2:19–33.
- [Conitzer & Sandholm, 2004] Conitzer, V., and Sandholm, T. 2004. Self-interested automated mechanism design and implications for optimal combinatorial auctions. In *ACM EC*, 132–141.
- [Groves, 1973] Groves, T. 1973. Incentives in Teams. *Econometrica* 41:617–631.
- [Guruswami et al., 2005] Guruswami, V.; Hartline, J. D.; Karlin, A. R.; Kempe, D.; Kenyon, C.; and McSherry, F. 2005. On profit-maximizing envy-free pricing. In *SODA*, 1164–1173.
- [Hartline & Roughgarden, 2009] Hartline, J. D., and Roughgarden, T. 2009. Simple versus optimal mechanisms. In *ACM EC*.
- [Iyengar & Kumar, 2006] Iyengar, G., and Kumar, A. 2006. Characterizing optimal keyword auctions. In *Workshop on Sponsored Search Auctions*.
- [Iyengar & Kumar, 2008] Iyengar, G., and Kumar, A. 2008. Optimal procurement mechanisms for divisible goods with capacitated suppliers. *Review of Economic Design* 12:129–154.
- [Jehiel, Meyer-Ter-Vehn, & Moldovanu, 2007] Jehiel, P.; Meyer-Ter-Vehn, M.; and Moldovanu, B. 2007. Mixed bundling auctions. *Journal of Economic Theory* 127(1):494–512.
- [Levin, 1997] Levin, J. 1997. An optimal auction for complements. *Games and Economic Behavior* 18(2):176–192.
- [Likhodedov & Sandholm, 2004] Likhodedov, A., and Sandholm, T. 2004. Methods for boosting revenue in combinatorial auctions. In *AAAI*, 232–237.
- [Likhodedov & Sandholm, 2005] Likhodedov, A., and Sandholm, T. 2005. Approximating revenue-maximizing combinatorial auctions. In *AAAI*, 267–273.
- [Maskin & Riley, 1989] Maskin, E., and Riley, J. 1989. Optimal multi-unit auctions. In *The Economics of Missing Markets, Information, and Games*. chapter 14, 312–335.
- [Myerson, 1981] Myerson, R. B. 1981. Optimal auction design. *Mathematics of Operations Research* 6(1):58–73.
- [Nisan & Ronen, 2007] Nisan, N., and Ronen, A. 2007. Computationally feasible VCG mechanisms. *JAIR* 29:19–47.
- [Nisan et al., 2007] Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V. 2007. *Algorithmic Game Theory*. Cambridge University Press.
- [Pai & Vohra, 2008] Pai, M. M., and Vohra, R. 2008. Optimal auctions with financially constrained bidders. Discussion Papers 1471, Northwestern University.
- [Vickrey, 1961] Vickrey, W. 1961. Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance* 8–37.