On the Complexity of Optimal Correlated Auctions and Reverse Auctions

Matthias Gerstgrasser*
Department of Computer Science, University of Oxford,
Wolfson Building, Parks Road,
OXFORD, OX1 3QD, UK
mgerst@cs.ox.ac.uk

ABSTRACT

We investigate the problem of finding a revenue-optimal auction with correlated bidders. We give an algorithm for the exact solution for two bidders, and for a $\frac{5}{3}$-approximation for many bidders, improving from $O(n^6)$ runtime to $O(n^3)$ for both problems by exploiting structural properties of this problem directly. We show that for correlated bidders, reverse auctions behave differently from auctions. For two bidders we discuss a constant-factor reduction in complexity. For $k \geq 3$ bidders, we show that the optimal reverse auction must sometimes buy k copies of the item.

ACM Reference Format:
Matthias Gerstgrasser. 2018. On the Complexity of Optimal Correlated Auctions and Reverse Auctions. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 9 pages.

1 INTRODUCTION

Within mechanism design, auctions are a major field of interest [12–14]. 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 [20] showed that with independent priors, (revenue-) optimal single-item auctions have a closed-form solution: In the deterministic case, in which we are solely interested here, the item is sold to the bidder with the highest "virtual valuation" and their payment is their critical bid. This also gives the more general theory: For a specific kind of "truthful" allocation functions together with uniquely determined payments, bidders are incentivised to reveal their true valuations to the auctioneer. We may therefore regard the problem as one of finding allocation functions that satisfy this truthfulness constraint. 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 [22]; but on the other hand both [22] as well as Dobzinski et al. [8] show that the optimal auction for *two* bidders can be computed in polynomial time; both approaches reduce the problem to generic known-polynomial problems.

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 being in exploring the structural properties of correlated auctions further. For the two-bidder auction, this allows us to construct a $O(n^3)$ algorithm, which improves on $O(n^6)$ of previous approaches. Ours is the first algorithm to exploit directly structural properties of the problem. For reverse auctions, we show that these behave differently than auctions, for any number of bidders; this raises interesting questions about their complexity.

1.1 Previous Work

For a single-item auction with bidders' valuations drawn from independent distributions, Myerson [20] shows that maximising revenue is equivalent to maximising virtual welfare. Several variations of optimal correlated auctions have been investigated. Most relevant to our discussion is the literature on the complexity of optimal correlated auctions in which the joint prior is given explicitly or as an oracle. Papadimitriou and Pierrakos [22] 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 $O(n^6)$ 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. [8] independently 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. The runtime of this approach depends on the LP-solver chosen; standard interior point methods give $O(n^7)$. Furthermore they investigate k-lookahead auctions, in which an

*Recipient of a DOC Fellowship of the Austrian Academy of Sciences.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

optimal auction is run on the highest k bidders' conditional distribution. They show that a polynomial-time algorithm for two bidders extends to a polynomial-time approximation algorithm for many bidders through the 2-lookahead auction. This builds on previous work by Ronen [23] and Ronen and Saberi [24]. Chen et al. [2] investigate the approximation ratio of the k-lookahead auction further. Caragiannis et al. [1] improve results by [8] on separation between deterministic vs. randomised expected revenue, and on the lower bound on the approximation ratio for the three-bidder auction by [22]. Diakonikolas et al. [7] show that an approximate trade-off between revenue and social welfare can be computed efficiently for two bidders. Esö [9] investigate optimal auctions for risk-averse buyers and sellers. A related setting with interdependent values has been investigated by several authors, see for instance [4, 15, 17]. Crémer and McLean [5, 6] discuss conditions for full surplus extraction with interim individual rationality. A similar line of inquiry to [8, 22] was pursued by Gerstgrasser et al. [10] for the setting of market intermediation with correlated priors, showing a polynomial-time algorithm for two bidders, and NP-hardness for one variant involving three bidders. Sections 3.2 and 3.3 in this paper are an extension of the results therein. In all these, as well as this paper, the focus is mainly on deterministic mechanisms, as [8] show that the randomised case is easy.

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 [16]. Chapter 13.5 by Hartline and Karlin [11] in Nisan et al. [21] discuss 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 [21]. 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 [18, 19], who discuss the more general setting of mechanism design for covering (as opposed to packing) problems. Conitzer and Sandholm [3] 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.

2 PRELIMINARIES

We consider a 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 $\sum_i x_i(\mathbf{v}) \leq 1$ for all $\mathbf{v}$. We assume quasilinear utilities, and require the usual notions of truthfulness / DSIC and individual rationality, as defined formally below. We therefore can assume that players' bids are equal to their valuations. The auctioneer's aim will be to maximise their expected revenue $\mathbb{E}[p_i(\mathbf{v})]$.

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

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

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

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

By Myerson [20], 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 $v_i' > v_i$; and their payment will be the smallest $v_i' \leq v_i$ 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)).

