

Incentive-Compatible Learning of Reserve Prices for Repeated Auctions*

Yash Kanoria

Graduate School of Business, Columbia University

ykanoria@columbia.edu

Hamid Nazerzadeh

Marshall School of Business, University of Southern California

hamidnz@marshall.usc.edu

A large fraction of online advertisements are sold via repeated second-price auctions. In these auctions, the reserve price is the main tool for the auctioneer to boost revenues. In this work, we investigate the following question: How can the auctioneer optimize the reserve prices by learning from the previous bids, while accounting for the long-term incentives and strategic behavior of the bidders? To this end, we consider a seller who repeatedly sells ex ante identical items via a second-price auction. Buyers' valuations for each item are drawn i.i.d. from a distribution F that is unknown to the seller. We find that if the seller attempts to dynamically update a common reserve price based on the bidding history, this creates an incentive for buyers to shade their bids, which can hurt revenue. When there is more than one buyer, incentive compatibility can be restored by using *personalized* reserve prices, where the personal reserve price for each buyer is set using the historical bids of *other* buyers. Such a mechanism asymptotically achieves the expected revenue obtained under the static Myerson optimal auction for F . Further, if valuation distributions differ across bidders, the loss relative to the Myerson benchmark is only quadratic in the size of such differences. We extend our results to a contextual setting where the valuations of the buyers depend on observed features of the items. When up-front fees are permitted, we show how the seller can determine such payments based on the bids of others to obtain an approximately incentive compatible-mechanism that extracts nearly all the surplus.

1. Introduction

Advertising is the main component of the monetization strategy of most Internet companies. A large fraction of online advertisements are sold via auctions where advertisers bid in real time for a chance to show their ads to users. Examples of such auction platforms, called advertisement

* Minor revision submitted to Operations Research on Oct 30, 2019.

exchanges (Muthukrishnan 2009, McAfee 2011), include Google’s Doubleclick (AdX), Facebook, AppNexus, and OpenX.

The second-price auction is a common mechanism used by advertisement exchanges. It is a simple mechanism that incentivizes advertisers to be truthful in a static setting. The second-price auction can maximize the social welfare (i.e., the value created in the system) by allocating the item to the highest bidder.

To maximize the revenue earned in a second-price auction, the auctioneer can set a reserve price and not make any allocations when the bids are low. In fact, under symmetry and regularity assumptions (see Section 2), the second-price auction with an appropriately chosen reserve price is optimal and maximizes revenue among all selling mechanisms (Myerson 1981, Riley and Samuelson 1981).

However, to set the reserve price effectively, the auctioneer requires information about the distribution of the valuations of the bidders. A natural idea, which is widely used in practice, is to construct these distributions using the history of the bids. This approach, though intuitive, raises a major concern with regard to the long-term (dynamic) incentives of the advertisers. Because the bid of an advertiser may determine the price she pays in future auctions, this approach may lead the advertisers to shade their bids and ultimately result in a loss of revenue for the auctioneer.

To understand the effects of changing reserve prices based on previous bids, we study a setting where the auctioneer sells impressions (advertisement space) via repeated second-price auctions. More specifically, in the main model we consider, the valuations of the bidders are drawn i.i.d. from a distribution. The bidders are strategic and aim to maximize their cumulative utility. We demonstrate that the long-term incentives of advertisers play an important role in the performance of these repeated auctions.

We show that natural mechanisms that set a common reserve price using the history of the bids may create substantial incentives for the buyers to shade their bids. On the other hand, we propose an *incentive-compatible* mechanism that sets a personalized reserve price for each agent based on the bids of *other* agents.¹ Our mechanism allocates the item to the highest bidder if his bid exceeds his personal reserve price. If the item is allocated, the price is equal to the maximum of the second-highest bid and the personal reserve price of the winner. This structure corresponds to mechanisms used in practice, as described in Paes Leme et al. (2016). By appropriately choosing the function that maps historical bids of others to a personal reserve price, we show that the expected revenue

¹ In case of unlimited supply the incentive compatibility directly follows if the price of each buyer only depends on the other buyers, cf., Balcan et al. (2008). With limited supply, obtaining incentive compatibility is more challenging because of “competition” among buyers.

per round is asymptotically as large as that under the static Myerson optimal auction that a-priori knows the distribution of the bids.²

We discussed earlier that using the bids from the others has a “first order effect” of preventing a bidder from lowering the reserve price he will see in future by misreporting his valuation. However, we show that despite only using bids of other agents, there is room for a “second order effect” by which a bidder could seek to benefit by affecting the reserve prices of others and thus indirectly himself. Hence, importantly, for addressing the second order effect, our mechanism is “lazy” (see more on this in the section on related work below), in that it allocates the item only to the highest bidder (if she clears her personal reserve); it is not enough for *some* bidder to clear her personal reserve. An “eager” variant would allocate the item to the highest bidder among those who clear their reserve; this would create an incentive for agents to overbid so as to increase the personal reserves of other agents in future, thereby increasing their likelihood of being eliminated.

As described earlier, our mechanism allocates the item to the highest bidder if her bid exceeds her personal reserve price. The personal reserve is chosen to maximize revenue for a distribution estimated using other agents’ bids. A natural concern with such an approach is that if agents’ valuation distributions differ from each other, it may lead to a lower personal reserve price for agents with a higher valuation distribution, and vice versa, thereby hurting revenue. We show that this issue is not significant when differences in valuation distributions are not too large (our notion of the distance between two distributions is the maximum absolute difference between their virtual value functions). In particular, we show that the loss relative to the Myerson benchmark is only quadratic in the size of such differences, and supplement this theoretical result with numerical examples.

We also generalize our result across another dimension. Namely, we extend our results to a contextual setting with heterogeneous items that are represented by a feature vector of covariates. The valuations of the buyers are linear in the feature vectors (with a-priori unknown coefficients) plus an idiosyncratic private component. We present a learning algorithm that determines the reserve price for each buyer using an ordinary least squares estimator for the feature coefficients. We show that the loss of revenue is sub-linear in the number of samples (previous auctions).

For the aforementioned results, we benchmarked the performance of the mechanisms with respect to the static Myerson optimal auction that knows the distribution of the bids in advance. However, we note that this static mechanism is not the optimal mechanism among the class of dynamic mechanisms. In fact, we present a mechanism that can extract (almost all of) the surplus of the

² From a technical perspective, we build on prior work that investigates how samples from a distribution can be used to set a near-optimal reserve price, cf., Dhangwatnotai et al. (2015).

agents. The basic idea is that using the bids of other agents, the seller can construct an estimate of the valuation distribution and hence of the expected utility per round of each agent when individual items are allocated using second-price auctions. Based on this estimate, the mechanism charges a surplus-extracting up-front payment at the beginning of each round. Since agents can influence the up-front payments of other agents, they may have an incentive to overbid to eliminate competing agents from future auctions. We propose a solution that asymptotically removes the incentive for agents to deviate from truthfulness: the mechanism *simulates* agents who choose not to pay the entrance fee. We show that under our mechanism, truthfulness constitutes an approximate equilibrium.

In each of our proposed mechanisms, we overcome incentive issues using the same two key ideas: (i) we eliminate incentives for lowering bids by individually choosing a pricing rule for each agent, based only on the bids of *other* agents, and (ii) we disincentivize *overbidding* by preventing an agent from benefiting from suppressing participation of other agents by raising the prices they face; this has been achieved in our setting by allocating the item only to the highest bidder in our mechanisms that achieve the Myerson benchmark, and by simulating non-participating agents in our surplus-extracting mechanism. In a setting where agents' valuation distributions are identical (or similar to each other), this approach enables the seller to obtain as much revenue as if he knew the valuation distribution F , while maintaining incentive compatibility. We believe that these design principles should be broadly applicable to overcome the lack of knowledge of F when there is competition between strategic agents/buyers; see Section 8 for further discussion.

Related Work

In this section, we briefly discuss work closely related to ours along two dimensions, behavior-based (personalized) pricing and reserve pricing optimization for online advertising.

Behavior-Based Pricing. Our work is closely related to the literature on behavior-based pricing strategies where the seller changes the prices for *one* buyer (or a segment of the buyers) based on her previous behavior. For instance, the seller may increase the price after a purchase or reduce the price in the case of no-purchase; see Fudenberg and Villas-Boas (2007) and Esteves (2009) for surveys.

The common insight from the literature is that the optimal pricing strategy is to commit to a single price over the length of the horizon (Stokey 1979, Salant 1989, Hart and Tirole 1988). In fact, when customers anticipate a future reduction in prices, dynamic pricing may hurt the seller's revenue (Taylor 2004, Villas-Boas 2004). Similar insights are obtained in environments where the goal is to sell a fixed initial inventory of products to unit-demand buyers who arrive over time (Aviv and Pazgal 2008, Dasu and Tong 2010, Aviv et al. 2013, Correa et al. 2016).

There has been renewed interest in behavior-based pricing strategies, mainly motivated by the development in e-commerce technologies that enable online retailers and other Internet companies to determine the price for the buyer based on her previous purchases. Acquisti and Varian (2005) show that when a sufficient proportion of customers are myopic or when the valuations of customers increase (by providing enhanced services), dynamic pricing may increase the revenue. Another setting where dynamic pricing can boost the revenue is when the seller is more patient than the buyer and discounts his utility over time at a lower rate than the buyer (Bikhchandani and McCardle 2012, Amin et al. 2013, Mohri and Medina 2014a, Chen and Wang 2016). See Taylor (2004) and Conitzer et al. (2012) for privacy issues and anonymization approaches in this context. In contrast with these works, our focus is on auction environments and we study the role of competition among strategic bidders who remain in the system over a long horizon. We observe that when there is competition, there is value in personalizing prices, in particular, when the valuations are drawn i.i.d. over time. In fact, the seller can extract nearly all the surplus.

The problem of learning the distribution of valuations and optimal pricing has also been studied in the context of revenue management and pricing for markets where each (infinitesimal) buyer does not have an effect on future prices and the demand curve can be learned with near-optimal regret (Segal 2003, Baliga and Vohra 2003, Besbes and Zeevi 2009, Harrison et al. 2012, Wang et al. 2014); see den Boer (2015) for a survey. In this work, we consider a setting where the goal is to learn the optimal reserve price with a small number of strategic and forward-looking buyers with multi-unit demand, where the action of each buyer can change the prices in the future.

Reserve Price Optimization. Several recent work have studied reserve price optimization. Most of the if focused on algorithmic issues but ignores strategic aspects and incentive-compatibility issues, cf., Cesa-Bianchi et al. (2013), Mohri and Medina (2014b), Cesa-Bianchi et al. (2015), Roughgarden and Wang (2016), Golrezaei et al. (2018a). Most closely related to our work is the work by Paes Leme et al. (2016) who compare different generalizations of the second-price auction with a personalized reserve. In their “lazy” version, the item is allocated only to the highest bidder. In their “eager” version, first, all the bidders below their personal reserve are eliminated, and then the item is allocated to the highest surviving bidder. From an optimization/learning perspective, they show that lazy reserves are easy to optimize and A/B test in production, whereas eager reserves lead to higher surplus, but their optimization is NP-complete, and naive A/B testing leads to incorrect conclusions. Whereas in their setting both eager and lazy are incentive compatible, this is not true in our setting. The mechanism we propose corresponds to their lazy version. We show how this mechanism – a lazy second-price auction with personalized reserves – can be used to optimize reserve prices in an incentive-compatible way by appropriately learning from the previous bids (in contrast, the eager version creates incentives to overbid).

Ostrovsky and Schwarz (2009) conducted a large-scale field experiment at Yahoo! and showed that choosing reserve prices guided by the theory of optimal auctions can significantly increase the revenue of sponsored search auctions. To mitigate the aforementioned incentive concerns, they drop the highest bid from each auction when estimating the distribution of the valuations. However, they do not formally discuss the consequences of this approach.

Another common solution offered to mitigate incentive constraints is to bundle a large number of impressions (or keywords) together so that the bid of each advertiser has little impact on the aggregate distribution learned from the history of bids. However, this approach may lead to significant estimation errors since a variety of different types of impressions fall into the same bundle, resulting in a suboptimal choice of reserve price, cf. Epasto et al. (2018). To the best of our knowledge, the present work is the first to rigorously study the long-term and dynamic incentive issues in repeated auctions with dynamic reserves.

Organization. The rest of the paper is organized as follows. We formally present our model in Section 2. In Section 3 we show that mechanisms which optimize a common reserve price suffer from incentive issues, and this may also hurt revenue. By contrast, in Section 4, we present truthful mechanisms with personalized reserve prices, where the reserves are optimized based on earlier bids by competing agents. In Section 5, we show robustness of our revenue guarantee to differences in valuation distributions across buyers. Then, in Section 6, generalize our result to the case of heterogeneous items. Finally, we present a truthful surplus-extracting mechanism in Section 7. Proofs are deferred to the appendices.

2. Model and Preliminaries

A seller, using a second-price auction, sells items over time to $n \geq 1$ agents. The valuation of agent $i \in \{1, \dots, n\}$ for an item at time t , denoted by v_{it} , is drawn independently and identically from distribution F . (Later, in Section 5, we will consider the case where different agents have different valuation distributions.) There is exactly one item for sale at each time $t = 1, 2, \dots$. In Section 6 we extend our results to a contextual setting with heterogeneous items. For the sake of simplicity, we assume that the length of the horizon is infinite and the seller and the agents respectively aim to maximize their average long-term revenue and utility. This is a reasonable assumption, given the very large number of impressions sold in practice.

More specifically, the average per-round revenue of the seller, denoted by REV , is equal to

$$\text{REV} = \lim_{T \rightarrow \infty} \left(\frac{1}{T} \times \mathbb{E} \left[\sum_{t=1}^T \sum_{i=1}^n p_{it} \right] \right), \quad (1)$$

where p_{it} denotes the payment of agent i at time t . Note that if the limit exists, then the average revenue is maximized. Otherwise, the seller aims to maximize the worst-case average revenue. Similarly, for the average per-round utility of buyer i , denoted by U_i , we have

$$U_i = \lim_{T \rightarrow \infty} \left(\frac{1}{T} \times \mathbb{E} \left[\sum_{t=1}^T (v_{it} q_{it} - p_{it}) \right] \right), \quad (2)$$

where $q_{it} = 1$ if the item at time t is allocated to agent i and otherwise it is equal to 0. The expectations are with respect to the realizations of the valuations of the agents and any randomization in the mechanism and agent strategies. Each agent aims to maximize the worst-case average utility. The mechanisms we will introduce and the corresponding equilibria/strategies of agents will be stationary in time (after an initial transient), hence the above limits will exist.

We assume that the distribution of valuations F is unknown to the auctioneer/seller, who may not even have a prior on F . The valuation v_{it} of agent i is privately known to agent i . To simplify the presentation, we assume that the distribution of the valuations F is common knowledge among the agents. (We discuss our informational assumptions later in this section.) We assume that F is a monotone hazard rate (MHR) distribution, i.e., the hazard rate $f(v)/(1 - F(v))$ is monotone non-decreasing in v . MHR distributions include all sufficiently light-tailed distributions, including uniform, exponential, and normal. For most of our results, we provide versions that apply to the larger class of *regular* distributions, i.e., distributions for which the virtual value function $\phi(v) = v - (1 - F(v))/f(v)$ is monotone increasing in v . (For instance, log-normal distributions are regular but not MHR.)

Let us now consider the seller's problem. The seller aims to maximize his expected revenue via a repeated second-price auction, despite his lack of knowledge of F . He can attempt to do this by dynamically updating the reserve price based on history of bids so far.

A “generic” dynamic second-price mechanism. At time 0, the auctioneer announces the reserve price function $\Omega : \mathcal{H} \rightarrow \mathbb{R}^+$ that maps the history observed by the mechanism to a reserve price. The history observed by the mechanism up to time τ , denoted by $H_{\Omega, \tau} \in \mathcal{H}$, consists of the reserve price, the agents participating in round t and their bids, and the allocation and payments for each round $t < \tau$. More precisely,

$$H_{\Omega, \tau} \triangleq \langle (r_1, b_1, q_1, p_1), \dots, (r_{\tau-1}, b_{\tau-1}, q_{\tau-1}, p_{\tau-1}) \rangle,$$

where

- r_t is the reserve price at time t .
- $b_t = \langle b_{1t}, \dots, b_{nt} \rangle$ where b_{it} denotes the bid of agent i at time t .

- q_t corresponds to the allocation vector. If all the bids are smaller than the reserve price r_t , the item is not allocated. Otherwise, the item is allocated to agent $i^* = \arg \max_i \{b_{it}\}$ and we have $q_{i^*t} = 1$; in case of a tie, the item is allocated to a uniformly random agent among those who bid highest. For all the agents who do not receive the item, q_{it} is equal to 0.

- p_t is the vector of payments. If $q_{it} = 0$, then $p_{it} = 0$; and if $q_{it} = 1$, then $p_{it} = \max \{\max_{j \neq i} \{b_{jt}\}, r_t\}$.

In our notation, Ω specifies a reserve price function for each period t . Note that the auctioneer commits beforehand to a reserve update rule Ω . It is well known that in the absence of commitment, the seller earns less revenue (see, e.g., Devanur et al. 2014).

An important subclass of the mechanisms above is *static* mechanisms where the reserve does not depend on the history or time. Another important subclass is *window-based* mechanisms, with window length W , which use only the bids received in the previous W periods to determine the reserve in the next period.³ A window-based mechanism is *stationary* if the rule that maps bids in the last W periods to the reserve in period t does not depend on t . When considering stationary window-based mechanisms, we call the function (a close cousin of Ω) that maps the history of bids in the last W periods to the reserve price in the next period the *reserve optimization function* (ROF).

The seller aims to choose a reserve price function Ω that maximizes the average revenue, defined in Eq. (1), when the buyers play an equilibrium with respect to the choice of Ω . To define the utility of the agents and the information available to them, let $H_{i,\tau}$ denote the history observed by agent i up to time τ , consisting of her valuations, bids, allocations and payments. Namely,

$$H_{i,\tau} = \langle (v_{i1}, b_{i1}, q_{i1}, p_{i1}), \dots, (v_{i,\tau-1}, b_{i,\tau-1}, q_{i,\tau-1}, p_{i,\tau-1}) \rangle.$$

We refer to $H_{i,\tau}$ as the *personal history* of agent i .

We assume that agents do not see the reserve price before they bid,⁴ but that they know the reserve price update rule Ω .

The bidding strategy $B_i : \mathcal{H}_i \times \mathbb{R} \rightarrow \mathbb{R}$ of agent i maps the valuation of the agent $v_{i\tau}$, the history $H_{i,\tau}$, and the reserve r_τ at time τ to a bid $b_{i\tau}$. Here $\mathcal{H}_i$ is the set of possible histories observed by agent i .

³ We allow such a mechanism to use only the bids and not the reserve prices (nor the allocations and payments), since the entire history can be “encoded” in the decimal representation of reserve r_τ , with vanishing impact on revenues, and this would defeat the purpose of defining window-length W .

⁴ This is similar to the common practice in ad exchanges, where the bidder may not see the reserve. Often the exchanges communicates a (possibly lower) reserve price, which may be different from the reserve price that is applied to the payments.

Finally, we define the history of the game up to time τ as

$$H_\tau = \langle (r_1, v_1, b_1, q_1, p_1), \dots, (r_{\tau-1}, v_{\tau-1}, b_{\tau-1}, q_{\tau-1}, p_{\tau-1}) \rangle.$$

Note that compared to $H_{\Omega, \tau}$, which is the history observed by the seller, H_τ also includes the valuations of the agents.

