



The Limits of an Information Intermediary in Auction Design

REZA ALIJANI, Duke University, USA

SIDDHARTHA BANERJEE, Cornell University, USA

KAMESH MUNAGALA, Duke University, USA

KANGNING WANG, Duke University, USA

We study the limits of an information intermediary in the classical Bayesian auction, where a revenue-maximizing seller sells one item to n buyers with independent private values. In addition, we have an intermediary who knows the buyers' private values, and can map these to a public signal so as to increase consumer surplus. This model generalizes the single-buyer setting proposed by Bergemann, Brooks, and Morris, who present a signaling scheme that raises the optimal consumer surplus, by guaranteeing that the item is always sold and the seller gets the same revenue as without signaling. Our work aims to understand how this result ports to the setting with multiple buyers.

We likewise define the benchmark for the optimal consumer surplus: one where the auction is efficient (i.e., the item is always sold to the highest-valued buyer) and the revenue of the seller is unchanged. We show that no signaling scheme can guarantee this benchmark even for $n = 2$ buyers with 2-point valuation distributions. Indeed, no signaling scheme can be efficient while preserving any non-trivial fraction of the original consumer surplus, and no signaling scheme can guarantee consumer surplus better than a factor of $\frac{1}{2}$ compared to the benchmark. These impossibility results are existential (beyond computational), and provide a sharp separation between the single and multi-buyer settings.

In light of this impossibility, we develop signaling schemes with good approximation guarantees to the benchmark. Our main technical result is an $O(1)$ -approximation for i.i.d. regular buyers, via signaling schemes that are conceptually simple and computable in polynomial time. We also present an extension to the case of general independent distributions.

CCS Concepts: • **Theory of computation** → **Algorithmic game theory**; • **Mathematics of computing** → **Probabilistic algorithms**.

Additional Key Words and Phrases: auctions, signaling, Bayesian persuasion

ACM Reference Format:

Reza Alijani, Siddhartha Banerjee, Kamesh Munagala, and Kangning Wang. 2022. The Limits of an Information Intermediary in Auction Design. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22), July 11–15, 2022, Boulder, CO, USA*. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3490486.3538370>

1 INTRODUCTION

Consider a seller selling an item to a buyer, whose private value V is drawn from some known distribution $\mathcal{D}$. The overall *social welfare* is maximized when the seller sells the item for \$0, assuming the seller has no cost for the item. In contrast, to maximize the (average) revenue, the seller's optimal strategy is to sell at a revenue-maximizing price, which may lead to welfare loss due to the item going unsold.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '22, July 11–15, 2022, Boulder, CO, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9150-4/22/07...\$15.00

<https://doi.org/10.1145/3490486.3538370>

More generally, in a single-item Bayesian auction with n buyers with independent private valuations, a welfare-optimal mechanism is the second-price (or VCG) auction, which always gives the item to the highest-valued buyer. In contrast, even when the buyers have i.i.d. regular valuations, the revenue-optimal mechanism was shown by Myerson [19] to be a second-price auction with a *reserve price*; this may lead to the item going unsold. The situation is more complex for non-regular distributions, and/or non-i.i.d. buyers, where the revenue-optimal mechanism may in addition sell the item to a buyer with lower value than the highest, leading to additional welfare loss. We visualize this via a *revenue-CS trade-off* diagram (Fig. 1b), where, for different mechanisms and value distributions, we plot expected consumer-surplus (i.e., value minus payment), denoted CS , versus expected seller-revenue, denoted by $\mathcal{R}$. Any welfare-maximizing mechanism including VCG (point V) lies on the $\mathcal{R} + \text{CS} = \mathcal{W}^*$ line. In contrast, Myerson's mechanism (point M) has revenue $\mathcal{R}^M$ greater than that under VCG, but can also lie below the maximum-welfare line.

Information Intermediary. Now consider the same setting, but with an additional *information intermediary*: a third-party who knows the *true* buyer values $\vec{V} = (V_1, V_2, \dots, V_n)$ and can provide a “signal” or side-information to the seller and the buyers. Both the signal and signaling scheme are common knowledge to all agents (buyers and seller), who can thus use Bayes' rule to update the prior over valuations given the signal. The signal “re-shapes” the joint prior over the buyer valuations in a *Bayes-plausible* manner (i.e., such that the posterior averaged over signals equals the prior). Though the intermediary can modulate information, it does not control the mechanism, which still resides with the seller. Such a setting is motivated by ad exchanges, where the platform (or intermediary) acts only as a clearinghouse, and does not itself run a mechanism. Therefore, given the signal, the seller then proposes the revenue-maximizing mechanism, and buyers bid optimally, under the posterior distribution. We illustrate this in Fig. 1a.

Formally, consider a setting where n buyers have independent private valuations $\vec{V}$ drawn from a distribution $\mathcal{D} = \mathcal{D}_1 \times \mathcal{D}_2 \times \dots \times \mathcal{D}_n$. The valuations $\vec{V}$ are known to the intermediary, who maps them to a signal σ via a public signaling scheme $\mathcal{Z}$. Given σ , all agents compute the posterior $\mathcal{S}$ over buyer values; note these can now be correlated. The seller then proposes a mechanism $\mathcal{M}_{\mathcal{S}}$ (comprising allocation and payment rules) which maximizes its expected revenue assuming buyers act in a manner which is ex-post incentive-compatible (IC) and interim individually-rational (IR) given $\mathcal{S}$. If σ is such that $\mathcal{S} = \mathcal{D}$, then $\mathcal{M}_{\mathcal{S}}$ is Myerson's auction (point M in Fig. 1b); on the other hand, if the signal fully reveals $\vec{V}$, then the seller can extract full surplus (i.e., get revenue $\mathcal{W}^*$, point A in Fig. 1b). Moreover, the seller gets revenue at least $\mathcal{R}^M$ under any signaling scheme, as she can always ignore the signal (see Section 2). Thus any signaling scheme $\mathcal{Z}$ must give a point in the shaded triangle with consumer surplus $\text{CS}(\mathcal{Z})$ and revenue $\mathcal{R}(\mathcal{Z})$, and the maximum possible surplus OPT is achieved at point O in Fig. 1b. Now we can ask:

What revenue-CS trade-offs can an information intermediary achieve via signaling? More specifically, what is the maximum possible consumer surplus that is achievable?

In the single-buyer case, the seminal work of Bergemann, Brooks, and Morris [2] completely answer these questions by showing that the entire shaded region is always achievable. In particular, the point O is met by a simple signaling scheme where the revenue is *exactly* $\mathcal{R}^M$, and the item is *always* sold thus the mechanism is efficient.

In this work we study the effectiveness of an information intermediary in a multi-buyer (i.e., $n \geq 2$ buyers) Bayesian auction. In brief, we expose a sharp separation between the single and multi-buyer settings, as in the latter, no signaling scheme can guarantee more than a constant fraction of the optimal consumer surplus (OPT in Fig. 1b). On the positive side, we obtain a novel yet simple signaling scheme with strong approximation guarantees for a wide range of settings.




Fig. 1. (a) The auction with information intermediary setting, where the intermediary has full knowledge of valuations $\bar{V}$, and can use this to provide a signal σ to the seller and the buyers. The seller then uses the revenue-optimal mechanism M_S for the posterior distribution over valuations S given σ . (b) Two-dimensional space of seller revenue, $\mathcal{R}$, and consumer surplus, $\mathcal{CS}$, of different signaling schemes. The points M and V correspond to Myerson’s and VCG mechanisms, and the point A corresponds to selling to the highest-valued buyer at her value when the seller has full information.

While our main focus is on theoretical results, our work has broader practical relevance. Consider an agency like the FCC with privileged information about bidders in a spectrum auction, or a bid optimizer working for multiple competing clients in an ad-exchange. These intermediaries have private information about the buyers, and can selectively release it to influence the auction. For instance, in an ad exchange, the platform running the exchange (intermediary) typically uses machine learning and advertiser features to infer true valuations. However, the pricing rules are decided by the publishers (seller) and not the exchange. For its own long term viability, the platform clearly has incentives to make both parties – publisher and advertisers – as happy as possible, and would therefore like to release information selectively to the publisher in order to maximize advertiser happiness (consumer surplus) while keeping publisher happiness (revenue) at least what it is without its presence. In effect, we use the alternate view of the market segmentation problem in [2] as a special case of a signaling problem where a more informed intermediary works for the benefit of the buyers.

Our work also fits in a broader space of multi-criteria optimization where a third-party platform or government agency can release information about agents to a principal in charge of an activity such as admissions or hiring, so as to trade-off the principal’s objective such as maximizing quality of hire, with a societal objective such as fairness or diversity.

1.1 Our Results

We consider a single-item auction with n buyers with discrete valuations. We assume the buyer valuations are independent, so $\mathcal{D} = \mathcal{D}_1 \times \mathcal{D}_2 \times \dots \times \mathcal{D}_n$, where $\mathcal{D}_i$ has support size $\mathcal{K}_i$, and the size of the union of the supports is $\mathcal{K}$. We parametrize our results in terms of $n, \mathcal{K}_i$, and $\mathcal{K}$.

Our first set of results (Section 3) shows a sharp demarcation between the cases of $n = 1$ and $n \geq 2$ buyers. In contrast with the former case (where signaling achieves the entire shaded region in Fig. 1b), we show in the latter case, the entire segment BO is not achievable; indeed, the only achievable points on segment AO are arbitrarily close to A . Therefore, achieving full welfare requires sacrificing an arbitrarily large fraction of consumer surplus compared to the no-signaling baseline.

THEOREM 1.1 (PROVED IN SECTION 3). *For any given constant $\varepsilon > 0$, there are instances with $n = 2$ buyers each with $\mathcal{K}_i = 2$, where any signaling scheme $\mathcal{Z}$ under which the revenue-optimal auction obtains full welfare (i.e., allocates to highest-value buyer), has $\text{CS}(\mathcal{Z}) \leq \varepsilon \cdot \text{CS}(\mathcal{D})$, where $\text{CS}(\mathcal{D})$ is the consumer surplus of Myerson's auction without signaling.*

We next ask if we can sacrifice on welfare, but raise a consumer surplus arbitrarily close to OPT ? We again answer in the negative, and show a lower bound of 2 on the approximation ratio.

THEOREM 1.2 (PROVED IN SECTION 3). *For any constant $\varepsilon > 0$, there are problem instances with $n = 2$ buyers each with $\mathcal{K}_i = 2$, where any signaling scheme $\mathcal{Z}$ has $\text{CS}(\mathcal{Z}) \leq (\frac{1}{2} + \varepsilon) \cdot \text{OPT}$.*

We note that the above results are *existential impossibility results*, and do not depend on the complexity of the signaling scheme.¹ Overall, the negative results in Theorems 1.1 and 1.2 strongly suggest that in this setting, the focus should be on approximating the consumer surplus.

