

Does Information Revelation Improve Revenue?

Constantinos Daskalakis, CSAIL and EECS, MIT
Christos Papadimitriou, EECS, UC Berkeley
Christos Tzamos, CSAIL and EECS, MIT

We study the problem of optimal auction design in a valuation model, explicitly motivated by online ad auctions, in which there is two-way informational asymmetry, in the sense that private information is available to both the seller (the item type) and the bidders (their type), and the value of each bidder for the item depends both on his own and the item's type. Importantly, we allow arbitrary auction formats involving, potentially, several rounds of signaling from the seller and decisions by the bidders, and seek to find the optimum co-design of signaling and auction. We call this optimum the “optimum augmented auction” — previous work on signaling in auctions fixes one of the two components. We characterize exactly the optimum augmented auction for our valuation model by establishing its equivalence with a multi-item Bayesian auction with additive bidders. Surprisingly, in the optimum augmented auction there is no signaling whatsoever, and in fact the seller need not access the available information about the item type until after the bidder chooses his bid. Suboptimal solutions to this problem, which have appeared in the recent literature, are shown to correspond to well-studied ways to approximate multi-item auctions by simpler formats, such as grand-bundling (this corresponds to Myerson's auction without any information revelation), selling items separately (this corresponds to Myerson's auction preceded by full information revelation as in [Fu et al. 2012]), and fractional partitioning (this corresponds to Myerson's auction preceded by optimal signaling). Consequently, all these solutions are separated by large approximation gaps in revenue from the optimum.

CCS Concepts: •**Theory of computation** → **Algorithmic mechanism design**;

ACM Reference Format:

Constantinos Daskalakis, Christos Papadimitriou and Christos Tzamos, 2016. Does Information Revelation Improve Revenue? ACM V, N, Article A (February 2016), 22 pages.

DOI = 10.1145/XXXXXXX.XXXXXXX <http://doi.acm.org/10.1145/XXXXXXX.XXXXXXX>

1. INTRODUCTION

Ad auctions — the allocation, through competitive bidding, of the attention of world-wide web users to advertisers in the context of keywords by search engines, or in web documents by publishers — are paramount. They are the lifeblood of the Information Economy, one of the most striking, innovative, and technically intriguing economic phenomena to emerge in recent decades, let alone one of the defining problems of the field of Algorithmic Game Theory.

Naturally, auctions have been studied intensely and brilliantly in Economics for half a century. However, ad auctions differ in important aspects, to be articulated below, from the traditional auction paradigms, inspired by auctions for art objects, mining rights, or broadcast frequencies. In particular, the important papers by Myerson [1981] on optimal auction design, and by Milgrom and Weber [1982] on the role of signaling in auctions, which have been shining lights of auction theory since the early 1980s, do not capture ad auctions well.

Myerson's work models the uncertainty in the bidders' valuations in a Bayesian manner, assuming that a distribution from which the bidders' valuations are drawn is known to all — the auctioneer and all bidders. If independence is assumed, a remarkably elegant mathematical analysis yields a revenue-optimal auction.¹ In contrast with the standard Bayesian model, in ad auctions the value of each bidder (advertiser) for an item (user) depends in intriguing ways on the advertiser's precise business model, but also on detailed information about the user that is available, typically through cookies, to the auctioneer; that is to say, there is significant information asym-

¹Auctioneer revenue is our main objective in this paper.

metry in ad auctions. This is more pronounced in display ads than in search ads, since in the latter the auctioneer and the bidder do share at least one important piece of information: the keyword.

Information asymmetry in auctions is tackled head on in the paper by Milgrom and Weber [1982]. They assume that the auctioneer as well as each agent receive real-valued signals, which together with some extra signals not observed by anyone determine the value of each agent for the item. Everyone is assumed to have consistent priors about all signals that he does not observe, and each bidder's value is the same, non-decreasing function of his own signal, the set of all other bidders' signals, and the signals observed by the auctioneer and by nobody. The signals are also subtly correlated, in that their joint distribution is assumed to be *affiliated*, satisfying an FKG-type technical condition. Intuitively, this condition ensures that larger values for some signals render more likely larger values for other signals. Hence, a larger value for one agent renders more likely larger values for the others. Milgrom and Weber analyze in this model the four common auction types, and order them with respect to expected auctioneer revenue: The English auction fares best, followed by the second-price auction, with the Dutch auction and the first-price auction trailing; reserve prices and entry fees are also considered. Then they prove that in all these auctions, *the auctioneer can increase the expected revenue by adopting a policy of always revealing all available information, i.e. the signals he observes and any other affiliated signal he can obtain*. This remarkable result, compatible with age old practices by auction houses, is known as the *linkage principle*, mathematical articulation of the intuition that more information means that all bidders are subject to less risk and information costs, and can therefore be a little bolder in their bidding.

Note that Milgrom and Weber do not address the *combined problem*: What is the optimum *augmented auction*? By “augmented auction” we mean here a design that specifies both a signaling protocol and an allocation policy (potentially multi-round) to be used by the auctioneer. In optimization, it is well known that separately optimizing two aspects of a design can result in gross suboptimality. Hence, finding the optimum augmented auction requires an optimal co-design of the signaling and the allocation rule of the auction. The motivation of Milgrom and Weber was different. They aimed to understand the role of signaling in the most commonly used single-item auction formats, and they established linkage for a general model affiliation. In contrast, we aim to identify the optimal augmented auction: Here we consider a simple and novel model of bidder valuations explicitly motivated by ad auctions, and we identify the optimum augmented auction. Our main result is that *the optimal augmented auction requires no information revelation by the auctioneer, and corresponds to the optimal auction of a related multi-item additive bidder setting*.

We proceed to explain our model and results, after providing some additional context. Signaling in auctions has reemerged in the study of ad auctions in three recent papers [Emek et al. 2014; Bro Miltersen and Sheffet 2012; Fu et al. 2012]. All three papers have a valuation model close to ours, explained shortly, which is not surprising as they focus on ad auctions. In [Emek et al. 2014; Bro Miltersen and Sheffet 2012], each bidder's valuation is a function of the bidder's and the item's type, both types being discrete. The assumption is that the auctioneer knows the realized item type, each bidder knows his own type, and everyone knows the joint distribution of types of all bidders and the item. Given this distribution, the authors seek to optimize the signaling by the auctioneer preceding a second-price auction, and conclude that designing the optimal signaling scheme is computationally highly non-trivial in general.

Similarly, in [Fu et al. 2012], every bidder b has a valuation function $v_b(t_b, t_i)$, which depends on the item type t_i and the bidder type t_b . Importantly, the bidder type t_b is assumed scalar (i.e. the bidder is single-dimensional) and the function $v_b(\cdot, t_i)$ is

increasing for all t_i . Each bidder knows his own type, the auctioneer knows the item types, and all types are assumed independent from distributions known to everybody. Following the formulations of [Emek et al. 2014; Bro Miltersen and Sheffet 2012], which have their roots in the signaling literature, they consider signaling schemes where the auctioneer commits to sending some (potentially randomized) signal $\sigma(t_i)$ drawn from a distribution $Q(t_i) \in \Delta(\Sigma)$ that depends on the realized item type t_i , where Σ is some signal alphabet. Upon receiving the signal, the bidders are expected to update their beliefs about the item type. Hence, the auctioneer should run Myerson’s optimal auction for the distributions reflecting the bidders’ updated expectations for the item value, conditioning on the signal. For this two-stage format, they show that full revelation of information by the auctioneer, namely using $\sigma(t_i) = t_i$, is optimal, i.e. linkage holds. It is quite subtle and surprising this mechanism (of optimal signaling followed by Myerson) is a suboptimal solution of the signaling and auction co-design problem that we consider here, as the following example suggests.

Example 1.1. We have a single bidder and a single item. There are two possible item types, H and L, and three possible bidder types A, B and C; both item and bidder types are uniformly random. The bidder’s value for the item as a function of the item and the bidder types is shown in this table:

		Item Types	
		H	L
Bidder Types	A	6	6
	B	10	6
	C	20	16

Notice that the valuation defined by the table is non-decreasing in the bidder’s type, for each item type. Hence, as implied by [Fu et al. 2012], the optimal signaling followed by Myerson’s auction is to perfectly reveal the item type and charge 10 for item H and 6 for item L. The resulting revenue is $6 + \frac{1}{3} \approx 6.333$. In this auction, we “go after” types B and C for item H, and all three types for item L. In particular, type C always has a surplus of at least 10. How can we get some of that surplus? Can we target C, without losing revenue from types A and B?

