



Information elicitation mechanisms for Bayesian auctions

Jing Chen¹ · Bo Li² · Yingkai Li³

Accepted: 1 July 2025
© The Author(s) 2025

Abstract

In this paper we design information elicitation mechanisms for Bayesian auctions. While in Bayesian mechanism design the distributions of the players' private types are often assumed to be common knowledge, information elicitation considers the situation where the players know the distributions better than the decision maker. To weaken the information assumption in Bayesian auctions, we consider an information structure where the knowledge about the distributions is *arbitrarily scattered* among the players. In such an unstructured information setting, we design mechanisms for unit-demand auctions and additive auctions that *aggregate* the players' knowledge, generating revenue that are constant approximations to the optimal Bayesian mechanisms with a common prior. Our mechanisms are 2-step dominant-strategy truthful and the approximation ratios improve gracefully with the amount of knowledge the players collectively have.

Keywords Bayesian auctions · Information elicitation · Distributed knowledge · Removing common prior

1 Introduction

Bayesian auction design has been extremely flourishing in the areas of artificial intelligence, computational economics and operations research, since the seminal work of [1]. One of the main focuses is to generate revenue, by selling m heterogeneous items to n players. Each

✉ Bo Li
comp-bo.li@polyu.edu.hk

Jing Chen
jchencs@mail.tsinghua.edu.cn

Yingkai Li
yk.li@nus.edu.sg

¹ Department of Computer Science and Technology, Tsinghua University, Beijing, China

² Department of Computing, The Hong Kong Polytechnic University, Hong Kong, China

³ Department of Economics, National University of Singapore, Singapore, Singapore

player has a private valuation function describing how much she values each subset of the items, and the valuations are drawn from prior distributions. An important assumption in Bayesian mechanism design is that the distributions are commonly known by the seller and the players—the *common prior* assumption. However, as pointed out by another seminal work [2], such common knowledge is “rarely present in experiments and never in practice”, and “only by repeated weakening of common knowledge assumptions will the theory approximate reality.”

In this paper, we weaken the information assumption about the seller and the players by adopting an information elicitation approach [3]. We consider a framework for auctions where the knowledge about the players’ value distributions is *arbitrarily scattered* among the players and the seller. The seller must *aggregate* pieces of information from all players to gain a good understanding about the distributions, so as to decide how to sell the items.

As in information elicitation, the players get rewards for reporting their knowledge. However, different from classic information elicitation where a player’s utility is exactly her reward, in our model a player’s utility comes not only from her knowledge, but also from participating in the auction (i.e., from buying items). Moreover, information elicitation usually assumes the prior distribution is correlated: each player observes a private signal and reports their private information. This means every player has information about every other player. In our model, following the convention in multi-item auctions, the players’ value distributions for individual items are assumed to be independent. A player may be totally ignorant about some players and only partially know some other players’ distributions.

We focus on unit-demand auctions and additive auctions—two valuation types widely studied in the literature [4, 5]. In such auctions, a player’s valuation function is specified by m values, one for each item. For each player i and item j , the value v_{ij} is independently drawn from a distribution \mathcal{D}_{ij} . When all players are unit-demand (respectively, additive), we call such an auction a *unit-demand auction* (respectively, an *additive auction*) for short. Each player privately knows her own values and some (or none) of the distributions of some other players for some items, like long-time competitors in the market. There is no constraint about who knows which distributions. The seller may also know some of the distributions, but she does not know which player knows what. We introduce directed *knowledge graphs* to succinctly describe the players’ knowledge. Each player knows the distributions of her neighbors, different items’ knowledge graphs may be totally different, and the structures of the graphs are *not known* by anybody.¹

Below we briefly state our main results.

1.1 Main results

Our goal is to design *2-step dominant strategy truthful* (2-DST) information elicitation mechanisms whose expected revenue approximates that of the optimal Bayesian incentive compatible (BIC) mechanism, denoted by OPT .² In order for the seller to aggregate the players’ knowledge about the distributions, it is natural for the mechanism to ask each player to report her knowledge to the seller, together with her own values. A 2-DST mecha-

¹Interestingly, the intuition behind such an information structure has long been considered by philosophers. In [6], the author discussed a world where “everything in the world might be known by somebody, yet not everything by the same knower.”

²A Bayesian mechanism is BIC if it is a Bayesian Nash equilibrium for all players to report their true values.

nism [7] is such that, (1) no matter what knowledge the players may report about each other, it is dominant for each player to report her true values; and (2) given that all players report their true values, it is dominant for each player to report her true knowledge about others.

First, by exploring the knowledge graph's combinatorial structure, we have the following result for single-good auctions.

Main result 1. *When the knowledge graph is 2-connected,³ there is a 2-DST information elicitation mechanism for single-good auctions with revenue $\geq (1 - \frac{1}{n})OPT$.*

When everything is known by somebody For multi-item auctions, we first consider a setting when every distribution is known by at least $k \geq 1$ players. We show the approximation ratios of the revenue generated by information elicitation mechanisms improve gracefully together with the amount of knowledge. More precisely, for any integer $k \geq 1$, let $\tau_k = \frac{k}{(k+1)\frac{k+1}{k}}$. Note $\tau_1 = \frac{1}{4}$ and $\tau_k \rightarrow 1$ when k gets larger. We have the following result, formalized in Section 4.

Main result 2. $\forall k \in [n-1]$, *when each distribution is known by at least k players, there is a 2-DST information elicitation mechanism for unit-demand auctions with revenue $\geq \frac{\tau_k}{24} \cdot OPT$, and such a mechanism for additive auctions with revenue $\geq \max\{\frac{1}{11}, \frac{\tau_k}{6+2\tau_k}\}OPT$.*

Our mechanisms build on top of existing Bayesian mechanisms, such as the post-price mechanisms in [8] and the bundling mechanism in [9]. But it is worth pointing out that, although these mechanisms are used as a black-box, our mechanisms are not reductions from arbitrary Bayesian mechanisms.

Under arbitrary knowledge graphs *When the knowledge graphs are such that some distributions are not known by anybody, it is easy to see that no information elicitation mechanism can be a bounded approximation to OPT . Thus it is natural to consider the following benchmark: the optimal BIC mechanism applied to players and items for whom the distributions are indeed known by somebody, denoted by OPT_K . This is a natural benchmark when considering players with limited knowledge and, if every distribution is known by somebody, then it is exactly OPT . We have the following result, formalized in Section 5.*

Main result 3. *For any knowledge graph, there is a 2-DST information elicitation mechanism for unit-demand auctions with revenue $\geq \frac{OPT_K}{96}$, and such a mechanism for additive auctions with revenue $\geq \frac{OPT_K}{70}$.*

To prove this result, for unit-demand auctions, we actually show a general result: any Bayesian mechanism for unit-demand auctions that is a good approximation in the COPIES

³A directed graph is 2-connected if for any node i , the graph with i and all adjacent edges removed is still strongly connected

setting (formally defined in Section 5.2) can be converted to information elicitation mechanisms. This applies to a large class of Bayesian mechanisms, including the ones in [4, 8, 10].

For additive auctions, we have developed a novel approach by using the *adjusted revenue* [11]. Although this concept is very useful in Bayesian auctions, it was unexpected that we found an interesting and highly non-trivial way of using it to analyze information elicitation mechanisms.

The use of scoring rules *Since our mechanisms elicit the players' knowledge about each other's value distributions, we can use scoring rules (see, e.g., [12]) to reward the players for their reported knowledge, as typical in information elicitation. However, the use of scoring rules does not solve the main problems in our auctions. Indeed, because a player's utility comes both from the reward and from participating in the auction, the difficulties in designing information elicitation mechanisms are to guarantee that, even without rewarding the players for their knowledge, (1) it is dominant for each player to report her true values, (2) reporting her true knowledge never hurts him, and (3) the resulting revenue approximates the desired benchmark.*

Accordingly, in Sections 3 to 5, we focus on designing information elicitation mechanisms without rewarding the players. Scoring rules are used later to break the utility-ties and make it strictly better for a player to report her true knowledge. In Section 6, we show how to add scoring rules to our mechanisms under the assumption of no-bluff and remove the assumption when everything is known.

1.2 Related work

Information elicitation *Following [3], information elicitation has become an important research area of artificial intelligence in the past decade [13–15], where the mechanism asks each player to report her private signal about the prior distribution. The decision maker wants the mechanism to be BIC, and a player is rewarded based on her reported information and the other players' reported signals. Proper scoring rules are widely used [12, 16–19]. The set of proper scoring rules was initially characterized in [16] and [20], and the proper scoring rules for characterizing various statistics of the subjective distributions are characterized in [21]. In our paper, we apply their techniques for eliciting the agents' knowledge and then incorporate information elicitation into auction design by carefully designing the mechanisms to maximize the expected revenue without violating the properness constraints in scoring rules. Similar studies on the combination of information elicitation and mechanism design include [22–26]. The main difference is that those works focus on the direct optimization of scoring rules for incentivizing costly information acquisition without the additional concern of selling items for revenue maximization. Most studies on information elicitation require a common prior. Mechanisms without this assumption are considered by [27], and our work is information elicitation in auctions without a common prior. Moreover, information elicitation does not consider the players to have any cost for revealing their knowledge. We refer to a recent survey [28] for some recent trends in the field of information elicitation. It would be interesting to include such costs in the general model as well as in ours, to see how the mechanisms will change accordingly. In Section 6, we will formally define scoring rules and use them to reward the players for their knowledge.*

Bayesian auction design In his seminal work [1], Myerson introduced the first optimal Bayesian mechanism for single-good auctions, which also applies to many single-parameter settings [29]. Since then, there has been a huge literature on designing (approximately) optimal Bayesian mechanisms that are either BIC or dominant-strategy truthful (DST); see [30] for an introduction to this literature. Regarding mechanisms for multi-parameter settings, in [31], the authors characterize optimal BIC mechanisms for combinatorial auctions. For unit-demand auctions, [4, 8–10] construct DST Bayesian mechanisms that are constant approximations, and recently [32] designed benchmark-tight simple mechanism for a single unit-demand buyer. For additive auctions, [5, 9, 11, 33] provide logarithmic or constant approximations under different conditions. Finally, for general XoS and subadditive auctions, logarithmic and constant approximations are provided in [34–39].

Removing the common prior assumption Following [2], a lot of effort has been made to remove the common prior assumption. In DST Bayesian mechanisms it suffices to assume that the seller knows the prior distribution. In prior-free mechanisms [40–44] the distribution is unknown and the seller learns it from the values of randomly selected players. In [45–47] and [48, 49] the seller observes independent or targeted samples from the distribution before the auction begins. In [50, 51] the players have arbitrary possibilistic belief hierarchies about each other. There is another line of research that assumes certain statistics of distributions are known, [52–55]. In robust mechanism design [56] the players have arbitrary probabilistic belief hierarchies. In crowdsourced Bayesian auctions [7] each player privately knows all the distributions (or their refinements), which is a special case of our model. Indeed, all knowledge graphs will be complete graphs under their setting (that is, everybody knows everything), while we allow arbitrary knowledge graphs. In Section 7 and Appendix B, we further discuss how to elicit the players' knowledge refinements, and how to handle correlated distributions in a setting that is a special case of our model but is still more general than that of [7].

2 Preliminaries

In this work, we focus on multi-item auctions with n players (denoted by N) and m items (denoted by M). A player i 's value for an item j , v_{ij} , is independently drawn from a distribution \mathcal{D}_{ij} . Let $v_i = (v_{ij})_{j \in M}$, $\mathcal{D}_i = \times_{j \in M} \mathcal{D}_{ij}$ and $\mathcal{D} = \times_{i \in N} \mathcal{D}_i$. Player i 's value for a subset S of items is $\max_{j \in S} v_{ij}$ in *unit-demand* auctions, and is $\sum_{j \in S} v_{ij}$ in *additive* auctions. The players' utilities, denoted by u_i , are quasi-linear, and the players are risk-neutral.

Knowledge graphs It is illustrative to model the players' knowledge graphically.⁴ We consider a vector of knowledge graphs, $G = (G_j)_{j \in M}$, one for each item. Each G_j is a directed graph with n nodes, one for each player. For any $i \neq i'$, an edge (i, i') is in G_j if and only if player i knows $\mathcal{D}_{i'j}$. There is no constraint about the knowledge graphs: the same player's

⁴We could have defined the players' knowledge using the standard notion in epistemic game theory [57–59]: roughly speaking, the state space consists of all possible distributions of the valuation profile, and player i knows $\mathcal{D}_{i'j}$ if she is in an information set where all distributions have the (i', j) -th component equal to $\mathcal{D}_{i'j}$. However, the knowledge graph is a more succinct representation and is enough for the purpose of this work.

distributions for different items may be known by different players, different players' distributions for the same item may also be known by different players, and some distributions may not be known by anybody. Each player knows her own out-going edges, and neither the players nor the seller knows the whole graph.

We measure the amount of knowledge in the system by the number of players knowing each distribution. For any $k \in \{0, 1, \dots, n-1\}$, a knowledge graph is k -informed if each node has in-degree at least k : a player's distribution is known by at least k other players. The vector G is k -informed if all knowledge graphs are so. Note that every knowledge graph is 0-informed, and "everything is known by somebody" when $k \geq 1$. A common prior would imply all knowledge graphs are complete directed graphs, or $(n-1)$ -informed, which is the strongest condition in our model. The seller's knowledge can be naturally incorporated into the knowledge graphs by considering her as a special "player 0". All our mechanisms can easily utilize the seller's knowledge, and we will not further discuss this issue.

Information elicitation mechanisms Let $\hat{\mathcal{I}} = (N, M, \mathcal{D})$ be a Bayesian auction instance and $\mathcal{I} = (N, M, \mathcal{D}, G)$ a corresponding information elicitation instance, where G is a knowledge graph vector. Different from Bayesian mechanisms, which has \mathcal{D} as input, an information elicitation mechanism has neither \mathcal{D} nor G as input. Instead, it asks each player i to report a valuation $b_i = (b_{ij})_{j \in M}$ and a knowledge $K_i = \times_{i' \neq i, j \in M} \mathcal{D}_{i'j}^i$ —a distribution for the valuation subprofile v_{-i} . K_i may contain " \perp " at some places, indicating i does not know the corresponding distributions. K_i is i 's true knowledge if $\mathcal{D}_{i'j}^i = \mathcal{D}_{i'j}$ whenever $(i, i') \in G_j$, and $\mathcal{D}_{i'j}^i = \perp$ otherwise. An information elicitation mechanism maps a strategy profile $(b_i, K_i)_{i \in N}$ to an allocation and a price profile, and may be randomized. To distinguish whether a mechanism \mathcal{M} is a Bayesian or an information elicitation mechanism, we may explicitly write $\mathcal{M}(\hat{\mathcal{I}})$ or $\mathcal{M}(\mathcal{I})$. The (expected) revenue of \mathcal{M} is denoted by $\text{Rev}(\mathcal{M})$, and sometimes by $\mathbb{E}_{\mathcal{D}} \text{Rev}(\mathcal{M})$ to emphasize the distribution.

An information elicitation mechanism is 2-step dominant strategy truthful (2-DST) if

- (1) For any player i , true valuation v_i , valuation b_i , knowledge K_i , and strategy subprofile $s_{-i} = (b_j, K_j)_{j \neq i}$ of the other players, $u_i((v_i, K_i), s_{-i}) \geq u_i((b_i, K_i), s_{-i})$.
- (2) For any player i , true valuation v_i , true knowledge K_i , knowledge K'_i , and knowledge subprofile $K'_{-i}(v_{-i}) = (K'_j(v_j))_{j \neq i}$ of the other players, where each $K'_j(v_j)$ is a function of player j 's true valuation v_j , $\mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} u_i((v_i, K_i), (v_{-i}, K'_{-i}(v_{-i}))) \geq \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} u_i((v_i, K'_i), (v_{-i}, K'_{-i}(v_{-i})))$.⁵

It is not hard to see that 2-DST is a refined solution concept of the Bayesian Nash equilibrium. When the inequalities are all strict in the two conditions of 2-DST, it is called two-step strictly dominant strategy truthful (2-SDST). As noted in [7], any 2-SDST mechanism has a unique equilibrium where players reveal their true valuations and their true beliefs. Moreover, truthfulness is the only strategy that survives two steps of iterated elimination of strictly dominated strategies.

⁵In Section 6, we introduce scoring rules to our mechanisms so that the inequality is strict whenever $K'_i \neq K_i$.

3 Single-good auctions with 2-connected knowledge graphs

As a warm up, in this section, we construct an information elicitation mechanism that is *nearly optimal* for the special case of single-good auctions under a natural structure of the knowledge graph. More precisely, recall that a directed graph is *strongly connected* if there is a directed path from any node i to any other node i' . Intuitively, in a knowledge graph this means that for any two players Alice and Bob, Alice knows a guy who knows a guy ... who knows Bob. Also recall that a directed graph is *2-connected* if it remains strongly connected after removing any single node and the adjacent edges. In a knowledge graph, this means there does not exist a crucial player as an “information hub”, without whom the players will split into two parts, with one part having no information about the other. It is easy to see that a knowledge graph being strong connected (or 2-connected respectively) implies it being 1-informed (or 2-informed respectively), but not vice versa. In fact, a graph of n nodes can be $(\lfloor \frac{n}{2} \rfloor - 1)$ -informed without being connected.


When the knowledge graph is 2-connected, we construct the *information elicitation Myerson* mechanism \mathcal{M}_{IEM} in Mechanism Algorithm 1. Recall that Myerson’s mechanism maps each player i ’s reported value b_i to the (*ironed*) virtual value, $\phi_i(b_i; \mathcal{D}_i)$. It runs the second-price mechanism with reserve price 0 on virtual values and maps the resulting “virtual price” back to the winner’s value space, as her price.

Algorithm 1 \mathcal{M}_{IEM}

- 1: Each player i reports a value b_i and a knowledge $K_i = (\mathcal{D}_j)_{j \in N \setminus \{i\}}$.
 - 2: Randomly choose a player a , let $S = \{j \mid \mathcal{D}_j^a \neq \perp\}$, $N' = N \setminus (\{a\} \cup S)$.
 - 3: If $S = \emptyset$, the item is unsold, the mechanism sets price $p_i = 0$ for each $i \in N$ and stop here.
 - 4: Set $\mathcal{D}'_j = \mathcal{D}_j^a$ for all $j \in S$.
 - 5: Set $i^* = \operatorname{argmax}_{j \in S} \phi_j(b_j; \mathcal{D}'_j)$, with ties broken lexicographically.
 - 6: **while** $N' \neq \emptyset$ **do**
 - 7: Set $S' = \{j \mid j \in N', \exists i' \in S \setminus \{i^*\} \text{ s.t. } \mathcal{D}_j^{i'} \neq \perp\}$.
 - 8: If $S' = \emptyset$ then go to Step 13.
 - 9: For each $j \in S'$, set $\mathcal{D}'_j = \mathcal{D}_j^{i'}$, where i' is the player in $S \setminus \{i^*\}$ with smallest identity such that $\mathcal{D}_j^{i'} \neq \perp$.
 - 10: Set $S = \{i^*\} \cup S'$ and $N' = N' \setminus S'$.
 - 11: Set $i^* = \operatorname{argmax}_{j \in S} \phi_j(b_j; \mathcal{D}'_j)$, with ties broken lexicographically.
 - 12: **end while**
 - 13: Set $\phi_{second} = \max_{j \in N \setminus (\{a, i^*\} \cup N')} \phi_j(b_j; \mathcal{D}'_j)$ and the price $p_i = 0$ for each player i .
 - 14: If $\phi_{i^*}(b_{i^*}; \mathcal{D}'_{i^*}) < 0$ then the item is unsold; otherwise, the item is sold to player i^* and $p_{i^*} = \phi_{i^*}^{-1}(\max\{\phi_{second}, 0\}; \mathcal{D}'_{i^*})$.
-

To help understand our mechanism, we illustrate in Fig. 1 the sets of players involved in the first round of Mechanism \mathcal{M}_{IEM} . Informally, \mathcal{M}_{IEM} starts from a randomly selected player a and computes the virtual values of the players whose distributions are reported by a . In the following rounds, \mathcal{M}_{IEM} maintains a set of explored players S whose virtual values have been computed and a player i^* who has the highest virtual value among them. In each round, \mathcal{M}_{IEM} computes the virtual values of more players using the information reported by S excluding i^* , and updates S and i^* accordingly. As we will prove, since the knowledge graph is 2-connected, eventually all the players’ virtual values can be reported and used to compute the winner, excluding a . We have the following theorem.

Fig. 1 The sets of players involved in the first round of Mechanism \mathcal{M}_{IEM} . The edges in the figure correspond to distributions reported by the players. In each round, the mechanism keeps in S the player with the highest virtual value so far, drops all the other players from S , and adds the players whose distributions are reported for the first time by the dropped ones



Theorem 1 For any single-good auction instances $\hat{\mathcal{I}} = (N, M, \mathcal{D})$ and $\mathcal{I} = (N, M, \mathcal{D}, G)$ where G is 2-connected, \mathcal{M}_{IEM} is 2-DST and $Rev(\mathcal{M}_{IEM}(\mathcal{I})) \geq (1 - \frac{1}{n}) OPT(\hat{\mathcal{I}})$.