The situation improves in Section 3.3 when we restrict to $\mathcal{D}_i$ that are identical and (discrete-)regular. Here, we first circumvent Theorem 1.1 by showing a simple signaling scheme that achieves the point B (i.e., optimal welfare, and same consumer surplus as under Myerson's auction). One problem that remains, however, is that Myerson's auction may have arbitrarily poor CS: for example, if the $\mathcal{D}_i$ are regular, and chosen such that the reserve price is the highest value in the support, then $\text{CS} = 0$, while $\text{OPT} > 0$ (and so the approximation factor of Myerson's auction is unbounded). Indeed, even restricting to MHR priors, one can construct instances where the reserve price is close to the maximum value in the support, leading Myerson's auction to have vanishing CS relative to OPT . This is one reason why getting any non-trivial approximation to OPT is challenging, and we present more discussion in Section 4.1.

In Section 4, we present our main technical result, where we show that when buyers' valuations are drawn from i.i.d. regular distributions, then a simple signaling scheme achieves a constant-approximation to OPT . In more detail, our RANK_t signaling scheme is based on two simple but critical steps: First, the intermediary can use its knowledge of agent valuations to perform a *pre-screening* step that eliminates all but the top- t buyers (for a carefully chosen t). Second, given the top t buyers, it can then choose a uniform buyer among this set to serve as a *hold-out buyer*, who the seller can sell to in case she is unable to raise sufficient revenue from the remaining $t - 1$ buyers via an auction; this can be achieved by using the single-buyer signaling scheme of Bergemann et al. [3] on the chosen buyer. Using a combination of these two ideas, we get the following:

THEOREM 1.3 (PROVED IN SECTION 4). *When the $\mathcal{D}_i$'s are identical and regular, there is a signaling scheme achieves an $O(1)$ -approximation to the optimal consumer surplus OPT , and has computation time polynomial in n and $\mathcal{K}$.*

The nice feature of RANK_t is that it generates signals with posteriors that are (non-identical) product distributions, so that the seller's optimal auction is Myerson's auction [19], which is also ex-post IC and IR. This scheme also turns out to achieve OPT for the special case when $\mathcal{K} = 2$ and n is arbitrary.

In the full paper [1], we extend this scheme to when the buyers are independent, but not necessarily identical or regular. We obtain the following theorem for this case.

THEOREM 1.4 (PROVED IN THE FULL PAPER [1]). *When the $\mathcal{D}_i$'s are arbitrary, the RANK_t scheme achieves an $O(\min(n \log n, \mathcal{K}^2))$ -approximation to the optimal consumer surplus OPT , and has computation time polynomial in n and $\mathcal{K}$.*

¹Our proofs also imply the same lower bounds when the seller is constrained to run an *ex-post* IR mechanism, provided the intermediary's signals induce a product-form posterior distribution.

1.2 Intuition and Techniques

For any n , the optimal signaling scheme for maximizing surplus can be obtained via an infinite-sized linear program (see Eq. (2) in Section 3) with variables for every possible signal, i.e., every possible joint distribution over buyer valuations. Further, for each such signal, the quantity of interest is the consumer surplus of the revenue-optimal auction given the signal. For $n = 1$ case, Bergemann et al. show this LP has a special structure in that it admits a basis comprising of “equal-revenue distributions” containing the revenue-maximizing price (see Section 2.2). Our work shows that this breaks down for optimal auctions with signaling involving $n \geq 2$ buyers.

To understand why things change dramatically from $n = 1$ to $n \geq 2$ buyers, in the former case, the optimal mechanism is a simple posted price scheme and its revenue is continuous in the distribution $\mathcal{D}$. However, with multiple buyers, the optimal auction does not have simple structure even for independent buyers (see Algorithm 1), and we need to analyze the consumer surplus of this auction, which can be a discontinuous function of the prior. (See Section 3 for examples.) Further, for correlated buyers, the revenue of the auction itself may not be continuous in the prior! Indeed, a celebrated result of Crémer and McLean [8] shows that slightly perturbing an independent prior to a correlated one can discontinuously increase the revenue to $\mathcal{W}^*$, hence decreasing consumer surplus to 0. (See Theorem 2.1 in Section 2.) This makes it tricky to reason about the optimal signaling scheme, leading to the gap between our upper and lower bounds.

In more detail: Our proofs of Theorems 1.1 and 1.2 use a special case of the Crémer-McLean characterization [8]: for $n = 2$ buyers each with $\mathcal{K}_i = 2$, under any non-independent prior the seller can extract full social surplus as revenue. This lets us focus on signaling schemes where buyers’ posterior given each signal are product distributions. Using Myerson’s characterization of the optimal auction for discrete valuations [13], we show a structural characterization that reduces the space of optimal signals to a sufficiently simple form, yielding the desired counterexamples. Note that we still need to reason about a large space of possible product distributions as signals, which makes our constructions quite non-trivial.

The technically most interesting result in the paper is the $O(1)$ -approximation signaling scheme for i.i.d. buyers (Theorem 1.3 in Section 4). The challenge is the following: Even if we restrict the space of signals so that the posteriors are (non-identical) product distributions, this space is still infinite size, with CS being a discontinuous function in this space. Our signaling scheme in Section 4 balances the trade-off between revealing enough information about valuations so that the item is sold to a high-value buyer, and revealing too much information such that the seller extracts all the surplus. Balancing these is delicate; nevertheless, our final scheme is simple with polynomial computation time and signal complexity. We present more intuition in Section 4.1, where we argue that the guarantee in Theorem 1.3 cannot be achieved in a straightforward fashion.

1.3 Related Work

The general problem of information structure design considers how sharing additional information can influence the outcome of a mechanism. Different variants of this problem have been formulated and studied; we refer the reader to [4, 10] for surveys. Of particular importance to us is the *Bayesian persuasion problem* formulated by Kamenica and Gentzkow [16], where a receiver selects a utility-maximizing action based on incomplete information about the state of nature. A sender who knows the state of nature can signal side-information to the receiver so that the action taken by the receiver is utility-maximizing for the sender. This general problem has been widely studied in different domains such as monopoly pricing and advertising [2, 7, 15, 22]. For this problem, there is a distinction between existence and computational results, and the work of Dughmi and Xu [12]

studies the computational complexity of finding the optimal signaling scheme under different input models.

The restriction of our problem to one buyer is the monopoly pricing problem. Here, the intermediary is the sender whose utility is consumer surplus, and the seller is the receiver whose action space is take-it-or-leave-it prices and whose utility is revenue. Beginning with the work of Bergemann et al. [2], several works [6, 9, 11, 15, 17, 18, 20] have considered various extensions and modifications to this basic problem. Unlike monopoly pricing where the buyer is perfectly informed, in our setting, not only the seller, but also all the buyers are receivers, in the sense that they have imperfect knowledge of the true valuations of other buyers, and modify their respective bidding strategies in response to the intermediary's signal to maximize their own utilities. Our setting is therefore a Bayesian persuasion problem with multiple receivers, and this aspect makes it significantly more complex.

There has been work on signaling in auctions that cannot be modeled as Bayesian persuasion, i.e., in which the common signal is not generated by an intermediary who knows all the true values of the buyers. For instance, in the work of Bergemann and Pesendorfer [5], the auctioneer has perfect information about buyer valuations and controls the precision to which buyers can learn it, and in the work of Fu et al. [14], the seller's signal is drawn from a distribution that is correlated with the buyer's value. In both these works, the goal is to maximize seller revenue. Finally, Shen et al. [21] studies equilibria of optimal auctions when each buyer commits to a signaling scheme with imperfect knowledge of other buyers' valuations, while Bergemann et al. [3] studies equilibria in first price auctions when buyers are provided correlated signals about other buyers' valuations. In contrast with the former, our work considers a richer space of signals via an information intermediary, while compared to the latter, in our setting the seller's mechanism is not fixed, but is instead also a function of the information structure.

2 PRELIMINARIES

We consider Bayesian single-item auctions with n buyers, with independent private valuations $\vec{V} = (V_1, V_2, \dots, V_n)$ drawn from a known product distribution $\mathcal{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_n$. Unless otherwise stated, we present our results for the setting in which each $\mathcal{D}_i$ is discrete. We denote by $\mathcal{K}_i$ the size of the support of $\mathcal{D}_i$, and by $\mathcal{K}$ the size of the union of these supports.

For distribution $\mathcal{D}_i$, we use $f_{\mathcal{D}_i}$ to denote its probability mass function, and define $S_{\mathcal{D}_i}(x) = \Pr_{V_i \sim \mathcal{D}_i}[V_i \geq x]$ and $F_{\mathcal{D}_i}(x) = \Pr_{V_i \sim \mathcal{D}_i}[V_i \leq x]$. For a joint distribution $\mathcal{D}$ and vector $\vec{v}$, we use $\Pr[\mathcal{D} = \vec{v}] = f_{\mathcal{D}}(\vec{v})$ as shorthand for denoting the probability of $\vec{v}$ drawn from $\mathcal{D}$.

2.1 Revenue-Maximizing Auctions

Given any shared prior $\mathcal{D}'$ on the valuations of the buyers, which in the case of signaling, can be different from $\mathcal{D}$ and arbitrarily correlated, the seller runs an optimal (revenue maximizing) auction that satisfies ex-post incentive compatibility and *interim* individual rationality. The standard description of these constraints is relegated to the full paper [1].

For any prior $\mathcal{D}'$, let $(\mathcal{R}(\mathcal{D}'), \mathcal{W}(\mathcal{D}'), \text{CS}(\mathcal{D}'))$ denote the expected revenue, welfare (or *total surplus*) and consumer surplus under the revenue-maximizing auction. Then we have $\text{CS}(\mathcal{D}') = \mathcal{W}(\mathcal{D}') - \mathcal{R}(\mathcal{D}')$, and:

$$\mathcal{R}(\mathcal{D}') = \sum_{\vec{v}} \Pr[\mathcal{D}' = \vec{v}] \cdot \sum_i \theta_i^*(\vec{v}) \quad \text{and} \quad \mathcal{W}(\mathcal{D}') = \sum_{\vec{v}} \Pr[\mathcal{D}' = \vec{v}] \cdot \sum_i v_i x_i^*(\vec{v}),$$

where $x^*(\vec{v}) \geq 0$ and $\theta^*(\vec{v})$ are the allocation rule and the payment rule of the optimal auction given any realized valuation profile $\vec{v}$.

Our work builds on two special cases – independent valuations, and full surplus extraction.

Optimal auction for independent valuations. When $\mathcal{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_n$ is a product distribution, the optimal auction has a simple form given by Myerson [19]. For distribution $\mathcal{D}_i$ with support $z_1 < z_2 < \cdots < z_k$, its *virtual value* function $\varphi_{\mathcal{D}_i}$ is defined as:

$$\varphi_{\mathcal{D}_i}(z_k) = z_k \quad \text{and} \quad \varphi_{\mathcal{D}_i}(z_\ell) = z_\ell - (z_{\ell+1} - z_\ell) \frac{S_{\mathcal{D}_i}(z_{\ell+1})}{f_{\mathcal{D}_i}(z_\ell)}, \quad \forall \ell < k. \quad (1)$$

If buyer i is the only buyer in the system, the optimal auction sets a fixed price, and the buyer buys the item when her valuation is at least this price. The *reserve price* of $\mathcal{D}_i$, denoted $r_{\mathcal{D}_i}$ is the smallest value r in the support of $\mathcal{D}_i$ that maximizes the corresponding revenue $rS_{\mathcal{D}_i}(r)$. It is easy to check that $\varphi_{\mathcal{D}_i}(r_{\mathcal{D}_i}) \geq 0$.