Here is how: Before revealing the item type, we announce the prices (10 for H and 6 for L) and then ask the bidder either to choose now one of the two item types foregoing his right to buy the other, or pay 5 for the privilege of postponing this choice until after the item type is revealed. Then we reveal the item type, and sell it — provided the bidder has not lost his right to buy it. It is easy to see that, in this protocol, a bidder of type A will waive his right to H, and will therefore pay 3 in expectation. A bidder of type B will waive his right to L, thereby paying 5 in expectation. However, a bidder of type C will prefer to pay 5 upfront, and then pay another 8 in expectation, for a total of 13. Notice that even though the new auction decreased the expected payment of type B from 8 to 5, the expected payment of type C increased from 8 to 13. This auction, by deviating from [Fu et al. 2012]’s optimal signaling followed by Myerson, managed to “target” a particular bidder type and thus achieve a better expected revenue of 7². Therefore, full information disclosure and the linkage principle may not hold — even when the distributions are affiliated — when followed by auction formats other than the four classical ones considered by Milgrom and Weber. This concludes Example 1.1.

²An equivalent implementation of this auction is through *discount*: Announce the prices to be 20 and 16 and give the bidder a 10-off discount coupon to be applied to an item type of his choice prior to the realization of the item type.

In further recent work on signaling, [Dughmi et al. 2014] study the optimum signaling scheme to precede a second-price auction under restrictions on the signals that can be sent, while [Cheng et al. 2015] develop algorithmic tools for optimum signaling under general objectives and constraints. Besides the works on signaling in auctions outlined so far, there are several other papers addressing the problems of “matching” and “targeting” [Board 2009; Abraham et al. 2013; Ghosh et al. 2015; Bhawalkar et al. 2014; Hummel and McAfee 2015], variants of the information disclosure conundrum, in various models which yield interesting results. Linkage—that is, advantage of disclosure by the auctioneer—sometimes prevails, under assumptions, and sometimes does not. None of the papers in the literature solves or considers what we call “the augmented auction design problem,” the problem of co-designing the optimal signaling and allocation rule.

Finally, an interesting precursor of our result is the work in [Bergemann and Pendorfer 2007]. They study a problem conceptually quite close to ours involving the simultaneous optimization over both information structure and allocation policy, in a setting where the bidder has no private information and the seller can directly reveal to the bidder what the value is *without observing that value*; interestingly, the optimal combined mechanism involves a coarse revelation followed by a Myerson auction. [Esö and Szentes 2007] and [Bergemann and Wambach 2015]) extend that model to allow private information to the bidders and study more general sequential auctions. In contrast, in our model multiple rounds of interaction cannot improve revenue (Theorem 3.4).

Our Model and Results. In our model, an auctioneer is selling a single item,³ which may be of n possible types and whose realized type t_i is drawn from a publicly known distribution $\bar{\pi}$. There are also m bidders interested in the item, who receive types independently from publicly known distributions D_1, \dots, D_m . Each distribution D_i is an n -dimensional distribution. Every vector in the support of D_i is a possible type for bidder i and specifies the bidder’s value for each type of item. An important special case of our problem is when there is a single bidder, i.e. $m = 1$. In this case, the setting can be summarized in a table V_{ij} , whose rows and columns correspond respectively to possible types of the bidder and the item, together with a distribution over rows (bidder types) and a distribution over columns (item types). In this language, the following example discusses crucial aspects of ad auctions missed by many of the current auction models.

Example 1.2.

		Item Types	
		young professional (user)	executive (user)
Bidder Types	sportscar ad	30	20
	business jet ad	10	50

Notice the cycle in the graph with table entries as nodes, and an edge from (i, j) to (i, k) if $V_{ij} > V_{ik}$, or from (i, j) to (ℓ, j) if $V_{ij} < V_{\ell j}$. We believe that such cycles are the signature of ad auctions, the distinguishing feature that differentiates them from other auction paradigms. In the setting studied by [Fu et al. 2012], the rows (bidder types) can be permuted so that every column is non-decreasing: no cycles are possible. Linear order of rows or columns is evidence of a model’s single-dimensionality. The assumption of affiliated signal distributions also imputes a kind of monotonicity and one-dimensionality, this time on the item types. The augmented auction problem

³All our results extend to multi-item settings, as discussed in Section 6, but we focus our discussion to single-item settings for simplicity of the exposition.

with cyclic tables “feels” inherently multi-dimensional, and, as we shall see, this multi-dimensional nature is amply reflected in its optimal solution.

Let us consider another example.

Example 1.3. There are two item types, uniformly distributed in {red, blue}, and one buyer with 25 possible types, uniformly distributed on the set $\{0, 1, 2, 3, 4\} \times \{0, 1, 2, 3, 4\}$. In Section 3.2, we show that the optimal signaling-followed-by-Myerson auction achieves revenue 1.2. In fact, as we show in Section 3.1, just the plain old Myerson auction without any information revelation achieves revenue of 1.2. Nevertheless, as we show in Section 3.3, the optimal revenue is 1.28. One way to achieve this revenue is to give the bidder the following options, before revealing the item type:

- Pay 1.5 upfront for the option of getting the item if it is blue;
- Pay 1.5 upfront for the option of getting the item if it is red;
- Pay 2.5 upfront to get the item independently of its color.

It is easy to see that a buyer of types (3, 0), (3, 1), (4, 0) or (4, 1) will purchase the “red” option, and similarly for the “blue” option. In contrast, a buyer of type (2, 3), (2, 4), (3, 2), (3, 3), (3, 4), (4, 2), (4, 3) or (4, 4) will for the last option, paying 2.5. Hence, the revenue equals $\frac{8}{25} \cdot 1.5 + \frac{8}{25} \cdot 2.5 = 1.28$.

Another way to get the optimal revenue does not even require the auctioneer to learn the item type. Here is the auction: The auctioneer offers the buyer two options:

- Pay 2.5 to get the item;
- Pay 3 to get the item, with the additional option of returning it for full refund.

It is easy to check that the types who went after the third option in the previous implementation, will go for the first option now. Intuitively, these types are strong and they are seeking a “discount.” The weaker types, who went for the first two options in the previous implementation, will go after the second option now, as it represents “insurance.” This insurance stratagem for increasing revenue is reminiscent of the role of warranties for avoiding collapse in Akerlof’s work on informationally problematic markets [Akerlof 1970].

The optimal augmented auctions obtained in Examples 1.1 and 1.3 with their non-standard design principles — targeting, discounts, insurance, etc. — are manifestations of our main theorem, relating optimal augmented auction design with the well known multidimensional problem of optimal multi-item auctions.

INFORMAL THEOREM 1.4 (SEE THEOREM 3.4). *The optimal augmented auction design for a single item of n possible types and m bidders is equivalent to an optimal multi-item auction design with n items and m additive bidders. Moreover, there is an implementation of the optimal augmented auction in which the auctioneer reveals no information about the item type to the bidder prior to bidding.*

Thus our main contribution is the observation that the augmented auction problem (finding the best co-design of signaling and auction format) is an inherently multi-dimensional problem, and all the phenomena/pathologies/nonstandard design principles identified in the examples above are manifestations of this multi-dimensional nature of the problem. The intuition for our theorem is simple: an auctioneer who auctions a pen that is 50% red and 50% blue, as per Example 1.3, is in fact auctioning, to additive bidders, two items: One item that is a blue pen with probability 50% and nothing otherwise; and another item that is a red pen with probability 50% and nothing otherwise. Hence, the optimum solution to the augmented auction problem maps to the optimum two-item auction with additive bidders. If in the latter some bidder

pays p for the (fractional) red pen, this means that in the former the same bidder pays p for the option of getting the red pen, if the pen is red. Importantly, the optimum augmented auction has an implementation where the auctioneer reveals no information about the item type to the bidder. That is, *when both auction format and signaling are optimized simultaneously, the optimal signaling is the null signaling*. Hence, the linkage principle does not prevail. Stronger than null signaling, sometimes the optimal augmented auction can be turned into a *no peeking auction*: in Example 1.3, the auctioneer can derive maximum revenue from the item without ever finding out if it is blue or red.

We also show that the above statements on auction optimality are more than academic curiosities: We can leverage the connection to multi-item auctions with additive bidders to show that the optimum augmented auction may result in dramatically more revenue than the suboptimal auction paradigms currently in use. Theorem 4.1 fleshes out the connections between augmented auctions and multi-item auctions, relating different signaling-auction formats to well-studied multi-item auction formats. Pushing the connection to the multi-item setting even further, Tables I and II provide a revenue comparison between the following auction formats for single and multi-bidder settings respectively and for different kinds of item/bidder type distributions: (a) running Myerson’s auction with no information revelation by the auctioneer, (b) running Myerson’s auction preceded by the optimal information revelation by the auctioneer, and (c) running the optimal augmented auction. We show that the approximation gaps between (c) and (b) may be unbounded.