Before proving the revenue bound, we first prove \mathcal{M}_{IEM} is 2-DST.

Lemma 1 \mathcal{M}_{IEM} is 2-DST.

Proof By the definition of 2-DST, we prove the lemma in two steps.

Claim 1 For any player i , true value v_i , value b_i , knowledge K_i , and strategy subprofile $s_{-i} = (b_j, K_j)_{j \neq i}$ of the other players, $\mathbb{E}_{\mathcal{M}_{IEM}} u_i((v_i, K_i), s_{-i}) \geq \mathbb{E}_{\mathcal{M}_{IEM}} u_i((b_i, K_i), s_{-i})$, where the expectation is taken over the mechanism's random coins.

Proof First, conditioning on $a = i$, player i does not get the item and her reported value is not used by the mechanism. Thus $u_i((v_i, K_i), s_{-i}) = u_i((b_i, K_i), s_{-i}) = 0$ in this case.

Second, we compare the two utilities conditional on $a \neq i$. Notice that when $a \neq i$, whether or not player i 's distribution is reported—that is, whether or not \mathcal{D}'_i is defined—only depends on s_{-i} . Thus \mathcal{D}'_i is defined under (v_i, K_i) if and only if it is defined under (b_i, K_i) .

If \mathcal{D}'_i is not defined, then $i \in N'$ at the end of the mechanism, she does not get the item, and her reported value is not used. Therefore $u_i((v_i, K_i), s_{-i}) = u_i((b_i, K_i), s_{-i}) = 0$ again.

If \mathcal{D}'_i is defined, then it is defined in the same round of the mechanism under both (v_i, K_i) and (b_i, K_i) , which we refer to as round r . Also, \mathcal{D}'_i is the same in both cases and the mechanism's execution is the same till this round. Notice that

- (1) $\phi_i(\cdot; \mathcal{D}'_i)$ is monotone in its input;
- (2) i gets the item if and only if $i = i^*$ in all rounds ℓ with $\ell \geq r$ and her virtual value is at least 0; and
- (3) when i gets the item, K_i is never used by the mechanism and thus does not affect the execution of any round ℓ with $\ell \geq r$.

Accordingly, the mechanism is *monotone* in player i 's reported value: if i gets the item by reporting some value, then she still gets it by reporting a higher value. Moreover, when i gets the item, her price in Step 14 is the *threshold* payment. Following standard characterizations

of single-parameter DST mechanisms, it is the best for player i to report her true value v_i . That is, $u_i((v_i, K_i), s_{-i}) \geq u_i((b_i, K_i), s_{-i})$ when \mathcal{D}'_i is defined.

Combining the above cases together, we have $\mathbb{E}_{\mathcal{M}_{IEM}} u_i((v_i, K_i), s_{-i}) \geq \mathbb{E}_{\mathcal{M}_{IEM}} u_i((b_i, K_i), s_{-i})$ and Claim 1 holds. \square

Claim2 For any player i , true value v_i , true knowledge K_i , knowledge K'_i , and knowledge subprofile $K_{-i}(v_{-i}) = (K'_j(v_j))_{j \neq i}$ of the other players, where each $K'_j(v_j)$ is a function of player j 's true value v_j , $\mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} u_i((v_i, K_i), (v_{-i}, K'_{-i}(v_{-i}))) \geq \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} u_i((v_i, K'_i), (v_{-i}, K'_{-i}(v_{-i})))$.

Proof Similar to Claim 1, conditional on $a = i$, player i does not get the item no matter what knowledge she reports. Thus

$$\begin{aligned} & \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} [u_i((v_i, K_i), (v_{-i}, K'_{-i}(v_{-i}))) \mid a = i] \\ &= \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} [u_i((v_i, K'_i), (v_{-i}, K'_{-i}(v_{-i}))) \mid a = i] = 0. \end{aligned}$$

Next, we compare the two utilities conditional on $a \neq i$. Again similar to Claim 1, \mathcal{D}'_i is the same under both strategies of i , and the mechanism's execution is also the same till the round r where \mathcal{D}'_i is defined (or till the end if \mathcal{D}'_i is not defined). There are three cases:

- If \mathcal{D}'_i is not defined, then neither K_i nor K'_i is used by the mechanism, and i has utility 0 under both strategies.
- If \mathcal{D}'_i is defined and $i = i^*$ from round r to the end of the mechanism, then again K_i and K'_i are not used. Thus i has the same utility (maybe non-zero) under both strategies.
- If \mathcal{D}'_i is defined and $i \neq i^*$ starting from some round $r' \geq r$, then i does not get the item under either strategy, thus her utility is 0 under both of them.

In sum, $\mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} u_i((v_i, K_i), (v_{-i}, K'_{-i}(v_{-i}))) = \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} u_i((v_i, K'_i), (v_{-i}, K'_{-i}(v_{-i})))$ and reporting her true knowledge does not hurt player i . \square

Lemma 1 follows directly from Claims 1 and 2. \square

In Section 6, we add scoring rules to mechanism \mathcal{M}_{IEM} to reward the players' knowledge, so that a player's utility will be strictly larger when she reports her true knowledge than when she lies. To prove Theorem 1, it remains to show $\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEM}(\mathcal{I})) \geq (1 - \frac{1}{n}) \text{OPT}(\hat{\mathcal{I}})$ under the players' truthful strategies.

Proof of Theorem 1 The key is to explore the structure of the knowledge graph to make sure that the player with the highest virtual value is found by the mechanism with high probability.

More specifically, arbitrarily fix the player a chosen by the mechanism. Notice that throughout the mechanism, N' is the set of players $i \in N \setminus \{a\}$ such that \mathcal{D}'_i is not defined. We show that $N' = \emptyset$ at the end of the mechanism. Indeed, since G is 2-connected, the out-degree of a in G is at least 2: otherwise, either a cannot reach any other node in G , or this becomes the case after removing the unique node j with $(a, j) \in G$, contradicting 2-connectedness. Since a reports her true knowledge K_a , we have $|S| \geq 2$ in Step 3 and the mechanism does not stop there. Moreover, at the beginning of each round, we have $|S| \geq 2$

and thus $S \setminus \{i^*\} \neq \emptyset$: otherwise $S' = \emptyset$ in the previous round, and the mechanism would not have reached this round.

Assume, for the sake of contradiction that the mechanism finally reaches a round r where $N' \neq \emptyset$ at the beginning but $S' = \emptyset$ in Step 7. Since all players report their true knowledge, by the definition of S' we have that, in graph G , all neighbors of $S \setminus \{i^*\}$ are in $N \setminus N'$. Furthermore, for any player $i \in (N \setminus N') \setminus S$, all neighbors of i are also in $N \setminus N'$: indeed, i has been moved from N' to S and then dropped from S (except player a , whose neighbors are in $N \setminus N'$ by definition); and when i is dropped from S , all her neighbors in N' are moved to S . Accordingly, all the edges going from $N \setminus N'$ to N' are from player i^* , and G becomes disconnected after removing i^* , again contradicting 2-connectedness. Thus $S' \neq \emptyset$ in all rounds and $N' = \emptyset$ in the end, as we wanted to show.

Because all players report their true values and true knowledge, we have $\mathcal{D}'_{-a} = \mathcal{D}_{-a}$ and $\phi_i(v_i; \mathcal{D}'_i) = \phi_i(v_i; \mathcal{D}_i)$ for all $i \neq a$. Letting $\hat{\mathcal{T}}_a = (N \setminus \{a\}, M, \mathcal{D}_{-a})$, we claim

$$\mathbb{E}_{v \sim \mathcal{D}}[\text{Rev}(\mathcal{M}_{IEM}(\mathcal{I})) \mid a] = \text{OPT}(\hat{\mathcal{T}}_a). \quad (1)$$

To see why this is true, note that by construction, in each round the mechanism keeps the player with the highest virtual value in S . Thus, the final player i^* has the highest virtual value in $N \setminus \{a\}$, and ϕ_{second} is the second highest virtual value in $N \setminus \{a\}$. Accordingly, the outcome of Step 14 is the same as that of Myerson's mechanism on $\hat{\mathcal{T}}_a$, so is the revenue. Therefore (1) holds.

Finally, it remains to show that, by throwing away a random player a , the mechanism does not lose much revenue. For each player i , letting $P_i(\text{OPT}(\hat{\mathcal{T}}))$ be the expected price paid by i in Myerson's mechanism under $\hat{\mathcal{T}}$, we have $\text{OPT}(\hat{\mathcal{T}}) = \sum_{i \in N} P_i(\text{OPT}(\hat{\mathcal{T}}))$. Next, consider the following Bayesian mechanism \mathcal{M}' on $\hat{\mathcal{T}}_a$: it runs Myerson's mechanism on $\hat{\mathcal{T}}$ and then projects the outcome to players $N \setminus \{a\}$, i.e., the item is unsold and the payment is not collected if a gets the good; otherwise, the outcome is kept the same. It is easy to see that \mathcal{M}' is DST, thus it cannot generate more revenue than $\text{OPT}(\hat{\mathcal{T}}_a)$. As the expected revenue of \mathcal{M}' is $\sum_{i \neq a} P_i(\text{OPT}(\hat{\mathcal{T}}))$, then

$$\text{OPT}(\hat{\mathcal{T}}_a) \geq \mathbb{E}_{\mathcal{D}_{-a}} \text{Rev}(\mathcal{M}'(\hat{\mathcal{T}}_a)) = \sum_{i \neq a} P_i(\text{OPT}(\hat{\mathcal{T}})). \quad (2)$$

Combining (1) and (2), we have

$$\begin{aligned} \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEM}(\mathcal{I})) &= \sum_{a \in N} \frac{1}{n} \mathbb{E}_{v \sim \mathcal{D}}[\text{Rev}(\mathcal{M}_{IEM}(\mathcal{I})) \mid a] = \sum_{a \in N} \frac{1}{n} (\text{OPT}(\hat{\mathcal{T}}_a)) \\ &\geq \sum_{a \in N} \frac{1}{n} \left(\sum_{i \neq a} P_i(\text{OPT}(\hat{\mathcal{T}})) \right) = \left(\sum_{i \in N} \frac{n-1}{n} P_i(\text{OPT}(\hat{\mathcal{T}})) \right) = (1 - \frac{1}{n}) \text{OPT}(\hat{\mathcal{T}}). \end{aligned}$$

Combining Lemma 1, Theorem 1 holds. \square

4 When everything is known by somebody

When the knowledge graph vector G is k -informed with $k \geq 1$, both mechanisms in Section 5 of course apply, but we can do better when k gets larger. In fact, we show that, for both unit-demand auctions and additive auctions, as k gets larger, the approximation ratio approaches the best known approximation to OPT by DST Bayesian mechanisms [8, 9, 11].

4.1 Unit-demand auctions

For unit-demand auctions, sequential post-price mechanisms have been constructed by [8, 10]. For information elicitation, if the seller asks the players to report both their values and knowledge, and directly uses the reported distributions in these mechanisms, then a player may want to withhold her knowledge about the other players. By doing so, a player may prevent the seller from selling the items to the others, so that the items are still available when it is her turn to buy.

A simple idea is to partition the players into two groups: a set of *reporters* who will not receive any item and is only asked to report their knowledge; and a set of *potential buyers* whose knowledge is never used. It is possible that the reported knowledge may not contain a potential buyer's value distributions on all items, thus the technical part is to prove that the seller generates a good revenue even though the players' knowledge is only partially recovered.

Our mechanism $\mathcal{M}_{I\text{EUD}}^k$ is simple and intuitive; see Mechanism Algorithm 2, where \mathcal{M}_{UD} is the Bayesian mechanism of [8] (see Section 6 therein). The probability that each player is assigned to N_1 (reporters) is $q = 1 - (k+1)^{-\frac{1}{k}}$, and the probability to N_2 (potential buyers) is $1 - q$. The probability q is chosen to achieve the maximum probability for each distribution to be reported, and the latter is exactly $\tau_k = \frac{k}{(k+1)^{\frac{k+1}{k}}}$. It's worth pointing

out that, although mechanism \mathcal{M}_{UD} is used as a black-box, Mechanism Algorithm 2 is not a reduction from arbitrary Bayesian mechanisms. We summarize our result in Theorem 2.

Algorithm 2 $\mathcal{M}_{I\text{EUD}}^k$

-
- 1: Each player i reports to the seller a valuation $b_i = (b_{ij})_{j \in M}$ and a knowledge $K_i = (\mathcal{D}_{i'j}^i)_{i' \neq i, j \in M}$.
 - 2: Randomly partition the players into two sets, N_1 and N_2 , where each player is independently put to N_1 with probability $q = 1 - (k+1)^{-\frac{1}{k}}$ and to N_2 with probability $1 - q$.
 - 3: Set $N_3 = \emptyset$.
 - 4: **for** players $i \in N_1$ lexicographically **do**
 - 5: For each player $i' \in N_2$ and item $j \in M$, if $\mathcal{D}_{i'j}^{i'}$ has not been defined yet and $\mathcal{D}_{i'j}^i \neq \perp$, then set $\mathcal{D}_{i'j}^{i'} = \mathcal{D}_{i'j}^i$ and add player i' to N_3 .
 - 6: **end for**
 - 7: For each $i \in N_3$ and $j \in M$ such that \mathcal{D}_{ij}^i is not defined, set $\mathcal{D}_{ij}^i \equiv 0$ (i.e., 0 with probability 1) and $b_{ij} = 0$.
 - 8: Run mechanism \mathcal{M}_{UD} on the unit-demand Bayesian auction $(N_3, M, (\mathcal{D}_{ij}^i)_{i \in N_3, j \in M})$, with the players' values being $(b_{ij})_{i \in N_3, j \in M}$. Let $x' = (x'_{ij})_{i \in N_3, j \in M}$ be the resulting allocation where $x'_{ij} \in \{0, 1\}$, and let $p' = (p'_i)_{i \in N_3}$ be the prices. Without loss of generality, $x'_{ij} = 0$ if $\mathcal{D}_{ij}^i \equiv 0$.
 - 9: For each player $i \notin N_3$, i gets no item and her price is $p_i = 0$.
 - 10: For each player $i \in N_3$, i gets item j if $x'_{ij} = 1$, and her price is $p_i = p'_i$.
-

Theorem 2 $\forall k \in [n - 1]$, any unit-demand auction instances $\hat{\mathcal{I}} = (N, M, \mathcal{D})$ and $\mathcal{I} = (N, M, \mathcal{D}, G)$ where G is k -informed, mechanism \mathcal{M}_{IEUD}^k is 2-DST and $Rev(\mathcal{M}_{IEUD}^k(\mathcal{I})) \geq \frac{\tau_k}{24} \cdot OPT(\hat{\mathcal{I}})$.

Note that τ_k is increasing in k , $\tau_1 = \frac{1}{4}$ and $\tau_k \rightarrow 1$. As k gets larger, the approximation ratio approaches 24, the best known approximation to OPT by DST Bayesian mechanisms [8, 9].

To prove Theorem 2, the 2-DSTness follows from Lemma 2 below. Regarding the revenue bound, in the next section, we will prove a more general result where k may be 0 (see Theorem 4), and omit a formal proof of the bound in Theorem 2. To understand the approximation ratio, below we show that under the players' truthful strategies, the probability for each distribution \mathcal{D}_{ij} to be reported in mechanism \mathcal{M}_{IEUD}^k is at least τ_k . Then following the proof of Theorem 4, we have the revenue bound of the mechanism in Theorem 2.

Indeed, for any player i and item j ,

$$\begin{aligned} & \Pr(\mathcal{D}_{ij} \text{ is reported in the mechanism}) \\ &= \Pr(i \in N_2) \Pr(\exists i' \in N_1, (i, i') \in G_j \mid i \in N_2) \geq (1 - q)(1 - (1 - q)^k), \end{aligned}$$

where the inequality is because \mathcal{D}_{ij} is known by at least k players other than i , and the players are partitioned independently. Taking derivatives of the last term, we have that it is maximized when $q = 1 - (k + 1)^{-\frac{1}{k}}$ as in the mechanism, in which case

$$\Pr(\mathcal{D}_{ij} \text{ is reported in the mechanism}) \geq \frac{1}{(k + 1)^{\frac{1}{k}}} \cdot \frac{k}{k + 1} = \tau_k.$$

It remains to prove the 2-DSTness.

Lemma 2 Mechanism \mathcal{M}_{IEUD}^k is 2-DST.

Proof We start with the first requirement in the solution concept: it is best for a player to report her true values, no matter what knowledge she reports and what strategies the others use.

Claim 3 For any player i , true valuation v_i , valuation b_i , knowledge K_i , and strategy subprofile $s_{-i} = (b_j, K_j)_{j \neq i}$ of the other players, $\mathbb{E}_{\mathcal{M}_{IEUD}^k} u_i((v_i, K_i), s_{-i}) \geq \mathbb{E}_{\mathcal{M}_{IEUD}^k} u_i((b_i, K_i), s_{-i})$, where the expectation is taken over the mechanism's random coins.

Proof If player i is not in N_3 given the mechanism's randomness and the reported knowledge of all players, then her reported valuation is never used to compute his allocation or price, and she gets the same utility for reporting any valuation. Thus $u_i((v_i, K_i), s_{-i}) = u_i((b_i, K_i), s_{-i})$ conditional on $i \notin N_3$.

If player i is in N_3 , then her utility is determined by \mathcal{M}_{UD} . If \mathcal{D}_{ij}' is defined to be 0 with probability 1 in Step 7, then i 's reported value for item j is not given to \mathcal{M}_{UD} as input, and i gets the same utility for reporting any value for j , including v_{ij} . Moreover, because i does not get such an item j , his utility is the same as a dummy player \hat{i} whose valuation is the

same as i 's, except that the true value of \hat{i} for j is 0. Since \mathcal{M}_{UD} is DST, no matter what the distributions are and what values the other players report, it is the best for \hat{i} to report her true valuation. Accordingly, it is the best for i to report her true valuation v_i as well. That is, $u_i((v_i, K_i), s_{-i}) \geq u_i((b_i, K_i), s_{-i})$ conditional on $i \in N_3$.

Combining these two cases, Claim 3 holds. \square

We now prove the second requirement in the solution concept: given that all players report their true valuations, it is best for a player to report her true knowledge.

Claim 4 *For any player i , true valuation v_i , true knowledge K_i , knowledge K'_i , and knowledge subprofile $K'_{-i}(v_{-i}) = (K'_j(v_j))_{j \neq i}$ of the other players, where each $K'_j(v_j)$ is a function of player j 's true valuation v_j , $\mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} u_i((v_i, K_i), (v_{-i}, K'_{-i}(v_{-i}))) \geq \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} u_i((v_i, K'_i), (v_{-i}, K'_{-i}(v_{-i})))$.*

Proof If player i is in N_1 , then she is guaranteed to get no item and pay 0, so her utility is 0 no matter which knowledge she reports. If player i is in N_2 , then her knowledge is never used, and she again gets the same utility no matter which knowledge she reports. Thus $\mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} u_i((v_i, K_i), (v_{-i}, K'_{-i}(v_{-i}))) = \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} u_i((v_i, K'_i), (v_{-i}, K'_{-i}(v_{-i})))$ and Claim 4 holds. \square

Lemma 2 follows directly from Claims 3 and 4. \square

4.2 Additive auctions

Information elicitation mechanisms for additive auctions are harder to construct and analyze than for unit-demand auctions. First, randomly partitioning the players as before may cause a significant revenue loss, as the revenue of additive auctions may come from selling a subset of items as a *bundle* to a player i . Even when i 's value distribution for each item is reported with constant probability, the probability that her distributions for all items in the bundle are reported may be very low, thus the mechanism may rarely sell the bundle to i at the optimal price. Second, the seller can no longer “throw away” player-item pairs whose distributions are not reported and focus on the projected instance. When the players are not partitioned into reporters and potential buyers, doing so may cause a player to lie and withhold her knowledge about others, so that they are thrown away.

To simultaneously achieve truthfulness and a good revenue guarantee, our mechanism is very stingy and never throws away any information. Following [9] (see Section 7 therein), we can consider a mechanism consisting of a “bundling part” and an “individual sale part”. The former is referred to as the *Bundle VCG* mechanism, denoted by *BVCG*; and the latter is the *individual 1-lookahead* mechanism, denoted by \mathcal{M}_{1LA} , which sells each item separately using the 1-lookahead mechanism of [60]. Mechanism \mathcal{M}_{1LA} can also be replaced by the *individual Myerson* mechanism, denoted by *IM*, which sells each item separately using Myerson's mechanism. By choosing the mechanism that generates a higher expected revenue between *IM* and *BVCG*, [9] provides a Bayesian mechanism that is an 8-approximation to *OPT*.