Throughout this paper, we assume the distributions $\mathcal{D}_i$ are *regular*, so that $\varphi_{\mathcal{D}_i}(z)$ is a non-decreasing function of z . Therefore, for all $v < r_{\mathcal{D}_i}$, we have $\varphi_{\mathcal{D}_i}(v) < 0$. Our results for the non-i.i.d. case also hold when the distributions are non-regular, by using the non-decreasing ironed virtual value function [13, 19] instead.

For discrete regular distributions, Myerson's auction takes the form [13] in Algorithm 1. Note that this auction is also *ex-post* IC and IR.

ALGORITHM 1: Myerson's Auction with prior $\mathcal{D}$ and valuations $\vec{v}$.

Sort the buyers in decreasing order of $q_i = \varphi_{\mathcal{D}_i}(v_i)$. Assume no two values are identical (can be ensured by using a fixed tie-breaking rule).

Allocate to the bidder j with highest virtual value q_j , provided $q_j \geq 0$.

Let m be the bidder with second highest virtual value, and let $w = \max(0, q_m)$.

Charge j the smallest value z in the support of $\mathcal{D}_j$ such that $\varphi_{\mathcal{D}_j}(z) > w$.

Extracting full surplus as revenue. At the other extreme, a celebrated result of Crémer and McLean [8] shows that for distributions $\mathcal{D}'$ which are “sufficiently correlated”, the optimal auction extracts full surplus (i.e., the revenue equals the maximum valuation in each valuation profile). Formally, the result requires that for each agent, their conditional distribution over others' values given their own value is full rank; for our purposes, we require a restriction of their result to $n = 2$ buyers, each with two possible valuations.

THEOREM 2.1 ([8]). *For $n = 2$ buyers, where each buyer i has $\mathcal{K}_i = 2$ and the joint distribution over the valuations is $\mathcal{D}'$, the seller (who faces an interim IR constraint) can extract the entire social welfare (i.e. get expected revenue equal to the expected value of the maximum of the buyer's valuations) when $\mathcal{D}'$ is a correlated (i.e. not independent) distribution.*

2.2 Auctions with an Information Intermediary

We next formalize the model of an information intermediary illustrated in Fig. 1a. Since the effect of the intermediary's signal is captured by the resulting posterior distribution over valuations, for ease of notation, we henceforth use “signal” to refer to a distribution $\mathcal{S}$ over valuations.

A signaling scheme $\mathcal{Z} = \{\gamma_q, \mathcal{S}_q\}_{q \in [m]}$ comprises a collection of *signals* (i.e., joint distributions over valuations) $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_m$ and corresponding non-negative weights $\gamma_1, \gamma_2, \dots, \gamma_m$. The scheme $\mathcal{Z}$ is feasible (or “Bayes plausible” [16]) if it satisfies $\sum_q \gamma_q = 1$ and $\sum_q \gamma_q \mathcal{S}_q = \mathcal{D}$. The intermediary commits to scheme $\mathcal{Z}$ before the auction, and it is known to the seller and all buyers.

The intermediary maps observed valuation profile $\vec{v} \sim \mathcal{D}$ to signal $\mathcal{S}_q$ with probability $\frac{\gamma_q \Pr[\mathcal{S}_q = \vec{v}]}{\Pr[\mathcal{D} = \vec{v}]}$. The seller uses $\mathcal{S}_q$ as the shared prior and runs an optimal auction on the buyers. Note that though $\mathcal{D}$ is a product distribution, the $\{\mathcal{S}_q\}$ can be correlated. Abusing the notations introduced earlier in

Section 2.1, we denote the revenue generated by signaling scheme $\mathcal{Z}$ as $\mathcal{R}(\mathcal{Z}) = \sum_q \gamma_q \mathcal{R}(S_q)$, its consumer surplus by $\text{CS}(\mathcal{Z}) = \sum_q \gamma_q \text{CS}(S_q)$, and its welfare by $\mathcal{W}(\mathcal{Z}) = \sum_q \gamma_q \mathcal{W}(S_q)$.

When $\mathcal{D}$ is a product distribution, the revenue from any signaling scheme must be at least the optimal revenue of Myerson's auction without signaling, $\mathcal{R}(\mathcal{D})$. To see this, we note that Myerson's auction on $\mathcal{D}$ is ex-post IC and IR. This means that this allocation and payment rule is still a feasible (interim IC and IR) mechanism conditioned on receiving any signal, completing the argument. Therefore, the consumer surplus $\text{CS}(\mathcal{Z})$ under any signaling scheme $\mathcal{Z}$ is bounded by the difference of the maximum possible welfare $\mathcal{W}^* = \mathbb{E}_{\tilde{V} \sim \mathcal{D}}[\max_i V_i]$ and the maximum revenue without signaling $\mathcal{R}(\mathcal{D})$. We henceforth denote this bound as OPT , which is defined as follows:

$$\text{OPT} = \mathcal{W}^* - \mathcal{R}(\mathcal{D}).$$

We say that $\mathcal{Z}$ is a τ -approximation signaling scheme if $\text{CS}(\mathcal{Z}) \geq \frac{\text{OPT}}{\tau}$. Our goal is to find the best approximation factor τ via a signaling scheme whose computation time is polynomial in n and $\mathcal{K}$. In the rest of the paper, we omit the dependence on $\mathcal{D}$ when clear from context.

Optimal signaling for a single buyer. For $n = 1$ buyer, Bergemann et al. [2] present a signaling scheme with consumer surplus exactly equal to OPT (i.e., implementing the point O in Fig. 1b). Their signaling scheme constructs distributions (signals) $S_1, S_2, \dots, S_m$ and assigns weights $\gamma_1, \gamma_2, \dots, \gamma_m$ to them such that $\sum_q \gamma_q S_q = \mathcal{D}$.

Let prior $\mathcal{D}$ takes value v_i with probability η_i , where $0 < v_1 < \dots < v_k$. Let $\vec{\eta} = (\eta_1, \eta_2, \dots, \eta_k)$. In each iteration ℓ , the algorithm constructs an *equal revenue* distribution S_ℓ and subtracts it from the prior $\mathcal{D}$. This equal revenue distribution assigns positive probability $\eta_{i\ell}$ to v_i if $\eta_i > 0$ and assigns $\eta_{i\ell} = 0$ if $\eta_i = 0$. In S_ℓ , the seller raises equal revenue by setting the price to be any of the values v_i with $\eta_i > 0$. It is easy to see that the equal revenue condition specifies a unique distribution S_ℓ . Note that since this signal is equal revenue, (we may assume) the seller sets the lowest value as price, so that the item always sells and the consumer surplus is maximum possible.

Let $\vec{\eta}_\ell$ be the probability vector of S_ℓ . We set the largest weight γ_ℓ such that $\vec{\eta} - \gamma_\ell \vec{\eta}_\ell \geq 0$. We update $\mathcal{D}$ by setting $\vec{\eta}$ to $\vec{\eta} - \gamma_\ell \vec{\eta}_\ell$, and increase ℓ by one. We repeat this till the support of $\mathcal{D}$ becomes empty. The $\{\gamma_\ell, S_\ell\}$ specifies the signaling scheme. We illustrate this procedure by an example.

Example 2.2. Suppose the type space is $\{1, 2, 3\}$ and $\mathcal{D} = \langle \frac{1}{3}, \frac{1}{3}, \frac{1}{3} \rangle$ are the probabilities of these types. The monopoly price is $\theta = 2$ with revenue $\mathcal{R}(\mathcal{D}) = \frac{4}{3}$, while the point A in Fig. 1b has social welfare $\mathcal{R}(A) = \mathcal{W}^* = \mathbb{E}[\mathcal{D}] = 2$. Suppose $S_1 = \langle \frac{1}{2}, \frac{1}{6}, \frac{1}{3} \rangle$ with $\gamma_1 = \frac{2}{3}$; $S_2 = \langle 0, \frac{1}{3}, \frac{2}{3} \rangle$ with $\gamma_2 = \frac{1}{6}$; and $S_3 = \langle 0, 1, 0 \rangle$ with $\gamma_3 = \frac{1}{6}$. It is easy to check that the monopoly price for each signal is the lowest price in its support so that the item always sells, and $\sum \gamma_i \mathcal{R}(S_i) = \frac{4}{3}$. Therefore, $\sum \gamma_i \text{CS}(S_i) = 2 - \frac{4}{3} = \frac{2}{3} = \text{OPT}$, which corresponds to point O in Fig. 1b.

We henceforth use $\text{BBM}(v, D)$ to refer to this scheme when the buyer has valuation distribution D and the realized value is $v \sim D$. Below we state some critical properties of the BBM scheme which we use in our results.

LEMMA 2.3 (IMPLICIT IN [2]). *For a single buyer with value distribution $\mathcal{D}$ (with reserve price $r_{\mathcal{D}}$), the BBM mechanism satisfies the following properties:*

- (1) For any signal S_q , $\varphi_{S_q}(v) \geq 0$ for all v in the support of S_q .
- (2) $\text{CS}(\text{BBM}) = \text{OPT} \geq \Pr_{V \sim \mathcal{D}}[V < r_{\mathcal{D}}] \cdot \mathbb{E}_{V \sim \mathcal{D}}[V \mid V < r_{\mathcal{D}}] = \sum_{v < r_{\mathcal{D}}} v f_{\mathcal{D}}(v)$.

3 LOWER BOUND INSTANCES

We now prove Theorems 1.1 and 1.2. We note that though our lower bounds assume that the seller runs interim IR and ex-post IC mechanisms, and the intermediary can send arbitrary signals, the




Fig. 2. Illustrating the setting for Theorems 1.1 and 1.2: On the left (below the axis) we show the setting without signaling, where buyer 1 (blue) has values (a, b) and buyer 2 (red) has values (c, d) ; we also show the corresponding virtual values (above the axis). On the right, we show the two settings characterized by Theorem 3.2 under which a signal s has non-zero consumer surplus (the changed virtual values are highlighted).

same lower bounds hold even when the seller runs ex-post IC and IR mechanisms, provided we restrict the intermediary's signals to induce posteriors that are product distributions.

Our lower bounds are based on a 2-buyer instance illustrated in Fig. 2: given values $a > b > c > d$, buyer 1 has value $V_1 \in \{a, b\}$ with probabilities α and $1 - \alpha$ respectively, while buyer 2 has value $V_2 \in \{c, d\}$ with probabilities β and $1 - \beta$ respectively. We choose $\alpha a = b$ and $\beta c = d$; thus, the virtual values satisfy: $\varphi_1(a) = a$, $\varphi_2(c) = c$, and $\varphi_1(b) = \varphi_2(d) = 0$. Call this distribution $\mathcal{D}$.

