The Power of Concealing the Number of Bidders

Marios Mertzanidis

Purdue University, West Lafayette, IN 47907, USA
mmertzan@purdue.edu

Abstract. We consider the problem of a revenue-maximizing seller, who wants to sell m indivisible items to strategic bidders with additive preferences. A standard assumption in this setting is that when a bidder interacts with the seller's mechanism she knows her competitors (but perhaps doesn't know their private information). In this paper, we study the seller's problem without this assumption. In our setting, bidders first draw their types from some known prior distribution, and then a random set of bidders A arrives and participates in the mechanism. We prove that, even when bidders arrive independently and have independent valuations, the optimal mechanism in our setting can extract more revenue compared to the optimal mechanism in the standard setting. Furthermore, for dependent bidder valuations but independent arrivals, this gap can be made arbitrarily large. On the flip side, we give sufficient conditions under which the two settings are equivalent. Finally, we prove a generalized version of Border's theorem for our problem, which might be of independent interest.

Keywords: Algorithmic Game Theory, Mechanism Design, Revenue Maximization, Incentive Compatibility, Stochastic Arrivals.

1 Introduction

Revenue maximization is a fundamental problem in Algorithmic Economics, and finding mechanisms that optimize for revenue extraction has been a focal point of research interest in recent years. In the standard setup, a seller has a set of indivisible items to sell to a group of strategic bidders. The seller's goal is to maximize her (expected) revenue, while the bidders are interested in maximizing their (expected) utility. Importantly, when bidders interact with the seller's mechanism, it is assumed that they have perfect knowledge of who their competitors are. However, one can imagine many settings where this assumption is unrealistic and impractical. Specifically, in many practical situations of interest (e.g., ad auctions), when a bidder reports her type to the mechanism she only has stochastic information about the total number and the identities of participants.

In this paper, we adopt the stochastic model, introduced by McAfee et al. [31]. In this setting a pool of bidders B exists, and a subset of them, A, will actually participate in the auction with some probability q_A. Notice that subsets have arbitrary probabilities which induce a correlation among the arrival of bidders, which makes this setting quite general. We consider additive, quasi-linear bidders where each bidder's type, t_i, is drawn from a known probability distribution $\mathcal{D}_i$. In the vast majority of this paper, we will assume that bidders' values are drawn independently. The auctioneer aims to extract the maximum possible revenue, and to achieve this, selects a mechanism $\mathcal{M} = \{\mathcal{M}^A\}_{A \subseteq B}$ through which the items will be sold. Each $\mathcal{M}^A$ represents the exact mechanism that will be used when A is the active set of bidders. Once the auctioneer commits to a mechanism $\mathcal{M}$, bidders arrive. We assume that the choice of mechanism does not affect the arrivals. The active bidders then, report their types without knowing the identities of their competitors. Finally, the auctioneer announces the allocation of items and the required payments.

Our setting, calls for notions of incentive compatibility beyond BIC and DSIC. First, we consider Open Bayesian Incentive Compatibility (O-BIC), which assumes that information on who participates in the auction is publicly available, following the traditional approach. The second is Concealed Bayesian Incentive Compatibility (C-BIC), which characterizes the new setting we consider.

We assume that bidders may not be interested in the auctioned items, allowing them to bid zero and receive/pay nothing. An astute reader will realize that non-participating bidders are indistinguishable from participating bidders who are not interested in the auctioned items. Hence, the C-BIC setting can be viewed as the traditional setting with structured correlation between bidders' valuations. Similarly, the O-BIC setting can be seen as a mechanism design problem with structured correlation between valuations, where bidders have partial information on each other's valuations. Mechanism design with dependent priors is a challenging problem, as evidenced by cautionary tales such as the auction of Cremer and McLean [12] or the construction of Hart and Nisan [21]. The above observation makes the problem at hand technically even more interesting. It not only aligns with many modern-day auctions but also provides an opportunity to study mechanism design under structured dependence on valuations.

McAffee et. al. [31] demonstrated that in single-item auctions, the optimal O-BIC and C-BIC mechanisms are equivalent, leading to the natural question: is concealing bidders' participation beneficial for the auctioneer in *multi-item* settings, and if so, can we compute the optimal C-BIC mechanism? Furthermore, it is important to understand whether O-BIC and C-BIC are equivalent in other auction settings. In this paper, we aim to address these questions.

1.1 Our Contribution and Techniques

Our work is comprised of three distinct parts. First, we motivate the problem by showing that, even in very simple settings with independent arrivals and valuations, this reduction in the perceived "power" of bidders allows the auctioneer to increase her expected revenue. We further show that if valuations are drawn from a correlated prior distribution, the optimal C-BIC mechanism can outperform the optimal O-BIC by an arbitrary amount.

In the second part of this paper, we formalize sufficient conditions under which the two notions of incentive compatibility are equivalent. We employ Lagrangian Duality and extend the framework of Cai et. al. [7] in our setting.

Lagrangian Duality enables us to provide a framework on how to show the equivalence of O-BIC and C-BIC mechanisms. In fact, using this framework we recover the result of McAffee et. al. [31], that in single parameter auctions, O-BIC and C-BIC are equivalent. Additionally, we show the equivalence of O-BIC and C-BIC in the two-item setting examined by Yao [38]. Our analysis also implicitly recovers one of Yao's [38] primary results, a tight upper bound on the expected revenue of BIC mechanisms in this two-item setting. As a byproduct of our analysis, we also get that ex-post and ex-interim individual rationality are equivalent in the C-BIC setting.

Finally, we flesh out a generalization of Border's theorem [2] that applies to the C-BIC setting. Leveraging this result we prove a proposition that exponentially decreases the number of constraints needed to check the feasibility of an ex-interim allocation rule.

1.2 Other Related Work

The Economics community has extensively studied auctions with a stochastic arrival of bidders, typically focusing on endogenous entries where bidders may choose to refrain from participating in an auction due to an entry fee ([35], [18], [1], [32]). Previous literature has also investigated the equilibria of common auctions, including the first and second-price auction from both a theoretical ([31],[10], [24], [23], [22], [25], [20]) and empirical perspective ([10], [35], [1], [26]). Additionally, settings with a stochastic number of participants have been studied in contests where players compete for a prize by making an expenditure and winning with a probability determined by the contest ([27], [33], [16], [1]). Despite the considerable interest in these topics, unfortunately, they have received little to no attention from the Computer Science community.

Our work bears many similarities to a prominent research thread in the literature: the comparison of Bayesian Incentive Compatible (BIC) and Dominant-Strategy Incentive Compatible (DSIC) mechanisms. As demonstrated by Myerson's [34] seminal work the optimal DSIC and BIC mechanisms are equivalent in single-parameter settings. Beyond the single-dimensional setting, Yao [38] established a gap between BIC and DSIC mechanisms by computing explicit solutions for non-trivial families of two-item auctions. Additionally, Yao [37] showed that when valuations are independent, the optimal revenue of a DSIC mechanism is at least a constant fraction of that achievable by the optimal BIC mechanism. Fu et. al. [17] showed that when valuations are dependent then this gap can become arbitrarily large, however, this is not the case when prices are non-negative.

We will be using Lagrangian Duality which is a highly effective tool for tackling challenging questions in the broader field of mechanism design. Finding simple and approximately optimal mechanisms in various settings ([7], [8], [17], [9], [15], [4], [19]), bounding the competition complexity of auctions ([14], [28]), designing persuasion policies ([13]), building neural network architectures to find optimal auctions ([36]), are just a few examples of such topics.

On the computational front, originally Border [2] established necessary and sufficient conditions for the feasibility of bidder-symmetric reduced forms. This work was built upon the previous work of Maskin et. al. [29] and Matthews [30]. It was later generalized to asymmetric bidders by Border [3] and Che et. al. [11]. This characterization of Border's theorem enabled Cai et. al. [5] to develop an efficiently computable separation oracle for determining whether reduced forms are feasible. This separation oracle combined with ellipsoid solvers, produces a PTAS algorithm for finding the optimal mechanism. They also provided a different generalization of Border's Theorem for correlated bidders' valuations. Building upon this work Cai et. al. [6] were later able to create separation oracles for reduced forms with arbitrary feasibility constraints. It remains an open question whether our results can produce computationally efficient separations oracles for our problem.

