Learning and Information Design on an Auction Platform

Joshua D. Higbee*

The University of Chicago

joshuadhigbee@uchicago.edu

March 20, 2025

Click **here** for latest draft

Online platforms often do not directly control users' pricing strategy, and instead offer analytics and other information to help steer user behavior. I study how information provision by an auction platform to sellers shapes platform fees and outcomes using data from eBay auctions of children's toys. I present evidence that new sellers face uncertainty about how to set optimal reserve prices: they set lower reserve prices and earn higher revenues as they gain more experience. I develop a model where new sellers learn to set reserve prices on an auction platform with selective participation, and I show that sellers choose reserve prices to both extract surplus from bidders and attract additional bidders to their auction. I provide conditions under which new sellers' beliefs about the effect of reserve prices on bidder arrival are semiparametrically identified. Estimates from the learning model indicate that new sellers underestimate the effect of high reserve prices on deterring bidder entry, which leads to higher reserve prices and more items listed than for fully-informed sellers. Counterfactual simulations show that platform information provision can help new sellers learn the true bidder arrival process, which increases bidder entry as well as seller and platform profits. However, when jointly choosing information provision and platform fees, it may be optimal to not fully reveal information to sellers.

*I thank Ali Hortaçsu, Dennis Carlton, Giovanni Compiani, Günter Hitsch, Eric Richert, Brad Larsen, Avner Strulov-Shlain, Eric Budish, Chad Syverson, Devin Pope, Stephane Bonhomme, Kirill Ponomarev, Max Farrell, and seminar participants at the University of Chicago and BYU for their helpful comments. For part of the duration of this project, my research was supported in part by an unrestricted grant from Amazon Inc. that funded of a portion of my fellowship with BYU Law School, and I have a have no conflicts of interest.

1 Introduction

Two-sided platforms help buyers and sellers transact by both providing a marketplace and accompanying infrastructure to successfully match users. These platforms earn profits by charging fees to buyers and sellers, and therefore seek to both attract users and ensure that they are generating revenue. However, these firms play a limited role in the decision problem of sellers that use their site: individual users choose which items to sell, as well as how to price and promote their listings within the platform's interface. While many platforms have a decentralized approach to on-site transactions, they often provide information to sellers to help them track and optimize their business (e.g., eBay Seller Hub, Amazon Seller Central, AirBnB Smart Pricing). This information—which can range from transaction data to fully automated tools—directly affects sellers' strategic decisions, since sellers may be uninformed or uncertain about how to maximize profits in a new setting.

Despite the prevalence of large platforms and the information services they provide, however, standard models of platform design ignore the use of information services by platforms. Both buyers and sellers are often assumed to have perfect information about how the platform functions, which they use to respond rationally to the fees chosen by the platform (Rochet and Tirole 2003; Klein et al. 2005). However, recent empirical work shows potential gains to information in platform settings when it is used for more targeted advertising (Mela, Roos, and Sousa 2023), personalized pricing (Wu, Huang, and Li 2023), or suggested pricing (Fong, Manchanda, and Song 2023). Ensuring that sellers have better information affects their behavior and can increase the total transacted revenue on the platform, allowing the platform to earn more commissions. Thus, a profit-maximizing platform must consider how to jointly choose the information and fees for its users.

In this paper, I study how information provision to new sellers by an auction platform determines the optimal fee structure and user welfare. I extend a two-sided auction platform model from Marra (2019) to a setting where sellers learn from past transactions whether to list items for sale and what reserve prices to set. Since new sellers update their beliefs with every item they list for auction, their beliefs—and consequently, their entry and reserve price decisions—vary with the information they observe. I apply this model to a large dataset of eBay auctions and find that new sellers are initially overoptimistic about bidder entry into their auctions, which leads to higher

reserve prices and more entry. When the platform updates new sellers' beliefs toward the truth, however, it can increase bidder entry as well as platform and seller profits.

The question of optimal information provision is central to the platform's strategic problem. Within the auction platform, sellers face a problem similar to a standard monopoly pricing problem (Bulow and Klemperer 1996). In this context, the demand faced by each individual seller is determined by both how many bidders enter each auction and how much those who enter value the item being sold; sellers must know these features to maximize profits. If sellers have incomplete information about their demand curve, the platform may wish to correct sellers' beliefs and benefit from their more informed choices. However, the platform may benefit from information asymmetries among its users if these increase either transaction volume or users' willingness to pay higher fees.[1]. Thus, sellers' beliefs about the arrival process—and the platform's role in influencing these beliefs—shape outcomes for sellers, bidders, and the platform itself.

This analysis is motivated by new evidence that new sellers learn to act optimally in both their choice of reserve prices and decision to list items for auction, suggesting that they may be influenced by new information about the auction process. I use rich item-level data from one million eBay auctions for children's toys to show that new sellers set lower reserve prices and earn higher revenues as they gain more experience. These new sellers' reserve prices are also more strongly correlated with past revenue than more experienced sellers' are, which is consistent with stronger early responses to new information. Reduced-form evidence alone, however, cannot isolate how information affects sellers' beliefs and behavior since sellers' private values also determine their decisions. Nor can this evidence reveal how much information is optimal for the *platform* to give to sellers or how information affects optimal fees.

To understand how new information shapes seller beliefs, and in turn determines sellers' reserve price and participation decisions within the platform, I develop a model of selective entry on an auction platform with seller learning. In equilibrium, both bidders and sellers participate in auctions if they believe it is profitable to do so, since each faces a fixed cost of either bidding on or listing an item. Since high reserve prices reduce the amount of expected surplus from an auction, I show

[1]That asymmetries in platform design may be optimal is not a new concept: both theoretical (Rochet and Tirole 2003; Klein et al. 2005) and empirical (Gomes 2014; Marra 2019) work shows that different fees for two sides of a market may act as a form of cross-subsidization to induce optimal entry.

that bidders are less likely to enter auctions with high reserve prices.[2] Sellers also decide to list an item for auction if their expected surplus from that auction is positive. However, new sellers do not fully understand the bidder entry process, and instead learn about bidder entry—and its effect on their expected profit, if they list an item for auction—through repeated transactions. In contrast, I show that a fully-informed seller chooses a reserve price to both attract and extract surplus from bidders, a generalization of the Myerson (1981) reserve price rule.

I then estimate the model to quantify how bidders and sellers—particularly new, uninformed sellers—act on the platform and how receiving new information shapes sellers' beliefs. I use rich text descriptions for each item to flexibly model heterogeneity in item values with a neural network, and derive a likelihood that corrects for selection in observed bids on eBay. To avoid potential bias from this method due to the high dimensionality of the text data, I derive an orthogonal score from which I obtain consistent estimates of the parameters of interest (Farrell, Liang, and Misra 2020; Chernozhukov et al. 2022; Ichimura and Newey 2022). I then show that seller beliefs about bidders' entry process are semiparametrically identified from the distribution of reserve prices, using results from the literature on identification of random coefficients models (Fox et al. 2012), and estimate the path of new sellers' beliefs as they transact on the platform. Since sellers' beliefs are very high dimensional due to the lack of a conjugate prior structure, I make this problem computationally tractable by approximating Bayesian learning via sequential update steps that I implement with neural networks.

I find that new sellers underestimate the effect of reserve prices on bidder entry, leading them to initially set too-high reserve prices and enter more than they should. Higher reserve prices have a strong negative effect on the expected number of bidders because bidders face a time cost of entering auctions. New sellers initially underestimate this effect, which causes them to be overoptimistic about the number of expected bidders. This means that new sellers are willing to set higher reserve prices, expecting to earn a higher revenue, which in turn causes more entry from overoptimistic sellers. However, new sellers' priors are relatively uninformative, so they rely heavily on new information to inform their beliefs. This allows many sellers to quickly learn to set prices that are broadly consistent with accurate beliefs about the bidder arrival process.

[2]This effect is distinct from the mechanical effect of a high reserve price in an online auction, which excludes potential bidders from submitting lower bids even if they enter the auction.

These results have important strategic implications for the platform and its decision of how much information to provide. The platform can choose to provide new sellers with data from past auctions, from which new sellers can learn about the true bidder entry process. At the same time, the platform chooses what fees it will charge to both new and experienced sellers, who are already familiar with the bidder arrival process and set prices optimally. My estimates quantify how much and in what direction new sellers will adjust their beliefs about the bidder entry process in response to new information. When provided with a random sample of other auctions on the platform, the average new seller is able to update their beliefs toward the true parameters without incurring the time and monetary cost of starting to use the platform. This allows the platform to change its fee structure to increase platform profits from the baseline, though it may be optimal for the platform to further increase profits by exploiting new sellers' information gap.

Empirically estimating new sellers' beliefs about the bidder entry process is critical, as a platform's optimal information structure may be ambiguous without empirical evidence. First, while platforms may have increasing economies of scale in analyzing data, it is still costly to provide users with additional information. This may be exacerbated by users' low willingness to pay for additional information: sellers who do not anticipate changing their beliefs will see little value in purchasing additional data. Thus, despite the cost, platforms may prefer to offer such information for little to no fee, if they offer it at all. Second, it is possible for the platform to help one side of the market without harming the other: correcting sellers' beliefs yields improved profits for sellers while inducing more bidder entry. The effects of this may be magnified by the two-sidedness of online platforms: increasing the surplus of one side of the market may lead to more entry on both sides. Finally, since the users' decision problem may change with new information, the optimal fee structure on a two-sided platform (as studied under perfect information in e.g. Rochet and Tirole 2003 and Klein et al. 2005) may also change.

1.1 Related literature

This paper contributes to the large and growing literature on agent learning. Much of this literature focuses on learning-by-doing and the extent to which more experienced agents are able to leverage information to improve outcomes (Simonsohn 2010; Haggag, McManus, and Paci 2017; Strulov-

Shlain 2021; Tadelis et al. 2023). In particular, Huang, Ellickson, and Lovett (2020) shows how firms in a new market respond to market signals and set prices accordingly; I document similar behavior among new sellers on an auction platform. Cho and Rust (2010) also documents that car rental firms behave suboptimally in choosing rental prices, and as in my setting, lowering prices (at least for some cars) can increase overall profits. Other work addresses the theory of optimal behavior under uncertainty, especially in games with updating (Rothschild 1974; Keller and Rady 1999; Hitsch 2006; Doraszelski, Lewis, and Pakes 2018). I contribute to the literature on estimating agent beliefs and learning process, as in Erdem and Keane (1996) and Kim (2020); I also provide semiparametric identification results for beliefs under learning as in Lu (2019) and Wang and Yang (2024).

I also address the field of auction design and optimal pricing in auctions. Auctions are a particularly well-suited setting in which to study information design in platform markets due to the rich literature on optimal bidder and seller behavior, including empirical tools to test and quantify theoretical results. In particular, I adapt the endogenous auction platform model of Marra (2019), which relates to other models of auctions with endogenous entry such as Levin and Smith (1996), to a setting with seller uncertainty. My model also develops an insight from Engelbrecht-Wiggans (1987) (that optimal reserve prices should account for their effect on bidder arrival) into a new reserve price condition that nests that of Myerson (1981); this reserve price also shows the tradeoff between the benefit of attracting another bidder (as in Bulow and Klemperer 1996) and expected surplus extraction when arrival is both endogenous and stochastic. Existing literature further shows that auction participants may act suboptimally, whether in laboratory settings or as large, sophisticated firms (Davis, Katok, and Kwasnica 2011; Ostrovsky and Schwarz 2016); my application examines small firms in a real marketplace. I further contribute the literature on empirical estimation of auctions by incorporating unstructured text data, similar to Netzer et al. (2012), Bajari et al. (2023), and Compiani, Morozov, and Seiler (2023).

This paper also relates to the growing literature on two-sided markets. Existing work has studied the question of cross-subsidization in optimal platform design both theoretically and empirically (Rochet and Tirole 2003; Klein et al. 2005; Gomes 2014; Jullien, Pavan, and Rysman 2021; Marra 2019). Other work also examines the role of information provision and learning in online platforms

(Mela, Roos, and Sousa 2023; Wu, Huang, and Li 2023; Foroughifar 2023). In particular, Fong, Manchanda, and Song (2023) finds that new sellers on a used goods platform set higher prices than experienced sellers, and evaluates in an experimental setting the extent to which providing a suggested price to sellers shifts their pricing strategies. I consider the fee and information provision problems jointly to understand the interaction between them, and in particular allow sellers to incorporate their private valuations in their pricing decisions. More broadly, studies of search and recommendation systems also examine how user behavior is influenced by non-price mechanisms, though these generally focus on buyers (Bronnenberg, Kim, and Mela 2016; Compiani et al. 2022; Xu, Deng, and Mela 2022; Hodgson and Lewis 2023).