We say that an agent plays the always truthful strategy, or we simply call the agent *truthful*, if at every time t , we have $b_{it} = v_{it}$ irrespective of the history H_{it} and the reserve r_t . We now formalize our definition of incentive-compatibility. We define the *inf-utility* and *sup-utility* of agent i when each agent i' plays strategy $B_{i'}$ respectively as follows.

$$\begin{aligned} \underline{U}_i(B_i, B_{-i}) &= \liminf_{T \rightarrow \infty} \left(\frac{1}{T} \times \mathbb{E} \left[\sum_{t=1}^T v_{it} q_{it} - p_{it} \right] \right), \\ \overline{U}_i(B_i, B_{-i}) &= \limsup_{T \rightarrow \infty} \left(\frac{1}{T} \times \mathbb{E} \left[\sum_{t=1}^T v_{it} q_{it} - p_{it} \right] \right). \end{aligned}$$

We say that a mechanism is *incentive compatible* (IC) if, for each agent i , if other agents are always truthful, then the inf-utility under the always truthful strategy (weakly) exceeds the sup-utility under any other strategy. Formally, we require

$$\underline{U}_i(B_i^{\text{Tr}}, B_{-i}^{\text{Tr}}) \geq \overline{U}_i(B_i, B_{-i}^{\text{Tr}})$$

for any strategy B_i , where B_i^{Tr} denotes the truthful strategy. Intuitively, a mechanism is IC if the always truthful strategy by all agents constitutes a Nash equilibrium. We emphasize that since our environment and proposed mechanisms are time-invariant (after an initial transient), and always truthful is also a time-invariant strategy, the right-hand side of the definition of utility (2) has a limiting value as $T \rightarrow \infty$ when all agents are always truthful. More precisely, $\underline{U}_i(B_i^{\text{Tr}}, B_{-i}^{\text{Tr}})$ is well defined and equal to $\underline{U}_i(B_i^{\text{Tr}}, B_{-i}^{\text{Tr}})$.

The above notion of incentive-compatibility is static in the sense that the strategies that agents choose before the game starts define an equilibrium. We now define a stronger and dynamic notion. We say that a mechanism is *dynamic incentive compatible*, or more precisely *periodic ex-post incentive-compatible*, if at every time τ , for every history H_τ , each agent i 's best-response strategy to her personal history $H_{i, \tau}$ is to be truthful assuming that all the other agents will be truthful in future (Bergemann and Välimäki 2010). More precisely, define the future inf utility of an agent as

$$\underline{U}_{i, H_{i, \tau}}(B_i, B_{-i}^{\text{Tr}}) = \liminf_{T \rightarrow \infty} \left(\frac{1}{T} \times \mathbb{E}_{H_{i, \tau}} \left[\sum_{t=\tau}^T v_{it} q_{it} - p_{it} \right] \right), \quad (3)$$

i.e., it is the (worst-case) future per-auction utility of agent i at time τ , assuming all other agents will be truthful and agent i plays strategy B_i . Again, this limit will exist for the mechanisms we consider when $B_i = B_i^{\text{Tr}}$. Also define the future sup utility as

$$\overline{U}_{i,H_{i,\tau}}(B_i, B_{-i}^{\text{Tr}}) = \limsup_{T \rightarrow \infty} \left(\frac{1}{T} \times \mathbb{E}_{H_{i,\tau}} \left[\sum_{t=\tau}^T v_{it} q_{it} - p_{it} \right] \right). \quad (4)$$

A mechanism is *dynamic incentive compatible* if, for each agent i , we have

$$\underline{U}_{i,H_{i,\tau}}(B_i^{\text{Tr}}, B_{-i}^{\text{Tr}}) \geq \overline{U}_{i,H_{i,\tau}}(B_i, B_{-i}^{\text{Tr}})$$

for any time τ , any personal history $H_{i,\tau}$, and any strategy B_i where B_i^{Tr} denotes the truthful strategy. As discussed earlier, when all agents follow a truthful strategy in our setting we have

$$\underline{U}_{i,H_{i,\tau}}(B_i^{\text{Tr}}, B_{-i}^{\text{Tr}}) = U_{i,H_{i,\tau}}(B_i^{\text{Tr}}, B_{-i}^{\text{Tr}}) = \lim_{T \rightarrow \infty} \left(\frac{1}{T} \times \mathbb{E}_{H_{i,\tau}} \left[\sum_{t=\tau}^T v_{it} q_{it} - p_{it} \right] \right).$$

In Section 7, we present an approximate notion of dynamic incentive-compatibility.

Discussion on Informational Assumptions. Our results are not sensitive to our informational assumptions. Our main results (Theorems 1, 2 and 4) are for incentive compatible mechanisms, so they hold even if the agents do not have perfect information regarding the valuation distribution(s) and/or the reserve price function Ω . Theorems 1-4 (and their proofs) remain valid if agents obtain information regarding past reserve prices and past bids, allocations and payments of other agents.⁵ Regarding the seller, our mechanisms use prior-free learning algorithms. Of course, the revenue guarantees remain valid if the seller *does* know something about the valuation distribution(s).

Finally, consider our negative results in Section 3 (Example 1 and Proposition 1). Providing additional information to agents can only make things (weakly) worse. On the other hand, strategic bid-shading by an agent does rely on knowledge of the valuation distribution; note that if an agent initially lacks this knowledge, she can acquire it over time.

Benchmark. In the first part of the paper, we restrict ourselves to dynamic second-price mechanisms. We use as a benchmark the average revenue that could have been achieved via the optimal static mechanism if F had been known to the seller, i.e., the average revenue per round under the static Myerson auction with the optimal reserve for distribution F . (Note that since F is an MHR distribution, Myerson's result says that the optimal static mechanism is, in fact, a second-price auction with a reserve price. This extends to the case where F is a regular distribution.) Let

⁵ This informational robustness is in contrast with repeated first-price auction settings (e.g., see Bergemann and Horner 2010) where information revelation can significantly change the outcome.

REV_* denote the benchmark average revenue. We demonstrate an incentive-compatible second-price mechanism (with personalized reserve prices) that asymptotically achieves the benchmark revenue (see Section 4). Later, in Section 7, we go beyond dynamic second-price mechanisms to allow additional mechanism features such as upfront payments. We show how, using a modification of the same ideas, the seller can approximately achieve the largest possible revenue, namely, the revenue corresponding to full surplus extraction, while retaining (approximate) incentive-compatibility.

3. Incentive Problems with Learning a Common Reserve Price

In this section, we argue that if the seller attempts to learn a common reserve price using historical bids, this leads to incentive issues, specifically, agents may shade their bids in order to reduce the reserve prices they face in future, and in turn such shading may reduce the revenue earned by the seller.

For simplicity, we start by analyzing a simple reserve-price optimization approach, which we call the histogram method, that is the basis of a lot of non-parametric approaches used in practice (see Nazerzadeh et al. 2016), and find significant issues as above. Next, in Section A.2, we argue that these issues are typical in mechanisms that attempt to learn a common reserve price from historical bids. Throughout this section, we will consider stationary mechanisms and time-invariant strategies, and look for a non-truthful agent strategy such that if other agents are always truthful, the limiting agent’s utility (as defined in (2); note that the limit exists) is strictly larger than that resulting from being always truthful.

Histogram Method. For simplicity, we consider a setting with $n = 2$ bidders and demonstrate the issue with incentives and the resulting revenue impact. (The problem is even more acute when there is just one buyer/agent, in which case the agent can drive the reserve price, and hence seller revenue, down to zero, while still winning the item each time. We comment on this case later.)

Let $\hat{F}_t$ be the joint empirical distribution of all the bids submitted during the last W periods (we will consider the limit $W \rightarrow \infty$ in our analysis below). The *histogram method* is a window-based stationary second-price mechanism with a very simple reserve optimization function (ROF). The reserve price at time t is chosen to be the one that maximizes expected revenue when the bid vector is from $\hat{F}_t$. Formally,

$$r_t = \arg \max_r \left\{ \mathbb{E}_{(b_1, b_2) \sim \hat{F}_t} \left[\max \left\{ \left(\max\{b_1, b_2\} - \max\{r, \min\{b_1, b_2\}\} \right), 0 \right\} \right] \right\}; \quad (5)$$

in case of ties, r_t is the smallest reserve in the arg max.

As described in Section 2, the seller allocates the item to a buyer with the highest bid larger than the reserve. If no bid is above r_t , the item is not allocated. In case of a tie, the item is allocated at random to one of the highest bidders.

To convey intuition about the incentive issues associated with this approach to reserve price optimization, we start with a simple model. Assume that the valuations of the bidders are drawn i.i.d. from a $\text{Uniform}(0, 1)$ distribution.

Let us see how an agent may react in response to this mechanism. Intuitively, an agent may want to shade her bid. We present a simple shading strategy where an agent shades her bid if her valuation is between two parameters $\underline{r}$ and $\bar{r}$ and bids truthfully otherwise. More specifically,

- If $0 \leq v_t \leq \underline{r}$ or $\bar{r} \leq v_t \leq 1$, then $b_t = v_t$.
- If $\underline{r} < v_t < \bar{r}$, then $b_t = \underline{r}$.

We observe that by playing the above strategy, an agent can increase her utility by reducing the reserve, *even if the other agent is truthful*.⁶ More importantly, shading can significantly increase the agent’s utility. Further, such strategic shading by an agent reduces the revenue of the seller.⁷

Example 1 (Learning a common reserve price using the Histogram method is not IC)

Assume that the valuations of the agents are i.i.d. $\text{Uniform}(0, 1)$, and that the seller uses histogram-based reserve optimization, using bids from the last W rounds, and consider $W \rightarrow \infty$. If one of the agents follows the above shading strategy for values of $\underline{r} = 0.378$ and $\bar{r} = 2/3 = 0.667$, while the other agent is always truthful, then the reserve price converges to⁸ $\underline{r} = 0.378$, with the following consequences.

Revenue. *The limiting average revenue obtained by the seller is close to 0.383. By contrast, if both agents are truthful, the limiting average revenue is equal to $\frac{5}{12} = 0.417$. If the seller does not use any reserve price, the average revenue is equal to $\frac{1}{3} = 0.333$. Therefore, more than 40% of the benefit from reserve-price optimization is lost even if one of the agents shades her bid strategically.*

Incentives. *The limiting expected utility of the agent from the always truthful strategy is close to 0.083. On the other hand, the limiting expected utility from the aforementioned shading strategy is close to 0.109. Therefore, the agent can increase her utility by more than 30% via shading.*

See Appendix A.1 for details. A little reflection immediately reveals that in the absence of competition between agents, the incentive issues associated with the histogram method are even more acute.

⁶ If both agents shade, the resulting equilibrium (or limit cycle) may involve further loss of revenue for the seller.

⁷ A more general class of strategies involves bidding some $r_0 \in [\underline{r}, \bar{r}]$ for all valuations in $[\underline{r}, \bar{r}]$ and bidding truthfully otherwise. We expect that a best response in this class would yield a larger benefit from deviation, while still hurting the revenue earned by the seller.

⁸ The numbers in this example are rounded to three decimal points; see Appendix A.1 for details.

Remark 1 *If, instead, there is only $n = 1$ agent, then the agent can employ the above shading strategy with $\underline{r} = \epsilon \in (0, 1/2)$ and $\bar{r} = \infty$ (or, equivalently, $\bar{r} = 1$). Under such an agent strategy, as $W \rightarrow \infty$, the seller’s estimated $\hat{F}_t$ has an atom of mass exceeding $1/2$ at ϵ , leading to $\lim_{W \rightarrow \infty} r_t = \epsilon$ for all $t > W$. By choosing ϵ close enough to 0, the agent can win the item in almost all rounds, while making arbitrarily small payments; thus, this is a best response for the agent as $\epsilon \rightarrow 0^+$. The result is that the seller revenue is vanishing when she uses the histogram method when selling to a single strategic agent.*

We note that our example extends to general valuation distributions F .

Remark 2 *Though we fixed F to Uniform(0,1) in Example 1, the idea easily extends to general regular F with continuous density f , when the seller sets the reserve using the histogram method. Let r_* be the Myerson optimal static reserve. (As the window-length $W \rightarrow \infty$, under truthful bidding by all agents, the reserve set by the seller converges to r_* .) Suppose that agents other than i are always truthful. Then for sufficiently small ϵ , a shading strategy based on $\underline{r} = r_* - \epsilon$ and $\bar{r} = r_* + 1.1\epsilon$ constitutes a profitable deviation for agent i , by causing the seller to set a reserve of $\underline{r} = r_* - \epsilon$ with high probability in steady state. Such shading leads to a myopic loss of $O(\epsilon^2)$ in expected utility from the current round—due to losing the item though i would have won it under truthful bidding—which occurs with probability $O(\epsilon)$ and causes an $O(\epsilon)$ loss in utility in each case. However, there is a (larger) $\Omega(\epsilon)$ increase in expected utility due to the reserve being lower by ϵ , due to bid-shading in the past. This bid-shading by agent i causes a loss in revenue of $\Omega(\epsilon)$ to the seller.*

In the Online Appendix A.2, we show that incentive concerns apply not just to the histogram method, but to a broad class of dynamic reserve mechanisms that set a common reserve based on historical bids.

4. Incentive-Compatible Optimization of Personalized Reserve Prices

In the previous section, we identified significant incentive concerns associated with optimizing the reserve price when all agents face exactly the same reserve price, and bidders are strategic. Specifically, natural mechanisms for optimizing the reserve based on historical bids encourage bidders to shade their bids, which in turn reduces seller revenue in turn.

In this section, we present a mechanism that eliminates incentives for agents to misreport their valuations. As mentioned earlier, we overcome incentive issues using two key ideas: (i) we *personalize* reserves by choosing a pricing rule⁹ for each agent, based only on the bids of *other* agents; hence, agents do not benefit from under-bidding and (ii) by allocating the item only to the highest bidder, agents do not benefit from overbidding that can prevent others from participating in the auctions.

⁹ Formally, the reserve price function Ω now outputs an n -vector of reserve prices, one for each agent.

Highest-Only Self-Excluding-Reserve-Price (HO-SERP) Mechanisms. A second-price auction with personalized reserves is an HO-SERP mechanism if it satisfies the following two properties:

- **Highest-Only:** The mechanism allocates the item only to the highest bidder. If the highest bidder i does not meet his reserve ($b_{it} < r_{it}$), then the item is not allocated. If $b_{it} \geq r_{it}$, the highest bidder i is charged a price equal to $\max\{r_{it}, \max_{j \neq i}\{b_{jt}\}\}$.
- **Self-Excluding-Reserve-Price:** The reserve price for agent i is determined using only the bids of other bidders, and does not depend on the bids of agent i herself.¹⁰ Let $\hat{F}_{-i}$ be the empirical distribution of the bids by agents other than i in the relevant rounds (a window-based mechanism will consider the last W rounds).¹¹ Then the personal reserve r_{it} of agent i is set based on $\hat{F}_{-i}$.

Personalized reserve prices may appear more complex than a mechanism with a common reserve. But we note that they have been widely used in practice; e.g., see Paes Leme et al. (2016). Moreover, we establish strong incentive properties for HO-SERP mechanisms.

To (approximately) maximize the revenue earned, we set r_{it} to be the optimal monopoly price for costless goods when faced with buyers with this valuation distribution, i.e.,¹²

$$r_{it} = \arg \max_r r(1 - \hat{F}_{-i}(r)). \quad (6)$$

This allows us to approximately achieve the revenue benchmark. The latter is proved using convergence rate bounds from Dhangwatnotai et al. (2015); other related papers on learning the optimal reserve price from samples include Cole and Roughgarden (2014), Huang et al. (2015), Devanur et al. (2016).

Theorem 1 *Any HO-SERP mechanism is periodic ex-post incentive compatible. In particular, all agents following the always truthful strategy constitutes an equilibrium. Further, there exists $C < \infty$ that does not depend on the valuation distribution F , such that for any F that is MHR and any $\epsilon \in (0, 1)$, the HO-SERP mechanism with window length $W \geq C \log(1/\epsilon)/\epsilon^2$ and personal reserve prices set as per (6) achieves average revenue per round that is at least $(1 - \epsilon)\text{REV}_*$, where REV_* is the expected revenue under the optimal static mechanism, i.e., the second-price auction with Myerson-optimal reserve.*

¹⁰ Goldberg et al. (2001) and follow-up works broadly inspired this approach, though the setting and results are quite different; there, a digital good (which can be reproduced costlessly) is sold simultaneously to multiple buyers, and the seller does not know the valuation distribution.

¹¹ To clarify this definition, suppose that there are three bidders i , j , and k . Then the bids $b_{j\tau}$ and $b_{k\tau}$ for relevant $\tau < t$ are regarded as two separate, scalar, data points in the definition of $\hat{F}_{-i}$. Thus, if window-length W is used, the empirical distribution is based on $(n - 1)W$ data points/bids by other bidders during the last W rounds.

¹² We adopt the definition $F(r) = \Pr(v < r)$ with a strict inequality so that the $\arg \max$ exists.

Theorem 1 is proved in Appendix B. The rapid decay of revenue loss with window-length W suggests that our approach should do well with as few as thousands of items/impressions. We remark that a similar result can be established under the weaker requirement of a *regular* valuation distribution F , for a window-length bounded as¹³ $W \geq C \log(1/\epsilon)/\epsilon^3$.

In Appendix B, we provide a finite horizon version of Theorem 2 (Corollary 1); showing that the revenue loss under our HO-SERP mechanism (using all samples so far) is $O(\sqrt{T} \log T)$ over a horizon of length T for MHR F . We further show that the revenue loss under our mechanism is lower bounded as $\Omega(T^{1/3-\epsilon})$ (Theorem 5) for a standard (exponential) distribution.

Note that the HO-SERP mechanism makes use of all bids by agents other than i , including those that do not clear the agent’s own reserve. However, unlike static settings, using other agents’ bids to determine the payments may not be enough to yield *robust* incentive compatibility. Whereas truthfulness is a best response when agent i ’s valuation $v_{it} < r_{it}$, it is also a best response to submit any other bid $b_{it} \in [0, r_{it})$. In order to make truthfulness the unique best response strategy, we can tweak the mechanism such that with a small probability γ in each round, all the reserve prices are set to 0, i.i.d. across rounds.¹⁴ Agents are told of this tweak, but they do not know if the reserves are zero in the *current* round at the time they submit their bids. This makes truthfulness the unique dominant strategy best response to other agents following the always truthful strategy, a more robust form of incentive compatibility. The loss in expected revenue due to occasionally setting the reserve prices to zero is bounded above as being at most a γ fraction of the benchmark.

A seeming disadvantage of the HO-SERP mechanism is that if the highest bidder’s valuation does not exceed her reserve, the item goes unallocated even though there may be other bidders who meet their reserves. (The reserves may differ from each other due to statistical variation and/or differences in valuation distributions across bidders.) An “eager” variation of SERP is the following: allocate the item to the highest bidder among all the agents whose bids “survive” by being above their personal reserve,¹⁵ and charge her the larger of her personal reserve and the second highest surviving bid. Unfortunately, this variation of the SERP auction creates an incentive to deviate from truthfulness. The intuition is that an agent can benefit from increasing the likelihood that

¹³ In this case the mechanism should compute the so-called “guarded empirical reserve” from the empirical distribution of historical bids, that eliminates the largest bids from consideration as potential reserve prices, see Dhangwatnotai et al. (2015, (12) and Lemma 4.1).

¹⁴ Since F is an MHR distribution, it has positive density everywhere in the support, making truthful bidding the unique myopic best response whenever there are two or more bidders.

¹⁵ These variations are sometimes called *lazy* and *eager* respectively; see Dhangwatnotai et al. (2015), Paes Leme et al. (2016).

a competing agent in eliminated (due to not clearing her personal reserve), and this creates an incentive to overbid so as to raise the personal reserve faced by competing agents in future. The example below illustrates this phenomenon.

Example 2 (“Eager” SERP is not IC) *Let us consider a setting with two agents whose valuations are drawn i.i.d. from uniform distribution $[0, 1]$. The item is allocated as follows: first remove all the agents whose bid is less than their personal reserve, set as per (6). If no agents remain, the item will not be allocated. If only one agent survives, the item will be allocated at a price equal to her personal reserve. If two agents remain, the item will be allocated to the highest agent at a price equal to the maximum of her personal reserve and the other bid.*

Suppose that the first agent is truthful. We present a profitable deviation for the second agent as follows:

$$b_{it} = \begin{cases} v_{it} & 0 \leq v_{it} < \frac{1}{2} \\ 1 & \frac{1}{2} \leq v_{it} \leq 1 \end{cases}$$

Note that the second agent overbids if her valuation is larger than $\frac{1}{2}$ and is truthful otherwise. Hence, the limiting reserve price for the first agent is equal to 1. Therefore, the first agent would be eliminated from all the auctions. In the appendix, we present a family of profitable deviation strategies including the one above and show that the expected per-round utility of the second agent will be increased by $\frac{1}{24}$ under the strategy above. In other words, the second agent can increase her utility by 50% since her utility under the truthful strategy is equal to $\frac{1}{12}$.

