SCALABLE OPTIMAL ONLINE AUCTIONS

DOMINIC COEY, BRADLEY J. LARSEN, KANE SWEENEY, AND CAIO WAISMAN

ABSTRACT. This paper studies reserve prices computed to maximize the expected profit of the seller based on historical observations of the top two bids from online auctions in an asymmetric, correlated private values environment. This direct approach to computing reserve prices circumvents the need to fully recover distributions of bidder valuations. We specify precise conditions under which this approach is valid and derive asymptotic properties of the estimators. We demonstrate in Monte Carlo simulations that directly estimating reserve prices is faster and, outside of independent private values settings, more accurate than fully estimating the distribution of valuations. We apply the approach to e-commerce auction data for used smartphones from eBay, where we examine empirically the benefit of the optimal reserve as well as the size of data set required in practice to achieve that benefit. This simple approach to estimating reserves may be particularly useful for auction design in Big Data settings, where traditional empirical auctions methods may be costly to implement, whereas the approach we discuss is immediately scalable.

Date: December 2, 2020.
We thank the editor, the associate editor and two anonymous referees for their detailed comments and feedback, and Isa Chaves, Han Hong, Dan Quint, Evan Storms, and Anthony Zhang for helpful comments. This paper was previously circulated under the title "The Simple Empirics of Optimal Online Auctions." Coey and Sweeney were employees of eBay Research Labs while working on this project, and Larsen was a part-time contractor when the project was started.
Coey: Facebook, Core Data Science; coey@fb.com.
Larsen: Stanford University and NBER; bjlarsen@stanford.edu.
Sweeney: Uber, Marketplace Optimization; kane@uber.com.
Waisman: Northwestern University; caio.waisman@kellogg.northwestern.edu.

1. Introduction

Auctions are a key selling mechanism in online markets. Not only are they employed to sell physical goods through sites such as eBay and Tophatter, and many others, but also to offer advertisement space online, such as through Google's DoubleClick.[1] A large structural literature in economics has studied these auctions, but this literature offers little practical advice for sellers wishing to determine the optimal reserve price. The traditional approach to structurally analyzing auctions in the economics literature is to estimate the full distribution of bidder valuations and only then compute counterfactual objects of interest, such as the optimal reserve price or other instruments of auction design. While off-the-shelf tools exist to estimate valuation distributions in many first-price auction settings (e.g. Guerre et al. 2000), such tools do not immediately apply to online auction settings, where the researcher/seller may not observe all bids or the number of bidders. Existing tools for estimating valuation distributions in online auctions are computationally involved and have only been derived for limited settings, such as independent private values (IPV) settings with symmetric bidders (e.g. Song 2004). In this paper, we adopt an approach advocated in the computer science and statistical learning literature (e.g. Mohri and Medina 2016): directly estimating optimal reserve prices without fully estimating underlying distributions of bidder valuations. Such a direct approach is simple for sellers to execute and scalable to large datasets and real-time computation without fully estimating distributions of bidder valuations.

The primary information environment we consider is a private values setting with bidders values being *non-independent* and potentially *asymmetric*. Our main focus is on online auctions that follow a second-price-auction-like format but with *sequential arrival* of bidders, such as those conducted on eBay. In such auctions, a bidder who arrives at the auction after the standing bid has passed her valuation is not observed bidding. Hence,

[1]Given the large size of eBay alone, even with the decline in auction usage in recent years relative to posted prices or bargaining (Einav et al. 2018; Coey et al. 2020; Backus et al. 2020), auctions remain a popular mechanism among many eBay users; as documented in Coey et al. (2020), as of 2014, eBay auctions annually sold over 80 million new-in-box items via auctions annually, yielding 1.6 billion dollars in annual revenue. These numbers are even larger for used auctions. In some categories, such as event tickets, auctions remain the dominant selling mechanism (see Waisman 2020). In addition to eBay and Tophatter, other similar online auction sites include eBid, ShopGoodwill, Webstore, Auction Zip, Catawiki, and Bid For Wine. Online auctions are also used for some government surplus or government/police-apprehended items on sites such as GovdDeal, PropertyRoom, Municibid, and IRS auctions.

some bids, as well as the true number of would-be bidders, is unobserved to the econometrician.[2] This data limitation precludes the use of popular nonparametric approaches for point identifying or partially identifying valuation distributions in ascending auctions, which rely on inverting order statistic distributions (e.g., Haile and Tamer 2003; Athey and Haile 2007; Aradillas-López et al. 2013; Coey et al. 2017). Indeed, the full distribution of bidder valuations is not identified in the setting we study, but the optimal reserve price itself *is* identified. Specifically, in eBay-like auctions, a bid placed by the highest bidder is recorded (unlike in traditional ascending/English auctions, where the auction ends without the winner's drop-out price being recorded). This opens up a direct approach to obtaining the information required for computing optimal reserve prices. Our starting point is the observation that seller profit is a simple-to-compute function of the reserve price and the two highest bids—and thus profit can be computed from only the top two bids and without observing the number of bidders and without observing other losing bids.

The seller's primary instrument of auction design in the real world is typically a single reserve price.[3] We demonstrate that, under mild conditions, the optimal single reserve price can be computed by simply maximizing the seller expected profit function (as a function of the empirical distribution of historical observations of the two highest bids) with respect to the reserve price. Throughout the paper, we treat the seller as synonymous with the practitioner/researcher/econometrician.

We derive several properties of estimated reserve prices. We prove consistency and derive the asymptotic distribution of the estimator of the optimal reserve price, which is non-normal. Our interest in deriving these results is not just theoretical but also practical: we wish to provide a clear notion of *how many previous transactions* the practitioner needs to observe in order for it be the case that designing an auction based on estimated reserves is a good idea, as well as guidance to perform statistical inference. In this spirit, building on the work of Mohri and Medina (2016), we also derive an explicit lower bound, based

[2]A number of aspects of our analysis can also be applied to ad auctions, which have traditionally been run as second-price auctions (although a number of platforms have recently switched to first-price auctions) and where the true number of bidders may be unobserved due to the practice of marketing agencies bidding on behalf of multiple bidders at once (Decarolis et al. 2020).

[3]In theory, more complex auction design may be optimal—such as the Myerson (1981) auction for asymmetric independent private values (IPV) settings—involving more than just a simple reserve price. And, in practice, other features of the auction beyond the reserve price can also impact revenue, such as the number of bidders (Bulow and Klemperer 1996; Coey et al. 2019) or the skill of the auctioneer (Lacetera et al. 2016). We do not focus on those features here.

on only weak assumptions, for the number of auctions one would need to observe in order to guarantee that the revenue based on the estimated reserve price approximates the true optimal with a specified degree of accuracy.

In Monte Carlo simulations we demonstrate that the directly estimated optimal reserve prices perform well relative to the approach of fully estimating the distribution of bidder valuations (Song 2004). The latter approach is correct only if the true environment is one of symmetric independent private values (IPV), whereas our approach does not rely on this assumption. Furthermore, our approach is much faster and does not require specifying any parametric approximation for the valuation distribution.

We then take a step beyond asymptotic and learning theory and examine empirically the number of auctions one needs to observe in practice in order to achieve a level of revenue close to the true optimal reserve price revenue. We analyze a sample of popular smartphone products sold through eBay auctions. We find that implementing the optimal reserve price in these settings would raise expected profit very little (less than 1%) compared to an auction with a reserve price equal to the seller's value. In contrast, in a setting in which the seller plans to re-auction the item if the auction fails, the gains from optimal reserve pricing are much larger (22–44%, depending on the product). In the data we find that a historical dataset of less than 25 auctions is sufficient for estimating a reserve price that will achieve a level of revenue that is within 1% of the true optimal-reserve revenue. We also find that, from the perspective of a seller offering a limited quantity of inventory (250 phones in our example), the seller would find it optimal to run fewer than 25 auctions for the sole purpose of data collection before implementing a reserve price based on that data. In online auctions for advertising or in e-commerce auctions for popular products, such data requirements are not likely to be restrictive.

2. Related literature

Our work contributes to the literature studying empirical approaches for ascending-like auctions. These methodologies generally rely on exploiting order statistics relationships.[4]

[4]A separate branch of the empirical auctions literature studies first-price auctions, such as Guerre et al. (2000) and a large work that builds on their approach, which exploits the relationship between equilibrium bidding

For example, in a symmetric, independent private values (IPV) button ascending auction, the distribution of valuations is identified by inverting the distribution of the winning bid (the second order statistic of valuations), as described in Athey and Haile (2002, 2007). Performing such an inversion requires observing the number of would-be bidders. This object is unobserved in eBay auctions. Song (2004) overcomes this challenge by pointing out that, in a symmetric IPV online auction, *two* order statistics of bids are sufficient to identify the underlying distribution of bidders, because the distribution of one order statistic conditional on a lower order statistic does not depend on the number of bidders. This profound insight led Hortaçsu and Nielsen (2010) to argue that the Song (2004) result has long been "the standard to beat in the empirical online auctions literature" due to its distinct ability to identify the distribution of valuations when the number of bidders is unknown.

Like the Song (2004) approach, our method relies on observing two order statistics of bids (in our case, the first and second highest). Our results are less useful than those of Song (2004) in one sense, because we do not obtain identification of the distribution of values, but our results are more useful in another sense, because we obtain identification of one particular object of interest—the optimal reserve price—in environments that are more general than the symmetric IPV auction of Song (2004). In particular, our results yield identification of the optimal reserve price in online auctions with arbitrarily correlated (and asymmetric) private values.

Song (2004) is the only approach of which we are aware that yields identification of the full distribution of valuations when only two order statistics of bids are observed and the number of bidders in unobserved. Other methodologies require stronger assumptions on the information environment, such as assumptions on the form of any auction-level unobserved heterogeneity (Freyberger and Larsen 2019; Luo and Xiao 2020), or stronger assumptions on the data, such as partial knowledge of the distribution of the number of bidders (Larsen and Zhang 2018; Hernández et al. 2020; Larsen 2020), observability of exogenously varying reserve prices (Freyberger and Larsen 2019), observability of more than two bids or an instrument (Mbakop 2017), or observability of the number of bidders (Luo and Xiao 2020). Our approach does not require any of these assumptions, but in

functions and valuations in first-price auctions. These approaches do not apply to ascending-like auction environments.

exchange we only obtain identification of the optimal reserve price, not the full distribution of valuations.

In the economics literature, the recent work of Aradillas-López et al. (2013) provides identification arguments for *bounds* on the optimal reserve price in symmetric correlated private values environments using information on the number of bidders and only one bid (the transaction price), and Coey et al. (2017) extends these arguments to asymmetric environments. However, these approaches yield only *partial*—rather than *point*—identification, and require the econometrician to observe the number of bidders. In this paper we instead place a stronger data requirement on bids, requiring the *top two* bids be observed, and in exchange we obtain point identification.

Our paper also relates to other empirical auction work in economics that, rather than focusing on identification of the full distribution of valuations, provides direct inference on identification of specific objects of interest, such as optimal reserve prices, seller profits, bidder surplus, gains from optimal auction design, or losses from bidder mergers or collusion. These papers include Li et al. (2003), Haile and Tamer (2003), Tang (2011), Aradillas-López et al. (2013), Chawla et al. (2014), Coey et al. (2017), and Coey et al. (2019). In this sense, our work is also related to the broader movement in empirical economics work advocating for "sufficient statistics" for welfare analysis (Chetty 2009). This literature focuses on obtaining robust implications for welfare, optimality, or other objects of interest from simple empirical objects without requiring the type of detailed estimation typical of structural econometrics. In our setting, the sufficient statistics for the optimal reserve price are the marginal distributions of the first and second-highest bids.[5] Finally, our work is also related to that of Ostrovsky and Schwarz (2016), where the authors experimentally vary reserve prices in position ad auctions to measure the improvement in profits from choosing different reserve prices.

Our study contributes to the theoretical literature at the intersection of economics and computer science. For example, a number of papers, surveyed by Roughgarden (2014),

[5]We also relate to other empirical work focusing specifically on online e-commerce auctions, such as Platt (2017). In our analysis we consider only private values settings; in theory work, Abraham et al. (2020) model ad auctions with common values. Our analysis in the Online Appendix, where we extend our approach to address the Myerson (1981) auction empirically, is related to Celis et al. (2014), which addresses non-regular distributions and Myerson's ironing in ad auctions, while we focus only on regular distributions.

examine *approximately optimal* auctions. One of the motivations for the focus on approximate optimality is that auction features, most notably reserve prices, are often estimated from data and are therefore subject to sampling error. Thus, a lot of attention has been devoted to deriving theoretical learning guarantees to approximate optimal auctions from data as in Cole and Roughgarden (2014). Mohri and Medina (2016) propose algorithms to compute reserves from historical data, incorporating auction-level covariates/features, and these have been extended by Rudolph et al. (2016). A different approach consists of tailoring algorithms to continuously update estimated reserve prices while data are collected, an approach pioneered by Cesa-Bianchi et al. (2015) for symmetric IPV settings. Other continuously updating methods include Austin et al. (2016) and Rhuggenaath et al. (2019), both of which propose smooth approximations for the sellers profit function, deriving methods for handling high-dimensional auction-level observable features. We do not consider controlling for auction-level observable features or dynamic updating of reserve prices here. Kanoria and Nazerzadeh (2020) consider a model where bidders account for the fact that their bids may be used to compute future reserve prices; we do not explicitly model this possibility.[6]

The most closely related study to ours is that of Mohri and Medina (2016). Our approximation of profits and the corresponding optimal reserve prices are equivalent to theirs, as both papers estimate reserve prices using historical observations of the first- and second-highest bids. Both studies propose searching for the optimal reserve price using a grid search, and Mohri and Medina (2016) provide specific guidelines on how to perform this search quickly. In one important dimension, the work of Mohri and Medina (2016) is more general than ours, as the authors discuss controlling for auction-level observable covariates in estimating the optimal reserve price, and we do not consider auction-level covariates here; rather, we limit to commodity-like products and then obtain reserve prices that average over any remaining auction-level heterogeneity, observable or unobservable to the econometrician. Our approach is valid even in the presence of auction-level heterogeneity (observable or unobservable), but our approach averages over this heterogeneity rather than controlling for observable heterogeneity as in Mohri and Medina (2016).

[6]More broadly, our work also relates to the growing literature considering sophisticated pricing strategies and algorithms available to sellers with access to Big Data (e.g. Bounie et al. 2020; Kehoe et al. 2020; Ali et al. 2019).

Our contribution relative to Mohri and Medina (2016) is threefold. First, we offer a treatment of asymptotics and inference for the estimated reserve price, which Mohri and Medina (2016) do not. Second, we describe underlying conditions on the auction environment that are sufficient for the approach to be valid: a sequential-arrival second-price auction with private values. In contrast, Mohri and Medina (2016) provide their results in terms of bids and remain agnostic about the model generating those bids. In being explicit about these conditions, we are also able to offer examples of specific settings in which the approach is *not* valid, such as common values environments, dynamic auctions (i.e. settings in which bidders can participate in another auction if they fail to win the current auction), and certain types of endogenous entry. These are discussed in Section 3.6. Third, we offer a comparison to the alternative of fully estimating the distribution of valuations.

Algorithmic-based studies in computer science and empirical auction studies in economics tend to be quite distinct, and the two strands of literature do not often speak to one another. A contribution of our work relative to each of the above computer science studies is to help bridge this gap by bringing an econometrics and economic theory point of view to the computer science algorithmic literature.

3. Model and Empirical Approach

3.1. **Model Overview.** We consider bidders in an eBay-like auction, which we will refer to as a *sequential-arrival second-price auction*. Each auction has a number of bidders N that will arrive to the auction. N is a random variable potentially varying from auction to auction. We do not model bidders' choice of which of multiple auctions to participate in; rather, we treat bidders as being assigned to one and only one auction, and we assume that after the auction all N bidders assigned to that auction (the winner and the losers) exit the market. The seller of the auction sets a reserve price, $r \geqslant 0$.[7] In a given auction, bidder i, where $i \in \{1, ..., N\}$, has a valuation $V_i \geqslant 0$ for the item. Let $\mathcal{F}$ denote the joint distribution of bidder valuations from which N bids are drawn. Bidders arrive sequentially to the auction and have a single opportunity to bid.[8] Bidders see the current second-highest bid or, if no bids have been placed exceeding r, bidders see r. At some set time, the auction ends. As

[7] eBay uses the term "start price" to refer to a public reserve price and the term "reserve price" to refer to a secret reserve price. We will simply use the term "reserve price" except when explicitly discussing a feature unique to secret reserve prices, in which case we will use the term "secret reserve price".

[8] This single-opportunity-to-bid assumption can be relaxed following the intuition in the random-bidding-opportunities model of Song (2004).

long as at least one bidder bids above the reserve price, the highest bidder wins and pays the maximum of the second-highest bid and the reserve price.

We assume bidders' valuations are private, and we allow for bidders' valuations within a given auction to be arbitrarily correlated and to be potentially drawn from different marginal distributions.[9] Thus, the environment we consider is one of asymmetric correlated private values. In this private values environment it is a weakly dominant strategy for each bidder to bid her value when arriving at the auction if the current price is below her value, and we assume all bidders play this weakly dominant strategy. This abstracts away from any bid-sniping. We also abstract away from any cost of entering the auction and placing a bid, and from minimum bid increments.