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

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

Papadimitriou and Pierrakos [22] 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 $v_i \geq \alpha_i(\mathbf{v}_{-i})$. Consider now for each bidder i the region $A_i = \{\mathbf{v} : v_i \geq \alpha_i(\mathbf{v}_{-i})\}$ 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 $v_i' \geq v_i$. 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 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. Figure 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 tradeoff: smaller $\alpha_i(u)$ means higher probability of drawing $v_i \geq \alpha_i(u)$, but selling at a lower price if so; smaller $\alpha_i(u)$ 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; we will write $\alpha_i(\mathbf{v}_{-i}) = n + 1$, if none of the bid vectors $(v_i, \mathbf{v}_{-i})$ are in A_i. When defining a mechanism through the α_i, DSIC and IR are automatic. The 1-item constraint (1d) for two bidders can be restated as a non-crossing property (3) [22]. The expected revenue has a simple

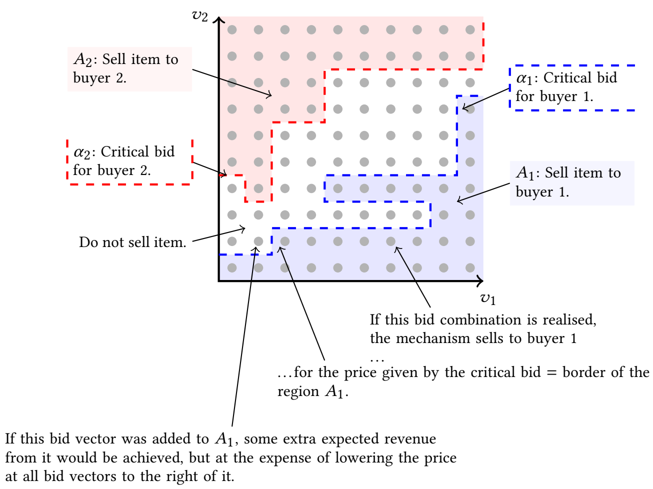

Figure 1: A mechanism as a partition of the bid space into regions of allocation, and the corresponding critical bid functions. Discrete prior support shown dotted, with critical bid functions and allocation regions drawn slightly larger for easier readability. (We are being slightly imprecise here: 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.)

closed form in terms of α_i given as equation (4) for two bidders.

$$\text{(Non-Crossing Property)} \quad v_1 \geq \alpha_1(v_2) \Rightarrow v_2 < \alpha_2(v_1) \quad \forall v_1, v_2 \tag{3}$$

$$R = \sum_{j=1}^{n} \left[\alpha_1(j) \cdot \sum_{\ell=\alpha_1(j)}^{n} F(\ell, j) \right] + \sum_{\ell=1}^{n} \left[\alpha_2(\ell) \cdot \sum_{j=\alpha_2(\ell)}^{n} F(\ell, j) \right] \tag{4}$$

In a reverse auction, 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.[1] 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).

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 THE $O(n^3)$ ALGORITHM FOR THE TWO-BIDDER AUCTION

We now show how to compute the optimal two-bidder auction in time $O(n^3)$. This immediately gives an $O(n^3)$ $\frac{5}{3}$-approximation algorithm for many bidders, as detailed in Dobzinski et al. [8]. We do this in three steps. First, we show that we only need to find the optimal allocation function for one of the bidders; finding the second bidder's allocation function is then very easy. Indeed, precomputing all possible optimal allocations for the second bidder is easy. Second, we show that the optimal critical bid for bidder 1 in each $v_2 = c$ "row" only depends in a limited way on the critical bids in other rows. Third, we use these results to construct a very simple bottom-up dynamic programming algorithm.

3.1 Step 1: Disentangling the Two Bidders from One Another

One source of complexity in finding the optimal auction with correlated priors is that the two bidders' allocations interact: taking still v_1 to be the horizontal axis and v_2 the vertical one; if at a bid vector $\mathbf{w}$ we allocate the item to bidder 1 (i.e. $w \in A_1$), then by monotonicity we also do so at all bid vectors to its right. In turn, this means that we cannot sell to bidder 2 at any of the bid vectors $\mathbf{u} : u_1 \geq w_1 \wedge u_2 \leq w_2$ to the bottom right of $\mathbf{w}$. Vice versa, if at $\mathbf{w}$ we sell to bidder 2, we cannot sell to bidder 1 at any of the points to its top left. This argument applies repeatedly: Allocation to bidder 1 influences potential (and thus also optimal) allocation to bidder 2, which in turn influences potential & optimal allocation to bidder 1. A simple lemma shows how to disentangle the two bidders' allocation: Suppose (say) bidder 1's allocation is fixed. How does the choice of a value $\alpha_2(\ell)$ now influence the optimal choices of all other values of α_2? Simple: It does not. While the inclusion or exclusion of (ℓ, v_2) in A_2 influences other points (ℓ, v_2') in the same "column" through monotonicity for bidder 2; the only way it could influence a point $(\ell', v_2'), \ell' \neq \ell$ in another "column" is through monotonicity for bidder 1. But by assumption bidder 1's allocation is already fixed. Therefore, each choice of $\alpha_2(\ell)$ is independent of all the others. Furthermore, clearly in column $v_1 = \ell$ the optimal thing to do is to run an optimal single-bidder auction for bidder 2 on all the bid vectors not inside or below A_1. This principle holds for many bidders; We state it rigorously for two bidders:

