Information Processing Letters 147 (2019) 49–54

Contents lists available at ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

Reverse auctions are different from auctions

Matthias Gerstgrasser [1]

Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford, OX1 3QD, UK

ARTICLE INFO

Article history:
Received 10 May 2018
Received in revised form 5 March 2019
Accepted 11 March 2019
Available online 16 March 2019
Communicated by Krishnendu Chatterjee

Keywords:
Auctions
Reverse auctions
Computational complexity
Correlated priors

ABSTRACT

Auctions are an important theoretical and practical tool in economics, and well studied in the literature. Their procurement siblings, reverse auctions, have received less attention, and are sometimes tacitly assumed to be exact counterparts. We show that for correlated bidders, reverse auctions behave differently from auctions. For two bidders we discuss a simplification of the problem of finding the optimal reverse auction. For $k \geq 3$ bidders, we show that the optimal reverse auction must sometimes buy k copies of the item (and discard all but one of them).

1. Introduction

Within mechanism design, auctions are a major field of interest [1–3]. In this, we consider a single auctioneer who wants to sell an item to one of several bidders, each of whom has a private valuation for the item which the auctioneer does not know. The challenge is to allocate the item according to some measure of optimality based on the private valuations. In addition to social welfare optimisation, in which we aim to allocate the item to the bidder who values it most, revenue maximisation (where we aim to maximise the auctioneer's expected profit) has received major attention. Myerson's seminal result [4] showed that with independent priors, (revenue-) optimal single-item auctions have a closed-form solution. For correlated priors, in contrast to the aforementioned independent-priors setting, this is an intricate computational problem. The case with three or more bidders has been shown to be intractable by Papadimitriou and Pierrakos [5]. On the other hand both [5] as well as Dobzinski et al. [6] show that the optimal auction for *two* bidders can be computed in polynomial time. Both approaches reduce the problem to problems known to be solvable in polynomial time.

In addition to selling an item, auctions may also be used by the auctioneer to buy an item or service from one of multiple sellers. These "reverse" or "procurement" auctions are widely used for instance to solicit bids for public projects. Many results from auctions carry over directly to the reverse auction case. For instance, the VCG mechanism for optimising social welfare works in a reverse auction, as do many other auction formats. So much do these cases appear to be mirror images of one another, that simple reverse counterparts of single-item auctions are rarely discussed explicitly in the literature. Most of the published results on reverse auctions investigate more complex scenarios such as differing quality or service levels from different sellers. To our knowledge, a significant distinction between an auction and its direct reverse counterpart has not been discussed in the literature before.

Our main interest is in exploring the structural properties of correlated reverse auctions. We show that these behave differently than auctions, for any number of bidders. Our results raise interesting questions about the complexity of reverse auctions.

The results in this paper have previously been presented at AAMAS 2018 as part of a larger text concerning the complexity of computing an optimal auction [7].

E-mail address: mgerst@cs.ox.ac.uk.

[1] Recipient of a DOC Fellowship of the Austrian Academy of Sciences.

https://doi.org/10.1016/j.ipl.2019.03.006

We restate them here with improved presentation, and focussing specifically on the asymmetry between the auction and reverse auction setting.

1.1. Previous work

Most of the literature on reverse or procurement auctions specifically seem to focus on more complex settings than the ones we are interested in. One major area of research is when sellers offer goods of differing qualities, see for instance Manelli and Vincent [8]. Chapter 13.5 by Hartline and Karlin [9] in Nisan et al. [10] discusses feasibility constraints in reverse auctions. Several chapters in the same book briefly mention that they consider reverse auctions to be covered by the model they use or similar, e.g. pages 220, 269, 332 therein [10].

To the best of our knowledge, almost no literature looks specifically at the simple reverse auction setting we are interested in. The main exception to this we are aware of is a paper by Minooei and Swamy [11,12], who discuss the more general setting of mechanism design for covering (as opposed to packing) problems. Conitzer and Sandholm [13] discuss collusion in combinatorial auctions and reverse auctions. They take the reverse setting to be a simple parallel of the forward case (as we do here), except for an explicit constraint on allowed allocations (in their case, for the VCG mechanism) which we also assume in this paper. They consider, among other results, the complexity of computing whether collusion is possible in a (forward or reverse) auction, showing that this is NP-hard even for 2 colluders.