2 Preliminaries

Multi-item Auctions with Stochastic Number of Bidders. We consider the problem of a revenue-maximizing seller with m indivisible items for sale. We assume that there exists a ground set B of $N = |B|$ bidders. For each subset of bidders $A \subseteq B$ there exists some probability q_A that A will be the actual set of bidders that will participate in the auction. Our goal is to create a generalized mechanism $\mathcal{M} = \{\mathcal{M}^A\}_{A \subseteq B}$ where each sub-mechanism $\mathcal{M}^A$ represents the mechanism that is going to be used if A is the set of active bidders that participates in the auction. The sub-mechanisms are defined by their allocation and payment rule (i.e. $\mathcal{M}^A = (x^A(\cdot), p^A(\cdot))$). The allocation rule $x^A(\cdot) : \mathbb{R}_+^{m|A|} \to [0,1]^{m|A|}$ will dictate who takes each item. Note that the allocation rule $x_{i,j}(\cdot)$ outputs values between 0, 1 which represent the probability with which bidder i will receive item j. The payment rule $p^A(\cdot) : \mathbb{R}_+^{m|A|} \to \mathbb{R}^{|A|}$ dictates how much each bidder will pay. Each bidder $i \in B$ will have a private type $t_i \in \mathcal{V}_i \subset \mathbb{R}_+^m$. Active bidders (i.e. $\in A$) will report a value (bid) $v_i \in \mathcal{V}_i$ to the mechanism. Since we will only consider truthful mechanism $t = v$, so from now on we will use bids and types interchangeably. Furthermore, define for each $A \subseteq B$, $\mathcal{V}_A = \{\mathcal{V}_i\}_{i \in A}$. We will consider additive, quasilinear bidders. More precisely, the utility of bidder i for sub-mechanism $\mathcal{M}^A$ is $u_i^A(v_i, v_{-i}) = \sum_{j \in [m]} v_{i,j} \cdot x_{i,j}^A(v_i, v_{-i}) - p_i^A(v_i, v_{-i})$ when she reports v_i and the rest bidders (i.e. $A/\{i\}$) report v_{-i} . We assume that the values of the bidders are drawn from a distribution $\mathcal{D}$ supported on $\mathbb{R}^{m \times N}$. Note that in the general case, the values of the bidders of the ground set can be arbitrarily correlated. Unless stated otherwise, we will assume that the value of each bidder i is drawn independently from a prior distribution $\mathcal{D}_i$ (thus $\mathcal{D} = \times_{i \in B} \mathcal{D}_i$). We will call $\mathcal{D}^A$ the probability distribution from which the valuations of bidders in A are drawn.

Revenue. Our main goal will be to create mechanisms that maximize the expected revenue the auctioneer extracts. The revenue of a sub-mechanism $\mathcal{M}^A$ under distribution $\mathcal{D}^A$ when A is the subset of bidders that participate in the auction, is simply $Rev_A(\mathcal{M}^A, \mathcal{D}^A) = \mathbb{E}_{v \sim \mathcal{D}^A}\left[\sum_{i \in A} p_i^A(v)\right]$. The revenue of mechanism $\mathcal{M}$ under distribution $\mathcal{D}$ is $Rev(\mathcal{M}, \mathcal{D}) = \mathbb{E}_A\left[Rev_A(\mathcal{M}^A, \mathcal{D}^A)\right]$ where the expectation is over the randomness of the set of active bidders.

Individual Rationality. We will only consider mechanisms that are ex-post individually rational. This means that if bidders are truthful then they cannot get negative utility. More formally, for each $A \subseteq B$, each bidder i, and each possible valuation profile $v \in \mathcal{V}^A$ it is true that $u_i^A(v) \geqslant 0$.

Incentive Compatibility. First, let's define an already existing notion of truthfulness for mechanisms where the identities of the participants are known. A mechanism is *Bayesian Incentive Compatible (BIC)* if for every bidder i, possible valuation $v_i \in \mathcal{V}_i$ and possible miss-report $v_i' \in \mathcal{V}_i$ we have that:

$$\mathop{\mathbb{E}}_{v_{-i} \sim \mathcal{D}_{-i}} \left[\sum_{j \in [m]} v_{i,j} \cdot x_{i,j}(v_i, v_{-i}) - p_i(v_i, v_{-i}) \right] \geqslant \mathop{\mathbb{E}}_{v_{-i} \sim \mathcal{D}_{-i}} \left[\sum_{j \in [m]} v_{i,j} \cdot x_{i,j}(v_i', v_{-i}) - p_i(v_i', v_{-i}) \right]$$

We will define two different notions of truthfulness in our stochastic setting, the Open Bayesian Incentive Compatibility (O-BIC) and the Concealed Bayesian Incentive Compatibility (C-BIC). In the O-BIC setting the bidders will know who their competitors are (the information is "open"). They do not know the valuations of their competitors. Thus the problem of finding an O-BIC mechanism reduces to finding a BIC mechanism for each feasible subset of bidders separately. In the C-BIC setting the bidders do not know who their competitors are when submitting their bid (the information is "concealed"). They only have the probabilities with which each subset of bidders arrives. The bidders, in this scenario, do not misreport if they cannot increase their expected utility (where the expectation is taken over the randomness of who their competitors are and what are their valuations). Thus, in this case, each sub-mechanism $\mathcal{M}^A$ need not be a BIC mechanism for the corresponding subset of bidders (as opposed to O-BIC).

More formally a mechanism is *Open Bayesian Incentive Compatible (O-BIC)* if for every set of active bidders $A \subseteq B$ for every bidder $i \in A$, every valuation $v_i \in \mathcal{V}$ and possible miss-report $v_i' \in \mathcal{V}$ we have that:

$$\mathop{\mathbb{E}}_{v_{-i} \sim \mathcal{D}^{A/\{i\}}} \left[\sum_{j \in [m]} v_{i,j} \cdot x_{i,j}^A(v_i, v_{-i}) - p_i^A(v_i, v_{-i}) \right] \geqslant \mathop{\mathbb{E}}_{v_{-i} \sim \mathcal{D}^{A/\{i\}}} \left[\sum_{j \in [m]} v_{i,j} \cdot x_{i,j}^A(v_i', v_{-i}) - p_i^A(v_i', v_{-i}) \right]$$

A mechanism is *Concealed Bayesian Incentive Compatible (C-BIC)* if for every bidder $i \in B$, every valuation $v_i \in \mathcal{V}$ and possible miss-report $v_i' \in \mathcal{V}$ we have that:

$$\mathop{\mathbb{E}}_{\substack{(A|i)\\ v_{-i}\sim\mathcal{D}^{A/\{i\}}}}\left[\sum_{j\in[m]} v_{i,j}\cdot x^A_{i,j}(v_i,v_{-i}) - p^A_i(v_i,v_{-i})\right] \geqslant \mathop{\mathbb{E}}_{\substack{(A|i)\\ v_{-i}\sim\mathcal{D}^{A/\{i\}}}}\left[\sum_{j\in[m]} v_{i,j}\cdot x^A_{i,j}(v'_i,v_{-i}) - p^A_i(v'_i,v_{-i})\right]$$

Where the expectation is now taken over both the randomness of the set of active bidders and the randomness of the valuations of the other bidders (i.e. $\mathbb{E}_{(A|i),v_{-i}\sim\mathcal{D}^{A/\{i\}}}[\cdot] = \mathbb{E}_{(A|i)}\left[\mathbb{E}_{v_{-i}\sim\mathcal{D}^{A/\{i\}}}[\cdot]\right]$). Notice that $\mathbb{P}[A|i]=0$ whenever $i\notin A$ and thus the above expression is well defined. It is easy to see that any O-BIC mechanism is also C-BIC and thus the optimal O-BIC mechanism cannot extract more revenue than the optimal C-BIC mechanism in expectation.

3 Motivating the problem

In this section, we will motivate our problem by demonstrating that even in very simple instances C-BIC extracts in expectation more revenue than O-BIC. We will furthermore show examples under which concealing the number of participants can yield arbitrarily larger revenue than revealing this information.

