

Revenue Maximization for Market Intermediation with Correlated Priors

Matthias Gerstgrasser, Paul W Goldberg, and Elias Koutsoupias

Department of Computer Science, University of Oxford,
Wolfson Building, Parks Road,
OXFORD, OX1 3QD, UK
{matthias.gerstgrasser,paul.goldberg,elias.koutsoupias}@cs.ox.ac.uk
<http://www.cs.ox.ac.uk/>

Abstract. We study the computational challenge faced by a intermediary who attempts to profit from trade with a small number of buyers and sellers of some item. In the version of the problem that we study, the number of buyers and sellers is constant, but their joint distribution of the item’s value may be complicated. We consider discretized distributions, where the complexity parameter is the support size, or number of different prices that may occur. We show that maximizing the expected revenue is computationally tractable (via an LP) if we are allowed to use randomized mechanisms. For the deterministic case, we show how an optimal mechanism can be efficiently computed for the one-seller/one-buyer case, but give a contrasting NP-completeness result for the one-seller/two-buyer case.

1 Introduction

We consider a double auction scenario from the perspective of a market intermediary, collecting bids from one or more sellers and buyers and determining payments and allocations. Real-world instances of this are manifold, including in electronic markets. Companies such as eBay or Amazon match sellers and buyers, and charge a fee for each successful transaction. Our aim is to maximize the intermediary’s profit in such settings.

There is an extensive literature on this challenge, some of which is discussed below, but it mostly considers the case of many buyers and/or sellers with independent priors. Our interest here is different, in that we assume only a constant number of buyers and sellers (in the simplest version, just one of each), and the complexity arises from their joint probability distribution of valuations for the item. In the simplest version of this, where there is just one buyer and one seller, the intermediary can profit from buying the item from the seller and selling at a higher price to the buyer. We assume their valuations for the item come from a known joint distribution, which is the input to the problem. We consider two versions: the “no short selling” version, with the natural constraint that we cannot sell more items than we buy; and the more restrictive “balanced inventory” version, where in addition we must sell all the items we buy. For multiple buyers

and sellers, we assume that the intermediary can buy and sell multiple items (but that each buyer/seller has unit demand/supply).

1.1 Related work

The problem of optimal mechanisms in a market intermediation setting was first studied by Myerson and Satterthwaite [11]. In addition to an impossibility result for ex-post efficiency in a bilateral trade setting without an intermediary, they show optimal intermediation mechanisms for both social welfare as well as the intermediary’s revenue in the case of one buyer and one seller, whose valuations are independent. Their revenue-maximization result is similar to Myerson’s seminal single auction result [10] in that it, too, uses virtual valuation functions, for both buyer and seller. Welfare maximization for multiple buyers and sellers has been further studied for instance by McAfee [9] or more recently by [4,5,1].

Our own interest is chiefly in the complexity of computing optimal (revenue-maximizing) or near-optimal mechanisms in the market intermediation setting. Prior work in this area has focused on the case where sellers’ and buyers’ valuations are independent. Deng et al. [2] show optimal and near-optimal mechanisms that can be computed in polynomial time for several variations of this setting, including continuous or discrete distributions and arbitrary or unlimited supply and demand. Niazadeh et al. [12] as well as Loertscher and Niedermayer [7,8,6] study a class of mechanisms called respectively *fee-setting mechanisms* or *affine fee schedules* in the independent setting. These are shown by Niazadeh et al. [12] to be able to extract a constant factor of the optimum revenue in the worst case, under certain assumptions on the buyer’s and seller’s distribution.

Here we are interested in potentially correlated distribution over buyers’ and sellers’ valuations. The complexity of this has been studied for (non-double) auctions. Papadimitriou and Pierrakos [13] show that for two buyers, an optimal mechanism (for a discrete joint distribution) can be found in polynomial time via a reduction to finding a maximum-weight independent set on a bipartite graph. For continuous distributions they give a FPTAS. For the case of three buyers, in contrast, they show that it is NP-hard to approximate the optimal auction to within a certain constant fraction. Dobzinski et al. [3] show a polynomial-time algorithm for the two-buyer auction through derandomization and give polynomial-time approximation mechanisms for the many-buyers correlated single auction problem, building on previous work by Ronen [14].

2 Preliminaries

2.1 Definitions, notation

We consider m buyers indexed by j , and k sellers indexed by i (where m, k are constants), each offering (respectively seeking) a single unit of an indivisible good. For fixed m, k , we use “ $m \times k$ ” as shorthand for the m buyers, k sellers case. They cannot trade with each other directly, and can only trade with the

intermediary. We assume each seller i has some valuation s_i and each buyer j has some valuation b_j for an item, and that these are drawn from a given joint probability distribution ψ . We focus on discrete distributions. For simplicity, we assume that the support of ψ is a grid of size n^{k+m} , with each player having possible valuations $\{i : 1 \leq i \leq n\}$ (as shown in Figure 1). The distribution ψ is assumed to be represented as a matrix of the probabilities on each of the grid points.

In this paper we focus on individually rational and incentive compatible mechanisms. Let \mathbf{b} and \mathbf{s} denote the vector of buyers' and sellers' bids received by the mechanism, and let \mathbf{b}_{-j} and \mathbf{s}_{-i} be (respectively) the bids of buyers other than j , and sellers other than i . Analogous to auctions, there are two equivalent ways in which we can define a deterministic, incentive compatible mechanism in this setting. Firstly, we can focus on allocations. For each seller / buyer we define a set $S_i / B_j (\subseteq \text{supp}(\psi))$ of bid vectors in which we buy an item from seller i / sell an item to buyer j . Incentive compatibility means monotonicity of allocations, meaning for each seller i , if $(s, b) \in S_i$, and $s'_i < s_i$, then $(s'_i, s_{-i}, b) \in S_i$. In words, if given everyone else's bids s_{-i}, b , seller i 's item would be bought by the intermediary if i bid s_i , then it would also be bought for any lower bid s'_i . For short we will say that S_i is "downward-closed" in the direction of s_i . Similarly for each buyer j , B_j needs to be upward-closed in the direction of b_j .