Most relevant to our discussion is the literature on the complexity of optimal correlated auctions. Papadimitriou and Pierrakos [5] show that for two bidders, a (revenue-) optimal auction can be found in polynomial time. Their algorithm reduces the problem to finding a maximum-weight independent set on a bipartite graph, with edges encoding allocation constraints of the auction. This yields an algorithm that runs in time $\mathcal{O}(n^6)$ for prior support size n^2 (each bidder's valuation taking one of n discrete values). For three or more bidders they show that it is NP-hard to approximate the optimal auction to within a factor of 1.0005. Dobzinski et al. [6] also give a polynomial algorithm for the two-bidder auction. They show that a truthful-in-expectation mechanism found via an LP can be derandomised. Furthermore they show that a polynomial-time algorithm for two bidders extends to a polynomial-time approximation algorithm for many bidders through a "2-lookahead" auction. This builds on previous work by Ronen [14] and Ronen and Saberi [15]. In all these, as well as this paper, the focus is mainly on deterministic mechanisms, as [6] shows that the randomised case is easy.

2. Preliminaries

We begin by considering the familiar single-item auction, in which an auctioneer wishes to sell one item to one of several bidders, numbered $1, \ldots, k$. We assume each bidder i has valuation v_i, which can take one of several discrete values. For ease of notation we take $v_i \in \{1, \ldots, n\} = [n]$. It is easily checked that none of our results depend on this. Let F denote the (joint) prior probability distribution over $\mathbf{v} = (v_1, \ldots, v_k)$. Our interest is only in deterministic mechanisms, which consist of allocation functions $x_i(\mathbf{v})$ together with payment functions $p_i(\mathbf{v})$ for each bidder. Let $x_i(\mathbf{v}) = 1$ if bidder i wins the item given bid vector $\mathbf{v}$, and $x_i(\mathbf{v}) = 0$ otherwise.

Given that we assume the auctioneer only has a single copy of the item to sell, we require $\sum_i x_i(\mathbf{v}) \le 1$ for all $\mathbf{v}$. We assume quasilinear utilities, and require the usual notions of dominant strategy incentive compatibility (DSIC) and individual rationality (IR), as defined formally below.

$$\text{(Utilities)} \quad u_i(\mathbf{v}) = v_i x_i(\mathbf{v}) - p_i(\mathbf{v}) \tag{1a}$$

$$\text{(DSIC)} \quad v_i x_i(\mathbf{v}) - p_i(\mathbf{v}) \ge v_i x_i(v_i', \mathbf{v}_{-i}) - p_i(v_i', \mathbf{v}_{-i}) \quad \forall i, \mathbf{v}, v_i' \tag{1b}$$

$$\text{(IR)} \quad u_i(\mathbf{v}) \ge 0 \quad \forall i, \mathbf{v} \tag{1c}$$

$$\text{(1-item)} \quad \sum_i x_i(\mathbf{v}) \le 1 \quad \forall \mathbf{v} \tag{1d}$$

We therefore can assume that players' bids are equal to their valuations. The auctioneer's aim will be to maximise their expected revenue $\mathbb{E}[p_i(\mathbf{v})]$. Formally, we may now define the optimal auction design problem as follows.

Definition 1. The optimal auction design problem takes as input a prior distribution F as defined above, given explicitly as a list of n^k values. The desired output is a pair of allocation and payment functions $\mathbf{x}()$ and $\mathbf{p}()$ that satisfy conditions (1a)-(1d) above, and which maximise the expected revenue $\mathbb{E}[p_i(\mathbf{v})]$. These may be output explicitly as a list of their kn^k values.

By Myerson [4], truthfulness in this domain for deterministic mechanisms is equivalent to monotone allocations, and the corresponding uniquely determined payments - the winner's critical bid. That is, if bidder i wins the auction given bid profile $\mathbf{v}$, then they also win the auction for bid profile $(v_i', \mathbf{v}_{-i})$, for any $v_i' > v_i$. Their payment will be the smallest $v_i' \le v_i$ such that they would still win the auction given bid profile $(v_i', \mathbf{v}_{-i})$. If bidder i does not win they pay nothing (by IR (1c)).

$$x_i(\mathbf{v}) = 1 \Rightarrow \forall v_i' \ge v_i : x_i(v_i', \mathbf{v}_{-i}) = 1 \tag{2a}$$

$$p_i(\mathbf{v}) = \min\left\{v_i' : x_i(v_i', \mathbf{v}_{-i}) = 1\right\} \quad \text{if } x_i(\mathbf{v}) = 1, \text{ else } p_i(\mathbf{v}) = 0 \tag{2b}$$