Establishing a gap between O-BIC and C-BIC. To solidify this intuition we have the following Theorem:

Theorem 1. *There exist instances (even with independent valuations and arrivals) where the optimal C-BIC extracts on expectation more revenue than the optimal O-BIC.*

Proof. We will prove the above theorem by providing a specific example. Consider an instance with two bidders and two items. The bidders will have independent valuations and will arrive with independent probabilities. Specifically, the first bidder will value each item independently and uniformly from $\{1,2\}$, and arrive with probability 1. The second bidder will always value each item at 1.5 and he will arrive with a probability of 1/2. More analytically, $\mathbb{P}[v_{11}=1]=1/2$, $\mathbb{P}[v_{11}=2]=1/2$, $\mathbb{P}[v_{12}=1]=1/2$, $\mathbb{P}[v_{12}=2]=1/2$, $\mathbb{P}[v_{21}=1.5]=1$, $\mathbb{P}[v_{22}=1.5]=1$, $q_{\{1\}}=1/2$, and $q_{\{1,2\}}=1/2$.

We will first examine the performance of the optimal O-BIC mechanism. As discussed earlier we essentially have to find the optimal BIC mechanism for each subgroup that may arrive. It is not difficult to see that when bidder 1 arrives on his own, selling both items as a bundle at a price of 3 is the optimal mechanism and gets an expected revenue of $9/4$ (we will implicitly prove this in Section 4). When both bidders arrive it is not difficult to see that we can extract the full social welfare. More precisely bidder 1 will get the items he had a bid of 2, for a price of 2, and bidder 2 will get the remaining items for a price of 1.5 each. Thus, on expectation, we extract 3 revenue. The total O-BIC mechanism extracts an expected revenue of 21/8.

Now we will present a feasible C-BIC mechanism in this setting that extracts on expectation, more than 21/8 revenue. In the case where both bidders arrive, we will keep the same mechanism that extracts the expected revenue of 3. In the case where bidder 1 arrives on his own, if he bids (1,1) he gets both for a price of 2. For any other bid, he gets both items for a price of 2.5. The expected revenue of this mechanism is 9.5/4. It is not difficult to check that the proposed mechanism is C-BIC. Essentially, bidder 1 has no incentive to lie because although when he is alone he can benefit by miss-reporting a bid of (1,1), half of the time he will lose both items. The expected revenue of this mechanism is 21.5/8.

□

Arbitrary large Gap: Dependent valuations. We now consider the case where the valuations for all bidders are drawn from some common probability distribution $\mathcal{D}$ (This will be the only case we consider such instances). Note that we only get to observe the bids (valuations) of the bidders that arrived.

Theorem 2. *When bidders' valuations are drawn from a common probability distribution, C-BIC can extract arbitrarily larger revenue than O-BIC (even in instances where bidders arrive independently).*

Proof. We will have two bidders and one item. Bidder's 1 valuation (v_i) is dictated by the equal revenue distribution. In other words, the CDF of bidder 1 is:

$$F(v_1) = \begin{cases} 1 - \frac{1}{v_1}, & \text{for } v_1 \in [1, H) \\ 1, & \text{if } v_1 = H \end{cases}$$

This bidder arrives with probability 1. Bidder 2 arrives with probability ε and values the item $v_2 = \delta \cdot v_1$. The idea here is that we will take $\delta \to 0$, thus bidder two will not contribute to the total revenue, however, when he arrives we will know exactly what bidder's 1 valuation is.

First, let's consider the O-BIC setting. It is not difficult to see that when bidder 1 arrives on his own the only available mechanisms are posted-price mechanisms and any posted price will extract expected revenue 1. When both bidders arrive we charge bidder 1 a price of $\frac{v_2}{\delta}$ for the item. This extracts on expectation $O(log\, H)$ revenue. Thus the optimal O-BIC mechanism extracts $O(1 + \varepsilon \cdot log\, H)$ revenue.

Now we will present a feasible C-BIC mechanism in this setting that extracts $\Omega((1-\varepsilon) log\, log\, H)$. When bidder 1 arrives on his own then he gets the item with probability $\frac{log\, v_1}{log\, H}$ and pays $\frac{v_1 \cdot log\, v_1}{log\, H}$ (i.e. $x_1^{\{1\}}(v_1) = \frac{log\, v_1}{log\, H}$ and $p_1^{\{1\}}(v_1) = \frac{v_1 \cdot log\, v_1}{log\, H}$). Thus bidder 1 gets 0 utility when he arrives on his own. When both bidders arrive then bidder 2 always gets nothing and pays nothing. If bidder 1 reports anything different than $\frac{v_2}{\delta}$ he gets nothing and pays nothing, if $v_1 = \frac{v_2}{\delta}$

then (define $g(z) = \max_y \frac{log\, y}{log\, H}(z-y)$), bidder 1 is payed an amount equal to $\frac{1-\varepsilon}{\varepsilon} g(v_1)$ (i.e. $p_1^{\{1,2\}}(v,v) = -\frac{1-\varepsilon}{\varepsilon} g(v)$). Obviously, this mechanism is IR. The mechanism is also C-BIC. This follows from the fact that when bidder 1 reports truthfully he extracts utility $(1-\varepsilon)\cdot 0 + \varepsilon \cdot \frac{1-\varepsilon}{\varepsilon} g(v) = (1-\varepsilon)\cdot g(v)$. When he misreports and bids w he has an expected utility of $(1-\varepsilon)\cdot(v\frac{log\, w}{log\, H} - \frac{w \cdot log\, w}{log\, H}) + \varepsilon \cdot 0 \leqslant (1-\varepsilon)\cdot g(v)$ where the last inequality is due to the construction of $g(v)$. □

The above example is an adaptation of an example proposed by Fu et. al. [17]. We have presented the main ideas for the sake of completion. For a detailed explanation of why this mechanism extracts $\Omega((1-\varepsilon) log\ log\ H)$ revenue on expectation, we refer the reader to [17].

4 On the Equivalence of O-BIC and C-BIC

Our main focus in this section will be to show how to use Lagrangian duality to prove that in certain settings the optimal O-BIC and C-BIC mechanisms are equivalent. As a byproduct of our analysis, we will be modernizing the proof of McAfee et. al. [31], and show how to derive some of Yao's [38] main results using these techniques.

For compactness of notation define $u_i^A(v_i' \leftarrow v_i, v_{-i}) = \sum_{j\in[m]} v_{i,j} \cdot x_{i,j}^A(v_i', v_{-i}) - p_i^A(v_i', v_{-i})$ which represents the utility of bidder i when his true valuations is v_i but he reports v_i', the rest of the bidders report v_{-i} and the set of active bidders is A. Now we can write the LP formulation of our C-BIC problem:

$$\begin{array}{llcll}
\max\limits_{x,p} & \sum_{A\subseteq B} q_A \sum_{v\in\mathcal{V}_A} f_A(v) \sum_{i\in A} p_i^A(v) & & & \\
\text{s.t.} & \mathbb{E}\left[u_i^A(v_i \leftarrow v_i, v_{-i})\right] & \geqslant & \mathbb{E}\left[u_i^A(v_i' \leftarrow v_i, v_{-i})\right], & (1)\\
 & u_i^A(v_i \leftarrow v_i, v_{-i}) & \geqslant & 0, & (2)\\
 & \sum_{i\in A} x_{i,j}^A(v) & \leqslant & 1, & (3)\\
 & x & \geqslant & 0 &
\end{array}$$

Where in (1) the expectations are taken over the randomness of the active set of bidders A, and the valuation of the other bidders v_{-i}. (1) is the C-BIC constraint and holds $\forall i \in B, (v_i, v_i') \in \mathcal{V}_i^2$. (2) is the ex-post IR constraint and holds $\forall A \subseteq B, i \in A, v_i \in \mathcal{V}_i, v_{-i} \in \mathcal{V}_{-i}$. Finally, (3) is the feasibility constraint and holds $\forall A \subseteq B, v \in \mathcal{V}_A$. We will explore the techniques presented by Cai et. al. [7] and derive an upper bound for the expected revenue of the optimal C-BIC mechanism by using Lagrangian duality. We postpone the exact expressions and calculations to Appendix A.