We design corresponding information elicitation mechanisms separately for mechanisms *BVCG* and *IM*. The resulting mechanisms are denoted by \mathcal{M}_{IEBVCG} and \mathcal{M}_{IEIM} , which are formally defined later. Because the seller does not know the prior \mathcal{D} , she cannot compute

the expected revenue of the two information elicitation mechanisms and choose the better one. Instead, we let her choose between the two mechanisms randomly, according to a probability distribution depending on k . It is worth pointing out that when $k \geq 1$, mechanism \mathcal{M}_{IEBVC} is able to recover *all* distributions in $\hat{\mathcal{I}}$, thus its revenue equals the corresponding Bayesian revenue, which is lower-bounded in [9]. For \mathcal{M}_{IEIM} , information elicitation is done by randomized partition depending on the value of k .

However, we can do even better. Indeed, although in Bayesian auctions the mechanism IM is optimal for individual item-sale and outperforms mechanism \mathcal{M}_{1LA} , in information elicitation auctions there is a tradeoff between the two. In order for the players to report their knowledge truthfully for mechanism IM , we need to randomly partition them into reporters and potential buyers, thus each distribution is only recovered with probability τ_k . In contrast, no partition is needed for aggregating the players' knowledge in mechanism \mathcal{M}_{1LA} , and we can recover all distributions simultaneously with probability 1. The resulting information elicitation mechanism, \mathcal{M}_{IE1LA} , is also defined later. As mechanism \mathcal{M}_{1LA} is a 2-approximation to mechanism IM , sometimes it is actually more advantageous to use \mathcal{M}_{IE1LA} rather than \mathcal{M}_{IEIM} , depending on the value of k .

Properly combining the above gadgets together, our mechanism \mathcal{M}'_{IEA} is defined as follows: when $k \leq 7$, it runs \mathcal{M}_{IEBVC} with probability $\frac{2}{11}$ and \mathcal{M}_{IE1LA} with probability $\frac{9}{11}$; when $k > 7$, it runs \mathcal{M}_{IEBVC} with probability $\frac{\tau_k}{3+\tau_k}$ and \mathcal{M}_{IEIM} with probability $\frac{3}{3+\tau_k}$. The choice of the two cases is to achieve the best approximation ratio for each k . We have the following theorem.

Theorem 3 $\forall k \in [n-1]$, any additive auction instances $\hat{\mathcal{I}} = (N, M, \mathcal{D})$ and $\mathcal{I} = (N, M, \mathcal{D}, G)$ where G is k -informed, \mathcal{M}'_{IEA} is 2-DST and $Rev(\mathcal{M}'_{IEA}(\mathcal{I})) \geq \max\{\frac{1}{11}, \frac{\tau_k}{6+2\tau_k}\} OPT(\hat{\mathcal{I}})$.

To prove Theorem 3, we first formally define mechanisms \mathcal{M}_{IEIM} , \mathcal{M}_{IE1LA} and \mathcal{M}_{IEBVC} .

The information elicitation individual myerson mechanism We start by introducing the information elicitation individual Myerson mechanism \mathcal{M}_{IEIM} , which runs the following mechanism $\mathcal{M}_{IEIM,j}$ for each item j separately. Mechanism $\mathcal{M}_{IEIM,j}$ is similar to \mathcal{M}_{IEUD}^k , thus we have omitted many details in the analysis.

Algorithm 3 $\mathcal{M}_{IEIM,j}$

-
- 1: Each player i reports a value b_{ij} and a knowledge $K_{ij} = (\mathcal{D}'_{i'j})_{i' \neq i}$.
 - 2: Randomly partition the players into two sets, N_1 and N_2 , where each player is independently put in N_1 with probability $q = 1 - (k+1)^{-\frac{1}{k}}$ and N_2 with probability $1 - q$.
 - 3: Let N_3 be the set of players in N_2 whose distributions are reported by some players in N_1 , and $\mathcal{D}'_{N_3,j}$ be the vector of reported distributions.
 - 4: Run Myerson's mechanism on the single-good Bayesian instance $\hat{\mathcal{I}}_{N_3,j} = (N_3, \{j\}, \mathcal{D}'_{N_3,j})$ with the values being $(b_{ij})_{i \in N_3}$; and use the resulting allocation and prices to sell to players in N_3 .
-

For each item j , let $v_j = (v_{ij})_{i \in N}$, $\mathcal{D}_j = (\mathcal{D}_{ij})_{i \in N}$, $\hat{\mathcal{I}}_j = (N, \{j\}, \mathcal{D}_j)$ be the corresponding single-good Bayesian instance, and $\mathcal{I}_j = (N, \{j\}, \mathcal{D}_j, G_j)$ be the corresponding single-good information elicitation instance. Lemma 3 below is similar to Lemma 2 and we provide its statement only.

Lemma 3 For any additive auction instances $\hat{\mathcal{I}} = (N, M, \mathcal{D})$ and $\mathcal{I} = (N, M, \mathcal{D}, G)$, mechanism $\mathcal{M}_{IEIM,j}$ is 2-DST for \mathcal{I}_j for each $j \in M$, and mechanism \mathcal{M}_{IEIM} is 2-DST for \mathcal{I} .

Next, we consider the expected revenue of \mathcal{M}_{IEIM} .

Lemma 4 $\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEIM}(\mathcal{I})) \geq \tau_k \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(IM(\hat{\mathcal{I}}))$.

Proof By definition,

$$\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEIM}(\mathcal{I})) = \sum_{j \in M} \mathbb{E}_{v_j \sim \mathcal{D}_j} \text{Rev}(\mathcal{M}_{IEIM,j}(\mathcal{I}_j))$$

and

$$\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(IM(\hat{\mathcal{I}})) = \sum_{j \in M} \text{OPT}(\hat{\mathcal{I}}_j).$$

Accordingly, it suffices to show that for each item j , $\mathbb{E}_{v_j \sim \mathcal{D}_j} \text{Rev}(\mathcal{M}_{IEIM,j}(\mathcal{I}_j)) \geq \tau_k \text{OPT}(\hat{\mathcal{I}}_j)$, and we have

$$\begin{aligned} & \mathbb{E}_{v_j \sim \mathcal{D}_j} \text{Rev}(\mathcal{M}_{IEIM,j}(\mathcal{I}_j)) \\ &= \mathbb{E}_{N_3} \mathbb{E}_{v_{N_3,j} \sim \mathcal{D}_{N_3,j}} \text{OPT}(\hat{\mathcal{I}}_{N_3,j}) \geq \mathbb{E}_{N_3} \mathbb{E}_{v_j \sim \mathcal{D}_j} \text{OPT}(\hat{\mathcal{I}}_j)_{N_3} \\ &= \mathbb{E}_{N_3} \mathbb{E}_{v_j \sim \mathcal{D}_j} \sum_{i \in N_3} P_i(\text{OPT}(\hat{\mathcal{I}}_j)) = \mathbb{E}_{v_j \sim \mathcal{D}_j} \mathbb{E}_{N_3} \sum_{i \in N_3} P_i(\text{OPT}(\hat{\mathcal{I}}_j)) \\ &= \mathbb{E}_{v_j \sim \mathcal{D}_j} \sum_i \Pr(i \in N_3) \cdot P_i(\text{OPT}(\hat{\mathcal{I}}_j)) \geq \tau_k \mathbb{E}_{v_j \sim \mathcal{D}_j} \sum_i P_i(\text{OPT}(\hat{\mathcal{I}}_j)) \\ &= \tau_k \text{OPT}(\hat{\mathcal{I}}_j), \end{aligned}$$

as desired. Thus Lemma 4 holds. \square

The information elicitation individual 1-lookahead mechanism Next, we introduce the information elicitation 1-lookahead mechanism \mathcal{M}_{IE1LA} , which runs the following mechanism $\mathcal{M}_{IE1LA,j}$ for each item j separately. We will show that the revenue of \mathcal{M}_{IE1LA} matches that of mechanism \mathcal{M}_{1LA} for any $k \geq 1$.

Algorithm 4 $\mathcal{M}_{IE1LA,j}$

-
- 1: Each player i reports a value b_{ij} and a knowledge $K_{ij} = (\mathcal{D}_{i'j}^i)_{i' \neq i}$.
 - 2: Set $i^* = \arg\max_i b_{ij}$ and $p_{second} = \max_{i \neq i^*} b_{ij}$.
 - 3: If i^* 's distribution is not reported,
sell item j to her at price p_{second} and halt here.
 - 4: Otherwise, let \mathcal{D}_{i^*j}' be the reported distribution for i^* (if there are many reported distributions for him, take the one by the lexicographically first reporter).
 - 5: Let $p_{i^*} = \max_p \Pr_{v_{i^*j} \sim \mathcal{D}_{i^*j}'}(v_{i^*j} \geq p \mid v_{i^*j} \geq p_{second}) \cdot p$. If $b_{i^*j} \geq p_{i^*}$ then sell item j to i^* at price p_{i^*} ; otherwise the item is unsold.
-

Note that $\mathcal{M}_{IE1LA,j}$ does not partition the players into two groups. Also, it is not exactly using the 1-lookahead mechanism as a blackbox, because it has to handle boundary cases where the players' distributions are not all reported. When the players all tell the truth, all true distributions will indeed be reported. However, for the mechanism to be well defined, it has to know what to do in all possible cases. Moreover, running the 1-lookahead mechanism on the set of players whose distributions are reported is not 2-DST: for example, if the player with the second highest value is the only one who knows the distribution for the player with the highest value, then she may choose not to report her knowledge about the latter, so that she herself has the highest value in the 1-Lookahead mechanism and gets a high utility. That is why the mechanism only tries to sell to the player with the highest value. We have the following two lemmas.

Lemma 5 *For any additive auction instances $\hat{\mathcal{I}} = (N, M, \mathcal{D})$ and $\mathcal{I} = (N, M, \mathcal{D}, G)$, mechanism $\mathcal{M}_{IE1LA,j}$ is 2-DST for each \mathcal{I}_j , and mechanism \mathcal{M}_{IE1LA} is 2-DST for \mathcal{I} .*

Proof As in mechanism \mathcal{M}_{IEA} , in each mechanism $\mathcal{M}_{IE1LA,j}$, the fact that it is dominant for each player i to report her true value no matter what knowledge the players report follows from the truthfulness of the second-price mechanism and that of the 1-lookahead mechanism. Given that all players report their true values, a player i 's reported knowledge does not affect whether she is i^* or not. It may affect the other players' utilities, but not her own. Thus reporting her true knowledge never hurts him, and mechanism $\mathcal{M}_{IE1LA,j}$ is 2-DST for \mathcal{I}_j .

Since the players have additive valuations and \mathcal{M}_{IE1LA} runs each mechanism $\mathcal{M}_{IE1LA,j}$ separately for item j , we have that \mathcal{M}_{IE1LA} is 2-DST for \mathcal{I} and Lemma 5 holds. \square

Lemma 6 $\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IE1LA}(\mathcal{I})) = \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{1LA}(\hat{\mathcal{I}})) \geq \frac{1}{2} \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\text{IM}(\hat{\mathcal{I}})).$

Proof When the players report their true values and true knowledge, the outcome of each $\mathcal{M}_{IE1LA,j}$ on the information elicitation instance \mathcal{I}_j is the same as that of mechanism \mathcal{M}_{1LA} on the Bayesian instance $\hat{\mathcal{I}}_j$, because the distribution for i^* is reported. Accordingly,

$$\begin{aligned} \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IE1LA}(\mathcal{I})) &= \sum_{j \in M} \mathbb{E}_{v_j \sim \mathcal{D}_j} \text{Rev}(\mathcal{M}_{IE1LA,j}(\mathcal{I}_j)) = \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{1LA}(\hat{\mathcal{I}})) \\ &\geq \sum_{j \in M} \frac{1}{2} \text{OPT}(\hat{\mathcal{I}}_j) = \frac{1}{2} \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\text{IM}(\hat{\mathcal{I}})), \end{aligned}$$

where the inequality is because the 1-lookahead mechanism is a 2-approximation to the optimal Bayesian mechanism for each item j [60]. Thus Lemma 6 holds. \square

Note that the approximation ratio of \mathcal{M}_{IE1LA} does not depend on the specific value of k , as long as $k \geq 1$.

The information elicitation BVCG mechanism The mechanism \mathcal{M}_{IEBVCG} is defined in Mechanism Algorithm 5. It is similar to \mathcal{M}_{IEA} and approximates mechanism BVCG in information elicitation settings. If a player i 's value distributions are not all reported,

\mathcal{M}_{IEBVC} throws i away and leaves his winning set unsold. This simplifies the instructions compared to \mathcal{M}_{IEA} and still ensures truthfulness. Doing so would seriously damage the revenue if the knowledge graphs can be totally arbitrary. However, when everything is known by somebody and when the players report their true knowledge, no player is actually thrown away. We have the following two lemmas.

Algorithm 5 \mathcal{M}_{IEBVC}

- 1: Each player i reports a valuation $b_i = (b_{ij})_{j \in M}$ and a knowledge $K_i = (\mathcal{D}'_{i'j})_{i' \neq i, j \in M}$.
 - 2: For each item j , set $i^*(j) = \operatorname{argmax}_i b_{ij}$ (ties broken lexicographically) and $p_j = \max_{i \neq i^*} b_{ij}$.
 - 3: **for** each player i **do**
 - 4: Let $M_i = \{j \mid i^*(j) = i\}$ be player i 's winning set.
 - 5: If not all m distributions of i 's values are reported, i gets no item and items in M_i are unsold.
 - 6: Otherwise, let \mathcal{D}'_i be the vector of reported distributions for i 's values (if there are more than one reporters for an item, take the lexicographically first).
 - 7: Compute the entry fee $e_i(\mathcal{D}'_i, b_{-i})$ using $BVCG$. Note that different from mechanism *Bund*, $BVCG$ does not compute extra reserve prices for i .
 - 8: Sell M_i to player i according to $BVCG$. That is, if $\sum_{j \in M_i} b_{ij} \geq e_i(\mathcal{D}'_i, b_{-i}) + \sum_{j \in M_i} p_j$ then i gets M_i with price $e_i(\mathcal{D}'_i, b_{-i}) + \sum_{j \in M_i} p_j$; otherwise i gets no item and the items in M_i are unsold.
 - 9: **end for**
-

Lemma 7 Mechanism \mathcal{M}_{IEBVC} is 2-DST.

Proof Arbitrarily fix a player i , a strategy subprofile of the other players, and a knowledge of i . If not all m distributions of i 's values are reported by the others, then i gets nothing and pays nothing, so it does not matter what valuation she reports about himself. Otherwise, \mathcal{M}_{IEBVC} sells to player i in the same way as $BVCG$: using the other players' highest reported value as the reserve price for each item, either player i gets the whole set of items for which her value passes the reserve price (i.e., her winning set), or she gets nothing and those items are unsold to anybody. Following [11], it is dominant for i to report her true values given any entry fee that does not depend on her reported values, so is it when the entry fee is computed based on \mathcal{D}'_i and b_{-i} .

Moreover, a player i 's reported knowledge K_i about others affects neither M_i nor e_i , nor the reserve prices for him. Thus reporting her true knowledge never hurts her and Lemma 7 holds. \square

Lemma 8 $\mathbb{E}_{v \sim \mathcal{D}} \operatorname{Rev}(\mathcal{M}_{IEBVC}(\mathcal{I})) = \mathbb{E}_{v \sim \mathcal{D}} \operatorname{Rev}(BVCG(\hat{\mathcal{I}})).$

Proof Since \mathcal{M}_{IEBVC} retrieves the whole distribution \mathcal{D} from the players, its outcome is exactly the same as that of $BVCG$ under the Bayesian instance $\hat{\mathcal{I}}$. \square

Remark Similar to \mathcal{M}_{IE1LA} , the revenue of \mathcal{M}_{IEBVC} does not depend on the specific value of k , as long as $k \geq 1$. Indeed, notice the special structures of the two Bayesian mechanisms \mathcal{M}_{1LA} and $BVCG$: the winning set of a player i solely depends on the players' values; the distribution \mathcal{D}_i is only used to compute better reserve prices or entry fee to increase revenue; and the distribution \mathcal{D}_{-i} is irrelevant to i . Therefore, in the information elicitation setting we can allow a player to be both a reporter about the others' distributions and a

potential buyer of some items. In some other mechanisms such as Myerson's mechanism, all players' distributions are used both to choose the potential winner and to set her price, thus in the information elicitation setting we must separate the knowledge reporters and the potential winners.

We are now ready to prove Theorem 3.

Proof of Theorem 3 Recall that mechanism \mathcal{M}'_{IEA} is defined as follows: when $k \leq 7$, it runs \mathcal{M}_{IEBVCg} with probability $\frac{2}{11}$ and \mathcal{M}_{IE1LA} with probability $\frac{9}{11}$; when $k > 7$, it runs \mathcal{M}_{IEBVCg} with probability $\frac{\tau_k}{3+\tau_k}$ and \mathcal{M}_{IEIM} with probability $\frac{3}{3+\tau_k}$.

The mechanism \mathcal{M}'_{IEA} is clearly 2-DST, since all the sub-mechanisms are 2-DST, and which mechanism is chosen does not depend on the players' strategies. When $k \leq 7$, we have $\frac{\tau_k}{6+2\tau_k} < \frac{1}{11}$. By Lemmas 6 and 8,

$$\begin{aligned} \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}'_{IEA}(\mathcal{I})) &= \frac{2}{11} \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEBVCg}(\mathcal{I})) + \frac{9}{11} \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IE1LA}(\mathcal{I})) \\ &\geq \frac{2}{11} \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(BVCg(\hat{\mathcal{I}})) + \frac{3}{11} \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(IM(\hat{\mathcal{I}})) + \frac{3}{11} \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{1LA}(\hat{\mathcal{I}})). \end{aligned} \quad (3)$$

When $k > 7$, we have $\frac{\tau_k}{6+2\tau_k} > \frac{1}{11}$. By Lemmas 4 and 8,

$$\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}'_{IEA}(\mathcal{I})) \quad (4)$$

$$\begin{aligned} &= \frac{\tau_k}{3+\tau_k} \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEBVCg}(\mathcal{I})) + \frac{3}{3+\tau_k} \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEIM}(\mathcal{I})) \\ &\geq \frac{\tau_k}{3+\tau_k} \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(BVCg(\hat{\mathcal{I}})) + \frac{3\tau_k}{3+\tau_k} \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(IM(\hat{\mathcal{I}})). \end{aligned} \quad (5)$$

By [9],

$$\begin{aligned} OPT(\hat{\mathcal{I}}) &\leq 2 \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(BVCg(\hat{\mathcal{I}})) + 3 \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(IM(\hat{\mathcal{I}})) + 3 \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{1LA}(\hat{\mathcal{I}})) \\ &\leq 2 \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(BVCg(\hat{\mathcal{I}})) + 6 \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(IM(\hat{\mathcal{I}})). \end{aligned}$$

Combined the first inequality with Inequality (3) and the second with Inequality (5), we have

$$\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}'_{IEA}(\mathcal{I})) \geq \max \left\{ \frac{1}{11}, \frac{\tau_k}{6+2\tau_k} \right\} \cdot OPT(\hat{\mathcal{I}}),$$

and Theorem 3 holds. \square

Note that when the knowledge graphs are 2-connected, instead of using mechanism \mathcal{M}_{IE1LA} or \mathcal{M}_{IEIM} , one can use \mathcal{M}_{IEM} for each item j . We thus have the following

corollary, where the mechanism \mathcal{M}_{IEA}'' runs \mathcal{M}_{IEM} with probability $\frac{3}{4}$ and \mathcal{M}_{IEBVCG} with probability $\frac{1}{4}$.

Corollary 1 *For any additive auction instances $\hat{\mathcal{I}} = (N, M, \mathcal{D})$ and $\mathcal{I} = (N, M, \mathcal{D}, G)$ where each G_j is 2-connected, mechanism \mathcal{M}_{IEA}'' is 2-DST and $\mathbb{E}_{\mathcal{D}} \text{Rev}(\mathcal{M}_{IEA}''(\mathcal{I})) \geq \frac{1}{8}(1 - \frac{1}{n}) \text{OPT}(\hat{\mathcal{I}})$.*

5 Under arbitrary knowledge graphs

5.1 Knowledge-based revenue benchmark

When the knowledge graphs can be totally arbitrary, some distributions may not be known by anybody. It is not hard to see that in this case, no information elicitation mechanism can be a bounded approximation to OPT . Indeed, if all but one value distributions of the players are constantly 0, and if the only non-zero distribution, denoted by \mathcal{D}_{ij} , is unknown by anybody, then a Bayesian mechanism can find the optimal reserve price based on \mathcal{D}_{ij} , while an information elicitation mechanism can only set the price for player i based on the reported values of the other players, which are all 0.