Papadimitriou and Pierrakos [5] give a very elegant geometric representation of this condition: For each bidder i, their critical bid is given by a function $\alpha_i(\mathbf{v}_{-i})$ of the other bidders' bids, where $x_i(\mathbf{v}) = 1$ iff $v_i \ge \alpha_i(\mathbf{v}_{-i})$. Consider now for each bidder i the region $A_i = \{\mathbf{v} : v_i \ge \alpha_i(\mathbf{v}_{-i})\}$ of all bid vectors for which i wins the item. Clearly this is bordered by $\alpha_i(\mathbf{v}_{-i})$. Furthermore, if $(v_i, \mathbf{v}_{-i}) \in A_i$, then also $(v_i', \mathbf{v}_{-i}) \in A_i$ for all $v_i' \ge v_i$. This follows both from the definition of A_i as the region bounded below by the graph of a function of v_{-i}, as well as directly from monotonicity. We will also say that A_i is "upward-closed in direction v_i" for this. The 1-item constraint (1d) entails

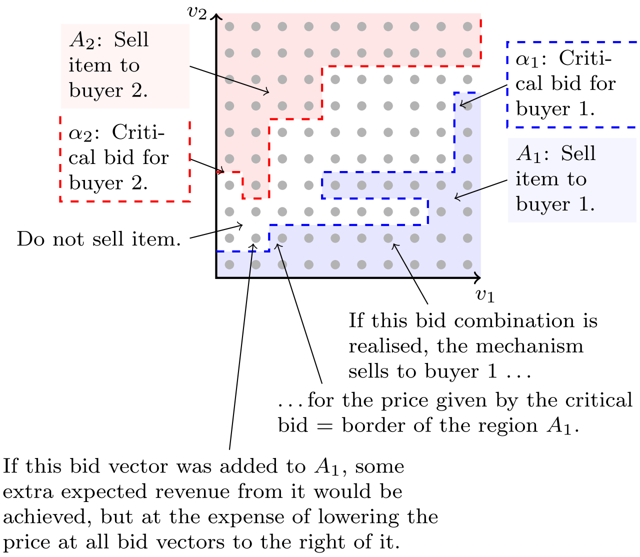

Fig. 1. A mechanism as a partition of the bid space into regions of allocation, and the corresponding critical bid functions. Discrete prior support shown dotted, with critical bid functions and allocation regions drawn slightly larger for easier readability. (We are being slightly imprecise here: The graph of α_1 is the vertical part of the dashed blue line; The graph of α_2 is the horizontal part of the dashed red line.) (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

that any two A_i must be disjoint. In summary, the picture we get is that looking for the optimal k-bidder auction is looking for a partition of the space of possible bid combinations into $k+1$ regions: k regions where the item is sold to each of the buyers (each upward-closed in the corresponding direction), and one where the item is not sold. Fig. 1 shows this picture for the two-bidder case. Taking bidder 1's bid to be on the x-axis and bidder 2's on the y-axis, we are looking for A_1 to be rightward-closed, and A_2 to be upward-closed. There is a two-fold trade-off: smaller $\alpha_i(v_{-i})$ means higher probability of drawing $v_i \geq \alpha_i(v_{-i})$, but selling at a lower price if so. Smaller $\alpha_i(v_{-i})$ also means "blocking" more bid vectors for the other bidder. We will often identify a mechanism through either the regions A_i or the functions α_i. When defining a mechanism this way, DSIC and IR are automatic. The 1-item constraint (1d) for two bidders can be restated as a non-crossing property [5]. Equivalently it may be stated as the disjointness requirement that no two A_i may overlap.

In the reverse auction setting we are interested in, again a single auctioneer faces k bidders having their valuations drawn from a joint distribution F supported on $[n]^k$. We assume that each bidder holds one copy of a single type of item, each bidder's copy identical to all others', and that the auctioneer wishes to procure one copy. For simplicity we now write $x_i(\mathbf{v}) = -1$ if the mechanism buys a copy of the item from bidder i. This allows us to leave the definitions of utilities, DSIC and IR in equations (1a)-(1c) unchanged.[2] It is easy to see that in this context it makes little sense to require that the mechanism buys *at most* one copy of the item. Instead we require the mechanism to always buy *at least* one copy of the item, i.e. we require that $\sum_i x_i(\mathbf{v}) \leq -1$, replacing the corresponding constraint (1d). Formally, we may define the optimal reverse auction design problem as follows.