One could object that we have not solved the problem of optimal augmented auctions, but instead we have *reduced it* to multi-item auctions with additive bidders, an important open problem in auction theory, notorious for its analytical difficulty. This is all true, and yet our reduction has many and potent consequences. First, it is important to know the nature of the true optimum, to which an auction designer should aspire, because this way one understands better how existing solutions are in reality compromises to the true goal, and one can start exploring ways to improve. Second, there is recent progress in multi-item auctions; for example, there are novel and promising linear characterizations of the problem [Alaei et al. 2012, 2013; Cai et al. 2012a,b, 2013; Haghpanah and Hartline 2015] and the one-bidder case is by now reasonably well understood [Daskalakis et al. 2013; Giannakopoulos and Koutsoupias 2014; Daskalakis et al. 2016]. Hence our reduction opens new avenues for improving on the status quo.

Another possible objection is that even the one-bidder solution of the multi-dimensional problem is in general quite complicated and of a stochastic nature, and therefore awkward to use in practice (the bidder chooses among a possibly long menu of complex lotteries involving subsets of the items). However, one can interpret the one-bidder optimal solution as *a set of possible contracts* offered by the auction designer, as an alternative to the auction. In a situation where there are millions of items of each type, and thousands of bidders of each type — not an unrealistic assumption for a day of ad auctions, assuming some coarse enough clustering of both sides — the probabilities of both kinds of types, as well as the probabilities in the lotteries, are multiplied by large enough numbers and are realized as magnitudes that are reasonably accurately known a priori. Such *large market interpretation* of the one-bidder solution may provide a new basis for understanding how ad contracts should be designed.

2. PRELIMINARIES

We consider an auction setting with a single item for sale⁴. The item can be of one of n different types, and there is a publicly known prior probability distribution $\vec{\pi}$ on

⁴Multiple items can be dealt with similarly, see the discussion in Section 6.

which type it is. There are m buyers interested in the item and every buyer i has an n -dimensional distribution D_i which is independent⁵ of the distributions of other buyers. Each vector in the support of D_i corresponds to a possible type of bidder i , specifying his value for each item type.

The seller is running an auction for the item, seeking to maximize revenue. Neither the seller nor the buyers know anything about the item type, other than its prior probability distribution.

2.1. Signaling Schemes

The seller may commit to a policy of revealing, prior to the auction, certain information about the item type to the buyers with the goal of increasing revenue. A *signaling scheme* $Q : [n] \rightarrow \Delta(S)$ maps each item type to a distribution over a set of signals S . Upon receiving a signal $s \in S$, each buyer updates his prior for the item $\vec{\pi}$ to a distribution $\vec{\Pi}(s)$, where

$$\vec{\Pi}_j(s) = \frac{\text{Prob}[\text{signal is } s | \text{item is } j] \cdot \text{Prob}[\text{item is } j]}{\text{Prob}[\text{signal is } s]} = \frac{Q_s(j)\pi_j}{\sum_{j'} Q_s(j')\pi_{j'}}.$$

Therefore, each signal $s \in S$ gives an updated prior $\vec{\Pi}(s)$. Given a collection of vectors $\vec{\Pi}(s)$ and their corresponding probabilities $P(s) = \text{Prob}[\text{signal is } s] = \sum_{j'} Q_s(j')\pi_{j'}$, they must satisfy

$$\sum_{s \in S} P(s)\vec{\Pi}(s) = \vec{\pi}$$

Conversely, given a collection of vectors $\vec{\Pi}(s)$ with corresponding probabilities $P(s)$, there exists a signaling scheme Q producing these belief vectors with these probabilities if the above equality holds. This is easy to see by setting $Q_j(s) = P(s)\frac{\vec{\Pi}_j(s)}{\pi_j}$.

We can thus use an alternative parametrization of a signaling scheme in terms of $\vec{\Pi}$ and P .

2.2. Auction Formats

We consider the following three auction formats and information structures.

Myerson (No Information Revelation). Neither the buyer nor the seller has any information about the item type. Every buyer i with value vector $\vec{v}_i \sim D_i$ has value $\vec{v}_i \cdot \vec{\pi}$ for the item and the setting becomes single dimensional. Here, the standard Myerson auction gives us the revenue-optimal mechanism.

Myerson with Signaling [Fu et al. 2012]. The seller can observe the item type and may engage in signaling, which may cause the buyers to update their prior about the item. Then the seller is restricted to run a Myerson auction.

Optimal Augmented Auction. The seller can observe the item type, engage in signaling to the buyer (in fact, in our result we will allow more general signaling schemes than the kind defined above), and then proceed to auction the item, so that the choice of the combination of signaling scheme and auction optimizes revenue and is implementable either as an ex-post or a Bayes-Nash equilibrium.

Notice, in particular, that the optimal mechanisms in Examples 1.1 and 1.3 were not expressible in the Myerson-with-signaling auction format.

⁵For an extension of the model allowing correlations, see the discussion in Section 6.

3. OPTIMAL AUCTIONS

3.1. Myerson (No Information Revelation)

When no information is available about the item type, the setting becomes one dimensional with bidder value distributions $\vec{\pi} \cdot D_i$. By $\vec{\pi} \cdot D_i$ we denote the distribution obtained by first sampling a vector $\vec{v}_i \sim D_i$ and then computing the inner product with the prior belief vector $\vec{\pi}$. Myerson gives us the optimal auction in this one dimensional setting which we denote by $\text{MYERSON}(\vec{\pi})$.

Recall example 1.3, where there are two item types ($n = 2$) and one buyer ($m = 1$) with D_1 being the uniform distribution over $\{0, 1, 2, 3, 4\} \times \{0, 1, 2, 3, 4\}$. The belief about the item type is uniform, i.e. $\vec{\pi} = (0.5, 0.5)$. If no information about the item type is revealed, the buyer can be collapsed to single-dimensional with value distribution $\vec{\pi} \cdot D_1$, which takes values in the set $\{0, 0.5, 1, 1.5, 2, 2.5, 3, 3.5, 4\}$. The optimal mechanism offers the item at a price of 2 which is accepted with probability 60%. This gives a revenue of 1.20.

3.2. Myerson with Signaling

Let us denote by $\text{MSIGNALING}(\vec{\pi})$ the maximum revenue the seller can achieve by fixing a signaling scheme and then running Myerson, where beliefs $\vec{\Pi}(s)$ occur with probability $P(s)$, i.e.:

$$\text{MSIGNALING}(\vec{\pi}) = \max_{(\vec{\Pi}, P)} \sum_s P(s) \cdot \text{MYERSON}(\vec{\Pi}(s))$$

where the maximum is computed over signaling schemes such that $\sum_{s \in S} p(s) \vec{\Pi}(s) = \vec{\pi}$.

We can show that, without loss of generality, we can restrict ourselves to signaling schemes with at most n signals (this is a fairly common phenomenon in the signaling literature).

LEMMA 3.1. *Consider a signaling scheme $(\vec{\Pi}, P)$. There exists a signaling scheme $(\vec{\Pi}', P')$ that achieves at least as much revenue in the Myerson with Signaling setting which uses at most n signals.*

PROOF. By reverse induction. Assume there exists a signaling scheme with $|S| > n$ signals. We will construct a signaling scheme with $|S| - 1$ signals achieving the same revenue. Consider the vectors $\vec{\Pi}(s)$ for all $s \in S$. Since $|S| > n$ the vectors are not linearly independent. Thus, there exist constants α_s which are not all 0 such that $\sum_s \alpha_s \vec{\Pi}(s) = \vec{0}$. Splitting the summation terms into positive and negative we get that $\sum_{s: \alpha_s > 0} |\alpha_s| \vec{\Pi}(s) = \sum_{s: \alpha_s < 0} |\alpha_s| \vec{\Pi}(s)$.

Now consider the corresponding difference in the revenue by the positive and negative sides, i.e. $\sum_s \alpha_s \text{MYERSON}(\vec{\Pi}(s))$. Without loss of generality, we can assume that this quantity is non-negative, since otherwise we can pick a new vector $\vec{\alpha}' = -\vec{\alpha}$. Furthermore, let $\lambda = \min_{s: \alpha_s < 0} \left(\frac{P(s)}{|\alpha_s|} \right)$ and s^* be the corresponding index where this is minimized.