Thus, for arbitrary knowledge graphs, we define a natural revenue benchmark: the optimal Bayesian revenue *on players and items for which the distributions are known in the information elicitation setting*. More precisely, let $\hat{\mathcal{I}} = (N, M, \mathcal{D})$ be a Bayesian instance and $\mathcal{I} = (N, M, \mathcal{D}, G)$ a corresponding information elicitation instance. Let $\mathcal{D}' = \times_{i \in N, j \in M} \mathcal{D}'_{ij}$ be such that $\mathcal{D}'_{ij} = \mathcal{D}_{ij}$ if there exists a player i' with $(i', i) \in G_j$, and \mathcal{D}'_{ij} is constantly 0 otherwise. We refer to \mathcal{D}' as \mathcal{D} *projected on G* . Letting $\mathcal{I}' = (N, M, \mathcal{D}')$ be the resulting Bayesian instance, the *knowledge-based revenue benchmark* is $\text{OPT}_K(\mathcal{I}) \triangleq \text{OPT}(\mathcal{I}')$, the optimal BIC revenue on \mathcal{I}' . This is a demanding benchmark in information elicitation settings: it takes into consideration the knowledge of *all* players, no matter who knows what. When everything is known by somebody, even if G is only 1-informed, we will have $\mathcal{I}' = \hat{\mathcal{I}}$ and $\text{OPT}_K(\mathcal{I}) = \text{OPT}(\hat{\mathcal{I}})$.

5.2 Unit-demand auctions

When some distributions may not be known by anybody, we design mechanism \mathcal{M}_{IEUD} which is very similar as \mathcal{M}_{IEUD}^k ; see Mechanism Algorithm 6, where \mathcal{M}_{UD} is also the Bayesian mechanism of [8]. It's worth pointing out that, although mechanism \mathcal{M}_{UD} is used as a black-box, Mechanism Algorithm 6 is not a reduction from arbitrary Bayesian mechanisms. Instead, we will prove a *projection lemma* that allows such a reduction from an important class of Bayesian mechanisms, where mechanism \mathcal{M}_{UD} is an important example. We have the following theorem.

Algorithm 6 \mathcal{M}_{IEUD}

-
- 1: Each player i reports to the seller a valuation $b_i = (b_{ij})_{j \in M}$ and a knowledge $K_i = (\mathcal{D}_{i'j}^i)_{i' \neq i, j \in M}$.
 - 2: Randomly partition the players into two sets, N_1 and N_2 , where each player is independently put in each set with probability $\frac{1}{2}$.
 - 3: Set $N_3 = \emptyset$.
 - 4: **for** players $i \in N_1$ lexicographically **do**
 - 5: For each player $i' \in N_2$ and item $j \in M$, if $\mathcal{D}_{i'j}^{i'}$ has not been defined yet and $\mathcal{D}_{i'j}^i \neq \perp$, then set $\mathcal{D}_{i'j}^{i'} = \mathcal{D}_{i'j}^i$ and add player i' to N_3 .
 - 6: **end for**
 - 7: For each $i \in N_3$ and $j \in M$ such that $\mathcal{D}_{ij}^{i'}$ is not defined, set $\mathcal{D}_{ij}^{i'} \equiv 0$ (i.e., 0 with probability 1) and $b_{ij} = 0$.
 - 8: Run mechanism \mathcal{M}_{UD} on the unit-demand Bayesian auction $(N_3, M, (\mathcal{D}_{ij}^{i'})_{i \in N_3, j \in M})$, with the players' values being $(b_{ij})_{i \in N_3, j \in M}$. Let $x' = (x'_{ij})_{i \in N_3, j \in M}$ be the resulting allocation where $x'_{ij} \in \{0, 1\}$, and let $p' = (p'_i)_{i \in N_3}$ be the prices. Without loss of generality, $x'_{ij} = 0$ if $\mathcal{D}_{ij}^{i'} \equiv 0$.
 - 9: For each player $i \notin N_3$, i gets no item and her price is $p_i = 0$.
 - 10: For each player $i \in N_3$, i gets item j if $x'_{ij} = 1$, and her price is $p_i = p'_i$.
-

Theorem 4 Mechanism \mathcal{M}_{IEUD} for unit-demand auctions is 2-DST and, for any instances $\hat{\mathcal{I}} = (N, M, \mathcal{D})$ and $\mathcal{I} = (N, M, \mathcal{D}, G)$, $\text{Rev}(\mathcal{M}_{IEUD}(\mathcal{I})) \geq \frac{\text{OPT}_K(\mathcal{I})}{96}$.

The 2-DSTness follows from Lemma 2, and in the following we prove the revenue bound.

The COPIES setting Before analyzing the revenue of mechanism \mathcal{M}_{IEUD} , we first recall the COPIES setting for reducing multi-parameter settings to single-parameter settings [10]. Given a unit-demand auction instance $\hat{\mathcal{I}} = (N, M, \mathcal{D})$, the corresponding COPIES instance is constructed as follows. We make m copies for each player, called player copies, and denote the resulting player set by $N^{CP} = N \times M$. We make n copies for each item, called item copies, and denote the resulting item set by $M^{CP} = M \times N$. Each player copy (i, j) has value $v_{ij} \sim \mathcal{D}_{ij}$ for the item copy (j, i) , and 0 for all the other item copies.

The set of feasible allocations in the original unit-demand auction naturally defines the set of feasible allocations in the COPIES auction: for any feasible allocation A in the original setting, if player i gets item j , then in the corresponding allocation in the COPIES setting, player i 's copy j gets item j 's copy i , and all other copies of i get nothing. Since in the original setting each item is sold to at most one player, in the COPIES setting, for all copies of the same item, at most one of them is sold. Moreover, since each player gets one item in the original setting, in the COPIES setting, for all copies of the same player, at most one of them gets an item copy. We denote by $\hat{\mathcal{I}}^{CP} = (N^{CP}, M^{CP}, \mathcal{D}^{CP})$ the COPIES instance, with $\mathcal{D}^{CP} = \times_{(i,j) \in N^{CP}} \mathcal{D}_{ij}$.

The projected setting Next, we consider the optimal Bayesian revenue for the COPIES setting when “projected” to smaller instances. Let $NM \subseteq N \times M$ be a subset of player-item pairs and N' be NM projected to N . The unit-demand instance $\hat{\mathcal{I}}_{NM} = (N', M, \mathcal{D}'_{N'})$, referred to as $\hat{\mathcal{I}}$ projected to NM , is such that $\mathcal{D}'_{ij} = \mathcal{D}_{ij}$ if $(i, j) \in NM$, and $\mathcal{D}'_{ij} \equiv 0$ otherwise.

Let $\hat{\mathcal{I}}_{NM}^{CP} = (N'^{CP}, M^{CP}, \mathcal{D}'_{N'}^{CP})$ be the COPIES instance corresponding to $\hat{\mathcal{I}}_{NM}$. It can also be considered as $\hat{\mathcal{I}}^{CP}$ projected to NM . That is, when projecting a COPIES instance to a set of player-item pairs, we still want the resulting instance to be a COPIES instance, thus we patch it up with the missing player-item pairs but with values constantly 0. The relations




Fig. 2 Relations of the COPIES and the projected instances

of these instances are illustrated in Fig. 2. We are interested in the optimal Bayesian revenue under $\hat{\mathcal{I}}_{NM}^{CP}$, $OPT(\hat{\mathcal{I}}_{NM}^{CP})$.

Let $OPT(\hat{\mathcal{I}}^{CP})_{NM}$ be the revenue of the optimal Bayesian mechanism for $\hat{\mathcal{I}}^{CP}$ obtained from NM: that is, for any pair $(i, j) \notin NM$, the price paid by the j -th copy of player i is not counted, even though this player copy may get the i -th copy of item j according to the optimal mechanism for $\hat{\mathcal{I}}^{CP}$. We have the following lemma, where we explicitly write out the distributions for different Bayesian instances.

Lemma 9 (The projection lemma) For any $\hat{\mathcal{I}}$ and $NM \subseteq N \times M$, $OPT(\hat{\mathcal{I}}_{NM}^{CP}) \geq OPT(\hat{\mathcal{I}}^{CP})_{NM}$.

Proof Given the instance $\hat{\mathcal{I}}_{NM}^{CP}$, consider a Bayesian mechanism \mathcal{M}' as follows:

- Patch up the instance to be exactly $\hat{\mathcal{I}}^{CP}$, including changing the distribution \mathcal{D}'_{ij} from 0 to \mathcal{D}_{ij} for any (i, j) with $i \in N'$ and $(i, j) \notin NM$;
- For any $(i, j) \in NM$, let i 's copy j report her value, while for any $(i, j) \notin NM$, sample the value of i 's copy j from \mathcal{D}_{ij} ;
- Run the optimal DST Bayesian mechanism on $\hat{\mathcal{I}}^{CP}$, with the reported and the sampled values;⁶
- Project the resulting outcome to NM .

The key is to show that \mathcal{M}' is a DST Bayesian mechanism for $\hat{\mathcal{I}}_{NM}^{CP}$. Indeed, for any (i, j) such that $i \in N'$ and $(i, j) \notin NM$, the value of i 's copy j in $\hat{\mathcal{I}}_{NM}^{CP}$ is 0 with probability 1, his reported value is not used by \mathcal{M}' , and at the end this player copy gets nothing and pays 0. Therefore it is dominant for this player copy to report her true value 0. For any $(i, j) \in NM$, the utility of i 's copy j under \mathcal{M}' is the same as that under the optimal DST Bayesian mechanism on $\hat{\mathcal{I}}^{CP}$. As it is dominant for this player copy to report her true value in the latter, so it is in \mathcal{M}' . Accordingly,

$$\mathbb{E}_{\mathcal{D}'_{N'}^{CP}} OPT(\hat{\mathcal{I}}_{NM}^{CP}) \geq \mathbb{E}_{\mathcal{D}'_{N'}^{CP}} Rev(\mathcal{M}'(\hat{\mathcal{I}}_{NM}^{CP})),$$

by the definition of OPT . By construction we have

$$\mathbb{E}_{\mathcal{D}'_{N'}^{CP}} Rev(\mathcal{M}'(\hat{\mathcal{I}}_{NM}^{CP})) = \mathbb{E}_{\mathcal{D}^{CP}} OPT(\hat{\mathcal{I}}^{CP})_{NM},$$

thus Lemma 9 holds. \square

⁶Since COPIES is a single-parameter setting, Myerson's mechanism is both the optimal BIC mechanism and the optimal DST mechanism here. In particular, the optimal BIC revenue equals the optimal DST revenue.

Note that the projection lemma is not true for non-COPIES settings in general, because the corresponding mechanism \mathcal{M}' is not DST. Indeed, when the same player i has $(i, j) \in NM$ and $(i, j') \notin NM$ for some items j and j' , she prefers receiving j to receiving j' in the patched-up auction, even if the former leads to a smaller utility: her projected utility will be 0 under the latter.

Now we are ready to finish the proof of Theorem 4.

Proof of Theorem 4 Letting $\mathcal{I}' = (N, M, \mathcal{D}')$ be the Bayesian instance where \mathcal{D}' is \mathcal{D} projected on the knowledge graph G , the COPIES instance \mathcal{I}'^{CP} is respectively defined. Following Lemma 2 and by the definition of OPT_K , it remains to show

$$\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{I\text{EUD}}(\mathcal{I})) \geq \frac{OPT(\mathcal{I}')}{96}.$$

For any player i and item j with \mathcal{D}_{ij} known by some other players, we have

$$\begin{aligned} & \Pr(\mathcal{D}_{ij} \text{ is reported in the mechanism}) \\ &= \Pr(i \in N_2) \Pr(\exists i' \in N_1 \text{ s.t. } (i', i) \in G_j \mid i \in N_2) = \frac{1}{4}, \end{aligned}$$

where the equality is because there exists at least one player i' with $(i', i) \in G_j$, and players i and i' are partitioned independently. Below we use NM_3 to denote the set of player-item pairs whose distribution is reported in the mechanism: that is, the set of players N_3 together with their reported items. Accordingly,

$$\Pr((i, j) \in NM_3) \geq \frac{1}{4}. \quad (6)$$

Also, we use $\mathcal{D}_{NM_3} = \times_{(i,j) \in NM_3} \mathcal{D}_{ij}$ and $v_{NM_3} = (v_{ij})_{(i,j) \in NM_3}$ to denote the vector of distributions and the vector of true values for NM_3 , respectively. Let $\hat{\mathcal{I}}_{NM_3} = (N_3, M, \mathcal{D}'_{N_3})$ be the unit-demand instance given to \mathcal{M}_{UD} in Step 8. Note that it is exactly \mathcal{I}' projected to NM_3 : that is, $\mathcal{D}'_{ij} = \mathcal{D}_{ij}$ for any $(i, j) \in NM_3$ and $\mathcal{D}'_{ij} \equiv 0$ for any $(i, j) \notin NM_3$. Accordingly,

$$\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{I\text{EUD}}(\mathcal{I})) = \mathbb{E}_{NM_3} \mathbb{E}_{v_{NM_3} \sim \mathcal{D}_{NM_3}} \text{Rev}(\mathcal{M}_{UD}(\hat{\mathcal{I}}_{NM_3})). \quad (7)$$

Let $\hat{\mathcal{I}}_{NM_3}^{CP} = (N_3^{CP}, M^{CP}, \mathcal{D}'_{N_3}^{CP})$ be the COPIES instance corresponding to $\hat{\mathcal{I}}_{NM_3}$. Following [8], given any set NM_3 , the revenue of \mathcal{M}_{UD} under the unit-demand Bayesian instance $\hat{\mathcal{I}}_{NM_3}$ is a 6-approximation to the optimal revenue for the COPIES instance. That is, for any NM_3 ,

$$\mathbb{E}_{v_{NM_3} \sim \mathcal{D}_{NM_3}} \text{Rev}(\mathcal{M}_{UD}(\hat{\mathcal{I}}_{NM_3})) \geq \frac{1}{6} \mathbb{E}_{\mathcal{D}'_{N_3}^{CP}} OPT(\hat{\mathcal{I}}_{NM_3}^{CP}). \quad (8)$$

Combining Inequalities (7) and (8), we have

$$\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEUD}(\mathcal{I})) \geq \frac{1}{6} \mathbb{E}_{NM_3} \mathbb{E}_{\mathcal{D}'_{N_3}{}^{CP}} \text{OPT}(\hat{\mathcal{I}}_{NM_3}^{CP}). \quad (9)$$

By Lemma 9,

$$\mathbb{E}_{\mathcal{D}'_{N_3}{}^{CP}} \text{OPT}(\hat{\mathcal{I}}_{NM_3}^{CP}) \geq \mathbb{E}_{\mathcal{D}^{CP}} \text{OPT}(\mathcal{I}'^{CP})_{NM_3},$$

thus

$$\mathbb{E}_{NM_3} \mathbb{E}_{\mathcal{D}'_{N_3}{}^{CP}} \text{OPT}(\hat{\mathcal{I}}_{NM_3}^{CP}) \geq \mathbb{E}_{NM_3} \mathbb{E}_{\mathcal{D}^{CP}} \text{OPT}(\mathcal{I}'^{CP})_{NM_3}. \quad (10)$$

Let $P_{ij}(\text{OPT}(\mathcal{I}'^{CP}))$ be the price paid by player i 's copy j under the optimal mechanism for \mathcal{I}'^{CP} . We can rewrite the right-hand side of (10) as follows.

$$\begin{aligned} \mathbb{E}_{NM_3} \mathbb{E}_{\mathcal{D}^{CP}} \text{OPT}(\mathcal{I}'^{CP})_{NM_3} &= \mathbb{E}_{NM_3} \mathbb{E}_{\mathcal{D}^{CP}} \sum_{(i,j) \in NM_3} P_{ij}(\text{OPT}(\mathcal{I}'^{CP})) \\ &= \mathbb{E}_{\mathcal{D}^{CP}} \mathbb{E}_{NM_3} \sum_{(i,j) \in NM_3} P_{ij}(\text{OPT}(\mathcal{I}'^{CP})) \\ &= \mathbb{E}_{\mathcal{D}^{CP}} \sum_{(i,j) \in N \times M} \Pr((i,j) \in NM_3) \cdot P_{ij}(\text{OPT}(\mathcal{I}'^{CP})), \end{aligned} \quad (11)$$

where the first equality is by the definition of the projection, the second is because sampling from \mathcal{D}^{CP} is done independently from NM_3 , and the third is because $P_{ij}(\text{OPT}(\hat{\mathcal{I}}^{CP}))$ does not depend on NM_3 . We can further lower-bound the last term of (11) as follows:

$$\begin{aligned} &\mathbb{E}_{\mathcal{D}^{CP}} \sum_{(i,j) \in N \times M} \Pr((i,j) \in NM_3) \cdot P_{ij}(\text{OPT}(\mathcal{I}'^{CP})) \\ &= \mathbb{E}_{\mathcal{D}^{CP}} \sum_{(i,j) \in N \times M: \exists i' \text{ s.t. } (i',i) \in G_j} \Pr((i,j) \in NM_3) \cdot P_{ij}(\text{OPT}(\mathcal{I}'^{CP})) \\ &\geq \frac{1}{4} \mathbb{E}_{\mathcal{D}^{CP}} \sum_{(i,j) \in N \times M: \exists i' \text{ s.t. } (i',i) \in G_j} P_{ij}(\text{OPT}(\mathcal{I}'^{CP})) \\ &= \frac{1}{4} \text{OPT}(\mathcal{I}'^{CP}) \geq \frac{1}{16} \text{OPT}(\mathcal{I}'). \end{aligned} \quad (12)$$

Here the first equality is because $P_{ij}(\text{OPT}(\mathcal{I}'^{CP})) = 0$ for every (i,j) such that \mathcal{D}_{ij} is unknown, the first inequality is by (6), the second equality is by the definition of revenue, and the second inequality is because $\text{OPT}(\mathcal{I}'^{CP}) \geq \frac{\text{OPT}(\mathcal{I}')}{4}$ by [9].

Combining (9), (10), (11) and (12), we have

$$\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEUD}(\mathcal{I})) \geq \frac{1}{6} \cdot \frac{1}{16} \cdot \text{OPT}(\mathcal{I}') = \frac{\text{OPT}(\mathcal{I}')}{96},$$

and Theorem 4 holds. \square

Note that Lemma 9 is only concerned with COPIES instances. Using this lemma and similar to our proof of Theorem 4, any Bayesian mechanism \mathcal{M} whose revenue can be properly lower-bounded by the COPIES instance can be converted to an information elicitation mechanism in a black-box way. We have the following theorem, with the proof omitted.

Theorem 5 *Let \mathcal{M} be any DST Bayesian mechanism such that $\text{Rev}(\mathcal{M}(\hat{\mathcal{I}})) \geq \alpha \text{OPT}(\hat{\mathcal{I}}^{CP})$ for some $\alpha > 0$. There exists a 2-DST information elicitation mechanism that uses \mathcal{M} as a black-box and is a $\frac{\alpha}{16}$ -approximation to OPT_K .*

By Theorem 5, the mechanisms in [4] and [10] automatically imply information elicitation mechanisms. For single-good auctions, replacing mechanism \mathcal{M}_{UD} with Myerson's mechanism, the information elicitation mechanism is a 4-approximation to OPT_K .

5.3 Additive auctions

For additive auctions, recall that we can never throw away any information. If a player i 's value distribution for an item j is reported by others, then j may be sold to i via the β -Bundling mechanism of [11] (see Sections 2 and 3 therein), denoted by $Bund$. If i 's distribution for j is not reported, then j may still be sold to i via the second-price mechanism. Indeed, our mechanism handles the players neither solely based on the original Bayesian instance $\hat{\mathcal{I}}$ nor solely based on the projected instance \mathcal{I}' . Rather, it works on a *hybrid* of the two.

Our mechanism \mathcal{M}_{IEA} is still simple; see Mechanism Algorithm 7. However, significant effort is needed to analyze its revenue. Indeed, note that in Mechanism Algorithm 7, each M_i – the *winning set* of agent i – is defined according to the original Bayesian instance $\hat{\mathcal{I}}$, while the partition of M is done according to the knowledge graphs in the information elicitation instance \mathcal{I} , where M_i^1 contains the items for which i 's distribution is reported and M_i^2 contains the ones that have not been reported by anyone. The mechanism $Bund$ is run on a hybrid instance, where β_i is based on $\hat{\mathcal{I}}$ and \mathcal{D}'_i is based on \mathcal{I} . The mechanism $Bund$ then returns the optimal entry fee e_i and reserve prices $(p'_j)_{j \in M_i^1}$. If $e_i > 0$, the items in $M_i^1 \cap M_i$ are sold to player i with price $e_i + \sum_{j \in M_i^1 \cap M_i} p'_j$ if $\sum_{j \in M_i^1 \cap M_i} b_{ij} \geq e_i + \sum_{j \in M_i^1 \cap M_i} p'_j$; otherwise the items in $M_i^1 \cap M_i$ are not sold. If $e_i = 0$, for each item $j \in M_i^1 \cap M_i$, player i gets item j with price p'_j if $b_{ij} \geq p'_j$; otherwise item j is not sold. Finally, the items in $M_i^2 \cap M_i$ are sold using second price.