From duality, we know that for any choice of λ, μ the revenue of our C-BIC mechanism is upper bounded by $\max_{x,p} \mathcal{L}(x,p,\lambda,\mu)$ (subject to the constraints we have not Lagrangified). In fact, from strong duality, we know that the optimal solution of our C-BIC mechanism is equal to $\min_{\lambda,\mu} \max_{x,p} \mathcal{L}(x,p,\lambda,\mu)$. By observing that $p_i^A(v_i, v_{-i})$ is an unrestricted variable, in order to have a meaningful bound, it must be the case that for all $A \subseteq B : q_A \neq 0$, $v_{-i} \in \mathcal{V}_{A/\{i\}}$: $f_{A/\{i\}}(v_{-i}) \neq 0$, $i \in A$, $v_i \in \mathcal{V}_i$:

$$f_i(v_i) + \sum_{v_i' \in \mathcal{V}_i} \lambda^i(v_i', v_i) - \sum_{v_i' \in \mathcal{V}_i} \lambda^i(v_i, v_i') - \frac{\mu_A^i(v_i, v_{-i})}{q_A \cdot f_{A/\{i\}}(v_{-i})} = 0$$

Since this is true for all A, v_{-i} it must be the case that $\frac{\mu_A^i(v_i,v_{-i})}{q_A \cdot f_{A/\{i\}}(v_{-i})}$ must not depend on A, v_{-i} [1]. So from now on we simplify $\frac{\mu_A^i(v_i,v_{-i})}{q_A \cdot f_{A/\{i\}}(v_{-i})} = \mu^i(v_i)$. The set of equations of the form $f_i(v_i) + \sum_{v_i' \in \mathcal{V}_i} \lambda^i(v_i', v_i) - \sum_{v_i' \in \mathcal{V}_i} \lambda^i(v_i, v_i') - \mu^i(v_i) = 0$ characterize for each i a flow where each v_i is a vertex that receives flow $f_i(v_i)$ from the source, receives $\lambda^i(v_i', v_i)$ flow from node v_i', sends $\lambda^i(v_i, v_i')$ flow to node v_i', and sends $\mu^i(v_i)$ flow to the sink. Using the above we can simplify our Lagrangian expression $\mathcal{L}(x, p, \lambda, \mu)$:

$$\sum_{\substack{A \subseteq B, i \in A, \\ v_i \in \mathcal{V}_i, v_{-i} \in \mathcal{V}_{A/\{i\}}, \\ j \in [m]}} q_A \cdot f_A(v_i, v_{-i}) \cdot x_{i,j}^A(v_i, v_{-i}) \left(v_{i,j} - \frac{1}{f_i(v_i)} \sum_{v_i' \in \mathcal{V}_i} \lambda^i(v_i', v_i)(v_{i,j}' - v_{i,j}) \right)$$

Thus by choosing any flow, we can derive an upper bound on the optimal C-BIC revenue. We continue by stating the following observation which was hinted at during the above analysis:

Observation 1. *The optimal ex-interim IR C-BIC mechanism is equivalent to the optimal ex-post IR C-BIC mechanism.*

Proof. Here we present a proof sketch of the above observation. An ex-post IR, C-BIC mechanism is a feasible ex-interim IR, C-BIC mechanism. Thus the revenue of the optimal ex-interim IR, C-BIC mechanism is equal or greater than the revenue of the optimal ex-post IR, C-BIC mechanism. On the other hand, if we were to write the Lagrangian Dual of ex-interim IR, C-BIC mechanism the only difference with the ex-post variation would be that instead of $\frac{\mu_A^i(v_i,v_{-i})}{q_A \cdot f_{A/\{i\}}(v_{-i})}$ at the flow constraint we would have a $\mu^i(v_i)$ term. However, as discussed above in the ex-post case the term $\frac{\mu_A^i(v_i,v_{-i})}{q_A \cdot f_{A/\{i\}}(v_{-i})}$ cannot depend on A or v_{-i}. Thus any flow of the ex-post Lagrangian is a feasible flow for the ex-interim one. Using this argument for the optimal ex-post Lagrangian flow we get that the revenue of the optimal ex-post IR, C-BIC mechanism is an upper bound on the revenue of the optimal ex-interim IR, C-BIC mechanism. Combining the two observations we get that the optimal revenue of the two variants is the same which concludes the proof.

□

In the next two sections we will prove that in single-parameter settings and in the setting studied by Yao [38], O-BIC and C-BIC are equivalent.

[1] This effectively proves that ex-post and ex-ante IR in our setting our equivalent.

4.1 Singe Item Auctions

Theorem 3. *In single-item settings with arbitrarily dependent arrivals, the optimal C-BIC mechanism is equivalent to the optimal O-BIC mechanism.*

Proof. For single-item auctions the Lagrangian $\mathcal{L}(x, p, \lambda, \mu)$ simplifies even further to:

$$\sum_{\substack{A \subseteq B, i \in A, \\ v_i \in \mathcal{V}_i, v_{-i} \in \mathcal{V}_{A/\{i\}}}} q_A \cdot f_A(v_i, v_{-i}) \cdot x_i^A(v_i, v_{-i}) \left(v_i - \frac{1}{f_i(v_i)} \sum_{v_i' \in \mathcal{V}_i} \lambda^i(v_i', v_i)(v_i' - v_i) \right)$$

For any valuation $v_i \in \mathcal{V}_i$ define as $v_i^- = sup\{x : x < v_i, x \in \mathcal{V}_i\}$. Essentially v_i^- is the next lower number in $\mathcal{V}_i$ after v_i. If that does not exist then set $v_i^- = \varnothing$. Similarly define $v_i^+ = inf\{x : x > v_i, x \in \mathcal{V}_i\}$. We define the flow as follows, every node v_i sends all the flow it receives at node v_i^-. If $v_i^- = \varnothing$ send all the flow to the sink. This is obviously a feasible flow and simplifies our expression to:

$$\sum_{\substack{A \subseteq B, i \in A, \\ v_i \in \mathcal{V}_i, v_{-i} \in \mathcal{V}_{A/\{i\}}}} q_A \cdot f_A(v_i, v_{-i}) \cdot x_i^A(v_i, v_{-i}) \left(v_i - \frac{\sum_{v_i' \in \mathcal{V}_i : v_i' > v_i} f_i(v_i)}{f_i(v_i)} (v_i^+ - v_i) \right)$$

However observe that $v_i - \frac{\sum_{v_i' \in \mathcal{V}_i : v_i' > v_i} f_i(v_i)}{f_i(v_i)} (v_i^+ - v_i)$ is the discrete analog of Myerson's [34] virtual welfare which we will represent as $\phi_i(v_i)$. We implicitly assume regular distributions, however, this assumption comes without loss of generality since we can always create loops as described by Cai et. al. [7] to iron the distribution. If we denote by OPT_C the optimal revenue achieved in our C-BIC setting, we have concluded that:

$$OPT_C \leqslant \sum_{A \subseteq B} q_A \cdot \sum_{v \in \mathcal{V}_A} f_A(v) \sum_{i \in A} x_i^A(v) \phi_i(v) = \sum_{A \subseteq B} q_A \cdot \mathop{\mathbb{E}}_{v \sim \mathcal{D}^A} \left[x_i^A(v) \phi_i(v_i) \right]$$

However, this is exactly the revenue we would achieve if we run Myerson's auction for all possible subsets of active bidders. This mechanism is obviously C-BIC and O-BIC, thus, from the fact that C-BIC extracts more revenue than O-BIC, we conclude that for single-item auctions the two settings are equivalent.

□

4.2 Bi-valued i.i.d. bidders

In this section, we will examine the setting of Yao [38]. The reason why this setting is particularly interesting is that it is the setting that established a gap between BIC and DSIC mechanisms. In this setting all bidders are i.i.d. . We have two items and each bidder independently values each item at a with probability p and b with probability $1 - p$ (i.e. $\mathbb{P}[v_{i,j} = a] = p$, $\mathbb{P}[v_{i,j} = b] = 1 - p$). From now on we will refer to this setting as the "Bi-Valued" setting.

Definition 1. *Let $p_0 = 1 - p^n$, $p_1 = p^{2n}$ and $p_2 = p^n(1-p^n)$ then:*

$$r_{OB}(n) = 2 \cdot p_0 \cdot b + 2 \cdot p_1 \left[a - \frac{1-p^2}{2p^2}(b-a)\right]_+ + 2 \cdot p_2 \left[a - \frac{1-p}{2p}(b-a)\right]_+$$

where $[a]_+ = max\{a, 0\}$.