Definition 2. The optimal reverse auction design problem takes as input a prior distribution F as defined above, given explicitly as a list of n^k values. The desired output is a pair of allocation and payment functions $\mathbf{x}()$ and $\mathbf{p}()$ that satisfy conditions (3a)-(3d), and which minimise the expected revenue $\mathbb{E}[p_i(\mathbf{v})]$. These may be output explicitly as a list of their kn^k values.

$$\text{(Utilities)}\quad u_i(\mathbf{v}) = v_i x_i(\mathbf{v}) - p_i(\mathbf{v}) \tag{3a}$$

$$\text{(DSIC)}\quad v_i x_i(\mathbf{v}) - p_i(\mathbf{v}) \geq v_i x_i(v_i', \mathbf{v}_{-i}) - p_i(v_i', \mathbf{v}_{-i}) \quad \forall i, \mathbf{v}, v_i' \tag{3b}$$

$$\text{(IR)}\quad u_i(\mathbf{v}) \geq 0 \quad \forall i, \mathbf{v} \tag{3c}$$

$$\text{(1-item)}\quad \sum_i x_i(\mathbf{v}) \leq -1 \quad \forall \mathbf{v} \tag{3d}$$

Geometrically, we get a very similar picture of regions A_i in which the mechanism buys from bidder i. However, they now need to be downward-closed in direction of v_i. In the two-bidder case, A_1 ought to be leftward-closed, and A_2 downward-closed. Secondly, two or more of the A_i may now overlap (when the mechanism buys two or more copies). The constraint that $\sum_i x_i(\mathbf{v}) \leq -1$ means that the union of all A_i must cover all of the bid space.

3. Our results

Recall that we are looking for regions A_i which are downward-closed in direction v_i, may overlap, and must cover all of the bid space. In an auction, a large part of

[2] Taking instead v_i to be nonpositive, or adjusting equations (1a)-(1c) is equivalent.

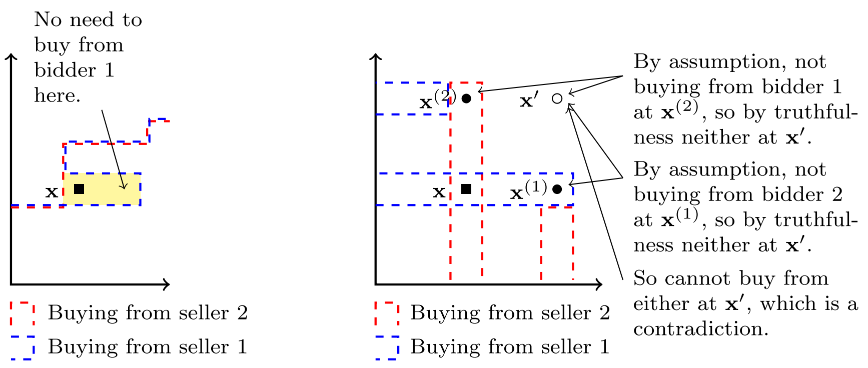

Fig. 2. From left to right: (a) In the first case of Theorem 1, if there is no point to the right of $\mathbf{x}$ in which we buy only from seller 1, we can improve our cost by not buying from seller 1 in the shaded region, i.e. moving the blue line so it coincides with the red one. (b) Buying from either seller is blocked at $\mathbf{x}'$ due to truthfulness, in the second case of Theorem 1. By assumption we do not buy from seller 1 at $\mathbf{x}^{(2)}$, and thus by truthfulness cannot buy from seller 1 at $\mathbf{x}'$. Vice versa we assume we do not buy from seller 2 at $\mathbf{x}^{(1)}$, and so cannot buy from them at $\mathbf{x}'$ either. That leaves us with noone to buy the item from at $\mathbf{x}'$, violating our feasibility constraint.

the auctioneer's power comes from the option of not selling the item. Indeed, reserve prices below which the item is not sold are at the heart of Myerson's seminal result [4]. For a single bidder, not selling is even *all* the power the auctioneer has to achieve any revenue. In a reverse auction, the counterpart of this is to buy multiple copies of the item from multiple sellers. In a way, both of these cases are suboptimal locally, but allow for higher expected revenue globally: If for a given bid vector $\mathbf{v}$ the auctioneer does not sell in the auction this clearly foregoes some potential contribution to expected revenue arising from selling at $\mathbf{v}$. But, it may allow the auctioneer to generate a higher contribution to expected revenue (through higher prices) at some other bid vectors. Similarly in the reverse auction, buying from multiple bidders at a bid vector $\mathbf{v}$ clearly incurs a double or multiple contribution to expected cost arising from $\mathbf{v}$. But, it may allow the auctioneer to achieve a lower expected cost elsewhere in the bid space as a result.