Algorithm 7 \mathcal{M}_{IEA}

-
- 1: Each player i reports a valuation $b_i = (b_{ij})_{j \in M}$ and a knowledge $K_i = (\mathcal{D}_{i'j}^i)_{i' \neq i, j \in M}$.
 - 2: For each item j , set $i^*(j) = \operatorname{argmax}_i b_{ij}$ (ties broken lexicographically) and $p_j = \max_{i \neq i^*} b_{ij}$.
 - 3: **for** each player i **do**
 - 4: Let $M_i = \{j \mid i^*(j) = i\}$ be player i 's *winning set*.
 - 5: Partition M into M_i^1 and M_i^2 as follows: for all $j \in M_i^1$, some i' has reported $\mathcal{D}_{ij}^{i'} \neq \perp$ (if there are more than one reporters, take the lexicographically first); and for all $j \in M_i^2$, $\mathcal{D}_{ij}^{i'} = \perp$ for all i' .
 - 6: $\forall j \in M_i^1$, set $\mathcal{D}_{ij}^{i'} = \mathcal{D}_{ij}^{i'}$; and $\forall j \in M_i^2$, set $\mathcal{D}_{ij}^{i'} \equiv 0$.
 - 7: Compute the optimal entry fee e_i and reserve prices $(p'_j)_{j \in M_i^1}$ according to mechanism *Bund* with respect to $(\mathcal{D}_{ij}^{i'}, \beta_j)$, where $\beta_{ij} = \max_{i' \neq i} b_{i'j} \ \forall j \in M$. By the definition of *Bund*, we always have $p'_j \geq \beta_{ij}$ for each j . If $e_i = 0$ then it is possible that $p'_j > \beta_{ij}$ for some j ; while if $e_i > 0$ then $p'_j = \beta_{ij}$ for every j .
 - 8: Sell $M_i^1 \cap M_i$ to player i according to *Bund*. That is, if $e_i > 0$ then do the following: if $\sum_{j \in M_i^1 \cap M_i} b_{ij} \geq e_i + \sum_{j \in M_i^1 \cap M_i} p'_j$, player i gets $M_i^1 \cap M_i$ with price $e_i + \sum_{j \in M_i^1 \cap M_i} p'_j$; otherwise the items in $M_i^1 \cap M_i$ are not sold. If $e_i = 0$ then do the following: for each item $j \in M_i^1 \cap M_i$, if $b_{ij} \geq p'_j$, player i gets item j with price p'_j ; otherwise item j is not sold.
 - 9: In addition, sell each item j in $M_i^2 \cap M_i$ to player i with price $p_j (= \beta_{ij})$.
 - 10: **end for**
-

Although running *Bund* on \mathcal{I}' achieves a constant approximation to $OPT(\mathcal{I}')$, some items sold by *Bund* under \mathcal{I}' may end up being sold by \mathcal{M}_{IEA} using second-price, and the revenue of \mathcal{M}_{IEA} cannot be lower-bounded by that of *Bund* on \mathcal{I}' . To overcome this difficulty, we develop a novel way to use the *adjusted revenue* [11] in our analysis; see Lemmas 11 and 13, where we also recall the related definitions. As we show there, the adjusted revenue in a hybrid information setting, combined with the revenue of the second-price sale, eventually provides a desirable lower-bound to the revenue of \mathcal{M}_{IEA} . We have the following theorem.

Theorem 6 *Mechanism \mathcal{M}_{IEA} for additive auctions is 2-DST and, for any instances $\hat{\mathcal{I}} = (N, M, \mathcal{D})$ and $\mathcal{I} = (N, M, \mathcal{D}, G)$, $\operatorname{Rev}(\mathcal{M}_{IEA}(\mathcal{I})) \geq \frac{OPT_K(\mathcal{I})}{70}$.*

Lemma 10 *Mechanism \mathcal{M}_{IEA} is 2-DST.*

Proof The structure of a detailed proof for Lemma 10 will be following the two requirements in the solution concept and similar to that of Lemma 2. Thus we only highlight the key points here.

For each player i , given the distributions and values reported by the other players, the subsets M_i^1 and M_i^2 do not depend on player i 's strategy, neither do the entry fee e_i and reserve prices $(p'_j)_{j \in M_i^1}$. As far as player i is concerned, the other players' values are always taken to be b_{-i} , even if some of their value distributions are not reported. Computing i 's winning sets $M_i \cap M_i^1$ and $M_i \cap M_i^2$ is just part of the two mechanisms, *Bund* and second-price, again with the other players' values taken to be b_{-i} . Thus, it is dominant for i to report her true values for M_i^1 , following the truthfulness of *Bund*; and it is dominant for her to report her true values for M_i^2 , following the truthfulness of the second-price mechanism.

Moreover, given that all players truthfully report their values, for each player i , reporting her true knowledge never hurts him, no matter what knowledge the other players report. Indeed, the winning set M_i only depends on the players' reported values. Player i 's reported knowledge may affect how M is partitioned into M_i^1 and M_i^2 for another player i' , but does not affect the sets M_i^1 and M_i^2 , or whether she gets some items or not, or the prices she pays. Thus \mathcal{M}_{IEA} is 2-DST as desired. \square

The adjusted revenue To lower-bound the expected revenue of \mathcal{M}_{IEA} , we first introduce several important concepts following [11].

For any single-player Bayesian instance $\hat{\mathcal{I}}_i = (\{i\}, M, \mathcal{D}_i)$ and any non-negative reserve-price vector $\beta_i = (\beta_{ij})_{j \in M}$, a single-player DST Bayesian mechanism is β_i -exclusive if it never sells an item j to i whenever her bid for j is no larger than β_{ij} . Denote by $Rev^X(\hat{\mathcal{I}}_i, \beta_i)$ the optimal β_i -exclusive revenue for $\hat{\mathcal{I}}_i$: that is, the supremum over the revenue of β_i -exclusive mechanisms.

For any single-player DST Bayesian mechanism, its β_i -adjusted revenue on $\hat{\mathcal{I}}_i$ is its revenue minus its social welfare generated from player-item pairs (i, j) such that i 's bid for j is no larger than β_{ij} . Denote by $Rev^A(\hat{\mathcal{I}}_i, \beta_i)$ the optimal β_i -adjusted revenue for $\hat{\mathcal{I}}_i$: that is, the supremum over the β_i -adjusted revenue of all single-player DST Bayesian mechanisms. Note that if a mechanism is β_i -exclusive, then its β_i -adjusted revenue is exactly its revenue.

Given a Bayesian instance $\hat{\mathcal{I}} = (N, M, \mathcal{D})$, for each player i and valuation subprofile $v_{-i} \sim \mathcal{D}_{-i}$, let $\beta_i(v_{-i}) = (\beta_{ij}(v_{-i}))_{j \in M}$ be such that $\beta_{ij}(v_{-i}) = \max_{i' \neq i} v_{i'j}$ for each item j . Note that we are slightly abusing notations here: each β_{ij} is now a function rather than a value. The optimal β -adjusted revenue for $\hat{\mathcal{I}}$ is

$$\begin{aligned} \mathbb{E}_{\mathcal{D}} Rev^A(\hat{\mathcal{I}}, \beta) &= \mathbb{E}_{v \sim \mathcal{D}} \sum_i Rev^A(\hat{\mathcal{I}}_i, \beta_i(v_{-i})) \\ &= \sum_i \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} \mathbb{E}_{v_i \sim \mathcal{D}_i} Rev^A(\hat{\mathcal{I}}_i, \beta_i(v_{-i})). \end{aligned}$$

When v_{-i} is clear from the context, we may simply write β_i and β_{ij} .

Because we also consider the projected Bayesian instance $\mathcal{I}' = (N, M, \mathcal{D}')$, let v' be v projected on the knowledge graph G : that is, $v'_{ij} = v_{ij}$ if there exists a player i' with $(i', i) \in G_j$, and $v'_{ij} = 0$ otherwise. As v is distributed according to \mathcal{D} , v' is distributed according to \mathcal{D}' . Thus we sometimes directly sample $v' \sim \mathcal{D}'$ rather than sample v first and then map it to v' . Let $\beta'_i(v'_{-i}) = \max_{i' \neq i} v'_{i'j}$ for each player i and item j . The optimal β' -adjusted revenue for \mathcal{I}' is defined respectively:

$$\begin{aligned} \mathbb{E}_{\mathcal{D}'} Rev^A(\mathcal{I}', \beta') &= \mathbb{E}_{v' \sim \mathcal{D}'} \sum_i Rev^A(\mathcal{I}'_i, \beta'_i(v'_{-i})) \\ &= \sum_i \mathbb{E}_{v'_{-i} \sim \mathcal{D}'_{-i}} \mathbb{E}_{v'_i \sim \mathcal{D}'_i} Rev^A(\mathcal{I}'_i, \beta'_i(v'_{-i})). \end{aligned}$$

It is important to emphasize that, given v_{-i} and $\beta_i(v_{-i})$, the optimal β_i -exclusive revenue and the optimal β_i -adjusted revenue are well defined for any Bayesian instance for player i , whether it is $\hat{\mathcal{I}}_i$ or \mathcal{I}'_i . In particular, we will consider $Rev^X(\mathcal{I}'_i, \beta_i)$ and $Rev^A(\mathcal{I}'_i, \beta_i)$, which

are on the hybrid of $\hat{\mathcal{I}}_i$ and \mathcal{I}'_i . The optimal β -adjusted revenue on the hybrid of $\hat{\mathcal{I}}$ and \mathcal{I}' are similarly defined:

$$\mathbb{E}_{\mathcal{D}} \text{Rev}^A(\mathcal{I}', \beta) = \sum_i \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} \mathbb{E}_{v'_i \sim \mathcal{D}'_i} \text{Rev}^A(\mathcal{I}'_i, \beta_i(v_{-i})).$$

To highlight that in the inner expectation, player i 's value is v'_i even if it is obtained by first sampling $v_i \sim \mathcal{D}_i$ and then projecting on G , we may also write it as $\mathbb{E}_{v_i \sim \mathcal{D}_i} \text{Rev}^A(\mathcal{I}'_i, \beta_i; v'_i)$ and

$$\begin{aligned} \mathbb{E}_{\mathcal{D}} \text{Rev}^A(\mathcal{I}', \beta) &= \sum_i \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} \mathbb{E}_{v_i \sim \mathcal{D}_i} \text{Rev}^A(\mathcal{I}'_i, \beta_i(v_{-i}); v'_i) \\ &= \mathbb{E}_{v \sim \mathcal{D}} \sum_i \text{Rev}^A(\mathcal{I}'_i, \beta_i(v_{-i}); v'_i). \end{aligned}$$

Finally, denote by IM the *individual Myerson* mechanism, which sells each item separately using Myerson's mechanism [1]. The revenue of IM under the projected Bayesian instance \mathcal{I}' , denoted by $IM(\mathcal{I}')$, is thus the optimal revenue by selling each item separately.

Having defined the notions and notations needed in our proof, we prove Theorem 6 via the following two technical lemmas, whose proofs can be found in Appendix C. For each lemma, note that on the left-hand side the values are drawn from the original Bayesian instance, and on the right-hand side the values are drawn from the projected instance.

Lemma 11 $\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEA}(\mathcal{I})) \geq \frac{1}{68} \mathbb{E}_{v' \sim \mathcal{D}'} \text{Rev}^A(\mathcal{I}', \beta').$

Lemma 12 $\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEA}(\mathcal{I})) \geq \frac{1}{2} \mathbb{E}_{v' \sim \mathcal{D}'} IM(\mathcal{I}').$

Proof of Theorem 6 By Theorem 8.1 of [11], $\mathbb{E}_{v' \sim \mathcal{D}'} \text{Rev}^A(\mathcal{I}', \beta') + \mathbb{E}_{v' \sim \mathcal{D}'} IM(\mathcal{I}') \geq OPT(\mathcal{I}')$. Combining this inequality with Lemmas 11 and 12, we have $\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEA}(\mathcal{I})) \geq \frac{OPT(\mathcal{I}')}{70} = \frac{OPT_K(\mathcal{I})}{70}$, and Theorem 6 holds. \square

6 Using scoring rules to buy knowledge from players

In this section we use proper scoring rules to reward the players for their knowledge, so that it is strictly better for them to report truthfully. More precisely, a *scoring rule* is a function f that takes as inputs a distribution \mathcal{D}' over a state space Ω and a random sample ω from an underlying true distribution \mathcal{D} over Ω , and outputs a real number. Scoring rule f is *proper* if

$$\mathbb{E}_{\omega \sim \mathcal{D}} f(\mathcal{D}, \omega) \geq \mathbb{E}_{\omega \sim \mathcal{D}} f(\mathcal{D}', \omega)$$

for any \mathcal{D} and \mathcal{D}' , and *strictly proper* if the inequality is strict for any $\mathcal{D}' \neq \mathcal{D}$. Moreover, f is *bounded* if there exist constants c_1, c_2 such that $c_1 \leq f(\mathcal{D}', \omega) \leq c_2$ for any \mathcal{D}' and ω . Our

mechanisms can use any strictly proper scoring rules that are bounded. For concreteness, we use Brier's scoring rule [12]:

$$BSR(\mathcal{D}', \omega) = 2 - \left(\sum_{s \in \Omega} (\delta_{\omega, s} - \mathcal{D}'(s))^2 \right) = 2\mathcal{D}'(\omega) - \|\mathcal{D}'\|_2^2 + 1,$$

where $\mathcal{D}'(s)$ is the probability of s according to \mathcal{D}' , and $\delta_{\omega, s}$ is the indicator for $\omega = s$. Note that $BSR(\mathcal{D}', \omega) \in [0, 2]$ for any \mathcal{D}' and ω .⁷ We remark that Brier's scoring rule can be replaced by any strictly proper and bounded scoring rules. However, scoring rules like the logarithmic scoring rule [18] are not suitable, as they are unbounded.

In all our mechanisms, when player i reports $\mathcal{D}_{i'j}^i \neq \perp$ for player i' and item j , the seller rewards i based on $BSR(\mathcal{D}_{i'j}^i, b_{i'j})$. If there are more than one reporters for the same distribution, the seller can either reward all of them or randomly choose one. We can scale the reward for each distribution so that the total reward given to the players is at most some constant ϵ , which is an ϵ additive loss to the revenue. For example, in Mechanism \mathcal{M}_{IEUD} , the seller could reward each player i with

$$r_{i'j}^i = \frac{\epsilon}{2mn^2} BSR(\mathcal{D}_{i'j}^i, b_{i'j})$$

for reporting the value distribution of player i' on item j .

Although scoring rules help breaking utility-ties, they cause another problem: a player who does not know a distribution may report something she made up, just to receive a reward. Therefore we start by considering our mechanisms under the *no-bluff* assumption: that is, a player will not report anything about a distribution that she does not know. More precisely, a player i in an information elicitation auction is *no-bluff* if, for any knowledge graph G_j and player i' with $(i, i') \notin G_j$, and for any strategy (b_i, K_i) of i , i reports \perp for the corresponding distribution of i' . Note that for a player i' with $(i, i') \in G_j$, i may report any distribution about i' , including \perp . In some sense, the no-bluff assumption is the analogy of the no-overbidding assumption adopted in budget-constrained auctions: a player will not bid higher than her true value or budget, even if doing so may not lead to a price higher than the latter.

For all our mechanisms, it is easy to see that the reward will only affect the players' incentives for reporting their knowledge, not their incentives for reporting their values. Accordingly, it is still dominant for the players to report their true values, no matter what knowledge they report. Given that the players all report their true values, the reported values are distributed according to the prior. Thus reporting her true knowledge is now strictly better than lying for a player, because it maximizes her reward. Rather than restating all our previous theorems, we summarize them in the theorem below.

Theorem 7 *Under the no-bluff assumption, for any information elicitation mechanism in previous sections, the revised mechanism with proper scoring rules is 2-DST, and reporting her true knowledge is strictly better than lying for each player i . Moreover, the mechanism's revenue is the same as before with an ϵ additive loss.*

⁷The original version of *BSR* is bounded by -1 and 1 , and we have shifted it up by 2 .

Next, we show how to remove the no-bluff assumption when everything is known. Without this assumption, player i may report a distribution for another player i' 's value for an item j , even if $(i, i') \notin G_j$. However, if there exists a third player \hat{i} who knows i' 's distribution $\mathcal{D}_{i'j}$, and if player i is also rewarded for player \hat{i} 's report, then intuitively player i would have no incentive to bluff about i' . That is, *not only a player is paid for reporting the distributions she knows, but she is also paid for keeping quiet about the distributions she does not know and letting the experts speak*. Surely reporting $\mathcal{D}_{i'j}$ maximizes player i 's expected reward, but she does not have the information to decide what $\mathcal{D}_{i'j}$ is.⁸ Therefore, as long as reporting “ \perp ” gives player i the same utility as the unknown strategy of reporting $\mathcal{D}_{i'j}$, and as long as reporting any distribution other than $\mathcal{D}_{i'j}$ gives her a strictly smaller utility, player i will report “ \perp ” about $\mathcal{D}_{i'j}$.

Taking mechanism $\mathcal{M}'_{I\text{EUD}}$ in Section 4 as an example, the players are rewarded as follows:

- For each player i' and item j , let $R_{i'j}$ be the set of players who did not report “ \perp ” about the value of i' for j . Randomly select a player \hat{i} from $R_{i'j}$ and let $r_{i'j} = \text{BSR}(\mathcal{D}_{\hat{i}j}, b_{i'j})$. Reward every player $i \neq i'$ using $r_{i'j}$, properly scaled.

Note that the reward $r_{i'j}$ is given to player i even if she has reported $\mathcal{D}_{i'j} = \perp$.

It is easy to see that, for any k -informed information elicitation instance with $k \geq 1$, if all players except i report their true values and true knowledge, then player i 's best strategy is to tell the truth about her own. Indeed, reporting her true values is still dominant no matter what knowledge the players' report. Moreover, for each player i' and item j , there are two ways for player i to maximize the reward $r_{i'j}$ she receives: (1) reporting $\mathcal{D}_{i'j} = \perp$, so that \hat{i} is chosen with probability 1 from the set of players who actually know $\mathcal{D}_{i'j}$; or (2) successfully guessing $\mathcal{D}_{i'j}$ and reporting it, so that \hat{i} 's report is still $\mathcal{D}_{i'j}$ with probability 1. Indeed, the resulting information elicitation mechanism is Bayesian incentive compatibility (BIC): that is, all players reporting their true values and true knowledge is a Bayesian Nash equilibrium. Note that option (2) is not a well-defined Bayesian strategy, because i does not have enough information to carry it out. In fact, option (1) is the only Bayesian Nash equilibrium in the mechanism, besides the unachievable ones where each player reports the unknown true distributions. We again summarize our results in the theorem below.

Theorem 8 *For any information elicitation mechanism in previous sections where the knowledge graph is at least 1-informed, the revised mechanism does not rely on the no-bluff assumption, and all players reporting their true values and true knowledge is the unique Bayesian Nash equilibrium. Moreover, the mechanism's revenue is the same as before with an ϵ additive loss.*

When not everything is known and the players may bluff, a player may not report “ \perp ” about a distribution she does not know, because she may still get some reward in case nobody knows that distribution. It is an interesting open problem to design information elicitation mechanisms when not everything is known and without the no-bluff assumption.

⁸Using the standard language from epistemic game theory, player i 's information set contains at least two different distributions for i' 's value for j .

7 Information elicitation mechanisms for combinatorial auctions

In our main results, the knowledge graphs for different items can be totally different from each other: player 1's value distribution for item 2 may be known by player 3, while her value distribution for item 4 may be known by player 5, etc. When all the knowledge graphs are the same, we say that the auction has *player-wise information*: all value distributions of a player have the same “knower”.

With player-wise information, only one knowledge graph G is needed: an edge (i, i') is in G if and only if player i knows the distribution $\mathcal{D}_{i'} = \mathcal{D}_{i'1} \times \cdots \times \mathcal{D}_{i'm}$. Such an information setting can also model arbitrary *combinatorial auctions*, where each player i 's valuation function v_i maps each subset of items to a non-negative real, with $v_i(\emptyset) = 0$. Thus the distribution \mathcal{D}_i is over such functions, and i 's values for items within a subset can be arbitrarily correlated. Given a combinatorial Bayesian auction instance $\hat{\mathcal{I}} = (N, M, \mathcal{D} = \times_{i \in N} \mathcal{D}_i)$, a corresponding information elicitation instance is denoted by $\mathcal{I} = (N, M, \mathcal{D}, G)$, where G is a single knowledge graph rather than a vector of graphs.