We create a new signaling scheme with corresponding probabilities $P'(s) = P(s) + \lambda \alpha_s$. This makes the generated revenue at least as much as it was before:

$$\begin{aligned} \sum_s (P(s) + \lambda \alpha_s) \text{MYERSON}(\vec{\Pi}(s)) &\geq \sum_s P(s) \text{MYERSON}(\vec{\Pi}(s)) + \lambda \sum_s \alpha_s \text{MYERSON}(\vec{\Pi}(s)) \\ &\geq \sum_s P(s) \text{MYERSON}(\vec{\Pi}(s)) \end{aligned}$$

Moreover, $P'(s^*) = 0$ since $P(s^*) + \frac{P(s^*)}{|\alpha_{s^*}|} \alpha_{s^*} = 0$ which means that the new signaling scheme uses at most $|S| - 1$ signals. By induction, we conclude that repeating this process we can obtain a signaling scheme that uses at most n signals which achieves at least as much revenue as the original signaling scheme. \square

We can also show an upper bound on the revenue of the optimal Myerson auction with signaling as follows:

LEMMA 3.2. *The maximum revenue achievable in the Myerson with Signaling setting is at most*

$$\text{MSIGNALING}(\vec{\pi}) \leq \max_{\vec{p} \geq 0: |\vec{p}|=1} \text{MYERSON}(\vec{p})$$

Moreover, the above holds with equality if all distributions are item-symmetric and the prior probabilities $\vec{\pi}$ are uniform.

PROOF. For any signaling scheme $(\vec{\Pi}, P)$ it holds that:

$$\begin{aligned} \text{MSIGNALING}(\vec{\pi}) &= \max_{(\vec{\Pi}, P)} \sum_s P(s) \text{MYERSON}(\vec{\Pi}(s)) \leq \max_{(\vec{\Pi}, P)} \max_s \text{MYERSON}(\vec{\Pi}(s)) \\ &\leq \max_{\vec{p} \geq 0: |\vec{p}|=1} \text{MYERSON}(\vec{p}) \end{aligned}$$

Now, let \vec{p}^* be the vector that maximizes the right-hand side. In an item-symmetric setting with uniform prior probabilities on the item types, we can construct a feasible signaling scheme $(\vec{\Pi}, P)$ by considering all cyclic rotations of the vector \vec{p}^* with equal probability. Since, all distributions are item symmetric, all cyclic rotations give revenue as $\text{MYERSON}(\vec{p}^*)$. This proves that in the symmetric setting, the inequality shown above is tight. \square

We revisit Example 1.3. Applying Lemma 3.2, we get the optimal revenue obtainable in the Myerson with signaling setting is upper bounded by

$$\max_{0 \leq p \leq 1} \text{MYERSON}(p, 1 - p)$$

which is maximized without loss of generality for some

$$p \in \left\{ \frac{a}{b} \mid 0 \leq a \leq b \leq 4 \text{ for } a, b \in \mathbb{Z} \right\}$$

Checking each of these values separately we see that the revenue is at most 1.2. This upper bound can be achieved by always revealing the item type to the buyer. Once the item type is revealed, the buyer's value is known to be distributed uniformly in $\{0, 1, 2, 3, 4\}$ and the optimal auction is to set a single price p that maximizes the expected revenue against this distribution. This is achieved by selling the item at 2 which is accepted with 60% probability. The same revenue can also be achieved by giving no information as in the previous section. We conclude that both signaling schemes are optimal.

3.3. Optimal Augmented Auction

We now study the problem of designing the combination of signaling scheme and subsequent auction that maximizes revenue and is either ex-post or Bayes-Nash implementable. We show that this is equivalent to the following multidimensional mechanism design problem:

Definition 3.3 (Optimal Multi-Item Auction). We have m additive bidders and n items. Bidder i 's value for item j is equal to $v_{ij}\pi_j$ where the vector \vec{v}_i comes from bidder

distribution D_i and $\vec{\pi}$ is a common parameter for all agents. We denote by $\text{REV}(\vec{\pi})$ the optimal revenue achievable by any incentive-compatible mechanism, given the parameter $\vec{\pi}$.

THEOREM 3.4. *The optimal augmented auction for a given $\vec{\pi}$ and D_1, \dots, D_m is equivalent to a multi-item auction with n items and m additive bidders as specified in Definition 3.3. In particular, for any ex-post implementable augmented auction, there exists an incentive-compatible multi-item auction that achieves the same revenue. Similarly, for any Bayes-Nash implementable augmented auction, there exists a Bayesian incentive-compatible multi-item auction that achieves the same revenue.*

Moreover, there is an implementation of the optimal augmented auction in which the auctioneer reveals no information about the item type to the bidder prior to bidding.

The proof of the theorem is quite simple, following by (i) observing that the revelation principle holds in settings with information asymmetry on both sides, and (ii) characterizing the auctions resulting from the revelation principle. What is subtle and noteworthy here is that the optimum is not a Myerson-with-signaling auction, even in the case of a single-dimensional bidder considered in [Fu et al. 2012]. Still, the optimum auction can be expressed succinctly through a one-to-one correspondence with a related multi-item auction with additive bidders.

PROOF. Consider a mechanism M that is ex-post implementable. The mechanism M may use multiple rounds of communication and information revelation to the buyers. For each $i \in [m]$, let $A_i(\vec{v}_i)$ be the (possibly randomized) equilibrium strategy of buyer i when his type is \vec{v}_i , i.e. his value for every item type j is v_{ij} . Each strategy $A_i(\vec{v}_i)$ specifies the actions that bidder i takes in the mechanism M given information about the item type and the actions of other buyers, in order to maximize buyer i 's expected payoff.

Let $X_i(j, \vec{A})$ be an indicator random variable that indicates whether buyer i gets the item when buyers choose strategies $\vec{A} = (A_1, \dots, A_m)$ and the realized item type is $j \in [n]$. Similarly, let $C_i(j, \vec{A})$ denote the price buyer i is asked to pay. The expected utility of buyer i is then equal to

$$E_{j \sim \pi}[E[X_i(j, \vec{A})v_{ij} - C_i(j, \vec{A})]]$$

where the first expectation is with respect to the randomness of the item type, while the second expectation is with respect to the randomness in the choices of the mechanism, the information revealed and the actions of the buyers under the mixed strategies \vec{A} .

We now examine the equilibrium conditions for the buyer strategies. For all possible types $\vec{V} = (\vec{v}_1, \dots, \vec{v}_m)$ and all possible misreports \vec{v}'_i for buyer i , we let $\vec{A}(\vec{V}) = (A_1(\vec{v}_1), \dots, A_m(\vec{v}_m))$ and $\vec{A}(\vec{v}'_i, \vec{V}_{-i}) = (A_1(\vec{v}_1), \dots, A_i(\vec{v}'_i), \dots, A_m(\vec{v}_m))$. The equilibrium conditions for the strategy profile $\vec{A}(\vec{V})$ imply that:

$$\begin{aligned} E_{j \sim \pi}[E[X_i(j, \vec{A}(\vec{V}))v_{ij} - C_i(j, \vec{A}(\vec{V}))]] &\geq \\ E_{j \sim \pi}[E[X_i(j, \vec{A}(\vec{v}'_i, \vec{V}_{-i}))v_{ij} - C_i(j, \vec{A}(\vec{v}'_i, \vec{V}_{-i}))]] &\end{aligned}$$

Now set the variables $x_{ij}(\vec{V}) = E[X_i(j, \vec{A}(\vec{V}))]$ and $c_i(\vec{V}) = E_{j \sim \pi}[E[C_i(j, \vec{A}(\vec{V}))]]$. The equation above can be rewritten for all i, \vec{V} and potential misreports \vec{v}'_i as:

$$\sum_{j=1}^n x_{ij}(\vec{V})\pi_j v_{ij} - c_i(\vec{V}) \geq \sum_{j=1}^n x_{ij}(\vec{v}'_i, \vec{V}_{-i})\pi_j v_{ij} - c_i(\vec{v}'_i, \vec{V}_{-i}) \quad (\text{IC})$$

Moreover, the mechanism must satisfy the individual rationality constraint in order for the bidders to participate

$$\sum_{j=1}^n x_{ij}(\vec{V})\pi_j v_{ij} - c_i(\vec{V}) \geq 0 \quad (\text{IR})$$

and that for all $j \in [n]$: $\sum_{i=1}^m x_{ij}(\vec{V}) \leq 1$. A mechanism that satisfies all these constraints and achieves expected revenue

$$E_{\vec{V} \sim D_1 \times \dots \times D_m} \left[\sum_{i=1}^m c_i(\vec{V}) \right]$$

can be viewed as a multi-item direct auction for m additive bidders with n items where item j is simply π_j units of item type j as in Definition 3.3. The direct mechanism allocates the item j to bidder i with probability $x_{ij}(\vec{V})$ when all bidders have types \vec{V} . The (IC) constraint implies that the auction is incentive compatible while the (IR) constraint ensures bidder participation. Any direct mechanism M for the multi-item setting is a feasible augmented auction since the seller can use M to presell all item types before observing the real item type and then decide on the final allocation after the type is observed.