1.2 Roadmap

The rest of this paper proceeds as follows. Section 2 presents the problem of optimal platform design. Section 3 describes the eBay data used in this analysis and presents descriptive evidence that new sellers learn to set reserve prices. Section 4, combines a model of two-sided selective entry on an auction platform with a model of seller learning. Section 5 presents the identification and estimation strategies and estimates the structural model. Section 6 revisits the platform's optimal information and fee structure in light of the model estimates, and section 7 concludes.

2 A simple model of platform fees and information provision

To fix ideas, I present a simple model of an auction platform that can choose both fees and what information to provide to potential sellers. This stylized model highlights the important potential tradeoffs (or complementarities) between fees and information, which motivates a more detailed empirical study of these forces but cannot itself determine the platform's optimal decision. I assume sellers have potentially-biased beliefs about how to act optimally on the platform, and that receiving additional information makes sellers' beliefs and actions monotonically approach those that would hold under the full-information benchmark.

I consider three main decision variables that can be chosen by the platform: the amount of

information a to provide to sellers, an insertion fee c^I for listing items, and a revenue fee c^R charged as a portion of successful transaction. Information provision $a \in \mathbb{R}_+$ is an index of how much the platform can shift sellers' beliefs, where $a = 0$ represents no information provision (i.e., sellers maintain their initial beliefs) and $a = \infty$ represents the limiting case of perfect information provision. For exposition, I also abstract from the two-sided nature of most platform fee structures and do not consider the incidence of either c^I or c^R on bidders or sellers.

The platform's profits depend on both whether sellers list items and the revenue they generate conditional on their participation. Both features are determined in equilibrium by bidder and seller behavior, so they are each functions of information a and fees c^I and c^R. I define the reduced-form objects $\mathcal{P}(a, c^I, c^R)$ as the probability that an item is listed and $\mathcal{R}(a, c^I, c^R)$ as the expected revenue from listed items. The platform's profit maximization problem is therefore

$$\max_{a,\, c^I,\, c^R} \quad \mathcal{P}(a, c^I, c^R) \cdot [c^I + c^R \cdot \mathcal{R}(a, c^I, c^R)]$$

For simplicity, I assume the platform faces zero marginal cost for facilitating auctions and for providing information, and that all other fixed costs are sunk.

Providing information affects seller behavior on both the extensive margin of participation and the intensive margin of expected revenue. The optimal platform design is characterized in part by the following first-order condition with respect to the amount of information provision:

$$0 = \underbrace{\mathcal{P}_a(a, c^I, c^R) \cdot [c^I + c^R \cdot \mathcal{R}(a, c^I, c^R)]}_{\text{Extensive margin}} + \underbrace{\mathcal{P}(a, c^I, c^R) \cdot c_S^R \cdot \mathcal{R}_a(a, c^I, c^R)}_{\text{Intensive margin}}$$

The signs of both $\mathcal{P}_a$ and $\mathcal{R}_a$ depend on sellers' initial beliefs about optimal participation and reserve prices, and have important implications for the platform's choice of how much information to provide to new sellers. I consider several cases informally here.

(i) *Overly pessimistic sellers.* Sellers may underestimate their expected profit per auction, listing fewer items and setting high reserve prices to avoid low sale prices. Here $\mathcal{P}_a > 0$ and $\mathcal{R}_a > 0$, and the platform optimally chooses to share information with sellers to increase the number and profitability of transactions.

(ii) *Overly optimistic sellers.* Sellers' optimism may cause them to list many items and set very low reserve prices that they do not expect to bind. Here $\mathcal{P}_a < 0$ and $\mathcal{R}_a < 0$, meaning that the platform may wish to take advantage of sellers' biased beliefs instead of providing them with helpful information.

(iii) *Uncertainty about the optimal reserve price.* Sellers may choose to set reserve prices to be either too high or too low. Here the signs of $\mathcal{P}_a$ and $\mathcal{R}_a$ may differ, so the platform must improve either the quantity of or revenue from transactions at the expense of the other.

While the details of the platform problem are important, these cases illustrate the potential factors that may be considered when determining optimal information provision.

Information provision also affects platform outcomes by determining how sellers respond to changes in fees. Consider as an example the case where $c^R = 0$, so the platform only earns revenue from charging insertion fees. Rearranging the platform's first-order condition with respect to insertion fees characterizes the optimal insertion fee as

$$c^I = \frac{-\mathcal{P}(a, c^I, 0)}{\mathcal{P}_{c^I}(a, c^I, 0)}$$

Though in general no closed-form solution exists for c^I, this equation provides some intuition for the platform's problem. Sellers who are likely to list items for auction—perhaps due to overoptimistic beliefs about their expected surplus—are willing to pay higher insertion fees since they believe they will recoup their loss. Alternatively, sellers who believe listing items is unprofitable must be enticed to participate with lower insertion fees. The effect of changing sellers' initial beliefs through information provision depends on both what these initial beliefs are and how they affect sellers' behavior.

Ultimately, this simple model cannot solve the joint problem of determining information provision and platform fees. Both platform participation and expected revenue per auction are determined by bidders' and sellers' equilibrium behavior. Thus, taking this problem seriously requires a more careful look "under the hood" of the auction platform equilibrium. With this in mind, I turn to the setting of interest to highlight key features and trends in new seller behavior.

3 Setting and auction data

I use data from eBay, a well-known auction platform, to study how new sellers act and respond to new information. This data is a subset of the eBay auctions used in Resnick and Zeckhauser (2002), and spans from January to June 1999. During this time, eBay was relatively new (having been started in late 1995) and auctions were the only mechanism used to sell items, making this an ideal dataset in which to study how new sellers learn to operate in an unfamiliar setting. This data also preceeds the introduction of eBay's data analytics service "Seller Hub" in 2016.

To fix ideas, I present a simplified outline of the eBay auction process. Sellers choose to list an item for auction, and choose a starting minimum bid and (if desired) a secret reserve price along with an item description. Prospective bidders can find listed items on a search page, including the current minimum bid, and then choose to click into the item page. Bidders may then observe additional item details along with seller information and an indicator for whether the secret reserve price (if any exists) has been met. Bidders submit their bids to eBay, which proceeds as a second-price ascending auction where the current minimum bid is the second-highest of existing bids and the initial minimum bid. Examples of the search and item pages are presented in Figure 1.

Figure 1: Example of eBay search and item pages



Notes: This figure was retrieved from the Wayback machine, and has been edited to conserve space. The search and item pages are from 2001 and 2005, respectively. The sellers' ability to feature their item can be seen here - featured items are listed at the top of the page, while other items are highlighted and/or have bold titles. The seller's total feedback score and positive feedback rating are visible, as is the current bid, starting bid, and and indicator for the reserve not being met by the current bid.

3.1 Auction listing and bid data

Due to the prevalence of antique and custom items on eBay, I restrict attention to one of the more popular categories: a brand of stuffed animals called Beanie Babies (BBs). There are approximately 1 million BB auctions in the dataset, with about 2.7 million bids. All distinct varieties of BBs were produced by a single company (TY), which ensures some level of homogeneity among the listed items. Table A1.1 presents various summary statistics for the analysis dataset.

Several features make this dataset attractive for empirical analysis. First, both the highest bids and secret reserve prices are recorded for each item in the dataset. Online auction datasets frequently impute secret reserve prices from observable indicators such as the "reserve not met" sign in Figure 1; since my focus is on sellers' choice of reserve price, it is helpful to obtain accurate measurements of this choice variable. Highest bids are similarly unobserved in many studies of online auctions due to the ascending minimum bid only depending on the second-highest bid. As will be shown later, the first-highest bid being observable is helpful for identifying and estimating seller beliefs. Other data such as the time of the auction, any promotional choices made by the seller, and seller-provided item descriptions are also included in the dataset.

Additionally, supplemental data on users' feedback reveals all positive, negative, and neutral ratings for accounts. Importantly, this feedback history dates back to the beginning of eBay in 1995; this allows me to construct the reputation variables that are visible to prospective bidders, specifically feedback scores (defined as the number of all feedback events) and rating (defined as the percent of feedback events which are positive). Unless otherwise noted, I use the inverse hyperbolic sine of feedback scores to address the significant skew in user feedback. While feedback scores are not a perfect measurement of the number of auctions in which sellers have participated, they are often used as a proxy for user experience on eBay (see e.g. Simonsohn 2010).

3.2 Evidence for learning among new sellers

I now present evidence that new sellers are learning to run auctions as they gain more experience, and therefore can be influenced by observing information. I define new sellers as all accounts who have no recorded feedback before the start of the dataset and who list at least one item for sale.

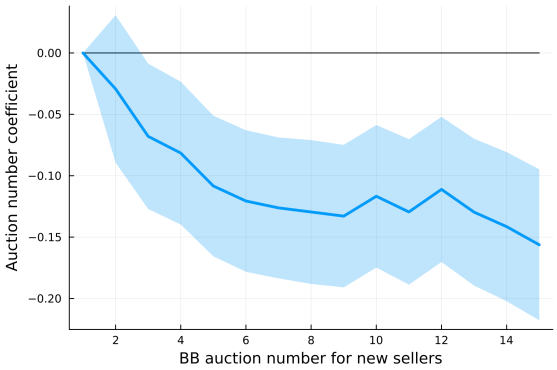
Similar to Kim (2020), I examine trends in new sellers' choices and outcomes as a reduced-form test for whether new sellers' behavior varies with experience. I also document how sellers' choice of reserve price becomes less strongly correlated with lagged revenues as they gain more experience, and that seller choices and outcomes are related to how many items they list.

I first show that as new sellers gain more experience, they change in both their choice of reserve prices and the outcomes they face. I plot time trends in seller revenue net of fees, the *effective reserve price* (defined as the maximum of the secret reserve price—if one exists—and the starting minimum bid), and the number of bidders per auction, while controlling for predicted item values, seller feedback, seller rating, and seller and month fixed effects. Throughout, I will homogenize variables such as reserve prices and revenues by dividing by the predicted item value. The process for estimating item values is detailed in section 5.1; additional plots in Appendix A1 replicate these trends without depending on the predicted values.

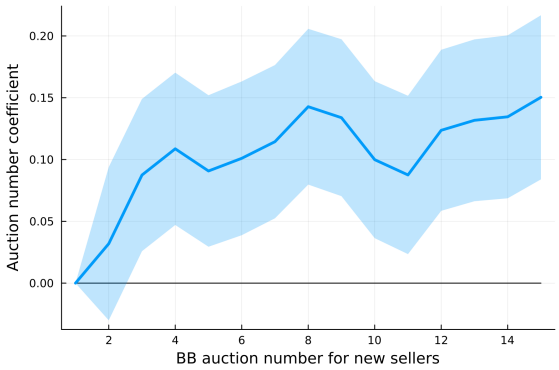
Figure 2 plots time trends when restricting the sample of new sellers to those with at least 15 auctions in the data.[3] These new sellers initially set higher prices, earn lower revenues, and attract fewer bidders than they do in later auctions. This pattern is consistent with seller learning when initial beliefs are biased toward setting higher reserve prices. Further, by examining only those who remain on the platform for at least 15 auctions and controlling for persistent seller heterogeneity, these trends do not simply reflect early exit by high-value sellers.

[3]I chose 15 auctions to avoid including sellers who may have few items in their possession and no interest in long-term trading, as well as to not have too short a panel for estimating seller fixed effects. The trends are similar for different windows.