It is easy to see that the possibility of buying from multiple bidders generates a much richer space of potential outcomes than in the auction. Whereas in the auction there is $k+1$ possible allocations for each bid vector (selling to each of the bidders, plus selling to none of them), in the reverse auction we potentially have to deal with 2^k-1 possible allocation (buying from any combination of bidders, except from none of them). The question we deal with in this section is whether all of these are actually relevant to the problem of finding the optimal reverse auction. That is, will all of these occur in an optimal mechanism? The answer is surprising: "Yes", for $k \geq 3$ bidders, but "No" for 2 bidders. So, in the former case, the reverse auction is clearly structurally different than the corresponding auction. The latter is surprising in itself, as a priori both the 2-bidder auction as well as the 2-bidder reverse auction potentially have three valid allocations. As it turns out, not even these two cases are structurally the same.

Theorem 1. *In the single-item reverse auction with two correlated sellers, the optimal mechanism will never buy from both bidders.*

Proof. Suppose for bid vector $\mathbf{x}$ we buy from both sellers. We consider two cases. Firstly, suppose that for one bidder, wlog for bidder 1, there exists no point $\mathbf{x}' = (x_1', x_2)$ with $x_1' > x_1$ we buy only from seller 1. Then we could strictly improve our cost if we did not buy from 1 at $\mathbf{x}$ and all those bid vectors $\mathbf{x}' = (x_1', x_2)$ with $x_1' > x_1$. Thus the mechanism was not optimal. See Fig. 2 (a) for an illustration. And similarly, if there were no bid vector $\mathbf{x}' = (x_1, x_2')$ with $x_2' > x_2$ where we bought only from seller 2, the mechanism was not optimal.

So assume that for both sellers there exists such a bid vector as above. That is, for bidder 1, there exists a bid vector $\mathbf{x^{(1)}} = (\mathbf{x_1'}, \mathbf{x_2})$, with $x_1' > x_1$, so that we buy only from seller 1 at $x^{(1)}$. And for bidder 2, there exists a bid vector $\mathbf{x^{(2)}} = (\mathbf{x_1}, \mathbf{x_2'})$, with $x_2' > x_2$, so that we buy only from seller 2 at $x^{(2)}$. But then by truthfulness it follows that at $\mathbf{x}' = (x_1', x_2')$ we cannot buy from either of the sellers. (If we bought from seller 1 at $\mathbf{x}'$, we would also buy from seller 1 at $\mathbf{x^{(2)}}$ by truthfulness, but that contradicts our assumption. Vice versa for seller 2.) But not buying at all at $\mathbf{x}'$ is not a valid mechanism by definition. Fig. 2 (b) shows this case.

Thus, the optimal (valid) reverse auction can never buy from both bidders at once. □

An immediate consequence of this result is that the optimal mechanism design problem in this setting is simpler than in the auction setting: We are now only looking for a partition of the bid space into *two* regions A_1 and $A_2 = A_1^c$.

Surprisingly, for $k \geq 3$ bidders the opposite holds: It is possible to construct instances in which it is optimal to buy all k copies of the item.

Theorem 2. *For three or more bidders, the optimal reverse auction may buy from all sellers.*

Proof. To show this, we will construct an instance. Our main gadget will be of the following form: Consider points $\mathbf{p}_1 = (c_L, c_M, c_H)$ and $\mathbf{p}_2 = (c_M, c_L, c_H)$ with high probability weight, and a third point $\mathbf{q} = (c_M, c_M, c_H)$ with very low probability weight, for some constants $c_L << c_M << c_H$.