Characterization of optimal signaling. By Theorem 2.1, we know any signal that correlates the buyers raises zero consumer surplus. Therefore, the only signals $\mathcal{S}$ of interest are those under which buyer values are independent. Abusing notation we denote such a signal as $s = (\alpha', \beta')$, where $\Pr[v_1 = a] = \alpha'$ and $\Pr[v_2 = c] = \beta'$. Note that in this instance, for a signal to get maximum welfare the resulting optimal mechanism must always award buyer 1, and for non-zero consumer surplus it must award the item to buyer 1 at price b , or buyer 2 at price d .

Let $\text{CS}(s)$ denote the consumer surplus under any such a signal s , and $\varphi_1(b|s)$ and $\varphi_2(d|s)$ denote the new virtual values (note that by definition, $\varphi_1(a|s) = a$ and $\varphi_2(c|s) = c$ under any signal s with $\alpha', \beta' > 0$). We can use Myerson's characterization (Section 2.1) to exhaustively characterize the resulting optimal mechanisms as a function of $(\varphi_1(b|s), \varphi_2(d|s))$:

PROPOSITION 3.1. *Conditioned on receiving a signal s , we have the following cases:*

- (1) If $\varphi_1(b) \geq c$, then the optimal mechanism is to sell to Buyer 1 at price b . $\text{CS}(s) = (a - b)\alpha'$.
- (2) If $\varphi_2(d) \geq \max(0, \varphi_1(b))$, then the optimal mechanism is to try selling to Buyer 1 at price a then to Buyer 2 at price d . $\text{CS}(s) = (1 - \alpha')\beta'(c - d)$.
- (3) If $\varphi_1(b) \leq 0$ and $\varphi_2(d) \leq 0$, then the optimal mechanism is to try selling to Buyer 1 at price a then to Buyer 2 at price c . $\text{CS}(s) = 0$.
- (4) If $0 \leq \varphi_1(b) \leq c$ and $\varphi_2(d) \leq \varphi_1(b)$, then the optimal mechanism is to sell to Buyer 1 at price b if Buyer 2 has valuation d ; otherwise, it tries selling to Buyer 1 at price a then to Buyer 2 at price c . $\text{CS}(s) = \alpha'(1 - \beta')(a - b)$.

Our main insight, however, is that the setting can be further simplified to get the following structural property for the optimal signaling scheme.

THEOREM 3.2 (STRUCTURAL THEOREM). *In an optimal signaling scheme, the only signals $s = (\alpha', \beta')$ that raise non-zero consumer surplus have the following form:*

- (1') Under signal s , $\varphi_1(b|s) = \varphi_2(c|s) = c$ and $\text{CS}(s) = \alpha'(a - b)$.
- (2') Under signal s , $\varphi_2(d|s) = \varphi_1(b|s) \geq 0$ and $\text{CS}(s) = \alpha'(1 - \beta')(a - b)$.

PROOF. Recall that we restrict ourselves to signals $\mathcal{S}$ under which the buyer valuations remain independent. Any such signal can be alternately written as $s = (\alpha', \beta')$ where $\alpha' = \Pr[v_1 = a]$ and

$\beta' = \Pr[v_2 = c]$. For ease of notation, we henceforth drop the conditioning of virtual valuations on signal s (i.e., write $\varphi(\cdot)$ for $\varphi(\cdot|s)$) when clear from context.

Next, let γ_s denote the weight of any signal $s = (\alpha', \beta')$. The signaling scheme that maximizes consumer surplus is the solution to the following linear program written over signals $s = (\alpha', \beta')$:

$$\begin{aligned}
 & \text{Maximize} && \sum_s \gamma_s \text{CS}(s) \\
 & \text{Subject to} && \sum_{s=(\alpha', \beta')} \gamma_s \alpha' \beta' \leq \alpha \beta \\
 & && \sum_{s=(\alpha', \beta')} \gamma_s \alpha' (1 - \beta') \leq \alpha (1 - \beta) \\
 & && \sum_{s=(\alpha', \beta')} \gamma_s (1 - \alpha') \beta' \leq (1 - \alpha) \beta \\
 & && \sum_{s=(\alpha', \beta')} \gamma_s (1 - \alpha') (1 - \beta') \leq (1 - \alpha) (1 - \beta) \\
 & && \gamma_s \geq 0 \quad \forall s
 \end{aligned} \tag{2}$$

We examine the cases in Proposition 3.1 with positive consumer surplus, and characterize the optimal solution:

- In Case (1), we have $\varphi_1(b) = c$. To see this, consider any signal s with $\varphi_1(b) > c$. Suppose we increase α' and decrease γ_s while preserving the product $\alpha' \gamma_s$. Since $\gamma_s \text{CS}(s) = \gamma_s \alpha' (a - b)$, this is preserved by the change. Therefore, the objective of LP (2) is preserved, and so are the first two constraints. Further, since $(1 - \alpha')$ and γ_s decrease, this only makes the third and fourth constraints more feasible. This transformation decreases $\varphi_1(b)$.
- In Case (2) and (4), we have $\varphi_2(d) = \varphi_1(b)$. It does not help to make them unequal by a similar argument as above: In case (2), if $\varphi_2(d) > \varphi_1(b)$, we can increase β' while preserving $\gamma_s \beta'$. Since $\gamma_s \text{CS}(s) = \gamma_s (1 - \alpha') \beta' (c - d)$, this does not change the contribution to the objective of LP (2), and preserves all constraints. This transformation decreases $\varphi_2(d)$. In case (4), if $\varphi_2(d) < \varphi_1(b)$, we can increase α' while preserving $\gamma_s \alpha'$. Since $\gamma_s \text{CS}(s) = \gamma_s \alpha' (1 - \beta') (a - b)$, this does not change the contribution to the objective of LP (2), and preserves all constraints. This transformation decreases $\varphi_1(b)$.

Therefore, the only two types of signals s that give positive CS are

- (1') If $\varphi_1(b) = c$, then $\text{CS}(s) = \alpha' (a - b)$.
- (2') If $\varphi_2(d) = \varphi_1(b) \geq 0$, then $\text{CS}(s) = \max((1 - \alpha') \beta' (c - d), \alpha' (1 - \beta') (a - b))$.

As $(1 - \beta')(b - d) \geq 0$, we have

$$\left(b - d + \frac{\beta'}{1 - \beta'} (c - d) \right) (1 - \beta') \geq \beta' (c - d).$$

Notice that in Case (2'), we have $\varphi_1(b) = b - \frac{\alpha'}{1 - \alpha'} (a - b) = d - \frac{\beta'}{1 - \beta'} (c - d) = \varphi_2(d)$. This gives

$$\alpha' (1 - \beta') (a - b) \geq (1 - \alpha') \beta' (c - d).$$

Thus, the two types of signals s that give positive CS become

- (1') If $\varphi_1(b) = c$, then $\text{CS}(s) = \alpha' (a - b)$.
- (2') If $\varphi_2(d) = \varphi_1(b) \geq 0$, then $\text{CS}(s) = \alpha' (1 - \beta') (a - b)$. □

Using the above structural theorem, the proofs of Theorems 1.1 and 1.2 follow by different choices of the parameters (a, b, c, d) . Suppose the virtual values of b and d are slightly above zero with $\varphi_1(b) > \varphi_2(d)$ so that Case (4) in Proposition 3.1 is uniquely optimal for the seller. The optimal auction generates consumer surplus $\text{CS}(\mathcal{D}) = \alpha(1 - \beta)(a - b) = \frac{b}{c} \cdot \frac{c - d}{a} \cdot (a - b)$ according to Proposition 3.1.

3.1 Proof of Theorem 1.1

To prove Theorem 1.1, we set $b \rightarrow c^+$.² Now in Proposition 3.1, in Case (1), we must have $\alpha' \rightarrow 0^+$ since $b \rightarrow c^+$, so that $CS \rightarrow 0$. Also if $\alpha' = 1$ in a signal then $CS = 0$ here. The only other signal where the item is allocated to the higher bidder is in Case (4) when $\beta' = 0$. Let γ denote the probability of the signal of this type $s = (\alpha', 0)$. (Having multiple signals of this form gives the same CS as having a single signal as their average.) Since $\varphi_1(b) \geq \varphi_2(d)$, we have $\alpha' \leq \frac{b-d}{a-d}$.

By the constraints of LP (2), we have:

$$\Pr[v_1 = b \wedge v_2 = d] = \gamma(1 - \alpha') \leq (1 - \alpha)(1 - \beta),$$

which simplifies to $\gamma \leq \frac{(a-d) \cdot (c-d)}{ac}$. The consumer surplus in this case is therefore:

$$CS = \gamma CS(s) = \gamma \alpha' (a - b) \leq \frac{(b-d) \cdot (c-d)}{ac} \cdot (a-b) \leq \frac{b-d}{b} \cdot CS(\mathcal{D}).$$

Setting $d = (1 - \frac{\varepsilon}{2})b$ and combining with the fact that $CS \rightarrow 0$ in Case (1), we have the consumer surplus of any efficient signaling, $CS \rightarrow \frac{\varepsilon}{2} \cdot CS(\mathcal{D})$ so $CS < \varepsilon \cdot CS(\mathcal{D})$.

3.2 Proof of Theorem 1.2

Without signaling, $E[\max v_i] = \alpha a + (1 - \alpha)b$ and $\mathcal{R}(\mathcal{D}) = \alpha a + (1 - \alpha)\beta c$. (Case (2), (3) and (4) in Proposition 3.1 give the same revenue $\mathcal{R}(\mathcal{D})$.) Therefore

$$\text{OPT} = E[\max v_i] - \mathcal{R}(\mathcal{D}) = (1 - \alpha)(b - \beta c).$$

Now we assign the values as: $\alpha = \beta = 1 - \delta$; $a = \frac{1}{(1-\delta)^2}$, $b = \frac{1}{1-\delta}$, $c = 1$, $d = 1 - \delta$ with $\delta \rightarrow 0^+$. Then we plug in the values and the two possible types of signals s in Theorem 3.2 become

(1') If $\alpha' = \frac{1-\delta}{2-\delta} < \frac{1}{2}$, then $CS(s) \leq \frac{1}{2}\delta(1 + o(1))$.

(2') If $\alpha' \leq 1 - \delta$; $\beta' = \frac{1-3(1-\alpha')+3\delta(1-\alpha')-\delta^2(1-\alpha')}{1-2(1-\alpha')+\delta(1-\alpha')} > \frac{1-3(1-\alpha')}{1-2(1-\alpha')}$, then $CS(s) \leq \alpha'(1 - \beta')\delta(1 + o(1))$.