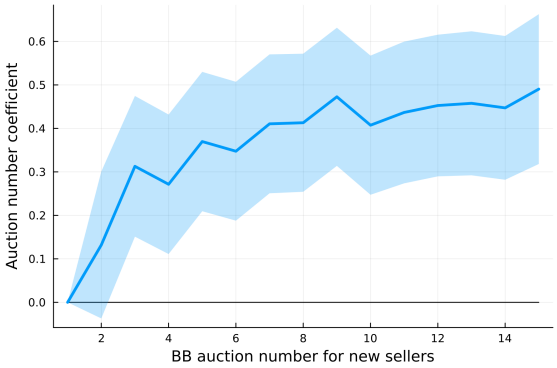
Figure 2: Trends in variables of interest as new sellers gain experience



(a) Effective reserve (homogenized)



(b) Net revenue (homogenized)

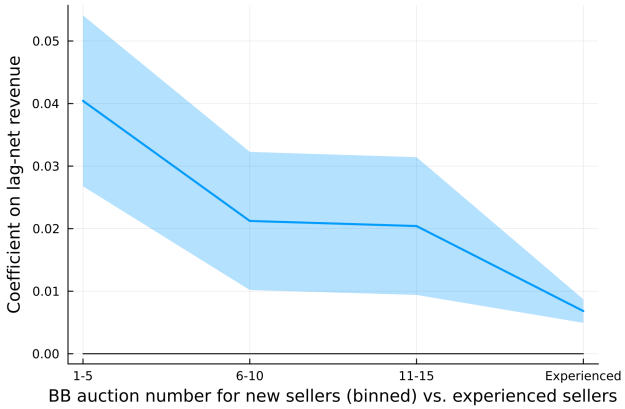


(c) Number of observed bidders

Notes: These regressions pool 1,639 new sellers' first 15 auctions with all auctions by 5,165 experienced sellers (defined as those with a feedback count of at least $\geq$47 at the start of the data, which is the 75th percentile of initial feedback count). The sample is limited to sellers with at least 15 auctions in the data; the results are similar when using different number of auctions. I plot 95% confidence intervals for the estimated coefficients.

Standard learning models also predict that the value of additional information decreases as sellers obtain more experience, and that seller beliefs converge toward toward the true parameter. To examine variability in seller choices over time, I regress current-auction reserve prices on lagged net revenue and lagged reserve prices (again conditioning on sellers with at least 15 auctions in the data). Figure 3 plots the coefficients of lagged net revenue from this regression, binned by the auction number among new sellers. New sellers' current prices are correlated with past revenue signals, and the magnitude of the lag coefficient is larger in early auctions. This is consistent with subsequent auctions containing relatively less information for more experienced sellers.

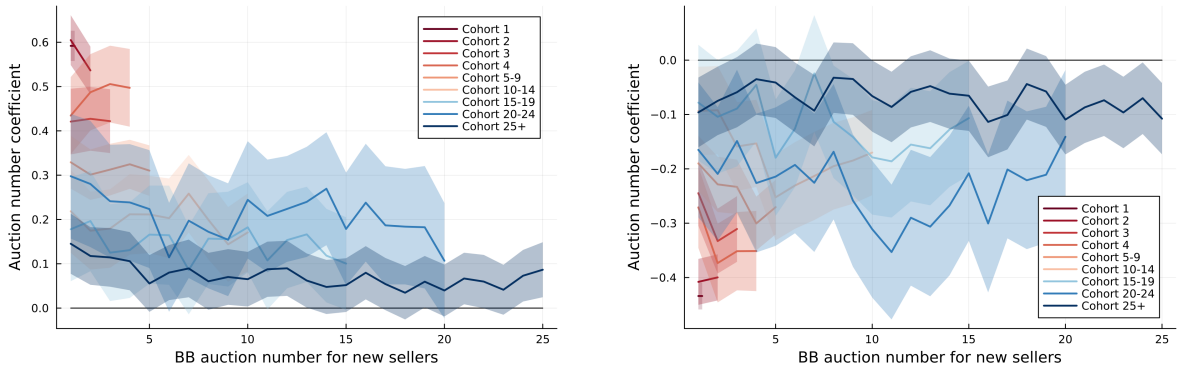
Figure 3: Coefficients from regressing reserve prices on lagged revenues, by binned seller experience



Notes: These are the coefficients when regressing current-period effective reserve price on lagged revenue, multiplied by indicator functions for new sellers being in the first 1-5, 6-10, and 11-15 auctions in the data (among new sellers with at least 15 auctions in the data and experienced sellers with >75th percentile of experience at the start of the data); these are compared with the coefficient of lagged revenue interacted with an indicator function for experienced sellers. The controls include month fixed effects, feedback percentage, predicted item value, and interactions of each experience indicator function with lagged effective reserve price. I plot 95% confidence intervals for the estimated coefficients.

Similar trends hold when examining new sellers more broadly, though the estimates are far noisier. To examine the effect of selection on new seller outcomes, I separate new sellers into cohorts based on the number of items they list. In Figure 4, I plot the average difference between new sellers' and experienced sellers' reserve prices and net revenue for each of new bidders' first k auctions. Consistent with selection, new sellers who list relatively few items are also those with significantly higher reserve prices and lower revenues.

Figure 4: New seller trends by cohorts of the number of items listed (difference from experienced seller averages)



(a) Effective reserve (homogenized)

(b) Net revenue (homogenized)

Notes: These figures represent a simple difference in means between new sellers in each cohort and experienced sellers (defined as sellers with >75th percentile of experience at the start of the data). This does not include any controls other than month fixed effects. I plot 95% confidence intervals for the estimated coefficients.

Additional results in Appendix A1 highlight other patterns that are suggestive of sellers learning to set reserve prices. While sellers exhibit trends in some non-reserve price choice varibles (such as featuring items and timing of auctions), the magnitude of these trends are generally small, which motivates sellers' reserve pricing decision as the focus of my analysis. I also show that, while there are some trends in item descriptions, some of the most prominent focus on the item reserve price (in particular, noting the lack of a secret reserve price).

3.3 Alternative explanations for new seller behavior

I also consider mechanisms that might drive the trends shown above other than sellers responding to information about auctions. These other mechanisms include behaving strategically to sell higher-value items first, facing higher-value bidders for whom higher reserve price may be optimal, having systematically higher private values for all items when entering the platform, and having an initial endowment effect that decreases with selling experience. I discuss each in turn, considering either theoretical explanations or empirical tests for each possible mechanism.

First, sellers might strategically choose the order of their listings when starting their account, perhaps starting with items they value more highly conditional on an item's book value. However, this strategy runs counter to any dynamic considerations like those discussed in Foster, Haltiwanger, and Syverson (2016). Upon entering the platform, sellers could instead list items they value less (and set a correspondingly lower reserve price) and yield higher sale probabilities. This would allow sellers to increase their reputation scores and potentially earn more in subsequent auctions. Thus, to the extent that items are listed in decreasing order of sellers' private value, this is unlikely to be done to optimize dynamic profits from listing auctions.

Sellers may face a different set of bidders in their initial auctions, for whom it may be optimal to set higher prices than for later auctions. It is difficult to directly analyze the distribution of bidder values without a model, since the increasing minimum bid and endogenous entry of bidders create a selection problem. To overcome this, I restrict attention to the first and second highest bids (where they exist) and control for the number of observed bids with fixed effects. Figure A1.4(a) shows generally flat trends in the first and second highest bids in new sellers' auctions, with a slight

but statistically insignificant increase for later auctions. This suggests that initially higher reserve prices are not driven by differentially higher bids.

Additionally, sellers may have systematically higher values for all items upon entering the platform. In this dataset, I can observe sellers whenever they bid for other items and test whether their bids change as they gain more experience on the platform. Figure A1.4(b) shows the bids of new sellers for other listings, and illustrates that new bidders do not place higher (or lower) bids for other items upon entering the platform. This is consistent with the underlying distribution of seller valuations remaining constant throughout new sellers' first auctions on the platform.

Finally, new sellers may have an endowment effect that diminishes as they gain additional experience. This story is consistent with several experimental studies in which participating in more transactions decreases the endowment effect (List 2003, 2004, 2011; Tong et al. 2016). In this observational data, I do not directly observe both measures of willingness to pay and willingness to accept for the same item; further, reserve prices may not reflect a true willingness to accept because sellers try to extract surplus from bidders. However, Figure A1.5 shows the distributions of reserve prices for experienced and inexperienced sellers are quite similar for the approximately 50% of listings with lower standardized reserve prices. This is suggestive evidence against new sellers having a strictly higher willingness to accept for all their items. Figure A1.6 also shows that new sellers' current reserve prices are more strongly (and negatively) correlated with failing to sell the previous item, while selling the previous item is uncorrelated with current reserve prices. Given evidence from Tong et al. (2016) that the endowment effect dissipates with successful transactions, this pattern does not seem to be driven by a diminishing endowment effect.

4 Auction platform model with seller learning

To directly study the extent to which new sellers are influenced by new information, I develop an auction platform model that allows for seller learning. The model adapts the two-sided endogenous entry model from Marra (2019) by allowing seller actions to vary with their individual beliefs about their profit function. I begin by introducing the following notation and assumptions.

A large number $\mathbf{N}_S$ of sellers and $\mathbf{N}_B$ of bidders can choose to participate on a monopoly auction platform. Each potential bidder i has valuation v_{ij} for each item j, where $v_{ij} \sim F_B$. Potential sellers s have outside option values v_{0sj} for each item j they possess, where $v_{0sj} \sim F_S$. Each item has auction-level observables X_j. I assume all prospective sellers and bidders know the valuation distributions F_B and F_S, which satisfy the following assumption.

> *Assumption* 1. The value distributions F_B and F_S are absolutely continuous and have connected support. Bidder values v_{ij} are independent from $v_{i'j}$ for all $i \neq i' \in \{1, ..., \mathbf{N}_B\}$, and seller values $v_{0\ell j}$ are independent from v_{ij} for all $i \in \{1, ..., \mathbf{N}_B\}$ and $\ell \in \{1, ..., \mathbf{N}_S\}$.[4] Further, the bidder value distribution F_B satisfies the strict monotone hazard rate property (i.e., $\frac{f_B(x)}{1-F_B(x)}$ is strictly increasing in x on the support of F_B). Finally, dependence on item j's characteristics takes the form $v \cdot \exp(\gamma(X_j))$, where X_j is exogenous to all values v, which are drawn from the players' respective distributions (F_B or F_S).

Throughout the discussion of the model, I focus on homogenized item values, which is equivalent to setting $\gamma(X_{jt}) = 0$ for all items. Since each item j is sold by a single seller s, I omit dependence of values and other terms on j and s where possible.

I assume all sellers know their values v_0 for the item they own. Sellers can choose to list the item for auction after incurring an item-specific entry cost $c_S^E \overset{\text{i.i.d.}}{\sim} F_{c_S^E}$, which is independent from sellers' private values v_0. After deciding to list the item, sellers also choose an effective reserve price r and minimum bid m. I assume m is exogenously drawn between 0 and r, and unlike Marra (2019), I assume that r is observable to all prospective bidders.[5] Potential bidders can see the item on a listing page and decide whether or not to enter the auction; if they do so, they incur entry cost c_B^E and only then learn their value v_i for the item and costlessly submit a bid. The bidder

[4] Quint (2017) explores the relationship between reserve prices and outcomes in private and common values settings, and shows that in settings with a common value component the reserve price may be lower than under pure private values. I examine the assumption of independent private values in more detail in Appendix A2. While this setting (as with many auction settings) likely has both a common value and private value component, I simplify the bidding model to allow more tractability in the higher-level problem of interest, i.e. the platform's decision of how much information to provide to shift seller behavior. In particular, this facilitates a more flexible specification of mean item utilities $\gamma(X_j)$ in estimation (see section 5 for more details).

[5] In this dataset, only 23.9% of auctions have a secret reserve price, while 89% have an minimum bid higher than the eBay default of \$1. Thus, this simplifying assumption is reasonable in the present setting. The "reserve not met" indicator disappears when at least one bid passes the reserve price, at which point the minimum bid has often increased from its starting value. Further, Katkar and Reiley (2007) documents that the "reserve not met" indicator still affects bidder entry even when the reserve price is still secret. Higher reserve prices are less likely to be met, which on averages makes this indicator last for a longer period. A broader literature studies the choice between public and secret reserve prices (Hasker and Sickles 2010).

with the highest bid exceeding the reserve price wins the item, and the transaction price p is equal to the highest of the effective reserve price and the second highest bid.