In the next section, we will show robustness of the near optimality of the HO-SERP mechanism to small differences in the valuation distributions across agents. In particular, this will imply that revenue losses caused by the item going unallocated even though some agents meet their reserve, are small in expectation when valuation distributions are similar across agents.

5. Robustness to Asymmetry Among Bidders

We have so far assumed that the agents have the same distributions of valuations. In this section, we discuss the robustness of our results with respect to the asymmetry among bidders. We first like to note that when the valuations are heterogenous, then the second price auction, even with optimized personalized reserver prices may not be the optimal static mechanism and revenue maximizing Myerson auction takes a bit more complicated form when the item is allocated to an agent with the highest virtual value (defined below).¹⁶ Nevertheless, we show that when agents have different

¹⁶ See Golrezaei et al. (2018b) for a discussion on challenges of implementing Myerson auction in practical settings.

valuation distributions, the loss of limiting revenue per round of HO-SERP compared to the static Myerson optimal auction can be bounded.

Consider a case with two agents i and j . Suppose that agent i has a higher valuation distribution than that of j , in the sense that the optimal monopoly price for F_i is larger than the optimal monopoly price for F_j . Then, an HO-SERP mechanism with personal reserve prices set as per (6) sets $r_j > r_i$ instead. As a result, losses are incurred due to two reasons. (i) The reserve for each agent is not suitable for the valuation distribution of that agent. (Further, the static Myerson optimal auction allocates to the bidder with the highest *virtual* value,¹⁷ which HO-SERP does not do.) (ii) The fact that the reserve prices for the two agents are different from each other means that the realized pair of valuations in a round could be such that $r_j > v_j > v_i > r_i$. If this occurs, the item is not allocated because the highest valuation agent (agent j) did not clear her reserve, though a different agent (agent i) did clear his reserve. In this section, we will show that if the valuation distributions are not too different from each other, the loss of revenue under our mechanism relative to the static Myerson-optimal mechanism is small (specifically, it is quadratic in the size of the difference between valuation distributions).

Consider a setting with two agents, whose valuation distributions are δ different from each other. (We formally define a notion of distance below.) We claim that the loss in revenue, relative to repeating the Myerson optimal mechanism that knows the valuation distributions, is typically $O(\delta^2)$. The rough reason is that each of the two problems causes a loss of this order. Having a reserve for each agent that is wrong by $O(\delta)$, or, related to this, *not* mapping the reported valuation appropriately to a virtual value, causes a loss of order $O(\delta^2)$, since we are at a distance $O(\delta)$ from the global maximum of a well-behaved optimization problem. The chance that i clears his reserve but the item is not allocated to anyone is bounded by

$$\Pr(r_2 > v_2 > v_1 > r_1) < \Pr(v_1 \in (r_1, r_2) \text{ AND } v_2 \in (r_1, r_2)) = O(\delta) \cdot O(\delta) = O(\delta^2).$$

Hence, this issue also causes a loss of order $O(\delta^2)$.

Let us begin with an example, before we make this rigorous.

Example 3 Suppose that agent 1 has a $\text{Uniform}(0, 1)$ valuation distribution, whereas agent 2 has a $\text{Uniform}(\delta, 1 + \delta)$ valuation distribution for some (small) $\delta > 0$. Then the mechanism we introduced above sets $r_1 = (1 + \delta)/2$ and $r_2 = 1/2$. The expected revenue it earns is

$$\begin{aligned} & E[\text{Revenue from agent 1}] + E[\text{Revenue from agent 2}] \\ &= \frac{5 - 6\delta - 3\delta^3 + 4\delta^3}{24} + \frac{5 + 15\delta}{24} \\ &= \frac{10 + 9\delta - 3\delta^2 + 4\delta^3}{24}. \end{aligned}$$

¹⁷ The virtual value for agent i is $\phi_i(v_i) = v_i - (1 - F_i(v_i))/f_i(v_i)$.

On the other hand, consider the Myerson-optimal mechanism, which uses virtual values $\phi_1(v_1) = 2v_1 - 1$ and $\phi_2(v_2) = 2v_2 - 1 - \delta$, allocating the item to the agent with the highest virtual value, if it is positive, and charges that agent the smallest bid/valuation for which she would still have been awarded the item. This mechanism produces revenue $\frac{10+9\delta+3\delta^2+3\delta^3}{24}$. It follows that the revenue under the Myerson optimal mechanism is $\frac{6\delta^2-\delta^3}{24} = O(\delta^2)$ more than that under our mechanism. For $\delta = 0.1$, the revenue loss is just 0.0025 or 0.54%. For $\delta = 0.2$, the revenue loss is just 0.0097 or 1.9% and even for large $\delta = 0.3$, the revenue loss is 0.021 or 3.9%.

We now formalize this. Let agent i have valuation distribution F_i , which is once again assumed to be MHR, i.e., to have an increasing hazard rate.¹⁸ We define the distance $\|F_i - F_j\|$ between distributions F_i and F_j as

$$\|F_i - F_j\| = \max_v |\phi_i(v) - \phi_j(v)|, \quad (7)$$

where $\phi_i(v) = v - (1 - F_i(v))/f_i(v)$ is the virtual value function.¹⁹

Theorem 2 *Consider a setting with n agents where agent i 's valuation distribution is F_i . Again, any HO-SERP mechanism is periodic ex-post incentive compatible. Suppose that for each agent i , the valuation distribution F_i is MHR and has density bounded above by $f_{\max}$. Suppose also that $\|F_i - F_j\| = \delta$ for all pairs of agents i and j , for some $\delta < \infty$. We have that the limiting average per round revenue under HO-SERP with personal reserve prices set as per (6) is at least $\text{REV}_* - 2(n-1)f_{\max}\delta^2$ as $W \rightarrow \infty$, where REV_* is the expected revenue achieved by the Myerson-optimal static mechanism. Equivalently, HO-SERP with these reserves achieves a fraction $(1 - 2(n-1)f_{\max}\delta^2/\text{REV}_*)$ of the benchmark revenue in the limit $W \rightarrow \infty$.*

Thus, if agent valuation distributions are not too different from each other, our proposed mechanism approximately achieves the benchmark revenue. The proof (see Appendix C) formalizes the intuition above by using Myerson's lemma (Myerson 1981), which says that the expected revenue of a truthful mechanism is equal to the expected virtual value of the winning bidder (defined as zero if the item is not allocated). The revenue-maximizing static mechanism allocates to the bidder with the largest virtual value, if this virtual value is non-negative. We show that our mechanism deviates from this allocation with probability no more than $2(n-1)f_{\max}\delta = O(\delta)$ and further chooses an

¹⁸ We should be able to extend to α -strongly regular distributions (Cole and Roughgarden 2014), where the virtual value functions increase at rate at least α everywhere in the support. The lower bound α on the rate of increase (which is currently 1, for MHR distributions), will be a part of the upper bound on revenue loss.

¹⁹ In fact, in definition (7) we can ignore values of v below $\min(r_1, r_2)$ (the smaller of the Myerson optimal reserve prices for F_i and F_j). Theorem 2 still holds, and the proof is unaffected.

allocation that is within δ of the ideal allocation in terms of virtual value in cases where it allocates wrongly. These bounds then enable us to obtain a $2(n-1)f_{\max}\delta^2 = O(\delta^2)$ bound on the loss in expected revenue.

As an illustration we can apply this result to the setting in Example 3. We have $n = 2$, $\bar{f} = 1$, and $\|F_1 - F_2\| = \delta$, and so we obtain from Theorem 2 that the revenue loss relative to the Myerson benchmark is bounded above by $2\delta^2$. The actual loss turns out to be $\frac{6\delta^2 - \delta^3}{24}$.

We conclude this section with a discussion of a byproduct of our results that could be of independent interest. Hartline and Roughgarden (2009) show that where the valuation of each agent is drawn independently from a different regular distribution, second-price auctions with a personalized reserve obtain a $\frac{1}{2}$ -approximation of the optimal revenue. As a corollary of the analysis leading to our Theorem 2, we obtain a complementary result: that using a second price auction in the asymmetric valuations case, the seller can obtain expected revenue within $O(\delta^2)$ of the optimal, where δ is the maximum “distance” between valuation distributions; see Remark 4 in the Online Appendix for details.

6. Heterogeneous Items

In this section we provide guidance on how heterogeneity between items can be incorporated in our proposed HO-SERP mechanism described in Theorem 1. (A similar approach can be used to extend the SESE mechanism from Section 7 to a heterogeneous items setting.) Our model of valuations may be interpreted as one way to incorporate correlation between agent valuations for an item, (cf. McAfee and Vincent 1992).

We generalize the model in Section 2 as follows. Each item has m attributes, where m is a fixed constant. We denote attributes of the t -th item by $x_t = (x_{t1}, x_{t2}, \dots, x_{tm})^T$, and henceforth call x_t the *context* at period t . We model the valuation v_i of each agent i for the t -th item as

$$v_i = \beta^T x_t + \tilde{v}_i, \quad (8)$$

where $\tilde{v}_i \sim F$ is drawn independently across agents and items, and $\beta \in \mathbb{R}^m$ is the vector of context coefficients (common across agents and items). Thus, the context causes an additive translation in valuations; the amount of translation has a linear functional form in the attributes and is common across agents. We assume that the contexts $(x_t)_{t=1}^\infty$ are drawn i.i.d. from some distribution G . Our technical development in this section draws upon the work of Golrezaei et al. (2018a) on contextual auctions. Two key high level differences from that paper are: (i) Our agents are patient, and hence to obtain good incentive properties, we stay with our proposal to choose a personalized price for agent i based on the past bids of *other* agents. In that paper, agents are impatient, and so the

mechanism is able to set a personalized price for i based on the past bids of agent i herself. (ii) We assume distribution F is time invariant and obtain revenue guarantees for any F in a class $\mathcal{F}$, whereas the other paper solves a robust optimization problem where F can vary arbitrarily over time within $\mathcal{F}$.

Assumptions on F , G and β . We assume F is MHR as before, but in addition assume that F has bounded support $(-B_F, B_F)$ for some $B_F < \infty$. We absorb the mean of distribution F into $\beta^T x_\tau$ (we can include an attribute that always takes the value 1 so that its coefficient will be the intercept which includes the mean of F) and hence assume $\mathbb{E}_{\tilde{v} \sim F}[\tilde{v}] = 0$. (This implies $\mathbb{E}[v_{i'\tau}] = \beta^T x_\tau$.) We assume that distribution G has bounded support; without loss of generality we assume it is supported on $\{x : \|x\| \leq 1\}$ (we use the Euclidean norm throughout). We further assume G has a second moment matrix $\Sigma = \mathbb{E}_{x \sim G}[xx^T]$ which is strictly positive definite with smallest eigenvalue at least $1/B_\Sigma$ for some $B_\Sigma < \infty$. We also assume that $\|\beta\| \leq B_\beta$ for some $B_\beta < \infty$.

Auctioneer's knowledge. The auctioneer observes the context x_t before each period t and knows the bounds B_F , B_Σ and B_β but does not know F , G or β beforehand. As before, the auctioneer wants to set personalized reserve prices to maximize long run average revenue (1) while accounting for the strategic response of the agents who aim to maximize the average per-round utility (2). (All expectations now include expectation over the contexts $x_t \sim G$ i.i.d. across periods t .)

The reserve price function $\Omega : \mathcal{H} \rightarrow \mathbb{R}^n$ maps the history observed by the mechanism up to t including the current context x_t to reserve prices $(r_{it})_{i=1}^n$. The history observed by the mechanism, and by each agent, up to time t now include the contexts in each period so far including period t

$$H_{\Omega,t} = \langle (x_1, r_1, b_1, q_1, p_1), \dots, (x_{t-1}, r_{t-1}, b_{t-1}, q_{t-1}, p_{t-1}), x_t \rangle, \quad (9)$$

$$H_{i,t} = \langle (x_1, v_{i1}, b_{i1}, q_{i1}, p_{i1}), \dots, (x_{t-1}, v_{i,t-1}, b_{i,t-1}, q_{i,t-1}, p_{i,t-1}), x_t \rangle. \quad (10)$$

We include x_t in $H_{i,t}$ to clarify that the agents know the current context x_t before they submit their bids b_t . The auctioneer commits beforehand to Ω . The definition of dynamic incentive compatibility remains as before. As in Theorem 1, we will provide a stationary window-based Ω with good properties and average revenue approaching that under the static Myerson auction. Note that the reserve price of the benchmark static Myerson auction will be dependent on the context, and correspondingly, the reserve price set by our mechanism in period t will account for x_t .

Let $\text{REV}_*(x)$ be, for context $x \in \mathbb{R}^m$, the expected revenue under the Myerson optimal auction for agents with valuations drawn i.i.d. from the contextual valuation distribution F^x given by

$$F^x(v) \triangleq F(v - \beta^T x). \quad (11)$$

In an effort to obtain a revenue close to $\text{REV}_*(x_t)$ in period t , but while retaining incentive compatibility, our proposed mechanism proceeds as follows. Fix the window length W . For every agent i , the mechanism sets the personalized price r_{it} using (i) the current context x_t and (ii) the bids of other agents in the last W periods (treating those bids as truthful) and the contexts in those periods.

To properly account for the context x_t in the choice of reserve prices r_{it} , our mechanism needs to learn the coefficient vector β . We proceed as follows. Recall that the expected valuation of item τ is $E[v_{i'\tau}] = \beta^T x_\tau$ allowing us to treat each period τ bid $b_{i'\tau} = v_{i'\tau}$ by agent $i' \neq i$ as a noisy observation of $\beta^T x_\tau$, corrupted by zero mean “noise” $\tilde{v}_{i'\tau} \sim F$ that is i.i.d. across agents and periods. We use these observed bids to obtain an ordinary least squares (OLS) estimate of β

$$\hat{\beta}_{-i} \triangleq \arg \min_{\tilde{\beta}: \|\tilde{\beta}\| \leq B_\beta} \mathcal{L}_{-i}(\tilde{\beta}), \quad \text{where } \mathcal{L}_{-i}(\tilde{\beta}) \triangleq \frac{1}{(n-1)W} \sum_{i' \neq i} \sum_{\tau=t-W}^{t-1} (b_{i'\tau} - \tilde{\beta}^T x_\tau)^2.$$

This estimate converges rapidly to the true β .

Lemma 1 *Fix the constants m , B_F , B_Σ and B_β . There exists a constant $C_1 = C_1(m, B_F, B_\Sigma, B_\beta) < \infty$ such that, for any F , G and β (satisfying B_F , B_Σ and B_β respectively), for any window length $W > 1$, and each agent i , the estimated coefficients are close to the true ones: with probability $1 - 1/W$ we have*

$$\|\hat{\beta}_{-i} - \beta\| \leq C_1 \sqrt{\frac{\log W}{W}}. \quad (12)$$

We then deploy this estimate to “translate” the past bids to the current context x_t : a bid of $b_{i'\tau}$ submitted under context x_τ maps to the translated bid $\tilde{b}_{i'\tau, -i} \triangleq b_{i'\tau} + \hat{\beta}_{-i}^T (x_t - x_\tau)$. The empirical distribution $\hat{F}_{-i}^{x_t}$ of translated bids serves as an estimate for the true contextual distribution

$$F^{x_t}(v) \triangleq F(v - \beta^T x_t). \quad (13)$$

We need to be careful here because our estimate of $\hat{F}_{-i}^{x_t}(v)$ is imperfect for two reasons: First, as in Section 4, it is based on a finite number of samples. Second, and this is an issue we did not encounter before, our estimate $\hat{\beta}_{-i}$ is imperfect. As a result, the samples upon which $\hat{F}_{-i}^{x_t}$ is based are not drawn from F^{x_t} itself: instead, the samples based on bids in period τ correspond to sampling from F^{x_t} but then adding $(\hat{\beta}_{-i} - \beta)^T (x_t - x_\tau)$ to the realization. To ensure that this additional source of error does not inadvertently lead to a large reduction in the probability of selling (e.g., this could happen if F^{x_t} has an atom at the Myerson optimal reserve), our mechanism

sets the price by making a small reduction to the estimated optimal reserve price. Accordingly, we set the personalized reserve price as per the following modification of (6)

$$r_{it} = -\delta + \arg \max_r r(1 - \hat{F}_{-i}^{x_t}(r)). \quad (14)$$

Here, we set $\delta = 2C_1\sqrt{\log W/W}$, where C_1 is the constant in Lemma 1. Informally, this is to ensure that with probability $1 - 1/W$, errors in bid translation do not cause us to unintentionally price out an agent. As a result of this adjustment to the mechanism, we now obtain an additive approximation to the revenue instead of a multiplicative approximation.²⁰

Theorem 3 *Consider the setting with item attributes described above, with constants n, m, B_F, B_Σ and B_β . Any HO-SERP mechanism is periodic ex-post incentive compatible. In particular, all agents following the always truthful strategy constitutes an equilibrium. Further, there exists $C = C(n, m, B_F, B_\Sigma, B_\beta) < \infty$, such that for any F that is MHR, G and β (satisfying B_F, B_Σ and B_β respectively), any $\epsilon \in (0, 1)$ and any context $x_t : \|x_t\| \leq 1$, the HO-SERP mechanism with window length $W \geq C \log(1/\epsilon)/\epsilon^2$ and personal reserve prices set as per (14) achieves expected²¹ revenue in period t that is at least $\text{REV}_*(x_t) - \epsilon$, where $\text{REV}_*(x_t)$ is the expected revenue under the optimal static mechanism (second-price auction with Myerson-optimal reserve) for the true bid distribution F^{x_t} given by (13).*

The proofs for this section are presented in Appendix D.

7. Incentive-Compatible Surplus Extraction

Although the second-price auction can be revenue-maximizing in static settings, it may not be the optimal mechanism in dynamic environments. To convey intuition, let us first consider a setting with n agents and a horizon of length T where the seller *knows the distribution of the valuations* of agents. Consider the following mechanism: (i) The mechanism charges each agent i an *up-front payment* equal to $\sum_{t=1}^T \mathbb{E}[u_{it}]$, where u_{it} denotes the random variable correspond to the utility of agent i at time t ; namely,

$$u_{it} = \max\{v_{it} - \max_{j \neq i} \{b_{jt}\}, 0\}. \quad (15)$$

²⁰ Note that a multiplicative approximation would be a stronger result — given our boundedness assumptions a multiplicative approximation implies an additive approximation but not vice versa. However, as a result of estimation errors in learning β we only obtain an additive approximation here.

²¹ The expectation is over the past contexts, past valuations and period t valuations.

The expectation is calculated assuming that all agents are truthful. (ii) The mechanism runs a second-price auction (with no reserve) in each of the T rounds. Notice that Eq. (15) is consistent with this design.

Note that by using the up-front payments, the mechanism extracts the whole surplus of the buyers and obtains average revenue of $E[\max_j \{v_{jt}\}]$. Assuming only individual rationality on the part of agents, this is the maximum achievable average revenue per round for any mechanism. This mechanism, although revenue optimal, is not directly applicable to the current online ad markets because it charges an up-front payment; see Mirrokni and Nazerzadeh (2017).²² However, ignoring this practical consideration, we show how the ideas above can be used to design an essentially optimal mechanism in our setting.

The above surplus-extracting mechanism can also be implemented as follows (when the distribution of the valuations, F , is known); cf. Arrow (1979), d’Aspremont and Gérard-Varet (1979), Baron and Besanko (1984),      and Szentes (2007). At each round t , the mechanism charges an entrance fee of

$$\mu_i = E_F[u_i] = E_F[\max\{v_{it} - \max_{j \neq i} \{v_{jt}\}, 0\}]. \quad (16)$$

The agent may accept the entrance fee. Agents who pay the entrance fee, then learn their valuation v_i and can bid in the auction. The item is allocated via a second-price auction with no reserve and therefore the agents will bid truthfully. Note that at the equilibrium the agents are indifferent between participating or leaving but the mechanism can always nudge the agents to participate by slightly reducing the entrance fee. Building on these ideas, we propose the following mechanism.

Surplus-Extracting-Self-Excluding (SESE) Mechanism. The mechanism consists of two phases.