Each bidder who arrives at the auction before the current price exceeds her valuations places a bid, and that bid is recorded by the platform. Each bidder who arrives *after* the bidding passes her valuation places no bid, and the platform has no information on her valuation. Under our modeling assumptions, in any auction in which at least two bidders have valuations weakly higher than r, the top two highest valuations among the N bidders will always be recorded, but lower order statistics of valuations may not be. For example, the third-highest-value bidder will not be observed bidding in an auction in which the two highest-valuation bidders arrive and bid their valuations *before* the third-highest-valuation bidder arrives. Consider another example: $r = 0$, and 10 bidders arrive at the auction. The fourth-highest bidder arrives first, followed by the first-highest, followed by the second-highest, followed by all other bidders in any arbitrary order. Only three bids are recorded because all other bidders arrive after the bidding has passed their valuations. However, among those three recorded bids, only the top two represent corresponding order statistics of valuations: the third-highest bid in this example is equal to the *fourth-highest* valuation.

Our approach relies only on observing the top two bids, not other losing bids and not the underlying number of bidders N. We also remain agnostic as to the precise arrival order of the bidders. Under our above assumptions, the top two valuations will be correctly observed regardless of the other details of the arrival process of bidders.

[9] In the Online Appendix we consider the asymmetric IPV environment of Myerson (1981). In this environment, if, in addition to the two highest bids, the econometrician observes the identity of the highest bidder, then bidder-specific marginal revenue curves (as defined by Bulow and Roberts 1989) are identified and estimable. These objects are then sufficient to implement the Myerson (1981) optimal auction. We use simulated data to illustrate this approach and quantify the revenue gain from optimal auction design.

We denote the highest and second-highest values by $V^{(1)}$ and $V^{(2)}$. In any auction in which only one bidder arrives, we consider $V^{(2)}$ to be equal to 0, and in any auction in which no bidders arrive, we consider $V^{(1)} = V^{(2)} = 0$. Thus, $V^{(1)}$ and $V^{(2)}$ refer to the top two order statistics of valuations within a given auction, whereas the set $\{V_i\}_{i=1}^N$ is the unordered set of bidders' valuations.

The seller's profit from setting a reserve price r in such an auction is given by

$$\pi(V^{(1)}, V^{(2)}, v_0, r) = r\mathbb{1}(V^{(2)} < r \leqslant V^{(1)}) + V^{(2)}\mathbb{1}(r \leqslant V^{(2)}) + v_0\mathbb{1}(r > V^{(1)}), \tag{1}$$

where v_0 indicates the seller's value from keeping the good. Expected profits as a function of the reserve are given by:

$$p(r) \equiv \mathbb{E}\left[\pi(V^{(1)}, V^{(2)}, v_0, r)\right], \tag{2}$$

where we suppress dependence of $p(\cdot)$ on v_0 for notational simplicity. The expectation is taken with respect to the distribution of $V^{(1)}$ and $V^{(2)}$. We use the term *optimal reserve price* to refer to a reserve price maximizing $p(r)$, and denote such a reserve by r^*.

3.2. **Assumptions and Main Result.** As highlighted above, in our environment, it is a weakly dominant strategy for bidders to bid their valuations if their valuations are above the current auction price. We explicitly assume that this is the bidding strategy that bidders follow. We summarize this and all of our key model assumptions below:

Assumption 1.
(i) Bidders have private valuations.
(ii) Bidders play the weakly dominant strategy of bidding their valuations if their valuations are above the current auction price.
(iii) Bidders are assigned to one auction and have one opportunity to bid in that auction.
(iv) Bidders face no cost of entry or bidding in the auction.
(v) $\Pr(N = n|r) = \Pr(N = n)$ *for all r and all n, and $\mathcal{F}$, the joint distribution of valuations, does not vary with r.*

Conditions (i)–(iv) are discussed above. Condition (v) ensures that if the auction designer were to change the reserve price (from $r = 0$ to $r = r^*$, for example) this would not change the unobserved number of bidders matched to the auction or the distribution from which they draw their values. Under these assumptions, the following identification result is immediate:

Remark 1. *Under Assumption 1, the optimal reserve price, r^*, is identified from the marginal distributions of the first and second-highest bids from auctions with $r = 0$.*

We can instead state a weaker (and less explicit) assumption under which Remark 1 holds. First, we define the concept of a *would-be bid*, which is the bid a bidder places in an auction or, if she is unable to place a bid because of the current price or reserve price exceeding her valuation, it is the bid she *would have* placed if not prevented from doing so.

Assumption 2. *The marginal distributions of the first and second-highest would-be bids do not vary with r.*

Under this assumption, Remark 1 will also hold. This assumption means that the data generating process for the historical auctions observed by the econometrician prior to implementing the reserve price is the same as the data generating process for the would-be bids after the reserve price is implemented. What we need for the estimated reserve price to be valid is that the distributions of bids (more specifically, the top two bids) does not change on the *intensive margin* when the reserve price is implemented, only on the extensive margin. In other words, the reserve price can determine *whether* a bidder places their bid, but not *what* that bid is.[10] Settings where our approach is not valid (discussed in Section 3.6) are violations of this assumption.

We assume that the econometrician has access to a random sample of $j = 1, ..., J$ auctions with zero reserve prices.[11] For each auction j, we assume that the top two highest bids, $V_j^{(1)}$ and $V_j^{(2)}$, are observed. With these data, we construct the sample analog of $p(\cdot)$ and its maximizer, which is our estimator of r^*:

$$\hat{p}(r) \equiv \frac{1}{J} \sum_{j=1}^{J} \left[\pi(V_j^{(1)}, V_j^{(2)}, v_0, r) \right]. \tag{3}$$

$$\hat{r}_J \equiv \arg\max_r \hat{p}(r). \tag{4}$$

[10]This same assumption, although not explicitly stated, also underlies the optimal reserve price exercises in other empirical auction work, such as Haile and Tamer (2003) and Aradillas-López et al. (2013), and the computer science literature cited in Section 2.

[11]Hence, the data are assumed to be independently and identically distributed across auctions.

We will also use the plug-in estimator $\hat{p}(\hat{r}_J)$ as our estimator of $p(r^*)$.

The reserve price estimated by this maximizer is the correct optimal reserve price under the assumptions of our model. Mohri and Medina (2016) advocate this same direct approach: maximizing expected seller profit directly using historical observations of the first- and second-highest bids rather than fully estimating the underlying distribution of seller valuations.

3.3. **Unobserved Auction-Level Heterogeneity.** We now comment briefly on auction-level unobservable heterogeneity—features of bidder valuations that are observable to all bidders (and potentially the seller) and that are not observable to the econometrician. Our proposed approach allows for the presence of such unobservable heterogeneity, as it simply introduces correlation between $V^{(1)}$ and $V^{(2)}$. If such unobservables are present, the estimated reserve price from our approach simply averages over these unobservables, yielding the optimal *unconditional* reserve price. What our approach does *not* allow for is the estimation of the *full distribution* of bidder valuations or the *full distribution* of unobserved heterogeneity. Other approaches do exist in the literature for estimating the full distribution of bidder valuations and the full distribution of unobserved heterogeneity (such as Freyberger and Larsen 2019), typically relying on convolution theorem arguments to disentangle these distributions. These approaches allow the researcher to compute the optimal reserve price *conditional* on a realization of the unobserved heterogeneity component. However, these approaches rely on stronger assumptions on the information environment or on the data than our approach. For example, Freyberger and Larsen (2019) require data on reserve prices and require assuming a scalar unobserved heterogeneity component in bidder values that is additively or multiplicatively separable from the private-value component. Our approach does not require these assumptions.

The way our approach handles unobserved auction-level heterogeneity by computing an optimal reserve price that averages over realizations of unobserved heterogeneity is similar to the way unobserved heterogeneity is handled in the bounds approaches of Aradillas-López et al. (2013) and Coey et al. (2017), which, like our paper, focus on correlated private values settings.[12] The important point we wish to emphasize is that the

[12]The bounds in both of those previous papers apply in the case where the number of bidders and the second-highest bid are observed (but the first-highest bid is not), obtaining *partial identification* in those cases. Our paper, on the other hand, obtains *point identified reserve prices* for the case where the number of bidders is

presence of unobserved heterogeneity does not in any way invalidate our approach; it simply means that the reserve prices delivered by our approach represent reserve prices averaged over any unobservable heterogeneity. Like the approaches of Aradillas-López et al. (2013) and Coey et al. (2017), it is possible that the reserve prices from our approach may underperform relative to a case in which the seller observes the realization of auction-level unobserved heterogeneity and conditions on this realization in setting the reserve price.

3.4. **Public vs. Secret Reserve Prices.** While Remark 1 refers to historical auctions with no reserve price, our method is also valid using historical auctions with a positive *secret* reserve price, because a secret reserve price would still allow the econometrician to observe realizations of the first- and second-highest bids in each auction. In fact, there is a strong argument for also using a secret reserve price to *implement* the optimal reserve price once it has been estimated, as this would allow the econometrician to continue observing more realizations of the first- and second-highest bids, further improving the accuracy of the estimated reserve price.[13]

3.5. **The Seller's Valuation, v_0.** We do not propose any estimate of the parameter v_0, the seller's valuation. We see this as a parameter of the problem known to the practitioner/seller, and to be specified by her when estimating the optimal reserve price. For example, a seller expecting to obtain \$100 for an item if the current auction fails to close should compute the optimal reserve price with v_0 set to \$100.[14] A reasonable value for v_0

unobservable but both the first- and second-highest bids are observable. In either case, unobserved auction-level heterogeneity only introduces correlation in private values and does not invalidate the methods.

[13]An alternative but complementary argument underlies the seller's motivation for using secret reserve prices in Andreyanov and Caoui (2020): If bidders are better informed about auction-level heterogeneity than the seller, a seller may wish to observe a given auction first and then choose to accept or reject the auction price only after the bidding ends, effectively implementing a *secret* reserve price. In the environment we study, a seller could potentially benefit from using a secret reserve price to learn about bids for *future* auctions, rather than profiting from learning in the current auction as in Andreyanov and Caoui (2020).

[14]Importantly, it is unnecessary to make any assumptions regarding what the seller's valuation is in the historical auction data; the pieces of information the seller/econometrician needs from these historical auctions are the marginal distribution of the first- and second-highest bids; the parameter v_0 only matters for computing the optimal reserve price going forward, given these historical auction bids. Similarly, it is also unnecessary to make any assumption regarding correlation (or lack thereof) between v_0 and bidder valuations; indeed, for a given seller computing the optimal reserve price, v_0 is a *scalar constant*. Furthermore, the correlation structure between seller valuations and bidder valuations in the historical auctions has no effect on the bids because the private values environment ensures that each bidder plays the weakly dominant strategy of bidding her valuation whenever the current auction price does not exceed her valuation, regardless of the seller's valuation or the seller's choice of the reserve price.

might be the seller's estimate of the price at which the good might sell if the auction were to fail (such as the expected price from recent posted-price listings in the eBay context, which the seller can search on the eBay website). An alternative would be for the seller to consider re-auctioning the item if it fails to sell. The following recursive formula specifies the seller's dynamic payoff $\breve{\pi}(r)$ in this setting:

$$\breve{\pi}(r) = E[\max\{r, V^{(2)}\}] - r\Pr(V^{(1)} < r) + \beta\Pr(V^{(1)} < r)\breve{\pi}(r)$$

where β is the seller's discount factor. Rearranging the above expression yields

$$\breve{\pi}(r) = \frac{1}{1 - \beta\Pr(V^{(1)} < r)}\left(E[\max\{r, V^{(2)}\}] - r\Pr(V^{(1)} < r)\right) \tag{5}$$

Maximizing this expression yields an estimate of the optimal reserve price in this repeated auction setting.[15] This can be implemented empirically similarly to the main case of (3).

3.6. **Limitations.** There are several settings (ruled out by Assumption 1 or Assumption 2) in which our identification and estimation approach would *not* be valid because the distribution of bids in auctions with no reserve price (which Remark 1 relies on) would not be representative of the distribution of bids in auctions with a positive reserve price. First, the reserve price estimated using historical bid data from zero-reserve auctions will not necessarily be optimal in an interdependent/common values environment. In such an environment, the reserve price itself affects equilibrium bidding (as described, for example, in Milgrom and Weber 1982; Cai et al. 2007; Quint 2017), and can provide signaling value for bidders. Assumption 1(i) rules out this environment.

Second, our estimated reserve prices would not necessarily be valid if bidders can choose which (of several) auctions to participate in, or can pass on a particular auction and await a future auction. In such a multiple, simultaneous (or dynamic) auction environment, a bidder's choice of which auction to participate in would depend on the reserve price in each auction, and hence the distribution of the two highest bids in an auction would also depend on r. This would violate Assumption 2, and a seller using historical bid data from

[15]Note that we implicitly assume here a *stationary* dynamic environment, ruling out the idea that the good itself is perishable (Waisman 2020).

auctions with $r = 0$ using our approach would not infer the correct optimal reserve price.[16] This type of environment is ruled out by Assumption 1(*iii*).[17]

Third, our model does not allow for certain types of endogenous entry. We explicitly rule out entry costs in Assumption 1(*iv*) and we impose that the distribution of the number of bidders does not depend on N in Assumption 1(v). Models of auctions with endogenous entry typically consider a two-stage decision-making process: first, bidders decide whether to participate in the auction, and then, conditional on participating, what bid to place. Whether one assumes that valuations are known prior to the entry decision (Samuelson 1985) or after it (Levin and Smith 1994), the decision to participate is often given by a threshold strategy: a bidder participates if and only if his expected payoff from the auction exceeds his participation (or entry) cost. In both cases, this expected payoff directly depends on r, which therefore directly determines the expected number of bidders that participate in the auction. Our approach can only allow for endogenous entry (as in the aforementioned models) if bidders' entry decisions are made *before* bidders observe r.

4. Asymptotic Properties of Estimated Reserves and Revenue

In this section we discuss the properties of the estimators we propose for the optimal profit and optimal reserve price, defined in equations (3) and (4). For ease of notation, throughout this section we set $v_0 = 0$, but the results we now present do not require this. All proofs are found in the Appendix.

We first present asymptotic results for the estimators $\hat{r}_J$ and $\hat{p}(\hat{r}_J)$. We maintain the following technical assumptions:

Assumption 3.
(i) The joint distribution of $V^{(2)}$ and $V^{(1)}$ admits an absolutely continuous, Lebesgue-measurable

[16]Under such circumstances, determining the optimal reserve price requires re-solving the bidders' optimization problem for a new equilibrium as in Balseiro et al. (2015) and Choi and Mela (2019), for example. In turn, determining an equilibrium solution for the bidders' optimization problem in these dynamic environments requires stronger assumptions than Assumption 1. An even more general and complex model is that of Kanoria and Nazerzadeh (2020), who consider forward-looking bidders that not only anticipate future reserve prices, but also internalize that their current bids will be used by the seller to estimate and implement future reserve prices.

[17]A recent empirical and theoretical literature does allow for multiple, dynamic auctions and long-lived bidders, such as Zeithammer (2006), Hendricks and Sorensen (2018), Backus and Lewis (2020), Bodoh-Creed et al. (2020), and Coey et al. (2020).

density; (ii) $0 \leqslant V^{(2)} \leqslant V^{(1)} \leqslant \bar{\omega} < \infty$*; (iii) The function* $p(r)$ *is uniquely maximized at* r^**; (iv) The function* $\tilde{\pi}(\cdot, r) \equiv \pi(\cdot, r) - \pi(\cdot, r^*)$ *has* $\mathbb{E}[\tilde{\pi}(\cdot, r)]$ *twice differentiable.*

These assumptions are standard technical conditions in the auction methodology literature and the econometrics literature more broadly. Conditions (i) and (ii) guarantee continuity of $p(r)$. Condition (iii) indirectly imposes restrictions on the underlying distributions of valuations and bidder arrival process. In some information environments, it is easy to specify sufficient conditions for (iii) to be satisfied. For example, in a setting with symmetric independent private values, a monotone hazard rate for bidder valuations is a sufficient condition for $p(r)$ to be uniquely maximized at r^*. Deriving a general sufficient condition for (iii) to be satisfied is beyond the scope of this paper, and we therefore state it directly as an econometric assumption.[18] Condition (iv) (twice differentiability) will be satisfied provided the joint density of $V^{(2)}$ and $V^{(1)}$ is sufficiently smooth.

Under these assumptions, we derive a number of results. First, we demonstrate consistency of both $\hat{r}_J$ and $\hat{p}(\hat{r}_J)$:

Theorem 1.
If Assumptions 1 and 3 are satisfied, then $\hat{p}(\hat{r}_J) \xrightarrow{p} p(r^*)$ *and* $\hat{r}_J \xrightarrow{p} r^*$*.*

The proof, given in the Appendix, consists of showing that, under Assumption 3, $\sup_{r \in \mathcal{R}} |\hat{p}(r) - p(r)| \xrightarrow{p} 0$, where $\mathcal{R} \equiv [0, \bar{\omega}]$. Hence, all the requirements from Theorem 2.1 of Newey and McFadden (1994) are satisfied, which yields Theorem 1.

Having established consistency, we now derive the asymptotic distribution of the estimator $\hat{r}_J$, which is not standard. The estimator belongs to a class of estimators that converge at a cube-root rate, of which an example is the maximum score estimator proposed by Manski (1975, 1985). To demonstrate this and derive the asymptotic distribution, we show that the conditions in the main theorem of Kim and Pollard (1990) are satisfied.[19]

Theorem 2.
If Assumptions 1 and 3 are satisfied, and if r^* *is an interior point, then the process* $J^{2/3} \frac{1}{J} \sum_{j=1}^{J} \tilde{\pi}(\cdot, r^* +$

[18]See van den Berg (2007) for a more extensive study that addresses sufficient conditions for a unique optimum in monopolistic selling problems.