Player-wise information is a much stronger assumption than *(player, item)-wise information*, the information settings considered in Sections 3 to 4 of this paper. We mention this model here mainly for completeness and to facilitate the comparison of our results with the literature. Indeed, with player-wise information and by randomly partitioning the players into reporters and potential buyers as we have done in mechanisms \mathcal{M}_{IEUD} and \mathcal{M}'_{IEUD} , we have a simple blackbox reduction from Bayesian mechanisms to information elicitation mechanisms. Our mechanism \mathcal{M}_{IEB} is shown in Mechanism Algorithm 8, where \mathcal{M}_B can be any Bayesian mechanism. We have the following theorem, whose proof is similar to that of Theorem 2 and thus omitted.

Algorithm 8 \mathcal{M}_{IEB}

- 1: Each player i reports a valuation function b_i and a knowledge $K_i = (\mathcal{D}_{i'})_{i' \neq i}$.
 - 2: Randomly partition the players into two sets, N_1 and N_2 , where each player is independently put in N_1 with probability $q = 1 - (k+1)^{-\frac{1}{k}}$ and N_2 with probability $1 - q$.
 - 3: Let N_3 be the set of players in N_2 whose distributions are reported by some players in N_1 , and let \mathcal{D}'_{N_3} be the vector of reported distributions.
 - 4: Run \mathcal{M}_B on the Bayesian instance $\hat{\mathcal{I}}_{N_3} = (N_3, M, \mathcal{D}'_{N_3})$ and the valuation functions b_{N_3} ; and use the resulting allocation and prices to sell to players in N_3 .
-

Theorem 9 For any $k \in [n-1]$, for any combinatorial auction instances $\hat{\mathcal{I}} = (N, M, \mathcal{D})$ and $\mathcal{I} = (N, M, \mathcal{D}, G)$ where \mathcal{I} has player-wise information and G is k -informed, if \mathcal{M}_B is DST then \mathcal{M}_{IEB} is 2-DST; and if \mathcal{M}_B is BIC then \mathcal{M}_{IEB} is BIC. Moreover, if \mathcal{M}_B is a σ -approximation to OPT, then \mathcal{M}_{IEB} is a $\tau_k \sigma$ -approximation to OPT.

With player-wise information, a player i 's valuation distribution \mathcal{D}_i is either “completely known” or “completely unknown”. Thus for unit-demand auctions there is no need to adopt the COPIES setting to handle the scenario where only part of \mathcal{D}_i is reported. As before, the approximation ratio of \mathcal{M}_{IEB} increases as k gets larger and converges to that of the Bayesian mechanism. When $k = 1$, it is a 4-approximation to OPT using the optimal Bayesian mechanism as a black-box. Note that the model studied in [7] is a very special case even compared with the player-wise information setting: that is, G is the complete graph and $k = n - 1$. Since $\tau_{n-1} = \frac{n-1}{n^{n/(n-1)}} \rightarrow 1 - \frac{1}{n}$ when n gets larger, the revenue of our mecha-

nism essentially matches that of [7] for combinatorial auctions. Moreover, if the players can observe private signals and refine their knowledge about the prior, then our mechanism can also aggregate such refinements as in [7]. Finally, we have the following corollary for arbitrary knowledge graphs (i.e., k may be 0).

Corollary 2 *For any combinatorial auction instances $\hat{\mathcal{I}} = (N, M, \mathcal{D})$ and $\mathcal{I} = (N, M, \mathcal{D}, G)$ with player-wise information, mechanism \mathcal{M}_{IEB} with $q = \frac{1}{2}$ is a $\frac{\sigma}{4}$ -approximation to $OPT_K(\mathcal{I})$ if \mathcal{M}_B is a σ -approximation to OPT .*

8 Discussion and conclusion

In our main results, the seller asks the players to report the distributions in their entirety, without being concerned with the communication complexity for doing so. This is common in information elicitation and allows us to focus on the main difficulties in aggregating the players' knowledge, but maybe unrealistic in practice. In Appendix A, we consider the communication complexity from a theoretical point of view and show how to modify our mechanisms so that the players only report a small amount of information about the distributions. A systematic and more practical study of efficient communication is left for further research. Another interesting future direction is to incorporate the cost of reporting into the model and design incentives for the agents to report.

As Bayesian auctions require the seller (and the players under common-prior assumption) has correct knowledge about *all* distributions, in our main results we do not consider scenarios where players have “insider” knowledge. If the insider knowledge is correct (i.e., is a refinement of the prior), then our mechanisms' revenue increases; see Appendix B. Still, how to aggregate even the incorrect information that the players may have about each other is a very interesting question for future studies.

Another important direction is to elicit players' information for BIC mechanisms. For example, the BIC mechanisms in [31, 61] are optimal in their own settings, and it is unclear how to convert them to information elicitation mechanisms.

A Information elicitation mechanisms with efficient communication

To reduce the communication complexity of our mechanisms, rather than asking each player to report her known distributions in their entirety, the seller can make specific queries to the players about the distributions. Indeed, the *query complexity* of Bayesian auctions has been studied by [48], where the seller does not know the prior distributions but is given oracle accesses to them. More precisely, for any distribution D over reals, in a *value query* the seller sends a value v and the oracle returns the corresponding quantile $q(v) = \Pr_{x \sim D}[x \geq v]$. In a *quantile query*, the seller sends a quantile $q \in [0, 1]$ and the oracle returns the corresponding value $v(q)$ such that $\Pr_{x \sim D}[x \geq v(q)] = q$.

In information elicitation auctions, as the players have knowledge about the distributions, it is very natural for the seller to use them as oracles. However, it is important to ensure that the queries to the players do not destroy their incentives to be truthful: both to

report their true values and to report their true knowledge.⁹ Fortunately, truthfulness in our mechanisms can be easily guaranteed by *non-adaptive queries*, where all the queries are made together, before the players report their values. As shown by [48], when the players' value distributions are bounded within $[1, H]$ for a given value H , the number of non-adaptive queries sufficing to approximate OPT in Bayesian auctions is polynomial in m and n , but only logarithmic in H , which is very efficient. Moreover, only value queries are needed in this case. When the distributions have unbounded supports but satisfy small-tail assumptions, non-adaptive quantile queries are enough, and the query complexity is polynomial in m , n and logarithmic in the cut-off value of the tail. We make the same queries in our mechanisms as in [48].

Below we show how to revise our information elicitation mechanisms to query the players, using mechanism $\mathcal{M}_{I\text{EUD}}$ as an example and for bounded distributions. The mechanism now has a parameter $\epsilon > 0$, which affects its approximation ratio.

- In Step 1, given $\epsilon > 0$, let $k = \lceil \log_{1+\epsilon} H \rceil$ and $\nu = (\nu_0, \nu_1, \dots, \nu_{k-1}, \nu_k) = (1, (1+\epsilon), (1+\epsilon)^2, \dots, (1+\epsilon)^{k-1}, H)$. Each player i reports, for each player $i' \neq i$ and item j , either \perp or a non-increasing quantile vector $q_{i'j}^i = (q_{i'j;0}^i, \dots, q_{i'j;k}^i)$, where $q_{i'j;0}^i = 1$. Allegedly, $q_{i'j;l}^i = q_{i'j}(\nu_l)$ for each $l \in \{0, \dots, k\}$, where $q_{i'j}(\cdot)$ is defined by $\mathcal{D}_{i'j}$. That is, if $(i, i') \in G_j$ then player i answers the value queries for distribution $\mathcal{D}_{i'j}$ and value vector ν . Simultaneously, each player i also reports a valuation $b_i = (b_{ij})_{j \in M}$.
- In Step 5, if player $i \in N_1$ is the reporter for player $i' \in N_2$ and item j , then construct a discrete distribution $\mathcal{D}_{i'j}'$ as follows: $\mathcal{D}_{i'j}'(\nu_l) = q_{i'j;l}^i - q_{i'j;l+1}^i$ for every $l \in \{0, \dots, k\}$, where $q_{i'j;k+1}^i \triangleq 0$. The other parts of the mechanism remain unchanged.

In the revised mechanism, it is still dominant for the players to report their true values, no matter how the queries are answered. Indeed, the fact that distribution $\mathcal{D}_{i'j}'$ is now different from $\mathcal{D}_{i'j}$ does not affect the players' truthfulness in the Bayesian mechanism \mathcal{M}_{UD} . Moreover, having player i answer the value queries for $\mathcal{D}_{i'j}$ is equivalent to first having her report $\mathcal{D}_{i'j}'$ and then having the seller answer the value queries accordingly. In the latter, reporting $\mathcal{D}_{i'j}'$ truthfully never hurts player i , because $i \in N_1$ when her knowledge is used.

Thus answering the value queries truthfully never hurts i either, and the mechanism is still 2-DST. Because the value queries for a distribution $\mathcal{D}_{i'j}$ may be answered by all the other $n - 1$ players (when they all know $\mathcal{D}_{i'j}$), the query complexity and thus the communication complexity of our mechanisms have an extra factor n compared with the query complexity in [48].

More precisely, we state the following theorem for arbitrary knowledge graphs and bounded distributions. The proof is relatively easy following those for Section 5 and those in [48], thus has been omitted.

- Theorem 10 $\forall \epsilon > 0, H > 1$ and for any auction instances $\hat{\mathcal{I}} = (N, M, \mathcal{D})$ and $\mathcal{I} = (N, M, \mathcal{D}, G)$, where each \mathcal{D}_{ij} 's support is bounded in $[1, H]$, our revised information elicitation mechanisms are 2-DST and make non-adaptive value queries. Moreover, for single-good auctions, with $O(n^2 \log_{1+\epsilon} H)$ queries, the mechanism

⁹ Here a player reporting her true knowledge no longer means that she reports the true distributions, but that she answers the seller's queries truthfully.

achieves revenue at least $\frac{OPT_K(\mathcal{I})}{4(1+\epsilon)}$;

- for unit-demand auctions, with $O(mn^2 \log_{1+\epsilon} H)$ queries, the mechanism achieves revenue at least $\frac{OPT_K(\mathcal{I})}{96(1+\epsilon)}$; and
- for additive auctions, with $O(mn^2 \log_{1+\epsilon} H)$ queries, the mechanism achieves revenue at least $\frac{OPT_K(\mathcal{I})}{70(1+\epsilon)}$.

The case of unbounded distributions with small-tail assumptions, as well as the cases of k -informed knowledge graphs with $k \geq 1$ and bounded/unbounded distributions, are similar. Indeed, the approximation ratios of our main results and the query complexity of Bayesian auctions in [48] combine nicely with each other, resulting in information elicitation mechanisms with very efficient communication.

To further reduce the communication complexity of our mechanisms, the seller can change them into extensive-form mechanisms and ask each player i to first report a bit about each pair (i', j) , indicating whether i knows $\mathcal{D}_{i'j}$ or not. The seller then selects one reporter and only asks her to answer the oracle queries. By doing so, the extra factor n in the query complexity of our mechanisms can be dropped, with the players communicating $O(n^2 m)$ bits besides the queries.

If scoring rules are used to buy the players' knowledge, then we can use the following *value scoring rule* g_V to reward value queries, which follows directly from Brier's scoring rule [12]. More precisely, for any value query $v \in \mathbb{R}$, letting x be a sample from the underlying value distribution and q be the answer of a reporter to the query, then

$$g_V(x, q; v) \triangleq 1 + 2q\mathbf{I}_{x \leq v} - q^2.$$

To reward the players' answers to quantile queries, we define the following *quantile scoring rule* g_Q , which is a variant of the one in [63]. More precisely, for any quantile query $q \in [0, 1]$, letting x be a sample from the underlying value distribution and z be the answer of a reporter to the query, then

$$g_Q(x, z; q) \triangleq q \arctan z - (\arctan z - \arctan x)\mathbf{I}_{z \geq x}.$$

Both scoring rules are strictly proper scoring rules with bounded ranges.

B Aggregating the players' refined insider knowledge

As mentioned in the introduction of the paper, since the common prior assumption implies that every player has correct and exact (that is, no more, no less) knowledge about all the distributions, in Sections 3 to 5 of this paper we do not consider scenarios where the players have "insider" knowledge. Incorrect insider knowledge has been studied in [50, 51, 56, 64, 65] and is not the focus of this paper. However, sometimes each player may have *correct* insider knowledge about the other players' value distributions¹⁰: that is, her knowledge is a refinement of the prior.

Different players' knowledge, although all correct, may refine the prior in different ways. For example, when the prior distribution of a player i 's value for an item j is uniform over

¹⁰ A similar scenario in the literature of contracts was considered in [66] with different concerns.

$[0, 100]$, after v_{ij} is drawn, another player i' may observe whether $v_{ij} \geq 50$ or not, and a third player i'' may observe whether $v_{ij} \in [20, 80]$. Thus, player i' knows whether v_{ij} is uniform over $[0, 50]$ or $(50, 100]$, depending on his signal; and player i'' knows whether v_{ij} is uniform over $[20, 80]$ or $[0, 20) \cup (80, 100]$, depending on the signal i'' observes.

The players' correct knowledge must be consistent with each other and one can obtain an even better refinement of the prior by combining their knowledge together. In the example above, if player i' observes $v_{ij} < 50$ and player i'' observes $v_{ij} \notin [20, 80]$, then it must be that $v_{ij} \in [0, 20)$ and the posterior distribution is uniform in this range. However, neither i' nor i'' knows this fact.

Enhanced knowledge graphs *To model the players' insider knowledge, we equip the knowledge graphs in the information elicitation setting with information sets. To begin with, for any two players i, i' and item j such that $(i, i') \in G_j$, there is a partition $\mathcal{P}_{i',j}^i$ of the support of $\mathcal{D}_{i',j}$, representing the possible signals player i will observe about $\mathcal{D}_{i',j}$. After the true value $v_{i',j}$ is drawn, letting $S(v_{i',j})$ be the unique set in the partition that contains $v_{i',j}$, player i learns the fact that player i' 's true value for j falls into $S(v_{i',j})$ and the posterior distribution is $\mathcal{D}_{i',j} \mid S(v_{i',j})$. More generally, the partitions may even depend on player i 's own true valuation v_i , because v_i is part of the information she has. Because different values are independently drawn, given v_i and the information sets observed by i for different distributions of the other players, i considers their posterior distributions to be independent.*

All our mechanisms remain 2-DST with respect to the players' refined knowledge, where a player's true knowledge is now the posterior distributions known by him. Since the optimal Bayesian revenue increases when the distributions are refined [7], the expected revenue of our mechanisms also increases, where the expectation is further taken over the private signals observed by the players. However, the revenue benchmark is still defined as before: that is, with respect to the prior and the knowledge graphs, without considering the refinements. An interesting open problem is to design information elicitation mechanisms whose revenue approximates a more demanding benchmark — the optimal revenue based on the "aggregated refinement" obtained by combining all players' refinements together.

C Missing proofs for Section 5

C.1 Proof of Lemma 11

Proof of Lemma 11 For each player i , the set M_i^1 of items for which i 's distributions are reported and the set M_i^2 of items for which i 's distributions are not reported are uniquely determined by the knowledge graph G in \mathcal{I} . Let $\mathcal{I}'_{i,M_i^1} = (\{i\}, M_i^1, \mathcal{D}'_{i,M_i^1})$ be the single-player Bayesian instance obtained by projecting \mathcal{I} on i and M_i^1 , with $\mathcal{D}'_{i,M_i^1} = \times_{j \in M_i^1} \mathcal{D}'_{ij}$. We have

$$\begin{aligned}
& \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEA}(\mathcal{I})) \\
&= \mathbb{E}_{v \sim \mathcal{D}} \sum_i \left(\text{Rev} \left(\text{Bund}_i(\mathcal{D}'_i, v_i, M_i^1, \beta_i(v_{-i})) \right) + \sum_{j \in M_i^2 \cap M_i} \max_{i' \neq i} v_{i'j} \right) \\
&= \mathbb{E}_{v \sim \mathcal{D}} \sum_i \left(\text{Rev} \left(\text{Bund}_i(\mathcal{D}'_{i, M_i^1}, v_i, M_i^1, \beta_{i, M_i^1}) \right) + \sum_{j \in M_i^2 \cap M_i} \max_{i' \neq i} v_{i'j} \right), \tag{13}
\end{aligned}$$

where Bund_i is Bund applied to player i in Steps 7 and 8. The first equality is by the construction of \mathcal{M}_{IEA} . The second equality holds because in the execution of Bund_i the items in M_i^2 do not affect anything, as player i 's values for them are constantly 0 in \mathcal{D}'_i .

By Theorem 6.1 of [11], for any single-player Bayesian instance for a player i and any β_i , Bund_i is

an 8.5-approximation to the optimal β_i -exclusive revenue for this instance. Moreover, by Theorem 1 of [11], the optimal β_i -exclusive revenue is an 8-approximation to the optimal β_i -adjusted revenue for the same instance. Accordingly, we have

$$\begin{aligned}
& \mathbb{E}_{v \sim \mathcal{D}} \sum_i \text{Rev}(\text{Bund}_i(\mathcal{D}'_{i, M_i^1}, v_i, M_i^1, \beta_{i, M_i^1})) \\
&= \sum_i \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} \mathbb{E}_{v_i \sim \mathcal{D}_i} \text{Rev}(\text{Bund}_i(\mathcal{D}'_{i, M_i^1}, v_i, M_i^1, \beta_{i, M_i^1})) \\
&= \sum_i \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} \mathbb{E}_{v'_i \sim \mathcal{D}'_i} \text{Rev}(\text{Bund}_i(\mathcal{D}'_{i, M_i^1}, v'_i, M_i^1, \beta_{i, M_i^1})) \\
&= \sum_i \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} \mathbb{E}_{v'_i \sim \mathcal{D}'_i} \text{Rev}(\text{Bund}_i(\mathcal{I}'_{i, M_i^1}, \beta_{i, M_i^1})) \\
&\geq \frac{1}{8.5} \sum_i \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} \mathbb{E}_{v'_i \sim \mathcal{D}'_i} \text{Rev}^X(\mathcal{I}'_{i, M_i^1}, \beta_{i, M_i^1}) \\
&= \frac{1}{8.5} \sum_i \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} \mathbb{E}_{v'_i \sim \mathcal{D}'_i} \text{Rev}^X(\mathcal{I}'_i, \beta_i) \\
&\geq \frac{1}{8.5 \times 8} \sum_i \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} \mathbb{E}_{v'_i \sim \mathcal{D}'_i} \text{Rev}^A(\mathcal{I}'_i, \beta_i) \\
&= \frac{1}{68} \sum_i \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} \mathbb{E}_{v_i \sim \mathcal{D}_i} \text{Rev}^A(\mathcal{I}'_i, \beta_i; v'_i) \\
&= \frac{1}{68} \mathbb{E}_{v \sim \mathcal{D}} \sum_i \text{Rev}^A(\mathcal{I}'_i, \beta_i; v'_i) = \frac{1}{68} \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}^A(\mathcal{I}', \beta). \tag{14}
\end{aligned}$$

The second equality above is because \mathcal{D}'_i and \mathcal{D}_i are the same when only M_i^1 is concerned, and $v'_{ij} = v_{ij}$ for all $j \in M_i^1$. The first inequality is by the relation between Bund_i and Rev^X . The next equality is because i 's values for items in M_i^2 are always 0 under \mathcal{I}'_i , thus even when these items are available, a β -exclusive mechanism does not sell them to i anyway. The second inequality is by the relation between Rev^X and the Rev^A . The next equal-

ity is because, as mentioned before, when sampling v_i from \mathcal{D}_i and then projecting it on G , the resulting v'_i is distributed according to \mathcal{D}'_i .

Now we state the key lemma in our analysis, which connects the hybrid adjusted revenue with that for \mathcal{I}' . Recall that each β'_i is defined based on $v'_{-i} \sim \mathcal{D}'_{-i}$.

$$\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}^A(\mathcal{I}', \beta) + \mathbb{E}_{v \sim \mathcal{D}} \sum_i \sum_{j \in M_i^2 \cap M_i} \max_{i' \neq i} v_{i'j} \geq \mathbb{E}_{v' \sim \mathcal{D}'} \text{Rev}^A(\mathcal{I}', \beta').$$

Lemma 13

Before proving Lemma 13, note that combining it with (13) and (14) we have

$$\mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEA}(\mathcal{I})) \geq \frac{1}{68} \mathbb{E}_{v' \sim \mathcal{D}'} \text{Rev}^A(\mathcal{I}', \beta'). \quad (15)$$

Thus Lemma 11 holds. \square

Next we prove Lemma 13.