Bidder and seller behavior is also affected by the cost structure of using the platform, denoted by a vector c of all associated costs and fees. In particular, entry costs c_B^E and c_S^E for both bidders and sellers are decomposed into time costs c_B^T and c_S^T and insertion fees c_B^I and c_S^I, respectively. The insertion fees are paid directly to the platform when a seller lists an item or when a bidder enters an auction, regardless of whether the item is sold. The platform can also impose bidder and seller fees c_B^R and c_S^R. If the item is sold, the highest bidder pays $(1 + c_B^R)p$ and the seller receives $(1 - c_S^R)p$, so the platform also receives revenue $(c_B^R + c_S^R)p$ from each successful sale.

The rest of the section is divided into three parts. I first describe bidders' optimal strategy, conditioning on seller behavior. I then examine the seller strategies and how they depend on sellers' beliefs about bidder behavior. Finally, I review the conditions on both bidder and seller behavior that must hold simultaneously in equilibrium.

4.1 Bidder strategies

I first focus on the the bidding strategy of actual bidders who enter the auction; these bidders learn their private valuation for each item and then submit bids. This implies an expected continuation value of entering in an auction, which determines bidders' decision of whether to enter an auction.

4.1.i Bidding stage

After entering an auction and learning their private valuation for the listed item, each bidder's optimal strategy is to bid if they are able to do so. All $\tilde{N}$ bidders who have entered the auction face no cost to submitting their bid, but may be prevented from doing so by the minimum bid m. Following Vickrey (1961) and Marra (2019), all bidders with private value v_i will submit bids $\frac{v_i}{1+c_B^R}$ as long as their bid exceeds the current value of m. Any bidder with a sufficiently low private valuation ($v_i < (1 + c_B^R)m$) will not bid at all, since bidding a higher amount would yield negative expected utility.

I abstract from any potential learning and uncertainty by bidders about how to submit bids. This is because eBay provides an automatic bidding tool that increments the minimum bid up to the maximum value a bidder reveals they are willing to pay. Thus, eBay already implements the optimal bidding rule for all bidders via algorithm. This also allows for a more tractable model of bidder behavior, both for the researcher and for the sellers' mental model of bidder behavior.

4.1.ii Bidder entry stage

Potential bidders determine their expected surplus from entering any given auction, and they enter if there is positive expected surplus. In this setting, in contrast to Marra (2019), I assume the reserve price is public. Thus, bidders make their entry decision based in part on the fee-adjusted reserve price $r^B \equiv (1 + c_B^R)r^*$, which is the winning bidder's payment if the reserve price is binding.

For tractability, I assume that the number of bidders per auction is drawn from a Poisson distribution. This is formalized in the following assumption.

Assumption 2. Assume the expected number of bidders per auction is written as

$$\Lambda(r \mid \delta_0) = \exp(\delta_{0,1} + \delta_{0,2}\rho(r)) \tag{1}$$

for some vector δ_0. Further, the number of bidders per auction is Poisson with mean $\Lambda(r \mid \delta_0)$, yielding the following probability mass for n bidders entering an auction with reserve price r

$$p_n(r \mid \delta_0) = \frac{\Lambda(r \mid \delta_0)^n \exp(-\Lambda(r \mid \delta_0))}{n!} \tag{2}$$

The Poisson mean Λ, parameterized by the vector δ_0, characterizes bidder arrival for any auction as a function of its reserve price. Note that equation (1) is quite general: without any restrictions on ρ, this allows for any strictly positive Poisson mean, which fully characterizes the bidder arrival process. However, this parameterization will prove convenient in the following sections. While it is not conventional for the expected number of bidders in an auction to depend on the reserve price, I allow for this possibility and consider it formally below.

The equilibrium bidding strategy and the assumed bidder arrival process imply an expression

for a potential bidder's *ex ante* expected surplus from entering an auction. This is given by

$$\pi_B(r \mid \delta_0, c) = \sum_{n=1}^{\mathbf{N}_B - 1} \underbrace{\frac{1}{n}}_{\text{(i)}} \cdot \underbrace{\mathbb{E}\Big[v_{n:n} - (1 + c_B^R)\max\{v_{(n-1):n}, r^*\} \;\Big|\; v_{n:n} \geq r^B\Big]}_{\text{(ii)}} \cdot \underbrace{(1 - F_B(r^B)^n)}_{\text{(iii)}} \cdot \underbrace{p_n(r \mid \delta_0)}_{\text{(iv)}} \tag{3}$$

where $v_{k:n}$ is the kth highest out of n realizations of v_i. The four components of π_B are (i) the probability that any given bidder has the highest value, (ii) the expected surplus when the highest bidder wins, (iii) the probability that the highest bid exceeds the fee-adjusted reserve price, and (iv) the probability that n bidders enter at the auction.

Under the previous assumptions, the bidder entry game has a unique equilibrium in which bidder entry depends directly on each item's reserve price. The existence and uniqueness of equilibrium follows from Marra (2019), but the presence of a public reserve price yields a new result: the expected number of bidders decreases in the reserve price.

Proposition 1. Let assumptions 1 and 2 hold. Then the bidder entry equilibrium exists and is unique. Further, $\frac{\partial \Lambda(r|\delta_0)}{\partial r} < 0$. (Proof in Appendix A3)

The intuition for this result is as follows. Potential bidders have zero expected profit (net of their entry cost) from entering an auction, i.e.

$$0 = \pi_B(r \mid \delta_0, c) - c_B^E \tag{4}$$

since any positive profit will induce additional entry and negative profit will cause excess bidders to leave. The expected surplus in an any auction decreases in the number of competing bidders, since additional bidders both increase the expected price paid and lower the probability that a new bidder will win the item. Expected surplus is decreasing in the equilibrium number of expected bidders, Λ, and properties of the Poisson distribution imply that a unique value of Λ satisfies the expected zero profit condition for each reserve price r. Finally, expected bidder surplus declines with r, so fewer bidders enter when they expect more competition from the seller.

In light of Proposition 1, I revisit the functional form of the Poisson mean given in equation

(1). Specifically, the function ρ is now determined by the zero profit condition (4) for each possible reserve price r. Without loss of generality, I assume that ρ is a strictly increasing function. Proposition 1 therefore implies that the bidder arrival parameter corresponding to the reserve price, $\delta_{0,2}$, is strictly negative. That is, increasing the reserve price will reduce the expected number of bidders that enter the auction.

4.2 Seller strategies

I now examine sellers' equilibrium behavior. I first show how the platform's pricing structure and endogenous bidder entry shape sellers' choice of optimal reserve price when sellers face no uncertainty. I then extend this result to allow for seller uncertainty and learning, and conclude with sellers' entry problem. Throughout, I treat each seller's item-specific entry cost c_S^E as fixed for a given auction.

4.2.i Optimal reserve price on the platform under perfect information

I begin by writing the general form of the seller profit function to fix ideas. Sellers' expected profit, conditional on the bidder arrival parameter vector δ_0, is given by

$$\Pi(v_0, r \mid \delta_0, c) = (1 - c_S^P) \cdot \underbrace{R(r \mid \delta_0)}_{\mathbb{E}[\text{Revenue}|r,\delta_0]} + v_0 \cdot \underbrace{K(r \mid \delta_0)}_{\mathbb{P}[\text{Keep}|r,\delta_0]} - c_S^E$$

where the functional forms of the expected revenue R and and probability K that a seller keeps the item follow from the second-price auction literature and the Poisson arrival function.[6] The optimal interior reserve price r^* satisfies the first-order condition

$$\psi(r^* \mid \delta_0, c) \equiv \frac{-(1 - c_S^R) R_r(r^* \mid \delta_0)}{K_r(r^* \mid \delta_0)} = v_0$$

[6] Formally, these functions are written $R(r \mid \delta_0) \equiv \sum_{n=0}^{\mathbf{N}_B} p_n(r \mid \delta_0) R_n(r)$ and $K(r \mid \delta_0) \equiv \sum_{n=0}^{\mathbf{N}_B} p_n(r \mid \delta_0) K_n(r)$, where $r^B = (1 + c_B^R) r$ and

$$R_n(r) \equiv nr(1 - F_B(r^B)) F_B(r^B)^{n-1} + \frac{n(n-1)}{1 + c_B^R} \int_{r^B}^{\infty} z(1 - F_B(z)) F_B(z)^{n-2} f_B(z) dz$$

$$K_n(r) \equiv F_B(r^B)^n$$

where ψ is the virtual type function mapping bids to the space of seller values. Applying the Poisson assumption and the relevant functional forms, and then rearranging the terms of R_r and K_r, yields the following characterization of the optimal reserve price:

Proposition 2. Assume bidder arrival is Poisson with mean $\exp(\delta_{0,1} + \delta_{0,2}\rho(r))$. Then, recalling that the fee-adjusted reserve price faced by bidders is $r^B \equiv (1 + c_B^R)r^*$, the optimal interior reserve price satisfies

$$\frac{v_0}{1 - c_S^R} = \left[r - \frac{1 - F_B(r^B)}{(1 + c_B^R) \cdot f_B(r^B)} - \frac{W_R(r \mid \delta_0)}{f_{\max}^B(r \mid \delta_0)} \right] \frac{f_{\max}^B(r \mid \delta_0)}{f_{\max}^B(r \mid \delta_0) + W_K(r \mid \delta_0)} \tag{5}$$

where $f_{\max}^B(r \mid \delta_0) \equiv (1 + c_B^R) \cdot \sum_{n=1}^{\bar{N}_B} p_n(r \mid \delta_0) F_B(r^B)^{n-1} f_B(r^B) n$ is the scaled density of the highest bid at r given δ_0, and $W_R(r \mid \delta_0) \equiv \sum_{n=0}^{\mathbf{N}_B} (\frac{\partial p_n(r|\delta_0)}{\partial r}) R_n(r)$ and $W_K(r \mid \delta_0) \equiv \sum_{n=0}^{\mathbf{N}_B} (\frac{\partial p_n(r|\delta_0)}{\partial r}) K_n(r)$ are weighted averages of the expected revenue and keep probabilities for each n.

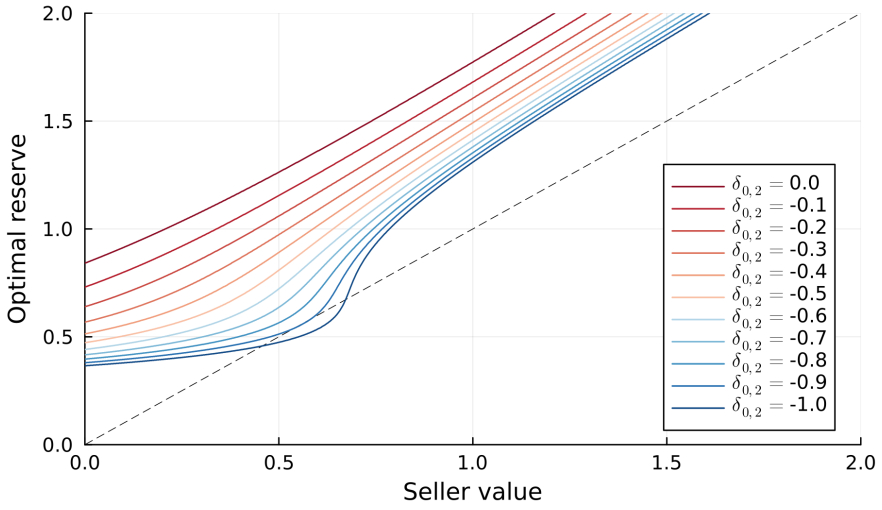
Note that when the bidder arrival process does not depend on the reserve price (that is, $\delta_{0,2} = 0$), it holds that $W_R = W_K = 0$. When the bidder and seller fees are also zero ($c_S^R = c_B^R = 0$), equation (5) reduces to the Myerson (1981) optimal reserve price formula.

To provide intuition for this pricing rule, Figure 5 plots the implied reserve price from Proposition 2 for different values of the bidder arrival coefficient $\delta_{0,2}$. The figure starts from the baseline case of $\delta_{0,2} = 0$ and shows increasingly negative values of $\delta_{0,2}$ in comparison. As $\delta_{0,2}$ decreases, sellers should optimally lower the markup in their reserve price. In this particular case, it may even be optimal to set a reserve price lower than the seller's own value if the bidder deterrence effect is sufficiently strong. The precise shape of the optimal reserve price function depends on the intercept $\delta_{0,1}$ of the log-mean of average bidders, as well as the shape of the bidder value distribution F_B and the function ρ that is determined by bidders' zero profit condition (4).