- In the first phase, which lasts for N rounds (where N is a parameter chosen by the seller), the item is allocated via a second-price auction with no reserve. At the end of the first phase, for each agent i , define $\hat{\mu}_i$ as follows:

$$\hat{\mu}_i = \frac{1}{n} \frac{1}{N/2} \sum_{k=1}^{N/2} z_k \quad (17)$$

and z_k ’s for $1 \leq k \leq N/2$ are constructed as follows. We repeatedly sample *without replacement* n bids from the set of bids in the first phase from all the bidders except agent i . Let Z_k be the k -th sampled set and let z_k be the difference between the highest and the second-highest bid in Z_k . Note that since $n \geq 2$, the total number of sampled bids is $nN/2 \leq (n-1)N$, ensuring feasibility.

²² Reservation (guaranteed-delivery) contracts for selling display advertising specify the number of impressions to be allocated under the contract in advance. The allocation is determined by the publisher and not via an auction.

- In each round $t > N$ in the second phase, the seller offers an entrance fee of $\left(\hat{\mu}_i - \sqrt{\frac{2 \log N}{N}}\right)$ to agent i . Note that the entrance fee is determined using the *other* agents' bids in the first phase.

The item is allocated using a second-price auction with no reserve. Let S be the set of agents who pay the entrance fee (and subsequently learn their valuation v_{it}) and let $\bar{S}$ represent the set of agents who refuse to participate in this round. The mechanism *simulates* the agents in $\bar{S}$. More specifically, the mechanism randomly chooses a round $\tau < N$ and uses the bids in that round for each agent $j \in \bar{S}$. At time t , if a simulated bid is the highest, the item will not be allocated. Otherwise, it will go to a highest bid at the price equal to the second-highest bid among agents in S and $\bar{S}$.

Here is the intuition behind the mechanism. Observe that by the above definition, when all the agents are truthful and have the same valuation distribution, we have

$$\mathbb{E}[z_k] = \mathbb{E}\left[\sum_{i=1}^n u_{it}\right] = \sum_{i=1}^n \mu_i = n\mu_i, \quad (18)$$

where μ_i denotes $\mathbb{E}[u_{it}]$ for agent i ; see Eq. (16). Hence we have $\mathbb{E}[\hat{\mu}_i] = \mu_i$.

Our mechanism achieves (approximate) incentive compatibility by leveraging the same two key ideas that led to Theorems 1 and 3: i) The entrance fee charged to each agent (in the second phase) depends only on the bids of the *other* agents in the first phase; thus an agent's bids do not affect the entrance fee that agent herself faces. We further deduce that the agents bid truthfully in the second phase, since their bids have no future impact whatsoever. Hence, they would pay the entrance fee if²³ $\mathbb{E}[u_i] \geq \hat{\mu}_i - \sqrt{\frac{2 \log N}{N}}$. ii) Using simulated bids, we bound the gain from overbidding for the agents: note that the bids of the agents in the first phase can influence the outcomes in the second phase. More specifically, agents can overbid and inflate the entrance fee of other agents, which may result in the latter's refusal to participate in the auctions in the second phase. Our mechanism that simulates non-participating agents' bids significantly lessens the benefit that may be obtained from such deviations.

Note that our mechanism that simulates non-participating agents does not entirely eliminate the incentive to deviate. For example, suppose that there are two agents and during the first (learning) phase, the first agent's bids are lower than usual. In this case, the second agent may prefer to compete against the "simulated version" of the first agent, and can ensure this by overbidding to force the first agent out of the auction. In addition, an agent may be eliminated by mistake.

²³ To simplify the presentation, we assume that the agents know the distribution of valuations, since agents may learn the distributions over time. Note that incentive compatibility clearly continues to hold even if agents do not know distributions over valuations.

Revisiting the scenario with two agents, suppose in the first phase the first agent's bids are higher than usual. This may result in a high entrance fee for the second bidder and may lead to elimination of the second bidder from all the subsequent auctions. We include a small slack in the chosen entrance fees to ensure that the likelihood of such mistaken elimination is small.

We can now state the main result of this section. Note that we do not need F to be an MHR or even a regular distribution. A bounded support suffices; any other conditions under which a Hoeffding-type bound holds uniformly would serve just as well.

Theorem 4 (Surplus-extracting mechanism) *Suppose that the valuations of all agents are drawn i.i.d. from distribution F over $[0, 1]$. Distribution F is a-priori unknown to the seller but it is known to the agents. If all the agents are truthful, the Surplus-Extracting-Self-Excluding mechanism with an exploration phase of length N obtains an expected per-auction revenue²⁴ of $E[\max_j \{v_{jt}\}] - O(\sqrt{\log N/N})$.*

In addition, under this mechanism, for any agent i and time t , if all the other agents are always truthful, then with probability $1 - O(N^{-2})$, the increase in per-auction utility that can be obtained by deviating from the truthful strategy is bounded by $O(\sqrt{\log N/N})$.

Note that the loss decreases as the length of the first phase increases. However, the mechanism loses revenue in the first phase. The above theorem shows that SESE is approximately incentive-compatible. In the proof, presented in Appendix E, we show that for any strategy B and every period τ , with probability $1 - O(N^{-2})$ when all agents are always truthful, the personal history $H_{i,\tau}$ seen by agent i so far is such that

$$\underline{U}_{i,H_{i,\tau}}(B_i^{\text{Tr}}, B_{-i}^{\text{Tr}}) + O(\sqrt{\log N/N}) \geq \bar{U}_{i,H_{i,\tau}}(B_i, B_{-i}^{\text{Tr}});$$

see (3) and (4). With the remaining probability, $O(N^{-2})$, the benefit from deviating might be larger but is nevertheless bounded by 1. Hence, the expected benefit of deviating from truthfulness is $O(\sqrt{\log N/N})$. In other words, truthfulness is an approximate best-response to the other agents being always truthful. The notion of approximate incentive compatibility implies that agents do not deviate from the truthful strategy when the benefit from such a deviation is insignificant. The notion of approximate incentive compatibility is appealing when characterizing or computing the best response strategy is challenging, and several works moreover use an additive notion of approximate IC similar to ours (Schummer 2004, McSherry and Talwar 2007, Daskalakis et al. 2009,

²⁴ Note that the limiting revenue (1) as well as the limiting per round utility (2) are well defined under SESE when agents are always truthful.

Nazerzadeh et al. 2013). In online ad auctions, finding profitable deviation strategies requires solving complicated dynamic programs in a highly uncertain environment. Thus, agents can plausibly be expected to bid truthfully under an approximately incentive-compatible mechanism.

We remark that our notion of approximate incentive-compatibility is additive in the sense that the absolute increase in utility from a deviation is small. An alternative definition would be multiplicative approximate incentive compatibility where the relative gain from a deviation is small. Note that these two notions differ when the utility of a bidder is small (close to zero). We believe that the additive notion may be a reasonable model for practical environments since figuring out profitable deviation strategies is costly a fractional increase of a gain which was small to begin with may not justify the required investment by the firm (bidder).²⁵

The first and second phases can be interpreted as exploration and exploitation phases, respectively. In an environment where valuations may change slightly over time, the seller can continue to explore occasionally in order to adjust for the change in valuations. For instance, with a small probability, any round $t > N$ can be designated an exploration round and the entrance fees can be set to zero. Stale exploration data can be discarded as new data is generated. (This will also ensure that the long-run average revenue converges to the ex-ante expected value with probability 1.)

8. Conclusion

Designing data-driven incentive-compatible mechanisms has become an important research agenda, motivated in part by the rapid growth of online marketplaces. In this work, we showed that the revenue of repeated auctions can be optimized when the valuations of each bidder can be estimated from the valuations of other bidders. The main goal of the paper was to study the tension between learning and incentive properties. The model is set up to study the hardest case of this tension, namely when all the bidders participate in all the auctions. If some bidders do not participate in an auction, their previous bids can be used for learning and setting prices without causing any incentive issues, in addition to previous bids by bidders who are participating. Even though we have not explicitly modeled participation, our results would extend to such environments since we proposed mechanisms based on the following two principles: (i) the personalized price for each agent should be based only on the historical bids of other agents, and (ii) an agent should not benefit from preventing other agents from participating by raising the prices they face.

We showed that our work can be practically useful by showing that there is only a small revenue loss in case of limited heterogeneity in bidder valuation distributions, and by extending our ideas to

²⁵ Technically, the mechanism can share some of the surplus with the bidders, hence maintaining both notions of approximate IC.

a contextual setting with heterogeneous items that allows for correlation between valuations of buyers. A natural research direction is to explore the optimal tradeoff between incentive compatibility and learning, as a function of heterogeneity among bidders, cf. Golrezaei et al. (2018b). Another interesting direction would be the case where the auctions are connected via budget constraints, cf., Balseiro and Gur (2019).

Furthermore, we believe that the ideas developed here can be applied to other repeated auctions mechanisms that were designed under the assumption that the valuation distributions are known. For instance, Balseiro et al. (2016) propose a repeated auction mechanism that is a hybrid of first-price and second-price auctions, and can extract almost the entire surplus of the buyers. We believe that similar incentive-compatible approximately surplus-extracting mechanisms can be constructed for an *unknown* distribution using our approach.

Acknowledgments. We would like to thank Tim Roughgarden and anonymous referees for their insightful comments and suggestions, along with seminar participants at the MDML workshop at WWW 2019, the INFORMS Annual meeting 2018 and 2019, Google, ACM EC 2017 and the Marketplace Innovation workshop 2017. This work was supported in part by Microsoft Research New England. The work of the second author was supported by a Google Faculty Research Award.

References

- Alessandro Acquisti and Hal R. Varian. Conditioning prices on purchase history. *Marketing Science*, 24(3):367–381, May 2005.
- Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Learning prices for repeated auctions with strategic buyers. In Christopher J. C. Burges, Léon Bottou, Zoubin Ghahramani, and Kilian Q. Weinberger, editors, *NIPS*, pages 1169–1177, 2013.
- Kenneth Arrow. *The property rights doctrine and demand revelation under incomplete information*. Economics and human welfare. New York Academic Press, 1979.
- Yossi Aviv and Amit Pazgal. Optimal Pricing of Seasonal Products in the Presence of Forward-looking Consumers. 10(3):339–359, 2008. ISSN 1526-5498.
- Yossi Aviv, Mingcheng Wei, and Fuqiang Zhang. Responsive pricing of fashion products: The effects of demand learning and strategic consumer behavior. 2013.
- Maria-Florina Balcan, Avrim Blum, Jason D Hartline, and Yishay Mansour. Reducing mechanism design to algorithm design via machine learning. *Journal of Computer and System Sciences*, 74(8):1245–1270, 2008.
- Sandeep Baliga and Rakesh Vohra. Market research and market design. *Advances in Theoretical Economics*, 3(1):1059–1059, 2003.

-
- Santiago Balseiro and Yonatan Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 2019.
- Santiago Balseiro, Vahab Mirrokni, and Renato Paes Leme. Dynamic mechanisms with martingale utilities. *Available at SSRN (2821261)*, 2016.
- David P. Baron and David Besanko. Regulation and information in a continuing relationship. *Information Economics and Policy*, 1(3):267–302, 1984.
- Dirk Bergemann and Johannes Horner. Should auctions be transparent? *Available at SSRN (1657947)*, 2010.
- Dirk Bergemann and Juuso Välimäki. The dynamic pivot mechanism. *Econometrica*, 78:771–789, 2010.
- Omar Besbes and Assaf Zeevi. Dynamic pricing without knowing the demand function: risk bounds and near-optimal algorithms. *Operations Research*, 57:1407–1420, 2009.
- Sushil Bikhchandani and Kevin McCardle. Behavior-based price discrimination by a patient seller. *B.E. Journals of Theoretical Economics*, 12, June 2012.
- Patrick Billingsley. *Convergence of probability measures*. John Wiley & Sons, 2013.
- Nicolò Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. Regret minimization for reserve prices in second-price auctions. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1190–1204, 2013.
- Nicolo Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. Regret minimization for reserve prices in second-price auctions. *IEEE Transactions on Information Theory*, 61(1):549–564, 2015.
- Xi Chen and Zizhuo Wang. Bayesian dynamic learning and pricing with strategic customers. *Available at SSRN (2715730)*, 2016.
- Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 243–252. ACM, 2014.
- Vincent Conitzer, Curtis R. Taylor, and Liad Wagman. Hide and seek: Costly consumer privacy in a market with repeat purchases. *Marketing Science*, 31(2):277–292, March 2012. ISSN 1526-548X.
- José Correa, Ricardo Montoya, and Charles Thraves. Contingent preannounced pricing policies with strategic consumers. *Operations Research*, 64(1):251–272, 2016.
- Constantinos Daskalakis, Aranyak Mehta, and Christos H. Papadimitriou. A note on approximate nash equilibria. *Theor. Comput. Sci.*, 410(17):1581–1588, 2009.
- Claude d’Aspremont and Louis-André Gérard-Varet. Incentives and incomplete information. *Journal of Public economics*, 11(1):25–45, 1979.
- Sriram Dasu and Chunyang Tong. Dynamic pricing when consumers are strategic: Analysis of posted and contingent pricing schemes. *European Journal of Operational Research*, 204(3):662–671, August 2010.
- Arnoud V den Boer. Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in operations research and management science*, 20(1):1–18, 2015.

-
- Nikhil R. Devanur, Yuval Peres, and Balasubramanian Sivan. Perfect bayesian equilibria in repeated sales. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 983–1002. SIAM, 2014.
- Nikhil R. Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. The sample complexity of auctions with side information. In *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing*, STOC '16, pages 426–439, 2016.
- Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, 2015.
- Alessandro Epasto, Mohammad Mahdian, Vahab Mirrokni, and Song Zuo. Incentive-aware learning for large markets. In *Proceedings of the 2018 World Wide Web Conference*, pages 1369–1378. International World Wide Web Conferences Steering Committee, 2018.
- Peter   so and Balazs Szentes. Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies*, 74(3):705–731, 2007.
- Rosa Branca Esteves. A survey on the economics of behaviour-based price discrimination. NIPE Working Papers 5/2009, NIPE - Universidade do Minho, 2009.
- Drew Fudenberg and J. Miguel Villas-Boas. *Behavior-Based Price Discrimination and Customer Recognition*. Elsevier Science, Oxford, 2007.
- Andrew V Goldberg, Jason D Hartline, and Andrew Wright. Competitive auctions and digital goods. In *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 735–744. Society for Industrial and Applied Mathematics, 2001.
- Negin Golrezaei, Adel Javanmard, and Vahab Mirrokni. Dynamic incentive-aware learning: Robust pricing in contextual auctions. 2018a.
- Negin Golrezaei, Max Lin, Vahab Mirrokni, and Hamid Nazerzadeh. Boosted second-price auctions for heterogeneous bidders. *Working paper*, 2018b.
- J. Michael Harrison, N. Bora Keskin, and Assaf Zeevi. Bayesian dynamic pricing policies: Learning and earning under a binary prior distribution. *Management Science*, 58(3):570–586, 2012.
- Oliver D. Hart and Jean Tirole. Contract renegotiation and coasian dynamics. *Review of Economic Studies*, 55:509–540, 1988.
- Jason Hartline, Vahab Mirrokni, and Mukund Sundararajan. Optimal marketing strategies over social networks. In *Proceedings of the 17th International Conference on World Wide Web*, WWW '08, pages 189–198, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-085-2. doi: 10.1145/1367497.1367524. URL <http://doi.acm.org/10.1145/1367497.1367524>.
- Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, pages 225–234, 2009.

-
- Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
- Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. Making the most of your samples. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 45–60. ACM, 2015.
- Preston McAfee. The design of advertising exchanges. *Review of Industrial Organization*, 39(3):169–185, 2011.
- R Preston McAfee and Daniel Vincent. Updating the reserve price in common-value auctions. *American Economic Review*, 82(2):512–18, May 1992.
- Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS’07. 48th Annual IEEE Symposium on*, pages 94–103. IEEE, 2007.
- Vahab Mirrokni and Hamid Nazerzadeh. Deals or no deals: Contract design for online advertising. In *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 2017.
- Mehryar Mohri and Andres Muñoz Medina. Revenue optimization in posted-price auctions with strategic buyers. *NIPS*, 2014a.
- Mehryar Mohri and Andres Muñoz Medina. Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *Proceedings of the 31th International Conference on Machine Learning, ICML*, pages 262–270, 2014b.
- S. Muthukrishnan. Ad exchanges: Research issues. In *Internet and Network Economics, 5th International Workshop (WINE)*, pages 1–12, 2009.
- Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- Hamid Nazerzadeh, Amin Saberi, and Rakesh Vohra. Dynamic cost-per-action mechanisms and applications to online advertising. *Operations Research*, 61(1):98–111, 2013.
- Hamid Nazerzadeh, Renato Paes Leme, Afshin Rostamizadeh, and Umar Syed. Where to sell: Simulating auctions from learning algorithms. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC ’16, pages 597–598, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-3936-0.
- Michael Ostrovsky and Michael Schwarz. Reserve prices in internet advertising auctions: A field experiment. ACM EC 2011, SSRN (1573947), 2009.
- Renato Paes Leme, Martin Pál, and Sergei Vassilvitskii. A field guide to personalized reserve prices. In *Proceedings of the 25th International Conference on World Wide Web*, pages 1093–1102. International World Wide Web Conferences Steering Committee, 2016.
- John G. Riley and William F. Samuelson. Optimal auctions. *American Economic Review*, 71(3):381–392, 1981.
- Tim Roughgarden and Joshua R Wang. Minimizing regret with multiple reserves. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 601–616. ACM, 2016.

- Stephen W Salant. When is inducing self-selection suboptimal for a monopolist? *The Quarterly Journal of Economics*, 104(2):391–97, May 1989.
- James Schummer. Almost-dominant strategy implementation: exchange economies. *Games and Economic Behavior*, 48(1):154–170, 2004.
- Ilya Segal. Optimal pricing mechanisms with unknown demand. *The American Economic Review*, 93(3):509–529, 2003.
- Nancy L Stokey. Intertemporal price discrimination. *The Quarterly Journal of Economics*, 93(3):355–71, August 1979.
- Curtis R. Taylor. Consumer privacy and the market for customer information. *RAND Journal of Economics*, 35(4):631–650, Winter 2004.
- J. Miguel Villas-Boas. Price cycles in markets with customer recognition. *RAND Journal of Economics*, 35(3):486–501, Autumn 2004.
- Zizhuo Wang, Shiming Deng, and Yinyu Ye. Close the gaps: A learning-while-doing algorithm for single-product revenue management problems. *Operations Research*, 62(2):318–331, 2014.

Appendix

A. Appendix to Section 3

Proof of Proposition 1. All agents other than agent i are always truthful. Consider the (myopic) loss in utility L in the current round experienced by agent i when she follows b_ϵ instead of bidding truthfully. This loss occurs in cases where i would have won the current round if she had been truthful, but loses under b_ϵ . The approach we adopt is to bound the probability density of L for positive values. Reveal valuations v_j for all bidders $j \neq i$. This fixes the largest valuation/bid $v_{-i}^{\max} = \max_{j \neq i} v_j$ among agents other than i . In order for agent i to lose L , it must be that the valuation $v_i = L + v_{-i}^{\max}$, and that $b_\epsilon(v_i) \leq v_{-i}^{\max} \Rightarrow v_i - v_{-i}^{\max} \leq d(b_\epsilon) \Rightarrow L \leq \epsilon$. This yields an upper bound of $f_{\max}$ on the probability density of L in the interval $[0, \epsilon]$ and an upper bound of 0 on the density for values exceeding ϵ . It follows that $E[L] \leq \int_0^\epsilon x f_{\max} dx = f_{\max} \epsilon^2 / 2$.

Next, consider the gain in utility for agent i resulting from the lowering of the reserve. Now, since b_ϵ has the $(C\epsilon)$ -reserve impact property, we have that the reserve is at most $r(\text{TR}) - C\epsilon$ when agent i uses b_ϵ . Since the ROF has the $(C\epsilon, \epsilon, q)$ -reserve matters property, and using symmetry of the valuation distribution across agents, we have that with probability at least q/n , agent i has valuation exceeding $r(\text{TR}) + \epsilon$ and all other agents have valuations below $r(\text{TR}) - C\epsilon$. Moreover, the bid of agent i is at least $r(\text{TR}) + \epsilon - d(b_\epsilon) \geq r(\text{TR})$, and hence agent i wins the item and pays at most $r(\text{TR}) - C\epsilon$. Thus, with probability at least q/n , agent i gains at least $C\epsilon$ in utility.

We deduce that the net gain in expected utility for agent i by following b_ϵ is at least

$$qC\epsilon/n - f_{\max}\epsilon^2/2 \geq qC\epsilon/(2n)$$

for $\epsilon \leq qC/(nf_{\max})$.