[19]This same result regarding the rate of convergence was obtained by Segal (2003) in the context of optimal mechanisms with unknown demand and by Prasad (2008) in the context of posted prices. While Segal (2003) only derived the rate of convergence, Prasad (2008) obtained the estimator's asymptotic distribution in the same way we do: by verifying that the conditions in Kim and Pollard (1990) are satisfied.

$\alpha J^{-1/3}$) *converges in distribution to a Gaussian process, and* $J^{1/3}\left(\hat{r}_J - r^*\right)$ *converges in distribution to the random maximizer of this process.*

For additional technical details and discussion, we refer the interested reader to the Appendix. The important implication we highlight here is the cube-root rate of convergence of $\hat{r}_J$, discussed in more detail in the Appendix. This is slower than the square-root rate typical to many econometric settings, and suggests a stronger data requirement to obtain a precise estimate of r^*. We examine the practical relevance of these theoretical results in our application in Section 6.

While accurately estimating the optimal reserve price itself may in theory require a large dataset, determining *how much the auction designer could gain* from optimally choosing the reserve price does not. In fact, $\hat{p}(\hat{r}_J)$ converges at a *square-root* rate to a normal distribution, which we present in the following theorem:

Theorem 3.
If the conditions from the previous theorems are satisfied, it then follows that
$\sqrt{J}\left[\hat{p}(\hat{r}_J) - p(r^*)\right] \xrightarrow{d} N\left(0, Var[\pi(\cdot, r^*)]\right)$.

We now discuss inference. Simulating the asymptotic distribution of $\hat{r}_J$ is impractical as the second derivative of the expected profit function depends upon the distributions of the two order statistics used to estimate r^*. Resampling methods are a useful alternative to simulation. Abrevaya and Huang (2005) showed that the most straightforward resampling method for inference, nonparametric bootstrap, is not valid for this cube-root class of estimators. Alternative resampling methods that may be used in this case include subsampling (Delgado et al., 2001), m out of n bootstrap (Lee and Pun, 2006), numerical bootstrap (Hong and Li, 2020), and rescaled bootstrap (Cattaneo et al., 2020). Another possibility is to replace the indicator in the objective function with a smoothed estimator, which, along with further assumptions, might restore asymptotic normality and achieve faster rates of convergence, akin to the smoothed maximum score estimator introduced by Horowitz (1992). We leave this possibility as an avenue for future research.

However, the nonparametric bootstrap *can* be used for inference on the object $\hat{p}(\hat{r}_J)$. Let $\hat{p}^b(r)$ be the objective function defined above calculated from a bootstrap sample of size J drawn with replacement from the original sample, and let $\hat{r}_J^b$ be the estimator calculated from this bootstrap sample. Results in Abrevaya and Huang (2005) yield $\hat{p}^b(\hat{r}_J^b) - \hat{p}^b(r^*) = O_P(J^{-2/3})$ and $\hat{r}_J^b - r^* = O_P(J^{-1/3})$, which imply that

$$\sqrt{J}\left[\hat{p}^b(\hat{r}_J^b) - \hat{p}(\hat{r}_J)\right] = \sqrt{J}\left[\hat{p}^b(r^*) - \hat{p}(r^*)\right] + o_P(1),$$

which, in turn, has the same limiting distribution as $\sqrt{J}\left[\hat{p}(\hat{r}_J) - p(r^*)\right]$ conditional on the data due to a standard result from bootstrap theory.

We now provide a theoretical argument establishing a probabilistic upper bound on the quantity $p(r^*) - p(\widehat{r}_J)$, which shrinks to zero at the rate $(\log J/J)^{-1/2}$. We seek an expression that, under very minimal assumptions on the distribution of valuations, can aid the practitioner in making an informed decision regarding the data collection process: how many absolute (no reserve) auctions to run to obtain the estimate $\hat{r}_J$.[20] This argument builds on arguments from Mohri and Medina (2016).[21] The bound relies on virtually no assumptions on the distribution of valuations (other than that they have a finite upper bound), and as such it is very conservative, as is typical in the algorithmic game theory literature in computer science. It yields a worst-case-scenario bound on revenue performance; as such, the result will be most useful as a guide to the conservative practitioner. A second use of the bound is that it allows us to provide a graphical illustration of how the accuracy of estimated reserves can improve with the size of the dataset.

Given a sample of J auctions, we refer to the seller's reserve price as being the *estimated* reserve price if the seller chooses the reserve that maximizes profit on past data, that is, in auction $J+1$ she chooses the reserve price $\widehat{r}_J = \arg\max_r \widehat{p}(r)$. In expectation (over the possible bids in the $J+1^{th}$ auction), this gives a profit of $p(\widehat{r}_J)$, which, by definition, is lower than the expected profit given by the optimal reserve price, $p(r^*)$. We study the size of this difference, and how it changes as the number of observed auctions becomes large. We state the bound in the following theorem. Its proof uses techniques from statistical learning to probabilistically bound the difference between $\widehat{p}(r)$ and $p(r)$ uniformly in r. Specifically, the *empirical Rademacher complexity* (defined in the Appendix) plays a key role in obtaining

[20]Note that this approach to pick J can also be used to tackle the so-called "cold start" problem in online adaptive methods that also require that a number of absolute auctions be run such as Cesa-Bianchi et al. (2015).

[21]A distinction between our work and Mohri and Medina (2016) is that we do not consider auction-level covariates, allowing us to derive an expression for the bound in terms of quantities that are known to (or can be assumed by) the econometrician. Note that our goal here differs from that of Mohri and Medina (2016), who focused on improved algorithms to solve the optimization problem in (3) and (4).

this bound. These techniques are developed in Koltchinskii (2001) and Koltchinskii and Panchenko (2002); Mohri et al. (2012) provide a textbook overview.

Theorem 4.

If Assumptions 1 and 3 are satisfied, for any $\delta > 0$, with probability at least $1 - \delta$ over the possible realizations of the J auctions, it holds that

$$\frac{p(r^*) - p(\hat{r}_J)}{\overline{\omega}} \leqslant \left(\frac{8\sqrt{\log 2}}{J} + 4\sqrt{\frac{2 + 2\log J}{J}} + 6\sqrt{\frac{\log \frac{4}{\delta}}{2J}} \right).$$

To interpret this result, we can define ε as the gap between the profit at the estimated reserve price and the profit at the true optimal reserve price, normalized by the upper bound of valuations, $\overline{\omega}$ (and thus, $\varepsilon \in [0, 1]$ and can be interpreted as a fraction of the maximum bidder valuation). Theorem 4 implies that, in order for the profit from the estimated reserve to be within ε of the profit from the optimal reserve with probability at least $1 - \delta$, the seller needs to have observed approximately J auctions such that the expression in the right-hand side of the expression in Theorem 4 equals ε. It can be shown that when $J = 1$ the expression is positive and that it is strictly decreasing in J, which, in principle, enables one to obtain the desired J for each (δ, ε) via a simple nonlinear solver. Note also that the right-hand side does not depend on *any* feature of the valuation distribution, and the theorem relates the optimal reserve, r^*, and the true expected profit function, $p(\cdot)$, without requiring these objects to be known; it is due to these weak requirements that the bound is conservative.

We illustrate the implications of Theorem 4 in Figure 1. For this illustration, we normalize $\overline{\omega} = 1$, and thus revenue is in units of fractions of the maximum willingness to pay. We then plot "iso-data" curves in (δ, ε) space, where each curve represents the possible combinations of ε and δ that are possible given a fixed history of observed auctions. In this figure, a curve located further to the southwest is preferable, as it represents a closer approximation to the true optimal revenue (i.e. a smaller ε) with a higher probability (i.e. a lower δ). The top line represents a sample size of $J = 1,000$, the middle line represents $J = 5,000$, and the bottom line represents $J = 10,000$. The middle line suggests that with a history of 5,000 auction realizations, one could guarantee a payoff within 0.348 (units of the maximum willingness to pay) of the optimal profit with probability 0.975; or, with the same size history, one could guarantee a payoff within 0.3447 of the optimal profit with probability 0.70. The larger sample, $J = 10,000$, can guarantee a payoff that is much closer

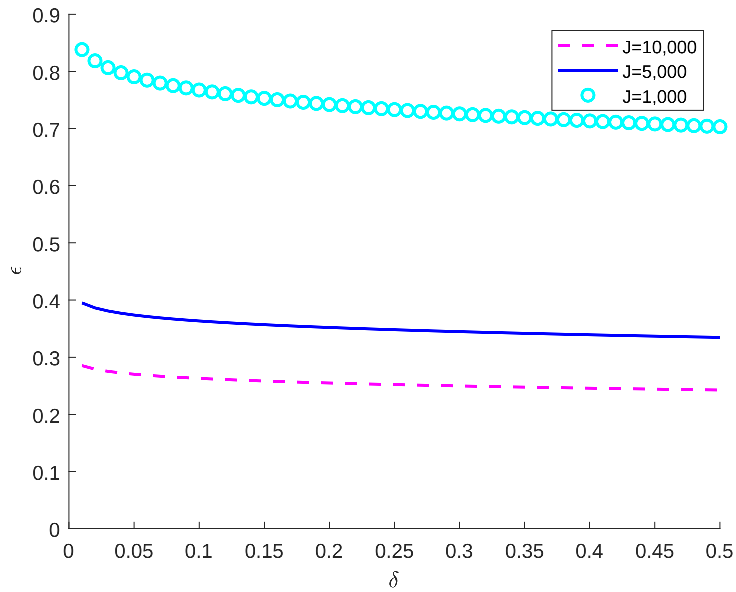

FIGURE 1. Iso-data curves in (δ, ε) space

Notes: Figure displays combinations of (δ, ε) that can be achieved given a fixed sample size J using the bound implied by Theorem 4. Top line represents $J = 1,000$, middle line represents $J = 5,000$, and bottom line represents $J = 10,000$.

to the optimal profit. Each iso-data curve is relatively flat in the δ dimension, reflecting the fact that the sample size requirements are more stringent for achieving a given level of ε closeness to the optimal profit, and are less stringent for achieving an improvement in δ (i.e. in the probability with which the revenue is reached).

We now relate the explicit bound obtained in this subsection to the asymptotic results obtained in the previous subsection. Theorem 2 implies that, for sufficiently large J, $p(\hat{r}_J) - p(r^*) = O_p(J^{-2/3})$.[22] By definition, therefore, for any $\delta > 0$ and sufficiently large J, there exists an $M > 0$ such that $\Pr(J^{2/3}|p(\hat{r}_J) - p(r^*)| < M) \geqslant 1 - \delta$. The fact that the convergence result in Theorem 2 is achieved at a $J^{-2/3}$ rate implies that the bound in Theorem 4 is conservative, as it is expressed as a function of $(\log J/J)^{-1/2}$. However, Theorem 2 does not allow one to explicitly compute the number of auctions required in order to approach the optimal revenue with a given probability; it simply states that such an M exists if J is large enough. The advantage of the bound in Theorem 4, on the other hand, is that it is explicit, allowing one to directly compute an estimate (albeit a very conservative estimate) of the

[22]Performing a second-order Taylor expansion yields $p(\hat{r}_J) - p(r^*) = p''(\tilde{r})(\hat{r}_J - r^*)^2 = O_p(J^{-2/3})$, where $\tilde{r}$ is an intermediate value between $\hat{r}_J$ and r^*.

number of auctions J which must be observed for estimated reserve prices to perform well without requiring knowledge of $p(\cdot)$ or r^*.

In our application in Section 6, we take a step beyond asymptotic and learning theory and demonstrate that each of these theoretical results may be quite conservative in practice, as we find that, even with a relatively small data set of previously observed auctions, revenue based on the estimated optimal reserve price can come quite close to the full optimal-reserve-auction revenue.

5. Monte Carlo Simulations

We now evaluate the finite-sample performance of our proposed procedure via Monte Carlo simulations. We evaluate two different scenarios. In each scenario, we simulate auctions with $N = 5$ bidders with valuations drawn from a $U[0,1]$, and for simplicity we set $v_0 = 0$. In the first scenario these valuations are independent and in the second these valuations are perfectly correlated (with $V_1 = ... = V_5$).[23] We report the difference in expected profits from using the estimated reserve price using a small sample of historical auction observations (samples of size $J = 10, 20, ..., 100$ auctions) relative to the expected profits from using the true optimal reserve price (computed based on a sample of 10,000 auctions). We use 500 replications of these simulations.

In addition to estimating reserve prices using our procedure, we also estimate reserve prices by first estimating the full valuation distribution F. As highlighted in Section 2, the only existing method for estimating the full distribution of valuations when only two order statistics of bids are observed and N is unobserved is the approach of Song (2004). This method relies on the novel insight that, in an *iid* environment, the distribution of a higher order statistic conditional on a lower order statistic does not depend on N. Song (2004) proposes using maximum likelihood estimation (MLE) to estimate F. The insight of Song (2004) holds in a symmetric IPV environment. In such an environment, the optimal reserve price can then be estimated in a number of possible ways once F is known. One valid approach would be to choose r^* as

$$\arg\max_r r(1 - F(r)). \tag{6}$$

[23]We choose these two cases to evaluate our approach, but intermediate cases can also capture the relationships we document here.

The benefit of the Song (2004) method, in addition to providing an estimate of F, is that it takes advantage of the strong assumption of symmetric IPV, and hence can yield more efficient estimates of the reserve price than our direct approach. A drawback of the Song (2004) approach is that it is not valid outside of symmetric IPV environment, unlike our approach. Other drawbacks are that it requires approximating F (which our approach does not), and that can be more costly in terms of computation time or ease of implementation (our approach can be implemented with a simple grid search without requiring a numerical optimization routine). We illustrate two approximations for F: a uniform distribution $U[1, \bar{w}]$, where $\bar{w}$ is the parameter we estimate, and a fifth-degree Hermite polynomial approximation for F. The former assumes much stronger knowledge about the shape of the underlying valuation distribution (knowledge a researcher is unlikely to have in practice). Details on the estimation approach based on Song (2004) are found in the Online Appendix.

Figures 2–4 displays the results of the simulation exercises. In each panel, solid lines indicate the average across 500 Monte Carlo replications and dashed lines indicate asymmetric 95% confidence intervals constructed from the 0.025 and 0.975 quantiles of the estimator across the 500 replications; note that for some panels these confidence intervals are so tight that they are indistinguishable from the solid lines. Panel A shows that, when data truly is generated by a symmetric IPV process, the Song (2004) approach (the red lines, "Est. full dist.") outperforms our approach (the blue lines, "Directly est. r^*") in terms of expected loss. In panel C, when we use a Hermite polynomial approximation to the valuation distribution instead of assuming knowledge that is a uniform distribution, the Song (2004) approach still outperforms ours in terms of expected loss but to a lesser degree. The performance of the directly estimated reserve prices improves as the sample size increases.

The real benefit of our approach is illustrated by panels B and D, where valuations arise from a correlated PV environment. Here the assumptions of Song (2004) are not satisfied, and the estimated reserve prices from that approach are clearly biased, with an average expected loss around 24% of the true optimal reserve profit in panel B (using Uniform MLE) and an average loss of about 3% in panel D (using Hermite MLE).

Figure 3 illustrates another strength of our approach, even in symmetric IPV environments: computation time. The average computation time for the Song (2004) approach is a full 4 seconds in panel B, which uses the more computationally burdensome Hermite

FIGURE 2. Monte Carlo Simulations: Expected Loss

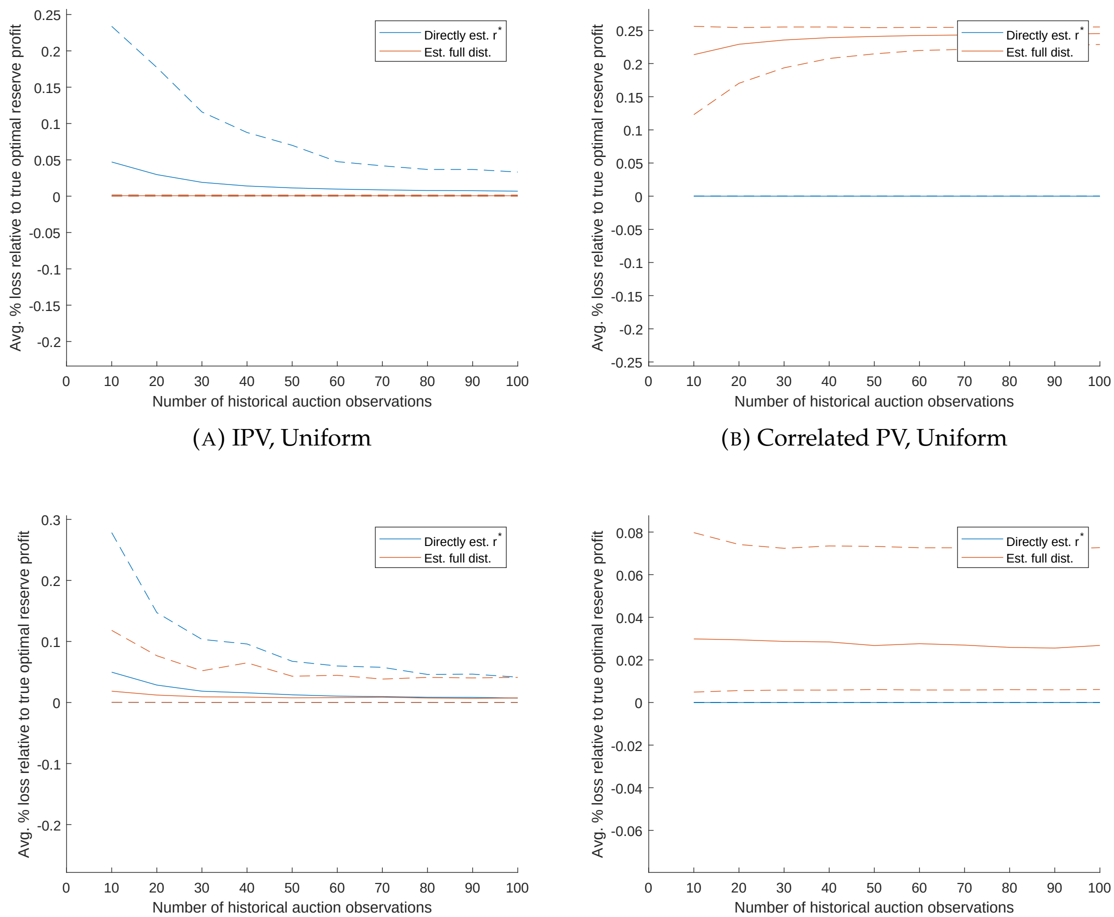

(A) IPV, Uniform

(B) Correlated PV, Uniform

(C) IPV, Hermite

(D) Correlated PV, Hermite

Notes: Panel A shows average (across 500 Monte Carlo replications) of the expected loss from using reserve prices estimated from a sample of $J = 10, 20, ...,$ or 100 observations relative to true optimal reserve profit. "Directly use $r*$" refers to reserve prices estimated using our approach, and "Est. full dist." refers to results from estimating the full distribution of valuations using MLE following Song (2004) (using a uniform distribution in panels A and B and using Hermite polynomials in panels C and D). Panels A and C show the symmetric IPV setting and panels B and D the correlated PV setting. Asymmetric 95% confidence bands are shown with dashed lines.

approximation for F. In the correlated PV environment, our approach also performs faster than the Song (2004) approach. In each panel of Figure 3 the average computation time for the directly estimated reserve is less than 0.002 seconds.