The consumer surplus maximizing signaling scheme should use t signals $S_{2,i}$ of type (2'), with $\alpha'_{2,i}$, $\beta'_{2,i}$ and weight $w(S_{2,i})$. There is an additional signal S_1 (with weight $w(S_1)$) of type (1') with α'_1 and β'_1 . (Having multiple signals of type (1') gives the same CS as having a single signal as their average.) Denoting the valuation of the first buyer by v_1 and the second buyer by v_2 , the constraints in LP (2) imply the two constraints:

$$\Pr[v_1 = b] = (1 - \alpha'_1)w(S_1) + \sum_{i=1}^t (1 - \alpha'_{2,i}) \cdot w(S_{2,i}) \leq 1 - \alpha = \delta, \quad (\text{Constraint (I)})$$

$$\Pr[v_1 = b \wedge v_2 = d] = \sum_{i=1}^t (1 - \alpha'_{2,i})(1 - \beta'_{2,i}) \cdot w(S_{2,i}) \leq (1 - \alpha)(1 - \beta) = \delta^2. \quad (\text{Constraint (II)})$$

Note that $\text{OPT} = (1 - \alpha)(b - \beta c) = 2\delta^2(1 + o(1))$.

² $\varphi_1(b) - \varphi_2(d)$ and $\varphi_2(d)$ can be arbitrarily small as long as positive, so we take the limits for them first, i.e., we are calculating $\lim_{b \rightarrow c^+} \lim_{\varphi_2(d) \rightarrow 0^+, \varphi_1(b) \rightarrow \varphi_2(d)^+} CS$ in the following part of the proof. This allows us to treat $\alpha = \frac{b}{a}$ and $\beta = \frac{d}{c}$ in calculating $(1 - \alpha)(1 - \beta)$, as α and β are not infinitesimally close to 1 for any fixed ε . (We will define $d = (1 - \varepsilon/2)b$.)

The total consumer surplus therefore is:

$$\begin{aligned}
\text{CS} &\leq \frac{1}{2}\delta(1+o(1)) \cdot w(S_1) + \sum_{i=1}^t \alpha'_{2,i}(1-\beta'_{2,i})\delta(1+o(1)) \cdot w(S_{2,i}) \\
&\leq \frac{1}{2}\delta(1+o(1)) \cdot 2 \left(\delta - \sum_{i=1}^t (1-\alpha'_{2,i}) \cdot w(S_{2,i}) \right) + \delta(1+o(1)) \sum_{i=1}^t \alpha'_{2,i}(1-\beta'_{2,i}) \cdot w(S_{2,i}) \\
&= \delta^2(1+o(1)) + \delta(1+o(1)) \sum_{i=1}^t (\alpha'_{2,i}(1-\beta'_{2,i}) - (1-\alpha'_{2,i})) \cdot w(S_{2,i}) \\
&\leq \delta^2(1+o(1)) + \delta(1+o(1)) \sum_{i=1}^t (1-\alpha'_{2,i})(1-\beta'_{2,i}) \cdot w(S_{2,i}) \\
&\leq \delta^2(1+o(1)) + \delta(1+o(1)) \cdot \delta^2 = \delta^2(1+o(1)).
\end{aligned}$$

Here the second inequality follows from Constraint (I), and $\alpha'_1 < \frac{1}{2}$ by the condition of (1'). The third inequality uses the implication of $\varphi_2(d) = \varphi_1(b)$ that $\beta'_{2,i} > \frac{1-3(1-\alpha'_{2,i})}{1-2(1-\alpha'_{2,i})}$. The fourth inequality uses Constraint (II). This establishes a lower bound of 2, since $\text{OPT} = 2\delta^2(1+o(1))$.

3.3 Achieving the Pareto-Frontier in the I.I.D. Case

We now ask if there are cases where we can circumvent Theorem 1.1 and maximize welfare while ensuring at least as much surplus as Myerson's auction (that is, achieve a point on the line BO in Fig. 1b). Note that Theorem 1.1 rules this out for non-i.i.d. distributions. Surprisingly, however, for i.i.d. regular distributions $\mathcal{D}_i$, the following simple signaling scheme turns out to be sufficient for achieving point B in Fig. 1b. Morally this shows why our lower bound constructions are delicate.

Suppose the common reserve price of $\mathcal{D}_i$ is r , and the maximum value of any buyer is v_m .

- (1) If $v_m < r$, then the intermediary reveals v_m and the identity of the highest buyer to the seller, who then sells to this bidder at price v_m .
- (2) If $v_m \geq r$, the intermediary only reveals the information that some buyer has value $\geq r$ (but does not reveal either v_m or the identity of the highest bidder). In this case, though the posterior is not a product distribution anymore, it can be shown that the seller's optimal auction remains the second price auction with reserve r .³

It is easy to check that the item always sells to the highest buyer, and the CS is exactly the same as in Myerson's auction without signaling, thereby achieving point B . Note however that this scheme does not give any guarantees on approximating CS itself, since the surplus of Myerson's auction is not an approximation to OPT . The question of approximating OPT is much more challenging even for the i.i.d. regular setting, and this is what we focus on in the next section.

4 APPROXIMATING CONSUMER SURPLUS: THE I.I.D. CASE

In this section, we present our main result (Theorem 1.3): An $O(1)$ -approximation to OPT when the buyers' valuation distributions $\mathcal{D}_i$ are identical and regular.

Recall we start with a product distribution $\mathcal{D} = \mathcal{D}_1 \times \mathcal{D}_2 \times \cdots \times \mathcal{D}_n$. Let $x_1 < x_2 < \cdots < x_K$ be the supports of value distributions. $\mathcal{R}(\mathcal{D})$ denotes the revenue of the optimal auction (Algorithm 1) on $\mathcal{D}$, and $\text{OPT} = \mathbb{E}_{\tilde{v} \sim \mathcal{D}}[\max_i v_i] - \mathcal{R}(\mathcal{D})$.

³Note this is the only case when the seller gets non-zero revenue in the optimal auction for the original product distribution. Suppose for the purpose of contradiction that the seller can do better for this posterior, she can also do better than the optimal auction for the original distribution.

When $\mathcal{D}_i$'s are identical and regular, since the highest virtual-value and highest value buyers coincide, the optimal auction (Algorithm 1) assigns the item to the buyer with highest value if this value is above the common reserve price. Therefore, we decompose OPT into two components:

- Myerson's surplus: The CS generated by Myerson's auction, denoted by $\text{CS}(\mathcal{D})$.
- Non-allocation surplus: The loss in CS due to Myerson's auction not allocating the item, denoted by CS_0 .

We therefore have $\text{OPT} = \text{CS}(\mathcal{D}) + \text{CS}_0$. In the remainder, we will present an $O(1)$ -approximation for the non-allocation surplus CS_0 , which will imply Theorem 1.3 when $\text{CS}_0 \geq \text{CS}(\mathcal{D})$. (When $\text{CS}(\mathcal{D}) > \text{CS}_0$, sending no signal is already a 2-approximation.)

4.1 Preliminaries and Intuition

Our approximation bound for the non-allocation surplus will also apply when $\mathcal{D}_i$ are independent but not necessarily identical, which will be required for showing Theorem 1.4. Therefore, in the sequel, we will proceed assuming the more general case that $\mathcal{D}_i$'s are not necessarily identical, and derive signaling schemes that approximate the non-allocation surplus CS_0 for this case.

We denote a realization from $\mathcal{D}$ by $\vec{v} = \{v_i\}$. Let $p_i = \Pr_{v_i \sim \mathcal{D}_i}[v_i < r_{\mathcal{D}_i}]$ for any buyer i , where $r_{\mathcal{D}_i}$ is the reserve price of $\mathcal{D}_i$. Let $Y_i = \mathcal{D}_{i|<r_{\mathcal{D}_i}}$ denote the distribution of $\mathcal{D}_i$ conditioned on being smaller than $r_{\mathcal{D}_i}$. Suppose we draw a sample independently from each distribution Y_i . Let Z_ℓ denote the distribution for the ℓ^{th} largest value among these n draws.

We first derive an expression for CS_0 .

LEMMA 4.1. *Let $P = \prod_{i=1}^n p_i$. Then, $\text{CS}_0 = P \cdot \mathbb{E}[Z_1]$.*

PROOF. Note that CS_0 is the expected surplus lost due to not allocating the item in Myerson's mechanism. This happens only when all realized values are below their corresponding reserve price. In this case, the value lost is the maximum valuation, since this value contributes to the welfare, and the revenue raised is zero. Therefore, we have:

$$\text{CS}_0 = \left(\prod_{i=1}^n p_i \right) \cdot \mathbb{E} \left[\max_{i=1,2,\dots,n} v_i \mid \forall i, v_i < r_{\mathcal{D}_i} \right]$$

where the expectation is over $\vec{v} \sim \mathcal{D}$. This is equal to $P \cdot \mathbb{E}[Z_1]$. $\square$

Vanilla Signaling Schemes. Before presenting our general signaling scheme, it is instructive to first consider simpler schemes to understand the challenges posed by this problem. One possible scheme is to pick a random buyer and apply the single-buyer BBM signaling scheme defined in Section 2 to it. Denote this buyer by i . Such a single-buyer signaling scheme will construct the set of BBM signals for buyer i by decomposing $\mathcal{D}_i$ and pretending the other buyers don't exist. Given the valuation $v_i \sim \mathcal{D}_i$ of this buyer, the scheme will send a signal $\text{BBM}(v_i, \mathcal{D}_i)$ just as in single-buyer case, and reveal the identity of this buyer. There is no signal sent for the other buyers, so that the seller's information for $j \neq i$ is their prior $\mathcal{D}_j$.

The nice property of the BBM signaling scheme is that the virtual value of buyer i is always non-negative (Lemma 2.3). Therefore, in the event when all buyers $j \neq i$ have values $v_j < r_{\mathcal{D}_j}$ (so that their virtual values are negative), the seller allocates the item to i . Since v_i is independent of other buyers' values, the mechanism therefore behaves exactly as $\text{BBM}(v_i, \mathcal{D}_i)$ from the perspective of buyer i . In other words, with probability $\prod_{j \neq i} p_j$, we generate the single-buyer CS from Lemma 2.3, which is at least:

$$\text{CS}(\text{BBM}) \text{ from buyer } i \geq \Pr[v_i < r_{\mathcal{D}_i}] \mathbb{E}[v_i \mid v_i < r_{\mathcal{D}_i}] = p_i \mathbb{E}[Y_i].$$

Since buyer i is a randomly chosen buyer, the overall CS generated is:

$$\text{Overall CS} \geq \frac{1}{n} \sum_{i=1}^n \left(\prod_{j \neq i} p_j \right) p_i \mathbb{E}[Y_i] = \left(\prod_{j=1}^n p_j \right) \frac{\sum_{i=1}^n \mathbb{E}[Y_i]}{n} = P \frac{\sum_{\ell=1}^n \mathbb{E}[Z_\ell]}{n} \geq P \frac{\mathbb{E}[Z_1]}{n}$$

where $P = \prod_{j=1}^n p_j$. Therefore, comparing the expression above with that in Lemma 4.1, this scheme yields an n -approximation to CS_0 . Further, we can construct identical regular distributions for which the expected value of the max is comparable to the expected value of the sum, that is, $\sum_{\ell=1}^n \mathbb{E}[Z_\ell] = O(\mathbb{E}[Z_1])$. Therefore, this analysis cannot be improved.