Notice that in any realization such that the winning bidder has valuation exceeding $r(\text{Tr}) + \epsilon$ whereas other valuations are all below $r(\text{Tr}) - C\epsilon$, the seller loses at least $C\epsilon$ in revenue from the lower reserve price. Since such a realization occurs with probability at least q , we deduce that the expected decrease in revenue resulting from agent i following b_ϵ is at least $qC\epsilon$. $\square$

A.1. Appendix to Example 1

We first show that if one of the buyers follows a shading strategy for values of $\underline{r} \leq \frac{1}{2}$ and $\bar{r} \leq \frac{2}{3}$ to be determined later, while the other buyer is truthful, then the reserve price converges to $\underline{r}$.

Suppose the first bidder decides to shade her bid. Since buyers are playing time-invariant strategies, F_t would converge to a distribution we call $\hat{F}$. The seller chooses the revenue-maximizing reserve price as per Eq. (5). Let $g(r)$ be the expected revenue obtained from reserve r . Note that the reserve price will converge to $\underline{r}$ if $g(\underline{r}) = g(\bar{r})$, based on our decision to break ties by picking the smallest revenue-maximizing price.

For $r \in (\underline{r}, \bar{r}]$, the seller does not obtain any revenue if the valuation of the first buyer lies in $(\underline{r}, \bar{r}]$, whereas the second agent has valuation less than r . Therefore, we have

$$\begin{aligned} g(r) &= \bar{r}(1-r) \times r + (1-\bar{r})r \times r + (1-\bar{r})(1-\bar{r}) \times (\bar{r} + (1-\bar{r})/3) + (1-\bar{r})(\bar{r}-r) \times ((r+\bar{r})/2) \\ &= (\bar{r} + r - 2\bar{r}r)r + (1-\bar{r})^2((1+2\bar{r})/3) + (1-\bar{r})(\bar{r}^2 - r^2)/2. \end{aligned} \quad (19)$$

Taking the derivative with respect to r , we have

$$g'(r) = \bar{r} + 2(1-2\bar{r})r - (1-\bar{r})r = \bar{r} + (1-3\bar{r})r \geq \bar{r} + (1-3\bar{r})\bar{r} = (2-3\bar{r})\bar{r},$$

which is non-negative for all $r \in (\underline{r}, \bar{r}]$ for $\bar{r} \leq \frac{2}{3}$. Therefore, g takes its maximum in $(\underline{r}, \bar{r}]$ at $\bar{r}$ and this maximum value is

$$g(\bar{r}) = 2(1-\bar{r})\bar{r}^2 + (1-\bar{r})^2((1+2\bar{r})/3).$$

For $r = \underline{r}$, we have

$$g(\underline{r}) = (\bar{r} + \underline{r} - 2\bar{r}\underline{r})\underline{r} + (1-\bar{r})^2((1+2\bar{r})/3) + (1-\bar{r})(\bar{r}^2 - \underline{r}^2)/2 + (\bar{r} - \underline{r})\underline{r}^2.$$

The last term corresponds to the revenue obtained from the first buyer from realizations where that buyer bids exactly at reserve $r = \underline{r}$ (note that g is discontinuous at $\underline{r}$; this term is absent for $r > \underline{r}$ as per (19)).

We now show that for any $\underline{r} \in [0, 1/2)$, there exists $\bar{r}$ such that $g(\underline{r}) = g(\bar{r})$. To this end, let us define h for $z \in (\underline{r}, 1]$, which, fixing $\underline{r}$, is equal to the difference between $g(\underline{r})$ and $g(z = \bar{r})$. More precisely,

$$\begin{aligned} h_{\underline{r}}(z) &\triangleq (z + \underline{r} - 2z\underline{r})\underline{r} + (1-z)^2((1+2z)/3) + (1-z)(z^2 - \underline{r}^2)/2 + (z - \underline{r})\underline{r}^2 \\ &\quad - 2(1-z)z^2 - (1-z)^2((1+2z)/3) \\ &= (z + \underline{r} - 2z\underline{r})\underline{r} - (1-z)(3z^2 + \underline{r}^2)/2 + (z - \underline{r})\underline{r}^2. \end{aligned} \quad (20)$$

Plugging $z = \frac{1}{2}$ into Eq. (20), we get

$$h_{\underline{r}}(1/2) = \underline{r}/2 - (3/4 + \underline{r}^2)/4 + (1/2 - \underline{r})\underline{r}^2,$$

which is negative for $\underline{r} < 1/2$. (To see why $h_{\underline{r}}(1/2) < 0$, observe that $h_0(1/2) = -3/16 < 0$, $h_{1/2}(1/2) = 0$ and $dh_{\underline{r}}(1/2)/d\underline{r} = (1/2 - \underline{r})(1 + 3\underline{r}) \geq 0$ for $\underline{r} \in [0, 1/2]$.) On the other hand, note that $h_{\underline{r}}(z = 1) = (1 - \underline{r})(\underline{r} + \underline{r}^2) > 0$. Therefore, for some value of $z \in (\underline{r}, 1)$, $h_{\underline{r}}(z)$ is equal to 0. Let $\bar{r}$ be the lowest such value.

Note that we should have $\bar{r} \leq \frac{2}{3}$. Hence, solving for $h_{\underline{r}}(2/3) = 0$, we find that $\underline{r} = 0.379$. For these parameters, we have $g(\underline{r}) = 0.3827$.

Now suppose that, using shading strategy $(\underline{r}, \bar{r})$, the reserve is reduced to $\underline{r}$. Then, we can write the utility of the buyer as follows:

$$\begin{aligned} U(\underline{r}, \bar{r}) &= \int_{\underline{r}}^1 (v - \underline{r}) \underline{r} dv + \int_{\bar{r}}^1 \int_{\underline{r}}^v (v - y) dy dv \\ &= \frac{1}{2}(1 - \underline{r})^2 \underline{r} + \int_{\bar{r}}^1 \frac{1}{2}(v - \underline{r})^2 dv \\ &= \underline{r}(1 - \underline{r})^2/2 + (1 - \underline{r})^3/6 - (\bar{r} - \underline{r})^3/6. \end{aligned}$$

Note that the truthful strategy corresponds to $\underline{r} = \bar{r} = \frac{1}{2}$. Plugging in the numbers, we have $U(1/2, 1/2) = 0.08333$ and $U(0.379, 2/3) = 0.1090$.

A.2. General negative result for learning a common reserve price

We now argue that for a broad class of (window-based) stationary mechanisms under which the reserve changes if the bidding behavior of a particular bidder changes, truthful bidding does *not* constitute an approximate equilibrium. For simplicity, we consider sequences of ROFs with increasing window-length W such that when agents use fixed bid functions, as $W \rightarrow \infty$, the resulting reserve has a limiting value for $t \geq W$. (Note that the distribution of the reserve price is time-invariant for all $t \geq W$ since the ROF is time-invariant and agents are using fixed bid functions.)

Consider the perspective of agent i . Suppose that other bidders are following the always truthful strategy, i.e., bidding truthfully irrespective of the information they receive over time. Agent i considers what *bid function* she should use to map from valuations v to corresponding bids $b(v) \leq v$, i.e., she does not bid higher than her true valuation. (She could use some more complex strategy, that reacts to the current reserve price, her history, etc., and/or overbids, but we are looking for a profitable deviation in this space of simple, “shading-only” strategies.) We define the *magnitude of shading* of a bid function as $d(b) = \sup_v (v - b(v))$. Let $f_{\max} = \sup_v f(v)$.

We say that a bid function involving bid-shading, i.e., $b(v) \leq v$ for all v , satisfies the δ -reserve impact property for agent i if the following holds. Let $r(\text{Tr})$ be the limiting reserve price when agent i employs the always truthful strategy (which corresponds to the identity bid function), and let $r(b)$ be the limiting reserve price when agent i employs the bid function b always. Then $r(b) \leq r(\text{Tr}) - \delta$.

We say that an ROF satisfies the (δ, ϵ, q) -reserve matters property (for a particular F , n) if the following holds. With probability at least q in any round, we have that (i) the highest valuation is at least $r(\text{Tr}) + \max(\epsilon - \delta, 0)$, and (ii) each of the other agents has a valuation below $r(\text{Tr}) - \delta$. Intuitively, an agent can hope to benefit from a lower reserve price (with likelihood q/n , by an amount δ) if an ROF has the reserve matters property, even if the agent is shading her bid by up to ϵ .

Proposition 1 *Fix the ROF and suppose that all agents other than i are always truthful. Suppose that there exist $C \in (0, 1]$, $q \in (0, 1]$ and $\epsilon \in (0, qC/(nf_{\max}))$ and a bid function b_ϵ such that the following conditions hold:*

- *The bid function b_ϵ satisfies $b_\epsilon(v) \leq v$ for all v and $d(b_\epsilon) \leq \epsilon$ as well as the $(C\epsilon)$ -reserve impact property.*
- *The ROF satisfies the $(C\epsilon, \epsilon, q)$ -reserve matters property.*

Then, the steady-state average utility per auction for agent i exceeds that resulting from always being truthful by at least $qC\epsilon/(2n)$.

Further, if agent i adopts the non-truthful strategy b_ϵ , this causes a decrease in revenue per auction of at least²⁶ $qC\epsilon$.

Despite the technical nature of its statement, this result says something powerful and general. Whenever an agent, by using a bid function that involves shading by no more than (small) ϵ , can cause the seller to reduce the reserve price by order ϵ , and the reserve price matters, this is a beneficial deviation for the buyer and it causes revenue loss to the seller. The reason the buyer benefits from deviating is similar to that captured in Remark 2; namely, the myopic utility loss to the agent from this bid-shading (which causes the agent to sometimes lose the item) is only $O(\epsilon^2)$, whereas the gain in utility resulting from the lower reserve is $\Omega(\epsilon)$. This argument lies at the heart of the proof of the proposition (see Appendix A).

For many if not most ROFs, one would expect the existence of a suitable shading-only bid function that causes a reduction in the reserve price by an amount of the same order as the magnitude of shading. Here is an argument to justify this claim: One may expect a reasonable ROF to be *scale-invariant*, meaning that if all historical bids are multiplied by the same factor α , then the reserve price chosen by the ROF should also be multiplied by α . Consider $\alpha = 1 - \epsilon$ for some small ϵ , and $n = 2$ agents (for instance). If *both* agents bid a multiple $1 - \epsilon$ of their true valuation, the resulting reserve under a scale-invariant ROF is reduced to $1 - \epsilon$ times its original value. Then, for any scale-invariant ROF that is differentiable in its inputs (i.e., previous bids), it must be that if just *one* agent bids a multiple $1 - \epsilon$ of her true valuation while the other agent is truthful, the resulting reserve is $\Omega(\epsilon)$ below the original one for small enough ϵ . (One may expect this to hold also for many scale-invariant ROFs that are *not* differentiable in the inputs.)

We informally remark here that the ability of an agent to influence a common reserve learnt from historical bids (cf. the reserve impact property) is significant when there are a small number of bidders, and this is also the setting in which the reserve matters (cf. the reserve matters property), both for the revenue earned by the seller and also for the expected utility of the agent. Thus, Proposition 1 suggests that there may be significant incentive issues associated with learning a common reserve when there are a small number of bidders, i.e., exactly when the reserve price is important for boosting revenues.

We now provide an illustration of the use of this result. Again considering the setting in Example 1 with two bidders and Uniform(0, 1) valuations, we show how Proposition 1 implies that a bid-shading strategy with $\underline{r} = 0.484$ and $\bar{r} = 0.526$ is beneficial to an agent and results in revenue loss to the seller, due to the reserve

²⁶ This is the loss due to reduction in the reserve price; we ignore here the additional loss of revenue due to bid-shading by agent i .

going down to $\underline{r} = 0.484$ from $r_* = 0.5$. To apply the proposition, we set $\epsilon = \bar{r} - \underline{r} = 0.042$, $\delta = r_* - \underline{r} = 0.016$, $C = \delta/\epsilon = 0.381$, and $q = \underline{r}(1 - \bar{r}) = 0.229$, and observe that $f_{\max} = 1$. Thus, $qC/(2f_{\max}) = 0.044 > \epsilon$ as needed, the bid function has the $(C\epsilon)$ -reserve impact property, and one can check that ROF has the $(C\epsilon, \epsilon, q)$ -reserve matters property. Thus, Proposition 1 tells us that the agent can gain at least $qC\epsilon/(2n) = 0.00092$ or 1.1% in expected utility by this deviation, and the seller loses at least $qC\epsilon = 0.0037$ or 0.9% in expected revenue. These numbers are somewhat smaller than those captured in Example 1 due to the slack in Proposition 1.

An alternative to setting the reserve based on the joint distribution of historical bids as per (5), is to pool historical bids from all bidders into a single empirical distribution $\hat{F}$, and then set the reserve to be the optimal monopoly price $\arg\max_r r(1 - \hat{F}(r))$ for this valuation distribution (Baliga and Vohra 2003). This approach suffers from similar incentive issues and again, Proposition 1 captures this. In fact, we can consider exactly the same deviation as above with $\underline{r} = 0.484$ and $\bar{r} = 0.526$, and again deduce from the proposition that the agent's expected utility increases by at least 1.1%, whereas the expected revenue of the seller decreases by at least 0.9%. (These effects can be increased slightly by considering symmetric deviations $\underline{r} = 1/2 - \delta$ and $\bar{r} = 1/2 + \delta$. Again, there is some slack in using Proposition 1, but it captures the qualitative fact that incentives are a concern when learning a common reserve, and this affects revenues.) In the next section, we modify the monopoly pricing approach to appropriately set *personalized* reserves for buyers in a manner that overcomes incentive issues.

B. Appendix to Section 4

The following lemma controls the amount by which the empirical cdf based on i.i.d. samples can deviate from the true cdf, allowing us to bound the loss in revenue. We state a general version that allows an adversary to perturb each sample by up to $\delta \geq 0$ so that we can reuse the lemma later to analyze the contextual setting (i.e., in the proof of Theorem 3). In proving the bound on revenue loss in Theorem 1, we will simply set $\delta = 0$ when we invoke this lemma.

Lemma 2 *There exists $C_8 < \infty$ such that the following holds. Consider any $\delta \geq 0$ and an arbitrary distribution F , and let $v_1, v_2, \dots, v_N$ be N i.i.d. samples from F , except that an adversary may have arbitrarily modified (increased or decreased) each sample by an amount of at most δ . Let $\hat{F}$ be the empirical distribution of these perturbed samples. Then, with probability at least $1 - 1/N$, it holds that for all r ,*

$$\hat{F}(r) \in [F(r - \delta) - C_8\sqrt{\log N/N}, F(r + \delta) + C_8\sqrt{\log N/N}]. \quad (21)$$

Proof of Lemma 2. We draw inspiration from the proof of (Dhangwatnotai et al. 2015, Lemma 4.1 and Remark 4.2). Using Chernoff bounds followed by a union bound as in that proof, we obtain that with no adversarial perturbations, with probability at least $1 - 1/N$, it holds for all $i = 1, 2, \dots, N$ that

$$\begin{aligned} |\hat{F}_{\text{np}}(\tilde{v}_i -) - F(\tilde{v}_i -)| &\leq C_8\sqrt{\log N/N} \quad \text{and} \\ |\hat{F}_{\text{np}}(\tilde{v}_i +) - F(\tilde{v}_i +)| &\leq C_8\sqrt{\log N/N}, \end{aligned} \quad (22)$$

where $\hat{F}_{\text{np}}(\cdot)$ is the empirical distribution of unperturbed samples $\tilde{v}_1, \tilde{v}_2, \dots, \tilde{v}_N$. (We carry bounds on both the left- and right-hand limits throughout this proof.) Now consider the effect of adversarial perturbations. Since these perturbations are of size at most δ , i.e., $|\tilde{v}_i - v_i| \leq \delta \forall i$, it is immediate to deduce that for all r we have

$$\hat{F}(r) \in [\hat{F}_{\text{np}}(r - \delta), \hat{F}_{\text{np}}(r + \delta)]. \quad (23)$$

Combining (22) and (23) we obtain that with probability at least $1 - 1/N$, it holds for all $i = 1, 2, \dots, N$ that

$$\begin{aligned} \hat{F}(v_i -) &\in [F((v_i - \delta) -) - C_8 \sqrt{\log N/N}, F((v_i + \delta) -) + C_8 \sqrt{\log N/N}] \quad \text{and} \\ \hat{F}(v_i +) &\in [F((v_i - \delta) +) - C_8 \sqrt{\log N/N}, F((v_i + \delta) +) + C_8 \sqrt{\log N/N}]. \end{aligned} \quad (24)$$

We can now extend to all r as follows. Without loss of generality assume that the values of the samples $v_1, v_2, \dots, v_N$ are in descending order. If r is equal in value to one of these samples we are done since $\hat{F}(r) = \hat{F}(r -)$ and $F(r') = F(r' -) \forall r'$, as per the convention in this paper of defining cumulative distributions with a strict inequality $F(r) = \Pr(v < r)$. For convenience, define $v_0 \triangleq B_F + \delta$ and $v_{N+1} \triangleq -B_F - \delta$. Given that F is supported on $(-B_F, B_F)$, observe that $F(v_0) = \hat{F}(v_0) = 0$ and $F(B_F) = \hat{F}(v_0) = 1$, and hence that (24) extends to v_0 and v_{N+1} . For any r such that $|r| < B_F + \delta$, there exists $i \in \{0, 1, \dots, N\}$ such that $v_i > r > v_{i+1}$. By definition of $\hat{F}(\cdot)$ we have $\hat{F}(r) = \hat{F}(v_i -) = \hat{F}(v_{i+1} +)$. Using (24) and monotonicity of $F(\cdot)$, we have

$$\begin{aligned} \hat{F}(r) = \hat{F}(v_i -) &\geq F((v_i - \delta) -) - C_8 \sqrt{\log N/N} \geq F(r - \delta) - C_8 \sqrt{\log N/N} \quad \text{and} \\ \hat{F}(r) = \hat{F}(v_{i+1} +) &\leq F((v_{i+1} + \delta) +) + C_8 \sqrt{\log N/N} \leq F(r + \delta) + C_8 \sqrt{\log N/N}. \end{aligned}$$

Finally, for $r \geq B_F + \delta$ we have $\hat{F}(r) = F(r) = 1$ and the bounds clearly hold; and similarly for $r \leq -B_F - \delta$. This completes the proof of the lemma. $\square$

We will need one more lemma that will ensure that the empirical revenue maximizing reserve price is within a constant factor of the true revenue maximizing reserve price.

Lemma 3 *Consider any MHR distribution F , and let $\hat{F}$ be the empirical distribution of N i.i.d. samples from F . Let $\hat{r}_* \triangleq \max_r r(1 - \hat{F}(r))$ and $r_* = \max_r r(1 - F(r))$. Then there exists $C' < \infty$ such that, for all N , we have that $\hat{r}_* \leq 6r_*$ holds with probability $1 - C' \exp(-N^{0.4})$.*

Proof of Lemma 3. Let the samples, in descending order, be $v_1, v_2, \dots, v_N$. Note that $\hat{r}_*$ takes on one of these N values. For each $i \leq N/(6e)$, we will show that the probability that $\hat{r}_* = v_i$ is small by showing that there are unlikely to be i samples that take large enough values. We will then use a union bound to establish the lemma.