Lemma 3.1. *In the optimal mechanism the following holds.*

$$\alpha_2(\ell) = \underset{\alpha_2(\ell) > u(\ell)}{\operatorname{argmax}} \; \alpha_2(\ell) \cdot \sum_{j=\alpha_2(\ell)}^{n} F(\ell, j)$$

$$u(\ell) = \max \{ u : \ell \leq \alpha_1(u) \}$$

Proof. Let α_1 be fixed. Our aim is to find the α_2 so as to maximise the expected revenue (4) while maintaining the 1-item constraint (1d), which by Papadimitriou and Pierrakos [22] is equivalent to a non-crossing property of critical bid functions (3). Now, clearly the first sum in equation (4) is constant for fixed α_1. So, we are looking to solve

$$R_2 = \max_{\alpha_2} \sum_{\ell=1}^{n} \left[\alpha_2(\ell) \cdot \sum_{j=\alpha_2(\ell)}^{n} F(\ell, j) \right] \tag{5}$$

Furthermore, from (3) it follows that we require

$$\alpha_2(\ell) > \max \{ u : \alpha_1(u) \leq \ell \} =: u(\ell) \tag{6}$$

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

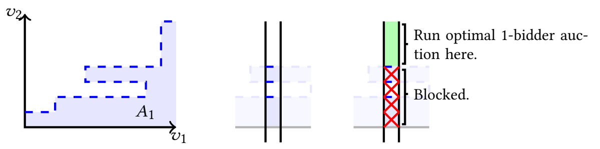

Figure 2: Finding the optimal allocation to bidder 2, given fixed allocation to bidder 1. From left to right: (a) For a given fixed allocation to bidder 1 ...
(b) ... look at each $v_1 = c$-"column" separately ...
(c) ... and run the optimal single-bidder auction for bidder 2, on the part that is not blocked by allocation to buyer 1.

Since this is the only constraint on α_2, it follows that we may interchange the maximum and sum:

$$R_2 = \sum_{\ell=1}^{n}\left[\max_{\alpha_2(\ell)>u(\ell)} \alpha_2(\ell)\cdot\sum_{j=\alpha_2(\ell)}^{n} F(\ell,j)\right] \tag{7}$$

And therefore, $\alpha_2(\ell)$ is as claimed.

$$\alpha_2(\ell) = \operatorname*{argmax}_{\alpha_2(\ell)>u(\ell)} \alpha_2(\ell)\cdot\sum_{j=\alpha_2(\ell)}^{n} F(\ell,j) \tag{8}$$

□

Figure 2 illustrates this lemma. This result tells us that if we already knew one player's allocation & critical bid function, it would be easy to calculate the optimal allocation & critical bids for the second player. In fact, we can calculate all possible ones: Notice that the optimal $\alpha_2(\ell)$ depends only on $u(\ell)$, and no other information on α_1 or α_2. We can iterate through all (n^2) possible values of ℓ and u, and calculate (in linear time each) the optimal $a_2(\ell,u) = \operatorname{argmax}_{m>u} m\cdot\sum_{j=m}^{n} F(\ell,j)$, taking time $O(n^3)$ total.

3.2 Step 2: Disentangling the Remaining Bidder's Allocation

Relying on the previous subsection, we can now extend an argument and algorithm that was first discussed by Gerstgrasser et al. [10]. Let us consider how the optimal choice for each $\alpha_1(j)$ depends on the value of α_1 in other rows. Let us assume that for some j the values of $\alpha_1(j+1),\ldots,\alpha_1(n)$ - i.e. the allocation to bidder 1 in rows above j - are fixed. A particular choice of $\alpha_1(j)$ contributes to the expected revenue of the mechanism in three ways:

- The contribution to expected revenue from the optimal $\alpha_1(1)$, $\ldots$, $\alpha_1(j-1)$, which may depend on $\alpha_1(j)$.
- In *row* j, the expected revenue from selling at points $(\alpha_1(j), j)$, $\ldots$, (n,j) at price $\alpha_1(j)$.
- For some *columns*, the choice of $\alpha_1(j)$ may entail that the mechanism may only sell to buyer 2 at points "above" row j. In particular, consider column $v_1=\ell$. If $\alpha_1(j)\le\ell$, then point $(\ell,j)\in A_1$, and thus (ℓ,j) and all points directly below cannot allocate to buyer 2. On the other hand, this is not necessarily influenced by the choice of $\alpha_1(j)$: If for some $m>j$, (ℓ,m) is in A_1, i.e. $\alpha_1(m)\le\ell$ (which is assumed to be fixed), then (ℓ,m),…, (ℓ,j),…, $(\ell,1)$ could not allocate to bidder 2 irrespective of the choice of $\alpha_1(j)$. This means that this contribution to expected revenue occurs exactly for columns ℓ with $\alpha_1(j)\le\ell<\min_{m>j}\alpha_1(m)$.

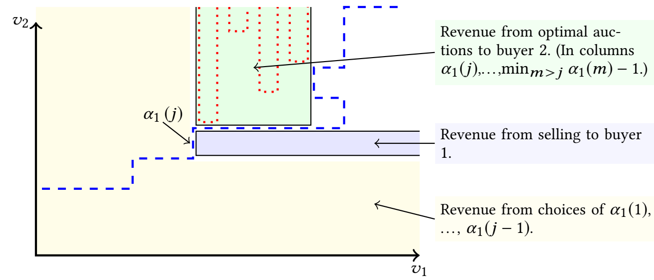

Figure 3: The contribution to the auctioneer's expected revenue due to the choice of $\alpha(j)$ as in Section 3.2. The blue area shows the bid profiles where we sell to buyer 1 due to this choice, the green area shows the bid profiles where we may sell to buyer 2 due to it (using the respective optimal single-bidder auctions, exemplified in red dotted lines).