On the other hand, the above scheme does achieve $\text{CS} = \text{OPT}$ for two-valued i.i.d. distributions ($\mathcal{K} = 2$ and arbitrary n). To see this, assume the support is $a < b$, and let $q = \Pr[\mathcal{D}_i = a]$. If the reserve price is a , Myerson's auction is already efficient, that is, $\text{CS}(\mathcal{D}) = \text{OPT}$; else Myerson's auction has $\text{CS}(\mathcal{D}) = 0$. We now have $Z_i = a$ for all i and $p_i = q$, so that the above scheme has surplus $q^n a = q^n \mathbb{E}[Z_1] = \text{CS}_0 = \text{OPT}$. Therefore, in either case, we extract $\text{CS} = \text{OPT}$. Interestingly, this also shows that our lower bounds in Theorems 1.1 and 1.2 do require non-i.i.d. distributions when each $\mathcal{K}_i = 2$ regardless of the number n of buyers.

Moving beyond the $\mathcal{K} = 2$ setting to general $\mathcal{K}$ and n , it is tempting to run the BBM signaling scheme directly on the buyer with highest value, hoping to extract surplus $P \mathbb{E}[Z_1]$. However, this requires revealing the identity of the highest buyer to the seller, since the signaling scheme itself is public knowledge. But if the seller knows the identity of the highest buyer, she can always increase the reserve price to be the second highest bid. In other words, the posterior of the highest buyer is *truncated* at Z_2 . This case needs a more careful construction of the signal and analysis, since the event of a buyer being the largest and hence BBM being applied to it is now correlated with the surplus this buyer generates in BBM. We perform this construction and analysis in Lemma 4.3. Intuitively, signaling using BBM on the largest buyer will only yield $\text{CS} \geq P \mathbb{E}[Z_1 - Z_2]$, which is again an $\Omega(n)$ -approximation to $P \mathbb{E}[Z_1]$.

Our signaling scheme in the next section chooses a middle ground between these extremes – we will choose a rank $t \in \{1, 2, \dots, n\}$ carefully, and choose a buyer whose value lies in the top t ranks at random. We will then perform the single-buyer BBM scheme on this buyer as we describe below. Surprisingly, this improves the naïve n -approximation to an $O(1)$ -approximation!

4.2 Ranking-Based Multi-Buyer Signaling Scheme

We now introduce the family of signaling schemes RANK_t . We will derive a lower bound for the CS obtained by these schemes in our key lemma, Lemma 4.3. As mentioned before, since this scheme will also form the basic subroutine for the non-i.i.d. case (Theorem 1.4), we present this scheme assuming $\mathcal{D}_i$ can be non-identical.

Recall the definitions of $p_i, r_{\mathcal{D}_i}, Y_i, Z_\ell$ from above. In order to define the signaling scheme, we need an additional definition. For agent i with $V_i \sim \mathcal{D}_i$, we use $\mathcal{D}_{i|>a}$ to denote the conditional distribution of V_i given $V_i > a$, and $\mathcal{D}_{i|<a}$ to denote the conditional distribution of V_i given $V_i < a$. Moreover, we use $\mathcal{D}_{i|>a} - b$ to denote the distribution of $V_i - b$ given $V_i > a$; we refer to it as the distribution of V_i truncated at a and reduced by b .

The following result relates the reserve price of the truncated and the original distributions.

LEMMA 4.2 (PROVED IN FULL PAPER [1]). *Let $\mathcal{D}'_i = \mathcal{D}_{i|>v^\circ} - v^\circ$ for any v° . Then we have $r_{\mathcal{D}'_i} \geq r_{\mathcal{D}_i} - v^\circ$, and moreover, for any $v > v^\circ$, we have $\varphi_{\mathcal{D}'_i}(v - v^\circ) = \varphi_{\mathcal{D}_i}(v) - v^\circ$.*

Note that $\mathcal{D}'_i = \mathcal{D}_{i|>v^\circ} - v^\circ$ represents buyer i 's *excess value* compared to v° . Lemma 4.2 shows that for any threshold v° and any buyer i , given the side-information that $V_i > v^\circ$, her new reserve price is greater than her original reserve price.

The RANK_t signaling scheme. We now present the family of signaling schemes RANK_t parameterized by the rank $t \in \{1, \dots, n\}$. For any realized joint valuation profile $\vec{v} = (v_1, v_2, \dots, v_n)$, the signal sent by RANK_t consists of two parts. In the first part, RANK_t observes $\vec{v}$ and outputs (v°, T) , where v° is the value of $(t+1)^{\text{st}}$ largest realized value (or 0 when $t = n$), and T is the subset of buyers with realized value strictly greater than v° . For the second part of the signal, RANK_t chooses a buyer j uniformly at random from T , and computes her excess distribution $\mathcal{D}_{j|>v^\circ} - v^\circ$. It then reveals both the identity of j , as well as the signal $\text{BBM}(v_j - v^\circ, \mathcal{D}_{j|>v^\circ} - v^\circ)$ generated by the single-buyer BBM scheme on a buyer with value distribution $\mathcal{D}_{j|>v^\circ} - v^\circ$. The scheme is formalized in Algorithm 2.

Optimal mechanism under RANK_t . Conditioned on receiving the signal generated by RANK_t , the seller is guaranteed a revenue of v° from the $(t+1)^{\text{st}}$ largest buyer, and knows that only buyers in T can pay more than v° . The seller can now charge at least v° to any buyer in T , and can further run an auction over the excess value of buyers in T , where for buyer $i \in T$, her excess value has distribution $\mathcal{D}'_i = \mathcal{D}_{i|>v^\circ} - v^\circ$. Note that for any buyer $i \in T$ except the randomly chosen buyer j , a value drawn from $\mathcal{D}'_i$ represents how much more than v° she is willing to pay. Moreover, distributions $\mathcal{D}'_i$ are independent, and also, since the identity of j is chosen uniformly at random, the BBM scheme modifies the distribution of buyer j in a fashion that is independent of $\mathcal{D}'_i$.

By Lemma 2.3, we know that the BBM scheme ensures the virtual value of buyer j is always non-negative. From the characterization of the optimal auction [13, 19], since the item is always allocated to the highest virtual value buyer as long as this value is non-negative, the item will always be allocated to buyer j if all other buyers $i \in T, i \neq j$ have excess values $v_i - v^\circ < r_{\mathcal{D}'_i}$ (and hence, negative virtual values).

ALGORITHM 2: $\text{RANK}_t(\vec{v}, \mathcal{D})$

```

 $v^\circ \leftarrow (t+1)^{\text{st}}$  largest value in  $\vec{v}$ 
 $T \leftarrow \{i : v_i > v^\circ\}$ 
if  $T \neq \emptyset$  then
     $j \leftarrow$  Buyer chosen uniformly at random from  $T$ 
     $s \leftarrow \text{BBM}(v_j - v^\circ, \mathcal{D}_{j|>v^\circ} - v^\circ)$ 
    return  $v^\circ, T, j$ , and  $s$ 
end
else
    return  $v^\circ$ , and  $T = \emptyset$ 
end

```

Consumer surplus under RANK_t . We require the following key lemma, that gives a lower bound for the consumer surplus generated under RANK_t . This lemma forms the crux of our subsequent analysis, helping us quantify how the BBM signal recovers much of CS lost by Myerson's auction. The difficulty in proving it arises because the random choice of buyer j in Algorithm 2 is correlated with its rank, which in turn is correlated with its winning the auction and the surplus it generates in BBM. We get around this correlation by carefully *coupling* the surplus generated when buyer values are above the reserve with the order statistics of buyer valuations below the reserve.

LEMMA 4.3. *For $1 \leq t \leq n$, the consumer surplus of RANK_t satisfies:*

$$\text{CS}(\text{RANK}_t) \geq \left(\prod_{i=1}^n p_i \right) \cdot \left(\left(\frac{1}{t} \cdot \sum_{\ell=1}^t \mathbb{E}[Z_\ell] \right) - \mathbb{E}[Z_{t+1}] \right).$$

PROOF. For convenience, denote $r_i = r_{\mathcal{D}_i}$. Fix a buyer b , and any valuation profile $\vec{v}_{-b} = \{v_i, i \neq b\}$ such that $v_i < r_i \forall i \neq b$. Define v_b^t as the t^{th} largest value in $\{v_i, i \neq b\}$. Now consider the event

$$Q(\vec{v}_{-b}, b, t) = \{V_b > v_b^t \text{ AND } b \text{ selected for BBM signaling}\}.$$

Conditioned on $Q(\vec{v}_{-b}, b, t)$, we have that the RANK_t scheme (Algorithm 2) with parameter t sets threshold value as $v^\circ = v_b^t$. By Lemma 4.2, we have that for every $i \in T, i \neq b$, their value v_i is smaller than their new reserve price $v^\circ + r_{\mathcal{D}'_i}$, since $v_i < r_{\mathcal{D}_i}$, and modifying $\mathcal{D}_i$ to $\mathcal{D}'_i = \mathcal{D}_{i|>v^\circ - v^\circ}$ does not decrease the reserve price. Therefore, conditioned on $Q(\vec{v}_{-b}, b, t)$, the auction behaves like the single item mechanism $\text{BBM}(v'_b, \mathcal{D}'_b)$, where $v'_b = v_b - v_b^t$, $\mathcal{D}'_b = \mathcal{D}_{b|>v'_b - v'_b}$. Let r'_b denote the reserve price of $\mathcal{D}'_b$; again using Lemma 4.2 we have $r'_b \geq r_b - v_b^t$. Now, using Lemma 2.3, we get that the expected consumer surplus generated by RANK_t under $Q(\vec{v}_{-b}, b, t)$ is at least:

$$\mathbb{E}[\text{CS}(\text{RANK}_t) \mid Q(\vec{v}_{-b}, b, t)] \geq \sum_{v'_b < r'_b} v'_b \Pr[\mathcal{D}'_b = v'_b] \geq \sum_{v'_b < v_b < r_b} (v_b - v_b^t) \frac{f_{\mathcal{D}_b}(v_b)}{S_{\mathcal{D}_b}(v_b^t)}.$$