This equation generalizes several results from the related auction literature. As previously noted, it nests the Myerson (1981) reserve price formula in the case where $\delta_{0,2} = 0$ and carries with it the same intuition: sellers may set a higher reserve to extract additional surplus from bidders. Sellers' influence over the arrival process means they explicitly weigh the expected benefit of surplus extraction against the expected benefit of a potential additional bidder (though an extra bidder

would be better in expectation, as in Bulow and Klemperer 1996, this arrival is not guaranteed). Thus, this equation echoes the argument in Engelbrecht-Wiggans (1987), in that sellers may wish to lower the reserve to attract more bidders. Since there is some probability that few bidders enter each auction, however, the seller may still "protect" some expected surplus by setting a non-zero reserve price.

Figure 5: Optimal reserve price for varying reserve price coefficients $\delta_{0,2}$



Notes: These figures show optimal reserve price functions for varying parameters of $\delta_{0,2}$, keeping $\delta_{0,1} = 1.0$ fixed, where log-bidder values are normally distributed with mean 0 and variance 0.4, and $\rho(r) = r \cdot F_B(r)$. The dashed 45-degree line represents all values where the seller value is equal to the optimal reserve price.

4.2.ii Optimal reserve price on the platform under seller uncertainty

I now examine sellers' strategies when they do not know the true vector δ_0 that parameterizes the bidder arrival process. This uncertainty arises because δ_0 is a part of bidders' equilibrium play in the "search" stage rather than "item" stage of the game. Any uncertainty about or inattention to the platform's search algorithm, bidders' search strategies, or bidders' search costs could therefore contribute to uncertainty about sellers' own effect on prospective bidders' search process. For example, Simonsohn (2010) documents that eBay sellers may not understand the impact of competition on their own profits and over-enter when market activity is high. This competition neglect is related to seller behavior in this setting, as sellers may set reserve prices too aggressively (in essence, competing) for their own items. These sellers may fail to realize the extent to which this behavior crowds out potential bidders, each of whom may choose another way to spend their

time instead of entering an auction with an uncertain payoff.

The previous derivations can be extended straightforwardly in the case where sellers have some belief density b about the true value of δ_0. In an abuse of notation, I define subjective expected profit as $\Pi(v_0, r \mid b, c) \equiv \int \Pi(v_0, r \mid \delta, c) b(\delta) d\delta$. The respective profit function components $R(r \mid b)$ and $K(r \mid b)$ are defined similarly, implying the subjective virtual type function $\psi(\cdot \mid b, c)$ from rearranging the first-order condition of $\Pi(v_0, r \mid b, c)$ with respect to r. For tractability, I assume all sellers are myopic, so they only maximize current-period profits conditional on their beliefs b and do not actively experiment.[7]

As is standard in models of learning, sellers also have a model of the true data-generating process, from which they learn about the unknown parameter δ_0. I assume sellers update their beliefs after every auction they run. All sellers believe profit draws y are generated by the process

$$y = \Pi(v_0, r \mid \delta_0, c) + \epsilon \tag{6}$$

where ϵ is drawn i.i.d. from some distribution $F_{\epsilon|r,v_0}$ that is known to sellers. This randomness comes from the stochastic nature of bidders' arrival and their idiosyncratic valuations for each item. Each auction, sellers observe the associated data $\mathbf{D} = \{y, v_0, r\}$, and use the likelihood function implied by equation (6) and the noise distribution $F_{\epsilon|r,v_0}$. After each auction, sellers use a deterministic transition function $\mathcal{T}$ (such as Bayes' rule) to update their prior beliefs b to the posterior b':

$$b'(\delta) = \mathcal{T}(b(\delta), \mathbf{D}) \tag{7}$$

Thus, the evolution of sellers' reserve price strategies depends entirely on their beliefs b as updated after each auction, which in turn are driven by variation in data observed after each auction. I note that while sellers may be Bayesian, this is not necessary for the conclusions of the model.

[7] This assumption is used by Huang, Ellickson, and Lovett (2020) to simplify the analysis of firms learning from demand signals. This also relates to the "anticipated utility model" discussed in Cogley, Colacito, and Sargent (2007); in their setting, the model without active experimentation is a good approximation for the fully dynamic Bayesian model. While this assumption is not strictly necessary, it is useful to reduce the computational burden of estimating this model in the presence of other setting-specific computational challenges, and so I impose it from the outset in building the model; see section 5 for more details on the estimation procedure and its implementation.

4.2.iii Seller entry stage

Prospective sellers will enter the platform as long as their net expected surplus (relative to never listing the item) is positive. This yields the following inequality for sellers with beliefs b:

$$\Pi(v_0, r^* \mid b, c) - v_0 \geq 0 \tag{8}$$

Sellers' expected surplus depends crucially on their choice of reserve price conditional on their beliefs and their private value for the item sellers may list. The following proposition provides a condition under which the optimal reserve price r^* is increasing in sellers' private value v_0 for any seller with beliefs b.

> **Proposition 3.** Assume $K_r(\cdot \mid b) > 0$. Then r^* is increasing in v_0 and the virtual type function $\psi(\cdot \mid b, c)$ is increasing. (Proof in Appendix A3)

This monotonicity is common and easily verified under the functional forms assumed in much of the related literature; I make this explicit because it depends on sellers' underlying beliefs b and it is important for subsequent propositions.

Proposition 3 implies a threshold rule for sellers, such that sellers with a private value above that threshold will not list an item for auction. Given the functional form of R and the monotonicity of r^* in v_0, sellers' gains from trade $\Pi(v_0, r^* \mid b, c) - v_0$ are strictly decreasing in v_0. This implies the following characterization of the seller entry problem.

> **Proposition 4.** There exists a unique threshold $\bar{v}(b, c)$ for each belief density b such that sellers with beliefs b will list their item for auction if and only if $v_0 \leq \bar{v}(b, c)$.

That is, sellers will only select into the platform if their values are sufficiently low given their beliefs, in which case it appears profitable to list items for auction. This holds regardless of heterogeneity across seller beliefs, since each auction's reserve price is public knowledge and bidder entry is determined by their zero profit condition. Further, heterogeneity in beliefs does not directly impact the bidder arrival problem: bidders are only impacted by the reserve price r^* instead of the underlying values of v_0 and b.

This threshold condition is similar to other results in the literature, and corresponds to Marra

(2019) when beliefs b are common across all sellers and a point mass on the true parameters. I weaken the assumption that this threshold is objectively correct under the true item values: it needs only be optimal according to each seller's private beliefs about the true arrival process. This ensures that a unique equilibrium exists for the seller entry game even under heterogeneous and potentially biased beliefs among sellers.

4.3 Equilibrium definition

Before proceeding, I review the conditions for both bidders and sellers that must hold in equilibrium. Given a distribution F_B of bidder values, a distribution F_S of seller outside option values, a fixed bidder entry cost c_B^E, a distribution $F_{c_S^E}$ of seller entry costs, bidder and seller transaction fees c_B^R and c_S^R, initial ($t = 0$) prior beliefs b_0 about the parameters of bidders' entry process, and an updating rule $\mathcal{T}$ by which sellers update their beliefs as in (7), equilibrium consists of

(i) the threshold bidding rule as defined in 4.1.i

(ii) the bidder arrival parameter δ_0 that determines mean bidder arrival $\Lambda(r \mid \delta_0)$ as in (1)

(iii) the seller reserve price rule defined by $\psi(r^* \mid b, c) = v_0$ for any seller beliefs b

(iv) the seller entry threshold $\bar{v}(b, c)$ implied by (8)

such that actual bidders maximize expected profit from participating in an auction, potential bidders earn zero expected profit from entering any auction, sellers maximize expected profits given their beliefs about the bidder entry process, and the marginal seller earns zero expected profit given these beliefs.

The equilibrium conditions highlight the importance of seller beliefs and information on the auction platform. Sellers' beliefs determine both their entry decision and their choice of reserve price conditional on entry. Further, the degree of dispersion in seller beliefs affects how they update their beliefs when receiving new information. In turn, bidders' entry decisions and whether they submit bid depend on sellers' choice of reserve price. Thus, all actions on the platform are shaped by sellers' beliefs and the information they receive about bidder behavior.

5 Identification and estimation

The model from the previous section highlights two important ways in which sellers' beliefs and information affect platform outcomes. First, sellers' beliefs determine both their decision of whether to list an item. Second, sellers' beliefs determine how they choose a reserve price to attract bidders and extract bidder surplus. I now estimate the model to evaluate the role these features play in the empirical patterns, and how sellers' information affects their actions and platform outcomes.

Estimation proceeds in two parts. First, I estimate the demand side of the model, consisting of bidder arrival parameters, the distribution of bidder valuations F_B, and heterogeneity in mean item valuations. To do this, I derive a likelihood function for the demand side that corrects for selection in which bids are observed due to the increasing minimum bid rule. I also use results from the literature on debiased machine learning to flexibly estimate observable item heterogeneity. Using bidders' zero profit condition and the estimated parameters, I then construct the expected surplus per auction and recover bidders' time cost.

Second, I estimate the supply side of the model using data from both experienced and inexperienced sellers. Assuming experienced sellers have more accurate beliefs about bidder arrival, I use their first-order condition to recover seller values for all listed items. I then combine these estimated seller values with the seller entry condition to obtain the distributions of seller values and entry costs. Using these estimated distributions and the distribution of new seller reserve prices, I identify and estimate new sellers' prior beliefs about the bidder arrival process.

5.1 Demand side estimation strategy

I construct a likelihood to estimate the parameters of the demand model, which include both a high-dimensional component and a low-dimensional component. The high-dimensional parameter γ accounts for observable heterogeneity in item values and is a flexible function of the item text descriptions. The low-dimensional parameter vector ϑ_d consists of the bidder arrival parameters, as well as the distributions F_B of bidder values distribution and the distribution F_r of reserve prices. I first describe key components of the likelihood, including how to correct for bias due to the presence of high-dimensional parameters, and then present the likelihood itself.

5.1.i Heterogeneity in item values

The high-dimensional parameter γ accounts for heterogeneity in item values as a function of each item's text description. Specifically, the mean of bidders' log-valuations for item j is written as $\gamma(X_j)$ for item-specific data X_j. The data X_j is a large collection of indicator variables for each month in which an item may be listed and for whether the text description includes one of the 5,368 words that appear in at least 10 item descriptions (more details on the data construction and estimation process may be found in Appendix A4). This data, though unstructured, provides an extremely detailed view of item characteristics that are relevant to potential bidders.

I model γ as a neural network, which can flexibly account for item heterogeneity in the demand model but also introduce bias. This bias appears because the function of interest may be very complex relative to the amount of data, and it is possible to overfit the neural network. While model selection and regularization can be used to avoid overfitting γ, this also adds bias in the estimated item heterogeneity. This bias distorts the demand-side likelihood—shifting where in the parameter space the likelihood is maximized—which in turn biases the estimates of low-dimensional demand parameters ϑ_d. In order to consistently estimate ϑ_d, it is necessary to correct the estimation problem to make it insensitive to errors in the high-dimensional parameter γ.

I derive an orthogonal score to both flexibly estimate item heterogeneity and avoid the resulting bias in the structural parameters of interest. Since it is known how the high-dimensional parameter γ enters the structural model, its effect on the likelihood can be explicitly solved for. As in Farrell, Liang, and Misra (2020), Chernozhukov et al. (2022), and Ichimura and Newey (2022), I derive an influence function that measures the effect of varying γ on the first-order condition that determines ϑ_d. This influence function is then used as a correction term to make the estimation process insensitive to errors in the estimated high-dimensional parameter.

The orthogonal score contains two parts. The first part is the original score, or the gradient of the likelihood with respect to ϑ_d. In standard settings, without complications from a high-dimensional term, this first-order condition for the likelihood is used to pin down the parameters of interest. The second part of the orthogonal score is an adjustment term, or the influence function. This term involves a projection of how the likelihood varies with γ onto the space of the original

score; this "partials out" the influence of γ when estimating ϑ_d in an analogue to the Frisch-Waugh-Lovell theorem. I use this orthogonal score to estimate ϑ_d via generalized method of moments. Further details on the debiasing procedure are shown in Appendix A5.