Figure 3 illustrates these three contributions. The crucial point here is in the last item, which we restate as a lemma due to its importance:

Lemma 3.2. *The optimal choice of $\alpha_1(1),\ldots,\alpha_1(j)$ depends only on the* minimum *of $\alpha_1(j+1),\ldots,\alpha_1(n)$, not all individual values.*

3.3 Step 3: The Dynamic Programming Algorithm

Using these two results, calculating the optimal auction is easy. First, let $r_2(\ell,u) = \max_{m>u}(m\cdot\sum_{j=m}^{n}F(\ell,j))$ be the expected revenue that can be obtained from bidder 2 on bid vectors $(\ell,u+1)$ and those directly above (cf. Lemma 3.1 and discussion after). As discussed in Section 3.1 we can compute these in time $O(n^3)$, which our algorithm does as a first step. Now, by the discussion in Section 3.2 and Lemma 3.2, we only need to consider the minimum of α_1 in the top $n-j$ rows in order to calculate the optimal α_1 in the first j rows. We will write $\alpha_1(j)=n+1$ if none of the points in row j are in A_1. Our algorithm is as follows: for each possible value of $\min\{\alpha_1(2),\ldots,\alpha_1(n)\}$ we calculate the optimal value of $\alpha_1(1)$ and the associated expected revenue arising from this choice. We save these values as $a_1(1,m)$ and $r_1(1,m)$ for $m=1,\ldots,n+1$. Then, we proceed upwards and for $j=2,\ldots,n$ calculate the optimal values of $\alpha_1(j)$ for each possible value of $m=\min\{\alpha_1(j+1),\ldots,\alpha_1(n)\}$ using the already stored values of $r_1(j-1,.)$. In recursion form:

$$a_1(j,m) = \operatorname*{argmax}_{1\le q\le n+1}\left\{q\sum_{\ell=q}^{n}F(\ell,j)+\sum_{\ell=q}^{m-1}r_2(\ell,j)+r_1(j-1,\min\{q,m\})\right\} \tag{9}$$

and $r_1(j,m)=\max_q\{\ldots\}$. The first term is the expected revenue from selling to bidder 1 in this row (blue in Figure 3); the second is the revenue from auctioning to bidder 2 in the appropriate columns (green); the third is the expected revenue from rows below j (yellow). In the third summand we use $\min\{x_j,\ldots,x_n\} = \min\{x_j,\min\{x_{j+1},\ldots,x_n\}\}$. It is easy to see that we do not need to calculate the sums inside the argmax from scratch for each value of q we consider; By memoising the partial sums we can evaluate the term inside the argmax in constant time. That then makes it possible to calculate all the a_1 and r_1 in time $O(n^3)$. $r_1(n,n+1)$ will then give the optimal expected revenue, and α_1 can be found by backtracking from $a_1(n, n+1)$. See Algorithm 1 for the non-memoised version. Furthermore, this extends also to a $\frac{5}{3}$ approximation for many bidders: Dobzinski et al. [8] give this bound for the 2-lookahead auction, which runs the optimal auction for the two highest bidders, with their priors conditioned on the remaining $k-2$ bidders' bids. Thus, any algorithm for the two-bidder auction also can be used to solve the 2-lookahead auction for many bidders, if given access to the conditional prior. The improvement our algorithm gives over existing approaches thus transfers immediately also to approximation algorithms for many bidders.

Algorithm 1 Optimal auction with two correlated bidders.

for $\ell = 1, \dots, n$ **do**
  for $u = 0, \dots, n$ **do**
    $r_2(\ell, u) = \max_{m>u}(m \cdot \sum_{r=m}^{n} F(\ell, r))$
for $m = 1, \dots, n+1$ **do**
  $r_1(0, m) = \sum_{q=1}^{m-1} r_2(0, q)$
for $j = 1, \dots, n$ **do**
  for $m = 1, \dots, n+1$ **do**
    $a_1(j, m) = \text{argmax}_{1 \le q \le n+1}$
    $\left\{q \sum_{\ell=q}^{n} F(\ell, j) + \sum_{\ell=q}^{m-1} r_2(\ell, j) + r_1(j-1, \min\{q, m\})\right\}$
    $r_1(j, m) = \max_{1 \le q \le n+1}$
    $\left\{q \sum_{\ell=q}^{n} F(\ell, j) + \sum_{\ell=q}^{m-1} r_2(\ell, j) + r_1(j-1, \min\{q, m\})\right\}$
  return $r_1(n, n+1)$

Theorem 3.3. *It is possible to calculate the optimal two-bidder auction with correlated priors in time $O(n^3)$, where the prior support is of size n^2. This also gives a $O(n^3)$ $\frac{5}{3}$-approximation algorithm for many bidders.*

Proof. The memoised $O(n^3)$ algorithm follows easily from the conceptual prototype presented in Algorithm 1. We defer the details to the full paper. To verify that the algorithm returns the correct value, consider that we can write the expected revenue as follows:

$$R = \sum_{j=1}^{n}\left[\alpha_1(j) \cdot \sum_{\ell=\alpha_1(j)}^{n} F(\ell, j)\right] + \sum_{\ell=1}^{n}\left[\alpha_2(\ell) \cdot \sum_{j=\alpha_2(\ell)}^{n} F(\ell, j)\right] \quad (10)$$

Here the first term gives the expected revenue from selling to bidder 1, the second term gives the expected revenue from selling to bidder 2: for $v_2 = j$ we sell to bidder 1 at price $\alpha_1(j)$ for bid combinations $\mathbf{v} = (\alpha_1(j), j), \dots, (n, j)$, giving probability $\sum_{\ell=\alpha_1(j)}^{n} F(\ell, j)$.