Equivalently, we may think of a mechanism in terms of critical bids. Myerson tells us that the unique payments that make a monotone allocation rule (as just defined via the regions S_i and B_j) are precisely the critical bids. That is, the lowest (highest) bid for which a buyer (seller) would still be allocated the item (the sale of their item) if everyone else's bids remained fixed. We write $\sigma_i(\mathbf{b}, \mathbf{s}_{-i})$ (respectively, $\beta_j(\mathbf{b}_{-j}, \mathbf{s})$) for these critical bids. (For simplicity sometimes just $\beta_j(\mathbf{b}, \mathbf{s})$ and $\sigma_i(\mathbf{b}, \mathbf{s})$.) If $s_i \leq \sigma_i(\mathbf{b}, \mathbf{s}_{-i})$ we buy an item from seller i (paying $\sigma_i(\mathbf{b}, \mathbf{s}_{-i})$), and similarly if $b_j \geq \beta_j(\mathbf{b}_{-j}, \mathbf{s})$ we sell an item to buyer j (charging $\beta_j(\mathbf{b}_{-j}, \mathbf{s})$). We write $\beta_j(\mathbf{b}_{-j}, \mathbf{s}) = n + 1$ to indicate that a mechanism does not sell to buyer j at all for this combination of others' bids, independently of j 's bid. Similarly $\sigma_i(\mathbf{b}, \mathbf{s}_{-i}) = 0$ to indicate not buying from seller i .

It is easy to see that these two yield equivalent definitions. Clearly S_i is simply the region "above" σ_i in the direction of s_i , (the graph of) which in turn is the boundary of S_i . Similarly B_j is the region below β_j in direction b_j . This is a slight generalization of the conceptually simpler picture in auctions. Here we have for each bidder a region B_j where they win the item, and a critical bid function β_j that gives their payment. If there is a single item to be sold, no two of the B_j may overlap. This constraint too generalizes to the market intermediation setting. Consider figure 2 in contrast with figure 1 to illustrate the difference. As mentioned above, we consider two variants. In the "no short-selling" setting, we must buy at least as many items from sellers as we sell to buyers; in the "balanced inventory" variant we must buy exactly as many as we sell. Formally in terms of critical bids: (Again these can be expressed equivalently in terms of

S_i and B_j .)

$$\begin{aligned} \text{No Short-Selling:} \\ \forall(\mathbf{b}, \mathbf{s}), |\{j : b_j \geq \beta_j(\mathbf{b}, \mathbf{s})\}| \leq |\{i : s_i \leq \sigma_i(\mathbf{b}, \mathbf{s})\}| \end{aligned} \quad (1a)$$

$$\begin{aligned} \text{Balanced Inventory:} \\ \forall(\mathbf{b}, \mathbf{s}), |\{j : b_j \geq \beta_j(\mathbf{b}, \mathbf{s})\}| = |\{i : s_i \leq \sigma_i(\mathbf{b}, \mathbf{s})\}| \end{aligned} \quad (1b)$$

2.2 The geometry of deterministic 1×1 market intermediation

If there is only a single buyer and a single seller, the constraints simplify significantly, most easily expressed in terms of now simply S and B . In the balanced-inventory case, constraint 1b simplifies to $B = S$. That is, a mechanism in this setting, with only one buyer and seller each, is determined only by a single region of bid-combinations that yield a successful transaction. In the no-short-selling case, constraint 1a simplifies to $B \subseteq S$. That is, S can potentially extend beyond B . However, we can say more, assuming optimality of the mechanism. Recall that by truthfulness, S is down-closed in the seller's direction and B is up-closed in the buyer's direction. If a mechanism is optimal, S must exactly be the down-closure (still in the seller's direction) of B . Firstly, it is easy to see that the down-closure of B must be contained in S : B is contained in S , and S is down-closed. Secondly, if S extended beyond the down-closure of B , we could strictly improve our revenue by removing this protruding part of S . (On the other hand, we may not elect to remove the part of $S - B$ that lies below any point in B due to truthfulness, i.e. down-closedness of S .) Figures 1 and 3 illustrate this. Note the contrast with a standard 2-bidder auction, where the shape of the region in which we sell to one buyer does not fully determine the region in which we sell to the other. In a way, in the 1×1 market intermediation setting, we have fewer degrees of freedom to consider than in a two-buyer auction setting.

3 The deterministic one seller, one buyer case

For the case of one seller and one buyer, we show how to compute an optimal deterministic solution using a dynamic programming approach. In Appendix A we show how to achieve this via modifications to known 2-bidder auctions in this setting, but the runtime guarantee of that approach, while still polynomial, is substantially worse. We represent a mechanism using values $\beta = (\beta(1), \dots, \beta(n))$ for $\beta(s) \in \{1, \dots, n+1\}$, where $\beta(s)$ signifies the leftmost point in row s that is a member of B . If $\beta(s) = 1$ then the entire row s is in B . We set $\beta(s) = n+1$ to signify that none of the points in row s are in B . We begin by noting that the contribution to the expected revenue that arises from the choice of one particular $\beta(s)$ does not depend on all the other β_{-s} simultaneously. Consider the expected revenue R for a given B , which is given by:

$$\mathbb{E}[R(\beta)] = \sum_{s=1}^n \beta(s) \sum_{b \geq \beta(s)} \psi_{bs} - \sum_{b=1}^n \sigma(b) \sum_{s \leq \sigma(b)} \psi_{bs} \quad (2)$$




Fig. 1: Example mechanism in the 1×1 case. Note that S contains B , to avoid short-selling. (In a balanced-inventory auction, B and S should coincide.) B and S lie below the diagonal $b = s$: any point above the diagonal is one where the buyer's bid is less than the seller's. The auction shown is sub-optimal: in most of the S region, the item is being bought without being sold. Note that we draw the outline of the regions slightly away from the points on the prior support for easier readability.




Fig. 2: Compare this to a two-bidder auction. Here B_1 and B_2 indicate where we sell to each of the two buyers. In the two-bidder auction B_1 and B_2 must be disjoint, as we cannot sell the item twice. In the market intermediation setting, B must be contained in S . Note also that in this setting both B_1 and B_2 are upward-closed in the direction of the respective buyer's bid.




Fig. 3: Removing the indicated area from S as in figure 1, the expected revenue of the mechanism cannot decrease. Below red line: Remaining region S , right of black line: Region B . For the remaining part of S that is not also in B , we still buy but not sell the item. This can be optimal, e.g. if there is very high probability weight on the two points indicated. Crucially, if at a point $(b, s) \in S - B$ an optimal mechanism buys but not sells, then there must exist a point $(b, s') \in S \cap B$ with $s' > s$ where it buys and sells. Truthfulness then dictates that it also needs to buy at (b, s) .