We use the simple fact, first noted in (Hartline et al. 2008, Lemma 4.1), that since F is MHR, we have

$$1 - F(r_*) \geq 1/e. \quad (25)$$

It follows using a Chernoff bound on $\text{Binomial}(N, 1/e)$ that event $E_0 \triangleq \{1 - \hat{F}(r_*) < 0.9/e\}$ occurs with probability at most $\exp(-(0.1)^2 N/(2e)) \leq \exp(-0.001N)$. Note that $1 - \hat{F}(v_i) = i/N$. Assume E_0 does not occur. Then in order to have $\hat{r}_* = v_i$, it must be that

$$v_i i/N \geq r_*(1 - \hat{F}(r_*)) \Rightarrow v_i \geq 0.9Nr_*/(ie), \quad (26)$$

using $1 - \hat{F}(r_*) \geq 0.9r_*/e$. Accordingly, define the event $E_i \triangleq \{v_i \geq 0.9Nr_*/(ie)\}$, i.e., E_i is the event that there are i or more samples with value at least $0.9Nr_*/(ie)$. We will show that E_i is unlikely. Let $h(r) \triangleq f(r)/(1 - F(r))$ denote the hazard rate of F . The first order condition corresponding to the definition r_* tells us that $h(r_*) = 1/r_*$. Using that F is MHR, for any $r > r_*$, we have

$$\begin{aligned} 1 - F(r) &= (1 - F(r_*)) \exp\left(-\int_{r_*}^r h(r')dr'\right) \\ &\leq (1 - F(r_*)) \exp(-(r - r_*)h(r_*)) = (1 - F(r_*)) \exp(-(r - r_*)/r_*) \\ &\leq \exp(-(r - r_*)/r_*) = \exp(1 - r/r_*). \end{aligned}$$

In particular,

$$1 - F(0.9Nr_*/(ie)) \leq q_i \triangleq \exp(1 - 0.9N/(ie)). \quad (27)$$

Thus, we have that

$$\Pr(E_i) \leq \Pr(\text{Binomial}(N, q_i) \geq i).$$

Note that for all $i < N/6e$, we have $eq_i < 0.6i$ and in particular $q_i < i$. (This can be seen by observing that $\ell(y) \triangleq y \exp(-y)$ is monotone decreasing for $y > 1$, that $\ell(5.4) = 0.544 \dots < 0.6$, and defining $y \triangleq 0.9N/(ie) > 5.4$.) Using a Chernoff bound for $i > N/6e$, the right-hand side is bounded above as

$$\Pr(\text{Binomial}(N, q_i) \geq i) \leq \frac{e^{i-q_i}}{(i/q_i)^i} = (eq_i/i)^i / e^{q_i} \leq (eq_i/i)^i.$$

For $i \geq \sqrt{N}$, using $eq_i < 0.6i$, we immediately get an upper bound of $0.6^i \leq \exp(-0.5\sqrt{N})$, using $\ln(0.6) < -0.5$. On the other hand, for $i < \sqrt{N}$, we have an upper bound of $eq_i/i = e \exp(1 - 0.9N/(ie))/i \leq \exp(2 - 0.3\sqrt{N})$ using $0.9/e > 0.3$. Thus, we have established that

$$\Pr(E_i) \leq \exp(2 - 0.3\sqrt{N}) \quad \text{for all } i < N/6e. \quad (28)$$

It follows using a union bound that probability that at least one of $E_0, E_1, \dots, E_{N/(6e)}$ occurs is bounded above by $N \exp(2 - 0.3\sqrt{N})/(6e) + \exp(-0.001N)$. Observe that for appropriately chosen $C' < \infty$, this quantity is further bounded above by $C' \exp(-N^{0.4})$ for all N .

Thus, with probability at least $1 - C' \exp(-N^{0.4})$, none of $E_0, E_1, \dots, E_{N/(6e)}$ occur. When none of these events occur, by the argument above, $\hat{r}_* < v_{N/(6e)} < 0.9Nr_*/(eN/(6e)) = 5.4r_* < 6r_*$. This completes the proof. $\square$

Proof of Theorem 1: Consider agent i at time t and assume that all other agents will be truthful (at time t and in the future). Note that for each agent i , the sequence of her reserve prices and her prices to win

the item (the maximum of r_{it} and the highest bids of other agents) do not depend at all on i 's own bids. Hence, the bid of agent i does not affect her utility in future rounds and myopically maximizing utility in each round is optimal for maximizing the long-term utility. Since truthful bidding is myopically a dominant strategy in each round, being truthful is a best response, and this establishes periodic ex-post incentive compatibility. In particular, all agents following the always truthful strategy constitutes an equilibrium.

The lower bound on revenue follows from Remark 4.2 from Dhangwatnotai et al. (2015), and the fact that it can be used to show a $(1 - \epsilon)$ -factor revenue optimality even when there are multiple bidders (see the proof of Theorem 4.3 in Dhangwatnotai et al. 2015). For completeness we provide a full proof.

By definition of r_{it} , we have

$$r_{it}(1 - \hat{F}_{-i}(r_{it})) \geq r_*(1 - \hat{F}_{-i}(r_*)). \quad (29)$$

For window length W , using Lemma 2 with $\delta = 0$, we have that with probability at least $1 - 1/W$,

$$\min(1, 1 - F(r) + C_8 \sqrt{\log W/W}) \geq 1 - \hat{F}_{-i}(r) \quad \text{for all } r, \quad \text{and} \quad (30)$$

$$1 - F(r_*) - C_8 \sqrt{\log W/W} \leq 1 - \hat{F}_{-i}(r_*). \quad (31)$$

where we also used $\hat{F}_{-i}(r) \geq 0 \Rightarrow 1 - \hat{F}_{-i}(r) \leq 1$ to obtain (30). Using (30) to upper bound the left-hand side of (29) and using (31) to lower bound the right-hand side of (29), we obtain that with probability at least $1 - 1/W$,

$$r_{it} \min(1, 1 - F(r_{it}) + C_8 \sqrt{\log W/W}) \geq r_*(1 - F(r_*) - C_8 \sqrt{\log W/W}). \quad (32)$$

Define the monopoly revenue $\text{MONREV}_r \triangleq r(1 - F(r))$. We deduce that

$$\text{MONREV}_{r_{it}} \geq \text{MONREV}_{r_*} - C_8(r_{it} + r_*)\sqrt{\log W/W}. \quad (33)$$

We now convert this into a multiplicative bound using Lemma 3. Using a union bound, the characterizations in both Lemma 2 and Lemma 3 simultaneously hold with probability at least $1 - 1/W - C' \exp(-W^{0.4})$. Lemma 3 tells us that $r_{it} < 6r_*$. Along with fact (25) that states $1 - F(r_*) \geq 1/e$, we obtain that

$$r_{it} + r_* < 7r_* = 7\text{MONREV}_{r_*}/(1 - F(r_*)) \leq 7e\text{MONREV}_{r_*}.$$

Plugging into (33), we obtain

$$\text{MONREV}_{r_{it}} \geq (1 - \tilde{\delta})\text{MONREV}_{r_*}, \quad \text{where } \tilde{\delta} \triangleq 7eC_8\sqrt{\log W/W}. \quad (34)$$

We now extend this multiplicative bound to the actual auction revenue with multiple bidders. As in the proof of (Dhangwatnotai et al. 2015, Theorem 4.3, top of page 332), we infer that

$$\text{MONREV}_{\max(r_{it}, b)} \geq (1 - \tilde{\delta})\text{MONREV}_{\max(r_*, b)} \quad \text{for all } b \in \mathbb{R}, \quad (35)$$

via a simple case analysis, using the property that monopoly revenue is a concave function of the price:

(i) If $b \leq r_*$, then the right-hand side is simply $(1 - \tilde{\delta})\text{MONREV}_{r_*}$. Since $\max(r_{it}, b) \in [r_{it}, r_*]$, and because the monopoly revenue is increasing in price to the left of r_* , the left-hand side is lower bounded by $\mathbb{E}[\text{MONREV}_{r_{it}}] \geq (1 - \tilde{\delta})\mathbb{E}[\text{MONREV}_{r_*}] = \text{RHS}$. Thus, we have established (35) for this case. (ii) If $b > r_*$,

then the right-hand side is simply $(1 - \tilde{\delta})\text{MONREV}_b$. The left-hand side is identical if $r_{it} \leq b$, in which case we are done. Else, we have $r_{it} > b$ and the LHS is $\text{MONREV}_{r_{it}} \geq (1 - \tilde{\delta})\text{MONREV}_{r_*} \geq (1 - \tilde{\delta})\text{MONREV}_b = \text{RHS}$. Again, we are done.

Now setting b to be the largest bid submitted by an agent other than i and taking expectation over the bids of the other agents, we get

$$\begin{aligned} & \mathbb{E}[\text{Revenue from } i \text{ under proposed mechanism}] \geq \mathbb{E}[\text{Revenue from } i \text{ under Myerson}](1 - \tilde{\delta}) \\ \Rightarrow & \mathbb{E}[\text{Revenue under proposed mechanism}] \geq (1 - \tilde{\delta})\text{REV}_*, \end{aligned}$$

where the second step follows by summing over i . Thus, incorporating the $1/W + C \exp(-W^{0.4})$ probability of failure of Lemma 2 and/or Lemma 3 due to an atypical empirical distribution of historical bids, in which case revenue loss can be up to REV_* , the overall loss in expected revenue under our mechanism (relative to Myerson) is bounded above by

$$\text{REV}_* (\tilde{\delta} + 1/W + C' \exp(-W^{0.4})) \leq \text{REV}_* C'' \sqrt{\log W/W}$$

for appropriate $C'' < \infty$, for all W . It follows that for appropriate $C < \infty$, for any $\epsilon > 0$, window length $W > C \log(1/\epsilon)/\epsilon^2$ ensures a revenue loss of at most ϵREV_* . The lower bound of $(1 - \epsilon)\text{REV}_*$ on revenue follows.

(In the case where F is a regular distribution – a weaker requirement than MHR – we need to use the so-called “guarded empirical reserve” that forbids the largest samples from acting as reserve prices, and to require that $W \geq C \log(1/\epsilon)/\epsilon^3$ in order to obtain $(1 - \epsilon)$ -factor optimality using Lemma 4.1 from Dhangwatnotai et al. (2015).) $\square$

Remark 3 *Our above proof of the lower bound on revenue leveraged results from Dhangwatnotai et al. (2015). We remark that using Lemma 2 (which we will establish en route to a proof of Theorem 3), the proof can be made self contained: Using Lemma 2, we see that with probability $1 - 1/N$, the following conditions hold*

We now provide a finite horizon version of Theorem 1. Suppose the seller is trying to maximize revenue over a horizon of T periods. Correspondingly, agents are trying to maximize their utility over T rounds. The definition of periodic ex-post incentive compatibility and HO-SERP mechanisms remains unchanged. We use the HO-SERP mechanism that uses all bids in periods 1 through $t - 1$ to determine the reserve prices in period t (rather than using bids in a window of length W), and provide an upper bound on the loss in average revenue relative to the Myerson benchmark.

Corollary 1 *Consider a finite horizon setting with horizon length T . Any HO-SERP mechanism is periodic ex-post incentive compatible. In particular, all agents following the always truthful strategy constitutes an equilibrium. Now consider the HO-SERP mechanism that sets personal reserve prices in period t as per (6) where $\hat{F}_{-i}(\cdot)$ is the empirical distribution of bids by other agents in periods 1, 2, $\dots$, $t - 1$. There exists $C < \infty$*

that does not depend on the valuation distribution F , such that for any F that is MHR, this mechanism achieves average revenue that is at least $(T - C\sqrt{T}\log T)\text{REV}_*$, where REV_* is the expected revenue under the optimal static mechanism, i.e., the second-price auction with Myerson-optimal reserve.

Proof of Corollary 1. The proof of periodic ex-post IC is as before. Consider the expected revenue in period $t \geq 2$. It is the same as the per period average revenue of the windowed HO-SERP mechanism with window length $t - 1$. Let $\hat{C}$ denote the constant in Theorem 1 for purposes of this proof. We can use Theorem 1 with ϵ such that $t - 1 = \hat{C}\log(1/\epsilon)/\epsilon^2$, i.e., $\epsilon \leq \tilde{C}\log t/\sqrt{t}$ for some $\tilde{C} < \infty$ and $t \geq 2$. This gives us a bound of $\tilde{C}\text{REV}_*\log t/\sqrt{t}$ on the expected revenue loss in period $t + 1$ for $t \geq 2$. In the first period the expected revenue loss is bounded by REV_* . It follows that the expected revenue loss over T periods is bounded by $\tilde{C}\text{REV}_*\left(1 + \sum_{t=2}^T \log t/\sqrt{t}\right) \leq C\text{REV}_*\sqrt{T}\log T$ for some $C < \infty$. This completes the proof. $\square$

Finally, we provide a lower bound on the revenue loss under our suggested mechanism. We choose to stay in the finite horizon setting to state this result. We show that the revenue loss is $\Omega(\text{REV}_*T^{0.33})$ under our proposed HO-SERP mechanism for a representative (exponential)²⁷ valuation distribution F .

Theorem 5 *Consider a finite horizon setting with horizon length T , and the HO-SERP mechanism in Corollary 1 that sets personal reserve prices in period t as per (6) where $\hat{F}_{-i}(\cdot)$ is the empirical distribution of bids by other agents in periods $1, 2, \dots, t - 1$. For any $\varepsilon > 0$ and $F(v) = e^{-v}$, the average revenue under this mechanism is bounded above as $T - \Omega(T^{1/3-\varepsilon})$.*

The key lemma to establish this result is the following.

Lemma 4 *Fix any $\varepsilon > 0$. Let $\hat{F}$ be the empirical distribution based on $t \in \mathbb{N}$ independent identically distributed samples from distribution $F(v) = e^{-v}$. Suppose we set*

$$\hat{r}_* = \arg \max_r r(1 - \hat{F}(r)). \quad (36)$$

Define $R(r) \triangleq r(1 - F(r))$. Note that $R(r)$ is maximized at $r_ = 1$. Then we have with probability $\Omega(1)$ that*

$$|\hat{r}_* - r_*| = \Omega(t^{-1/3-\varepsilon/2}), \quad \text{and} \quad (37)$$

$$R(\hat{r}_*) \leq R(r_*) - \Omega(t^{-2/3-\varepsilon}). \quad (38)$$

Proof of Lemma 4. For convenience, define $\delta \triangleq \min(\varepsilon/4, 1/50)$. Define $r_0 \triangleq r_* - t^{-1/3-\delta}$, $r_1 \triangleq r_* - t^{-1/3-2\delta}$ and $r_2 \triangleq r_* + t^{-1/3-2\delta}$. Define the empirically estimated counterpart of $R(\cdot)$,

$$\hat{R}(r) \triangleq r(1 - \hat{F}(r)). \quad (39)$$

²⁷ The analogous bound in the infinite horizon setting with the windowed HO-SERP mechanism is: to obtain average revenue loss per period of no more than ϵREV_* , the required minimum window length is $W = \Omega(1/\epsilon^{1.49})$. We omit this infinite horizon bound in the interest of space.

We will show that with probability $\Omega(1)$, we have that

$$\hat{R}(r_0) > \hat{R}(r) \quad \forall r \in [r_1, r_2] \Rightarrow \hat{r}_* \notin [r_1, r_2], \quad (40)$$

yielding (37). Throughout the proof, we omit the phrase “for large enough t ” though it is repeatedly invoked.

Of the t total samples, let $N(r) \triangleq t(1 - \hat{F}(r-))$ be the number of samples of value at least R . We define the following events

$$E_0 \triangleq \{N(r_0) - N(r_1) \geq te^{-1}(r_1 - r_0) + t^{1/3-\delta/2}\}, \quad (41)$$

$$E_1 \triangleq \{|N(r_1) - te^{-1}| \leq 2t^{2/3-2\delta}\}, \quad (42)$$

$$E_2 \triangleq \{N(r_1) - N(r) \geq (r - r_1)e^{-1}t - t^{1/3-2\delta/3} \text{ for all } r \in [r_1, r_2]\}. \quad (43)$$

The event E_0 captures “more than the expected number of samples in $[r_0, r_1]$ ” and occurs with probability $\Omega(1)$, for the following reason: The number of samples in $[r_0, r_1]$ is Binomial($t, F(r_0) - F(r_1)$). We can bound the mean from below as $t(F(r_0) - F(r_1)) \geq t(r_1 - r_0)f(r_*) = t(r_1 - r_0)e^{-1}$ using that $f(\cdot)$ is decreasing. Since the variance of the binomial is then $t^{2/3-\delta}(1 - O(t^{-1/3-\delta}))$, we have

$$\Omega(1) = P\left(N(r_0) - N(r_1) \geq te^{-1}(r_1 - r_0) + \sqrt{t^{2/3-\delta}}\right) = P(E_0)$$

using the central limit theorem.

The event E_1 captures “typical number of samples with value at least r_1 ” and occurs with high probability, i.e., with probability $1 - o(1)$. This is straightforward to see: the number of samples is Binomial(t, e^{-r_1}) and $e^{-r_1} = 1 + t^{-1/3-2\delta} + O(t^{-2/3-4\delta})$. The variance of the Binomial is less than t . The claim then follows from a standard Chernoff bound since the permitted slack of $t^{2/3-2\delta}$ grows faster than the standard deviation of $O(\sqrt{t})$.

The event E_2 captures, roughly that “the empirical distribution of samples in $[r_1, r_2]$ will not be much sparser than a typical realization” and occurs with high probability. To establish this, we argue as follows: First note that we can legitimately generate t i.i.d. samples from F via the following approach, which will facilitate our analysis. Simulate a Poisson process on $\mathbb{R}_+$ with intensity $tf(r)(1 - t^{-0.4})$ at r . With high probability (using a Chernoff bound), the total number of points in the realization will be $t' < t$. Then draw $t - t'$ additional i.i.d. samples from F to complete the set of t i.i.d. samples. What we gained here is that we “lower bounded” the point process of t samples from F via a Poisson process of intensity $tf(r)(1 - t^{-0.4})$. Now we characterize the typical realizations of the lower-bounding Poisson process on $[r_1, r_2]$. The interval has length $2t^{-1/3-2\delta}$. The Poisson density on this interval is everywhere at least $tf(r_2)(1 - t^{-0.4}) \geq te^{-1}(1 - 2t^{-1/3-2\delta})$. We now rescale r as well as the Poisson intensity by the same factor $t^{-1/3-2\delta}$. Let $s(r) \triangleq (r - r_1)/t^{-1/3-2\delta}$ for $r \in [r_1, r_2] \Leftrightarrow s \in [0, 2]$. Let $\tilde{N}(s)$ be the value of a Poisson process of uniform intensity $\lambda \triangleq te^{-1}(1 - 2t^{-1/3-2\delta}) \times t^{-1/3-2\delta} = e^{-1}(t^{2/3-2\delta} - 2t^{1/3-4\delta})$ on $s \in [0, 2]$. Note that we have retained the lower bounding property that, with high probability, $N(r) - N(r_1) \geq \tilde{N}(s)$ for all $r \in [r_1, r_2]$. Since the intensity λ of $\tilde{N}(s)$ scales up with t , the process converges to a standard Brownian motion (e.g., see Billingsley 2013, Theorem 37.8) after subtracting the mean and scaling by $\sqrt{\lambda}$, i.e.,

$$\frac{\tilde{N}(s) - \lambda s}{\sqrt{\lambda}} \xrightarrow{t \rightarrow \infty} W_s \quad \text{on } s \in [0, 2], \quad (44)$$

where W_s is standard Brownian motion. For standard Brownian motion in a horizon of length 2, we know that, with high probability, it does not ever become too small;

$$P\left(\min_{s \in [0,2]} W_s < -M\right) \xrightarrow{M \rightarrow \infty} 0. \quad (45)$$

Using $M \triangleq t^{\delta/3}$ and plugging back into (44) we obtain that, with high probability, for all $s \in [0, 2]$ we have²⁸

$$\tilde{N}(s) \geq s\lambda - (M+1)\sqrt{\lambda} \geq se^{-1}t^{2/3-2\delta} - t^{1/3-2\delta/3} = (r-r_1)e^{-1}t - t^{1/3-2\delta/3}. \quad (46)$$

The lower bounding property of $\tilde{N}(s)$ then implies that E_2 occurs with high probability.

Since E_0 occurs with probability $\Omega(1)$ whereas E_1 and E_2 occur with high probability, it follows that $E_0 \cap E_1 \cap E_2$ occurs with probability $\Omega(1)$. Henceforth in this proof we assume that all three events occur. Our goal is to establish (40). We will rely on E_0 and E_1 to show that $\hat{R}(r_0)$ is significantly larger than $\hat{R}(r_1)$, and then deduce (40) by using E_2 .

Observe that, by definition, we have

$$\hat{R}(r_0) - \hat{R}(r_1) = -N(r_1)(r_1 - r_0) + (N(r_0) - N(r_1))r_0$$

Using the upper bound on $N(r_1)$ in (42) and the lower bound on $N(r_0) - N(r_1)$ in (41), we deduce

$$\hat{R}(r_0) - \hat{R}(r_1) \geq t^{1/3-\delta/2}/2. \quad (47)$$

Now consider

$$\hat{R}(r) - \hat{R}(r_1) = -N(r_1)(r_1 - r) + (N(r) - N(r_1))r,$$

for $r \in [r_1, r_2]$. Using the lower bound on $N(r_1)$ in (42) for the first term and the lower bound on $N(r) - N(r_1)$ in (43) for the second term, we obtain

$$\hat{R}(r) - \hat{R}(r_1) \leq 2t^{1/3-2\delta/3} \quad (48)$$

for all $r \in [r_1, r_2]$.