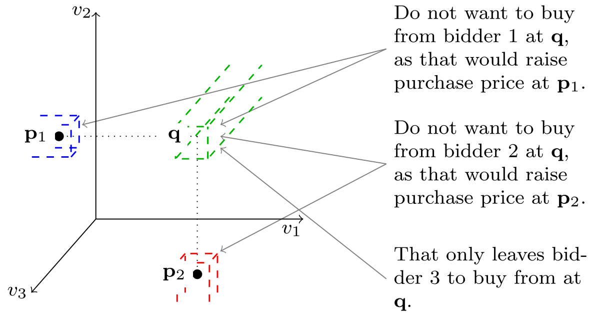

Fig. 3. The gadget we will use in the proof of Theorem 2. High probability weight on points $\mathbf{p}_1$ and $\mathbf{p}_2$ makes it optimal to buy from seller 3 (in green) at point $\mathbf{q}$. Buying from either of the other sellers at $\mathbf{q}$ would raise the purchase price at either $\mathbf{p}_1$ or $\mathbf{p}_2$, thus raising the expected cost. By monotonicity, the mechanism then must also buy from seller 3 at all points behind $\mathbf{q}$ in this view.

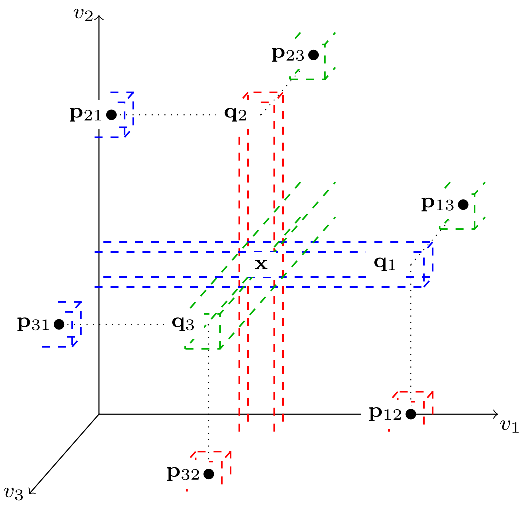

Fig. 4. Three gadgets make up the construction used in the proof of Theorem 2. Notice that for bid vector $\mathbf{x}$ (in the centre at the intersection of the three $\mathbf{q}_i$-segments), the mechanism will buy from all three sellers, due to monotonicity and the allocation at the points $\mathbf{q}_i$.

We will want to make this so that the optimal mechanism will want to buy at the point $\mathbf{p}_1$ cheaply from seller 1 - and thus cannot buy at point $\mathbf{q}$ from seller 1, as by monotonicity that would also raise the purchase price at $\mathbf{p}_1$. Similarly for buyer 2 and points $\mathbf{p}_2$ and $\mathbf{q}$. As a consequence, it will want to buy at $\mathbf{q}$ from seller 3. This will be at a very high purchase price, but if the probability weight on $\mathbf{q}$ is small enough, this will still be optimal in expectation. Fig. 3 illustrates this construction. By monotonicity it follows that if the mechanism buys from seller 3 at $\mathbf{q} = (c_\mathrm{M}, c_\mathrm{M}, c_\mathrm{H})$, it must also buy from seller 3 at all points $(c_\mathrm{M}, c_\mathrm{M}, v_3)$, $v_3 \le c_\mathrm{H}$.

By creating three such gadgets in the right places and rotated appropriately, we can then make it optimal to buy from all three sellers at the intersection of these $\mathbf{q}$-segments. Consider the construction in Fig. 4. In this we have one gadget consisting of $\mathbf{p}_{12} = (c_\mathrm{H}, c_\mathrm{L}, c_\mathrm{M})$, $\mathbf{p}_{13} = (c_\mathrm{H}, c_\mathrm{M}, c_\mathrm{L})$ and $\mathbf{q}_1 = (c_\mathrm{H}, c_\mathrm{M}, c_\mathrm{M})$ with the auctioneer buying from bidder 1 in the $\mathbf{q}_1$-segment, similarly one gadget consisting of $\mathbf{p}_{21}$, $\mathbf{p}_{23}$ and $\mathbf{q}_2$ for bidder 2, and a third one

$v_3=1$	$v_1=1$	$v_1=2$	$v_1=3$
$v_2=3$	1	**3**	3
$v_2=2$	3	3	**3**
$v_2=1$	2	2	2
$v_3=2$	$v_1=1$	$v_1=2$	$v_1=3$
$v_2=3$	**1**	2	3
$v_2=2$	1	1, 2, 3	1
$v_2=1$	2	2	**2**
$v_3=3$	$v_1=1$	$v_1=2$	$v_1=3$
$v_2=3$	1	1	3
$v_2=2$	**1**	3	2
$v_2=1$	2	**2**	2