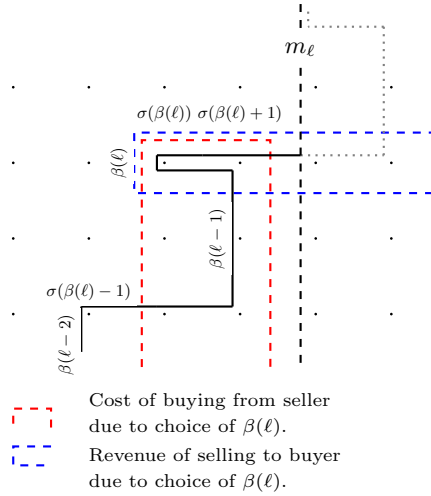


Fig. 4: For fixed m_ℓ and $\beta(1), \dots, \beta(\ell - 1)$, the choice of a particular $\beta(\ell)$ influences the expected revenue in two ways: On the one hand, the revenue from selling to the buyer at all points to the right of $\beta(\ell)$ in row ℓ . On the other hand, the cost of buying from the seller for all points below $\sigma(\beta(\ell)) = \ell$ in rows $\beta(\ell) \leq b < m_\ell$.

That is, the first sum gives the expected profit from selling at points to the right of each $(\beta(s), s)$, while the second sum gives the cost of buying points that are below each $(b, \sigma(b))$. The contribution of a particular choice for one single $\beta(s)$ to the first of these sums is easily seen to be simply the profit of selling the points in row s to the right of and including $\beta(s)$. The impact of a particular $\beta(s)$ on the second of the sums is slightly more intricate. There are two ways in which the choice of $\beta(s)$ impacts the cost of buying. Firstly, we may have to buy the item at some points in row s , where we would not buy the item otherwise. Consider the minimum of $\beta(s+1), \dots, \beta(n)$, say $\beta(t)$. We know that in row t , we buy and sell at points $(\beta(t), t), \dots, (n, t)$. So by truthfulness, we must also buy at points $(\beta(t), s), \dots, (n, s)$. This is regardless of our choice of $\beta(s)$. For points to the left of $(\beta(t), s)$, whether we buy the item or not does depend on $\beta(s)$. Secondly, in all those columns in which we buy due to $\beta(s)$, also affect the rows below s . We may increase the buying price from a lower value to s at those points, and (in the no short-selling case) we may have to buy the item (due to truthfulness) at points at which we would not otherwise buy it. The magnitude of this effect depends on all the $\beta(1), \dots, \beta(s-1)$. This suggests a bottom-up dynamic programming approach, which we develop in this section.

3.1 Algorithm for the no short-selling case

We next describe our dynamic programming algorithm first for the no short-selling setting. The idea is as follows: because the optimal choice of $\beta(1), \dots, \beta(\ell)$ depends only on the minimum of the $\beta(\ell+1), \dots, \beta(n)$, we can iteratively compute all the potential optimal values for row 1 given values of $\min\{\beta(2), \dots, \beta(n)\}$; then all optimal values of $\beta(1), \beta(2)$ given all possible values of $\min\{\beta(3), \dots, \beta(n)\}$. We do not need to consider all n^2 combinations of $\beta(2)$ and $\beta(1)$. Since given $\beta(2)$ and $\min\{\beta(3), \dots, \beta(n)\}$, we can immediately look up the best $\beta(1)$ using the information computed in the first step. We then proceed iteratively up the rows until we have computed to optimal values for β .

Let us start by defining $R(n, \beta(\ell), m_\ell)$ to be the expected revenue of the best deterministic mechanism that takes points $(\beta(\ell), \ell)$ and rightward in row ℓ , no points in rows $\ell+1$ and above, and does not have to pay for points in columns m_ℓ to n . We set $R(0, \dots) = 0$. The idea is that we want to capture the best possible revenue extractable from rows 1 to ℓ for a particular choice of $\beta(\ell)$, disregarding the cost of buying in columns m_ℓ to n . We can take $R(n, \beta(\ell), n+1)$ to denote the optimal revenue among mechanisms that have to pay for all rows. More precisely,

$$R(\ell, \beta(\ell), m_\ell) = \max_{\beta(1), \dots, \beta(\ell-1)} \sum_{s=1}^{\ell} \beta(s) \sum_{b \geq \beta(s)} \psi_{bs} - \sum_{b=1}^{m_\ell-1} \sigma(b) \sum_{s \leq \sigma(b)} \psi_{bs} \quad (3)$$

It is easy to see that $\max_{\beta(n)} R(n, \beta(n), n+1)$ gives the revenue of the optimal auction. Indeed, by definition this is the maximum expected revenue extractable from all rows, if we have to pay in all columns. We can then show how to

recursively compute the values of R , laying the groundwork for our dynamic programming algorithm.

Theorem 1 (Recursion for the no short-selling case). *The $R(\ell, \beta(\ell), m_\ell)$ as defined above satisfy the following recursion:*

$$R(\ell, \beta(\ell), m_\ell) = \max_{\beta(\ell-1)} R(\ell-1, \beta(\ell-1), \min\{\beta(\ell), m_\ell\}) + \beta(\ell) \sum_{b \geq \beta(\ell)} \psi_{b\ell} - \ell \sum_{\beta(\ell) \leq b < m_\ell} \sum_{s \leq \ell} \psi_{bs} \quad (4)$$

Proof. We can check this by splitting up the explicit formula for $R(\ell, \beta(\ell), m_\ell)$ into terms for rows below ℓ and row ℓ , and columns to the left of $\min(\beta(\ell), m_\ell)$ and those between the two.

$$R(\ell, \beta(\ell), m_\ell) = \sum_{s=1}^{\ell-1} \beta(s) \sum_{b \geq \beta(s)} \psi_{bs} + \beta(\ell) \sum_{b \geq \beta(\ell)} \psi_{b\ell} - \sum_{b=1}^{\min(\beta(\ell), m_\ell)-1} \sigma(b) \sum_{s \leq \sigma(b)} \psi_{bs} - \sum_{b=\beta(\ell)}^{m_\ell-1} \sigma(b) \sum_{s \leq \sigma(b)} \psi_{bs}$$