The same reduction goes through for converting any Bayes-Nash equilibrium to an equivalent Bayesian incentive-compatible direct mechanism that presells all the items. The difference is that in the (IC) and (IR) constraints above for bidder i , we take an expectation over all bidder types \vec{V}_{-i} other than i . \square

Therefore, the optimal augmented auction achieves revenue $\text{REV}(\vec{\pi})$, as defined in Definition 3.3, when the prior on the items is $\vec{\pi}$.

Remark 3.5. One way to convince somebody that no information revelation before running the optimal augmented auction can possibly increase revenue is to notice that the function $\text{REV}(\cdot)$ is concave.

LEMMA 3.6. *The optimal revenue function $\text{REV}(\vec{x})$, as per Definition 3.3, is concave for all $\vec{x} \in \mathbb{R}_+^n$.*

PROOF. We must have that for all $x, y \in \mathbb{R}_+^n$, $\frac{1}{2}\text{REV}(\vec{x}) + \frac{1}{2}\text{REV}(\vec{y}) \leq \text{REV}\left(\frac{\vec{x}+\vec{y}}{2}\right)$. Since the revenue function is homogeneous, this is equivalent to showing that $\text{REV}(\vec{x}) + \text{REV}(\vec{y}) \leq \text{REV}(\vec{x} + \vec{y})$. Let M_x and M_y be the optimal mechanisms that achieve $\text{REV}(\vec{x})$ and $\text{REV}(\vec{y})$. We can show that $\text{REV}(\vec{x} + \vec{y})$ achieves higher revenue by constructing a mechanism that sells to the buyers a fraction of the items using mechanism M_x and the rest using M_y . In particular, in M_x , we sell $\frac{x_j}{x_j+y_j}$ units of every item j while in M_y , we sell $\frac{y_j}{x_j+y_j}$. This achieves revenue $\text{REV}(\vec{x}) + \text{REV}(\vec{y})$. \square

So what if we sent a signal and then run an optimal augmented auction for the posterior of the item type? Suppose $(\vec{\Pi}, P)$ is a signaling scheme, as per Section 2.1. This signaling followed by the optimal augmented auction would result in revenue $\sum_s P(s)\text{REV}(\vec{\Pi}(s)) \leq \text{REV}(\sum_s P(s)\vec{\Pi}(s)) = \text{REV}(\vec{\pi})$. The inequality follows from Jensen's inequality and Lemma 3.6. This confirms that no signaling scheme can preface running the optimal augmented auction on the posterior to increase revenue (which of course we knew already from optimality, but the concavity illustrates one reason why it happens). This concludes Remark 3.5. \blacksquare

We now revisit Example 1.3 to compute the optimal augmented auction. It follows from Theorem 3.4 that this can be computed by solving an associated multi-item auction problem. The result⁶ is the auction that gives the following choices to the buyer:⁷

- Buy one item at price 1.5.
- Buy both items at price 2.5.

We can see that if the buyer has values $(3, 0)$, $(3, 1)$, $(4, 0)$ or $(4, 1)$, he purchases only item 1 and similarly for item 2. The bundle of both items is bought by types $(2, 3)$, $(2, 4)$, $(3, 2)$, $(3, 3)$, $(3, 4)$, $(4, 2)$, $(4, 3)$ and $(4, 4)$. This means that the revenue is equal to $\frac{8}{25} \cdot 1.5 + \frac{8}{25} \cdot 2.5 = 1.28$. This is strictly better than the revenue achievable by running Myerson with signaling.

4. CONNECTIONS OF SIGNALING TO MULTI-ITEM MECHANISM DESIGN

In Theorem 3.4, we showed that the combined problem of designing a signaling scheme together with an optimal mechanism to sell the item to the bidders is equivalent to the multi-item mechanism design problem in Definition 3.3. In particular, we showed that no signaling is required and that the seller can treat the mechanism design problem as a multi-item auction where all types of items are available but their values are discounted according to the prior probability vector. Then, it becomes clear that the seller should not reveal any information but instead run an auction preselling each of the item types. The auction format should be the one that maximizes revenue in the equivalent multi-item setting.

It is interesting to consider how the restricted formats of auctions we studied above translate to the multi-item setting.

We first look at the case where no information is given to the buyers and a simple Myerson auction is run for the item. In this setting, the item is given to the winner of the auction no matter what type it is. This is equivalent to a multi-item mechanism for the setting of Definition 3.3 that bundles all items together and sells the bundle as one. This bundling mechanism is a well-studied simple mechanism in the multi-item auction literature, denoted as BREV. It performs quite well compared to the optimal mechanism in cases where the bidder values for the bundle of all items are well-concentrated. One important case where this happens is when items are identical and the value distributions for every item are identical and independent. We consider such settings in the Section 5 where we make a detailed comparison of how well the optimal augmented auction is approximated by simple signaling schemes.

We now consider the case of running a simple Myerson auction but using signaling to increase revenue. In this restricted auction format, we can see that the optimal

⁶The optimality of the mechanism was verified by solving a small linear program that encodes all the incentive compatibility constraints. For an analytical derivation in a similar setting with item values distributed uniformly in $[0, 1]$, we refer the reader to [Daskalakis et al. 2016].

⁷We state the solution to the associated multi-item problem here. In Section 1, we have already described what this auction corresponds to for the optimal augmented problem, and gave two possible implementations of the optimal augmented auction.

revenue can be written as

$$\begin{aligned}
\text{MSIGNALING}(\vec{\pi}) &= \max_{(\vec{\Pi}, P)} \sum_s P(s) \cdot \text{MYERSON}(\vec{\Pi}(s)) \\
&= \max_{(\vec{\Pi}, P)} \sum_s \text{MYERSON}(P(s) \cdot \vec{\Pi}(s)) \\
&= \max_{\substack{\vec{x}_1 + \dots + \vec{x}_n = \vec{\pi} \\ \vec{x}_j \in \mathbb{R}_{\geq 0}^n}} \sum_{j=1}^n \text{MYERSON}(\vec{x}_j)
\end{aligned} \tag{1}$$

This partitions the items into at most n groups according to the vectors \vec{x}_j . In the full information setting where the item type is always revealed to the buyers, we have that $\vec{x}_j = \pi_j \vec{e}_j$, where $\vec{e}_1, \dots, \vec{e}_n$ are the standard basis vectors. When viewed as a multi-item auction as in Definition 3.3, this corresponds to selling each item separately, by running separate Myerson auctions. This is an important mechanism, denoted by SREV , which is studied in the multi-item literature since it is simple and can achieve significant fraction of the revenue in many cases. A notable such case is when the value distributions for every item are independent. It was shown in [Babaioff et al. 2014] that in such a case, this simple mechanism can achieve an $O(\log n)$ fraction of the optimal revenue, i.e. $\text{REV} = O(\log n) \cdot \text{SREV}$. It was also shown there, that $\max\{\text{SREV}, \text{BREV}\}$ can perform even better achieving a constant fraction approximation for the previous case when there is only one bidder. Translating this result to signaling, this immediately implies that in many situations a significant fraction of the optimal revenue can be achieved by the extreme cases of giving full-information to the buyers or no information at all.

Other than these two extreme cases, we can achieve even higher revenue in the Myerson with signaling auction by partitioning items with more general vectors \vec{x}_j . We can partition items into arbitrary disjoint groups and sell each group of items separately, a case studied in the multi-item literature and denoted as PREV . In the most general form, we are allowed to make arbitrary fractional partitions of the items, which we denote by FPREV .

We summarize these connections in the following theorem:

THEOREM 4.1. *The problem of selling a single item whose type is one of n different types occurring with some prior probability $\vec{\pi}$ is equivalent to a multi-item auction setting where there are n different items for sale as in Definition 3.3. Moreover, different signaling mechanisms are equivalent to particular multi-item auctions:*

- Myerson without Information Revelation is equivalent to Grand-Bundling auctions that sell all items as one (BREV).
- Myerson with Full Information Revelation is equivalent to Selling-Separately (SREV).
- Myerson with Arbitrary Signaling Schemes is equivalent to Fractional-Partition mechanisms (FPREV).
- Optimal Augmented Auction is equivalent to the Revenue-Optimal multi-item mechanism (REV).

The following section studies how each of these restricted auction forms compare to the optimal achievable revenue.

5. REVENUE COMPARISON

In this section, we explore the gaps in revenue achievable in each of the signaling auction formats, i.e. Myerson with no information revelation, Myerson with optimal

Table I. Comparison of the worst-case approximation gap among the different auction formats for the case of a **single bidder**.