Now, as a first step we partition the range of the outer sum in the second term along those indices for which $\ell = \alpha_1(j)$ for an $\alpha_1(j)$ with $\alpha_1(j) < \min\{\alpha_1(j+1), \dots, \alpha_1(n)\}$.

$$R = \sum_{j=0}^{n}\left[\alpha_1(j) \sum_{\ell=\alpha_1(j)}^{n} F(\ell, j) + \sum_{\ell=\alpha_1(j)}^{\min\{\alpha_1(j+1),\dots\}-1} \alpha_2(\ell) \cdot \sum_{j=\alpha_2(\ell)}^{n} F(\ell, j)\right] \quad (11)$$

Note that $\sum_{\ell=\alpha_1(j)}^{\min\{\alpha_1(j+1),\dots,\alpha_1(n)\}-1}[\dots]$ is indexing over the empty set and we take it to equal 0, if $\alpha_1(j) \ge \min\{\alpha_1(j+1), \dots, \alpha_1(n)\}$.

We also take $\alpha_1(0) = 1$ and $F(., 0) = 0$. The outer sum now iterates over all the "rows" $v_2 = 1, \dots, n$. Let now $r_2(\ell, u)$ be defined as before; $r_2(\ell, u) = \max_{m>u}(m \cdot \sum_{r=m}^{n} F(\ell, r))$, and $a_2(\ell, u) = \text{argmax}_{m>u}(m \cdot \sum_{r=m}^{n} F(\ell, r))$. By the discussion in Section 3.1 in the optimal auction $\alpha_2(\ell) = a_2(\ell, u(\ell))$ where $u(\ell) = \max\{u : \ell \le \alpha_1(u)\}$ is the topmost point in column j that is in A_1. And similarly $r_2(\ell, u(\ell))$ is precisely the contribution to expected revenue from bidder 2 in column ℓ. It is easy to see that each column ℓ is counted in the outer sum precisely in the summand $j = u(\ell)$. We can therefore rewrite the expected revenue in terms of r_2 instead of α_2.

$$R = \sum_{j=0}^{n}\left[\alpha_1(j) \sum_{\ell=\alpha_1(j)}^{n} F(\ell, j) + \sum_{\ell=\alpha_1(j)}^{\min\{\alpha_1(j+1),\dots,\alpha_1(n)\}-1} r_2(\ell, j)\right] \quad (12)$$

Let now $r_1'(j, m, q)$ be the j-th summand of this, with two free parameters:

$$r_1'(j, m, q) = q \sum_{\ell=q}^{n} F(\ell, j) + \sum_{\ell=q}^{m-1} r_2(\ell, j) \quad (13)$$

Then taking $\min \emptyset = n+1$, we can write the optimal revenue using r_1'.

$$R = \sum_{j=0}^{n}\left[r_1'\Big(j, \min\{\alpha_1(j+1), \dots, \alpha_1(n)\}, \alpha_1(j)\Big)\right] \quad (14)$$

So far we have only reasoned about the revenue given fixed α_1 and α_2, which when introducing r_2 we assumed to be optimal. Clearly for the optimal mechanism it holds, by definition, that that is optimal over all possibilities for α_1:

$$R = \max_{\substack{1 \le a_1(j) \le n+1; \\ a_1(0)=0}} \sum_{j=0}^{n}\left[r_1'\Big(j, \min\{a_1(j+1), \dots, a_1(n)\}, a_1(j)\Big)\right] \quad (15)$$

Since not all of the summands depend on all of the $a_1(j)$, and since they are all non-negative, we can interchange summation and the maximum operators. [2]

$$R = \max_{a_1(n)}\Bigg[r_1'\big(n, \min\emptyset, a_1(n)\big) + \max_{a_1(n-1)}\Big[r_1'\big(n-1, \min\{a_1(n)\}, a_1(n-1)\big) + \max_{a_1(n-2)}\Big[r_1'\big(n-2, \min\{a_1(n), a_1(n-1)\}, a_1(n-2)\big) + \cdots\Big]\Big]\Bigg] \quad (16)$$

[2]More generally, it is easy to check that for $f_i(x_i, \dots, x_n) \ge 0$ the following holds.

$$\max_{\mathbf{x}}\left[\sum_{i=1}^{n} f_i\right] = \max_{x_n}\left[f_n(x_n) + \max_{x_{n-1}}\left[f_{n-1}(x_{n-1}, x_n) + \left[\cdots + \max_{x_1} f_1(\mathbf{x}) \cdots\right]\right]\right]$$