Note also that $\Pr[Q(\vec{v}_{-b}, b, t)] = \frac{1}{t} \cdot \left(\prod_{i \neq b} f_{\mathcal{D}_i}(v_i) \mathbb{1}_{\{v_i < r_{\mathcal{D}_i}\}} \right) \cdot S_{\mathcal{D}_b}(v_b^t)$. Thus for any buyer b , and any valuation profile $\vec{v}_{-b}$ with $v_i < r_i$ for all $i \neq b$, we have

$$\begin{aligned} \mathbb{E}[\text{CS}(\text{RANK}_t) \cdot \mathbb{1}_{Q(\vec{v}_{-b}, b, t)}] &= \Pr[Q(\vec{v}_{-b}, b, t)] \cdot \mathbb{E}[\text{CS}(\text{RANK}_t) \mid Q(\vec{v}_{-b}, b, t)] \\ &\geq \frac{1}{t} \left(\prod_{i \neq b} f_{\mathcal{D}_i}(v_i) \right) \sum_{v'_b < v_b < r_b} (v_b - v_b^t) f_{\mathcal{D}_b}(v_b) \\ &= \frac{1}{t} \sum_{v_b < r_b} \left(\prod_i f_{\mathcal{D}_i}(v_i) \right) \max\{(v_b - v_b^t), 0\}. \end{aligned}$$

For any $\vec{v}$ let $v^{(t)}$ denote the t^{th} largest value, and $I_t(\vec{v})$ to be the indices corresponding to the top t values in $\vec{v}$. Summing up over all b , and all $\vec{v}_{-b}$ such that $v_i < r_i \forall i \neq b$, we have

$$\begin{aligned} \sum_b \sum_{\vec{v}_{-b}} \mathbb{E}[\text{CS}(\text{RANK}_t) \cdot \mathbb{1}_{Q(\vec{v}_{-b}, b, t)}] &\geq \sum_{\vec{v} | v_i < r_i} \frac{1}{t} \cdot \left(\prod_i f_{\mathcal{D}_i}(v_i) \right) \cdot \left(\sum_b \max\{(v_b - v_b^t), 0\} \right) \\ &= \sum_{\vec{v} | v_i < r_i} \left(\prod_i f_{\mathcal{D}_i}(v_i) \right) \left(\sum_{i \in I_t(\vec{v})} \frac{1}{t} (v_i - v^{(t+1)}) \right) \\ &= \sum_{\vec{v} | v_i < r_i} \left(\prod_i f_{\mathcal{D}_i}(v_i) \right) \left(\left(\sum_{i \in I_t(\vec{v})} \frac{v_i}{t} \right) - v^{(t+1)} \right). \end{aligned} \quad (3)$$

Let $\mathcal{D}_i'' = \mathcal{D}_{i|<r_{\mathcal{D}_i}}$ be the distribution of buyer i 's value conditioned on $V_i < r_{\mathcal{D}_i}$. Recall we define $p_i = \Pr_{v_i \sim \mathcal{D}_i}[v_i < r_{\mathcal{D}_i}]$; thus $f_{\mathcal{D}_i''}(v) = f_{\mathcal{D}_i}(v)/p_i$. Suppose we independently sample $Y_i \sim \mathcal{D}_{i|<r_{\mathcal{D}_i}}$ for each i , and define Z_ℓ as the ℓ^{th} largest value in $\{Y_i\}$. Then Eq. (3) can be written as

$$\begin{aligned} \sum_b \sum_{\vec{v}_{-b} | v_i < r_i \forall i \neq b} \mathbb{E}[\text{CS}(\text{RANK}_t) \cdot \mathbb{1}_{Q(\vec{v}_{-b}, b, t)}] &\geq \sum_{\vec{v} | v_i < r_i} \left(\prod_i p_i \right) f_{\mathcal{D}''}(\vec{v}) \left(\left(\sum_{i \in I_t(\vec{v})} \frac{v_i}{t} \right) - v^{(t+1)} \right) \\ &= \left(\prod_i p_i \right) \left(\left(\sum_{\ell=1}^t \frac{\mathbb{E}[Z_\ell]}{t} \right) - \mathbb{E}[Z_{t+1}] \right). \end{aligned}$$

Finally, noting that the $Q(\vec{v}_{-b}, b, t)$ events are all non-overlapping, we can write

$$\text{CS}(\text{RANK}_t) \geq \sum_b \sum_{\vec{v}_{-b} | v_i < r_i \forall i \neq b} \mathbb{E}[\text{CS}(\text{RANK}_t) \cdot \mathbb{1}_{Q(\vec{v}_{-b}, b, t)}],$$

thereby completing the proof. $\square$

4.3 Approximating Non-allocation Surplus CS_0 and Proof of Theorem 1.3

Given the above signaling scheme, approximating CS_0 (and hence showing Theorem 1.3) is now simple: We choose the parameter $t \in \{1, 2, \dots, n\}$ that maximizes $\text{CS}(\text{RANK}_t)$ and run RANK_t . We denote this scheme as $\{S_0(\vec{v}, \mathcal{D})\}_{\vec{v}}$, and present it in Algorithm 3.

ALGORITHM 3: $S_0(\vec{v}, \mathcal{D})$

Choose $t = \arg\max_{t'=1}^n \text{CS}(\text{RANK}_{t'})$. **return** $\text{RANK}_t(\vec{v}, \mathcal{D})$

The following theorem shows that this scheme approximates CS_0 when the $\mathcal{D}_i$ are a common regular distribution Θ . Since $\text{OPT} = \text{CS}(\mathcal{D}) + \text{CS}_0$ when $\mathcal{D}_i$ are regular and identical, Theorem 4.4 shows the better of no signaling and Algorithm 3 is an $O(1)$ -approximation to OPT , completing the proof of Theorem 1.3.

THEOREM 4.4. *The consumer surplus of the signaling scheme $S_0(\vec{v}, \mathcal{D})$ is an $O(1)$ -approximation to the non-allocation surplus, CS_0 .*

Let ALG denote the CS of $S_0(\vec{v}, \mathcal{D})$. The above theorem will follow from the following, since ALG is at least the LHS by Lemma 4.3 and CS_0 is equal to $P \cdot \mathbb{E}[Z_1]$ by Lemma 4.1.

THEOREM 4.5. *When each $\mathcal{D}_i$ is i.i.d. and regular with common distribution Θ , then we have:*

$$P \cdot \max_t \left(\left(\frac{1}{t} \sum_{i=1}^t \mathbb{E}[Z_i] \right) - \mathbb{E}[Z_{t+1}] \right) \geq \frac{1}{1900} \cdot P \cdot \mathbb{E}[Z_1].$$

The rest of this section is devoted to proving Theorem 4.5. Note that the above theorem is not true if the distribution Y (the conditional distribution of Θ below its reserve) on which the order statistics Z_i are defined, is a generic regular distribution such as an Exponential distribution. Consider the following example:

Example 4.6. Suppose Y is $\text{Exponential}(1)$. Then $\mathbb{E}[Z_i] = \sum_{j=i}^n \frac{1}{j} = H_n - H_{i-1}$, where $H_i = \sum_{j=1}^i \frac{1}{j}$. Therefore $\text{CS}_0 = \mathbb{E}[Z_i] = H_n$. It is easy to check that $(\frac{1}{t} \sum_{i=1}^t \mathbb{E}[Z_i]) - \mathbb{E}[Z_{t+1}] = 1 + H_n - H_t - (H_n - H_t) = 1$. Therefore, $\text{CS}_0 = \Omega(\log n) \cdot \text{ALG}$.

Note however that Y is of a more specific form: It is the conditional distribution of a regular distribution Θ below its reserve r . This means in particular that Y *cannot* be an Exponential distribution (or its discrete counterpart, the Geometric distribution). We will crucially use the property that the revenue of Θ when the price is set below the reserve r is a concave function of the quantile of the price, and further, this function is non-decreasing. Formally, we use the following.

LEMMA 4.7 (PROVED IN THE FULL PAPER [1]). *For any regular distribution Θ with $S(v) = \Pr[\Theta \geq v]$ and reserve r , and for any $u \leq v < r$ in the support of Θ , we have: (1) $u \cdot S(u) \geq \frac{1-S(u)}{1-S(v)} \cdot v \cdot S(v)$, and (2) $u \cdot S(u) \leq v \cdot S(v)$.*

The next idea in the analysis is to decompose the distribution Y into the “core” and “tail”, so that values v with $\Pr[Y \geq v] = O(1/n)$ lie in the tail, and the rest lie in the core. Roughly speaking, we will show that the expected value of Y in either the core or the tail is upper bounded by $\text{CS}(\text{RANK}_t)$ for suitable choices of t . Such choices are non-trivial and form the crux of our analysis.

4.3.1 Core-Tail Decomposition. To fix notation, recall that every $\mathcal{D}_i$ is the same regular distribution Θ . The values $x_1 < x_2 < \dots < x_K$ form the support of Θ and its reserve price is r . Recall $S(x) = \Pr[\Theta \geq x]$ and Y is the conditional distribution of Θ strictly below r .

First, if $n \leq 1000$, $\text{ALG} \geq \frac{1}{1000} \mathbb{E}[Z_1] = \frac{1}{1000} \text{CS}_0$. Therefore, we assume $n > 1000$ in the following analysis. We further assume $r > x_1$ since otherwise $\text{CS}_0 = 0$ and Myerson's auction itself raises optimal CS. Suppose $r = x_i$ with $i > 1$, then $\{x_1, \dots, x_{i-1}\}$ is the support of Y .

Let u^+ be the smallest value in the support of Y with $\Pr[Y \geq u^+] \leq \frac{4}{n}$ and let u^- be the largest value in the support of Y with $\Pr[Y \geq u^-] > \frac{4}{n}$. Both of these values exist since $\Pr[Y \geq x_1] = 1$ (we assumed $r > x_1$) and $\Pr[Y \geq r] = 0$. Further, u^- and u^+ are consecutive values in the support of Y .

We divide CS_0 into CORE and TAIL, based on whether there is a realized value at least u^+ .

Definition 4.8. $\text{CORE} := P \cdot u^-$, and $\text{TAIL} := P \cdot n \cdot \mathbb{E}[Y \cdot \mathbb{1}(Y \geq u^+)]$.

LEMMA 4.9. $\text{CS}_0 \leq \text{CORE} + \text{TAIL}$.

PROOF. This comes immediately from

$$\begin{aligned} \text{CS}_0 &= P \cdot (\mathbb{E}[Z_1 \cdot \mathbb{1}(Z_1 \leq u^-)] + \mathbb{E}[Z_1 \cdot \mathbb{1}(Z_1 \geq u^+)]) \\ &\leq P \cdot (u^- \cdot \Pr[Z_1 \leq u^-] + n \cdot \mathbb{E}[Y \cdot \mathbb{1}(Y \geq u^+)]) \\ &\leq \text{CORE} + \text{TAIL}. \end{aligned}$$

where the first inequality upper bounds the max of n i.i.d. samples from $\mathbb{1}(Y \geq u^+)$ by their sum. $\square$

4.3.2 Upper Bound on CORE and TAIL. We will separately bound CORE and $\text{TAIL} - O(1) \cdot \text{CORE}$ in terms of ALG. We first bound TAIL using the following lemma.

LEMMA 4.10. $\text{ALG} \geq \frac{1}{250} \cdot (\text{TAIL} - 4 \cdot \text{CORE})$.