Observe that for $b \geq \beta(\ell)$, $\sigma(b)$ will be equal to ℓ (in the $(\ell, \beta(\ell))$ -auction), so the last term in the above sum is precisely $\ell \sum_{b=\beta(\ell)}^{m_\ell-1} \sum_{s \leq \ell} \psi_{bs}$. Similarly, $\min(\beta(\ell), m_\ell)$ is precisely the $m_{\ell-1}$ we used in the recursion, and therefore the first and third term are precisely $R(\ell-1, \beta(\ell-1), m_{\ell-1})$. Putting these together, we get that:

$$R(\ell, \beta(\ell), m_\ell) = R(\ell-1, \beta(\ell-1), m_{\ell-1}) + \beta(\ell) \sum_{b \geq \beta(\ell)} \psi_{b\ell} - \ell \sum_{b=\beta(\ell)}^{m_\ell-1} \sum_{s \leq \ell} \psi_{bs} \quad (5)$$

i.e. precisely our claimed recursion. (The max follows from optimality of the auction.) The second term on the right hand side is the revenue from selling at points due to the choice of $\beta(\ell)$, while the third term accounts for the cost of buying at points due to this choice. Figure 4 illustrates these two terms. Note that if $\ell = 1$ then the first term vanishes since we defined $R(0, \dots) = 0$, and we are left with the explicit formula for $R(1, \dots)$. \square

We can therefore compute the $R(\ell, \beta(\ell), m_\ell)$ recursively, as claimed. This suggests the following algorithm, listed below as Algorithm 1. This can easily be augmented to keep track of the values used for the $\beta(s)$, and to return the optimal β together with its expected revenue. Therefore, we can compute the optimal region B and thereby the optimal mechanism in the no short-selling setting in time $\mathcal{O}(n^4)$.¹

¹ Careful analysis of the algorithm presented shows that the last summand in the recursion for $R()$ has $(m_\ell - \beta(\ell)) \cdot \ell$ summands. It is easy to see however that we need not recompute the inner sum from scratch in each iteration. We can thus easily make the computation of the recursion run in linear time, giving the overall running time stated.

Algorithm 1 Optimal revenue in the no short-selling setting

```

1: for  $\ell = 1, \dots, n$  do
2:   for  $\beta(\ell) = 1, \dots, n$  do
3:     for  $m_\ell = 1, \dots, n + 1$  do
4:       if  $\ell = 1$  then
5:          $R(1, \beta(1), m_1) \leftarrow \beta(1) \sum_{b \geq \beta(1)} \psi_{b1} - \sum_{\beta(1) \leq b < m_1} \psi_{b1}$ 
6:       else
7:         Compute  $R(\ell, \beta(\ell), m_\ell)$  using the recursion in theorem 1.
8:   return  $\max_{\beta(n)} R(n, \beta(n), n + 1)$ .
```

We can easily modify this algorithm to return the optimal mechanism that satisfies the balanced inventory property. We show the details in Appendix B. This modified algorithm runs in time $\mathcal{O}(n^3)$.

4 NP-hardness for the deterministic multiple buyers or sellers case

For three or more buyers, it follows from Papadimitriou and Pierrakos [13] that computing the optimal mechanism is NP-hard. We show that this is also true for the 2×1 case (i.e. two buyers, one seller) in the no-short-selling setting by reducing from Maximum Independent Set. The idea here is to place high probability weight on high-revenue points along a diagonal in the $s = 1$ plane for each vertex of a given instance of Independent Set. We then use appropriately placed high-probability points for each of the edges to “force” a higher buying price for (at least) one of any two points corresponding to adjacent vertices. We can do this in a way that ensures that in the optimal mechanism the number of vertex points with a low buying price is maximized and corresponds to the maximum independent set.

Theorem 2 (NP-hardness). *It is NP-hard to compute the optimal mechanism in the 1 seller, 2 buyers setting with no short selling.*

Proof. In the following we construct a prior distribution in (b_1, b_2, s) -space. We will “choose” points, and place equal probability mass $\frac{1}{|V|+|E|}$ on all of these chosen points. In the analysis we will omit these weights to simplify the algebra. We place probability 0 on all other points in the prior support. We use K_1 and K_2 as constants whose values we define at the end of the proof.

The construction Given a graph (V, E) with $|V| = n$, pick any order of vertices and begin by placing probability weight $\frac{1}{|V|+|E|}$ on point $(K_1 + \lfloor \frac{n}{2} \rfloor - i, K_1 - \lfloor \frac{n}{2} \rfloor + i, 1)$ for each vertex $0 \leq i < n$. Next, enumerate the edges e_j , $0 \leq j < |E|$. We will write each edge as $e_j = (e_{j1}, e_{j2})$, where $e_{j1} < e_{j2}$ in the order of vertices just picked. For each edge put probability weight $\frac{1}{|V|+|E|}$ on point $(K_1 + \lfloor \frac{n}{2} \rfloor - e_{j2}, K_1 - \lfloor \frac{n}{2} \rfloor + e_{j1}, K_2 + j)$. That is, we put probability weight

for each edge on a point that has the same b_1 -coordinate as the vertex point for its lower-numbered vertex and the same b_2 -coordinate as its higher vertex. We choose these edge points with a different s -coordinate each, and all of them with a higher s -coordinate than the vertex points. It is clear that if the mechanism wants to buy and sell at an edge point $(K_1 + \lfloor \frac{n}{2} \rfloor - e_{j2}, K_1 - \lfloor \frac{n}{2} \rfloor + e_{j1}, K_2 + j)$, it will also need to sell (and therefore buy by truthfulness) at one of the points $(K_1 + \lfloor \frac{n}{2} \rfloor - e_{j2}, K_1 - \lfloor \frac{n}{2} \rfloor + e_{j2}, K_2 + j)$ or $(K_1 + \lfloor \frac{n}{2} \rfloor - e_{j1}, K_1 - \lfloor \frac{n}{2} \rfloor + e_{j1}, K_2 + j)$, when it sells to buyer 1 or buyer 2, respectively. But by truthfulness this entails a raised purchase price of $K_2 + j$ at the corresponding vertex points directly below $((K_1 + \lfloor \frac{n}{2} \rfloor - e_{j2}, K_1 - \lfloor \frac{n}{2} \rfloor + e_{j2}, 1)$ or $(K_1 + \lfloor \frac{n}{2} \rfloor - e_{j1}, K_1 - \lfloor \frac{n}{2} \rfloor + e_{j1}, 1)$) where it had otherwise been 1. Figure 5 illustrates this construction.