Now we can write R as a recursion in r_1. The following coincides with the definition of r_1 in equation (9):

$$
\begin{aligned}
r_1(j,m) &= \max_q \left[r_1'(j,m,q) + r_1\left(j-1, \min\{q,m\}\right) \right] \\
&= \max_q \left[q \sum_{\ell=q}^{n} F(\ell,j) + \sum_{\ell=q}^{m-1} r_2(\ell,j) + r_1\left(j-1, \min\{q,m\}\right) \right]
\end{aligned} \tag{17}
$$

It is easy to check that (for the optimal value of m), $r_1(j,m)$ is exactly the sum of the first j summands in the expression for the optimal revenue (equations (14), (15)).

$$
r_1\left(j, \min\{\alpha_1(j+1),\ldots,\alpha_1(n)\}\right) = \sum_{s=0}^{j} \left[r_1'\left(s, \min\{\alpha_1(s+1),\ldots,\alpha_1(n)\}, \alpha_1(s)\right) \right] \tag{18}
$$

From this it follows immediately that $R = r_1(n, n+1)$, as desired. Algorithm 1 calculates this value $R = r_1(n, n+1)$ by construction. The algorithm also keeps track of the associated α_1. This can be retraced as $\alpha_1(n) = a_1(n, n+1)$, and $\alpha_1(j) = a_1\left(j, \min\{\alpha_1(j+1),\ldots,\alpha_1(n)\}\right)$. We can calculate α_2 as $\alpha_2(\ell) = a_2(\ell, \max\{j : \alpha_1(j) \le \ell\})$. □

4 THE REVERSE AUCTION CASE

We now turn to the reverse auction case. This is in many settings equivalent to the auction case; However with correlated priors unexpected things happen. Recall that here 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 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 [20]. For a single bidder, not selling is even *all* the power the auctioneer has to achieve any revenue. In a reverse auction, the equivalent 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 $k \ge 3$ bidders, so the reverse auction in these cases is

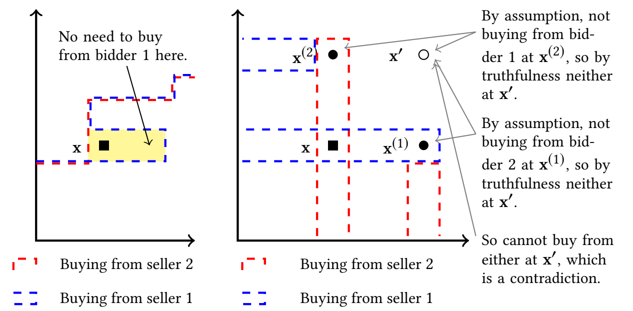

Figure 4: From left to right:
(a) In the first case of Theorem 4.1, if there is no point to the right of x in which we buy only from seller 1, we can improve our cost by not buying from seller 1 in the shaded region, i.e. moving the blue line so it coincides with the red one.
(b) Buying from either seller is blocked at $\mathbf{x}'$ due to truthfulness, in the second case of Theorem 4.1. 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.

clearly structurally different than the corresponding auction; but "No" for 2 bidders. 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 4.1. *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 there exists an i such that for no point $\mathbf{x}' = (x_i', x_{-i})$ with $x_i' > x_i$ we buy only from seller i. Then we could strictly improve our cost if we did not buy from i at $\mathbf{x}$ and all those bid vectors $\mathbf{x}' = (x_i', x_{-i})$ with $x_i' > x_i$. Thus the mechanism was not optimal. See Figure 4 (a) for an illustration.

So assume that for both i, there exists a bid vector $\mathbf{x}^{(\mathbf{i})} = (\mathbf{x}_\mathbf{i}', \mathbf{x}_{-\mathbf{i}})$, with $x_i' > x_i$, so that we buy only from seller i at $x^{(i)}$. 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. Figure 4 (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 $A_2 = A_1^c$. It is easy to check that this allows us to shave a constant multiplicative factor off the runtime of our two-bidder auction algorithm. We list this as Algorithm 2, again in non-memoised form for conciseness. To our knowledge this is the first algorithm specific to the reverse auction setting, exploiting structural arguments of this problem, and therefore also the first to show a lower runtime of this problem compared to the optimal correlated auction.

Algorithm 2 Optimal reverse auction with two correlated bidders.

$c_1(0,.) = 0$
for $j = 1, \ldots, n$ **do**
  for $m = 1, \ldots, n+1$ **do**
    $a_1(j,m) = \text{argmin}_{1 \le q \le m}$
    $\left\{ q \sum_{\ell=1}^{q} F(\ell, j) + j \sum_{\ell=q}^{m-1} \sum_{s=1}^{j} F(\ell, s) + c_1(j-1, q) \right\}$
    $c_1(j,m) = \min_{1 \le q \le m}$
    $\left\{ q \sum_{\ell=1}^{q} F(\ell, j) + j \sum_{\ell=q}^{m-1} \sum_{s=1}^{j} F(\ell, s) + c_1(j-1, q) \right\}$
  return $c_1(n, n+1)$

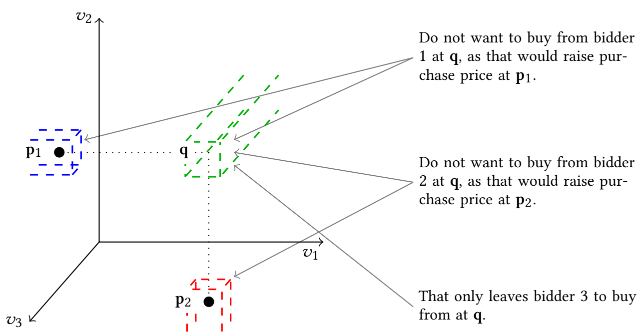

Figure 5: The gadget we will use in the proof of Theorem 4.2. High probability weight on points p_1 and p_2 makes it optimal to buy from seller 3 (in green) at point q. Buying from either of the other sellers at q would raise the purchase price at either p_1 or p_2, thus raising the expected cost. By monotonicity, the mechanism then must also buy from seller 3 at all points behind q in this view.