5.1.ii Bidder arrival process

I assume bidders' entry equilibrium varies with observable seller and item characteristics. Specifically, I parameterize the expected number of bidders to enter an auction as

$$\Lambda_j \equiv \Lambda(r_j, Z_j \mid \lambda, \delta_0) = \exp(\delta_{0,1} + Z_j'\lambda + \delta_{0,2}\rho(r_j)) \tag{9}$$

where Z_j contains the seller feedback score, seller ratings, and the average log item value $\gamma(X_j)$. For tractability, I assume the additional arrival parameter λ is known to all sellers, regardless of their level of experience. To streamline notation, I combine the known arrival shifter $Z_j'\lambda$ with the intercept to write the item-specific arrival parameter $\delta_{j,0,1} \equiv \delta_{0,1} + Z_j'\lambda$, with $\delta_{j,0} \equiv \{\delta_{j,0,1},\ \delta_{0,2}\}$. I also use similar notation for beliefs, where b_j represents sellers' item-specific beliefs about bidder arrival with $\mathbb{E}_{b_j}[\delta_1] = \mathbb{E}_b[\delta_1] + Z_j'\lambda$.

The function ρ is determined by the bidder zero-profit condition at each point in the support of r. Since evaluating the high-dimensional index γ is itself computationally challenging, I opt for a reduced-form representation of the entry process instead of solving the fixed-point problem for ρ while estimating γ. I set ρ to be the product of the identity function and the CDF of the bidder value distribution evaluated at r, i.e. $\rho(r) = rF_B(r)$. The intuition for this choice comes from the zero profit condition for bidder entry and Proposition 1. The expected bidder surplus from entering the auction depends on the probability that their bid will be below the reserve price, $F_B(r)$; it is also strictly decreasing in r, which is the price paid when the reserve is binding. I also estimate the model with alternative specifications for ρ and find similar results.

5.1.iii Likelihood function for demand side

eBay's use of an increasing minimum bid creates a selection problem in which not all bids are observed for any given item. Due to randomness in the order of bidder arrival, some bidders will enter the auction after others have placed bids, learn their valuation, and not submit a bid because the minimum bid is already higher than their value for that item. In this and other online auction settings, the minimum bid will increase to the second highest of the existing bids and "lock out" subsequent arrivals (Platt 2017; Freyberger and Larsen 2022). This implies that only the two highest bids in an auction are known to correspond to the two highest-value bidders.

To counteract this problem, I derive a likelihood that explicitly models the selection process (see Appendix A6 for a detailed derivation). I first denote $v_j^{(k)}$ as the kth highest homogenized log-bid. Using the Poisson assumption on bidder arrival, I model the distributions of the reserve price r_j and the highest two bids $v_j^{(k)}$ when these bids are observed.[8] This likelihood also conditions on the starting minimum bid m_j (which is binding when there are fewer than two bids), the number N_j of bids observed, arrival shifters Z_j, and components of observable heterogeneity X_j. Together, the demand likelihood contribution for a single auction is

$$\begin{aligned}\ell_j^{\text{demand}}(\vartheta_d, \gamma \mid N_j, \{v_j^{(k)}\}, X_j, Z_j, m_j, r_j) = \\ & f_r(r_j \mid \vartheta_d, \gamma) \cdot \left[e^{-\Lambda[1-F_B(m_j|\vartheta_d,\gamma)]}\right]^{\mathbb{1}[N_j=0]} \\ & \cdot \left[f_B(v_j^{(1)} \mid \vartheta_d, \gamma)\Lambda e^{-\Lambda_j[1-F_B(m_j|\vartheta_d,\gamma)]}\right]^{\mathbb{1}[N_j=1]} \\ & \cdot \left[f_B(v_j^{(1)} \mid \vartheta_d, \gamma) f_B(v_j^{(2)} \mid \vartheta_d, \gamma)\Lambda_j^2 e^{-\Lambda_j[1-F_B(v_j^{(2)}|\vartheta_d,\gamma)]}\right]^{\mathbb{1}[N_j\geq 2]}\end{aligned} \tag{10}$$

Appendix A6 shows how this likelihood approach, together with the orthogonalization step, performs with simulated data.

After estimating the bidder value distribution F_B and the other bidder arrival parameters, I estimate the bidder entry cost c_B^E. For each auction, I compute the ex ante expected surplus $\pi_B(r \mid \Lambda, c)$ and use the bidder zero-profit condition in (4) to estimate bidders' entry cost as the mean of the expected bidder surplus across all auctions. Since bidder insertion fees c_B^I are zero in

[8] Note that r_j and $v_j^{(k)}$ are not observed directly; rather, each is a residual representing its counterpart in the data after being homogenized using the item value index $\gamma(X_j)$.

the dataset, this implies that the full bidder entry cost is in fact the time cost c_B^T.

5.2 Supply side estimation strategy

I now turn to the problem of identifying and estimating the supply side of the model. The remaining parameters of interest are the distribution F_S of sellers' outside option values, sellers' entry cost distribution $F_{c_S^E}$, and new sellers' beliefs about the bidder arrival parameter δ_0. While seller values and costs have been estimated in many similar settings, estimating seller beliefs is key to understanding the role of information and learning on the auction platform.

5.2.i Seller value and entry cost distributions

I first estimate all non-belief seller parameters (denoted ϑ_s) using data from experienced sellers. I make the simplifying assumption that the most experienced sellers have perfect information about the arrival process, so plugging the experienced sellers' reserve prices into the virtual type function $\psi(\cdot \mid \delta_0, c)$ yields the imputed seller values $\hat{v}_{0j}$ for all listed items.

In order to estimate the seller parameters ϑ_s, I make several functional form assumptions. I assume sellers' entry costs are i.i.d. Exponential, their values v_0 are from a two-component i.i.d. Gaussian mixture, and the nuisance distribution of entry parameters $\delta_{j,0,1}$ are i.i.d. Gaussian.[9] I also introduce some additional notation. First, denote $\tilde{c}(z)$ as the vector of seller costs where c_S^E is replaced with z and $\Pi^*(v_0 \mid \delta_0, c)$ as the maximized profit given seller value v_0 and costs c. Further, denote by $\bar{v}(\delta_{j,0}, c)$ the experienced seller entry threshold with known bidder entry parameter $\delta_{j,0}$ and costs c.

The likelihood contribution of a single auction run by an experienced seller is the density of the observed data, multiplied by a selection correction term to account for sellers' entry problem. The observed data includes both the implied seller values $\hat{v}_{0j}$ and the baseline arrival parameter $\delta_{j,0,1}$. There is random truncation in which data is observed due variation in entry costs c_S^E: each

[9] Since neither sellers' entry costs nor the full population of potentially-listed items are observed, I rely on stronger parametric assumptions to pin down the right tail of seller values. I do not assume sellers and bidders have the same value distribution, since the bidder value distribution consists of all bidders with a weakly positive valuation for the items, and the sellers are a selected sample of individuals who have by some means already acquired items.

item is only listed if the entry cost is below the maximum expected seller surplus (that is, with zero entry cost). The probability that this occurs for a specific imputed item value $\hat{v}_{0j}$ is divided by the probability that any seller lists their item, for all possible entry costs and arrival parameters. Taken together, this yields the likelihood contribution

$$\ell_j^{\text{supply-e}}(\vartheta_s \mid \hat{v}_{0j}, \delta_{j,0,1}) = \frac{f_s(\hat{v}_{0j} \mid \vartheta_s) \cdot f_{\delta_{j,0,1}}(\delta_{j,0,1} \mid \vartheta_s) \cdot F_{c_S^E}(\Pi^*(\hat{v}_{0j} \mid \delta_{j,0}, \tilde{c}(0)) - \hat{v}_{0j} \mid \vartheta_s)}{\int_0^\infty \left[\int_{-\infty}^\infty F_S(\bar{v}(l, \tilde{c}(z)) \mid \vartheta_s) f_{\delta_{0,1}}(l \mid \vartheta_s) dl\right] f_{c_S^E}(z \mid \vartheta_s) dz}$$

Since seller entry costs are unobserved, I further assume that the mean of the seller entry cost distribution is equal to $c_S^I + c_B^T$, where the mean platform entry fees are observed in the data and bidders' time cost is identified from demand-side data. Variation in $\delta_{j,0,1}$ shifts the resulting distribution of sellers' participation thresholds and traces out the distribution of seller values.

5.2.ii Identifying new sellers' prior beliefs

Having obtained the sellers' value distribution, I use variation in new sellers' reserve prices to identify their underlying beliefs about the bidder arrival process. At a high level, a seller's choice of reserve price is determined by both their beliefs and their private underlying value for the item. New information, including from previous auctions, affects current reserve prices only through beliefs, so variation in sellers' private values for sellers who receive the same information yields useful variation in reserve prices. The distribution of reserve prices can then be inverted to recover the underlying distribution of seller beliefs.

The following proposition offers formal conditions under which seller beliefs b_{δ_2} about the effect of reserve prices on arrival can be semiparametrically identified from reserve price data.

Proposition 5. Denote history $\mathcal{H}$ as a collection of data $\mathbf{D}$ from auctions. Let Assumption 1 hold, assume F_S is known, and further assume

(i) The prior b is common to all sellers with a shared history $\mathcal{H}$.

(ii) b is composed of independent marginal densities b_{δ_1} and b_{δ_2}, where both the marginal b_{δ_1} corresponding to the intercept and the seller entry threshold $\bar{v}(b, c)$ are known.

(iii) b_{δ_2} satisfies the Carleman condition, i.e. the absolute moments of b_{δ_2} (written as $\mu_j = \int |\delta|^j b_{\delta_2}(\delta_2) d\delta_2$) are finite for all $j \geq 1$ and satisfy $\sum_{j=1}^\infty \mu_j^{-1/j} = \infty$.

(iv) The reserve price r has full support on a positive-measure interval including some x^* for which $\rho(x^*) = 0$ and $\psi(x^* \mid b, c) < \bar{v}(b, c)$.

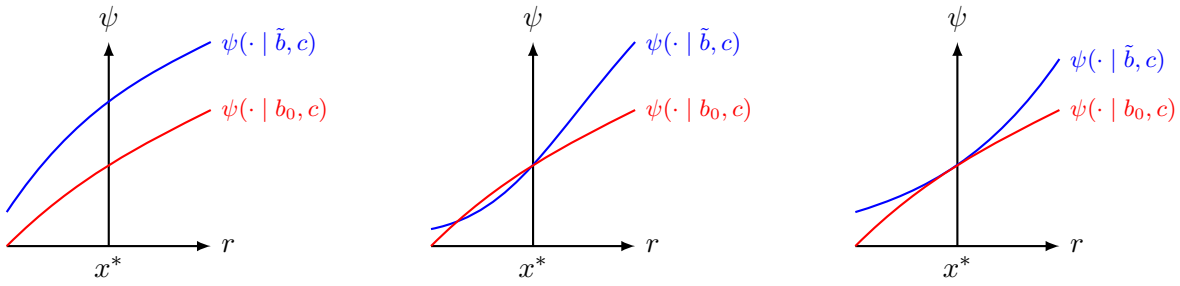
(v) For $\xi_k(\beta, x^*, c)$ defined as $\frac{\partial^{k-1}}{\partial x^{k-1}}\psi(x^* \mid b, c)$ with all terms of the form $\sum_{n=0}^{\infty} G_n(x^*)\frac{\partial^k}{\partial x^k}p_n(x^* \mid b)$ replaced with $\beta\sum_{n=0}^{\infty} G_n(x^*)\frac{\partial^k}{\partial x^k}p_n(x^* \mid b)$ for $G \in \{R, K\}$, ξ_k is invertible in β.

Then the marginal density b_{δ_2} of a seller with history $\mathcal{H}$ is identified up to its first $\bar{k}$ moments. (Proof in Appendix A3)

The proof proceeds in two parts, for which I give a heuristic explanation here. The first step is to recover the virtual type function ψ shared by all sellers with the same information and therefore same beliefs. All reserve prices satisfy the first-order condition $v_0 = \psi(r \mid b, c)$, where the distribution F_S of v_0 is known. Given sellers' participation threshold $\bar{v}(b, c)$, the distribution of observed reserve prices can be written as a known, invertible function of $\psi(\cdot \mid b, c)$.