	Myerson (No Information Revelation) vs Myerson with Optimal Signaling	Myerson with Optimal Signaling vs Optimal Augmented Auction
i.i.d.	$O(1)$ [Lemma 5.1]	$O(1)$ [Lemma 5.1]
$\bar{\pi}$ uniform	$\Omega(1)$ [Trivial]	$\Omega(1)$ [Trivial]
independent	$O(n)$ [Lemma 5.2]	$O(1)$ [Lemma 5.4]
$\bar{\pi}$ arbitrary	$\Omega(n)$ [Lemma 5.3]	$\Omega(1)$ [Trivial]
general	$O(n)$ [Lemma 5.2]	$\leq \infty$ [Trivial]
$\bar{\pi}$ arbitrary	$\Omega(n)$ [Lemma 5.3]	$\geq \infty$ [Lemma 5.5]

signaling and the optimal augmented auction. To do this, we exploit the connections to well-studied auction formats in multi-item auction theory presented in Theorem 4.1.

Using our understanding from Section 4, we can immediately get the following relation between the achievable revenue in each of the auction formats:

$$\text{MYERSON} = \text{BREV} \leq \text{PREV} \leq \text{FPREV} = \text{MSIGNALING} \leq \text{REV}$$

We study how large the gap in revenue can get between different auction formats. We examine three different cases that are of particular interest:

- *i.i.d case* - A case where all item types are equally likely and every distribution D_i is product with i.i.d. marginals.
- *independent case* - A case where each D_i is a product distribution with independent (but not necessarily identical) marginals.
- *general case* - A case with arbitrary distributions D_i and $\bar{\pi}$.

5.1. Single Bidder

We will show a series of lemmas that bound the worst case revenue gap between different auction formats. The results are summarized in Table I.

LEMMA 5.1. *In a single bidder setting where D_1 is a product distribution with identical marginals and the prior distribution on item types $\bar{\pi}$ is uniform, we have that $\text{REV}(\bar{\pi}) = O(1) \cdot \text{MYERSON}(\bar{\pi})$.*

PROOF. Since D_1 is a product distribution with identical marginals and the prior distribution on item types $\bar{\pi}$ is uniform, the corresponding multi-item auction of Definition 3.3 is a setting with n items and i.i.d. distributions. This case has been studied in [Li and Yao 2013] where it was shown that for a single bidder, bundling can achieve at least a constant fraction of the optimal revenue. This immediately implies that for the optimal augmented auction $\text{REV}(\bar{\pi}) = O(1) \cdot \text{MYERSON}(\bar{\pi})$. \square

LEMMA 5.2. *For any number of bidders, arbitrary distributions D_i and arbitrary prior distribution on item types $\bar{\pi}$, we have that $\text{MSIGNALING}(\bar{\pi}) = O(n) \cdot \text{MYERSON}(\bar{\pi})$.*

PROOF. As shown in Lemma 3.1, the optimal signaling scheme uses at most n signals. Using Equation 1, we get:

$$\text{MSIGNALING}(\bar{\pi}) = \max_{\substack{\bar{x}_1 + \dots + \bar{x}_n = \bar{\pi} \\ \bar{x}_j \in \mathbb{R}_{\geq 0}^n}} \sum_{j=1}^n \text{MYERSON}(\bar{x}_j) \leq n \cdot \text{MYERSON}(\bar{\pi})$$

The inequality follows by noting that $\text{MYERSON}(\bar{x}_j) \leq \text{MYERSON}(\bar{\pi})$ since $\bar{x}_j \leq \bar{\pi}$. \square

LEMMA 5.3. *There is a single-bidder instance with a product distribution D_1 , such that $\text{MSIGNALING}(\bar{\pi}) = \Omega(n) \cdot \text{MYERSON}(\bar{\pi})$.*

Table II. Comparison of the worst-case approximation gap among the different auction formats for the case of **many bidders**.

	Myerson (No Information Revelation) vs Myerson with Optimal Signaling		Myerson with Optimal Signaling vs Optimal Augmented Auction	
i.i.d.	$O(n)$	[Lemma 5.2]	$O(\log n)$	[Lemma 5.7]
$\bar{\pi}$ uniform	$\Omega(\frac{n}{\log n})$	[Lemma 5.6]	$\Omega(\frac{\log n}{\log \log n})$	[Lemma 5.8]
independent	$O(n)$	[Lemma 5.2]	$O(\log n)$	[Lemma 5.7]
$\bar{\pi}$ arbitrary	$\Omega(n)$	[Lemma 5.3]	$\Omega(\frac{\log n}{\log \log n})$	[Lemma 5.8]
general	$O(n)$	[Lemma 5.2]	$\leq \infty$	[Trivial]
$\bar{\pi}$ arbitrary	$\Omega(n)$	[Lemma 5.3]	$\geq \infty$	[Lemma 5.5]

PROOF. To show the lemma, we use a lower bound from [Hart and Nisan 2012] (Example 15). Consider a uniform prior distribution on item types. Take a large M and let D_1 have support $0, M^i$ with $Pr(M^i) = M^{-i}$. The revenue achieved by MYERSON (i.e. BREV) is at most $\frac{1}{n}(1 + \frac{1}{M-1})$. In contrast, by always revealing the item type to the buyer, we can sell every item type i at a price M^i getting revenue 1 in expectation. This proves that $MSIGNALING(\bar{\pi}) = \Omega(n) \cdot MYERSON(\bar{\pi})$. \square

LEMMA 5.4. *In the case of a single bidder with a product distribution D_1 , we have that $REV(\bar{\pi}) = O(1) \cdot MSIGNALING(\bar{\pi})$.*

PROOF. As shown in Theorem 4.1, Myerson with Signaling can achieve revenue at least SREV by full information revelation and BREV by providing no information about the item type. For a single additive bidder with independent distributions, it is shown in [Babaioff et al. 2014], that $REV \leq 6 \cdot \max\{SREV, BREV\}$. This immediately implies that $REV(\bar{\pi}) = O(1) \cdot MSIGNALING(\bar{\pi})$. \square

LEMMA 5.5. *For any $a \geq 1$, there exists a single-bidder instance with $n = 2$ item types such that:*

$$REV(\bar{\pi}) \geq a \cdot MSIGNALING(\bar{\pi})$$

PROOF. In Lemma 5.2, we showed that $MSIGNALING(\bar{\pi}) = O(n) \cdot MYERSON(\bar{\pi})$. Since $MYERSON = BREV$, we get that for two items $MSIGNALING = O(1) \cdot BREV$. The lemma follows since BREV and REV can have an arbitrary gap in revenue even for two items as shown in [Hart and Nisan 2013]. \square

5.2. Multiple Bidders

Working similarly for multiple bidders as in the case of one bidder, we show a series of lemmas that bound the worst case revenue gap between different auction formats. The results are summarized in Table II.

LEMMA 5.6. *There exists an instance where every D_i is a product distribution with identical marginals and the prior distribution on item types $\bar{\pi}$ is uniform, such that $MSIGNALING(\bar{\pi}) = \Omega(\frac{n}{\log n}) \cdot MYERSON(\bar{\pi})$.*

PROOF. Consider an instance with $m = n$ bidders such that each bidder has value 1 for item type j with probability $\frac{1}{n}$ and 0 otherwise. The distribution over item types is uniform. By a simple tail bound for the Binomial distribution, we can see that a single bidder has non-zero value for $\Omega(\log n)$ items types with probability at most $\frac{1}{n^2}$. Taking a union bound over all bidders, we get that with probability $1 - \frac{1}{n}$, every bidder has at most $O(\log n)$ non-zero values and thus the maximum revenue achievable in this case by MYERSON is $O(\frac{\log n}{n})$ (recall that each type appears with probability $\frac{1}{n}$). In contrast,

if the seller always reveals the type of the item, the achieved revenue is $O(1)$, since he can sell the item at price 1 which is bought with constant probability $\approx 1 - \frac{1}{e}$. \square

LEMMA 5.7. *In the case where all distributions D_i are identical product distributions, we have that $\text{REV}(\bar{\pi}) = O(\log n) \cdot \text{MSIGNALING}(\bar{\pi})$.*

PROOF. In this setting (independent case), [Babaioff et al. 2014] show that $\text{REV} \leq (\ln n + 6) \cdot \text{SREV}$. The lemma immediately follows as for the Myerson with Signaling auction, we have $\text{MSIGNALING} \geq \text{SREV}$. \square

LEMMA 5.8. *There exists an instance where every D_i is a product distribution with identical marginals and the prior distribution on item types $\bar{\pi}$ is uniform, such that $\text{REV}(\bar{\pi}) = \Omega\left(\frac{\log n}{\log \log n}\right) \cdot \text{MSIGNALING}(\bar{\pi})$.*

PROOF. To show the lower bound on optimal revenue, we will construct an example where the optimal multi-item auction gets revenue $\text{REV} = \Omega(n \log n)$, while the optimal fractional partition mechanism FPREV that corresponds to MSIGNALING , as shown in Theorem 4.1, gets revenue $\text{FPREV} = O(n \log \log n)$.