Reducing from Maximum Independent Set Now, in order to ensure that the optimal mechanism raises the purchasing price at all vertex points except those that are in an independent set of maximum size, we need to pick constants K_1, K_2 in a way that ensures that:

1. The optimal mechanism always buys and sells at all the edge points.
2. The optimal mechanism raises the purchasing price at as few vertex points as possible.

From condition 1: The worst possible selling price at any edge point is given by $K_1 - \lfloor \frac{n}{2} \rfloor$, and the highest possible purchase price is $K_2 + |E| \leq K_2 + n^2$, for a revenue that is at least $K_1 - \lfloor \frac{n}{2} \rfloor - K_2 - n^2$. On the other hand, buying and selling at an edge point could necessitate a higher purchasing price at a vertex point, raising it by an amount that is bounded above by $K_2 + n^2$ as well. The profit obtained from the edge point must outweigh this. So in order to ensure that the optimal mechanisms buys and sells at all edge points, we need to ensure:

$$K_1 - \lfloor \frac{n}{2} \rfloor - 2K_2 - 2n^2 > 0$$

From condition 2: We need to ensure that if for an edge point $(K_1 + \lfloor \frac{n}{2} \rfloor - e_{j2}, K_1 - \lfloor \frac{n}{2} \rfloor + e_{j1}, K_2 + j)$, only one of the two corresponding vertex points already has a purchase price of at least K_2 due to another edge, but the other is still 1, the optimal mechanism will always prefer to sell to the buyer whose corresponding vertex point already has a high price. In other words, we need to ensure that the potential difference in revenue from selling to one buyer over the other is outweighed by the required raise in the purchase price by (at least) $K_2 - 1$. But the highest difference in selling price is bounded by n , and so we required that $K_2 > n + 1$.

Combining the two we get our desired result: Set $K_2 = 2n$ and $K_1 = 4n^2$ in the above construction for a given instance of Maximum Independent Set. Since the optimal mechanism will buy and sell at all the edge points, it is clear that at most one vertex point corresponding to two adjacent vertices can have a purchase price of 1. On the other hand, in the optimal mechanism the number of vertex points with a raised purchase price will be minimized. Therefore, the

vertex points with purchase price 1 in the optimal mechanism correspond to the vertices of the maximum independent set in the graph. \square




Fig. 5: The construction for the reduction from Maximum Independent Set.

5 Truthful-in-expectation mechanisms

While in the preceding section we have shown that we cannot compute an optimal deterministic mechanism for the general case, we can however compute the optimal truthful-in-expectation mechanism for a fixed number of buyers and sellers. In single-item auctions a randomized mechanism is easily described by allocation probabilities $x_i(\mathbf{v})$ and expected payments $p_i(\mathbf{v})$ for all players for each possible bid vector. In the market intermediation setting with multiple buyers and sellers this is not obviously the case. For instance, there are many ways in which to make allocation probabilities of $\frac{1}{2}$ for each of two buyers and two sellers into a randomization over valid outcomes. The mechanism could flip a coin and buy from seller 1 and sell to buyer 1 on heads, seller 2 and buyer 2 on tails. It could not, however, independently flip four coins if we want to fulfill condition (1a) respectively (1b) ex-post. In the following we will consider the balanced inventory case. Our arguments easily extend to the no short-selling case. We first show that for our purposes, it is indeed sufficient to consider only the marginal allocation probabilities x_i, y_j and expected payments p_i, q_j . First, observe that any two randomized mechanisms that have the same marginal probabilities and expected payments will lead to identical expected utilities for players and expected revenue. It remains to show that any sensible vector of marginal allocation probabilities can be made into a probability distribution over valid outcomes (i.e. allocations which buy exactly as many items as they sell).

Theorem 3. *Let x, y be k -dimensional vectors of probabilities, i.e. $0 \leq x_i, y_i \leq 1$, with $\sum_i x_i = \sum_i y_i$. Then there exists a joint probability distribution over 2 k -dimensional 0/1 vectors $\{(a, b) \in \{0, 1\}^{2k} \mid \sum_i a_i = \sum_i b_i\}$ which satisfies $\Pr(a_i = 1) = x_i$ and $\Pr(b_i = 1) = y_i$.*

Proof. Let $H_{2k} = [0, 1]^{2k}$ be the $2k$ -dimensional hypercube, and $H_{2k}^* = \{0, 1\}^{2k}$ its vertices. Let $D_{2k} = \{(x, y) \in H_{2k} \mid \sum x_i = \sum y_i\}$ be the “generalized diagonal” of the hypercube. Let $D_{2k}^* = \{(x, y) \in \{0, 1\}^{2k} \mid \sum x_i = \sum y_i\}$ be the vertices of H_{2k} with as many x -coordinates set to 1 as y -coordinates. That is, this is the set of valid (deterministic) allocation vectors for k buyers and sellers. Then our claim is equivalent to saying that D_{2k} is (in) the convex hull of D_{2k}^* . By the Krein-Milman theorem a convex set S is exactly the convex hull of its extreme points. An extreme point $s \in S$ is any point in S which can not be written as a convex combination of points in $S \setminus s$. Clearly D_{2k} is convex. It remains to show that the extreme points of D_{2k} are precisely D_{2k}^* . Clearly $D_{2k}^* \subseteq D_{2k}$. So let $(\mathbf{x}, \mathbf{y}) \in D_{2k} \setminus D_{2k}^*$ be a point in D_{2k} that does not have all elements equal to 0 or 1. We show that (\mathbf{x}, \mathbf{y}) is not an extreme point of D_{2k} .