The next step of the proof is to invert the virtual type function to recover the marginal density b_{δ_2} of seller beliefs about the effect of the reserve price on bidder arrival. Figure 6 provides some intuition for this identification problem: the behavior of the virtual type function around a particular reference point $r = x^*$ reveals the moments of the belief density of interest. A key insight from Fox et al. (2012) is that evaluating a function and its derivatives at some carefully chosen point x^* (that is, where $\rho(x^*) = 0$) can greatly simplify these functions. This yields known functions of *only* model primitives that do not depend on b_0 (such as $F_B(x^*)$) and various raw moments of the true density b_0. These functions can then be inverted to recover the moments of the underlying belief density. Higher-order derivatives of ψ with respect to the reserve price can be inverted to obtain higher-order moments of b_{δ_2} and reject other densities that do not yield the same virtual type function.

Figure 6: Intuition for identification of beliefs b_0 through variation in ψ and its derivatives



(a) Identifying $\mathbb{E}_{b_0}[\delta_2]$ by evaluating $\psi(x^*|b, c)$

(b) Identifying $\mathbb{E}_{b_0}[\delta_2^2]$ by evaluating $\psi'(x^*|b, c)$

(c) Identifying $\mathbb{E}_{b_0}[\delta_2^3]$ by evaluating $\psi''(x^*|b, c)$

Notes: Each panel shows a thought exercise of comparing the virtual type function under the true belief density b_0 with the virtual type function under various candidate belief densities $\tilde{b}$ that share the same lower-order moments (i.e., in panel (b) the two densities have the same first moment, and in panel (c) the two densities have the same first two moments). Examining the virtual type function at x^* discards all candidates $\tilde{b}$ such that $\psi(x^* \mid \tilde{b}, c) \neq \psi(x^* \mid b_0, c)$, while examining the first derivative of the virtual type function of x^* discards all candidates $\tilde{b}$ such that $\psi'(x^* \mid \tilde{b}, c) \neq \psi'(x^* \mid b_0, c)$, and so on.

The identification result is general since an arbitrary number of moments of the marginal belief distribution are identified. However, it is important to note that these beliefs are only identified in the context of the larger parametric model. It is not necessary to assume that sellers are Bayesian: identification of b_{δ_2} is achieved with data from a single period for all sellers that observe the same data history $\mathcal{H}$. In principle this allows researchers to test whether sellers follow different candidate learning rules, though this is outside the scope of this paper.

Though Proposition 5 offers semiparametric identification of b_{δ_2}, the necessary assumptions are somewhat restrictive. First, Assumption 1 imposes that all sellers' outside option values are drawn i.i.d. from the same distribution, which rules out time-invariant heterogeneity in sellers' value distributions. Also, assumption (i) of the proposition rules out unobserved determinants of beliefs to allow comparisons between sellers that receive identical information. Assumption (ii) is also restrictive: other components of the sellers' decision problem must be known and separable in some sense (here, independence of marginal beliefs) to isolate the effect of beliefs about any one parameter. Assumptions (iii) and (iv) are similar to assumptions in the random coefficients literature, principally Fox et al. 2012; these use variation in a linear index to recover population densities, while I study an individual's belief density. Assumption (v) is also technical, and requires the derivatives of the virtual type function to be invertible in weighted sums of the derivatives of bidder arrival probabilities. This is a joint restriction on the seller beliefs about $\delta_{0,1}$ and the bidder

value distribution F_B at the point x^* from assumption (iv).

This result is related to others in the literature on identifying individual beliefs in structural models. Lu (2019) shows that state-dependent beliefs can be identified in a setting with finite support and Bayesian updating; in contrast, I do not require Bayesian updating and allow for absolutely continuous density functions. Wang et al. (2024) likewise adopts a finite-support approach with Bayesian updating, which is used to identify beliefs about time-varying macroeconomic trends; Wang and Yang (2024) offers more general results in finite-support settings for both myopic and forward-looking agents. Aguirregabiria and Magesan (2020) semiparametrically identifies firm beliefs within a game, and similarly relies on a finite support. While they do not require Bayesian updating, they require beliefs to align with the truth in some cases.

5.2.iii Estimating new sellers' prior beliefs

Though beliefs are semiparametrically identified for each history $\mathcal{H}$ of auction data, in practice I make several additional assumptions on the prior and updating process for computational tractability. The model as written has no conjugate prior structure, so the posterior depends on an increasingly large state space of all possible profit signals, reserve prices, private seller values, and covariates for all sellers' previous auctions. Further, solving for each seller's path of posterior distributions is computationally costly, especially when searching for the parameters of the true prior distribution.

I simplify the state space by using a parametric approximation to Bayesian updating, and I solve the computational bottleneck by approximating this updating process with neural networks. I assume sellers' initial beliefs about δ_0 are bivariate normal, with parameters jointly denoted as ϑ_b. I also assume beliefs are updated according to a modified Laplace approximation to Bayes' rule: each period, sellers' beliefs are a bivariate normal with mean equal to the maximum a posteriori estimate of the true Bayesian posterior and covariance matrix given by the curvature of the true Bayesian posterior at the maximum a posteriori estimate. This assumption helps solve a major computational challenge, since the prior is not conjugate with the posterior due to the nonlinear dependence of expected profit on the bidder arrival parameters.[10] Rather than solving each updating step for

[10] In addition to this nonlinearity, each sellers' path of beliefs about the arrival parameter can evolve differently

all individuals and auctions in the data (a prohibitively slow process), I simulate possible beliefs, covariates, and signals and fit a neural network to approximate the updating process for new sellers' beliefs. Appendix A7 discusses additional details of the estimation procedure, particularly the steps taken to approximate the updating rule and other associated functions.

I also exploit the structural model in estimation to determine other variables of interest. I use $\psi(r^* \mid b, c)$ as a control function for seller values v_0 since they are an unobservable but critical component of sellers' updating and decision processes. Assuming that all beliefs are bivariate normal helps pin down b_{δ_1}, which allows me to relax the independence assumption in Proposition 5. Together, parametric assumptions on the entry cost distribution and sellers' prior yield the distribution of sellers' entry threshold $\bar{v}(b, c)$.

Having established the identification of new sellers' prior parameters, I now explain the likelihood approach used for estimation. I use a change of variables to obtain the density of the new sellers' reserve prices from their underlying value distribution. The density of sellers' reserve prices, conditioning on beliefs b_t and costs $\tilde{c}(0)$,[11] is obtained from the distribution of seller values

$$\frac{\partial}{\partial x}\mathbb{P}[r \leq x] = \frac{\partial}{\partial x}\mathbb{P}[v_0 \leq \psi(x \mid b_t, \tilde{c}(0))] = f_s(\psi(x \mid b_t, \tilde{c}(0)) \mid \vartheta_s) \cdot \psi'(x \mid b_t, \tilde{c}(0))$$

where $\psi'(x \mid b_t, \tilde{c}(0))$ is the first derivative of the virtual type function with respect to x. As with the experienced-seller likelihood, there is a selection term to account for the probability of each item being listed by any given seller: the numerator is the probability that an item is listed conditional on $\hat{v}_{0j}$ (i.e., the seller has a favorable-enough entry cost to list the item) and the denominator is the probability that any seller with beliefs b_t lists an item. The likelihood contribution of a single

according to their signals. The Laplace approximation is one of several posterior approximations used in Bayesian statistics, and imposing that beliefs are updated in this manner ensures that learning follows a computationally tractable Markov process with a relatively low-dimensional state.

[11] The entry cost is irrelevant for the virtual type function and can be treated as zero; see 4.2.i for reference.

auction of item j run by an inexperienced seller in their tth auction is therefore

$$\ell_{jt}^{\text{supply-i}}(\vartheta_b \mid \mathbf{D}_{jt}, \vartheta_s) = f_s(\hat{v}_{0j} \mid \vartheta_s) \cdot \psi'(r_j \mid b_t, \tilde{c}(0)) \cdot \frac{F_{c_S^E}(\Pi^*(\hat{v}_{0j} \mid b_t, \tilde{c}(0)) - \hat{v}_{0j} \mid \vartheta_s)}{\int_0^\infty F_S(\bar{v}(b_t, \tilde{c}(z)) \mid \vartheta_s) \cdot f_{c_S^E}(z \mid \vartheta_s) dz}$$

$$s.t. \quad b_{t+1} = \mathcal{T}(b_t, \mathbf{D}_{jt} \mid \vartheta_b) \; \forall t$$

$$\hat{v}_{0j} = \psi(r_j \mid b_t, \tilde{c}(0))$$

Though beliefs in this model can be identified from data within a single period, additional variation across and within different sellers is pooled via the assumed learning rule to aid in estimation.

5.3 Estimates from the platform model

Table 1 presents the estimated bidder arrival parameters. As predicted by the model, the reserve price coefficient $\delta_{0,2}$ is negative, indicating that higher reserve prices deter potential bidders from entering auctions. The relationship between bidder entry and the covariates is intuitive as well: more bidders enter when sellers have better ratings and more feedback, as well as when the estimated mean item value is higher.

Table 1: Estimated bidder arrival parameters and entry cost

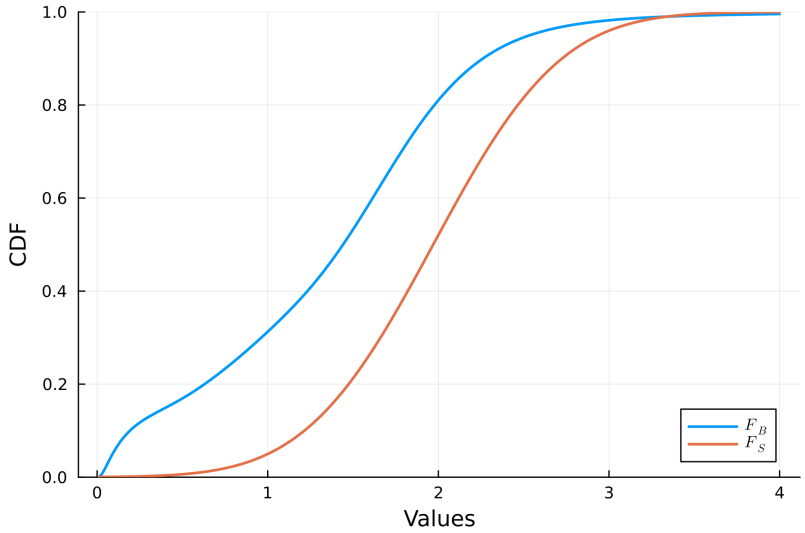
			λ_0				
	$\delta_{0,1}$	$\delta_{0,2}$	Rating	IHS(Feedback)	ln(Pred. Item Value)	c_B^T	$\mathbb{E}[c_S^E]$
Estimate	0.871	-0.245	0.055	0.146	0.409	0.054	0.131
Std. Err.	(0.059)	(0.007)	(0.026)	(0.010)	(0.022)	(0.004)	(0.004)

Notes: The estimated coefficients are obtained via debiased GMM; further details are presented in Appendix A5. Estimates and standard errors for all but the entry costs are computed using the optimal weighting matrix. The entry cost parameters are computed using mean insertion fees for sellers and the average expected bidder surplus per auction; standard errors for both are adjusted for first-step estimation of the other demand side parameters.

The estimated bidder entry cost comes from the equilibrium bidder entry condition and the other estimated bidder arrival parameters. Evaluating the expected zero profit condition in equation (4) yields an estimated homogenized bidder time cost c_B^T of 0.056, since there are no bidder insertion fees. For the average (median) item in the dataset, this is approximately \$0.48 (\$0.33). This represents a moderate but not prohibitive time cost to entering each auction and inspecting the listing. The average homogenized seller insertion fee is approximately $c_S^I = 0.075$, though sellers' overall entry costs are assumed to be heterogeneous for different items.

As described in the previous section and as is common in the auction literature, the estimated demand parameters yield the optimal reserve pricing rule for sellers with perfect information. I use the sample of reserve prices chosen by experienced sellers (defined as those in the top 25% of sellers by experience at the start of the data) to impute the sellers' outside option for each item, under the assumption that experienced sellers have perfect information about the bidder arrival process. Figure 7 plots the estimated value distribution F_S along with the bidder value distribution F_B obtained in the demand-side estimation.