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

THEOREM 4.2. *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 $\mathbf{p}_2 = (c_M, c_C, c_H)$ with high probability weight, and a third point $\mathbf{q} = (c_M, c_M, c_H)$ with very low probability weight, for some constants $c_L << c_M << c_H$. 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 $\mathbf{p}_2$ & $\mathbf{q}$. 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. Figure 5 illustrates this construction. By monotonicity it follows that if the mechanism buys from seller 3 at $\mathbf{q} = (c_M, c_M, c_H)$, it must also buy from seller 3 at all points (c_M, c_M, v_3), $v_3 \le c_H$.

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 Figure 6. In this we have one gadget consisting of $\mathbf{p}_{12} = (c_H, c_L, c_M)$, $\mathbf{p}_{13} = (c_H, c_M, c_L)$ and $\mathbf{q}_1 = (c_H, c_M, c_M)$ with

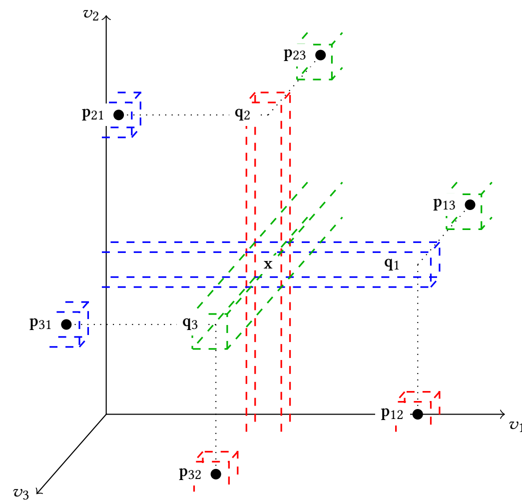

Figure 6: Three gadgets make up the construction used in the proof of Theorem 4.2. Notice that for bid vector x (in the centre at the intersection of the three q_i-segments), the mechanism will buy from all three sellers, due to monotonicity and the allocation at the points q_i.

the auctioneer buying from bidder 1 in the $\mathbf{q}_1$-segment; and similarly one gadget consisting of $\mathbf{p}_{21}$, $\mathbf{p}_{23}$ and $\mathbf{q}_2$ for bidder 2; and a third one 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 $\mathbf{x} = (c_M, c_M, c_M)$. 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 in the main text. To show this 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 is two things to check here; step 1(a), we show that a valid allocation exists that has this property. Step 1(b), we show that any allocation with this property has lower expected cost than any allocation without this property. 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 Figure 7. Each cell lists the bidder(s) that the item is bought from for this given bid vector; with the high probability bid vectors 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

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

Figure 7: The full allocation for the instance in Theorem 4.2. Each cell shows the bidder(s) the mechanism buys from at the given bid vector. High probability points are shown in bold face.

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.

□

5 DISCUSSION AND FUTURE WORK

Our $O(n^3)$ algorithm for the two-bidder auction is interesting not only in its own right - presenting a substantial improvement of the previously known $O(n^6)$ approaches, both for the exact solution for two bidders but also for a $\frac{5}{3}$-approximation for many bidders; but it is significant also because of the techniques used in arriving at it. Our algorithm is the first to exploit structural properties of the 2-bidder auction design problem, whereas previous approaches reduce the problem to a generic graph algorithm respectively a derandomised LP. In contrast, structural insights allow us to decompose the problem into a simple recursion. Generalised versions of these observations hold for many settings, and could potentially be useful there.

Lemma 3.1 gives a useful characterisation of the optimal two-bidder auction, and indeed this generalises to virtually all correlated mechanism design settings we can think of. For two-bidder settings, this gives a useful decomposition into optimal single-bidder mechanisms, that may be tractable in many cases. For more general cases, an analogue of Lemma 3.1 states — informally — that in an optimal mechanism, for any partition of bidders into two sets S_1, S_2, for a fixed bid vector $\mathbf{v}_1$ of bidders in S_1, the allocation to bidders in S_2 is simply that of the optimal $|S_2|$-bidder mechanism taking into account the (fixed) allocations to bidders in S_1. This may be a useful aid in exploring multi-bidder settings.