Proof of Lemma 13 Arbitrarily fixing a player i and $v'_{-i} \sim \mathcal{D}'_{-i}$, let \mathcal{M}^* be the single-player DST Bayesian mechanism with the optimal β'_i -adjusted revenue for \mathcal{I}'_i .¹¹ Accordingly, \mathcal{M}^* maximizes the following quantity:

$$\mathbb{E}_{v'_{-i} \sim \mathcal{D}'_{-i}} \mathbb{E}_{v'_i \sim \mathcal{D}'_i} [\text{Rev}(\mathcal{M}^*(\mathcal{I}'_i)) - \sum_{j: v'_{ij} \leq \beta'_{ij}} q_{ij} v'_{ij}],$$

where q_{ij} is the probability for i to get item j in \mathcal{M}^* under v'_i . By definition,

$$\mathbb{E}_{v' \sim \mathcal{D}'} \text{Rev}^A(\mathcal{I}'_i, \beta'_i) = \mathbb{E}_{v' \sim \mathcal{D}'} [\text{Rev}(\mathcal{M}^*(\mathcal{I}'_i)) - \sum_{j: v'_{ij} \leq \beta'_{ij}} q_{ij} v'_{ij}]. \quad (16)$$

Again because sampling v from \mathcal{D} and projecting it on G induces the same distribution for v' as \mathcal{D}' , we have

$$\mathbb{E}_{v' \sim \mathcal{D}'} [\text{Rev}(\mathcal{M}^*(\mathcal{I}'_i)) - \sum_{j: v'_{ij} \leq \beta'_{ij}} q_{ij} v'_{ij}] = \mathbb{E}_{v \sim \mathcal{D}} [\text{Rev}(\mathcal{M}^*(\mathcal{I}'_i; v'_i)) - \sum_{j: v'_{ij} \leq \beta'_{ij}} q_{ij} v'_{ij}], \quad (17)$$

where we explicitly include v'_i in the input of \mathcal{M}^* to emphasize the projection.

Arbitrarily fix v_{-i} and let $\mathbb{U}_i = \{j \mid j \in M_i^1, \beta'_{ij} \neq \beta_{ij}\}$. It is clear that $\beta'_{ij} < \beta_{ij}$ for each $j \in \mathbb{U}_i$, as $v'_{-i} \leq v_{-i}$ (component-wise). For any v_i and the corresponding v'_i , as $v'_{ij} = 0$ for any $j \notin M_i^1$, we have

¹¹ If the supremum is not achieved by any mechanism, one can take a sequence of mechanisms whose β'_i -adjusted revenue approaches the supremum in the limit.

$$\begin{aligned}
& \sum_{j: v'_{ij} \leq \beta'_{ij}} q_{ij} v'_{ij} = \sum_{j \in M_i^1: v'_{ij} \leq \beta'_{ij}} q_{ij} v'_{ij} = \sum_{j \in M_i^1 \setminus \mathbb{U}_i: v'_{ij} \leq \beta'_{ij}} q_{ij} v'_{ij} + \sum_{j \in \mathbb{U}_i: v'_{ij} \leq \beta'_{ij}} q_{ij} v'_{ij} \\
&= \sum_{j \in M_i^1 \setminus \mathbb{U}_i: v'_{ij} \leq \beta_{ij}} q_{ij} v'_{ij} + \sum_{j \in \mathbb{U}_i: v'_{ij} \leq \beta_{ij}} q_{ij} v'_{ij} - \sum_{j \in \mathbb{U}_i: \beta'_{ij} < v'_{ij} \leq \beta_{ij}} q_{ij} v'_{ij} \\
&= \sum_{j \in M_i^1: v'_{ij} \leq \beta_{ij}} q_{ij} v'_{ij} - \sum_{j \in \mathbb{U}_i: \beta'_{ij} < v'_{ij} \leq \beta_{ij}} q_{ij} v'_{ij} \\
&= \sum_{j: v'_{ij} \leq \beta_{ij}} q_{ij} v'_{ij} - \sum_{j \in \mathbb{U}_i: \beta'_{ij} < v'_{ij} \leq \beta_{ij}} q_{ij} v'_{ij}.
\end{aligned} \tag{18}$$

Combining Equations (16), (17), (18) and taking summation over all players, we have

$$\begin{aligned}
& \mathbb{E}_{v' \sim \mathcal{D}'} \text{Rev}^A(\mathcal{I}', \beta') = \sum_i \mathbb{E}_{v' \sim \mathcal{D}'} \text{Rev}^A(\mathcal{I}'_i, \beta'_i) \\
&= \sum_i \mathbb{E}_{v \sim \mathcal{D}} [\text{Rev}(\mathcal{M}^*(\mathcal{I}'_i; v'_i)) - \sum_{j: v'_{ij} \leq \beta'_{ij}} q_{ij} v'_{ij}] \\
&= \sum_i \mathbb{E}_{v \sim \mathcal{D}} \left(\text{Rev}(\mathcal{M}^*(\mathcal{I}'_i; v'_i)) - \sum_{j: v'_{ij} \leq \beta_{ij}} q_{ij} v'_{ij} + \sum_{j \in \mathbb{U}_i: \beta'_{ij} < v'_{ij} \leq \beta_{ij}} q_{ij} v'_{ij} \right) \\
&\leq \sum_i \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}^A(\mathcal{I}'_i, \beta_i; v'_i) + \sum_i \mathbb{E}_{v \sim \mathcal{D}} \sum_{j \in \mathbb{U}_i: \beta'_{ij} < v'_{ij} \leq \beta_{ij}} q_{ij} v'_{ij} \\
&= \sum_i \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}^A(\mathcal{I}'_i, \beta_i; v'_i) + \mathbb{E}_{v \sim \mathcal{D}} \sum_i \sum_{j \in \mathbb{U}_i: \beta'_{ij} < v_{ij} \leq \beta_{ij}} q_{ij} v_{ij} \\
&\leq \sum_i \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}^A(\mathcal{I}'_i, \beta_i; v'_i) + \mathbb{E}_{v \sim \mathcal{D}} \sum_i \sum_{j \in \mathbb{U}_i: \beta'_{ij} < v_{ij} \leq \beta_{ij}} v_{ij} \\
&= \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}^A(\mathcal{I}', \beta) + \mathbb{E}_{v \sim \mathcal{D}} \sum_i \sum_{j \in \mathbb{U}_i: \beta'_{ij} < v_{ij} \leq \beta_{ij}} v_{ij}.
\end{aligned} \tag{19}$$

The first inequality above is because the first two terms in the expectation is exactly the β -adjusted revenue of mechanism \mathcal{M}^* on \mathcal{I}'_i . The following equality holds because $v'_{ij} = v_{ij}$ for any $j \in M_i^1$. Finally, the second inequality is because $0 \leq q_{ij} \leq 1$ for any i, j .

Arbitrarily fix a valuation profile v and consider the term $\sum_i \sum_{j \in \mathbb{U}_i: \beta'_{ij} < v_{ij} \leq \beta_{ij}} v_{ij}$ at the

end of (19). We show that each item j appears in the summation at most once. Indeed, for each v_{ij} that appears in the summation, $j \in M_i^1$ and \mathcal{D}_{ij} is known in \mathcal{I} . As $v_{ij} > \beta'_{ij}$, player i has strictly higher value for j than any other player i' whose distribution for j is also known. For any such player i' , $\beta'_{i'j} = v_{ij} > v_{i'j}$ and j is not in the set $\{j \in \mathbb{U}_{i'} : \beta'_{i'j} < v_{i'j} \leq \beta_{i'j}\}$, which implies that $v_{i'j}$ does not appear in the summation. Moreover, for any player $i'' \neq i$ with $\mathcal{D}_{i''j}$ unknown, $j \notin M_{i''}^1$ and $v_{i''j}$ does not appear in the summation either. Therefore item j only appears once in the summation. Accordingly,

$$\sum_i \sum_{j \in \mathbb{U}_i: \beta'_{ij} < v_{ij} \leq \beta_{ij}} v_{ij} = \sum_{j: \exists! i \text{ s.t. } j \in \mathbb{U}_i, \beta'_{ij} < v_{ij} \leq \beta_{ij}} v_{ij}. \quad (20)$$

We now show that

$$\sum_{j: \exists! i \text{ s.t. } j \in \mathbb{U}_i, \beta'_{ij} < v_{ij} \leq \beta_{ij}} v_{ij} \leq \sum_i \sum_{j \in M_i^2 \cap M_i} \max_{i' \neq i} v_{i'j}, \quad (21)$$

where the right-hand side is exactly the revenue generated by mechanism \mathcal{M}_{IEA} in Step 9 and also the desired term in the statement of Lemma 13. Indeed, for each v_{ij} that appears in the left-hand side, because $v_{ij} \leq \beta_{ij}$, player i 's value for j is at most the second highest among all players. Letting $i^* = \operatorname{argmax}_{i' \neq i} v_{i'j}$ with ties broken lexicographically, we have $\beta_{ij} = v_{i^*j}$ and v_{i^*j} is the highest value for j among all players. Since $\beta'_{ij} < \beta_{ij}$, \mathcal{D}_{i^*j} is unknown and $j \in M_{i^*}^2$. Below we show $j \in M_{i^*}$, which then implies $j \in M_{i^*}^2 \cap M_{i^*}$ and the price paid by i^* in Step 9 for item j is $\max_{i' \neq i^*} v_{i'j} \geq v_{ij}$.

When the distributions are generic, there are no ties in the players' values and v_{i^*j} is the unique maximum value for j , thus $j \in M_{i^*}$. For arbitrary distributions, problems occur when $v_{ij} = \beta_{ij}$ and $i < i^*$, which implies $j \notin M_{i^*}$. To deal with this special case, consider the following tie-breaking method for the players: while the value β_{ij} is still defined to be v_{i^*j} , we denote it by β_{ij}^+ if $i^* < i$ and β_{ij}^- if $i^* > i$. When $v_{ij} = \beta_{ij}$, we treat v_{ij} as strictly smaller if facing β_{ij}^+ and strictly larger if facing β_{ij}^- . All results proved in [11] and above continue to hold with respect to this tie-breaking method. Now for any v_{ij} that appears in the summation, either we have $v_{ij} < \beta_{ij}$ or we have $v_{ij} = \beta_{ij}$ and $i^* < i$, thus it is always the case that $j \in M_{i^*}$.

Accordingly,

$$\begin{aligned} \sum_{j: \exists! i \text{ s.t. } j \in \mathbb{U}_i, \beta'_{ij} < v_{ij} \leq \beta_{ij}} v_{ij} &\leq \sum_{j: \exists! i^* \text{ s.t. } j \in M_{i^*}^2 \cap M_{i^*}} \max_{i' \neq i^*} v_{i'j} \\ &= \sum_i \sum_{j \in M_i^2 \cap M_i} \max_{i' \neq i} v_{i'j}, \end{aligned}$$

and (21) holds. Combining (19), (20) and (21), we have

$$\mathbb{E}_{v' \sim \mathcal{D}'} \operatorname{Rev}^A(\mathcal{I}', \beta') \leq \mathbb{E}_{v \sim \mathcal{D}} \operatorname{Rev}^A(\mathcal{I}', \beta) + \mathbb{E}_{v \sim \mathcal{D}} \sum_i \sum_{j \in M_i^2 \cap M_i} \max_{i' \neq i} v_{i'j},$$

and Lemma 13 holds. \square

C.2 Proof of Lemma 12

Proof of Lemma 12 Denote by \mathcal{M}_{1LA} the individual 1-lookahead mechanism, which sells each item separately using the 1-lookahead mechanism [60] and is a 2-approximation to Myerson's mechanism for each item. Similar to the proof of Lemma 11, the expected revenue generated by Bund_i in Steps 7 and 8 is

$$\begin{aligned}
& \mathbb{E}_{v \sim \mathcal{D}} \sum_i \text{Rev}(\text{Bund}_i(\mathcal{D}'_i, v_{i, M_i^1}, \beta_i(v_{-i}))) \\
&= \sum_i \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} \mathbb{E}_{v_i \sim \mathcal{D}_i} \text{Rev}(\text{Bund}_i(\mathcal{D}_{i, M_i^1}, v_{i, M_i^1}, \beta_{i, M_i^1})) \\
&\geq \sum_i \mathbb{E}_{v_{-i} \sim \mathcal{D}_{-i}} \mathbb{E}_{v_i \sim \mathcal{D}_i} \text{Rev}(\mathcal{M}_{1LA}(\mathcal{D}_{i, M_i^1}, v_{i, M_i^1}, \beta_{i, M_i^1})) \\
&= \sum_i \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{1LA}(\hat{\mathcal{I}}_{i, M_i^1}, \beta_{i, M_i^1})).
\end{aligned} \tag{22}$$

The first equality above is again because items in M_i^2 does not affect Bund_i given that player i 's values for them are constantly 0 according to \mathcal{D}'_i , and \mathcal{D}'_i and \mathcal{D}_i coincide when only items in M_i^1 are concerned. The inequality holds because given \mathcal{D}_{i, M_i^1} and β_{i, M_i^1} , mechanism Bund_i chooses between optimally selling to i the items as a bundle (i.e., $e_i > 0$) and optimally selling to i each item separately (i.e., $e_i = 0$), whichever generates higher expected revenue over $v_{i, M_i^1} \sim \mathcal{D}_{i, M_i^1}$; while \mathcal{M}_{1LA} is a particular mechanism that sells each item $j \in M_i^1$ to i separately based on \mathcal{D}_{ij} and β_{ij} . Accordingly, letting $\mathcal{D}_j = \times (\mathcal{D}_{ij})_{i \in N}$ and $v_j \sim \mathcal{D}_j$ for each item j , we have

$$\begin{aligned}
& \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEA}(\mathcal{I})) \\
&= \mathbb{E}_{v \sim \mathcal{D}} \sum_i \left(\text{Rev}(\text{Bund}_i(\mathcal{D}'_i, v_{i, M_i^1}, \beta_i(v_{-i}))) + \sum_{j \in M_i^2 \cap M_i} \max_{i' \neq i} v_{i'j} \right) \\
&\geq \sum_i \mathbb{E}_{v \sim \mathcal{D}} \left(\text{Rev}(\mathcal{M}_{1LA}(\hat{\mathcal{I}}_{i, M_i^1}, \beta_{i, M_i^1})) + \sum_{j \in M_i^2 \cap M_i} \max_{i' \neq i} v_{i'j} \right) \\
&= \sum_{j \in M} \sum_i \mathbb{E}_{v_j \sim \mathcal{D}_j} \left(\mathbf{I}_{j \in M_i^1} \cdot \text{Rev}(\mathcal{M}_{1LA}(\hat{\mathcal{I}}_{i, j}, \beta_{ij})) + \mathbf{I}_{j \in M_i^2 \cap M_i} \cdot \max_{i' \neq i} v_{i'j} \right).
\end{aligned} \tag{23}$$

Arbitrarily fixing an item j , we show that

$$\begin{aligned}
& \sum_i \mathbb{E}_{v_j \sim \mathcal{D}_j} \left(\mathbf{I}_{j \in M_i^1} \cdot \text{Rev}(\mathcal{M}_{1LA}(\hat{\mathcal{I}}_{i, j}, \beta_{ij})) + \mathbf{I}_{j \in M_i^2 \cap M_i} \cdot \max_{i' \neq i} v_{i'j} \right) \\
&\geq \sum_i \mathbb{E}_{v'_j \sim \mathcal{D}'_j} \mathbf{I}_{j \in M_i^1} \cdot \text{Rev}(\mathcal{M}_{1LA}(\mathcal{I}'_{i, j}, \beta'_{ij})).
\end{aligned} \tag{24}$$

To do so, by definition, if $j \in M_i^1$ for a player i then in the Bayesian instance $\mathcal{I}'_{i, j}$ and given the reserve price β'_{ij} , the 1-lookahead mechanism tries to sell j to i at price¹²

¹² Strictly speaking, under the Bayesian instance \mathcal{I}' , the individual 1-lookahead mechanism for item j works as follows. Given $v'_j \sim \mathcal{D}_j$, it finds the highest bidder i for j with ties broken lexicographically, as well as the second highest bid which is exactly β'_{ij} . It then tries to sell j to i at price r'_{ij} . For generic distributions, there are no ties in v'_j and the mechanism can be run on each instance \mathcal{I}'_{ij} and β'_{ij} separately, without ever selling j to more than one players. For arbitrary distributions, this can be done by introducing proper tie-breaking rules.

$$r'_{ij} = \operatorname{argmax}_x \Pr_{v'_{ij} \sim \mathcal{D}'_{ij}} (v'_{ij} \geq x \mid v'_{ij} \geq \beta'_{ij}).$$

Thus

$$\begin{aligned} & \sum_i \mathbb{E}_{v'_j \sim \mathcal{D}'_j} \mathbb{I}_{j \in M_i^1} \cdot \operatorname{Rev}(\mathcal{M}_{1LA}(\mathcal{I}'_{i,j}, \beta'_{ij})) \\ &= \sum_i \mathbb{E}_{v'_{-i,j} \sim \mathcal{D}'_{-i,j}} \mathbb{E}_{v'_{ij} \sim \mathcal{D}'_{ij}} \mathbb{I}_{j \in M_i^1} \cdot \mathbb{I}_{v'_{ij} \geq r'_{ij}} \cdot r'_{ij} \\ &= \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \mathbb{I}_{j \in M_i^1} \cdot \mathbb{I}_{v'_{ij} \geq r'_{ij}(v'_{-i,j})} \cdot r'_{ij}(v'_{-i,j}), \end{aligned} \quad (25)$$

where the second equality is again by sampling v_j from \mathcal{D}_j and projecting on G to get v'_j . We write r'_{ij} as $r'_{ij}(v'_{-i,j})$ to highlight this fact. Note that given $v'_{-i,j}$, $r'_{ij}(v'_{-i,j})$ depends on the distribution \mathcal{D}'_{ij} but not on any concrete v'_{ij} . Also note that $r'_{ij}(v'_{-i,j}) \geq \beta'_{ij}$. Next, we divide the last term in (25) into two parts, depending on whether $r'_{ij}(v'_{-i,j}) < \beta_{ij}$ or not. That is,

$$\begin{aligned} & \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \mathbb{I}_{j \in M_i^1} \cdot \mathbb{I}_{v'_{ij} \geq r'_{ij}(v'_{-i,j})} \cdot r'_{ij}(v'_{-i,j}) \\ &= \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \mathbb{I}_{j \in M_i^1} \cdot \mathbb{I}_{v'_{ij} \geq r'_{ij}(v'_{-i,j})} \cdot \\ & \quad \left(\mathbb{I}_{r'_{ij}(v'_{-i,j}) < \beta_{ij}} + \mathbb{I}_{r'_{ij}(v'_{-i,j}) \geq \beta_{ij}} \right) \cdot r'_{ij}(v'_{-i,j}). \end{aligned} \quad (26)$$

For the $r'_{ij}(v'_{-i,j}) < \beta_{ij}$ part in (26), we have

$$\begin{aligned} & \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \mathbb{I}_{j \in M_i^1} \cdot \mathbb{I}_{v'_{ij} \geq r'_{ij}(v'_{-i,j})} \cdot \mathbb{I}_{r'_{ij}(v'_{-i,j}) < \beta_{ij}} \cdot r'_{ij}(v'_{-i,j}) \\ &= \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \mathbb{I}_{j \in M_i^1} \cdot \left(\mathbb{I}_{v'_{ij} \geq \beta_{ij} > r'_{ij}(v'_{-i,j})} + \mathbb{I}_{\beta_{ij} > v'_{ij} \geq r'_{ij}(v'_{-i,j})} \right) \cdot r'_{ij}(v'_{-i,j}) \\ &= \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \mathbb{I}_{j \in M_i^1} \cdot \left(\mathbb{I}_{v_{ij} \geq \beta_{ij} > r'_{ij}(v'_{-i,j})} + \mathbb{I}_{\beta_{ij} > v_{ij} \geq r'_{ij}(v'_{-i,j})} \right) \cdot r'_{ij}(v'_{-i,j}) \\ &\leq \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \mathbb{I}_{j \in M_i^1} \cdot \left(\mathbb{I}_{v_{ij} \geq \beta_{ij} > r'_{ij}(v'_{-i,j})} \cdot \beta_{ij} + \mathbb{I}_{\beta_{ij} > v_{ij} \geq r'_{ij}(v'_{-i,j})} \cdot v_{ij} \right) \\ &= \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \mathbb{I}_{j \in M_i^1} \cdot \mathbb{I}_{\beta_{ij} > r'_{ij}(v'_{-i,j})} \cdot \mathbb{I}_{v_{ij} \geq \beta_{ij}} \cdot \beta_{ij} \\ & \quad + \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \mathbb{I}_{j \in M_i^1} \cdot \mathbb{I}_{\beta_{ij} > v_{ij} \geq r'_{ij}(v'_{-i,j})} \cdot v_{ij}. \end{aligned} \quad (27)$$

The first equality above is by distinguishing whether $\beta_{ij} \leq v'_{ij}$ or not. The second equality is because $v_{ij} = v'_{ij}$ whenever $j \in M_i^1$. The inequality is because $r'_{ij}(v'_{-i,j}) < \beta_{ij}$ following the indicator in the first term and $r'_{ij}(v'_{-i,j}) \leq v_{ij}$ following the indicator in the second

term. Finally, the last equality is because $I_{v_{ij} \geq \beta_{ij} > r'_{ij}(v'_{-i,j})} = I_{\beta_{ij} > r'_{ij}(v'_{-i,j})} \cdot I_{v_{ij} \geq \beta_{ij}}$ in the first term.