Note that:

1. p_0 is the probability that for item j there exists at least one bidder i^* such that $v_{i^*,j} = b$.
2. p_1 is the probability that all bidders have valuations (a, a).
3. p_2 is the probability that for item j and every bidder i, $v_{i,j} = a$ and for the other item j' there exists a bidder i^* such that $v_{i^*,j'} = b$.

Lemma 1. *The optimal O-BIC mechanism in this setting extracts expected revenue equal to:*

$$\sum_{A \subseteq B} q_A \cdot r_{OB}(|A|)$$

Proof. Yao [38] has proved that the optimal BIC mechanism, in this setting, where n bidders are participating is $r_{OB}(n)$. As we have discussed already, finding the optimal O-BIC mechanism boils down to finding the optimal BIC mechanism for each $A \subseteq B$. The result follows.

□

Theorem 4. *In the "Bi-Valued" setting with arbitrarily dependent arrivals, the optimal C-BIC mechanism is equivalent to the optimal O-BIC mechanism.*

Proof. In our "Lagrangian flow graph" we will have for each i a total of 4 nodes. Define the following flow. Node (b, b) receives $(1-p)^2$ flow from the source and sends $\frac{1}{2}(1-p)^2$ flow to node (a, b) and $\frac{1}{2}(1-p)^2$ flow to node (b, a). Both nodes (b, a) and (a, b) receive a total of $\frac{1}{2}(1-p^2)$ flow and send it to node (a, a). Finally, (a, a) sends all the flow it receives to the sink. The exact flow is depicted in Fig. 1.

Given this flow, Table 1 summarizes the values of $v_{i,j} - \frac{1}{f_i(v_i)} \sum_{v'_i \in \mathcal{V}_i} \lambda^i(v'_i, v_i)(v'_{i,j} - v_{i,j})$ for different items and different v_i. Notice that $b > a - \frac{1-p}{2p}(b-a) > a - \frac{1-p^2}{2p^2}(b-a)$. Combining this information with the feasibility constraint $\sum_{i \in A} x^A_{i,j}(v) \leqslant 1$, $\forall A \subseteq B, v \in \mathcal{V}_A$ we now can predict exactly the values of $x^A_{i,j}(v)$. For every $v \in \mathcal{V}_A$, if for item $j \in [2]$ there exists $i^* \in A$ such that $v_{i^*,j} = b$ then $x^A_{i^*,j}(v) = 1$ and for all $i \in A/\{i^*\}$, $x^A_{i,j}(v) = 0$. If $a - \frac{1-p}{2p}(b-a) \leqslant 0$ then all other $x^A_{i,j}(v) = 0$. If $a - \frac{1-p}{2p}(b-a) > 0$ then for any $v \in \mathcal{V}_A$ and item j such that $\nexists i^* : v_{i^*,j} = b$ but $\exists \hat{i} : v_{\hat{i},j'} = b$ (where j' is the other item) then $x^A_{\hat{i},j}(v) = 1$ and for all $i \in A/\{\hat{i}\}$ it must be that $x^A_{i,j}(v) = 0$. Lastly, if $a - \frac{1-p^2}{2p^2}(b-a) \leqslant 0$ then all other $x^A_{i,j}(v) = 0$. If $a - \frac{1-p}{2p}(b-a) > 0$ then for any

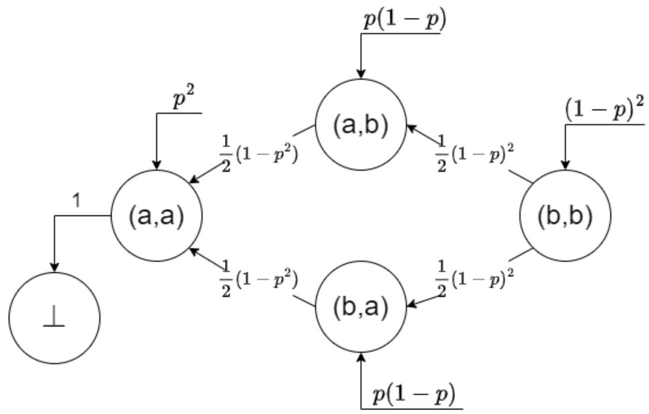

Fig. 1. Flow in Bi-valued setting.

$v \in \mathcal{V}_A$ where all valuations in v are of the form (a, a) then some (any choice is equivalent) $i^* \in A$, $x^A_{i^*,1}(v) = x^A_{i^*,2}(v) = 1$ and for all $i \in A/\{i^*\}$ it must be that $x^A_{i,1}(v) = x^A_{i,2}(v) = 0$. Inputting these observations into our Lagrangian we get an upper bound of $\sum_{A \subseteq B} q_A \cdot r_{OB}(|A|)$.

Profiles	Item 1	Item 2
(a,a)	$a - \frac{1-p^2}{2p^2}(b-a)$	$a - \frac{1-p^2}{2p^2}(b-a)$
(a,b)	$a - \frac{1-p}{2p}(b-a)$	b
(b,a)	b	$a - \frac{1-p}{2p}(b-a)$
(b,b)	b	b

Table 1. "Virtual Values" of described flow.

However, running the optimal BIC mechanism identified by Yao [38] is O-BIC and C-BIC and achieves the revenue dictated by the above upper bound.

□

5 Generalizing Border's Theorem

Now that we have established that running C-BIC mechanisms can be beneficial for the auctioneer we will provide a Generalization of Border's Theorem ([2]) that applies to our setting. Cai et. al. [5] have provided an explicit characterization of

Border's theorem for correlated items and bidders. Although, as discussed earlier, our problem can be viewed as a special case of correlated mechanism design, we flesh out a generalization that is much simpler and intuitive for our setting. This simpler characterization helps us prove Proposition 1 that decreases the number of constraints by an exponential factor. First, let's define the problem that will be the main focus of this section:

The C-BIC Mechanism Design Problem.
Given as input, a set of m indivisible items, a ground set of bidders B, for each subset $A \subseteq B$ a probability q_A, and for each bidder $i \in B$ a probability distribution $\mathcal{D}_i$ we want to output a mechanism $\mathcal{M}$ that is C-BIC w.r.t. $\{q_A\}_{A \subseteq B}$ and $\mathcal{D} = \bigtimes_{i \in B} \mathcal{D}_i$, ex-post IR, and maximizes the expected revenue of the auctioneer.

Definition 2 (Interim Allocation Rule). *Given a generalized mechanism $\mathcal{M} = \{\mathcal{M}^A\}_{A \subseteq B}$ and for each bidder i probability distributions $\mathcal{D}_i$, and for each subset $A \subseteq B$ probability q_A, then the interim allocation rule is defined as $\pi_{i,j}(v_i) = \mathbb{E}_{(A|i), v_{-i} \sim \mathcal{D}^{A/\{i\}}}\left[x^A_{i,j}(v_i, v_{-i})\right]$. We define the interim payment rule in an equivalent manner.*

For ease of notation lets define $f_i(v) = \mathbb{P}_{\hat{v} \sim \mathcal{D}_i}\left[\hat{v} = v\right]$. Similarly, for each $i \in B$ define $f_{-i}(v)$ and for each $A \subseteq B$ define $f_A(v)$. We further abuse notation by defining $q^i = \sum_{A \subseteq B : i \in A} q_A$ which represents the probability of bidder i becoming active.

From now on we will consider instances with one item. Generalizing too many items will be trivial. Just observe that whether an interim allocation rule is feasible for the allocation of one item, is independent of whether it is feasible for another item. Thus from now on we will omit the subscript j that denotes the item we are considering.

Theorem 5. *Given an interim allocation rule π for the C-BIC mechanism design problem, π is feasible iff for all possible sets $S = \{S_i\}_{i \in B}$ (where $S_i \subseteq \mathcal{V}_i$ is a subset if bidders i possible valuations) it is true that :*

$$\sum_{i=1}^{N} \sum_{v_i \in S_i} q^i f_i(v_i) \pi_i(v_i) \leqslant 1 - \sum_{A \subseteq B} q_A \prod_{i \in A} \left(1 - \sum_{v_i \in S_i} f_i(v_i)\right)$$

where q^i is the probability that bidder i is an active bidder, $f_i(v_i)$ is the probability that bidder i valuation is v_i (when drawn independently from other bidders' valuations), and q_A is the probability that A is the active set of bidders.