PROOF. If $\Pr[Y \geq u^+] = 0$, then the inequality trivially holds since $\text{TAIL} = 0$. Otherwise,

$$\begin{aligned} \text{ALG} &\geq \text{CS}(\text{RANK}_1) = P \cdot \mathbb{E}[Z_1 - Z_2] \\ &\geq P \cdot \Pr[Z_1 \geq u^+ \wedge Z_2 \leq u^-] \cdot \mathbb{E}[Z_1 - Z_2 \mid Z_1 \geq u^+ \wedge Z_2 \leq u^-] \\ &\geq P \cdot \Pr[Z_1 \geq u^+ \wedge Z_2 \leq u^-] \cdot (\mathbb{E}[Z_1 \mid Z_1 \geq u^+ \wedge Z_2 \leq u^-] - u^-). \end{aligned}$$

Now, to bound the term $\Pr[Z_1 \geq u^+ \wedge Z_2 \leq u^-]$, we have the following. Note here u^- and u^+ are consecutive values in the support of Y .

$$\begin{aligned} \Pr[Z_1 \geq u^+ \wedge Z_2 \leq u^-] &= \Pr[Z_1 \geq u^+] \cdot \frac{\Pr[Z_1 \geq u^+ \wedge Z_2 \leq u^-]}{\Pr[Z_1 \geq u^+]} \\ &= \Pr[Z_1 \geq u^+] \cdot \frac{n \cdot (1 - \Pr[Y \leq u^-]) \cdot \Pr[Y \leq u^-]^{n-1}}{1 - \Pr[Y \leq u^-]^n} \\ &= \Pr[Z_1 \geq u^+] \cdot \frac{n}{\sum_{i=0}^{n-1} \Pr[Y \leq u^-]^i} \cdot \Pr[Y \leq u^-]^{n-1} \\ &\geq \Pr[Z_1 \geq u^+] \cdot \Pr[Y \leq u^-]^{n-1} \\ &\geq \Pr[Y \leq u^-]^n = \Pr[Z_1 \geq u^+] \cdot \Pr[Z_1 \leq u^-]. \end{aligned}$$

We now bound $\Pr[Z_1 \geq u^+]$ and $\Pr[Z_1 \leq u^-]$ separately. To bound $\Pr[Z_1 \geq u^+]$, note that

$$\Pr[Z_1 \geq u^+] = 1 - (1 - \Pr[Y \geq u^+])^n \geq 1 - \exp(-n \cdot \Pr[Y \geq u^+]) \geq \frac{1 - e^{-4}}{4} \cdot n \cdot \Pr[Y \geq u^+]$$

where the final inequality follows since the function $\psi(x) = \frac{1 - e^{-4x}}{(1 - e^{-4})x} \geq 1$ when $x = \frac{1}{4} \cdot n \cdot \Pr[Y \geq u^+] \leq 1$. To bound $\Pr[Z_1 \leq u^-]$, note that:

$$\Pr[Z_1 \leq u^-] = (1 - \Pr[Y \geq u^+])^n \geq (1 - 4/n)^n \geq 0.018$$

where we have used $n \geq 1000$ in the last inequality. Therefore, we have

$$\Pr[Z_1 \geq u^+ \wedge Z_2 \leq u^-] \geq \Pr[Z_1 \geq u^+] \cdot \Pr[Z_1 \leq u^-] \geq 0.004 \cdot n \cdot \Pr[Y \geq u^+].$$

To bound ALG, we also need to simplify the term $\mathbb{E}[Z_1 \mid Z_1 \geq u^+ \wedge Z_2 \leq u^-]$. For this, let $Y_1, \dots, Y_n$ be independent draws from Y . We have:

$$\begin{aligned} \mathbb{E}[Z_1 \mid Z_1 \geq u^+ \wedge Z_2 \leq u^-] &= \sum_{i=1}^n \frac{1}{n} \cdot \mathbb{E}[Y_i \mid Y_i \geq u^+ \wedge Y_j \leq u^-, \forall j \neq i] \\ &= \mathbb{E}[Y_1 \mid Y_1 \geq u^+ \wedge Y_j \leq u^-, \forall j \neq 1] \\ &= \mathbb{E}[Y_1 \mid Y_1 \geq u^+] = \mathbb{E}[Y \mid Y \geq u^+]. \end{aligned}$$

Here, the first equality follows since any Y_i is equally likely to be the maximum value, and the third equality follows by the independence of Y_i 's. Putting all this together, we bound ALG as:

$$\begin{aligned} \text{ALG} &\geq P \cdot \Pr[Z_1 \geq u^+ \wedge Z_2 \leq u^-] \cdot (\mathbb{E}[Z_1 \mid Z_1 \geq u^+ \wedge Z_2 \leq u^-] - u^-) \\ &\geq 0.004P \cdot n \cdot \Pr[Y \geq u^+] \cdot (\mathbb{E}[Y \mid Y \geq u^+] - u^-) \\ &= 0.004P \cdot (n \cdot \mathbb{E}[Y \cdot \mathbb{1}(Y \geq u^+)] - n \cdot \Pr[Y \geq u^+] \cdot u^-) \\ &\geq 0.004 \cdot (\text{TAIL} - 4 \cdot \text{CORE}). \end{aligned}$$

where the last inequality uses $n \cdot \Pr[Y \geq u^+] \leq 4$. □

We next bound CORE and the proof is relegated to the full paper [1]. This proof will crucially use the regularity of Θ (via Lemma 4.7).

LEMMA 4.11 (PROVED IN FULL PAPER [1]). $\text{ALG} \geq \frac{1}{330} \cdot \text{CORE}$.

4.3.3 *Completing Proof of Theorem 4.5.* Using Lemma 4.9, Lemma 4.11, and Lemma 4.10 we obtain:

$$\text{CS}_0 \leq \text{CORE} + \text{TAIL} = 5 \cdot \text{CORE} + (\text{TAIL} - 4 \cdot \text{CORE}) \leq 1650 \cdot \text{ALG} + 250 \cdot \text{ALG} = 1900 \cdot \text{ALG}.$$

This completes the proof of Theorem 4.5 and hence, Theorem 4.4 and Theorem 1.3.

5 CONCLUSION AND OPEN QUESTIONS

Note that our RANK_t mechanism can be viewed as a screening procedure – the intermediary only allows a fixed number of high-value bidders to bid. When the intermediary is an independent (typically, governmental) agency, such screening would map to “pre-certifying” bidders entering into private auctions. Similarly, when real-estate agencies have agents representing both sellers and buyers, they could (and often do) recommend a particular listing only to a chosen set of buyers based on better knowing their utilities. Therefore, as a side-effect, our procedures yield realistic mechanisms for an intermediary to increase surplus for both buyers and the seller.

In terms of open questions, beyond improving the lower and upper bounds in our specific setting (both existence and computational), it would be interesting to explore the equilibria in optimal auctions when the intermediary can send different signals to the seller and to the buyers, much like in [3, 21]. At an even higher level, our work can be considered a special case of a larger problem of information intermediaries for multi-agent mechanisms. As mentioned before, in our case, the optimal auction is the mechanism, and the intermediary can change the information to this mechanism in order to achieve “fairness” between producer and consumer surplus. It would be interesting to explore the question of achieving fairness by selectively regulating information to a black-box optimizer or mechanism in more general settings.

Acknowledgment. This work is supported by NSF grants ECCS-1847393, DMS-1839346, CNS-1955997, and CCF-2113798, and and ONR award N00014-19-1-2268. Reza Alijani is now at Google.

REFERENCES

- [1] Reza Alijani, Siddhartha Banerjee, Kamesh Munagala, and Kangning Wang. 2020. The Limits of an Information Intermediary in Auction Design. *CoRR* abs/2009.11841 (2020). <https://arxiv.org/abs/2009.11841>
- [2] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. 2015. The limits of price discrimination. *American Economic Review* 105, 3 (2015), 921–957.
- [3] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. 2017. First-Price Auctions With General Information Structures: Implications for Bidding and Revenue. *Econometrica* 85, 1 (2017), 107–143.
- [4] Dirk Bergemann and Stephen Morris. 2019. Information design: A unified perspective. *Journal of Economic Literature* 57, 1 (2019), 44–95.
- [5] Dirk Bergemann and Martin Pesendorfer. 2007. Information structures in optimal auctions. *Journal of Economic Theory* 137, 1 (2007), 580–609.
- [6] Yang Cai, Federico Echenique, Hu Fu, Katrina Ligett, Adam Wierman, and Juba Ziani. 2020. Third-Party Data Providers Ruin Simple Mechanisms. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 4, 1 (2020), 1–31.
- [7] Archishman Chakraborty and Rick Harbaugh. 2014. Persuasive puffery. *Marketing Science* 33, 3 (2014), 382–400.
- [8] Jacques Crémer and Richard P. McLean. 1988. Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions. *Econometrica* 56, 6 (1988), 1247–1257.
- [9] Rachel Cummings, Nikhil R Devanur, Zhiyi Huang, and Xiangning Wang. 2020. Algorithmic price discrimination. In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*. SIAM, 2432–2451.
- [10] Shaddin Dughmi. 2017. Algorithmic information structure design: A survey. *ACM SIGecom Exchanges* 15, 2 (2017), 2–24.
- [11] Shaddin Dughmi, David Kempe, and Ruixin Qiang. 2016. Persuasion with limited communication. In *Proceedings of the 17th ACM Conference on Economics and Computation, EC*. 663–680.
- [12] Shaddin Dughmi and Haifeng Xu. 2019. Algorithmic bayesian persuasion. *SIAM J. Comput.* (2019), 16–68.
- [13] Edith Elkind. 2007. Designing and learning optimal finite support auctions. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*. SIAM, 736–745.
- [14] Hu Fu, Christopher Liaw, Pinyan Lu, and Zhihao Gavin Tang. 2018. The Value of Information Concealment. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*. 2533–2544.
- [15] Nima Haghpahan and R. Siegel. 2022. The limits of multi-product price discrimination. *American Economic Review* (2022).
- [16] Emir Kamenica and Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review* 101, 6 (2011), 2590–2615.
- [17] Shao-Heng Ko and Kamesh Munagala. 2022. Optimal Price Discrimination for Randomized Mechanisms. In *Proceedings of the 23rd ACM Conference on Economics and Computation, EC*.
- [18] Jieming Mao, Renato Paes Leme, and Kangning Wang. 2022. Interactive Communication in Bilateral Trade. In *13th Innovations in Theoretical Computer Science Conference, ITCS*. 105:1–105:21.
- [19] Roger B Myerson. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
- [20] Weiran Shen, Pingzhong Tang, and Yulong Zeng. 2018. A closed-form characterization of buyer signaling schemes in monopoly pricing. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS*. 1531–1539.
- [21] Weiran Shen, Pingzhong Tang, and Yulong Zeng. 2019. Buyer Signaling Games in Auctions. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS*. 1591–1599.
- [22] Haifeng Xu, Zinovi Rabinovich, Shaddin Dughmi, and Milind Tambe. 2015. Exploring information asymmetry in two-stage security games. In *29th AAAI Conference on Artificial Intelligence, AAAI*.