If there is exactly one x_i with $0 < x_i < 1$, then there must be at least one y_j with $0 < y_j < 1$. (Otherwise $\sum x_i \notin \mathbb{N}$, but $\sum y_j \in \mathbb{N}$, which contradicts the assumption that $\sum x_i = \sum y_j$.) Then for $0 < \epsilon < \min\{x_i, 1 - x_i, y_j, 1 - y_j\}$, we have that $(x_i + \epsilon, y_j + \epsilon, \mathbf{x}_{-i}, \mathbf{y}_{-j}) \in D_{2k}$, and also $(x_i - \epsilon, y_j - \epsilon, \mathbf{x}_{-i}, \mathbf{y}_{-j}) \in D_{2k}$. Clearly (\mathbf{x}, \mathbf{y}) is a convex combination of these two. If there is at least two distinct $0 < x_i, x_\ell < 1, i \neq \ell$, then for $0 < \epsilon < \min\{x_i, 1 - x_i, x_\ell, 1 - x_\ell\}$, we have that $(x_i + \epsilon, x_\ell - \epsilon, \mathbf{x}_{-i\ell}, \mathbf{y}) \in D_{2k}$, and also $(x_i - \epsilon, x_\ell + \epsilon, \mathbf{x}_{-i\ell}, \mathbf{y}) \in D_{2k}$. Again, clearly (\mathbf{x}, \mathbf{y}) is a convex combination of these two. Similarly, if there is no $0 < x_i < 1$ there is at least two such y_j, y_ℓ . So D_{2k} is the convex hull of D_{2k}^* . This shows our claim. \square

From this it follows immediately that we need only concern ourselves with the marginal allocation probabilities in computing an optimal randomized mechanism. Therefore we can write this as a LP following the approach of Dobzinski et al. [3] for auctions. We defer the proof of this theorem to appendix C.

Theorem 4 (The optimal randomized mechanism as an LP). *For a fixed number of buyers and sellers, we can compute the optimal truthful-in-expectation mechanism using a linear program that is polynomial in the size of the prior.*

6 Discussion and further work

One question raised by our results is that of the relation between single seller, single buyer market intermediation and two-bidder auctions. As discussed in appendix A the graph algorithm of Papadimitriou and Pierrakos [13] can be used to solve the no short-selling 1×1 market intermediation case, and the derandomization in Dobzinski et al. [3] applies immediately to both this and the 1×1 balanced inventory setting. These give running times of $\mathcal{O}(n^6)$ and $\mathcal{O}(n^7)$ in contrast to a running time of $\mathcal{O}(n^4)$, respectively $\mathcal{O}(n^3)$ in the balanced inventory

case, for our approach in the market intermediation setting. It is not clear immediately that the 2-bidder auction design problem could in turn be solved using a modified version of this algorithm, given the additional complexity of two interdependent regions for each seller. We suspect that there might indeed be a gap between the complexity of these two problems. Furthermore, we believe that an optimal 2-bidder reverse auction can be computed using our dynamic program for the balanced inventory case. Thus a gap between auctions and market intermediation would imply an asymmetry between auctions and reverse auctions. An immediate follow-up question is whether we can give good approximations in polynomial time. In appendix D we show that no good multiplicative guarantees are possible using prior-independent mechanisms.

References

1. Colini-Baldeschi, R., de Keijzer, B., Leonardi, S., Turchetta, S.: Approximately efficient double auctions with strong budget balance. In: Proceedings of SODA. ACM-SIAM (2016)
2. Deng, X., Goldberg, P., Tang, B., Zhang, J.: Revenue maximization in a bayesian double auction market. *Theoretical Computer Science* 539, 1–12 (2014)
3. Dobzinski, S., Fu, H., Kleinberg, R.D.: Optimal auctions with correlated bidders are easy. In: Proceedings of the forty-third annual ACM symposium on Theory of computing. pp. 129–138. ACM (2011)
4. Dütting, P., Roughgarden, T., Talgam-Cohen, I.: Modularity and greed in double auctions. In: Proceedings of the fifteenth ACM conference on Economics and computation. pp. 241–258. ACM (2014)
5. Feldman, M., Gravin, N., Lucier, B.: Combinatorial auctions via posted prices. In: Proceedings of SODA. pp. 123–135. ACM-SIAM (2015)
6. Loertscher, S., Niedermayer, A.: When is seller price setting with linear fees optimal for intermediaries? Tech. rep., Discussion Papers, Department of Economics, Universität Bern (2007)
7. Loertscher, S., Niedermayer, A.: Fee setting intermediaries: on real estate agents, stock brokers, and auction houses. Tech. rep., Discussion paper//Center for Mathematical Studies in Economics and Management Science (2008)
8. Loertscher, S., Niedermayer, A.: Fee-setting mechanisms: On optimal pricing by intermediaries and indirect taxation. Tech. rep., SFB/TR 15 Discussion Paper (2012)
9. McAfee, R.P.: A dominant strategy double auction. *Journal of economic Theory* 56(2), 434–450 (1992)
10. Myerson, R.B.: Optimal auction design. *Mathematics of operations research* 6(1), 58–73 (1981)
11. Myerson, R.B., Satterthwaite, M.A.: Efficient mechanisms for bilateral trading. *Journal of economic theory* 29(2), 265–281 (1983)
12. Niazadeh, R., Yuan, Y., Kleinberg, R.D.: Simple and near-optimal mechanisms for market intermediation. arXiv preprint arXiv:1409.2597 (2014)
13. Papadimitriou, C., Pierrakos, G.: Optimal deterministic auctions with correlated priors. arXiv preprint arXiv:1011.1279 (2010), <http://arxiv.org/abs/1011.1279>
14. Ronen, A.: On approximating optimal auctions. In: Proceedings of the 3rd ACM conference on Electronic Commerce. pp. 11–17. ACM (2001)

A Reduction from the 1×1 case to known algorithms for 2-bidder auctions

For the two bidder single auction setting, Papadimitriou and Pierrakos [13] take a different approach and reduce the problem to that of finding a maximum-weight independent set on a bipartite graph. We can similarly reduce the no short-selling market intermediation problem to this, with only a few steps, but at the cost of a larger polynomial runtime bound. Firstly, in the single auction setting we have a strong exclusivity constraint: We cannot allocate the item to both bidders at the same time, so any valid mechanisms must satisfy $A_1 \cap A_2 = \emptyset$. In the intermediation setting we do not have such a constraint directly; rather we must have B contained in S . However, we can consider \bar{S} the complement of S , and indeed have the same constraint $\bar{S} \cap B = \emptyset$. In an analogous fashion to their construction we can define marginal revenue contributions f_b for the buyer and f_s for the seller, the latter in our setting referring to the marginal profit contribution of *not buying* at point (i, j) . It is easy to see that the maximum-weight independent set on the so constructed bipartite graph will give the optimal mechanism in the no short-selling setting. This approach runs in time $\mathcal{O}(n^6)$ for a support size of n^2 .