FIGURE 3. Monte Carlo Simulations: Average Computation Time

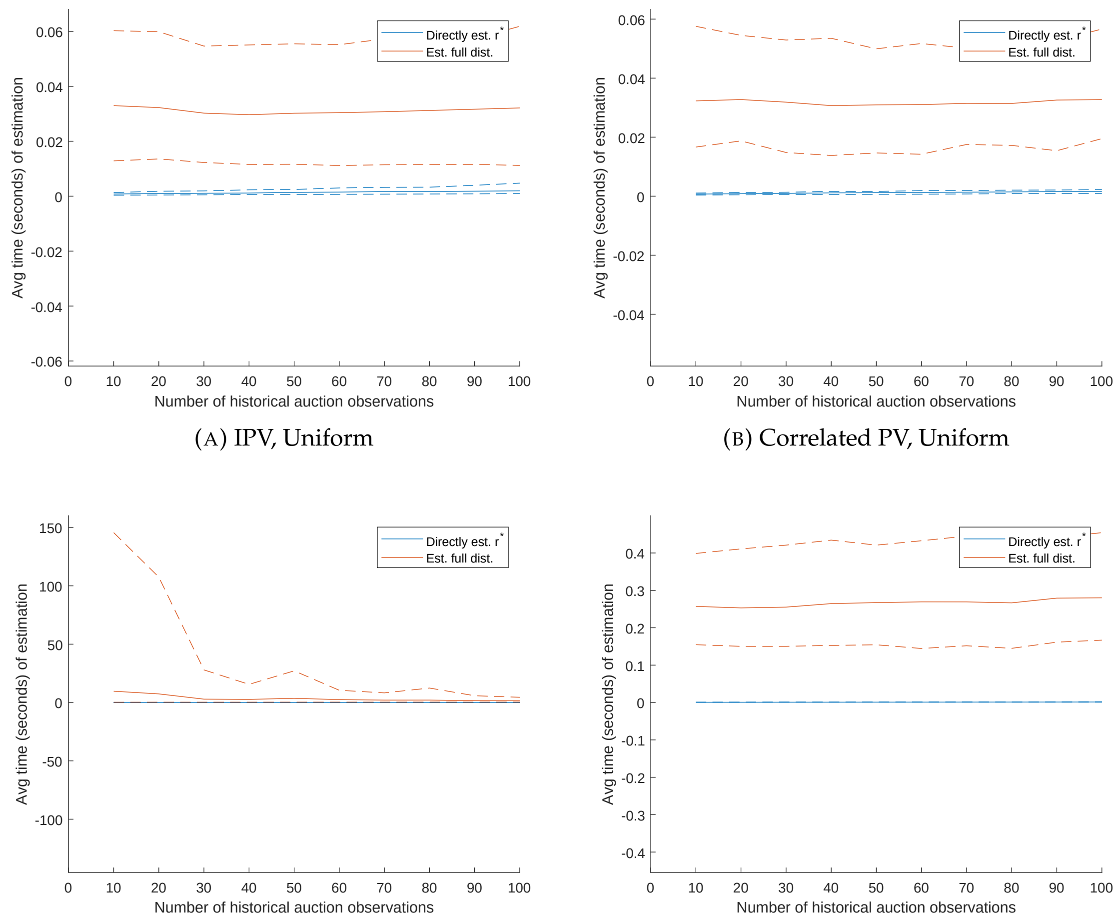

(A) IPV, Uniform

(B) Correlated PV, Uniform

(C) IPV, Hermite

(D) Correlated PV, Hermite

Notes: Panel A shows average (across 500 Monte Carlo replications) computation time in seconds for estimated reserve prices using a sample of $J = 10, 20, ...,$ or 100 observations. "Directly use r^*" refers to reserve prices estimated using our approach, and "Est. full dist." refers to results from estimating the full distribution of valuations using MLE following Song (2004) (using a uniform distribution in panels A and B and using Hermite polynomials in panels C and D). Panels A and C show the symmetric IPV setting and panels B and D the correlated PV setting. Asymmetric 95% confidence bands are shown with dashed lines.

In Figure 4, panel A compares the expected loss from our approach to what could be achieved if the econometrician were to observed *all* bids from the auction, yielding an immediate estimate of F that can be plugged into (6) to obtain an estimate of r^* in a symmetric IPV environment. Panel A indicates that this kind of data would indeed be beneficial relative to our estimation approach in a symmetric IPV setting. Panel B demonstrates, however, that this approach would be biased in a correlated PV setting. This is because, in a general correlated private value function, all information about the seller's profit (and hence the optimal reserve price) is contained in the marginal distributions of the first- and second-highest bids; lower order statistics of bids contain no additional information for the seller's profit optimization problem.[24]

We also highlight here that observing all bids is not possible in the sequential-arrival auction we model, where some bidders can arrive after the current bid has passed their valuation. Panel C of Figure 4 demonstrates that erroneously using the *observed* bids (used for the estimates show in green) as though they represent *all* bids (used for the estimates shown in purple) will lead to biased reserve prices, even in a symmetric IPV environment.

Panel D demonstrates the sensitivity of our approach to Assumption 2, which requires that the data generating process for would-be bids be the same the historical auctions and the future auctions. In panel D the blue line indicates results using our approach and the yellow line indicates the expected loss when we instead artificially inflate historical bids by 1%, holding fixed the true distribution of bids. These inflated historical bids lead to a clear bias in estimated reserve prices that does not disappear as the sample size increases.

6. Computing Optimal Reserve Prices in E-commerce Auctions

We apply our methodology to a dataset of eBay auctions selling commodity-like products, which we define as those products which are cataloged in one of several commercially available product catalogs. Examples of commodity products include "Microsoft Xbox One, 500 GB Black Console", "Chanel No.5 3.4oz Women's Eau de Parfum", and "The Sopranos - The Complete Series (DVD, 2009)". We will refer to each distinct product as a "product" or "product-category." Within each product, the items sold are relatively homogeneous. For this exercise, we select popular iPhone products listed through auctions from 2011–2015. We consider only auctions with no reserve price; specifically, we only include auctions for which the start price was less than or equal to $0.99, the default start price recommendation on eBay. We omit auctions in which the highest bid is in the top 1% of all highest bids for that product and limit to products that are auctioned at least 1,000 times in our sample.

[24] Athey and Haile (2007) and Aradillas-López et al. (2013) discuss how, in the general correlated private values case, the seller gains no additional benefit in computing reserve prices by observing additional bids other than the two highest.

FIGURE 4. Monte Carlo Simulations: All Bids vs. Observed Bids vs. Biased Bids

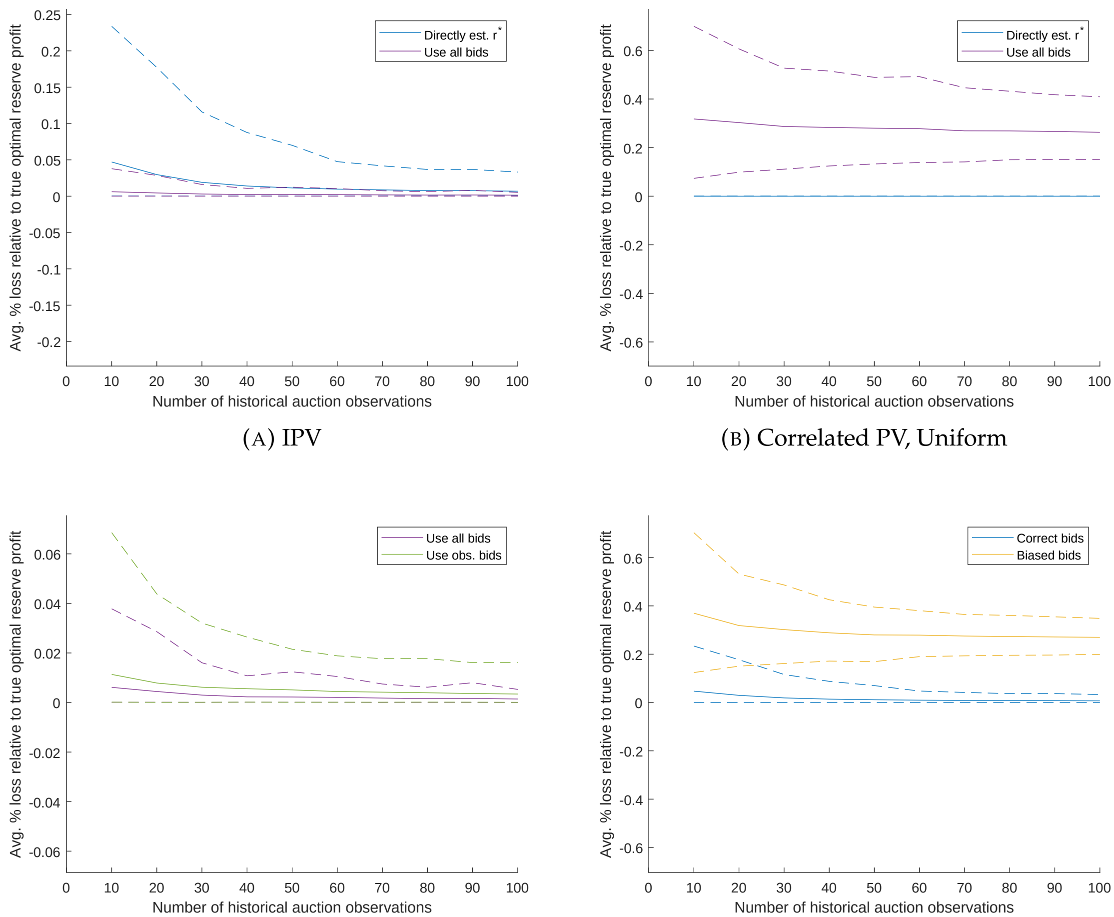

(A) IPV

(B) Correlated PV, Uniform

(C) IPV

(D) IPV

Notes: Panel A shows average (across 500 Monte Carlo replications) of the expected loss from using reserve prices estimated from a sample of $J = 10, 20, ...,$ or 100 observations relative to true optimal reserve profit. "Directly use $r*$" (and "Correct bids" in panel D) refers to reserve prices estimated using our approach; "Use all bids" refers to results from estimating the full distribution of valuations from all bids; "Use obs. bids" refers to results from estimating the full distribution of valuations from only those bids that are not censored by bidders' sequential arrival; "Biased bids" in panel D refers to results using bids historical bids that are 1% higher than the true simulated draws. Panels A, C, and D show the symmetric IPV setting and panel B the correlated PV setting. Asymmetric 95% confidence bands are shown with dashed lines.

Table 1 shows summary statistics at the product level. There are 22 distinct iPhone products in our sample, with the number of auctions per product ranging from 1,144 to 11,733. Table 1 shows evidence of price dispersion within a product across auctions. Our modeling framework rationalizes such dispersion through different numbers of bidders arriving to different auctions—more bidders will result in a higher auction price—and

TABLE 1. Product-Level Descriptive Statistics

		Highest Bid		Second Highest Bid	
Product	# Obs	Average($)	Std Dev ($)	Average ($)	Std Dev ($)
1. iP3GS, 16GB, AT&T	3,907	147.42	68.50	135.00	67.03
2. iP3GS, 32GB, AT&T	1,197	170.72	70.62	155.51	69.11
3. iP3GS, 8GB, AT&T	2,705	97.28	55.29	86.62	53.00
4. iP3G, 8GB, AT&T	3,545	95.43	47.95	85.16	45.86
5. iP4S, 16GB, AT&T	5,761	243.94	130.22	221.95	121.39
6. iP4S, 16GB, Sprint	2,978	205.74	110.19	185.80	104.08
7. iP4S, 16GB, Unlocked	1,199	297.08	148.26	267.59	124.25
8. iP4S, 16GB, Verizon	4,096	211.89	115.28	190.95	109.04
9. iP4S, 32GB, AT&T	1,493	292.21	133.64	266.66	128.57
10. iP4, 16GB, AT&T	11,733	231.28	105.86	213.13	101.24
11. iP4, 16GB, Unlocked	1,590	258.09	121.76	235.89	110.22
12. iP4, 16GB, Verizon	6,698	161.37	89.05	145.96	85.39
13. iP4, 32GB, AT&T	4,245	259.57	110.11	239.14	107.50
14. iP4, 32GB, Verizon	1,600	187.68	97.59	169.75	94.47
15. iP4, 8GB, AT&T	2,302	150.01	86.51	134.80	80.84
16. iP4, 8GB, Sprint	2,198	116.88	69.30	103.19	66.16
17. iP4, 8GB, Verizon	3,003	112.51	69.17	99.70	64.84
18. iP5, 16GB, AT&T	1,553	348.83	179.91	311.52	155.60
19. iP5, 16GB, T-Mobile	1,284	267.09	175.58	235.52	145.14
20. iP5, 16GB, Unlocked	3,381	288.81	164.58	257.38	149.70
21. iP5, 16GB, Verizon	1,605	307.24	152.28	271.54	132.95
22. iP5, 64GB, AT&T	1,144	231.86	133.43	206.87	121.88

Notes: Table displays, for each product, the number of auctions recorded and the average and standard deviation of the first and second highest bids.

different realizations of bidders' valuations in different auctions. Given that reserve prices increase revenue only when they lie between the highest and second highest bids, the size of the gap between these bids is of particular interest. This gap ranges from \$10.27 (12% of the mean second highest bid) for product 4 to \$37 for product 18 (also 12% of this product's the mean second highest bid). This gap is not large, suggesting that these products may have relatively competitive markets on eBay and that reserve prices may only be able to increase revenue slightly over a no-reserve auction for these products.

We estimate optimal reserves separately in each of the 22 products. Figures 5 displays the gain from using the optimal reserve price, where we evaluate these gains in three different ways. In Panel A we consider the case where $v_0 = 0$ and we report the percentage gain in revenue from using the estimated optimal reserve price relative to an auction with no reserve price (or a reserve price of zero). In panel A, we find that for $v_0 = 0$ it is not clear whether *any* positive reserve is beneficial: the expected percentage gain is less than 0.01 for each product and not statistically significantly different from zero for most products. Note that treating $v_0 = 0$ (regardless of the seller's true valuation) will yield the reserve price that will maximize the expected payment of the winning bidder to the seller, which may be the quantity that the *auction platform* (here, eBay) is most interested in maximizing, as platform fees are typically proportional to this payment. Our results in panel A therefore suggest that eBay would prefer a zero reserve for these products, consistent with eBay's practice of recommending low reserve prices (0.99) to sellers.

In panel B we set $v_0 = E[V^{(2)}]$ and consider the expected gain from using the optimal reserve price relative to an auction with $r = v_0$. We consider the $r = v_0$ auction as our benchmark because the seller would clearly find it suboptimal to sell the good at a price lower than her valuation v_0. Our specific choice of setting v_0 to the expected second-highest bid is a form of capturing the seller's perceived value of keeping the good herself. It also captures in a simple way what the seller might expect from re-auctioning the item. Here we find again small gains from implementing the optimal reserve price (less than 0.01), but the gains are statistically significantly different from zero for most products.

In panel C we consider the case where the seller re-auctions the item if it fails to sell, using Equation (5) with a discount factor of $\beta = 0.9$.[25] We evaluate the gain from using the optimal reserve price in this scenario, and here the gains are large and significant, ranging from 22% to 44%. These gains are larger than in panels A and B because here we evaluate the gains relative to a scenario in which the seller uses *no* reserve price. The reason for this choice of benchmark is the following. Unlike in the static auction case, where a benchmark reserve price related to the seller's valuation, v_0, is natural, there is no role for v_0 in the infinite-horizon, repeated-auction case. Rather, the seller's outside option when a given auction fails is defined recursively as the option to relist the item. Any arbitrary choice of an alternative, sub-optimal reserve price would also work as a benchmark, but we see the no-reserve benchmark as the most reasonable choice for this case.