The example is the same as in [Babaioff et al. 2014] (Proposition 8), where it was used to show that $\text{REV} = \Omega(\log n) \text{PREV}$. There are n items and \sqrt{n} bidders and the value of each bidder for each item is sampled with probability $\frac{1}{\sqrt{n}}$ from a distribution F with CDF⁸

$$F(x) = \begin{cases} 1 - x^{-1} & \text{if } x \in [1, n^{1/8}] \\ 1 & \text{if } x > n^{1/8} \end{cases}$$

and with the remaining probability it is 0. In [Babaioff et al. 2014], it is shown that $\text{REV} = \Omega(n \log n)$. We will show that $\text{FPREV} = O(n \log \log n)$ by showing that for any fractional bundle \vec{f} that contains f_j units of item j , selling the bundle gives revenue $\text{BREV}(\vec{f}) = O(\log \log n) \cdot \sum_{j=1}^n f_j$.

Without loss of generality, assume that $\max_j f_j = 1$. Now consider the fractional bundle \vec{f}' by rounding up all values of \vec{f} to the smallest value 2^{-a} for a in the set $\{0, 1, \dots, \log n\}$. It is clear that the revenue from the fractional bundle \vec{f}' is at least as much as the revenue from the fractional bundle \vec{f} since all fractional values increased. Moreover, we have that $\sum_j f'_j = \sum_{j: f_j > n^{-1}} f'_j + \sum_{j: f_j \leq n^{-1}} f'_j \leq \sum_{j: f_j > n^{-1}} 2f_j + 1$. So $\sum_j f'_j = O(1) \cdot \sum_j f_j$.

We now bound the revenue of the fractional bundle \vec{f}' . The vector \vec{f}' consists of at most $\log n$ different values. For every such value $a \in \{0, 1, \dots, \log n\}$, we let S_a be the set of all items j with $f'_j = 2^{-a}$. We now consider a different setting with $\log n$ items, where item a corresponds to 2^{-a} units of the items in S_a . In this setting, there are $\log n$ items and we know that selling the items separately is a $O(\log \log n)$ approximation to the optimal revenue and in particular a $O(\log \log n)$ approximation to selling the grand bundle which corresponds to the revenue achieved by the fractional bundle \vec{f}' in the original setting.

This implies that $\text{BREV}(\vec{f}') \leq O(\log \log n) \sum_{a=0}^{\log n} 2^{-a} \text{BREV}(S_a)$ where $\text{BREV}(S_a)$ is the revenue achievable by selling the bundle of items in S_a . As shown in [Babaioff et al. 2014] (Proposition 8), $\text{BREV}(S_a) \leq O(1)|S_a|$ which implies that $\text{BREV}(\vec{f}') \leq \text{BREV}(\vec{f}) \leq O(\log \log n)|\vec{f}'| = O(\log \log n)|\vec{f}|$ as required. \square

⁸ F is an Equal-Revenue distribution with all mass above $n^{1/8}$ moved to an atom at $n^{1/8}$.

6. DISCUSSION AND EXTENSIONS

Our work studies the problem of signaling through the lens of multidimensional mechanism design in a model of valuations inspired by ad auctions. We show that the optimal co-design of signaling and auction maps to an optimal multi-item auction design in an associated multi-item setting with additive bidders, as specified in Definition 3.3. Our reduction allows us to leverage our understanding of multi-item auctions to (i) design optimal augmented auctions; (ii) understand how stratagems such as targeting, discount and insurance can be used to increase revenue; and (iii) quantify the suboptimality of signaling schemes studied in the literature by relating them to suboptimal multi-item auction designs in the associated multi-item setting. In the rest of the section we discuss additional benefits from our reduction, relating to multi-item extensions, computation and structure.

Multiple-Items. Since our mapping to multi-item auctions is already multi-dimensional, we can easily accommodate multiple items in our framework for augmented auction design as follows: in a setting with m bidders, k items with n types each, the optimal augmented auction design reduces to a multi-item auction design problem with m additive bidders and $k \cdot n$ items. (The proof of this is identical to that of Theorem 3.4.) Multi-item settings are important in applications of online ads where there are several ad-slots on a webpage and advertisers may have different values for each of these slots, which may also depend on the type of user that visits the page. For example, an advertiser can have values $(5, 3)$ for the slots on the top and the bottom of the page if the user is under 25 years old and values $(8, 2)$ otherwise.

Structure and Computation. By reducing the augmented auction design problem to that of revenue-optimal multi-item auctions, we get immediate implications about the structure of the problem and its computational complexity. The results of [Alaei et al. 2012; Cai et al. 2012a,b] imply that finding the optimal BIC augmented auction can be done in time polynomial in the support size of all bidder distributions D_i . At the same time, when the distributions D_i are product distributions and are given implicitly by listing their marginals (as in the i.i.d. and independent settings discussed in Section 5) the support size becomes exponential in the number of item types, and we can leverage our reduction to the multi-item case and the results of [Daskalakis et al. 2014] to show that the problem is #P-hard. Similarly, our connection to the multi-item case (in particular Theorem 4.1) and the results of [Daskalakis et al. 2012] imply that Myerson with No Information Revelation is also #P-hard. Still, when the distributions D_i are i.i.d. or independent across item types, we can leverage recent work of [Babaioff et al. 2014; Yao 2015] for approximately optimal multi-item auctions to obtain an augmented auction achieving a constant fraction of the optimal revenue. Lastly, there has been much recent effort in characterizing the structure of the optimal multi-item auction in the single-bidder case [Daskalakis et al. 2013; Giannakopoulos and Koutsoupias 2014; Daskalakis et al. 2016]. We can leverage these results to immediately get structural characterizations for the optimal single-bidder augmented auction.

Correlation among bidder values. Throughout the paper, we have assumed that the bidder values are independent and that the prior on the item types is the same for all bidders. A more general model, that does not make any independence assumptions, allows correlation of the bidder types. Our reduction can be easily extended to capture such situations by reducing the problem again to a multi-item auction where bidders have correlated values. See Appendix A for a more detailed discussion.

REFERENCES

- ABRAHAM, I., ATHEY, S., BABAIOFF, M., AND GRUBB, M. 2013. Peaches, lemons, and cookies: designing auction markets with dispersed information. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 7–8.
- AKERLOF, G. A. 1970. The market for “lemons”: Quality uncertainty and the market mechanism. *The quarterly journal of economics*, 488–500.
- ALAEI, S., FU, H., HAGHPANAH, N., AND HARTLINE, J. 2013. The Simple Economics of Approximately Optimal Auctions. In *the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- ALAEI, S., FU, H., HAGHPANAH, N., HARTLINE, J., AND MALEKIAN, A. 2012. Bayesian Optimal Auctions via Multi- to Single-agent Reduction. In *the 13th ACM Conference on Electronic Commerce (EC)*.
- BABAIOFF, M., IMMORLICA, N., LUCIER, B., AND WEINBERG, S. M. 2014. A simple and approximately optimal mechanism for an additive buyer. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*. IEEE, 21–30.
- BERGEMANN, D. AND PESENDORFER, M. 2007. Information structures in optimal auctions. *Journal of Economic Theory* 137, 1, 580–609.
- BERGEMANN, D. AND WAMBACH, A. 2015. Sequential information disclosure in auctions. *Journal of Economic Theory* 159, 1074–1095.
- BHAWALKAR, K., HUMMEL, P., AND VASSILVITSKII, S. 2014. Value of targeting. In *Algorithmic Game Theory*. Springer, 194–205.
- BOARD, S. 2009. Revealing information in auctions: the allocation effect. *Economic Theory* 38, 1, 125–135.
- BRO MILTERSEN, P. AND SHEFFET, O. 2012. Send mixed signals: earn more, work less. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 234–247.
- CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2012a. An Algorithmic Characterization of Multi-Dimensional Mechanisms. In *the 44th Annual ACM Symposium on Theory of Computing (STOC)*.
- CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2012b. Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization. In *the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2013. Reducing Revenue to Welfare Maximization: Approximation Algorithms and other Generalizations. In *the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. <http://arxiv.org/pdf/1305.4000v1.pdf>.
- CHENG, Y., CHEUNG, H. Y., DUGHMI, S., EMAMJOMEH-ZADEH, E., HAN, L., AND TENG, S.-H. 2015. Mixture selection, mechanism design, and signaling. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*. IEEE, 1426–1445.
- DASKALAKIS, C., DECKELBAUM, A., AND TZAMOS, C. 2012. Optimal pricing is hard. In *Internet and Network Economics: 8th International Workshop, WINE 2012, Singapore, December 11-14, 2012. Proceedings*. Vol. 7695. Springer, 298.
- DASKALAKIS, C., DECKELBAUM, A., AND TZAMOS, C. 2013. Mechanism Design via Optimal Transport. In *the 14th ACM Conference on Electronic Commerce (EC)*.
- DASKALAKIS, C., DECKELBAUM, A., AND TZAMOS, C. 2014. The Complexity of Optimal Mechanism Design. In *the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- DASKALAKIS, C., DECKELBAUM, A., AND TZAMOS, C. 2016. Strong Duality for a Multiple-Good Monopolist. *Econometrica (accepted)*.
- DUGHMI, S., IMMORLICA, N., AND ROTH, A. 2014. Constrained signaling in auction