It is furthermore possible to obtain an optimal deterministic mechanism via the derandomization approach described in Dobzinski et al. [3]. It is easy to see that the optimal truthful-in-expectation mechanism can be computed through a LP in either market intermediation setting. The derandomization applies to both without modification. A cursory analysis of this approach gives a runtime of $\mathcal{O}(n^7)$ for a support size of n^2 . (The LP involves $2n^2$ variables. The exact time complexity depends on the LP solver chosen; Karmarkar's algorithm runs in time $\mathcal{O}(n^{3.5})$ for n variables.)

B DP for the balanced inventory case

We show how to modify the dynamic programming approach for the balanced inventory setting, which we omitted in section 3. Let $R'(\ell, \beta(\ell), \beta(\ell + 1))$ be the revenue of the optimal mechanism that satisfies the balanced inventory property, sells at points $(\beta(\ell), \ell)$ and rightward, and does not have to pay for columns $\beta(\ell + 1)$ to n . Then the value of R' is given by the following explicit formula analogous to the no short-selling case.

$$R'(\ell, \beta(\ell), \beta(\ell + 1)) = \max_{\beta(1) \leq \dots \leq \beta(\ell-1)} \sum_{s=1}^{\ell} \beta(s) \sum_{b \geq \beta(s)} \psi_{bs} - \sum_{b=1}^{\beta(\ell+1)-1} \sigma(b) \sum_{s \leq \sigma(b)} \psi_{bs} \quad (6)$$

And we can prove a very similar recursion.

Theorem 5 (Recursion for the balanced inventory setting).

The $R'(\ell, \beta(\ell), \beta(\ell + 1))$ as defined above satisfy the following recursion:

$$R'(\ell, \beta(\ell), \beta(\ell + 1)) = \max_{\beta(\ell-1) \leq \beta(\ell)} R'(\ell - 1, \beta(\ell - 1), \beta(\ell)) + \beta(\ell) \sum_{b \geq \beta(\ell)} \psi_{b\ell} - \ell \sum_{\beta(\ell) \leq b < \beta(\ell+1)} \sum_{s \leq \ell} \psi_{bs} \quad (7)$$

Proof. The proof of theorem 1 applies mutatis mutandis. □

This suggests a slight modification of the no short-selling algorithm, listed as algorithm 2.

Algorithm 2 Optimal revenue in the balanced inventory setting

```

1: for  $\ell = 1, \dots, n$  do
2:   for  $\beta(\ell) = 1, \dots, n$  do
3:     for  $\beta(\ell + 1) = \beta(\ell), \dots, n + 1$  do
4:       if  $\ell = 1$  then
5:          $R'(1, \beta(1), \beta(2)) \leftarrow \beta(1) \sum_{b \geq \beta(1)} \psi_{b1} - \sum_{\beta(1) \leq b < \beta(2)} \psi_{b1}$ 
6:       else
7:         Compute  $R'(\ell, \beta(\ell), \beta(\ell + 1))$  using the recursion in theorem 5.
   return  $\max_{\beta(n)} R(n, \beta(n), n + 1)$ .
```

This algorithm obviously has slightly lower runtime than the one in the no short-selling setting, since we only need to loop through values $\beta(\ell + 1) \geq \beta(\ell)$, and only consider values $\beta(\ell - 1) \leq \beta(\ell)$ in the maximum in the recursion. It is straightforward to modify the definition of R' to get rid of the third parameter altogether, by accounting for the cost of buying in all columns. The recursion then needs to be modified to add back the cost of buying in columns $\beta(\ell)$ to n that has already been accounted for in $R'(\ell - 1, \cdot)$. This gives an algorithm that runs in time $\mathcal{O}(n^3)$.

C Proof of theorem 4

Proof. Consider the following LP, where $\psi_{\mathbf{b},\mathbf{s}}$ denotes the prior probability of bid vectors (\mathbf{b}, \mathbf{s}) .

$$\min \sum_{\mathbf{b},\mathbf{s}} \psi_{\mathbf{b},\mathbf{s}} \left(\sum_i -p_i(\mathbf{b}, \mathbf{s}) + \sum_j q_j(\mathbf{b}, \mathbf{s}) \right) \text{ s.t.} \quad (8a)$$

$$\sum_i x_i(\mathbf{b}, \mathbf{s}) = \sum_j y_j(\mathbf{b}, \mathbf{s}) \quad \forall \mathbf{b}, \mathbf{s} \quad (8b)$$

$$p_i(\mathbf{b}, \mathbf{s}) \geq 0, \quad 0 \leq x_i(\mathbf{b}, \mathbf{s}) \leq 1 \quad \forall i, \mathbf{b}, \mathbf{s} \quad (8c)$$

$$q_j(\mathbf{b}, \mathbf{s}) \geq 0, \quad 0 \leq y_j(\mathbf{b}, \mathbf{s}) \leq 1 \quad \forall j, \mathbf{b}, \mathbf{s} \quad (8d)$$

$$\begin{aligned} x_i(\mathbf{b}, \mathbf{s}'_i, \mathbf{s}_{-i})s_i + p_i(\mathbf{b}, \mathbf{s}'_i, \mathbf{s}_{-i}) &\leq \\ &\leq -x_i(\mathbf{b}, \mathbf{s})s_i + p_i(\mathbf{b}, \mathbf{s}) \end{aligned} \quad \forall i, \mathbf{s}'_i, \mathbf{b}, \mathbf{s} \quad (8e)$$

$$\begin{aligned} y_j(b_j, \mathbf{b}_{-j}, \mathbf{s})b_j + p_i(b_j, \mathbf{b}_{-j}, \mathbf{s}) &\leq \\ &\leq y_j(\mathbf{b}, \mathbf{s})b_j - q_j(\mathbf{b}, \mathbf{s}) \end{aligned} \quad \forall j, b'_j, \mathbf{b}, \mathbf{s} \quad (8f)$$

$$x_i(\mathbf{b}, s_{\max}, \mathbf{s}_{-i}) = 0 \quad \forall i, \mathbf{s}_{-i}, \mathbf{b} \quad (8g)$$

$$q_j(0, \mathbf{b}_{-j}, \mathbf{s}) = 0 \quad \forall j, \mathbf{b}, \mathbf{s}_{-j} \quad (8h)$$