For the first term in (27), because the indicators $I_{j \in M_i^1}$ and $I_{\beta_{ij} > r'_{ij}(v'_{-i,j})}$ does not depend on v_{ij} , we have

$$\begin{aligned} & \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} I_{j \in M_i^1} \cdot I_{\beta_{ij} > r'_{ij}(v'_{-i,j})} \cdot I_{v_{ij} \geq \beta_{ij}} \cdot \beta_{ij} \\ &= \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} I_{j \in M_i^1} \cdot I_{\beta_{ij} > r'_{ij}(v'_{-i,j})} \left(\mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} I_{v_{ij} \geq \beta_{ij}} \cdot \beta_{ij} \right) \\ &\leq \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} I_{j \in M_i^1} \cdot I_{\beta_{ij} > r'_{ij}(v'_{-i,j})} \left(\mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \text{Rev}(\mathcal{M}_{1LA}(\hat{\mathcal{I}}_{i,j}, \beta_{ij})) \right), \end{aligned}$$

where the inequality is because on the Bayesian instance $\hat{\mathcal{I}}_{i,j}$ and given the reserve price β_{ij} , mechanism \mathcal{M}_{1LA} chooses the optimal price r_{ij} to maximize the expected revenue and is no worse than simply setting $r_{ij} = \beta_{ij}$.

For the second term in (27), for any valuation profile v , if there exists a player i is such that the two indicators both equal to 1, then we have $\beta_{ij} > v_{ij} \geq \beta'_{ij}$, because $\beta_{ij} > v_{ij} \geq r'_{ij}(v'_{-i,j})$ and $r'_{ij}(v'_{-i,j}) \geq \beta'_{ij}$. Thus the (lexicographically first) highest bidder i^* for j in v_j has her distribution unknown, and $j \in M_{i^*}^2 \cap M_{i^*}$; and player i is the (lexicographically first) highest bidder for j in v'_j , which is unique. Accordingly,

$$\begin{aligned} & \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} I_{j \in M_i^1} \cdot I_{\beta_{ij} > v_{ij} \geq r'_{ij}(v'_{-i,j})} \cdot v_{ij} \\ &= \mathbb{E}_{v_j \sim \mathcal{D}_j} \sum_i I_{j \in M_i^1} \cdot I_{\beta_{ij} > v_{ij} \geq r'_{ij}(v'_{-i,j})} \cdot v_{ij} \\ &\leq \mathbb{E}_{v_j \sim \mathcal{D}_j} I_{j \in M_{i^*}^2 \cap M_{i^*}} \cdot \max_{i: j \in M_i^1} v_{ij} \leq \mathbb{E}_{v_j \sim \mathcal{D}_j} I_{j \in M_{i^*}^2 \cap M_{i^*}} \cdot \max_{i' \neq i^*} v_{i'j} \\ &= \mathbb{E}_{v_j \sim \mathcal{D}_j} \sum_i I_{j \in M_i^2 \cap M_i} \cdot \max_{i' \neq i} v_{i'j}, \end{aligned}$$

where the last equality is because i^* is the unique player such that the indicator is 1.

Combining the above two equations with (27), for the $r'_{ij}(v'_{-i,j}) < \beta_{ij}$ part in (26) we have

$$\begin{aligned} & \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} I_{j \in M_i^1} \cdot I_{v'_{ij} \geq r'_{ij}(v'_{-i,j})} \cdot I_{r'_{ij}(v'_{-i,j}) < \beta_{ij}} \cdot r'_{ij}(v'_{-i,j}) \\ &\leq \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} I_{j \in M_i^1} \cdot I_{\beta_{ij} > r'_{ij}(v'_{-i,j})} \left(\mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \text{Rev}(\mathcal{M}_{1LA}(\hat{\mathcal{I}}_{i,j}, \beta_{ij})) \right) \\ &\quad + \mathbb{E}_{v_j \sim \mathcal{D}_j} \sum_i I_{j \in M_i^2 \cap M_i} \cdot \max_{i' \neq i} v_{i'j} \\ &= \sum_i \mathbb{E}_{v_j \sim \mathcal{D}_j} \left(I_{j \in M_i^1} \cdot I_{\beta_{ij} > r'_{ij}(v'_{-i,j})} \cdot \text{Rev}(\mathcal{M}_{1LA}(\hat{\mathcal{I}}_{i,j}, \beta_{ij})) + I_{j \in M_i^2 \cap M_i} \cdot \max_{i' \neq i} v_{i'j} \right). \end{aligned}$$

For the $r'_{ij}(v'_{-i,j}) \geq \beta_{ij}$ part in (26), we immediately have $r'_{ij}(v'_{-i,j}) = r_{ij}(v_{-i,j})$ when the indicators are 1, because $\mathcal{D}'_{ij} = \mathcal{D}_{ij}$ and $v'_{ij} = v_{ij}$ whenever $j \in M_i^1$, and $r_{ij}(v_{-i,j})$ maximizes the expected revenue over \mathcal{D}_{ij} conditional on $v_{ij} \geq \beta_{ij}$. Thus

$$\begin{aligned} & \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \mathbb{I}_{j \in M_i^1} \cdot \mathbb{I}_{v'_{ij} \geq r'_{ij}(v'_{-i,j})} \cdot \mathbb{I}_{r'_{ij}(v'_{-i,j}) \geq \beta_{ij}} \cdot r'_{ij}(v'_{-i,j}) \\ &= \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \mathbb{I}_{j \in M_i^1} \cdot \mathbb{I}_{v_{ij} \geq r_{ij}(v_{-i,j})} \cdot \mathbb{I}_{r_{ij}(v_{-i,j}) \geq \beta_{ij}} \cdot r_{ij}(v_{-i,j}) \\ &= \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{I}_{j \in M_i^1} \cdot \mathbb{I}_{r'_{ij}(v'_{-i,j}) \geq \beta_{ij}} \left(\mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \mathbb{I}_{v_{ij} \geq r_{ij}(v_{-i,j})} \cdot r_{ij}(v_{-i,j}) \right) \\ &= \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{I}_{j \in M_i^1} \cdot \mathbb{I}_{r'_{ij}(v'_{-i,j}) \geq \beta_{ij}} \left(\mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \text{Rev}(\mathcal{M}_{1LA}(\hat{\mathcal{I}}_{i,j}, \beta_{ij})) \right). \end{aligned}$$

Combining the above two equations with (26) and then (25), we have

$$\begin{aligned} & \sum_i \mathbb{E}_{v'_j \sim \mathcal{D}'_j} \mathbb{I}_{j \in M_i^1} \cdot \text{Rev}(\mathcal{M}_{1LA}(\mathcal{I}'_{i,j}, \beta'_{ij})) \\ &\leq \sum_i \mathbb{E}_{v_j \sim \mathcal{D}_j} \left(\mathbb{I}_{j \in M_i^1} \cdot \mathbb{I}_{\beta_{ij} > r'_{ij}(v'_{-i,j})} \cdot \text{Rev}(\mathcal{M}_{1LA}(\hat{\mathcal{I}}_{i,j}, \beta_{ij})) + \mathbb{I}_{j \in M_i^2 \cap M_i} \cdot \max_{i' \neq i} v_{i'j} \right) \\ &\quad + \sum_i \mathbb{E}_{v_{-i,j} \sim \mathcal{D}_{-i,j}} \mathbb{I}_{j \in M_i^1} \cdot \mathbb{I}_{r'_{ij}(v'_{-i,j}) \geq \beta_{ij}} \left(\mathbb{E}_{v_{ij} \sim \mathcal{D}_{ij}} \text{Rev}(\mathcal{M}_{1LA}(\hat{\mathcal{I}}_{i,j}, \beta_{ij})) \right) \\ &= \sum_i \mathbb{E}_{v_j \sim \mathcal{D}_j} \left(\mathbb{I}_{j \in M_i^1} \cdot \text{Rev}(\mathcal{M}_{1LA}(\hat{\mathcal{I}}_{i,j}, \beta_{ij})) + \mathbb{I}_{j \in M_i^2 \cap M_i} \cdot \max_{i' \neq i} v_{i'j} \right), \end{aligned}$$

and (24) holds.

Taking summation over all items j on both sides of (24) and combining with (23), we have

$$\begin{aligned} & \mathbb{E}_{v \sim \mathcal{D}} \text{Rev}(\mathcal{M}_{IEA}(\mathcal{I})) \geq \sum_{j \in M} \sum_i \mathbb{E}_{v'_j \sim \mathcal{D}'_j} \mathbb{I}_{j \in M_i^1} \cdot \text{Rev}(\mathcal{M}_{1LA}(\mathcal{I}'_{i,j}, \beta'_{ij})) \\ &= \mathbb{E}_{v' \sim \mathcal{D}'} \sum_i \sum_{j \in M_i^1} \text{Rev}(\mathcal{M}_{1LA}(\mathcal{I}'_{i,j}, \beta'_{ij})) = \mathbb{E}_{v' \sim \mathcal{D}'} \text{Rev}(\mathcal{M}_{1LA}(\mathcal{I}')) \geq \frac{1}{2} \mathbb{E}_{v' \sim \mathcal{D}'} \text{IM}(\mathcal{I}'), \end{aligned}$$

where the inequality is because \mathcal{M}_{1LA} is a 2-approximation to Myerson's mechanism for each item [60]. Thus Lemma 12 holds. \square

Acknowledgements An extended abstract of this paper appeared in Proceedings of the 11th International Symposium on Algorithmic Game Theory (SAGT 2018) [62].

Author contributions All three authors designed the mechanisms, proved the theorems, and prepared the manuscript together.

Funding Open access funding provided by The Hong Kong Polytechnic University. Bo Li is funded by the Hong Kong SAR Research Grants Council (No. PolyU 25211321). and the Guangdong Basic and Applied Basic Research Foundation (No. 2024A1515011524). Yingkai Li is funded by the NUS Start-up Grant.

Declarations

Competing interests The authors declare no competing interests

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research.*, 6(1), 58–73.
2. Wilson, R. (1987). Game-theoretic analysis of trading processes. In: T. Bewley (Ed.), *Advances in economic theory: 5th world congress*(pp. 33–77).
3. Miller, N., Resnick, P., & Zeckhauser, R. (2005). Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9), 1359–1373.
4. Chawla, S., Hartline, J., & Kleinberg, R. (2007). Algorithmic pricing via virtual valuations. In: *8th Conference on electronic commerce (EC'07)*(pp. 243–251).
5. Hart, S., & Nisan, N. (2012). Approximate revenue maximization with multiple items. In: *13th Conference on electronic commerce (EC'12)*(pp. 656–656).
6. James, W. (1979). *Some Problems of Philosophy*. Cambridge: Harvard University Press.
7. Azar, P., Chen, J., & Micali, S. (2012). Crowdsourced Bayesian auctions. In: *3rd Innovations in theoretical computer science conference (ITCS'12)*(pp. 236–248).
8. Kleinberg, R., & Weinberg, S. M. (2019). Matroid prophet inequalities and applications to multi-dimensional mechanism design. *Games and Economic Behavior.*, 113, 97–115.
9. Cai, Y., Devanur, N. R., & Weinberg, S. M. (2019). A duality-based unified approach to bayesian mechanism design. *SIAM Journal on Computing.*, 50(3), STOC16–160–STOC16–200.
10. Chawla, S., Hartline, J., Malec, D., & Sivan, B. (2010). Multi-parameter mechanism design and sequential posted pricing. In: *43th symposium on theory of computing (STOC'10)*(pp. 311–320).
11. Yao, A. (2015). An n-to-1 bidder reduction for multi-item auctions and its applications. In: *Symposium on discrete algorithms (SODA'15)*(pp. 92–109).
12. Brier, G. W. (1950). Verification of Forecasts Expressed in Terms of Probability. *Monthly Weather Review.*, 78(1), 1–3.
13. Radanovic, G., & Faltings, B. (2013). A robust Bayesian truth serum for non-binary signals. In: *27th AAAI conference on artificial intelligence (AAAI'13)*(pp. 833–839).
14. Zhang, P., & Chen, Y. (2014). Elicitability and knowledge-free elicitation with peer prediction. In: *13th International conference on autonomous agents and multigent systems (AAMAS'14)*(pp. 245–252).
15. Kong, Y., Ligett, K., & Schoenebeck, G. (2016). Putting peer prediction under the micro (economic) scope and making truth-telling focal. In: *12th International conference on web and internet economics (WINE'16)*(pp. 251–264).
16. McCarthy, J. (1956). Measures of the value of information. *Proceedings of the National Academy of Sciences of the United States of America.*, 42(9), 654.
17. Cooke, R. (1991). Experts in uncertainty: Opinion and subjective probability in science. Oxford University Press on Demand
18. Hanson, R. (2007). Logarithmic markets coring rules for modular combinatorial information aggregation. *The Journal of Prediction Markets.*, 1(1), 3–15.
19. Gneiting, T., & Raftery, A. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association.*, 102(477), 359–378.

20. Savage, L. J. (1971). Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association.*, 66(336), 783–801.
21. Lambert, N. S. (2022). Elicitation and evaluation of statistical forecasts. Revised and resubmitted at *Econometrica*
22. Li, Y., Hartline, J. D., Shank L., & Wu, Y. (2022). Optimization of scoring rules. In: *23rd ACM conference on economics and computation (EC'22)*(pp. 988–989).
23. Neyman, E., Noarov, G., & Weinberg, S. M. (2021). Binary scoring rules that incentivize precision. In: *22nd ACM conference on economics and computation (EC'21)*(pp. 718–733).
24. Hartline, J. D., Shan, L., Li, Y., & Wu, Y. (2023). Optimal scoring rules for multi-dimensional effort. In: *36th Annual conference on learning theory (COLT'23)*(pp. 2624–2650).
25. Chen, Y., Yu, F. Y. (2021). Optimal scoring rule design. Working Paper
26. Xu, S., Zhang, Y., Resnick, P., & Schoenebeck, G. (2024). Spot check equivalence: an interpretable metric for information elicitation mechanisms. In: *ACM web conference 2024 (WWW'24)*(pp. 276–287).
27. Witkowski, J., & Parkes, D. (2012). Peer prediction without a common prior. In: *13th Conference on electronic commerce (EC'12)*(pp. 964–981).
28. Frongillo, R., & Waggoner, B. (2024). Recent Trends in Information Elicitation. *ACM SIGecom Exchanges.*, 22(1), 122–134.
29. Archer, A., Tardos, É. (2001). Truthful mechanisms for one-parameter agents. In: *42nd Symposium on foundations of computer science (FOCS'01)*(pp. 482–491).
30. Hartline, J., & Karlin, A. (2007). Profit maximization in mechanism design. In N. Nisan, T. Roughgarden, É. Tardos, & V. Vazirani (Eds.), *Algorithmic Game Theory* (pp. 331–361). Cambridge
31. Cai, Y., Daskalakis, C., & Weinberg, M. (2012). Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In: *53rd Symposium on foundations of computer science (FOCS'12)*(pp. 130–139).
32. Jin, Y., & Lu, P. (2024). Benchmark-tight approximation ratio of simple mechanism for a unit-demand buyer. In: *2024 IEEE 65th annual symposium on foundations of computer science (FOCS)*(pp. 1251–1259).
33. Li, X., & 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.
34. Babaioff, M., Immorlica, N., Lucier, B., & Weinberg, S. M. (2014). A simple and approximately optimal mechanism for an additive buyer. In: *55rd Symposium on foundations of computer science (FOCS'14)*(pp. 21–30).
35. Chawla, S., Miller, J. B. (2016). Mechanism design for subadditive agents via an ex ante relaxation. In: *16th Conference on economics and computation (EC'16)*(pp. 579–596). ACM
36. Cai, Y., & Zhao, M. (2017). Simple mechanisms for subadditive buyers via duality. In: *49th Symposium on theory of computing (STOC'17)*(pp. 170–183).
37. Yao, A. C. C. (2018). On revenue monotonicity in combinatorial auctions. In: *11th International symposium on algorithmic game theory*(pp. 1–11). Springer
38. Correa, J., & Cristi, A. (2023). A constant factor prophet inequality for online combinatorial auctions. In: *55th annual ACM symposium on theory of computing (STOC'23)*(pp. 686–697).
39. Cai, Y., Chen, Z., & Wu, J. (2023). Simultaneous auctions are approximately revenue-optimal for subadditive bidders. In: *2023 IEEE 64th annual symposium on foundations of computer science (FOCS)*(pp. 134–147).
40. Hartline, J., & Roughgarden, T. (2008). Optimal mechanism design and money burning. In: *40th Symposium on theory of computing (STOC'08)*(pp. 75–84).
41. Devanur, N., & Hartline, J. (2009). Limited and online supply and the Bayesian foundations of prior-free mechanism design. In: *10th Conference on electronic commerce (EC'09)*(pp. 41–50).
42. Feng, Y., Hartline, J. D., & Li, Y. (2021). Revelation gap for pricing from samples. In: *53rd Annual ACM SIGACT symposium on theory of computing (STOC'21)*(pp. 1438–1451).
43. Fu, H., Haghpahan, N., Hartline, J., & Kleinberg, R. (2021). Full surplus extraction from samples. *Journal of Economic Theory.*, 193, Article 105230.
44. Allouah, A., Bahamou, A., & Besbes, O. (2022). Pricing with samples. *Operations Research.*, 70(2), 1088–1104.
45. Cole, R., & Roughgarden, T. (2014). The sample complexity of revenue maximization. In: *46th Symposium on theory of computing (STOC'14)*(pp. 243–252).
46. Devanur, N., Huang, Z., Psomas, C. A. (2016). The sample complexity of auctions with side information. In: *48th Symposium on theory of computing (STOC'16)*(pp. 426–439).
47. Morgenstern, J., & Roughgarden, T. (2016). Learning simple auctions. In: *29th Conference on learning theory (COLT'16)*(pp. 1298–1318).
48. Chen, J., Li, B., Li, Y., & Lu, P. (2022). Bayesian auctions with efficient queries. *Artificial Intelligence.*, 303,

49. Hu, Y., Huang, Z., Shen, Y., Wang, X. (2021). Targeting makes sample efficiency in auction design. In: *22nd ACM conference on economics and computation (EC'21)*(pp. 610–629).
50. Chen, J., & Micali, S. (2013). Mechanism design with possibilistic beliefs. *Journal of Economic Theory*, 156, 77–102.
51. Chen, J., Micali, S., & Pass, R. (2015). Tight revenue bounds with possibilistic beliefs and level-k rationality. *Econometrica*, 83(4), 1619–1639.
52. Azar P, Kleinberg, R., & Weinberg, M. (2014). Prophet inequalities with limited information. In: *25th Symposium on discrete algorithms (SODA'14)*(pp. 1358–1377).
53. Suzdaltsev, A. (2022). Distributionally robust pricing in independent private value auctions. *Journal of Economic Theory*, 206, Article 105555.
54. Allouah, A., Bahamou, A., & Besbes, O. (2023). Optimal pricing with a single point. *Management Science*, 69(10), 5866–5882.
55. Anunrojwong, J., Balseiro, S. R., & Besbes, O. (2024). On the robustness of second-price auctions in prior-independent mechanism design. *Operations Research*
56. Bergemann, D., & Morris, S. (2012). *Robust Mechanism Design*. Hackensack, NJ: World Scientific.
57. Harsanyi, J. (1967). Games with incomplete information played by “Bayesian” players, I–III. Part I: The basic model. *Management Science*, 14(3), 159–182.
58. Aumann, R. (1976). Agreeing to disagree. *The Annals of Statistics*, 4(6), 1236–1239.
59. Fagin, R., Halpern, J. Y., Moses, Y., & Vardi, M. (2004). Reasoning about knowledge. MIT press
60. Ronen, A. (2001). On approximating optimal auctions. In: *3rd Conference on electronic commerce (EC'01)*(pp. 11–17).
61. Cremer, J., & McLean, R. (1988). Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica*, 56(6), 1247–1257.
62. Chen, J., Li, B., Li, Y. (2018). Information elicitation for bayesian auctions. In: *11th International symposium on algorithmic game theory (SAGT'18)* (vol. 11059, pp. 43–55).
63. Cervera, J., & Munoz, J. (1996). Proper scoring rules for fractiles. *Bayesian Statistics*, 5, 513–519.
64. Artemov, G., Kunitomo, T., & Serrano, R. (2013). Robust virtual implementation: Toward a reinterpretation of the Wilson doctrine. *Journal of Economic Theory*, 148(2), 424–447.
65. Bergemann, D., & Morris, S. (2014). Informational robustness and solution concepts. Princeton University William S. Dietrich II Economic Theory Center Research Paper No. 067_2014.
66. Carroll, G., & Meng, D. (2016). Robust contracting with additive noise. *Journal of Economic Theory*, 166, 586–604.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.