[25]This choice of β is only for illustrative purposes, and it implies that the seller discounts the future by more than would be implied by the time value of money alone; we consider this discount factor as also capturing other unmodeled features that may result in impatience on the seller's part or that may prohibit the seller from quickly re-auctioning the item. The results can easily be generated with other values of β.

FIGURE 5. Revenue Increase from Optimal Reserve Price

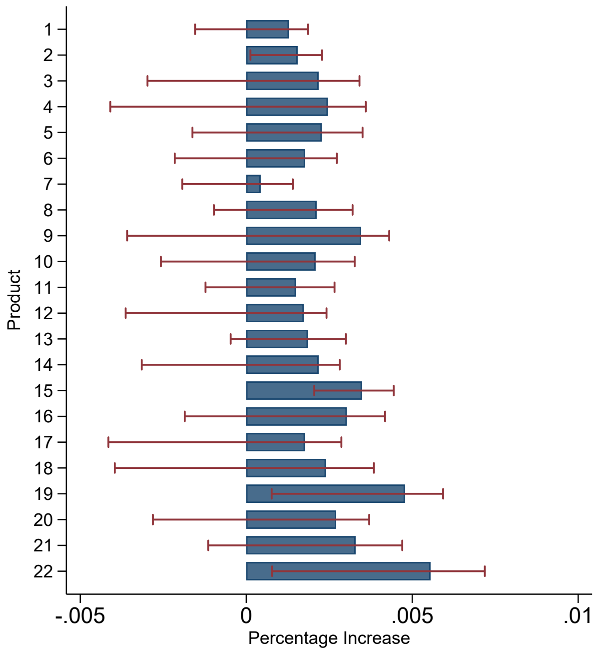

(A) $v_0 = 0$

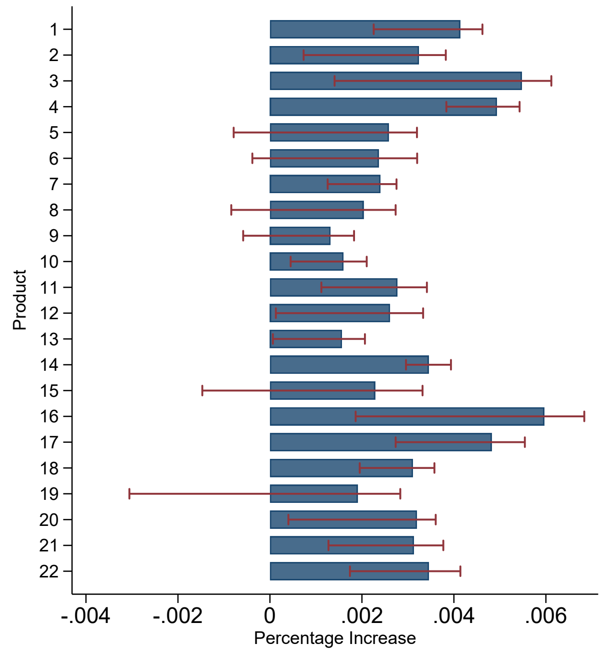

(B) $v_0 = E(V^{(2)})$

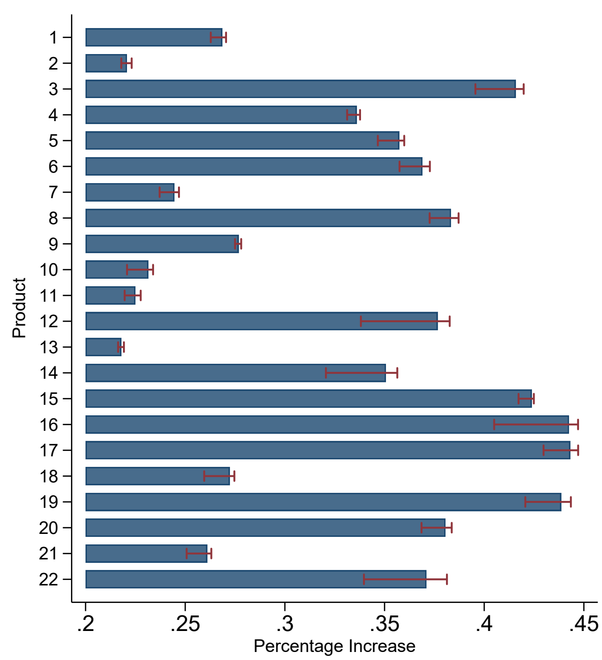

(C) Repeated Auction

Notes: Expected revenue increase from using estimated optimal reserve price relative to using no reserve price in panel A (where $v_0 = 0$) and relative to a reserve price of $r = E[V^{(2)}]$ in panel B (where $v_0 = E[V^{(2)}]$) and relative to no reserve price in panel C, where the seller's outside option is to re-auction the item. 95% confidence interval is shown in red, based on 1,000 subsample draws of size 250 from the full sample, separately for each product.

In Figure 6, we select one specific product in our sample, product #17, which is the 8GB version of the Apple iPhone 4 locked to Verizon. In the figure we plot expected profit as a function of the reserve price given the empirical distribution of the first and second highest bids. Panel A considers the setting where $v_0 = 0$, panel B considers the setting where v_0 is the average second order statistic, and panel C consider the repeated auction case. Unsurprisingly for a product supplied elastically on other online or offline platforms, the figure shows that there is a sharp drop-off in profit for reserves beyond a certain point. This large drop off illustrates a point also discussed in Ostrovsky and Schwarz (2016) and Kim (2013): the loss from setting a non-optimal reserve price is asymmetric, such that overshooting the optimal reserve results in a much larger loss in magnitude than undershooting it. The vertical line represents the estimated optimal reserve price.

We now turn to the question of how close optimal reserve prices will be to those estimated using a finite history of first and second-highest bids. The theoretical guarantee of Theorem 4 assures us that estimated reserve prices will eventually perform close to optimally. We assess this feature through a simulation exercise.

For each product, we draw 1,000 sequences, each of length 250, at random with replacement from the empirical distribution of all auctions observed for that product over the sample period. Within each sequence, we then estimate the reserve price suggested by our approach using only the first τ observations in the sequence, doing so separately for each $\tau \in \{2, ..., 250\}$. Thus, we begin with only 2 historical auction observation, then 3, then 4, and so on, for each drawn sequence. Next, at each of these estimated reserve prices, using the *full sample* of historical observations for the product, we compute the expected profit the seller would receive from using this computed reserve price. Therefore, for this exercise we treat the empirical distribution of auctions in our sample as representing the "true" distribution of first and second highest bids, and we treat sellers as only having information on a history of τ auctions drawn at random from the full empirical distribution.

Figure 7 shows the results of this exercise for the same product as in Figure 6. In panels A, C, and E, the quantities on the y-axis are expressed as the expected loss in profit from using reserve prices estimated from a given small sample size relative to the profit from using the "true" optimal reserve price. The solid line represents the average loss, averaged across the 1,000 subsample draws, and the dashed lines represent pointwise 95% confidence intervals (the 0.025 and 0.975 quantiles from the simulations). In panels B, D, and F

FIGURE 6. Profit Under Different Reserve Prices for Apple iPhone 4 8GB Verizon

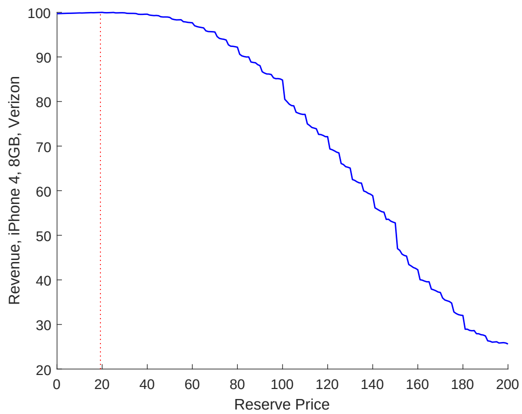

(A) $v_0 = 0$

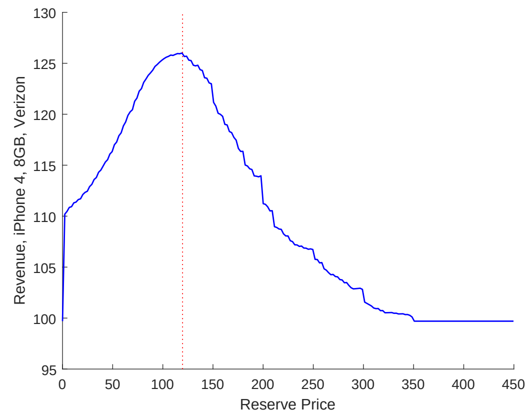

(B) $v_0 = E(V^{(2)})$

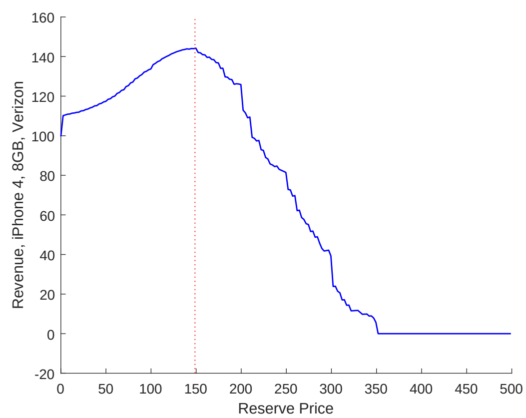

(C) Repeated Auction

Notes: Seller expected profit as a function of the reserve price, given the empirical distribution of first and second highest bids, for Apple iPhone 4 8GB Verizon (using all observations for this product). Panel A sets $v_0 = 0$, panel B sets v_0 to the average second highest bid, and panel C considers a repeated auction, where the outside option is to re-auction the item. Vertical line displays location of optimal reserve price.

we plot the actual estimated reserve price from each of these samples. Thus, in each panel, the x-axis represents the sample size used to compute the optimal reserve price (for panels on the right) and the corresponding loss in profit from using this reserve price (for panels on the left).

As expected given Theorem 4, the loss does indeed converge to zero (i.e. the profit converges to the true optimal reserve profit level). In the initial phases, estimated reserve prices can be seriously suboptimal, even compared to setting a reserve of $r = v_0$. However,

FIGURE 7. Expected Loss and Average Reserves From Computation with Different Numbers of Observed Auctions for Apple iPhone 4 8GB Verizon

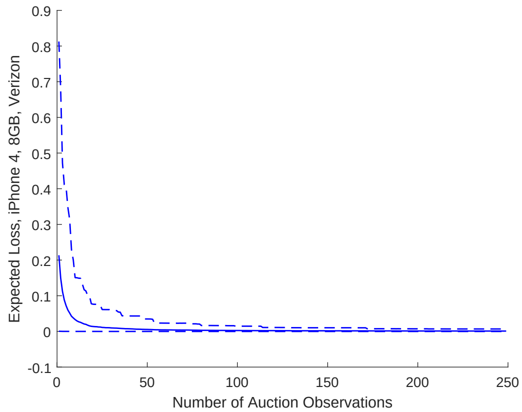

(A) Exp. Loss, $v_0 = 0$

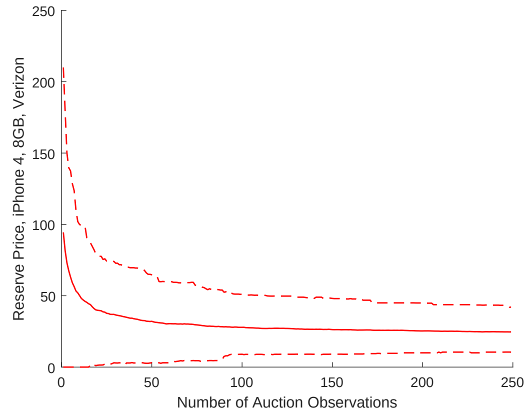

(B) Avg. Reserve, $v_0 = 0$

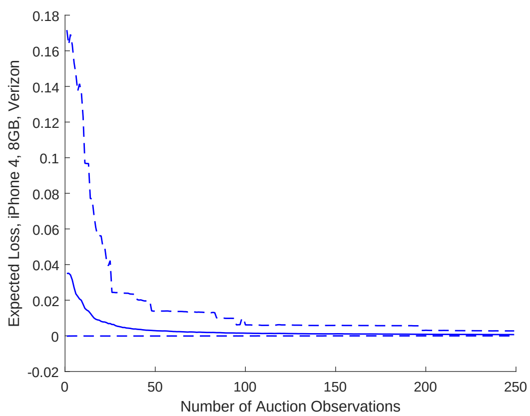

(C) Exp. Loss, $v_0 = E(V^{(2)})$

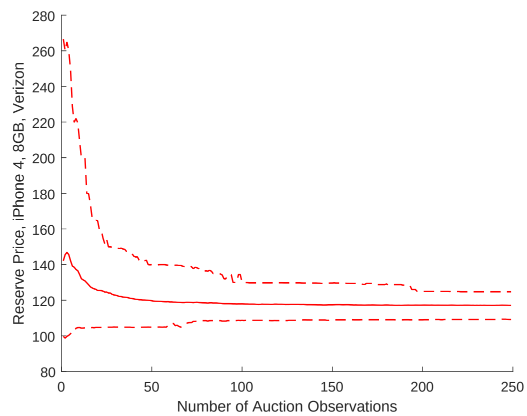

(D) Avg. Reserve, $v_0 = E(V^{(2)})$

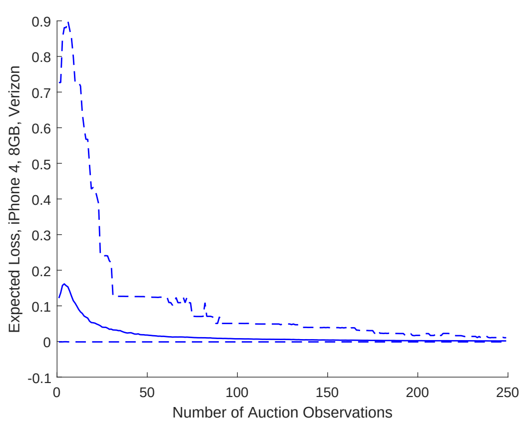

(E) Exp. Loss, Repeated Auc.

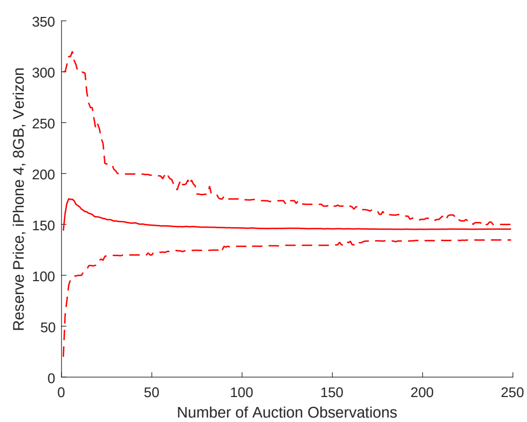

(F) Avg. Reserve, Repeated Auc.

Notes: Panels on left show expected loss (as a fraction of the true optimal expected profit) from using estimated reserve prices as a function of the number of auctions observed, where simulations are conducted by drawing sequences of auctions from the empirical distribution of first and second highest bids from auctions for Apple iPhone 4 8GB Verizon and computing the estimated expected profit progressively adding each auction one at a time. Panels on the right show the average estimated optimal reserve price. Solid line represents the average across the 1,000 simulation replications and dashed lines represent 95% confidence intervals. Panels A and B set $v_0 = 0$, panels B and C set v_0 to the average second highest bid, and panels D and E consider a repeated auction, where the outside option is to re-auction the item.

convergence to the optimal level appears to occur quite quickly. The optimal reserve prices themselves also converge relatively quickly. Initially, with especially small sample sizes, the estimated reserves are, on average, too high. The lower confidence bands in panels B, D, and F demonstrate that the estimated reserve prices can also be too low in some samples. These confidence bands shrink quickly as the sample size grows. This quick convergence to the truth is quite robust regardless of what outside option the seller chooses to consider. Figures corresponding to Figures 6 and 7 for all products are found in the Online Appendix.

We now address the question of how many auction observations a practitioner may wish to collect before implementing the estimated optimal reserve price. Table 2 displays results using the same simulation exercise described above, evaluated separately for each product. The first three columns show (for our three scenarios for the seller's outside option) the median number of historical auctions (across the 1,000 simulation draws) required to achieve a revenue level that is within 1% of the true optimal-reserve-auction revenue. For the $v_0 = 0$ case, this ranges from 6–10 auctions. For the $v_0 = E[V^{(2)}]$ case, as few as 1–2 auctions can achieve a revenue within 1% of the truth. For the repeated auction case, 4–22 auctions are required.

The last three columns of Table 2 address this question differently. We consider a scenario in which a seller has an inventory of 250 iPhones to sell, and she wishes to run J^* no-reserve auctions to collect data, and then in the remaining $(250 - J^*)$ auctions she will implement the optimal reserve price estimated from the first J^* auctions. The seller solves

$$J^* = \arg\max_J J p(0) + (250 - J) p(\hat{r}_J). \tag{7}$$

This problem is not actually feasible in practice because it depends on knowing the true profit function $p(\cdot)$.[26] Nonetheless, we consider it an interesting thought experiment. We report the resulting J^* from solving (7) separately for each product in the last three columns of Table 2. Here we find quantities that are of a similar order of magnitude to those in the first three columns, suggesting again that estimating and implementing a reserve price from even small samples (less than 25) of historical auctions can be profitable. These results also suggest that, if a mechanism designer concerned with changes in demand wishes to

[26]The results of this exercise also depend crucially on the size of the inventory; here, 250. A different inventory size would results in a different data-collection threshold.