Here, conditions (b)-(d) ensure feasibility of the (class of) mechanisms. (e) and (f) ensure truthfulness on the sellers' and buyers' side, respectively. Together with (g) and (h) they give IR. The proof is the same as in Dobzinski et al. [3] for auctions. This gives us a non-empty class of valid mechanisms, all of which are optimal and have the same expected utilities for the players. Finding an arbitrary member of this class is easy.² \square

D Prior-independent mechanisms with approximation guarantees

We consider simple mechanisms that do not use knowledge of the prior, other than that all prices lie in $[0, 1]$. In the following let D_x be the diagonal $b = s + x$, and let M_y be the mechanism that buys and sells iff the valuations are to the right of D_y , i.e. $b \geq s + y$. It is easy to see that for $x > y$, and any point on D_x , the revenue obtained by M_y is $2y - x$. (For point $(s, s + x)$, M_y buys from the seller for $(s + x) - y$ and sells to the buyer for $y + s$.)

Theorem 6 (Prior-independent mechanism for the 1×1 case). *There is a prior-independent deterministic mechanism that extracts revenue at least*

² It is equivalent to finding for a given point in $(\mathbf{x}, \mathbf{y}) \in D_{2k}$ coefficients for a convex combination of vertices in D_{2k}^* . While the number of extreme points is exponential in the k , by Carathéodory only $2k + 1$ of the vertices are needed. To find them, we could simply pick any $2k$ vertices of D_{2k}^* to partition D_{2k} into two sets. Discard all the vertices not contained in the same partition as (\mathbf{x}, \mathbf{y}) , and repeat with the remaining vertices.

$\text{OPT} - \frac{2}{3}n$ in the single seller, single buyer setting with valuations on the unit interval. Furthermore, this is optimal among prior-independent mechanisms in this setting.

Proof. Consider the mechanism $M_{2n/3}$ that buys and sells iff $b \geq s + \frac{2n}{3}$. This achieves expected revenue of at least $\text{OPT} - \frac{2n}{3}$ for any point (b, s) : Clearly for any b and s , the optimal revenue will be at most n . On the other hand, if $s > b - \frac{2n}{3}$, then the revenue is bounded by $\frac{2n}{3}$. $M_{2n/3}$ will extract revenue at least $\frac{n}{3}$ for $s \leq b - \frac{2n}{3}$ and 0 otherwise.

Moreover, this is optimal among prior-independent mechanisms. Consider the points $(b, s) = (\frac{2n}{3}, 0)$, $(n, \frac{n}{3})$ and $(n, 0)$. If a mechanism does not buy and sell at $(\frac{2n}{3}, 0)$, then for the instance that has all the probability concentrated at this point, it generates revenue 0, while the optimal revenue would clearly be $\frac{2n}{3}$. Similarly for $(n, \frac{n}{3})$. Thus, we have to buy and sell at both those points if we want to generate profit at least $\text{OPT} - \frac{2n}{3}$ in the worst case. On the other hand, if we buy and sell at both those points, it follows that by truthfulness we cannot charge more than $\frac{2n}{3}$ to the buyer nor offer less than $\frac{n}{3}$ to the seller at $(n, 0)$ either. Therefore we generate profit $\frac{n}{3}$ at $(n, 0)$, as opposed to a maximum revenue of n . Thus no deterministic prior-independent mechanisms can generate revenue greater than $\text{OPT} - \frac{2n}{3}$ for all instances. \square

By allowing randomization we can do even better.

Theorem 7 (Prior-independent randomized mechanism for the 1×1 case). *There is a prior-independent mechanism that extracts revenue at least $\text{OPT} - \frac{n}{2}$ in the single seller, single buyer setting with valuations on the unit interval.*

Proof. Take the mechanism M that randomizes uniformly amongst $\{M_y : \frac{n}{2} \leq y \leq n\}$. Then for any point on a diagonal $D_{\frac{n}{2}+x}$, the expected revenue in M is given by $\int_0^x 2 \cdot [2(\frac{n}{2} + t) - (\frac{n}{2} + x)] dt = x$. But the optimal revenue for any point on $D_{\frac{n}{2}+x}$ is $\frac{n}{2} + x$. Clearly, for any diagonal to the left of $D_{\frac{n}{2}}$, the revenue is bounded by $\frac{n}{2}$. Thus, M gives expected revenue of at least $\text{OPT} - \frac{n}{2}$. \square

On the other hand, for multiplicative guarantees knowledge of the prior is essential, as it is easy to see that we can give at best a $\frac{1}{n}$ worst-case guarantee, dependent on the granularity of the prior. We give a (weak) guarantee that depends on the granularity of the prior support, and show that this is the best we can do. It immediately follows that for arbitrary continuous distributions we could not give any guarantee.

Theorem 8 (Prior-independent randomized mechanism for the 1×1 case, multiplicative). *There is a prior-independent mechanism that extracts revenue at least $\frac{1}{n} \cdot \text{OPT}$ in the single seller, single buyer setting with valuations on $\{i : 0 \leq i \leq n\}$. Furthermore, this is optimal among prior-independent mechanisms in this setting.*

Proof. As before let D_i be the diagonal $s = b + i$, and let M_j be the mechanism that buys iff the valuations are to the right of D_j , i.e. $s \geq b + j$. It is easy to see that for $i > j$, and any point on D_j , the revenue obtained by M_i is $2j - i$.

Take the mechanism M that randomizes uniformly amongst $\{M_j : 1 \leq j \leq n\}$. Then for any point on a diagonal D_i it is easy to see that the expected revenue arising from the M_1, \dots, M_{i-1} equals 0. With probability $\frac{1}{n}$ the mechanism selects M_i , giving optimal revenue for any point on this diagonal. In expectation the mechanism thus generates revenue $\frac{1}{n} \cdot \text{OPT}$.

To see that this is tight, suppose we want to extract a c -fraction of the optimal revenue. In order to generate a c -fraction of the optimal revenue for any point on D_1 , we need to place at least c probability mass there. It is easy to see that this gives net revenue 0 for any point on D_2 , so we need to place an additional c probability, for a total allocation probability of $2c$. By induction it follows that we need to buy and sell with probability nc at $(n, 0)$, and thus $c \leq \frac{1}{n}$. \square