Proof. The proof is a generalization of the max-flow/min-cut proof of Border's Theorem. We an analysis inspired by the network flow argument presented by Che et. al. [11]. In the first part of the proof we will argue that for any feasible generalized allocation rule $x = \{x^A(.)\}_{A \subseteq B}$, the induced interim allocation rule satisfies the proposed inequality. Then we will show that given an interim allocation rule that satisfies this inequality, we can always compute a feasible generalized mechanism.

First, let's compute what is the probability that the winner (i^*) has a valuation in S_{i^*}. It is not difficult to see that this probability is described by the following expression:

$$\begin{aligned}\mathbb{P}\left[v_{i^*} \in S_{i^*}\right] &= \sum_{A \subseteq B} q_A \cdot \sum_{i \in A} \sum_{v_i \in S_i} \sum_{v_{-i} \in \mathcal{V}_{-i}} f_i(v_i) \cdot f_{-i}(v_{-i}) \cdot x_i^A(v_i, v_{-i}) \\ &= \sum_{i=1}^{N} \sum_{v_i \in S_i} \sum_{v_{-i} \in \mathcal{V}_{-i}} \sum_{A \subseteq B : i \in A} q_A \cdot f_i(v_i) \cdot f_{-i}(v_{-i}) \cdot x_i^A(v_i, v_{-i}) \\ &= \sum_{i=1}^{N} \sum_{v_i \in S_i} \sum_{v_{-i} \in \mathcal{V}_{-i}} \sum_{A \subseteq B} q_{A,i} \cdot f_i(v_i) \cdot f_{-i}(v_{-i}) \cdot x_i^A(v_i, v_{-i}) \\ &= \sum_{i=1}^{N} \sum_{v_i \in S_i} \sum_{v_{-i} \in \mathcal{V}_{-i}} \sum_{A \subseteq B} q_{(A|i)} \cdot q^i \cdot f_i(v_i) \cdot f_{-i}(v_{-i}) \cdot x_i^A(v_i, v_{-i}) \\ &= \sum_{i=1}^{N} \sum_{v_i \in S_i} q^i f_i(v_i) \pi_i(v_i)\end{aligned}$$

Where $q_{(A,i)}$ is the probability of A being the subset of active bidders and i being an active bidder. Obviously for a feasible generalized mechanism the probability of a bidder winning with a valuation in S is less than the probability of at least one bidder having a valuation in S. This probability is exactly expressed by the right-hand side of our inequality. Thus we have proved the first part of our analysis.

Now consider the following s-t directed graph. We first have the source node s. We also have a set of nodes X that contains a node for each subset of bidders $A \subseteq B$ and each possible set of valuations $v \in \mathcal{V}_A$ for those bidders (represented as (A, v)). We then have a set of nodes Y that contains a node for each bidder $i \in B$ and all possible valuations $v_i \in \mathcal{V}_i$ for this bidder (represented as (i, v_i)). Y also contains an extra "no sell" node (represents the scenario where we do not sell the item), which we represent as $\perp$. Lastly, we have the sink node t. For each node in $(A, v) \in X$ there exists a directed edge from the source to that node with a capacity equal to $q_A \cdot f_A(v)$. Every node in $(A, v) \in X$ (where $v = \{v_i\}_{i \in A}$) has a directed edge of infinite capacity towards all nodes (i, v_i) of Y and $\perp$. From each node $(i, v_i) \in Y$ there exists a directed edge towards t with capacity $q^i \pi_i(v_i) f_i(v_i)$. Lastly, we have a directed edge from $\perp$ to t with capacity $1 - \sum_{i=1}^{N} \sum_{v_i \in \mathcal{V}_i} q^i \pi_i(v_i) f_i(v_i)$. Since we are assuming that our inequality holds then we know that this capacity is non-negative (simply set $\forall i : S_i = \mathcal{V}_i$). The capacity of all directed edges from s to X has a total capacity of 1. When we have a flow of 1 it is easy to see our interim allocation is feasible. Simply set $x_i^A(v)$ to be $\frac{1}{q_A f_A(v)}$ times the amount of flow that passes through the edge that connects (A, v) to (i, v_i). The above allocation rule induces our desired interim allocation rule. The only thing left to prove is that the maximum flow in this graph is 1, which is equivalent to showing that the minimum cut of the graph is 1. Consider a cut that includes nodes s, $Z \subseteq X$, and $T \subseteq Y$. Define $S_i \subseteq \mathcal{V}_i$ to be the set of valuations of bidder i that are not in any valuation profile in Z. All nodes from s to X/Z are cut. Their capacities are at least $1 - \sum_{A \subseteq B} q_A \prod_{i \in A} \left(1 - \sum_{v_i \in S_i} f_i(v_i)\right)$, since as discussed earlier this represents the probability that at least one of the valuations of the bidders come from S. Since edges connecting X to Y have infinite capacity then none of those edges must be cut. So we know that all nodes of the form (i, v_i) with $v_i \notin S_i$ are in T and thus their edge with t must be cut. Also, $\perp$ must be in T since if that was not the case, either the cut would have infinite capacity, or $Z = \varnothing$ in which case the cut is obviously at least 1. The total capacity of edges from Y to t is 1, and thus the total capacity that has been cut is 1 minus the capacity of edges to t that have not been cut. The capacity of those edges as discussed is at most $\sum_{i=1}^{N} \sum_{v_i \in S_i} q^i f_i(v_i) \pi_i(v_i)$. Combining the above we can see that the cut has a capacity of at least 1, concluding our proof.

□

Definition 3 (Constricting Sets). *Let π be an infeasible interim allocation rule for the C-BIC mechanism design problem. Then there must exist a set $S = \{S_i\}_{i \in B}$ (where $S_i \subseteq \mathcal{V}_i$) such that the inequality of Theorem 5 is violated. We will call such a set S a Constricting Set.*

Proposition 1. *If π is an infeasible interim allocation rule for the C-BIC mechanism design problem, then there must exist a constricting set of the form $S = \{S_{x_i}\}_{i \in B}$ where $S_{x_i} = \{v_i \in \mathcal{V}_i : \pi_i(v_i) \geqslant x_i\}$.*

We postpone the proof of Proposition 1 to Appendix B.

6 Conclusion and Future Work

The aim of this paper is to initiate a constructive discussion within the algorithmic game theory community regarding this problem. Our findings reveal some unexpected results. Notably, we have demonstrated that in the Bi-Valued setting, if bidders are i.i.d. but their arrivals are arbitrary, the optimal O-BIC mechanism is equivalent to the optimal C-BIC mechanism. Conversely, if bidders are not i.i.d., but their valuations and arrivals are independent, a gap exists between the two settings. Several fascinating questions remain unanswered:

1. Are there necessary and sufficient conditions under which O-BIC and C-BIC are equivalent?
2. When valuations are independent, is the gap between the two settings bounded?
3. Is there a simple C-BIC mechanism that extracts more revenue than the optimal O-BIC?

Acknowledgements I would like to thank Prof. Alex Psomas (Purdue University) for useful discussions and feedback he provided.