Fig. 5. The full allocation for the instance in Theorem 2. Each cell shows the bidder(s) the mechanism buys from at the given bid vector. High probability points are shown in bold face.

comprising $\mathbf{p}_{31}$, $\mathbf{p}_{32}$ and $\mathbf{q}_3$ for bidder 3. Again let there be very high probability weight on the $\mathbf{p}_{ij}$, and very small probability weight ϵ on the $\mathbf{q}_i$ (and everywhere else). The $\mathbf{q}_i$ are placed such that the $\mathbf{q}_i$-segments intersect at the point $\mathbf{x} = (c_\mathrm{M}, c_\mathrm{M}, c_\mathrm{M})$. It is easy to check that the optimal mechanism will indeed buy from bidder i in each $\mathbf{q}_i$-segment, for ϵ small enough. It will thus buy from all three bidders at $\mathbf{x}$.

To show this formally, wlog we take the prior support to be $[3]^3$, and $c_L = 1, c_M = 2, c_H = 3$. It is easy to see that the following arguments work for any other choice of these constants. Let there be probability weight $\frac{1-\epsilon}{6}$ on points $\mathbf{p}_{12} = (3, 1, 2)$, $\mathbf{p}_{13} = (3, 2, 1)$, $\mathbf{p}_{21} = (1, 3, 2)$, $\mathbf{p}_{23} = (2, 3, 1)$, $\mathbf{p}_{31} = (1, 2, 3)$, $\mathbf{p}_{32} = (2, 1, 3)$, and probability weight $\frac{\epsilon}{21}$ on each of the remaining 21 points of the prior support. We will denote by $\mathbf{q}_1 = (3, 2, 2)$, $\mathbf{q}_2 = (2, 3, 2)$, $\mathbf{q}_3 = (2, 2, 3)$ among these. Notice how for $i = 1, 2, 3$ each of these sets of two $\mathbf{p}_{ij}$ and one $\mathbf{q}_i$ forms one of the gadgets discussed at the start of this proof.

We will proceed as follows. First, we show that the optimal mechanism has the property that the auctioneer buys each of the points $\mathbf{p}_{ij}$ for price 1 from seller j (and only seller j). There are two things to check here: In step 1(a), we show that a valid allocation exists that has this property. In step 1(b), we show that any allocation with this property has lower expected cost than any allocation without this property. Second, in step 2, we deduce from this that the optimal mechanism buys from all three sellers at point $\mathbf{x} = (2, 2, 2)$.

Step 1(a): There is a valid mechanism that buys from seller j (and only seller j) at each point $\mathbf{p}_{ij}$, for price 1: This is easy to see. We show one such mechanism in Fig. 5. Each cell lists the bidder(s) that the item is bought from for this given bid vector. The high probability bid vectors are shown in bold face.

Step 1(b): Any mechanism that allocates at the points $\mathbf{p}_{ij}$ in this manner has lower expected cost than any mechanism that does not. We show this by giving first an upper bound on the expected cost of any mechanism with this property. This consists of an exact expression for the contribution to expected cost incurred at the $\mathbf{p}_{ij}$, plus an upper bound on the contribution at all the other points. Second, we give a lower bound of the expected cost of any mechanism which does not have this property. For this it suffices to lower bound the expected cost incurred at the $\mathbf{p}_{ij}$.

So, assume a mechanism allocates at the $\mathbf{p}_{ij}$ in the manner claimed. Then the expected cost can be (very crudely) bounded above by $6 \cdot 1 \cdot 1 \cdot (\frac{1-\epsilon}{6}) + 21 \cdot 3 \cdot 3 \cdot (\frac{\epsilon}{21}) = 1 + 8\epsilon$. The first term is the contribution from the 6 points $\mathbf{p}_{ij}$ where we buy at price 1 from exactly 1 seller with probability $\frac{1-\epsilon}{6}$ each, the second term a bound from the 21 remaining points, where we buy from at most from 3 sellers, for at most a price of 3, with probability $(\frac{\epsilon}{21})$ each.

On the other hand, if a mechanism allocated at any of the $\mathbf{p}_{ij}$ differently (while maintaining monotonicity), that would mean either raising the purchase price to at least 2 at a $\mathbf{p}_{ij}$ (either due to buying from the same bidder at a higher price, or buying from a different bidder at price ≥ 2), or buying from more than one buyer at a $\mathbf{p}_{ij}$. Either way we would incur at least an extra $(\frac{1-\epsilon}{6})$ expected cost at one of the p_{ij}. The resulting total expected cost of the mechanism would thus also be at least $7 \cdot (\frac{1-\epsilon}{6})$.