Combining (47) and (48) and using $1/3 - 2\delta/3 > 1/3 - \delta/2$, we deduce (40). (38) follows: by definition of r_* , we have $R'(r_*) = 0$, and one can further verify that $R'(r) = -rf(r) + 1 - F(r)$ is positive for $r < r_*$ and negative for $r > r_*$. Further $R''(r) = e^{-r}(r-2) < -0.5e^{-1.5}$ for all $r \in [0.5, 1.5] \supseteq [r_1, r_2]$. It follows that $\sup_{r \notin [r_1, r_2]} R(r) \leq \max(R(r_1), R(r_2)) \leq R(r_*) - (1/2)(0.5e^{-1.5})(t^{-1/3-2\delta})^2 = R(r_*) - \Omega(t^{-2/3-4\delta})$, where the first inequality follows from the unimodality of $R(\cdot)$ and the second inequality follows from Taylor's theorem using the bound on $R''(r)$ and $|r_1 - r_*| = |r_2 - r_*| = t^{-1/3-2\delta}$. Hence, $r \notin [r_1, r_2]$ implies (38). $\square$

Proof of Theorem 5. Lemma 4 tells us that the expected (additive) revenue loss relative to REV_* in period t under the proposed HO-SERP mechanism is $\Omega(t^{-2/3-\epsilon})$. Summing over $t = 1, 2, \dots, T$ gives the result. $\square$

²⁸ Since the left-hand side in (44) converges to the right-hand side, the two stochastic process do not differ by 1 unit at any s , and in particular, the minimum value of the left-hand process is smaller than the minimum value of right-hand process by at most 1.

Our lower bound in Lemma 4 contributes to the agenda pursued in Dhangwatnotai et al. (2015) regarding choosing a price to optimize revenue based on a limited number of samples from the valuation distribution. The reader will notice a gap between our lower bound of $\Omega(\text{REV}_* T^{0.33})$ and our upper bound of $O(\text{REV}_* \sqrt{T} \log T)$ on the expected revenue loss. We leave it as an interesting open problem to close this gap. We have a (weak) conjecture that our lower bound on revenue is nearly tight; the worst case expected revenue loss may indeed be²⁹ $\Theta(T^{1/3})$.

Proof for Example 2. As mentioned in the example, suppose that the first buyer is truthful. We present a profitable deviation for the second buyer, parameterized by Δ and later we show that the deviation is most profitable at $\Delta = \frac{1}{2}$. The second buyer, for a parameter $\Delta \leq \frac{1}{2}$, bids as follows:

$$b_{it} = \begin{cases} v_{it} & 0 \leq v_{it} < \frac{1}{2} \\ \frac{1}{2} + \Delta & \frac{1}{2} \leq v_{it} < \frac{1}{2} + \Delta \\ v_{it} & \frac{1}{2} + \Delta \leq v_{it} \leq 1 \end{cases}$$

Note that the second bidder overbids if her valuation is in $[\frac{1}{2}, \frac{1}{2} + \Delta)$ and is truthful otherwise. Therefore, the limiting reserve price for the *first bidder* (him) is $\frac{1}{2} + \Delta$. Since the first bidder is truthful, the limiting reserve price for the second bidder is $\frac{1}{2}$.

We now calculate the per-round increase in profit from this deviation. Observe that the deviating second agent obtains higher utility only when the price reduces because the first bidder is eliminated on account of not clearing his personal reserve price, whereas he would have cleared the reserve of $1/2$ resulting from truthful bidding by the second agent. Thus, the elimination happens only if the truthful bidder's valuation lies in $[\frac{1}{2}, \frac{1}{2} + \Delta)$ which occurs with probability Δ . If the deviating agent's valuation is above $\frac{1}{2} + \Delta$, she benefits from such elimination via a price that is $\Delta/2$ lower on average. This case contributes

$$\Pr(v_1 \in [1/2, 1/2 + \Delta)) \Pr(v_2 \in [1/2 + \Delta, 1]) \Delta/2 = \Delta(1/2 - \Delta) \times \Delta/2 = \Delta^2/4 - \Delta^3/2$$

to the gain in expected utility. In addition, if $v_i \in [1/2, 1/2 + \Delta)$ for both $i = 1, 2$, this leads to a $\min(v_1, v_2) - 1/2$ gain in utility for the second bidder due to agent 1 not clearing his reserve. The resulting contribution to the gain in expected utility is

$$\int_{1/2}^{1/2+\Delta} \int_{1/2}^{1/2+\Delta} (\min(v_1, v_2) - 1/2) f(v_1) f(v_2) dv_1 dv_2 = \Delta^3/3.$$

Hence, the overall gain in expected utility from bid-shading is $\Delta^2/4 - \Delta^3/6$.

Note that in the limiting steady state, overbidding does not lead to utility loss for the second agent. The reason is that when she overbids, her bid is $1/2 + \Delta$, but this bid wins only if the first agent does not clear his reserve of $1/2 + \Delta$. Thus, the price paid by the second agent is only her reserve of $1/2$, which is less than her true valuation.

The gain in expected utility from deviation of $\Delta^2/4 - \Delta^3/6$ is non-negative for all $\Delta \in [0, 1/2]$. The benefit of deviation is maximized at $\Delta = 1/2$, which can increase the utility of the second agent, compared to the

²⁹ This conjecture arose from a discussion with Amine Allouah.

truthful strategy, by $1/24 = 0.0417$ or 50% (since her utility under the truthful strategy is $1/12 = 0.0833$). Note how this effectively prevents the first agent from ever being allocated the item since $r_{1t} = 1$ under the deviation. $\square$

C. Appendix to Section 5

We first provide a fact we will use to establish Theorem 2 for $n > 2$.

Fact 5 *Let F_a and F_b be distributions with virtual value functions ϕ_a and ϕ_b respectively. Then, for any $\lambda \in [0, 1]$, the convex combination of the two distributions $F_\lambda = \lambda F_a + (1 - \lambda)F_b$ has virtual value function ϕ_λ which is between the virtual value function of the individual distributions, i.e., for all v ,*

$$\phi_\lambda(v) \in \begin{cases} [\phi_a(v), \phi_b(v)] & \text{if } \phi_a(v) \leq \phi_b(v), \\ [\phi_b(v), \phi_a(v)] & \text{otherwise.} \end{cases} \quad (49)$$

Proof of Fact 5. Recall that the virtual value function for any distribution F is defined as $\phi(v) = v - (1 - F(v))/f(v)$. Hence, in order to show that ϕ_λ is between ϕ_a and ϕ_b , it suffices to show that $\rho_\lambda(v) \triangleq (1 - F_\lambda(v))/f_\lambda(v)$ is between $\rho_a(v) \triangleq (1 - F_a(v))/f_a(v)$ and $\rho_b(v) \triangleq (1 - F_b(v))/f_b(v)$ for all v . Observe that from the definition of F_λ we have $f_\lambda(v) = \lambda f_a(v) + (1 - \lambda)f_b(v)$. Without loss of generality, suppose $\rho_a(v) \geq \rho_b(v)$. We immediately deduce that

$$\begin{aligned} \rho_\lambda(v) &= \frac{\lambda(1 - F_a(v)) + (1 - \lambda)(1 - F_b(v))}{\lambda f_a(v) + (1 - \lambda)f_b(v)} \\ &= \frac{\lambda \rho_a(v) f_a(v) + (1 - \lambda) \rho_b(v) f_b(v)}{\lambda f_a(v) + (1 - \lambda)f_b(v)} && \text{(using } 1 - F_i(v) = \rho_i(v) f_i(v) \text{ for } i = a, b) \\ &\geq \frac{\lambda \rho_b(v) f_a(v) + (1 - \lambda) \rho_b(v) f_b(v)}{\lambda f_a(v) + (1 - \lambda)f_b(v)} && \text{(using } \rho_a(v) \geq \rho_b(v)) \\ &= \rho_b(v). \end{aligned}$$

A very similar argument establishes $\rho_\lambda(v) \leq \rho_a(v)$, thus yielding $\rho_\lambda(v) \in [\rho_b(v), \rho_a(v)]$ as needed. $\square$

We remark that if F_a and F_b are MHR, the convex combination F_λ may still not be MHR (for example, the convex combination of two exponential distributions with different parameters is not MHR) and in fact, may not even be regular (we omit an example in the interest of space). However, lack of regularity will not create difficulties for us in proving our robustness result. (As an aside for the reader who is familiar with auction design theory: If F_a and F_b are regular, the “ironed” virtual value function (Myerson 1981, (6.5)) of the convex combination $\tilde{\phi}_\lambda$ can also be shown to lie between ϕ_a and ϕ_b . The function $\tilde{\phi}_\lambda$ is monotone non-decreasing and we could have chosen to work with it in place of the virtual value function. For instance, $r_* = \arg \max_r r(1 - F_\lambda(r))$ is the value at which $\tilde{\phi}_\lambda(r_*) = 0$. However, the first order condition also tells us that $\phi_\lambda(r_*) = 0$ for the virtual value function without ironing (though this function may have multiple zeros).

It turns out that $\phi_\lambda(r_*) = 0$ is the only property we need for our proof, and so we do not pass to the ironed virtual value function in our formal analysis.)

We are now ready to prove Theorem 2.

Proof of Theorem 2. The proof of periodic ex-post incentive compatibility in Theorem 1 remains valid when agents have heterogeneous valuation distributions.

We now prove that the loss in average per round seller revenue, relative to the Myerson-optimal static benchmark, is no more than $2(n-1)f_{\max}\delta^2$. Note that since the agent valuation distributions are MHR, the virtual value functions are monotone increasing, and that we have

$$\phi_i(v+z) - \phi_i(v) \geq z \quad (50)$$

for all $z > 0$. We begin by considering $n = 2$ agents and later generalize to an arbitrary number of agents.

Proof for $n = 2$. Consider our mechanism after learning has occurred. Our mechanism (as well as Myerson’s mechanism) is truthful. As a result, using Myerson’s lemma (Myerson 1981), the revenue of the mechanism can be written purely in terms of allocation (a_1, a_2) , where $a_1 = a_1(v_1, v_2)$ is the likelihood (0 or 1 for our mechanism, which is deterministic) that agent 1 is allocated the item, and similarly for $a_2 = a_2(v_1, v_2)$. The revenue is

$$\text{REV} = \int_0^\infty \int_0^\infty (a_1(v_1, v_2)\phi_1(v_1) + a_2(v_1, v_2)\phi_2(v_2)) f_1(v_1)f_2(v_2)dv_1dv_2. \quad (51)$$

The revenue under Myerson’s mechanism (which allocates the item to the agent with the highest virtual value, if that virtual value is non-negative) is simply

$$\text{REV}_* = \int_0^\infty \int_0^\infty \max(\phi_1(v_1), \phi_2(v_2), 0) f_1(v_1)f_2(v_2)dv_1dv_2. \quad (52)$$

Preliminary observations. We make some preliminary observations before proceeding to bound $\text{REV}_* - \text{REV}$. Note that if $v_1 \geq v_2$, we have

$$\phi_1(v_1) \geq \phi_2(v_1) - \delta \geq \phi_2(v_2) - \delta + v_1 - v_2, \quad (53)$$

where the first inequality follows from the assumption $\|F_i - F_j\| \leq \delta$, and the second inequality follows from (50). In particular, we deduce

$$\phi_1(v_1) \geq \phi_2(v_2) - \delta \quad \text{for all } v_1 \geq v_2, \quad \text{and} \quad (54)$$

$$\phi_1(v_1) \geq \phi_2(v_2) \quad \text{for all } v_1 \geq v_2 + \delta. \quad (55)$$

Let r_2 be the unique value such that $\phi_1(r_2) = 0$ and let r_1 be the unique value such that $\phi_2(r_1) = 0$. Then it follows from (55) that $|r_2 - r_1| \leq \delta$. Under our mechanism personal reserve prices are set as per (6), and in the limit $W \rightarrow \infty$, the personalized reserve for agent 1 is r_1 and for agent 2 it is r_2 . Without loss of generality, suppose that $r_1 \geq r_2$.

We now derive two identities regarding quantities that will come up when we compute the revenue loss. First, note that when agent 1 wins the item under our mechanism, i.e., when $a_1 = 1$, agent 1 clears her reserve $v_1 \geq r_1$, and so $\phi_1(v_1) \geq \phi_1(r_1) \geq \phi_1(r_2) = 0$. We deduce that

$$\mathbb{I}(a_1 = 1) \Rightarrow \phi_1(v_1) \geq 0 \Rightarrow (z_+ - \phi_1(v_1))_+ = (z - \phi_1(v_1))\mathbb{I}(\phi_1(v_1) < z) \quad \text{for any } z \in \mathbb{R}, \quad (56)$$

where $\tilde{z}_+ \triangleq \max(\tilde{z}, 0)$ for $\tilde{z} \in \mathbb{R}$.

Similarly, if the item is not allocated $a_1 = a_2 = 0$, then the highest bidder does not clear her personal reserve and we can rewrite $\max(\phi_1(v_1), \phi_2(v_2), 0)$ — If agent 1 is the highest bidder, we have $v_2 < v_1 < r_1$ and so $\phi_2(v_2) < \phi_2(r_1) = 0$, and if $\phi_1(v_1) > 0$ then $v_1 > r_2$; overall $\max(\phi_1(v_1), \phi_2(v_2), 0) = \phi_1(v_1)\mathbb{I}(v_1 \in (\max(r_2, v_2), r_1))$. On the other hand if agent 2 is the highest bidder, we have $v_1 \leq v_2 < r_2$ and so $\phi_2(v_2) < \phi_2(r_2) \leq \phi_2(r_1) = 0$, and $\phi_1(v_1) < \phi_1(r_2) = 0$; overall $\max(\phi_1(v_1), \phi_2(v_2), 0) = 0$. Again $\max(\phi_1(v_1), \phi_2(v_2), 0) = \phi_1(v_1)\mathbb{I}(v_1 \in (\max(r_2, v_2), r_1))$ remains valid since $\mathbb{I}(v_1 \in (\max(r_2, v_2), r_1)) = 0$. Thus, we have established that

$$\mathbb{I}(a_1 = a_2 = 0) \Rightarrow \max(\phi_1(v_1), \phi_2(v_2), 0) = \phi_1(v_1)\mathbb{I}(v_1 \in (\max(r_2, v_2), r_1)).$$

In fact, we can write

$$\mathbb{I}(a_1 = a_2 = 0) \max(\phi_1(v_1), \phi_2(v_2), 0) = \phi_1(v_1)\mathbb{I}(v_1 \in (\max(r_2, v_2), r_1)), \quad (57)$$

since $\mathbb{I}(a_1 = a_2 = 0) = 0 \Rightarrow \mathbb{I}(v_1 \in (\max(r_2, v_2), r_1)) = 0$.

Bounding the revenue loss. We now proceed to bound the loss in revenue under our mechanism, relative to Myerson. Subtracting (51) from (52) and using that the allocation is deterministic under our mechanism, we obtain

$$\begin{aligned} \text{REV}_* - \text{REV} &= \int_0^\infty \int_0^\infty \left[\mathbb{I}(a_1 = 1) \left((\phi_2(v_2))_+ - \phi_1(v_1) \right)_+ \right. \\ &\quad \left. + \mathbb{I}(a_2 = 1) \left((\phi_1(v_1))_+ - \phi_2(v_2) \right)_+ \right. \\ &\quad \left. + \max(\phi_1(v_1), \phi_2(v_2), 0) \mathbb{I}(a_1 = a_2 = 0) \right] f_1(v_1) f_2(v_2) dv_1 dv_2, \\ &= \mathbb{E} \left[\mathbb{I}(\phi_1(v_1) < \phi_2(v_2) \wedge a_1 = 1) (\phi_2(v_2) - \phi_1(v_1)) \right. \\ &\quad \left. + \mathbb{I}(\phi_2(v_2) < (\phi_1(v_1))_+ \wedge a_2 = 1) ((\phi_1(v_1))_+ - \phi_2(v_2)) \right. \\ &\quad \left. + \phi_1(v_1) \mathbb{I}(v_1 \in (\max(r_2, v_2), r_1)) \right], \end{aligned}$$

where we made use of (56) to rewrite the first term and (57) to rewrite the third term. The first term captures revenue loss from the item being allocated to agent 1 though Myerson would not have (the revenue loss is $\phi_2(v_2) - \phi_1(v_1) = (\phi_2(v_2))_+ - \phi_1(v_1)$ when $\phi_2(v_2) > \max(\phi_1(v_1), 0)$ and so the item should instead have been allocated to agent 2, the loss is $-\phi_1(v_1) = (\phi_2(v_2))_+ - \phi_1(v_1) > 0$ when $\phi_1(v_1) < 0, \phi_2(v_2) < 0$ and so the item should not have been allocated at all, and there is no loss otherwise when $\phi_1(v_1) \geq \max(\phi_2(v_2), 0)$). The second term similarly captures revenue loss from wrongly allocating the item to agent 2. The third term captures revenue loss from not allocating the item, whereas Myerson would have allocated it.

Let us bound each of the terms. We start by considering the first two terms. For the first term to be positive we need event $E_1 \triangleq (\phi_1(v_1) < \phi_2(v_2) \wedge a_1 = 1)$ to occur. Now if E_1 occurs then $a_1 = 1 \Rightarrow v_1 > v_2$ so it must be that $\phi_2(v_1) > \phi_2(v_2) > \phi_1(v_1)$, and it also must be that $v_2 \in (v_1 - \delta, v_1)$ using (55), and of course $v_1 \geq r_1$. Similarly, for the second term to be positive we need event $E_2 \triangleq (\phi_2(v_2) < (\phi_1(v_1))_+ \wedge a_2 = 1)$. If E_2 occurs then $v_2 \geq \max(v_1, r_2)$ and one of the following two cases must arise:

1. $v_1 \geq r_2 \Rightarrow \phi_1(v_1) \geq 0$ in which case $\phi_2(v_1) \leq \phi_2(v_2) < \phi_1(v_1)$ and $v_2 \in (v_1, v_1 + \delta]$ using (55).
2. $v_1 < r_2 \Rightarrow \phi_1(v_1) < 0$ and hence $\phi_2(v_2) < 0$, implying $v_2 \in [r_2, r_1]$. Note that

$$\phi_1(v_2) \geq \phi_1(r_2) = 0 \Rightarrow \phi_2(v_2) \geq -\delta \quad (58)$$

follows, using $\|F_1 - F_2\| \leq \delta$.

Observe that the above events are mutually exclusive $E_1 \cap E_2 = \emptyset$, since they result in different agents being allocated the item. We will now show that $P(E_1 \cup E_2 | v_1) \leq f_{\max} \delta$ for any v_1 . If $v_1 < r_2$ this rules out event E_1 since $v_1 < r_2 < r_1$ and so $a_1 = 0$. For event E_2 to occur requires case 2 to arise, in particular $v_2 \in [r_2, r_1]$, which occurs with probability at most $f_{\max} \delta$ using $r_1 - r_2 \leq \delta$. Now consider the complementary case $v_1 \geq r_2$. Note that we need case 1 to arise in order for event E_2 to occur. Comparing $\phi_2(v_1)$ with $\phi_1(v_1)$ rules out one of the two events; if $\phi_2(v_1) \leq \phi_1(v_1)$ this rules out event E_1 , whereas $\phi_2(v_1) > \phi_1(v_1)$ rules out event E_2 . The event which has not yet been ruled out then requires v_2 to lie in an interval of length δ , and this occurs with probability at most $f_{\max} \delta$. Thus we established that $P(E_1 \cup E_2 | v_1) \leq f_{\max} \delta$ for all v_1 . We are now in a position to bound the expectation of the first two terms above. Each of them is pointwise bounded above by δ using (54) and (58). Hence, the sum of the expectations of the first two terms is bounded above (recalling $E_1 \cap E_2 = \emptyset$) by

$$\mathbb{E}_{v_1}[\mathbb{E}_{v_2}[\mathbb{I}(E_1 \cup E_2)\delta]] \leq \mathbb{E}_{v_1}[\delta P(E_1 \cup E_2 | v_1)] \leq \delta \mathbb{E}_{v_1}[f_{\max} \delta] = f_{\max} \delta^2.$$

We are left with the third term which requires $\mathbb{I}(v_1 \in (\max(r_2, v_2), r_1)) = 1$ to be non-zero. This term again is pointwise bounded above by δ since $\phi_2(v_1) \leq \phi_2(r_1) = 0 \Rightarrow \phi_1(v_1) \leq \delta$ similar to (58). Further, we need that $v_1 \in (r_2, r_1)$, which occurs with likelihood bounded above by $f_{\max} \delta$ using $r_1 - r_2 \leq \delta$, and hence the expectation of the third term is again bounded above by $f_{\max} \delta^2$.