References

1. Boosey, L., Brookins, P., Ryvkin, D.: Information disclosure in contests with endogenous entry: An experiment. Management Science **66**(11), 5128–5150 (2020)
2. Border, K.C.: Implementation of reduced form auctions: A geometric approach. Econometrica: Journal of the Econometric Society pp. 1175–1187 (1991)
3. Border, K.C.: Reduced form auctions revisited. Economic Theory **31**(1), 167–181 (2007)
4. Brustle, J., Cai, Y., Wu, F., Zhao, M.: Approximating gains from trade in two-sided markets via simple mechanisms. In: Proceedings of the 2017 ACM Conference on Economics and Computation. pp. 589–590 (2017)
5. Cai, Y., Daskalakis, C., Weinberg, S.M.: An algorithmic characterization of multi-dimensional mechanisms. In: Proceedings of the forty-fourth annual ACM symposium on Theory of computing. pp. 459–478 (2012)
6. Cai, Y., Daskalakis, C., Weinberg, S.M.: Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In: Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science. p. 130–139. FOCS '12, IEEE Computer Society, USA (2012)
7. Cai, Y., Devanur, N.R., Weinberg, S.M.: A duality-based unified approach to bayesian mechanism design. SIGecom Exch. **15**(1), 71–77 (sep 2016)
8. Cai, Y., Oikonomou, A.: On simple mechanisms for dependent items. In: Proceedings of the 22nd ACM Conference on Economics and Computation. p. 242–262. EC '21, Association for Computing Machinery, New York, NY, USA (2021)
9. Cai, Y., Zhao, M.: Simple mechanisms for subadditive buyers via duality. In: Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing. pp. 170–183 (2017)
10. Canals-Cerdá, J.J., Pearcy, J.: Arriving in time: Estimation of english auctions with a stochastic number of bidders. Journal of Business & Economic Statistics **31**(2), 125–135 (2013)
11. Che, Y.K., Kim, J., Mierendorff, K.: Generalized reduced-form auctions: A network-flow approach. Econometrica **81**(6), 2487–2520 (2013)
12. Crémer, J., McLean, R.P.: Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. Econometrica **53**(2), 345–361 (1985)
13. Dughmi, S., Niazadeh, R., Psomas, A., Weinberg, S.M.: Persuasion and incentives through the lens of duality. In: Web and Internet Economics: 15th International Conference, WINE 2019, New York, NY, USA, December 10–12, 2019, Proceedings 15. pp. 142–155. Springer (2019)
14. Eden, A., Feldman, M., Friedler, O., Talgam-Cohen, I., Weinberg, S.M.: The competition complexity of auctions: A bulow-klemperer result for multi-dimensional bidders. arXiv preprint arXiv:1612.08821 (2016)
15. Eden, A., Feldman, M., Friedler, O., Talgam-Cohen, I., Weinberg, S.M.: A simple and approximately optimal mechanism for a buyer with complements. Operations Research **69**(1), 188–206 (2021)
16. Ewerhart, C., Quartieri, F.: Unique equilibrium in contests with incomplete information. Economic theory **70**(1), 243–271 (2020)
17. Fu, H., Liaw, C., Lu, P., Tang, Z.G.: The value of information concealment. In: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 2533–2544. SIAM (2018)

18. Fu, Q., Jiao, Q., Lu, J.: Contests with endogenous entry. International Journal of Game Theory **44**, 387–424 (2015)
19. Haghpanah, N., Hartline, J.: Reverse mechanism design. In: Proceedings of the Sixteenth ACM Conference on Economics and Computation. pp. 757–758 (2015)
20. Harstad, R.M., Kagel, J.H., Levin, D.: Equilibrium bid functions for auctions with an uncertain number of bidders. Economics Letters **33**(1), 35–40 (1990)
21. Hart, S., Nisan, N., et al.: The menu-size complexity of auctions. Center for the Study of Rationality (2013)
22. Janssen, M., Rasmusen, E.: Bertrand competition under uncertainty. The Journal of Industrial Economics **50**(1), 11–21 (2002)
23. Kong, Y.: Not knowing the competition: evidence and implications for auction design. The RAND Journal of Economics **51**(3), 840–867 (2020)
24. Levin, D., Ozdenoren, E.: Auctions with uncertain numbers of bidders. Journal of Economic Theory **118**(2), 229–251 (2004)
25. Levin, D., Smith, J.L.: Equilibrium in auctions with entry. The American Economic Review **84**(3), 585–599 (1994)
26. Li, T., Zheng, X.: Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions. The Review of Economic Studies **76**(4), 1397–1429 (10 2009)
27. Lim, W., Matros, A.: Contests with a stochastic number of players. Games and Economic Behavior **67**(2), 584–597 (2009)
28. Liu, S., Psomas, C.A.: On the competition complexity of dynamic mechanism design. In: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 2008–2025. SIAM (2018)
29. Maskin, E., Riley, J.: Optimal auctions with risk averse buyers. Econometrica: Journal of the Econometric Society pp. 1473–1518 (1984)
30. Matthews, S.A.: On the implementability of reduced form auctions. Econometrica: Journal of the Econometric Society pp. 1519–1522 (1984)
31. McAfee, R., McMillan, J.: Auctions with a stochastic number of bidders. Journal of Economic Theory **43**(1), 1–19 (1987)
32. Menezes, F.M., Monteiro, P.K.: Auctions with endogenous participation. Review of Economic Design **5**, 71–89 (2000)
33. Münster, J.: Contests with an unknown number of contestants. Public Choice **129**(3-4), 353–368 (2006)
34. Myerson, R.B.: Optimal auction design. Mathematics of Operations Research **6**(1), 58–73 (1981)
35. Palfrey, T.R., Pevnitskaya, S.: Endogenous entry and self-selection in private value auctions: An experimental study. Journal of Economic Behavior & Organization **66**(3), 731–747 (2008)
36. Rahme, J., Jelassi, S., Weinberg, S.M.: Auction learning as a two-player game. arXiv preprint arXiv:2006.05684 (2020)
37. Yao, A.C.C.: An n-to-1 bidder reduction for multi-item auctions and its applications. p. 92–109. SODA '15, Society for Industrial and Applied Mathematics, USA (2015)
38. Yao, A.C.C.: Dominant-strategy versus bayesian multi-item auctions: Maximum revenue determination and comparison. In: Proceedings of the 2017 ACM Conference on Economics and Computation. p. 3–20. EC '17, Association for Computing Machinery, New York, NY, USA (2017)

A Ommited details from Section 4

First, let's rewrite our LP constraint in a more convenient manner:

$$\mathop{\mathbb{E}}_{(A|i),v_{-i}}\left[u_i^A(v_i \leftarrow v_i, v_{-i})\right] \geqslant \mathop{\mathbb{E}}_{(A|i),v_{-i}}\left[u_i^A(v_i' \leftarrow v_i, v_{-i})\right]$$

Opening up the expression we get:

$$\sum_{A\subseteq B}\frac{q_{(A,i)}}{q^i}\sum_{v_{-i}\in\mathcal{V}_{-i}} f_{-i}(v_{-i})u_i^A(v_i \leftarrow v_i, v_{-i}) \geqslant \sum_{A\subseteq B}\frac{q_{(A,i)}}{q^i}\sum_{v_{-i}\in\mathcal{V}_{-i}} f_{-i}(v_{-i})u_i^A(v_i' \leftarrow v_i, v_{-i})$$

By rearranging we get:

$$\sum_{A\subseteq B: i\in A}\sum_{v_{-i}\in\mathcal{V}_{-i}}\left(q_A f_{-i}(v_{-i})u_i^A(v_i \leftarrow v_i, v_{-i}) - q_A f_{-i}(v_{-i})u_i^A(v_i' \leftarrow v_i, v_{-i})\right) \geqslant 0$$

Now we can "Lagrangify" ($\mathcal{L}(x,p,\lambda,\mu)$) the first and second constraints of our LP:

$$\sum_{A\subseteq B} q_A \sum_{v\in\mathcal{V}_A} f_A(v)\sum_{i\in A} p_i^A(v) + \sum_{A\subseteq B}\sum_{i\in A}\sum_{v_i\in\mathcal{V}_i}\sum_{\substack{v_{-i}\in\\ \mathcal{V}_{A/\{i\}}}} \mu_A^i(v_i, v_{-i})u_i^A(v_i \leftarrow v_i, v_{-i}) +$$

$$\sum_{i\in B}\sum_{v_i\in\mathcal{V}_i}\sum_{v_i'\in\mathcal{V}_i}\lambda^i(v_i, v_i')\sum_{\substack{A\subseteq B:\\ i\in A}}\sum_{\substack{v_{-i}\in\\ \mathcal{V}_{A/\{i\}}}}\mathbb{P}\left[A, v_{-i}\right]\left(u_i^A(v_i \leftarrow v_i, v_{-i}) - u_i^A(v_i' \leftarrow v_i, v_{-i})\right) =$$