TABLE 2. How Many No-Reserve Auctions to Run?

	Within 1% Opt. Rev.			Maximize Rev. 250 Sales		
	$v_0 = 0$	$v_0 = E[V^{(2)}]$	Repeat	$v_0 = 0$	$v_0 = E[V^{(2)}]$	Repeat
1. iP3GS, 16GB, AT&T	8	1	5	12	1	12
2. iP3GS, 32GB, AT&T	10	1	4	10	1	12
3. iP3GS, 8GB, AT&T	6	1	8	11	2	18
4. iP3G, 8GB, AT&T	7	1	9	11	1	15
5. iP4S, 16GB, AT&T	8	1	7	10	1	24
6. iP4S, 16GB, Sprint	9	1	8	10	1	18
7. iP4S, 16GB, Unlocked	7	1	5	10	1	1
8. iP4S, 16GB, Verizon	8	1	9	11	1	21
9. iP4S, 32GB, AT&T	7	1	7	11	1	20
10. iP4, 16GB, AT&T	7	2	4	10	2	1
11. iP4, 16GB, Unlocked	8	1	5	10	1	1
12. iP4, 16GB, Verizon	7	1	15	10	1	16
13. iP4, 32GB, AT&T	9	2	5	12	1	18
14. iP4, 32GB, Verizon	8	1	11	10	1	16
15. iP4, 8GB, AT&T	6	1	9	11	1	21
16. iP4, 8GB, Sprint	7	1	16	10	1	18
17. iP4, 8GB, Verizon	8	1	10	11	1	2
18. iP5, 16GB, AT&T	7	1	7	11	1	1
19. iP5, 16GB, T-Mobile	7	1	10	11	1	2
20. iP5, 16GB, Unlocked	7	1	8	9	1	2
21. iP5, 16GB, Verizon	7	1	8	10	1	1
22. iP5, 64GB, AT&T	9	1	22	11	1	2

Notes: For each product, the first three columns displays (for the three different seller outside option scenarios) the median number of auctions required for the computed reserve price to yield an expected profit that is within 1% of the true optimal reserve auction profit. The median is taken across 1,000 simulated samples drawn with replacement from the full set of observations for that product. The last three columns display the optimal number of no-reserve auctions to run prior to estimating and implementing the optimal reserve on the remainder of the auctions in the inventory, where the inventory size in this exercise is 250 items.

"reset" or "update" the optimal reserve price based on recently observed auctions, doing so may not be very costly, as each instance of setting a new reserve may only require a small sample of recent auctions.

7. Conclusion

We study a computationally simple approach for estimating optimal reserve prices in asymmetric, correlated private values settings. The approach applies to settings with incomplete bidding data where only the top two bids are observed and where the number of bidders is unknown. These data requirements are frequently met in online (advertising or e-commerce) settings. We also derive a bound on the number of auction records one needs to observe in order for realized revenue based on estimated reserve prices to approximate the optimal revenue closely. We illustrate the approach using eBay auctions of used iPhones, and illustrate that revenue could potentially increase if optimal reserve prices were employed in practice. We examine the empirical relevance of our theoretical results and find that fewer than 25 auctions need to be recorded prior to estimating reserve prices in order for the estimated reserve price to yield an expected loss of less than 1% relative to the true optimal reserve revenue.

While the approach abstracts away from a number of information settings or real-world details (such as common values or inter-auction dynamics), we believe the virtue of the approach is its simplicity, providing a tractable and scalable approach to computing reserve prices even in large, unwieldy datasets where typical computationally demanding empirical auction approaches would be infeasible.

References

Abraham, I., Athey, S., Babaioff, M., and Grubb, M. D. (2020). Peaches, lemons, and cookies: Designing auction markets with dispersed information. *Games and Economic Behavior*, 124:454–477.

Abrevaya, J. and Huang, J. (2005). On the bootstrap of the maximum score estimator. *Econometrica*, 73(4):1175–1204.

Ali, S. N., Lewis, G., and Vasserman, S. (2019). Voluntary disclosure and personalized pricing. NBER Working Paper 26592.

Andreyanov, P. and Caoui, E. (2020). Secret reserve prices by uninformed sellers. Working paper, University of Toronto.

Aradillas-López, A., Gandhi, A., and Quint, D. (2013). Identification and inference in ascending auctions with correlated private values. *Econometrica*, 81(2):489–534.

Athey, S. and Haile, P. A. (2002). Identification of standard auction models. *Econometrica*, 70(6):2107–2140.

Athey, S. and Haile, P. A. (2007). Nonparametric approaches to auctions. In Heckman, J. J. and Leamer, E. E., editors, *Handbook of Econometrics, Vol. 6A*, pages 3847–3965. North-Holland.

Austin, D., Seljan, S., Moreno, J., and Tzeng, S. (2016). Reserve price optimization at scale. In *Proceedings of the 2016 IEEE International Conference on Data Science and Advanced Analytics*, pages 528–536.

Backus, M., Blake, T., Larsen, B., and Tadelis, S. (2020). Sequential bargaining in the field: Evidence from millions of online bargaining interactions. *Quarterly Journal of Economics*, 135(3):1319–1361.

Backus, M. and Lewis, G. (2020). Dynamic demand estimation in auction markets. Working paper, Columbia University.

Balseiro, S. R., Besbes, O., and Weintraub, G. Y. (2015). Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science*, 61(4):864–888.

Bodoh-Creed, A., Boehnke, J., and Hickman, B. (2020). How efficient are decentralized auction platforms? *Review of Economic Studies*, forthcoming.

Bounie, D., Dubus, A., and Waelbroeck, P. (2020). Selling strategic information in digital competitive markets. *RAND Journal of Economics*, forthcoming.

Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, 86(1):180–194.

Bulow, J. and Roberts, J. (1989). The simple economics of optimal auctions. *Journal of Political Economy*, 97(5):1060–1090.

Cai, H., Riley, J., and Ye, L. (2007). Reserve price signaling. *Journal of Economic Theory*, 135(1):253–268.

Cattaneo, M. D., Jansson, M., and Nagasawa, K. (2020). Bootstrap-based inference for cube root asymptotics. *Econometrica*, 88(5):2203–2219.

Celis, L. E., Lewis, G., Mobius, M., and Nazerzadeh, H. (2014). Buy-it-now or take-a-chance: Price discrimination through randomization auctions. *Management Science*, 60(12):2927–2948.

Cesa-Bianchi, N., Gentile, C., and Mansour, Y. (2015). Regret-minimization for reserve prices in second-price auctions. *IEEE Transactions on Information Theory*, 61(1):549–564.

Chawla, S., Hartline, J., and Nekipelov, D. (2014). Mechanism design for data science. arXiv preprint arXiv:1404.5971.

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, 1:451–488.

Choi, H. and Mela, C. F. (2019). Display advertising pricing in exchange markets. Working paper, University of Rochester.

Coey, D., Larsen, B., and Sweeney, K. (2019). The bidder exclusion effect. *RAND Journal of Economics*, 50(1):93–120.

Coey, D., Larsen, B., Sweeney, K., and Waisman, C. (2017). Ascending auctions with bidder asymmetries. *Quantitative Economics*, 8(1):181–200.

Coey, D., Larsen, B. J., and Platt, B. (2020). Discounts and deadlines in consumer search. *American Economic Review*, 110(12):3748–3785.

Cole, R. and Roughgarden, T. (2014). The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 243–252.

Decarolis, F., Goldmanis, M., and Penta, A. (2020). Marketing agencies and collusive bidding in online ad auctions. *Management Science*, 66(10):4433–4454.

Delgado, M. A., Rodriguez-Poo, J. M., and Wolf, M. (2001). Subsampling inference in cube root asymptotics with an application to Manski's maximum score estimator. *Economics*

Letters, 73(2):241–250.

Einav, L., Farronato, C., Levin, J. D., and Sundaresan, N. (2018). Auctions versus posted prices in online markets. *Journal of Political Economy*, 126(1):178–215.

Freyberger, J. and Larsen, B. (2019). Identification in ascending auctions, with an application to digital rights management. Working paper, Stanford University.

Guerre, E., Perrigne, I., and Vuong, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3):525–574.

Haile, P. A. and Tamer, E. (2003). Inference with an incomplete model of English auctions. *Journal of Political Economy*, 111(1):1–51.

Hendricks, K. and Sorensen, A. (2018). Dynamics and efficiency in decentralized online auction markets. NBER Working Paper 25002.

Hernández, C., Quint, D., and Turansick, C. (2020). Estimation in English auctions with unobserved heterogeneity. *RAND Journal of Economics*, 51(3):868–904.

Hong, H. and Li, J. (2020). The numerical bootstrap. *Annals of Statistics*, 48(1):397–412.

Horowitz, J. L. (1992). A smoothed maximum score estimator for the binary response model. *Econometrica*, 60(3):505–531.

Hortaçsu, A. and Nielsen, E. R. (2010). Commentary: Do bids equal values on eBay? *Marketing Science*, 29(6):994–997.

Kanoria, Y. and Nazerzadeh, H. (2020). Incentive-compatible learning of reserve prices for repeated auctions. *Operations Research*, forthcoming.

Kehoe, P. J., Larsen, B. J., and Pastorino, E. (2020). Dynamic competition in the era of big data. Working paper, Stanford University.

Kim, D.-H. (2013). Optimal choice of a reserve price under uncertainty. *International Journal of Industrial Organization*, 31(5):587–602.

Kim, J. and Pollard, D. (1990). Cube root asymptotics. *Annals of Statistics*, 18(1):191–219.

Koltchinskii, V. (2001). Rademacher penalties and structural risk minimization. *IEEE Transactions on Information Theory*, 47(5):1902–1914.

Koltchinskii, V. and Panchenko, D. (2002). Empirical margin distributions and bounding the generalization error of combined classifiers. *Annals of Statistics*, 30(1):1–50.

Lacetera, N., Larsen, B. J., Pope, D. G., and Sydnor, J. R. (2016). Bid takers or market makers? The effect of auctioneers on auction outcomes. *American Economic Journal: Microeconomics*, 8(4):195–229.

Larsen, B. (2020). The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. *Review of Economic Studies*, forthcoming.

Larsen, B. and Zhang, A. (2018). A mechanism design approach to identification and estimation. NBER Working Paper 24837.

Lee, S. M. S. and Pun, M. C. (2006). On m out of n bootstrapping for nonstandard M-estimation with nuisance parameters. *Journal of the American Statistical Association*, 101(475):1185–1197.

Levin, D. and Smith, J. L. (1994). Equilibrium in auctions with entry. *American Economic Review*, 84(3):585–599.

Li, T., Perrigne, I., and Vuong, Q. (2003). Semiparametric estimation of the optimal reserve price in first-price auctions. *Journal of Business & Economic Statistics*, 21(1):53–64.

Luo, Y. and Xiao, R. (2020). Identification of auction models using order statistics. Working paper, University of Toronto.

Manski, C. F. (1975). Maximum score estimation of the stochastic utility model of choice. *Journal of Econometrics*, 3(3):205–228.

Manski, C. F. (1985). Semiparametric analysis of discrete response: Asymptotic properties of the maximum score estimator. *Journal of Econometrics*, 27(3):313–333.

Mbakop, E. (2017). Identification of auctions with incomplete bid data in the presence of unobserved heterogeneity. Working paper, University of Calgary.

Milgrom, P. R. and Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122.

Mohri, M. and Medina, A. M. (2016). Learning algorithms for second-price auctions with reserve. *Journal of Machine Learning Research*, 17:1–25.

Mohri, M., Rostamizadeh, A., and Talwalkar, A. (2012). *Foundations of Machine Learning*. MIT press.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.

Newey, W. K. and McFadden, D. (1994). Large sample estimation and hypothesis testing. In Engle, R. F. and McFadden, D. L., editors, *Handbook of Econometrics, Vol. 4*, pages 2111–2245. North-Holland.

Ostrovsky, M. and Schwarz, M. (2016). Reserve prices in internet advertising auctions: A field experiment. Working paper, Stanford University.

Platt, B. C. (2017). Inferring ascending auction participation from observed bidders. *International Journal of Industrial Organization*, 54:65–88.

Pollard, D. (1989). Asymptotics via empirical processes. *Statistical Science*, 4(4):341–354.

Prasad, K. (2008). Price asymptotics. *Review of Economic Design*, 12(1):21–32.

Quint, D. (2017). Common values and low reserve prices. *Journal of Industrial Economics*, LXV(2):363–396.

Rhuggenaath, J., Akcay, A., Zhang, Y., and Kaymak, U. (2019). Optimizing reserve prices for publishers in online ad auctions. In *Proceedings of the 2019 IEEE Conference on Computational Inteliggence for Financial Engineering and Economics*, pages 1–8.

Roughgarden, T. (2014). Approximately optimal mechanism design: Motivation, examples, and lessons learned. *ACM SIGEcom Exchanges*, 13(2):4–20.

Rudolph, M. R., Ellis, J. G., and Blei, D. M. (2016). Objective variables for probabilistic revenue maximization in second-price auctions with reserve. In *Proceedings of the 25th International Conference on World Wide Web*, pages 1113–1122.

Samuelson, W. F. (1985). Competitive bidding with entry costs. *Economics Letters*, 17(1):53–57.

Segal, I. (2003). Optimal pricing mechanisms with unknown demand. *American Economic Review*, 93(3):509–529.

Song, U. (2004). Nonparametric estimation of an eBay auction model with an unknown number of bidders. Working paper, University of British Columbia.

Tang, X. (2011). Bounds on revenue distributions in counterfactual auctions with reserve prices. *RAND Journal of Economics*, 42(1):175–203.

van den Berg, G. J. (2007). On the uniqueness of optimal prices set by monopolistic sellers. *Journal of Econometrics*, 141(2):482–491.

van der Vaart, A. W. and Wellner, J. A. (1996). *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer.

Waisman, C. (2020). Selling mechanisms for perishable goods: An empirical analysis of an online resale market for event tickets. Working paper, Northwestern University.

Zeithammer, R. (2006). Forward-looking bidding in online auctions. *Journal of Marketing Research*, 43(3):462–476.

APPENDIX A. PROOFS

Proof of Theorem 1

Proof. We will show that $\hat{p}(r)$ converges uniformly in probability to $p(r)$. Note that $\pi(V_j^{(1)}, V_j^{(2)}, v_0, r)$ can alternatively be written as $\max\{r, V_j^{(2)}\} + (v_0 - r)\mathbb{1}(V_j^{(1)} < r)$ (see, for example, Aradillas-López et al. 2013). For simplicity, assume that $v_0 = 0$ and let $\hat{p}_1(r) = \frac{1}{J}\sum_{j=1}^{J} \max\left\{r, V_j^{(2)}\right\}$ and $\hat{p}_2(r) = -\frac{1}{J}\sum_{j=1}^{J} r\mathbb{1}\left\{V_j^{(1)} < r\right\}$, so that $\hat{p}(r) = \hat{p}_1(r) + \hat{p}_2(r)$. Notice that $\hat{p}_1(r)$ is Lipschitz continuous because, for any r_1 and r_2, it follows that $|\hat{p}_1(r_1) - \hat{p}_1(r_2)| \leqslant |r_1 - r_2|$. Furthermore, for any r, it follows by the law of large numbers that $\hat{p}_1(r) \xrightarrow{p} p_1(r)$. Thus, we can invoke Lemma 2.9 in Newey and McFadden (1994) to obtain $\sup_{r\in\mathcal{R}} |\hat{p}_1(r) - p_1(r)| \xrightarrow{p} 0$. Finally, it is straightforward to check that the function $f(x, r) = -r\mathbb{1}\{x \leqslant r\}$ belongs to a VC subgraph class (see, for example, van der Vaart and Wellner 1996), which guarantees uniform convergence of $\hat{p}_2(\cdot)$. Consequently, we have $\sup_{r\in\mathcal{R}} |\hat{p}(r) - p(r)| \xrightarrow{p} 0$, which guarantees that $\hat{r}_J \xrightarrow{p} r^*$ and $\hat{p}(\hat{r}_J) \xrightarrow{p} p(r^*)$. □

Proof of Theorem 2

Proof. We first note that the Kim and Pollard (1990) result requires the following assumption, the definition of which is somewhat involved. Let $\tilde{\pi}(\cdot, r) \equiv \pi(\cdot, r) - \pi(\cdot, r^*)$ and let $P_R(\cdot)$ be defined as the supremum of $|\tilde{\pi}(\cdot, r)|$ over the class $\mathcal{P}_R \equiv \{\tilde{\pi}(\cdot, r) : |r - r^*| \leqslant R\}$. The functions $P_R(\cdot)$ are referred to as *envelopes* of the classes $\mathcal{P}_R$. The class $\mathcal{P}_R$ is referred to as *manageable* for the envelope P_R (Def. 4.1 of Pollard 1989) if there exists a decreasing function $D(\cdot)$ for which (i) $\int_0^1 (\log D(x))^{1/2} dx < \infty$ and (ii) for every measure Q with finite support, $D_2(\epsilon(QP_R^2)^{1/2}, \mathcal{P}_R, Q) < D(\epsilon)$ for $0 < \epsilon < 1$, where the term $D_2(\epsilon, \mathcal{P}_R, Q)$ equals the largest J for which there are functions $P_{R,1}, ..., P_{R,J}$ in $\mathcal{P}_R$ with $\Pr_j |P_{R,k} - P_{R,j}|^2 > \epsilon^2$ for $k \neq j$; and where $\Pr_j$ is the empirical probability measure. The class $\mathcal{P}_R$ is then referred to as *uniformly manageable* if there exists a function that bounds every subclass of $\mathcal{P}_R$; the precise properties this bounding function must satisfy are described in detail in Section 3 of Kim and Pollard (1990). The condition required by Kim and Pollard (1990), applied to our context, is as follows:

Assumption 4.
The classes $\mathcal{P}_R$, for R near 0, are uniformly manageable for the envelopes P_R.