It is easy to check that $1 + 8\epsilon$ is less than $7(\frac{1-\epsilon}{6})$ if $\epsilon < \frac{1}{55}$. So, for any such ϵ the optimal mechanism will have the property that the auctioneer buys each of the points $\mathbf{p}_{ij}$ for price 1 from seller j (and only seller j).

Step 2: Since the optimal mechanism buys from seller j for price 1 at each $\mathbf{p}_{ij}$, it follows that it buys from bidder i at each $\mathbf{q}_i$, as buying from either of the other bidders would contradict the low buying price at a $\mathbf{p}_{ij}$. Therefore by monotonicity, it will buy from all three sellers at $\mathbf{x} = (2, 2, 2)$. For k bidders, this construction easily generalises. Use k gadgets, each with $k - 1$ points $\mathbf{p}_{ij}$ forcing the mechanism to buy point $\mathbf{q}_i$ from the remaining bidder. This shows our claim. □

4. Discussion and future work

Our results on correlated reverse auctions for the first time (to our knowledge) show an asymmetry between auctions and reverse auctions. For two bidders, a further structural analysis allows us to show a small reduction in complexity compared to the auction. Our result for three or more bidders is surprising, as it shows a much higher dimensional space of possible outcomes - exponential (in the number of bidders) compared to linear in an auction. This suggests an exciting uncertainty regarding the complexity of the optimal reverse auction design problem. More generally, we take our results as evidence that the reverse auction case is interesting to consider as a separate problem from the standard auction model. We have shown that at least for the correlated prior case they behave structurally different. A main question for future work is to investigate differences in other settings.

Acknowledgement

We thank Paul Goldberg, Elias Koutsoupias and especially Maria Kyropoulou and Aris Filos-Ratsikas for helpful feedback and discussions.

References

[1] P. Klemperer, Auctions: Theory and Practice, Princeton University Press, 2004, https://books.google.co.uk/books?id=E56FQgAACAAJ.
[2] P. Klemperer, Auction theory: a guide to the literature, J. Econ. Surv. 13 (3) (1999) 227–286.
[3] V. Krishna, Auction Theory, Academic Press, 2009.
[4] R.B. Myerson, Optimal auction design, Math. Oper. Res. 6 (1) (1981) 58–73.
[5] C.H. Papadimitriou, G. Pierrakos, On optimal single-item auctions, in: Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing, ACM, 2011, pp. 119–128.
[6] S. Dobzinski, H. Fu, R.D. Kleinberg, Optimal auctions with correlated bidders are easy, in: Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing, ACM, 2011, pp. 129–138.
[7] M. Gerstgrasser, On the complexity of optimal correlated auctions and reverse auctions, in: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, International Foundation for Autonomous Agents and Multiagent Systems, 2018, pp. 605–613.
[8] A.M. Manelli, D.R. Vincent, Optimal procurement mechanisms, Econometrica (1995) 591–620.
[9] J. Hartline, A. Karlin, Profit maximization in mechanism design, in: N. Nisan, T. Rougarden, E. Tardos, V.V. Vazirani (Eds.), Algorithmic Game Theory, Cambridge University Press, 2007, p. 331.
[10] N. Nisan, T. Roughgarden, E. Tardos, V.V. Vazirani, Algorithmic Game Theory, Cambridge University Press, 2007.
[11] H. Minooei, C. Swamy, Near-optimal and robust mechanism design for covering problems with correlated players, in: International Conference on Web and Internet Economics, Springer, 2013, pp. 377–390.
[12] H. Minooei, C. Swamy, Near-optimal and robust mechanism design for covering problems with correlated players, ACM Trans. Econ. Comput. 3 (4) (2015) 19.
[13] V. Conitzer, T. Sandholm, Failures of the VCG mechanism in combinatorial auctions and exchanges, in: Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems, ACM, 2006, pp. 521–528.
[14] A. Ronen, On approximating optimal auctions, in: Proceedings of the 3rd ACM Conference on Electronic Commerce, ACM, 2001, pp. 11–17.
[15] A. Ronen, A. Saberi, On the hardness of optimal auctions, in: Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science, IEEE, 2002, pp. 396–405.