$$\sum_{A\subseteq B}\sum_{i\in A}\sum_{v_i\in\mathcal{V}_i}\sum_{\substack{v_{-i}\in\\ \mathcal{V}_{A/\{i\}}}} f_i(v_i)f_{A/\{i\}}(v_{-i})\cdot q_A \cdot p_i^A(v_i, v_{-i}) +$$

$$\sum_{A\subseteq B}\sum_{i\in A}\sum_{v_i\in\mathcal{V}_i}\sum_{\substack{v_{-i}\in\\ \mathcal{V}_{A/\{i\}}}} \mu_A^i(v_i, v_{-i})\left(\sum_{j\in[m]} v_{i,j}\cdot x_{i,j}^A(v_i, v_{-i}) - p_i^A(v_i, v_{-i})\right) +$$

$$\sum_{A\subseteq B}\sum_{i\in A}\sum_{v_i\in\mathcal{V}_i}\sum_{\substack{v_{-i}\in\\ \mathcal{V}_{A/\{i\}}}}\sum_{v_i'\in\mathcal{V}_i}\lambda^i(v_i, v_i')\cdot\mathbb{P}\left[A, v_{-i}\right]\left(\sum_{j\in[m]} v_{i,j}\cdot x_{i,j}^A(v_i, v_{-i}) - p_i^A(v_i, v_{-i})\right) -$$

$$\sum_{A\subseteq B}\sum_{i\in A}\sum_{v_i\in\mathcal{V}_i}\sum_{\substack{v_{-i}\in\\ \mathcal{V}_{A/\{i\}}}}\sum_{v_i'\in\mathcal{V}_i}\lambda^i(v_i, v_i')\cdot\mathbb{P}\left[A, v_{-i}\right]\left(\sum_{j\in[m]} v_{i,j}\cdot x_{i,j}^A(v_i', v_{-i}) - p_i^A(v_i', v_{-i})\right) =$$

$$\sum_{\substack{A\subseteq B, i\in A,\\ v_i\in\mathcal{V}_i, v_{-i}\in\mathcal{V}_{A/\{i\}}}}\mathbb{P}\left[A, v_{-i}\right]\cdot p_i^A(v_i, v_{-i})\left(f_i(v_i) + \sum_{v_i'\in\mathcal{V}_i}\lambda^i(v_i', v_i) - \sum_{v_i'\in\mathcal{V}_i}\lambda^i(v_i, v_i') - \frac{\mu_A^i(v_i, v_{-i})}{q_A\cdot f_{A/\{i\}}(v_{-i})}\right) +$$

$$\sum_{\substack{A\subseteq B, i\in A,\\ v_i\in\mathcal{V}_i, v_{-i}\in\mathcal{V}_{A/\{i\}}\\ j\in[m]}} \mathbb{P}\left[A, v_{-i}\right]\cdot x^A_{i,j}(v_i, v_{-i}) \left(\sum_{v'_i\in\mathcal{V}_i} v_{i,j}\cdot\lambda^i(v_i, v'_i) - \sum_{v'_i\in\mathcal{V}_i} v'_{i,j}\cdot\lambda^i(v'_i, v_i) + v_{i,j}\cdot \frac{\mu^i_A(v_i, v_{-i})}{q_A\cdot f_{A/\{i\}}(v_{-i})}\right)$$

For ease of notation we replaced $q_A \cdot f_{A/\{i\}}(v_{-i})$ with $\mathbb{P}\left[A, v_{-i}\right]$.

B Omitted details from Section 5

Proof (Proposition 1). We will prove this claim using contradiction. Let's assume that there exists no constricting set of the desired form. Since our interim allocation rule π is infeasible, there must be a constricting set $S = \{S_i\}_{i\in B}$ such that, for some $k \in B$, S_k is not of the desired form and cannot be replaced with a subset of that form. Now let's fix $S_{-k} = \{S_i\}_{i\in B/\{k\}}$. It must be the case that $S' = \{S_{-k}, S_{x_k}\}$ is not a constricting set for any x_k. Now recall the expression of Theorem 5:

$$\sum_{i=1}^{N}\sum_{v_i\in S_i} q^i f_i(v_i)\pi_i(v_i) \leqslant 1 - \sum_{A\subseteq B} q_A \prod_{i\in A}\left(1 - \sum_{v_i\in S_i} f_i(v_i)\right)$$

The LHS can be broken down into two parts:

$$\sum_{i=1}^{N}\sum_{v_i\in S_i} q^i f_i(v_i)\pi_i(v_i) = \sum_{i\in B/\{k\}}\sum_{v_i\in S_i} q^i f_i(v_i)\pi_i(v_i) + \sum_{v_k\in S_k} q^k f_i(v_k)\pi_k(v_k)$$

The RHS can also be broken into several parts:

$$\sum_{A\subseteq B} q_A \prod_{i\in A}\left(1 - \sum_{v_i\in S_i} f_i(v_i)\right) = \sum_{\substack{A\subseteq B\\ k\in A}} q_A \prod_{i\in A}\left(1 - \sum_{v_i\in S_i} f_i(v_i)\right) + \sum_{\substack{A\subseteq B\\ k\notin A}} q_A \prod_{i\in A}\left(1 - \sum_{v_i\in S_i} f_i(v_i)\right)$$

Furthermore, observe that:

$$\begin{aligned}\sum_{\substack{A\subseteq B\\ k\in A}} q_A \prod_{i\in A}\left(1 - \sum_{v_i\in S_i} f_i(v_i)\right) &= \left(1 - \sum_{v_k\in S_k} f_k(v_k)\right)\cdot \sum_{\substack{A\subseteq B\\ k\in A}} q_A \prod_{\substack{i\in A\\ i\neq k}}\left(1 - \sum_{v_i\in S_i} f_i(v_i)\right)\\ &= \sum_{\substack{A\subseteq B\\ k\in A}} q_A \prod_{\substack{i\in A\\ i\neq k}}\left(1 - \sum_{v_i\in S_i} f_i(v_i)\right) - \sum_{v_k\in S_k} f_k(v_k) \sum_{\substack{A\subseteq B\\ k\in A}} q_A \prod_{\substack{i\in A\\ i\neq k}}\left(1 - \sum_{v_i\in S_i} f_i(v_i)\right)\end{aligned}$$

Finally, observe that:

$$\sum_{\substack{A\subseteq B\\ k\in A}} q_A \prod_{\substack{i\in A\\ i\neq k}}\left(1 - \sum_{v_i\in S_i} f_i(v_i)\right) + \sum_{\substack{A\subseteq B\\ k\notin A}} q_A \prod_{i\in A}\left(1 - \sum_{v_i\in S_i} f_i(v_i)\right) = \sum_{A\subseteq B} q_A \prod_{\substack{i\in A\\ i\neq k}}\left(1 - \sum_{v_i\in S_i} f_i(v_i)\right)$$

Now by re-arranging all the terms of the original expression, we get that:

$$\sum_{v_k \in S_k} \left[q^k \cdot \pi_k(v_k) - \sum_{\substack{A \subseteq B \\ k \in A}} q_A \prod_{\substack{i \in A \\ i \neq k}} \left(1 - \sum_{v_i \in S_i} f_i(v_i) \right) \right] f_k(v_k) \leqslant$$

$$1 - \sum_{i \in B/\{k\}} \sum_{v_i \in S_i} q^i f_i(v_i) \pi_i(v_i) - \sum_{A \subseteq B} q_A \prod_{\substack{i \in A \\ i \neq k}} \left(1 - \sum_{v_i \in S_i} f_i(v_i) \right)$$

Notice that the RHS and $\sum_{\substack{A \subseteq B \\ k \in A}} q_A \prod_{\substack{i \in A \\ i \neq k}} \left(1 - \sum_{v_i \in S_i} f_i(v_i)\right)$ do not rely on our choice of S_k. The choice of S_k that maximizes the LHS is:

$$S'_k = \left\{ v_k \in \mathcal{V}_k : \pi_k(v_k) \geqslant \frac{1}{q^k} \sum_{\substack{A \subseteq B \\ k \in A}} q_A \prod_{\substack{i \in A \\ i \neq k}} \left(1 - \sum_{v_i \in S_i} f_i(v_i) \right) \right\}$$

Thus since S is a constricting set then $\{S'_k, S_{-k}\}$ must be a constricting set which contradicts our initial assumption.

□