It is straightforward, though tedious, to show that this assumption of uniformly manageability is satisfied in our context. We omit the proof here.

Throughout the rest of the proof of our Theorem 2, we use r_1 and r_2 such that, without loss of generality, $r_1 > r_2 > r^*$. We also denote the marginal densities of $V^{(1)}$ and $V^{(2)}$ by $f_1(\cdot)$ and $f_2(\cdot)$, respectively. Also, let $-\Sigma$ denote the second derivative of $\mathbb{E}[\tilde{\pi}(\cdot, r)]$ evaluated at r^*.

By the main theorem of Kim and Pollard (1990), if Assumptions 3 and 4 are satisfied, and r^* is an interior point, and if the following conditions hold

(1) $H(\beta, \alpha) = \lim_{t\to 0} \frac{1}{t}\mathbb{E}[\tilde{\pi}(\cdot, r^* + \beta t)\tilde{\pi}(\cdot, r^* + \alpha t)]$ exists for each β, α in $\mathbb{R}$ and

$$\lim_{t\to 0} \frac{1}{t}\mathbb{E}[\tilde{\pi}(\cdot, r^* + \alpha t)^2 \mathbb{1}\{|\tilde{\pi}(\cdot, r^* + \alpha t)| > \epsilon/t\}] = 0$$

for each $\epsilon > 0$ and α in $\mathbb{R}$;

(2) $\mathbb{E}[P_R^2] = O(R)$ as $R \to 0$ and for each $\epsilon > 0$ there is a constant K such that $\mathbb{E}[P_R^2 \mathbb{1}\{P_R > K\}] < \epsilon R$ for R near 0;

(3) $\mathbb{E}[|\tilde{\pi}(\cdot, r_1) - \tilde{\pi}(\cdot, r_2)|] = O(|r_1 - r_2|)$ near r^*;

then the process $J^{2/3}\frac{1}{J}\sum_{j=1}^{J} \tilde{\pi}(\cdot, r^* + \alpha J^{-1/3})$ converges in distribution to a Gaussian process $Z(\alpha)$ with continuous sample paths, expected value $-\frac{1}{2}\alpha^2\Sigma$, and covariance kernel H, where $H(\beta, \alpha) = \lim_{t\to 0} \frac{1}{t}\mathbb{E}[\tilde{\pi}(\cdot, r^* + \beta t)\tilde{\pi}(\cdot, r^* + \alpha t)]$ for any β, α in $\mathbb{R}$. Furthermore, if Z has nondegenerate increments, then $J^{1/3}(\hat{r}_J - r^*)$ converges in distribution to the random maximizer of Z. We now prove that conditions (1)–(3) above are satisfied.

To establish condition (1), we first characterize the limiting behavior of $\frac{1}{t}\mathbb{E}[\tilde{\pi}(\cdot, r^* + \beta t)\tilde{\pi}(\cdot, r^* + \alpha t)]$ as $t \to 0$. By the definition of $\tilde{\pi}(\cdot, r)$, this amounts to studying the behavior of four different terms, which we conduct separately below. Let $r_1 = r^* + \alpha t$ and $r_2 = r^* + \beta t$. First, we consider the term

$$h_1 \equiv \left(\max\{V^{(2)}, r_1\} - \max\{V^{(2)}, r^*\}\right)\left(\max\{V^{(2)}, r_2\} - \max\{V^{(2)}, r^*\}\right).$$

Notice that when $V^{(2)} > r_2$ then $h_1 = 0$; when $V^{(2)} < r^*$ then $h_1 = (r_1 - r^*)(r_2 - r^*)$; and when $r^* < V^{(2)} < r_2$ then $h_1 = (r_1 - V^{(2)})(r_2 - V^{(2)})$. Therefore,

$$\frac{1}{t}\mathbb{E}\left[\left(\max\{V^{(2)}, r_1\} - \max\{V^{(2)}, r^*\}\right)\left(\max\{V^{(2)}, r_2\} - \max\{V^{(2)}, r^*\}\right)\right]$$
$$= \frac{1}{t}\left\{(r_1 - r^*)(r_2 - r^*)Pr(V^{(2)} < r^*)\right.$$

$$+\mathbb{E}\left[(r_1 - V^{(2)})(r_2 - V^{(2)})|r^* < V^{(2)} < r_2\right] Pr(r^* < V^{(2)} < r_2)\Big\}$$

$$= \frac{1}{t}\left\{\alpha\beta t^2 \int_0^{r^*} f_2(u)du + r_1 r_2 \int_{r^*}^{r_2} f_2(u)du - (r_1 + r_2)\int_{r^*}^{r_2} u f_2(u)du + \int_{r^*}^{r_2} u^2 f_2(u)du\right\}$$

$$= \frac{1}{t}\left\{\alpha\beta t^2 F_2(r^*) + \left[(r^*)^2 + (\alpha+\beta)t + \alpha\beta t^2\right]\left[f_2(r^*)(r_2 - r^*) + o(r_2 - r^*)\right]\right.$$

$$\left. - \left[2r^* + (\alpha+\beta)t\right]\left[r^* f_2(r^*)(r_2 - r^*) + o(r_2 - r^*)\right] + \left[(r^*)^2 f_2(r^*)(r_2 - r^*) + o(r_2 - r^*)\right]\right\}$$

$$= \frac{1}{t}\left\{\alpha\beta t^2 F_2(r^*) + \alpha\beta t^2 f_2(r^*)\beta t + o(t)\right\}$$

$$= \frac{1}{t} o(t) = o(1).$$

The second term we consider is

$$h_2 \equiv \left(r^* \mathbb{1}\{V^{(1)} < r^*\} - r_2 \mathbb{1}\{V^{(1)} < r_2\}\right)\left(\max\{V^{(2)}, r_1\} - \max\{V^{(2)}, r^*\}\right).$$

When $V^{(1)} < r^*$, $h_2 = (r^* - r_2)(r_1 - r^*)$; when $r^* < V^{(1)} < r_2$ and $V^{(2)} < r^*$, $h_2 = -r_2(r_1 - r^*)$; and when $r^* < V^{(1)} < r_2$ and $r^* < V^{(2)} < V^{(1)}$, $h_2 = -r_2(r_1 - V^{(2)})$. Hence,

$$\frac{1}{t}\mathbb{E}\left[\left(r^* \mathbb{1}\{V^{(1)} < r^*\} - r_2 \mathbb{1}\{V^{(1)} < r_2\}\right)\left(\max\{V^{(2)}, r_1\} - \max\{V^{(2)}, r^*\}\right)\right]$$

$$= \frac{1}{t}\Big\{(r^* - r_2)(r_1 - r^*) Pr(V^{(1)} < r^*) - r_2(r_1 - r^*) Pr(V^{(2)} < r^* < V^{(1)} < r_2)$$

$$- r_2 r_1 Pr(r^* < V^{(2)} < V^{(1)} < r_2)$$

$$+\mathbb{E}[V^{(2)}|r^* < V^{(2)} < V^{(1)} < r_2] Pr(r^* < V^{(2)} < V^{(1)} < r_2)\Big\}$$

$$= \frac{1}{t}\left\{-\alpha\beta t^2 F_1(r^*) - r_2 \alpha t \int_{r^*}^{r_2}\int_0^{r^*} f_{1,2}(u,v)dudv - r_1 r_2 \int_{r^*}^{r_2}\int_{r^*}^{v} f_{1,2}(u,v)dudv\right.$$

$$\left. + r_2 \int_{r^*}^{r_2}\int_{r^*}^{v} u f_{1,2}(u,v)dudv\right\}$$

$$= \frac{1}{t}\left\{-\alpha\beta t^2 F_1(r^*) - r_2 \alpha t \int_{r^*}^{r_2} F_v(r^*, v)dv - r_1 r_2 \int_{r^*}^{r_2}[F_v(v,v) - F_v(r^*, v)]dv\right.$$

$$\left. + r_2 \int_{r^*}^{r_2}[F_{1,2}(v,v) - F_{1,2}(v, r^*)]dv\right\}$$

$$= \frac{1}{t}\left\{-\alpha\beta t^2 F_1(r^*) - r_2 \alpha t[\beta t F_v(r^*, r^*) + o(t)] - r_1 r_2 o(t) + r_2 o(t)\right\}$$

$$= \frac{1}{t} o(t) = o(1).$$

The third term is

$$h_3 \equiv \left(r^*\mathbb{1}\{V^{(1)} < r^*\} - r_1\mathbb{1}\{V^{(1)} < r_1\}\right)\left(\max\{V^{(2)}, r_2\} - \max\{V^{(2)}, r^*\}\right).$$

The term $h_3 = 0$ when $V^{(1)} > r_1$ or $V^{(2)} > r_2$; $h_3 = -r_1(r_2 - r^*)$ when $V^{(2)} < r^* < V^{(1)} < r_1$; $h_3 = -r_1(r_2 - V^{(2)})$ when $r^* < V^{(1)} < r_1$ and $r^* < V^{(2)} < r_2$; and $h_3 = (r^* - r_1)(r_2 - r^*)$ when $V^{(1)} < r^*$. Consequently,

$$\begin{aligned}
&\frac{1}{t}\mathbb{E}\left[\left(r^*\mathbb{1}\{V^{(1)} < r^*\} - r_1\mathbb{1}\{V^{(1)} < r_1\}\right)\left(\max\{V^{(2)}, r_2\} - \max\{V^{(2)}, r^*\}\right)\right] \\
&= \frac{1}{t}\Big\{-r_1(r_2 - r^*)Pr(V^{(2)} < r^* < V^{(1)} < r_1) \\
&\quad -r_1\left(r_2 - \mathbb{E}[V^{(2)}|r^* < V^{(1)} < r_1, r^* < V^{(2)} < r_2]\right)Pr(r^* < V^{(1)} < r_1, r^* < V^{(2)} < r_2)\Big\} \\
&= \frac{1}{t}\Big\{-r_1\beta t\int_{r^*}^{r_1}\int_0^{r^*} f_{1,2}(u,v)dudv - \alpha\beta t^2 F_1(r^*) - r_1r_2\int_{r^*}^{r_2}\int_{r^*}^{v} f_{1,2}(u,v)dudv \\
&\quad -r_1r_2\int_{r_2}^{r_1}\int_{r^*}^{r_2} f_{1,2}(u,v)dudv \\
&\quad +r_1\int_{r^*}^{r_2}\int_{r^*}^{v} uf_{1,2}(u,v)dudv + r_1\int_{r_2}^{r_1}\int_{r^*}^{r_2} uf_{1,2}(u,v)dudv\Big\} \\
&= \frac{1}{t}\Big\{-r_1\beta t\int_{r^*}^{r_1} F_v(r^*, v)dv - \alpha\beta t^2 F_1(r^*) - r_1r_2\int_{r^*}^{r_2}[F_v(v,v) - F_v(r^*, v)]dv \\
&\quad -r_1r_2\int_{r_2}^{r_1} f_{1,2}(r^*, v)(r_2 - r^*)dv \\
&\quad +r_1\int_{r^*}^{r_2}[F_{1,2}(v,v) - F_{1,2}(v,r^*)]dv + r_1r^*\int_{r_2}^{r_1} f_{1,2}(r^*, v)(r_2 - r^*)dv + o[(r_2 - r^*)^2]\Big\} \\
&= \frac{1}{t}\big\{-r_1\alpha\beta t^2 F_v(r^*, r^*) - \alpha\beta t^2 F_1(r^*) - r_1r_2 o(r_1 - r^*) \\
&\quad +r_1 o(r_2 - r^*) - r_1\beta^2t^2\int_{r_2}^{r_1} f_{1,2}(r^*, v)dv + o[(r_2 - r^*)^2]\Big\} \\
&= \frac{1}{t}o(t) = o(1).
\end{aligned}$$

The fourth and final term is

$$h_4 \equiv \left(r^*\mathbb{1}\{V^{(1)} < r^*\} - r_2\mathbb{1}\{V^{(1)} < r_2\}\right)\left(r^*\mathbb{1}\{V^{(1)} < r^*\} - r_1\mathbb{1}\{V^{(1)} < r_1\}\right).$$

The term $h_4 = 0$ if $V^{(1)} > r_2$; $h_4 = r_1 r_2$ if $r^* < V^{(1)} < r_2$; and $h_4 = (r^* - r_2)(r^* - r_1)$ if $V^{(1)} < r^*$. Thus,

$$\begin{aligned}
&\frac{1}{t}\mathbb{E}\left[\left(r^*\mathbb{1}\{V^{(1)} < r^*\} - r_2\mathbb{1}\{V^{(1)} < r_2\}\right)\left(r^*\mathbb{1}\{V^{(1)} < r^*\} - r_1\mathbb{1}\{V^{(1)} < r_1\}\right)\right] \\
&= \frac{1}{t}\left\{r_1 r_2 Pr(r^* < V^{(1)} < r_2) + (r^* - r_2)(r^* - r_1)Pr(V^{(1)} < r^*)\right\} \\
&= \frac{1}{t}\left\{r_1 r_2 \int_{r^*}^{r_2} f_1(v)dv + \alpha\beta t^2 F_1(r^*)\right\} \\
&= \frac{1}{t}\left\{r_1 r_2[f_1(r^*)(r_2 - r^*) + o(r_2 - r^*)] + \alpha\beta t^2 F_1(r^*)\right\} \\
&\frac{1}{t}\left\{(r^*)^2 f_1(r^*)\beta t + o(t)\right\} \\
&= (r^*)^2 f_1(r^*)\beta + o(1).
\end{aligned}$$

These four results show that the limit $H(\beta, \alpha)$ is well defined for each β, α in $\mathbb{R}$, which establishes the first part of (1). For the second part of (1), notice that $|\tilde{\pi}(\cdot, r^* + \alpha t)|$ is bounded for any t, which means that there exists a $\underline{t} < \infty$ such that the indicator would be 0 for all $t < \underline{t}$, establishing the result.

We now establish (2). Let $R > 0$ and $\check{r}$ be the maximizer of $\tilde{\pi}(\cdot, r)$ such that $|r - r^*| < R$. We first need to show that $\mathbb{E}[\tilde{\pi}(\cdot, \check{r})^2] = O(R)$. We will split the analysis into three terms. First, notice that for the first term:

$$(\max\{V^{(2)}, \check{r}\} - \max\{V^{(2)}, r^*\})^2 = (|\max\{V^{(2)}, \check{r}\} - \max\{V^{(2)}, r^*\}|)^2 \quad \leqslant (|\check{r} - r^*|)^2 < R^2$$

which implies that its expected value is $O(R^2)$. Moving to the next two terms we will assume that $\check{r} > r^*$, as the calculations for the opposite case are analogous. For the second one,

$$(r^*\mathbb{1}\{V^{(1)} < r^*\} - \check{r}\mathbb{1}\{V^{(1)} < \check{r}\})^2 = (r^*)^2\mathbb{1}\{V^{(1)} < r^*\} + \check{r}^2\mathbb{1}\{V^{(1)} < \check{r}\})^2 - 2r^*\check{r}\mathbb{1}\{V^{(1)} < r^*\}$$

so taking expectations yields:

$$\begin{aligned}
&(r^*)^2\int_0^{r^*} f_1(v)dv - r^*\check{r}\int_0^{r^*} f_1(v)dv + +\check{r}^2\int_{r^*}^{\check{r}} f_1(v)dv + \check{r}^2\int_0^{r^*} f_1(v)dv - r^*\check{r}\int_0^{r^*} f_1(v)dv \\
&= (\check{r} - r^*)^2 F_1(r^*) + \check{r}^2[f_1(r^*)(\check{r} - r^*) + o(\check{r} - r^*)] \\
&= O(\check{r} - r^*) + o(\check{r} - r^*) = O(\check{r} - r^*) < O(R).
\end{aligned}$$

The third and last term is given by:

$$
\begin{aligned}
&\mathbb{E}\left[(\max\{V^{(2)},\check{r}\}-\max\{V^{(2)},r^*\})(r^*\mathbb{1}\{V^{(1)}<r^*\}-\check{r}\mathbb{1}\{V^{(1)}<\check{r}\})\right]\\
&=\mathbb{E}\left[(\max\{V^{(2)},r^*\}r^*\mathbb{1}\{V^{(1)}<r^*\}\right]-\mathbb{E}\left[(\max\{V^{(2)},\check{r}\}\check{r}\mathbb{1}\{V^{(1)}<\check{r}\}\right]\\
&\quad+\mathbb{E}\left[(\max\{V^{(2)},r^*\}\check{r}\mathbb{1}\{V^{(1)}<\check{r}\}\right]-\mathbb{E}\left[(\max\{V^{(2)},\check{r}\}r^*\mathbb{1}\{V^{(1)}<r^*\}\right]\\
&=(r^*)^2\int_0^{r^*}\int_0^{v}f_{1,2}(u,v)dudv-\check{r}^2\int_0^{\check{r}}\int_0^{v}f_{1,2}(u,v)dudv\\
&\quad+\check{r}\int_0^{\check{r}}\int_0^{v}\max\{u,r^*\}f_{1,2}(u,v)dudv-r^*\check{r}\int_0^{r^*}\int_0^{v}f_{1,2}(u,v)dudv\\
&=[(r^*)^2-\check{r}^2]\int_0^{r^*}\int_0^{v}f_{1,2}(u,v)dudv-\check{r}^2\int_{r^*}^{\check{r}}\int_0^{v}f_{1,2}(u,v)dudv\\
&\quad+\check{r}\int_{r^*}^{\check{r}}\int_0^{v}\max\{u,r^*\}f_{1,2}(u,v)dudv\\
&=(r^*-\check{r})(r^*+\check{r})F_{1,2}(r^*,r^*)-\check{r}^2\int_{r^*}^{\check{r}}F_v(v,v)dv+r^*\check{r}\int_{r^*}^{\check{r}}\int_0^{r^*}f_{1,2}(u,v)dudv\\
&\quad+\check{r}\int_{r^*}^{\check{r}}\int_{r^*}^{v}uf_{1,2}(u,v)dudv\\
&=O(r^*-\check{r})-\check{r}(\check{r}-r^*)[F_v(r^*,r^*)(r^*-\check{r})+o(r^*-\check{r})]+\check{r}\int_{r^*}^{\check{r}}G(v)dv\\
&=O(r^*-\check{r})+\check{r}[G(r^*)(\check{r}-r^*)+o(\check{r}-r^*)]=O(\check{r}-r^*)<O(R),
\end{aligned}
$$

which, along with the results for the previous two terms, establishes the first part of condition (2). The second part of condition (2) follows directly from the integrability of $\tilde{\pi}(\cdot,\check{r})^2$.