- design. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 1341–1357.
- EMEK, Y., FELDMAN, M., GAMZU, I., PAESLEME, R., AND TENNENHOLTZ, M. 2014. Signaling schemes for revenue maximization. *ACM Transactions on Economics and Computation* 2, 2, 5.
- ESÖ, P. AND SZENTES, B. 2007. Optimal information disclosure in auctions and the handicap auction. *The Review of Economic Studies* 74, 3, 705–731.
- FU, H., JORDAN, P., MAHDIAN, M., NADAV, U., TALGAM-COHEN, I., AND VASSILVITSKII, S. 2012. Ad auctions with data. In *Symposium on Algorithmic Game Theory (SAGT)*.
- GHOSH, A., MAHDIAN, M., MCAFEE, R. P., AND VASSILVITSKII, S. 2015. To match or not to match: Economics of cookie matching in online advertising. *ACM Transactions on Economics and Computation* 3, 2, 12.
- GIANNAKOPOULOS, Y. AND KOUTSOPIAS, E. 2014. Duality and optimality of auctions for uniform distributions. In *Proceedings of the 15th ACM conference on Economics and computation*.
- HAGHPANAH, N. AND HARTLINE, J. 2015. Reverse mechanism design. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, 757–758.
- HART, S. AND NISAN, N. 2012. Approximate revenue maximization with multiple items. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 656–656.
- HART, S. AND NISAN, N. 2013. The menu-size complexity of auctions. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 565–566.
- HUMMEL, P. AND MCAFEE, R. P. 2015. When does improved targeting increase revenue? In *Proceedings of the 24th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 462–472.
- LI, X. AND YAO, A. C.-C. 2013. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences* 110, 28, 11232–11237.
- MILGROM, P. R. AND WEBER, R. J. 1982. A theory of auctions and competitive bidding. *Econometrica*, 1089–1122.
- MYERSON, R. B. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1, 58–73.
- YAO, A. C. 2015. An n -to-1 Bidder Reduction for Multi-item Auctions and its Applications. In *the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

A. CORRELATED BIDDER VALUES

Throughout the paper, we have assumed that the bidder values are independent and that the prior on the item types is the same for all bidders. A more general model, that does not make any independence assumptions, allows correlation of the bidder types. Our reduction can be easily extended to capture such situations by reducing the problem again to a multi-item auction where bidders have correlated values. In particular, we can have an arbitrarily correlated distribution $D \in \Delta(\mathbb{R}^{nm})$ over vectors $(\vec{v}_1, \vec{v}_2, \dots, \vec{v}_m)$ that specifies how the bidder values $\vec{v}_1, \vec{v}_2, \dots, \vec{v}_m$ are distributed.

Similar to Theorem 3.4, this setting is equivalent to a full-information multi-item auction setting with additive bidders whose values come from a correlated distribution. A bidder i with type \vec{v}_i in the original setting has value $\pi_j v_{ij}$ for every item j in the multi-item auction. The proof of this equivalence uses the revelation principle and is identical to the proof of Theorem 3.4. We repeat it here with the appropriate changes for completeness.

PROOF. Consider a mechanism M that is ex-post implementable. The mechanism M may use multiple rounds of communication and information revelation to the buyers. For each $i \in [m]$, let $A_i(\vec{v}_i)$ be the (possibly randomized) equilibrium strategy of buyer i when his type is \vec{v}_i , i.e. his value for every item type j is v_{ij} . Each strategy $A_i(\vec{v}_i)$ specifies the actions that bidder i takes in the mechanism M given information about the item type and the actions of other buyers, in order to maximize buyer i 's expected payoff.

Let $X_i(j, \vec{A})$ be an indicator random variable that indicates whether buyer i gets the item when buyers choose strategies $\vec{A} = (A_1, \dots, A_m)$ and the realized item type is $j \in [n]$. Similarly, let $C_i(j, \vec{A})$ denote the price buyer i is asked to pay. The expected utility of buyer i is then equal to

$$E_{j \sim \pi}[E[X_i(j, \vec{A})v_{ij} - C_i(j, \vec{A})]]$$

where the first expectation is with respect to the randomness of the item type according to bidder i , while the second expectation is with respect to the randomness in the choices of the mechanism, the information revealed and the actions of the buyers under the mixed strategies \vec{A} .

We now examine the equilibrium conditions for the buyer strategies. For all possible types $\vec{V} = (\vec{v}_1, \dots, \vec{v}_m)$ and all possible misreports \vec{v}'_i for buyer i , we let $\vec{A}(\vec{V}) = (A_1(\vec{v}_1), \dots, A_m(\vec{v}_m))$ and $\vec{A}(\vec{v}'_i, \vec{V}_{-i}) = (A_1(\vec{v}_1), \dots, A_i(\vec{v}'_i), \dots, A_m(\vec{v}_m))$. The equilibrium conditions for the strategy profile $\vec{A}(\vec{V})$ imply that:

$$\begin{aligned} E_{j \sim \pi}[E[X_i(j, \vec{A}(\vec{V}))v_{ij} - C_i(j, \vec{A}(\vec{V}))]] &\geq \\ E_{j \sim \pi}[E[X_i(j, \vec{A}(\vec{v}'_i, \vec{V}_{-i}))v_{ij} - C_i(j, \vec{A}(\vec{v}'_i, \vec{V}_{-i}))]] & \end{aligned}$$

Now set the variables $x_{ij}(\vec{V}) = E[X_i(j, \vec{A}(\vec{V}))]$ and $c_i(\vec{V}) = E_{j \sim \pi}[E[C_i(j, \vec{A}(\vec{V}))]]$. The equation above can be rewritten for all i, \vec{V} and potential misreports \vec{v}'_i as:

$$\sum_{j=1}^n x_{ij}(\vec{V})\pi_j v_{ij} - c_i(\vec{V}) \geq \sum_{j=1}^n x_{ij}(\vec{v}'_i, \vec{V}_{-i})\pi_j v_{ij} - c_i(\vec{v}'_i, \vec{V}_{-i}) \quad (\text{IC})$$

Moreover, the mechanism must satisfy the individual rationality constraint in order for the bidders to participate

$$\sum_{j=1}^n x_{ij}(\vec{V})\pi_j v_{ij} - c_i(\vec{V}) \geq 0 \quad (\text{IR})$$

and that for all $j \in [n]$: $\sum_{i=1}^m x_{ij}(\vec{V}) \leq 1$. A mechanism that satisfies all these constraints and achieves expected revenue

$$E_{\vec{V} \sim D} \left[\sum_{i=1}^m c_i(\vec{V}) \right]$$

can be viewed as a multi-item direct auction for m additive bidders with n items where the value of a bidder with type \vec{i} for item j is simply $\pi_j v_{ij}$. The direct mechanism allocates the item j to bidder i with probability $x_{ij}(\vec{V})$ when all bidders have types \vec{V} . The (IC) constraint implies that the auction is incentive compatible while the (IR) constraint ensures bidder participation. Any direct mechanism M for the multi-item setting is a feasible augmented auction since the seller can use M to presell all item types before observing the real item type and then decide on the final allocation after the type is observed.

The same reduction goes through for converting any Bayes-Nash equilibrium to an equivalent Bayesian incentive-compatible direct mechanism that presells all the items. The difference is that in the (IC) and (IR) constraints above for bidder i , we take an expectation over all bidder types \vec{V}_{-i} other than i . \square

A further extension to the model could allow for correlation among bidder types and item types. The revelation principle is applicable to this setting to show that direct mechanisms which don't reveal any information to the bidders are optimal. However, the resulting direct mechanism is equivalent to a more complicated multi-item auction that allows the seller to charge bidders based on the realization of the item types.