It is interesting to consider why this generalisation cannot be used to show a polynomial-time algorithm for the three-bidder auction. For this setting, we get two characterisations depending on whether we partition into $|S_1| = 2$, $|S_2| = 1$, or vice versa. If we take $S_1 = \{b_1, b_2\}$, $S_2 = \{b_3\}$, we get that for any fixed v_1, v_2, the allocation to bidder 3 in the induced affine linear v_3-subspace (i.e. in each "vertical line" in the bid space) is that of the optimal single-bidder auction to bidder 3 on all bid vectors that are above any allocated to bidders 1 and 2 (i.e. on those not blocked by the allocation to bidders 1 and 2). Vice versa, if we take $S_1 = \{b_3\}$, $S_2 = \{b_1, b_2\}$, we get that for fixed v_3, the allocation to bidders 1 and 2 on the induced v_1, v_2-subspace (each horizontal plane in the bid space) is that of the optimal two-bidder mechanism on all bid vectors not blocked by the allocation to bidder 3. While these are interesting properties, what is missing is an analogue of Lemma 3.2, which allows us to rewrite the optimal revenue as a tractable recursion. In the three-bidder case, we do not have an efficient way of computing the optimal allocation to two or even one bidder. Even in such settings where this result does not lead to an efficient algorithm, it may still be useful in characterising the optimal mechanism. One open question that is of particular interest here is whether Lemma 3.1 and its generalisation also hold as an "if and only if" statement. That is, it would be nice if they would give sufficient conditions for a mechanism to be optimal.

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, and to devise the first algorithm specific to the correlated reverse auction setting. 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. We take this as evidence that the reverse auction case is interesting to consider as a separate problem from the standard auction model.

ACKNOWLEDGEMENTS

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

REFERENCES

[1] Ioannis Caragiannis, Christos Kaklamanis, and Maria Kyropoulou. 2016. Limitations of deterministic auction design for correlated bidders. *ACM Transactions on Computation Theory (TOCT)* 8, 4 (2016), 13.

[2] Xue Chen, Guangda Hu, Pinyan Lu, and Lei Wang. 2011. On the Approximation Ratio of k-Lookahead Auction. In *International Workshop on Internet and Network Economics*. Springer, 61–71.

[3] Vincent Conitzer and Tuomas Sandholm. 2006. 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, 521–528.

[4] Florin Constantin, Takayuki Ito, and David C Parkes. 2007. Online auctions for bidders with interdependent values. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*. ACM, 110.

[5] Jacques Crémer and Richard P. McLean. 1985. Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent. *Econometrica* 53, 2 (1985), 345–361. http://www.jstor.org/stable/1911240

[6] Jacques Crémer and Richard P. McLean. 1988. Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions. *Econometrica* 56, 6 (1988), 1247–1257. http://www.jstor.org/stable/1913096

[7] Ilias Diakonikolas, Christos Papadimitriou, George Pierrakos, and Yaron Singer. 2012. Efficiency-revenue trade-offs in auctions. In *International Colloquium on Automata, Languages, and Programming*. Springer, 488–499.

[8] Shahar Dobzinski, Hu Fu, and Robert D Kleinberg. 2011. Optimal auctions with correlated bidders are easy. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*. ACM, 129–138.

[9] Péter Eső. 2005. An optimal auction with correlated values and risk aversion. *Journal of Economic Theory* 125, 1 (2005), 78–89.

[10] Matthias Gerstgrasser, Paul W Goldberg, and Elias Koutsoupias. 2016. Revenue Maximization for Market Intermediation with Correlated Priors. In *International Symposium on Algorithmic Game Theory*. Springer, 273–285.

[11] Jason Hartline and Anna Karlin. 2007. Profit maximization in mechanism design. In *Algorithmic Game Theory*, Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani (Eds.). Cambridge University Press, 331.

[12] Paul Klemperer. 1999. Auction theory: A guide to the literature. *Journal of economic surveys* 13, 3 (1999), 227–286.

[13] P. Klemperer. 2004. *Auctions: Theory and Practice*. Princeton University Press. https://books.google.co.uk/books?id=E56FQgAACAAJ

[14] Vijay Krishna. 2009. *Auction theory*. Academic press.

[15] Dan Levin and James L Smith. 1996. Optimal reservation prices in auctions. *The Economic Journal* (1996), 1271–1283.

[16] Alejandro M Manelli and Daniel R Vincent. 1995. Optimal procurement mechanisms. *Econometrica: Journal of the Econometric Society* (1995), 591–620.

[17] Paul R Milgrom and Robert J Weber. 1982. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society* (1982), 1089–1122.

[18] Hadi Minooei and Chaitanya Swamy. 2013. Near-Optimal and Robust Mechanism Design for Covering Problems with Correlated Players. In *International Conference on Web and Internet Economics*. Springer, 377–390.

[19] Hadi Minooei and Chaitanya Swamy. 2015. Near-optimal and robust mechanism design for covering problems with correlated players. *ACM Transactions on Economics and Computation* 3, 4 (2015), 19.

[20] Roger B Myerson. 1981. Optimal auction design. *Mathematics of operations research* 6, 1 (1981), 58–73.

[21] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. 2007. *Algorithmic game theory*. Cambridge University Press.

[22] Christos H Papadimitriou and George Pierrakos. 2011. On optimal single-item auctions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*. ACM, 119–128.

[23] Amir Ronen. 2001. On approximating optimal auctions. In *Proceedings of the 3rd ACM conference on Electronic Commerce*. ACM, 11–17.

[24] Amir Ronen and Amin Saberi. 2002. On the hardness of optimal auctions. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*. IEEE, 396–405.