Summing gives the bound of $2f_{\max} \delta^2$ on revenue loss, as stated in the proposition.

Extension to general n . We start with some preliminary observations. First, use Fact 5 to observe that the virtual value function corresponding to the average valuation distribution of any combination of agents must lie between the largest and smallest virtual value function for any of the individual agents. In particular, ϕ_{-i} lies between the largest and smallest ϕ_j for $j \neq i$. Using the hypothesis that $\|F_i - F_j\| \leq \delta$ for all agent pairs i, j , we then deduce that for any agent i we have

$$|\phi_i(v) - \phi_{-i}(v)| \leq \max_{i' \neq i} |\phi_i(v) - \phi_{i'}(v)| \leq \delta,$$

i.e., $\|F_i - F_{-i}\| \leq \delta$. For the limiting reserve price $r_i = \arg \max_r r(1 - F_{-i}(r))$ for agent i as $W \rightarrow \infty$, the first order condition tells us that $\phi_{-i}(r_i) = 0$. We deduce the analog of (58), namely, that for any $v_i \geq r_i$, we have

$$\phi_i(v_i) \geq \phi_i(r_i) \geq \phi_{-i}(r_i) - \delta = 0 - \delta = -\delta. \quad (59)$$

Define $\tilde{r}_i$ to be the Myerson optimal reserve price for agent i , i.e., the unique value satisfying $\phi_i(\tilde{r}_i) = 0$. Using (50) and $|\phi_i(r_i)| \leq \delta$, we are guaranteed

$$|r_i - \tilde{r}_i| \leq \delta. \quad (60)$$

Observe that (54) and (55) hold as before.

We now proceed to bound the loss in revenue. In bounding the contribution due to the item being wrongly allocated to some agent, we argue as follows. Reveal the valuations sequentially, starting with the agent $i_{\min} = \arg \min_i \tilde{r}_i$. We have $r_{i_{\min}} \geq \tilde{r}_{i_{\min}}$, which holds since $\phi_{-i_{\min}}(r_{i_{\min}}) = 0$ whereas for all $v < \tilde{r}_{i_{\min}}$ we have $\phi_j(v) < 0 \ \forall j \in i_{\min} \Rightarrow \phi_{-i_{\min}}(v) < 0$ by definition of $i_{\min}$ and using Fact 5. Note that agent $i_{\min}$ is never allocated the item when her virtual value is negative. For each valuation revealed after that of $i_{\min}$, declare a failure if there is a possibility of revenue loss based on the value just revealed as follows. Let v be the highest valuation so far, achieved by agent i_* . Suppose that the next valuation to be revealed is that of the j -th agent. If $\phi_{i_*}(v) < 0$, then we declare failure if $v_j \in [r_j, \tilde{r}_j]$ (if this occurs the item may be allocated to j , even though $\phi_j(v_j) < 0$, causing revenue loss). Using (60) we observe that this occurs with probability at most $f_{\max}\delta$. Else, if $\phi_{i_*}(v) \geq 0$, compare $\phi_j(v)$ with $\phi_{i_*}(v)$. If $\phi_j(v) > \phi_{i_*}(v)$, then it must be³⁰ that $v_j \in (v - \delta, v]$ for there to be a contribution to revenue loss (from not giving the item to j though Myerson would have allocated it to j), whereas if $\phi_j(v) \leq \phi_{i_*}(v)$, then it must be that $v_j \in [v, v + \delta]$ for there to be a contribution to revenue loss (from allocating j the item though Myerson would have allocated it to some other agent). Either way, the problematic case arises with probability at most $f_{\max}\delta$, and the revenue loss is at most δ using (54) and (59). Using a union bound over the $n - 1$ successive valuations revealed after the first one, we find that the probability of failure during the entire revelation process is at most $(n - 1)f_{\max}\delta$. A failure contributes a revenue loss of at most δ ; hence, the loss in expected revenue due to misallocation is bounded above by $(n - 1)f_{\max}\delta^2$.

Next, we can bound the loss in revenue due to the item not being allocated, though it should have been. This never occurs with the agent $i_{\max} = \arg \max_i \tilde{r}_i$ being the one who should have been allocated the item. The reason is that $r_{i_{\max}} \leq \tilde{r}_{i_{\max}}$, which holds since $\phi_{-i_{\max}}(r_{i_{\max}}) = 0$ whereas for all $v > \tilde{r}_{i_{\max}}$ we have $\phi_j(v) > 0 \ \forall j \in i_{\max} \Rightarrow \phi_{-i_{\max}}(v) > 0$ by definition of $i_{\max}$ and using Fact 5. For each of the other agents $j \neq i_{\max}$, this source of revenue loss requires the valuation v_j of the agent to satisfy $v_j \in [\tilde{r}_j, r_j]$, an interval of length at most δ by (60). Hence, this kind of failure occurs with probability at most $f_{\max}\delta$ for each of these agents, and probability at most $(n - 1)f_{\max}\delta$ overall. When this failure occurs, again the contribution to revenue loss is at most δ , since, for $v_j < r_j$, we have $\phi_j(v_j) < \phi_j(r_j) \leq \phi_{-j}(r_j) + \delta = 0 + \delta = \delta$ using $\|F_j - F_{-j}\| \leq \delta$. Hence, the revenue loss due to the item not being allocated, though it should have been, is bounded above by $(n - 1)f_{\max}\delta^2$.

Combining the above bounds, we obtain the required bound of $2(n - 1)f_{\max}\delta^2$ on the total additive revenue loss. $\square$

Remark 4 *A very similar analysis to the above shows that in a static setting with known asymmetric valuation distributions separated by at most δ , using a second-price auction with a common reserve (which can be set, e.g., according to any one of the valuation distributions F_i), the revenue obtained is (additively) within*

³⁰ We need $v_j \leq v$ else agent j would have been given the item over agent i_* , and there would be no revenue loss. We further use (55) and $\phi_j(v) > \phi_{i_*}(v)$ to infer that $v_j > v - \delta$.

$(n-1)f_{\max}\delta^2$ of the optimal revenue. This complements the worst-case result of Hartline and Roughgarden (2009), who show that with arbitrary asymmetric regular valuations, second-price auctions (with personalized reserves) suffice to obtain at least $1/2$ the optimal revenue.

D. Appendix to Section 6

Proof of Lemma 1. We follow the approach in Golrezaei et al. (2018a). The gradient and the Hessian of the loss function $\mathcal{L}(\cdot) = \mathcal{L}_{-i}(\cdot)$ are given by

$$\nabla \mathcal{L}(\tilde{\beta}) = \frac{1}{(n-1)W} \sum_{i' \neq i} \sum_{\tau=t-W}^{t-1} \mu_{i'\tau}(\tilde{\beta}) x_{\tau}, \quad \nabla^2 \mathcal{L}(\tilde{\beta}) = \frac{1}{W} \sum_{\tau=t-W}^{t-1} x_{\tau} x_{\tau}^T, \quad (61)$$

where $\mu_{i'\tau}(\tilde{\beta}) \triangleq 2(x_{\tau}^T \tilde{\beta} - b_{i'\tau})$. Since $\hat{\beta}_{-i}$ minimizes $\mathcal{L}(\cdot)$, we have $\mathcal{L}(\hat{\beta}_{-i}) \leq \mathcal{L}(\beta)$. Substituting a Taylor expansion of $\mathcal{L}$ around β on the left-hand side then gives

$$(1/2) (\hat{\beta}_{-i} - \beta)^T \nabla^2 \mathcal{L}(\tilde{\beta}) (\hat{\beta}_{-i} - \beta) \leq -(\nabla \mathcal{L}(\beta))^T (\hat{\beta}_{-i} - \beta), \quad (62)$$

for some $\tilde{\beta}$ on the segment between β and $\hat{\beta}_{-i}$. Notice that $\mu_{i'\tau}(\beta)$ are i.i.d., zero mean and bounded. Analogous to Golrezaei et al. (2018a, Lemma 7), we can use the Azuma-Hoeffding inequality (or Matrix Freedman inequality) to control $\nabla \mathcal{L}(\beta)$. In particular, there exists a constant $C_2 < \infty$ such that, with probability $1 - 1/(2W)$, it holds that

$$\|\nabla \mathcal{L}(\beta)\| \leq C_2 \sqrt{\frac{\log W}{W}}. \quad (63)$$

Next, notice that $\nabla^2 \mathcal{L}(\tilde{\beta})$ does not depend on $\tilde{\beta}$ and, in fact, is nothing but the sample covariance matrix of distribution G , which moreover has subgaussian rows since $\|x_{\tau}\| \leq 1$ for all τ . Following the steps Golrezaei et al. (2018a, (71)-(73)), we can show (relying upon convergence of the sample covariance matrix to the true covariance matrix Σ) that there exists $C_3 < \infty$ such that for any $W > C_3$, with probability at least $1 - 1/(2W)$, we have

$$(1/2) (\hat{\beta}_{-i} - \beta)^T \nabla^2 \mathcal{L}(\tilde{\beta}) (\hat{\beta}_{-i} - \beta) \geq \|\hat{\beta}_{-i} - \beta\|^2 / C_3. \quad (64)$$

Upper bounding the right-hand side of (62) by $\|\nabla \mathcal{L}(\beta)\| \|\hat{\beta}_{-i} - \beta\|$ and substituting (63) and (64), we obtain via a union bound that for any $W > C_3$, with probability at least $1 - 1/W$ we have

$$\|\hat{\beta}_{-i} - \beta\|^2 / C_3 \leq C_2 \sqrt{\frac{\log W}{W}} \|\hat{\beta}_{-i} - \beta\|. \quad (65)$$

The lemma follows from dividing both sides by $\|\hat{\beta}_{-i} - \beta\|/C_3$ and then choosing $C \triangleq \max(C_2 C_3, 2B_{\beta} \sqrt{C_3 / \log C_3})$ (the second term in the max ensures that the bound holds for all $W \leq C_3$ as well, since $\|\hat{\beta}_{-i} - \beta\| \leq 2B_{\beta}$ given $\|\hat{\beta}_{-i}\| \leq B_{\beta}, \|\beta\| \leq B_{\beta}$). $\square$

Proof of Theorem 3. We make use of Lemma 1 to conclude that with probability $1 - 1/W$, (12) holds for agent i , i.e., our estimate $\hat{\beta}_{-i}$ is close to β . The loss in expected revenue from the possibility that (12) fails

for one or more agents is at most $E_{x_t}[\text{REV}_*(x_t)]/W \leq (B_\beta + B_F)n/W$, since the probability of failure is at most n/W , and $\text{REV}_*(x_t) \leq B_\beta + B_F$ for all $x_t : \|x_t\| \leq 1$. Henceforth, assume that (12) holds. Then $\hat{F}_{-i}^{x_t}$ is the empirical distribution of translated samples, each of which deviate by at most $\delta = 2C_1\sqrt{\log W/W}$ from being true samples from F^{x_t} , since the deviation of each sample based on bids in period τ is $(\hat{\beta}_{-i} - \beta)^T(x_t - x_\tau) \leq 2\|\hat{\beta}_{-i} - \beta\| \leq \delta$. Let

$$r_{t,*} \triangleq \arg \max_r r(1 - F^{x_t}(r)),$$

$$\hat{r}_{it,*} \triangleq \arg \max_r r(1 - \hat{F}_{-i}^{x_t}(r)).$$

By definition of $\hat{r}_{it,*}$, we have

$$\hat{r}_{it,*}(1 - \hat{F}_{-i}^{x_t}(\hat{r}_{it,*})) \geq (r_{t,*} - \delta)(1 - \hat{F}_{-i}^{x_t}(r_{t,*} - \delta)). \quad (66)$$

Using Lemma 2, we have that with probability at least $1 - 1/W$,

$$\min(1, 1 - F^{x_t}(\hat{r}_{it,*} - \delta) + C_8\sqrt{\log W/W}) \geq 1 - \hat{F}_{-i}^{x_t}(\hat{r}_{it,*}), \quad \text{and} \quad (67)$$

$$1 - F^{x_t}(r_{t,*}) - C_8\sqrt{\log W/W} \leq 1 - \hat{F}_{-i}^{x_t}(\hat{r}_{it,*} - \delta), \quad (68)$$

where we also used $\hat{F}_{-i}^{x_t}(\hat{r}_{it,*}) \geq 0 \Rightarrow 1 - \hat{F}_{-i}^{x_t}(\hat{r}_{it,*}) \leq 1$ to obtain (67). Using (67) to upper bound the left-hand side of (66) and using (68) to lower bound the right-hand side of (66), we obtain that with probability at least $1 - 1/W$,

$$\hat{r}_{it,*} \min(1, 1 - F^{x_t}(\hat{r}_{it,*} - \delta) + C_8\sqrt{\log W/W}) \geq (r_{t,*} - \delta)(1 - F^{x_t}(r_{t,*}) - C_8\sqrt{\log W/W}). \quad (69)$$

Observe that the left-hand side of (69) is at most

$$\begin{aligned} & (\hat{r}_{it,*} - \delta)(1 - F^{x_t}(\hat{r}_{it,*} - \delta) + C_8\sqrt{\log W/W}) + \delta \\ &= \text{MONREV}_{r_{it}}(x_t) + (\hat{r}_{it,*} - \delta)C_8\sqrt{\log W/W} + \delta \quad (\text{recalling } r_{it} = \hat{r}_{it,*} - \delta) \\ &\leq \text{MONREV}_{r_{it}}(x_t) + (B_F + B_\beta)C_8\sqrt{\log W/W} + \delta \quad (\text{since } \hat{r}_{it,*} < B_F + B_\beta + \delta), \end{aligned}$$

where $\text{MONREV}_r(x_t) \triangleq r(1 - F^{x_t}(r))$ is the expected revenue of a monopolist who sells to a single buyer using price r (conditioned on the past bids). The right-hand side of (69) is at least $r_{t,*}(1 - F^{x_t}(r_{t,*})) - \delta - (B_F + B_\beta)C_8\sqrt{\log W/W} = \text{MONREV}_{r_{t,*}}(x_t) - \delta - (B_F + B_\beta)C_8\sqrt{\log W/W}$, by definition of $r_{t,*}$. It follows immediately that

$$\begin{aligned} & \text{MONREV}_{r_{it}}(x_t) + \delta + (B_F + B_\beta)C_8\sqrt{\log W/W} \geq \text{MONREV}_{r_{t,*}}(x_t) - \delta - (B_F + B_\beta)C_8\sqrt{\log W/W} \\ \Rightarrow & \text{MONREV}_{r_{it}}(x_t) \geq \text{MONREV}_{r_{t,*}}(x_t)(1 - \tilde{\delta}) \\ & \text{where } \tilde{\delta} \triangleq 2(\delta + (B_F + B_\beta)C_8\sqrt{\log W/W})/\text{MONREV}_{r_{t,*}}(x_t), \end{aligned}$$

as long as (12) holds. Exactly as we showed (35) in the proof of Theorem 1, we extend this multiplicative bound to the actual revenue under multiple bidders:

$$\text{MONREV}_{\max(r_{it}, b)}(x_t) \geq \text{MONREV}_{\max(r_{t,*}, b)}(x_t)(1 - \tilde{\delta}) \quad \text{for all } b \in \mathbb{R}. \quad (70)$$

Now setting b to be the largest bid submitted by an agent other than i and taking expectation over the bids of the other agents, we get

$$\begin{aligned} \mathbb{E}[\text{Revenue from } i \text{ under proposed mechanism}] &\geq \mathbb{E}[\text{Revenue from } i \text{ under Myerson}](1 - \tilde{\delta}) \\ \Rightarrow \mathbb{E}[\text{Revenue under proposed mechanism}] &\geq \text{REV}_*(x_t)(1 - \tilde{\delta}), \end{aligned}$$

where the second step follows by summing over i . Thus, incorporating the $1/W$ probability of failure of (67) and/or (68), the overall additive loss in expected revenue under our mechanism (relative to Myerson) is bounded above by

$$\begin{aligned} \text{REV}_*(x_t)\tilde{\delta} + (B_\beta + B_F)n/W &= 2(\delta + (B_F + B_\beta)C_8\sqrt{\log W/W})\text{REV}_*(x_t)/\text{MONREV}_{r_{t,*}}(x_t) + (B_\beta + B_F)n/W \\ &\leq 2(\delta + (B_F + B_\beta)C_8\sqrt{\log W/W})n + (B_\beta + B_F)n/W. \end{aligned}$$

The theorem follows. $\square$

E. Appendix to Section 7

Proof of Theorem 4. To prove the first part, let us assume that all agents are truthful. Note that the loss of revenue from the first phase is constant, and so its per-auction contribution vanishes in the long run. We now show that the loss in revenue from agents refusing to pay the up-front payment is $O(N^{-2})$. Using Eq. (18), since z_{it}/n are distributed i.i.d. and are bounded between 0 and 1, by Hoeffding’s inequality (Hoeffding 1963), we have

$$\Pr \left[\left| \mu_i - \hat{\mu}_i \right| \geq \sqrt{\frac{2 \log N}{N}} \right] = O(N^{-2}).$$

Therefore, with probability $1 - O(N^{-2})$, in each round all agents pay their entrance fees and the per-round loss from “refusals” is bounded by $O(N^{-2})$. Also, for the auctions where all the agents pay their up-front payment, the per-round loss of revenue relative to the optimal is bounded by the slack in the entrance fees, which is $O\left(\sqrt{\frac{\log N}{N}}\right)$. This yields an overall $O\left(\sqrt{\frac{\log N}{N}}\right)$ bound on loss in revenue relative to the optimal.

We now prove the incentive compatibility result. As mentioned before, the only benefit from deviating from the truthful strategy would be to eliminate some of the bidders from the second-phase auctions and to instead compete with their simulated versions. To establish the incentive compatibility properties of our mechanism, let us consider the following augmented action space for an agent i . Suppose that at the beginning of each round, bidder i can choose sets S and $\bar{S}$. Namely, for each agent i' , she can decide whether she wants to compete against the actual bidder or against her simulated version that uses bids from the first phase. This setting provides an upper bound on the benefit from deviating for each agent i . Nevertheless, we show that the upper bound is $O(\sqrt{\log N/N})$ per round.

For $t \leq N$, let us define $y_{i,t,\bar{S}} = \mathbb{E}_{\bar{S},\bar{S}}[\max\{V_{it} - \max\{\max_{i' \in S \setminus \{i\}}\{V_{i't}\}, \max_{i' \in \bar{S}}\{b_{i't}\}, 0\}\}]$, where for agents $i' \in S$ (including i), random variable $V_{i't}$ is drawn i.i.d. from distribution F . For $i' \in \bar{S}$, we use their bids at

period t . Note that $E[y_{it}] = \mu_i$ when the expectation is over the bids in $\bar{S}$ and all the agents in $\bar{S}$ are truthful. Therefore, using Hoeffding's inequality, we get

$$\Pr \left[\left| \mu_i - \frac{1}{N} \sum_{i=1}^N y_{i,t,\bar{S}} \right| \geq \sqrt{\frac{2 \log N}{N}} \right] = O(N^{-2}).$$

If this event $\left| \mu_i - \frac{1}{N} \sum_{i=1}^N y_{i,t,\bar{S}} \right| \geq \sqrt{\frac{2 \log N}{N}}$ occurs, then dropping agents in $\bar{S}$ may increase utility by more than $\sqrt{\frac{2 \log N}{N}}$ per round, but not otherwise. Note that there are at most 2^n choices for $\bar{S}$. We can use a simple union bound since n is a constant to show that with probability $1 - O(N^{-2})$, the per-round benefit from deviating from the truthful strategy is bounded by $O\left(\sqrt{\frac{\log N}{N}}\right)$. $\square$