To verify that (3) holds, notice that:

$$
\begin{aligned}
&|\tilde{\pi}(\xi_j,r_1)-\tilde{\pi}(\xi_j,r_2)|\\
&=|\pi(\xi_j,r_1)-\pi(\xi_j,r_2)|\\
&=\left|\max\left\{V_j^{(2)},r_1\right\}-r_1\mathbb{1}\left\{V_j^{(1)}<r_1\right\}-\max\left\{V_j^{(2)},r_2\right\}+r_2\mathbb{1}\left\{V_j^{(1)}<r_2\right\}\right|\\
&\leqslant\left|\max\left\{V_j^{(2)},r_1\right\}-\max\left\{V_j^{(2)},r_2\right\}\right|+\left|r_2\mathbb{1}\left\{V_j^{(1)}<r_2\right\}-r_1\mathbb{1}\left\{V_j^{(1)}<r_1\right\}\right|\\
&\leqslant|r_2-r_1|+\mathbb{1}\left\{V_j^{(1)}<r_2\right\}|r_2-r_1|+r_1\left|\mathbb{1}\left\{V_j^{(1)}<r_2\right\}-\mathbb{1}\left\{V_j^{(1)}<r_1\right\}\right|
\end{aligned}
$$

Taking the expectation, we obtain

$$
\mathbb{E}\left|\tilde{\pi}(\xi_j,r_1)-\tilde{\pi}(\xi_j,r_2)\right|\leqslant|r_2-r_1|+|r_2-r_1|Pr(V_j^{(1)}<r_2)+r_1Pr(r_2<V_j^{(1)}<r_1)
$$

$$= O(|r_2 - r_1|) + f_1(r_2)(r_1 - r_2) + o(r_1 - r_2) = O(|r_2 - r_1|).$$

This establishes (3).

Finally, we derive Σ. Notice that:

$$\begin{aligned}
\mathbb{E}[\tilde{\pi}(\cdot, r)] &= \mathbb{E}[\pi(\cdot, r)] - \mathbb{E}[\pi(\cdot, r^*)] \\
&= \int_0^{\bar{\omega}} \max\{r, V^{(2)}\} f_2(u) du - \int_0^{\bar{\omega}} r \mathbb{1}\{V^{(1)} < r\} f_1(v) dv - \mathbb{E}[\pi(\cdot, r^*)] \\
&= r \int_0^r f_2(u) du + \int_r^{\bar{\omega}} u f_2(u) du - r \int_0^r f_1(v) dv - \mathbb{E}[\pi(\cdot, r^*)]
\end{aligned}$$

Differentiating this expression with respect to r, we obtain

$$\begin{aligned}
\frac{\partial \mathbb{E}[\tilde{\pi}(\cdot, r)]}{\partial r} &= \int_0^r f_2(u) du + r f_2(r) - r f_2(r) - \int_0^r f_1(v) dv - r f_1(r) \\
&= \int_0^r f_2(u) du - \int_0^r f_1(v) dv - r f_1(r)
\end{aligned}$$

Taking the second derivative and evaluating at r^*, we obtain

$$\Sigma = -f_2(r^*) + 2 f_1(r^*) + r^* f_1'(r^*).$$

□

Proof of Theorem 3

Proof. This theorem follows directly from Theorem 2. First, notice that $\hat{p}(\hat{r}_J) - p(r^*) = \frac{1}{J} \sum_{j=1}^J \tilde{\pi}(\xi_j, \hat{r}_J) + \hat{p}(r^*) - p(r^*)$. Since $\hat{r}_J - r^* = O_P(J^{-1/3})$ and $\frac{1}{J} \sum_{j=1}^J \tilde{\pi}(\cdot, r^* + \alpha J^{-1/3}) = O_P(J^{-2/3})$, the term $\frac{1}{J} \sum_{j=1}^J \tilde{\pi}(\xi_j, \hat{r}_J)$ is also $O_P(J^{-2/3})$. A simple application of the Central Limit Theorem establishes the result. □

Proof of Theorem 4 We will use slightly different notation in this proof than elsewhere in the paper, letting $S_J = (z_1, \ldots, z_J)$ be a fixed sample of size J and denoting quantities estimated on this sample by a subscript S_J.

We start with the definition of *empirical Rademacher complexity*:

Definition. *Let G be a family of functions from Z to $[a, b]$, and $S_J = (z_1, \ldots, z_J)$ a fixed sample of size J with elements in Z. Then the empirical Rademacher complexity of G with respect to S_J is* *defined as:*

$$\widehat{\mathfrak{R}}_{S_J}(G) = E_{\sigma}\left[\sup_{g \in G} \frac{1}{J}\sum_{j=1}^{J} \sigma_j g(z_j)\right]$$

where $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_J)$*, and each* σ_j *is an independent uniform random variable with values in* $\{-1, +1\}$.

We state a useful preliminary result.

Theorem 5. *Let* G *be a family of functions mapping* Z *to* $[0, 1]$*. Then for any* $\delta > 0$*, with probability at least* $1 - \delta$*, for all* $g \in G$*:*

$$E[g(z)] - \frac{1}{J}\sum_{j=1}^{J} g(z_j) \leqslant 2\widehat{\mathfrak{R}}_{S_J}(G) + 3\sqrt{\frac{\log \frac{2}{\delta}}{2J}}$$

Proof. See Mohri et al. (2012), Theorem 3.1. □

This result can be straightforwardly adapted to obtain a two-sided bound.

Corollary 1. *Let* G *be a family of functions mapping* Z *to* $[0, 1]$*. Then for any* $\delta > 0$*, with probability at least* $1 - \delta$*, for all* $g \in G$*:*

$$\left| E[g(z)] - \frac{1}{J}\sum_{j=1}^{J} g(z_j) \right| \leqslant 2\widehat{\mathfrak{R}}_{S_J}(G) + 3\sqrt{\frac{\log \frac{4}{\delta}}{2J}}$$

Proof. Applying Theorem 5 above to $G' = \{-g + 1 : g \in G\}$ and noting that $\widehat{\mathfrak{R}}_{S_J}(G) = \widehat{\mathfrak{R}}_{S_J}(G')$, we obtain the result that for any $\delta/2 > 0$, with probability at least $1 - \delta/2$, for all $g \in G$:

$$\frac{1}{J}\sum_{j=1}^{J} g(z_j) - E[g(z)] \leqslant 2\widehat{\mathfrak{R}}_{S_J}(G) + 3\sqrt{\frac{\log \frac{4}{\delta}}{2J}}.$$

Theorem 5 also implies that for any $\delta/2 > 0$, with probability at least $1 - \delta/2$, for all $g \in G$:

$$E[g(z)] - \frac{1}{J}\sum_{j=1}^{J} g(z_j) \leqslant 2\widehat{\mathfrak{R}}_{S_J}(G) + 3\sqrt{\frac{\log \frac{4}{\delta}}{2J}}.$$

Combining these two results and applying the union bound gives the desired result. □

Now for simplicity let $v_0 = 0$ so that $\pi(V_j^{(1)}, V_j^{(2)}, r) = r\mathbb{1}(V_j^{(2)} < r \leqslant V_j^{(1)}) + V_j^{(2)}\mathbb{1}(r \leqslant V_j^{(2)})$ and $G \equiv \{\pi(\cdot,\cdot,r) : r \in [0,\overline{\omega}]\}$. We can now prove an upper bound on $p(r^*) - p(\widehat{r}_{S_J})$ in terms of the empirical Rademacher complexity of G.

Lemma 1. *Let $0 \leqslant V_j^{(1)} \leqslant \overline{\omega} < \infty$. For any $\delta > 0$, with probability at least $1-\delta$, it holds that $p(r^*) - p(\widehat{r}_{S_J}) \leqslant 4\widehat{\mathfrak{R}}_{S_J}(G) + 6\overline{\omega}\sqrt{\frac{\log\frac{4}{\delta}}{2J}}$.*

Proof.

$$\begin{aligned}
p(r^*) - p(\widehat{r}_{S_J}) &= p(r^*) - \widehat{p}_{S_J}(\widehat{r}_{S_J}) + \widehat{p}_{S_J}(\widehat{r}_{S_J}) - p(\widehat{r}_{S_J}) \\
&\leqslant p(r^*) - \widehat{p}_{S_J}(r^*) + \widehat{p}_{S_J}(\widehat{r}_{S_J}) - p(\widehat{r}_{S_J}). \\
&\leqslant 2\sup_{r\in\mathcal{R}} |p(r) - \widehat{p}_{S_J}(r)|.
\end{aligned}$$

The first inequality follows because $\widehat{r}_{S_J}$ maximizes $\widehat{p}_{S_J}$ by definition. Applying Corollary 1 with $z_j = (V_j^{(1)}, V_j^{(2)})$ we have that for any $\delta > 0$, with probability at least $1-\delta$:

$$\sup_{r\in[0,\overline{\omega}]}\left|\frac{1}{\overline{\omega}}E[\pi(V_j^{(1)}, V_j^{(2)}, r)] - \frac{1}{J\overline{\omega}}\sum_{j=1}^{J}\pi(V_j^{(1)}, V_j^{(2)}, r)\right| \leqslant \frac{2}{\overline{\omega}}\widehat{\mathfrak{R}}_{S_J}(G) + 3\sqrt{\frac{\log\frac{4}{\delta}}{2J}},$$

or equivalently,

$$\sup_{r\in\mathcal{R}}|p(r) - \widehat{p}_{S_J}(r)| \leqslant 2\widehat{\mathfrak{R}}_{S_J}(G) + 3\overline{\omega}\sqrt{\frac{\log\frac{4}{\delta}}{2J}}.$$

Therefore for any $\delta > 0$, with probability at least $1-\delta$:

$$p(r^*) - p(\widehat{r}_{S_J}) \leqslant 4\widehat{\mathfrak{R}}_{S_J}(G) + 6\overline{\omega}\sqrt{\frac{\log\frac{4}{\delta}}{2J}}.$$

□

Following Mohri and Medina (2016), define $\pi_1(V_j^{(1)}, V_j^{(2)}, r) = V_j^{(2)}\mathbb{1}(r \leqslant V_j^{(2)}) + r\mathbb{1}(V_j^{(2)} < r \leqslant V_j^{(1)}) + V_j^{(1)}\mathbb{1}(V_j^{(1)} < r)$ and $\pi_2(V_j^{(1)}, r) = -V_j^{(1)}\mathbb{1}(V_j^{(1)} < r)$, so that $\pi(V_j^{(1)}, V_j^{(2)}, r) = \pi_1(V_j^{(1)}, V_j^{(2)}, r) + \pi_2(V_j^{(1)}, r)$. Define also $G_1 = \{\pi_1(\cdot,\cdot,r) : r \in [0,\overline{\omega}]\}$ and $G_2 = \{\pi_2(\cdot, r) : r \in [0,\overline{\omega}]\}$. The following lemma is useful:

Lemma 2. *Let H be a set of functions mapping $\mathcal{X}$ to $\mathbb{R}$ and let $\Psi_1, \ldots, \Psi_J$ be μ-Lipschitz functions for some $\mu > 0$. Then for any sample S_J of J points $x_1, \ldots, x_J \in \mathcal{X}$, the following inequality holds:*

$$\frac{1}{J} E_\sigma \left[\sup_{h \in H} \sum_{j=1}^{J} \sigma_j (\Psi_j \circ h)(x_j) \right] \leqslant \frac{\mu}{J} E_\sigma \left[\sup_{h \in H} \sum_{j=1}^{J} \sigma_j h(x_j) \right].$$

Proof. See Lemma 14 in Mohri and Medina (2016). □

We now find an upper bound for the right hand side of Lemma 1, which is not expressed in terms of Rademacher complexity and which makes the asymptotic behavior of the term $p(r^*) - p(\widehat{r}_{S_J})$ clear. This will lead to Theorem 4.

Lemma 3. *Let $0 \leqslant V_i^{(1)} \leqslant \overline{\omega} < \infty$. Then $\widehat{\mathfrak{R}}_{S_J}(G) \leqslant \frac{2\overline{\omega}\sqrt{\log 2}}{J} + \overline{\omega}\sqrt{\frac{2+2\log J}{J}}$.*

Proof. Note that $\widehat{\mathfrak{R}}_{S_J}(G) \leqslant \widehat{\mathfrak{R}}_{S_J}(G_1) + \widehat{\mathfrak{R}}_{S_J}(G_2)$, as the supremum of a sum is less than the sum of suprema. We give upper bounds on both of these terms. For the first term, we have:

$$\begin{aligned}
\widehat{\mathfrak{R}}_{S_J}(G_1) &\equiv E_\sigma \left[\sup_{r \in [0,\overline{\omega}]} \frac{1}{J} \sum_{j=1}^{J} \sigma_j [V_j^{(2)} \mathbb{1}(r \leqslant V_j^{(2)}) + r \mathbb{1}(V_j^{(2)} < r \leqslant V_j^{(1)}) + V_j^{(1)} \mathbb{1}(V_j^{(1)} < r)] \right] \\
&\leqslant \frac{1}{J} E_\sigma \left[\sup_{r \in \mathcal{R}} \sum_{j=1}^{J} \sigma_j r \right] \\
&= \frac{1}{J} E_\sigma \left[\sup_{r \in \{0,\overline{\omega}\}} \sum_{j=1}^{J} \sigma_i r \right] \\
&\leqslant \frac{2\overline{\omega}\sqrt{\log 2}}{J}
\end{aligned}$$

The first inequality follows from applying Lemma 2 with $\Psi_j(x) \equiv \pi_1(V_j^{(1)}, V_j^{(2)}, x)$ and $h(x) \equiv x$, and the observation that the functions $\Psi_j(r)$ are $1-$Lipschitz for all j. The equality follows because the supremum will always be attained at $r = 0$ (if $\sum_{j=1}^{J} \sigma_i \leqslant 0$) or at $r = \overline{\omega}$ (if $\sum_{j=1}^{J} \sigma_j > 0$). The final inequality is an application of Massart's lemma (see, for example, Mohri et al. 2012).

For the second term, we have

$$\widehat{\mathfrak{R}}_{S_J}(G_2) \equiv E_\sigma \left[\sup_{r \in [0,\overline{\omega}]} \frac{1}{J} \sum_{j=1}^{J} -\sigma_j V_j^{(1)} \mathbb{1}(V_j^{(1)} < r) \right]$$

$$\leqslant \frac{\overline{\omega}}{J} E_\sigma \left[\sup_{r\in[0,\overline{\omega}]} \sum_{j=1}^{J} -\sigma_j \mathbb{1}(V_j^{(1)} < r) \right]$$
$$= \frac{\overline{\omega}}{J} E_\sigma \left[\sup_{r\in[0,\overline{\omega}]} \sum_{j=1}^{J} \sigma_j \mathbb{1}(V_j^{(1)} < r) \right]$$
$$\leqslant \overline{\omega}\sqrt{\frac{2+2\log J}{J}}.$$

The first inequality follows from applying Lemma 2 with $\Psi_j(x) \equiv V_j^{(1)} x$ and $h(x) \equiv \mathbb{1}(V_j^{(1)} < x)$, noting that $\Psi_j(x)$ are $\overline{\omega}$-Lipschitz for all j. The equality follows because the distributions of σ_j and $-\sigma_j$ are identical. Finally, the last inequality follows from Massart's lemma (see Proposition 2 in Mohri and Medina 2016).

Putting the bounds on $\widehat{\mathfrak{R}}_{S_J}(G_1)$ and $\widehat{\mathfrak{R}}_{S_J}(G_2)$ together, we have:

$$\widehat{\mathfrak{R}}_{S_J}(G) \leqslant \frac{2\overline{\omega}\sqrt{\log 2}}{J} + \overline{\omega}\sqrt{\frac{2+2\log J}{J}}.$$

□

This leads immediately to Theorem 4:

Theorem. *Let $0 \leqslant V_j^{(1)} \leqslant \overline{\omega} < \infty$. For any $\delta > 0$, with probability at least $1-\delta$ over the possible realizations of S_J, it holds that*

$$\frac{p(r^*) - p(\hat{r}_{S_J})}{\overline{\omega}} \leqslant \left(\frac{8\sqrt{\log 2}}{J} + 4\sqrt{\frac{2+2\log J}{J}} + 6\sqrt{\frac{\log \frac{4}{\delta}}{2J}} \right).$$

Proof. Combine Lemmas 1 and 3. □