Figure 7: Estimated value distributions for auction participants



Notes: The bidder value distribution F_B is estimated via the maximum likelihood approach in 5.1 using a 5-component Gaussian mixture model for log values. The seller outside option distribution F_S is fit to imputed seller values among experienced sellers, as in 5.2, and uses a 2-component Gaussian mixture model for seller values.

The estimated seller value distribution in Figure 7 largely first-order stochastically dominates the estimated bidder value distribution. Since the population of sellers is the group of users who have previously acquired Beanie Babies, it is reasonable for them to have a higher value distribution for these items than any random bidder. However, this difference in value distributions is not unreasonably large: the seller value distribution largely falls between distributions of the maximum value of two bidders and that of three bidders (which are not plotted here), so it may be profitable in expectation for a seller with a large $\hat{v}_{0j}$ to list an item for sale.

I now estimate new sellers' beliefs about the bidder arrival process. I estimate the model on a sample of all sellers with at least 5 auctions in the data, to restrict attention to "serious" sellers. I

also limit the sample to the first 5 auctions of all such sellers to focus on the early beliefs of new sellers. The prior mean for the reserve price coefficient $\delta_{0,2}$ is higher than the estimated parameter -0.245, implying that new sellers' beliefs about bidder arrival are upwardly biased, and particularly so for auctions with high reserve prices. The variance of beliefs about $\delta_{0,2}$ is moderately large, allowing sellers to update their beliefs about the effect of reserve prices on bidder entry.

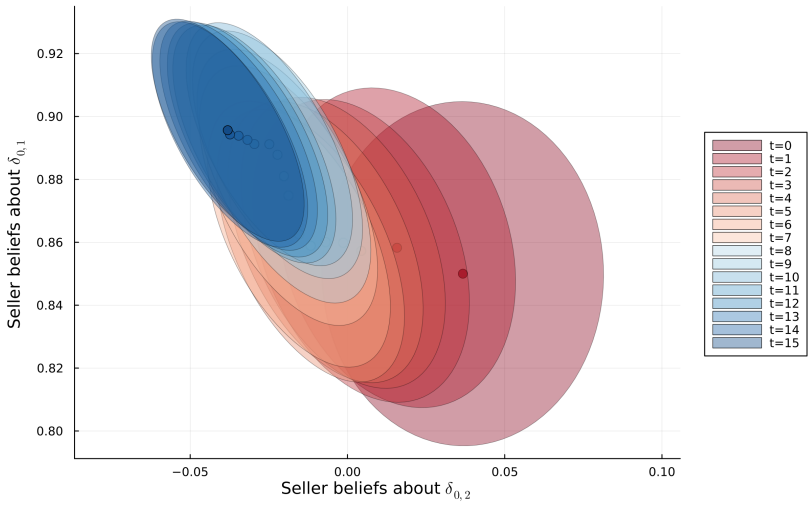
Table 2: Estimated new seller priors about the bidder arrival process

	$\mathbb{E}[\delta_{0,1} \mid b_0]$	$\mathbb{E}[\delta_{0,2} \mid b_0]$	StdDev$(\delta_{0,1} \mid b_0)$	StdDev$(\delta_{0,2} \mid b_0)$	Cor$(\delta_{0,1}, \delta_{0,2} \mid b_0)$
Estimate	0.85	0.037	0.547	0.448	-0.006
Std. Error	(3e-5)	(2e-5)	(3e-5)	(3e-5)	(4e-5)

Notes: The model is estimated on the first 5 auctions (where applicable) of all 3,975 new sellers that list at least 5 auctions for sale. Standard errors are naive standard errors, treating seller value distribution parameters as known (*standard error correction for two-step estimation in progress*).

To help interpret the estimated prior parameters in Table 2, I evaluate the implied path of new sellers' average beliefs about the unknown parameter δ_0. Figure 8 plots ellipses corresponding to the estimated beliefs of these new sellers for up to their first 15 auctions; the contours represent 0.1-standard deviations around the mean. The prior mean shifts with successive auctions, moving toward a lower arrival coefficient $\delta_{0,2}$.

Figure 8: Estimated path of average new seller beliefs

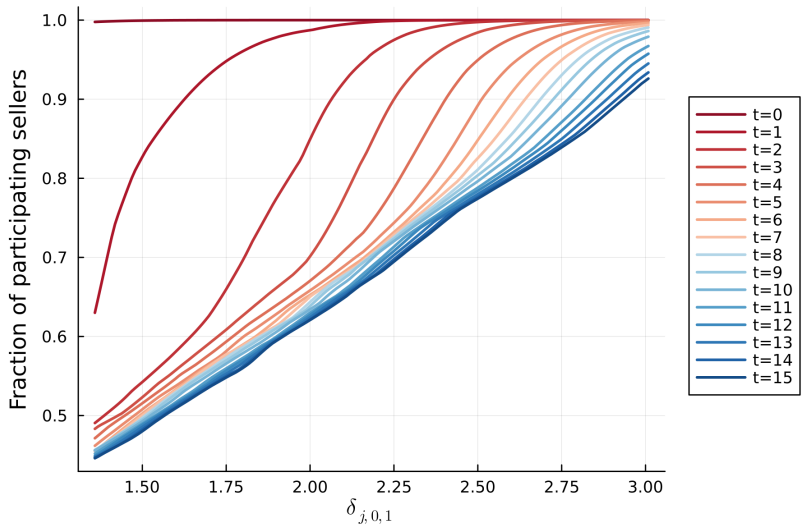


Notes: The ellipses represent the average beliefs of new sellers in their first 15 auctions, conditioning on all 3,975 new sellers who list at least 5 auctions. The contours represent a 0.1-standard deviation around the mean of the belief distribution.

Finally, I directly show the effect of seller beliefs on selective entry by plotting the average seller

entry threshold for new sellers relative to the experienced seller entry threshold. Figure 9 shows how the experienced sellers' entry threshold is increasing in the expected arrival rate, where the entry threshold is evaluated at the mean of the estimated entry cost distribution $F_{c_S^E}$. The entry threshold is significantly higher among new sellers in their first auctions. This is consistent with Table 2: in spite of the lower prior mean for $\delta_{0,1}$ (relative to the true parameter), its high variance combined with the higher prior mean for $\delta_{0,2}$ makes entry attractive to new sellers. This relative ordering of entry thresholds (most visible for items with a high baseline entry parameter $\delta_{j,0,1}$) is consistent with the selection pattern in Figure 4, where new sellers that exit early also set higher reserve prices in their first auctions.

Figure 9: Estimated average seller entry threshold conditional on baseline bidder arrival parameter $\delta_{j,0,1}$



Notes: The estimated entry threshold is evaluated at the mean of the estimated entry cost distribution, conditional on the log of the expected number of bidders for $r = 0$. I plot these entry thresholds between the 2.5% and 97.5% quantiles of the observed baseline entry parameters $\delta_{j,0,1}$.

6 Counterfactual platform design with information provision

I now study how information provision by the platform affects the optimal fee structure, platform profits, and both bidder and seller welfare. I simulate platform outcomes under alternative information structures and show that the platform can improve its own profits, as well as bidder entry and seller surplus, from providing information. I also estimate the difference between new sellers' expected return to information and the true return to information, which helps explain the

prevalence of free information services provided by platforms.

6.1 Revisiting the platform's problem of fees and information provision

The estimated structural model characterizes the full platform problem, taking into account how bidders and sellers act in equilibrium. As in the expositional model in section 2, sellers choose whether to participate on the platform and what reserve prices to set as a function of their beliefs and the fees they face. I assume the platform faces both inexperienced sellers (with mass ω) that update their beliefs when they receive additional information, and experienced sellers (with mass $1-\omega$) that are not affected by information provision. The choice of fees affects sellers of both types, and when the platform cannot provide more information to inexperienced sellers, the optimal fee structure depends on the fraction of sellers who are new.

The auction platform chooses information provision and fees jointly to maximize profits. The platform sets seller-facing fees[12] c_S^I and c_S^R, where the former shifts the average entry cost to potential sellers.[13] The platform can provide information in the form of a sample auctions for each seller, comprising a dataset $\mathbf{D}_a$ that is drawn i.i.d. from the distribution $F_{\mathbf{D}}$ of auctions run by experienced sellers under the true data-generating process. All new sellers use this dataset to update their priors from b_0 to $\mathcal{T}(b_0, \mathbf{D}_a)$ before their first auction, and all sellers have the option to list T items.[14] Formally, the platform's profit maximization problem for each potentially listed

[12] I treat bidder-facing fees c_B^I and c_B^R as being fixed at zero. This is motivated both by the fee being equal to zero for bidders in the data, as well as analysis by Marra (2019) on a wine auction platform indicating that increasing revenue while maintaining transaction volume requires setting $c_B^R < 0$.

[13] Recall that the distribution of entry costs is i.i.d. Exponential, with mean $\hat{c}_S^I + c_B^T$ where $\hat{c}_S^I$ is the average insertion fee in the dataset. I assume that choosing $c_S^I \neq \hat{c}_S^I$ affects only the mean of the shifted Exponential distribution $F_{c_S^E}(\cdot \mid c_S^I)$; all other moments are determined by the original fee structure.

[14] This assumption simplifies the seller arrival process over time. In practice, sellers may differ in how many items they wish to sell, and new waves of sellers may arrive at different times. This particular setup highlights the tradeoffs between focusing on new sellers that learn over time and experienced sellers that are unaffected by information provision.

item is

$$\max_{a,\, c_S^I,\, c_S^R} \int\int \Bigg(\omega \cdot \sum_{t=1}^{T} \frac{1}{T} \mathbb{E}_{b_{t-1}} \Bigg[\int \underbrace{\mathbb{1}[v_0 \leq \bar{v}(\mathcal{T}(b_{t-1}, \mathbf{D}_a), \tilde{c}(z))]}_{\mathbb{P}[\text{Entry} \mid \text{inexperienced}]} \cdot \underbrace{\left[c_S^I + c_S^R \cdot \int R(r^*(v_0 \mid \delta, c_S^R) \mid \delta_0) \mathcal{T}(b_{t-1}(\delta), \mathbf{D}_a) d\delta \right]}_{\text{Platform revenue} \mid \text{entry, inexperienced}} \Bigg| \, b_0,\ \mathbf{D}_a \Bigg] dF_{\mathbf{D}}(\mathbf{D}_a) + (1-\omega) \cdot \underbrace{\mathbb{1}[v_0 \leq \bar{v}(\delta_0, \tilde{c}(z))]}_{\mathbb{P}[\text{Entry} \mid \text{experienced}]} \cdot \underbrace{\left[c_S^I + c_S^R \cdot R(r^*(v_0 \mid \delta_0, c_S^R) \mid \delta_0) \right]}_{\text{Platform revenue} \mid \text{entry, experienced}} \Bigg) dF_S(v_0) dF_{c_S^E}(z \mid c_S^I) \tag{11}$$

where b_t evolves from b_0 according to the updating rule $\mathcal{T}$ and the path of observed profit signals. As in the simple model, I assume there is zero marginal cost to providing data and hosting auctions on the platform. Though it is likely that there is some fixed cost in providing sellers with data (moving from $a = 0$ to $a > 0$), I assume this cost is sunk and not relevant for the platform's problem. I continue to use standardized average item values, abstracting away from heterogeneity in the dollar value of each potentially-listed item.

The maximization problem in (11) highlights the importance of sellers' initial beliefs in the platform's choice of fees and information provision. The first two lines correspond to the new sellers on the platform, each of whom observes a dataset $\mathbf{D}_a$ with a auctions before listing their first item. Each seller only chooses to list an item if their updated beliefs $\mathcal{T}(b_0, \mathbf{D}_a)$ about the bidder entry process imply it is optimal to do so; conditional on entry, they will also choose a reserve price according to their beliefs. Thus, knowledge of b_0 is critical for understanding how any additional information $\mathbf{D}_a$ will shift seller beliefs and therefore behavior. The third line corresponds to experienced sellers, whose selective entry and choice of reserve price are informed by their perfect information about the bidder entry process. While these sellers are not directly affected by the platform's information provision, they must pay fees c_S^I and c_S^R which apply to all sellers regardless of their beliefs and are chosen by the platform jointly with its information provision a.

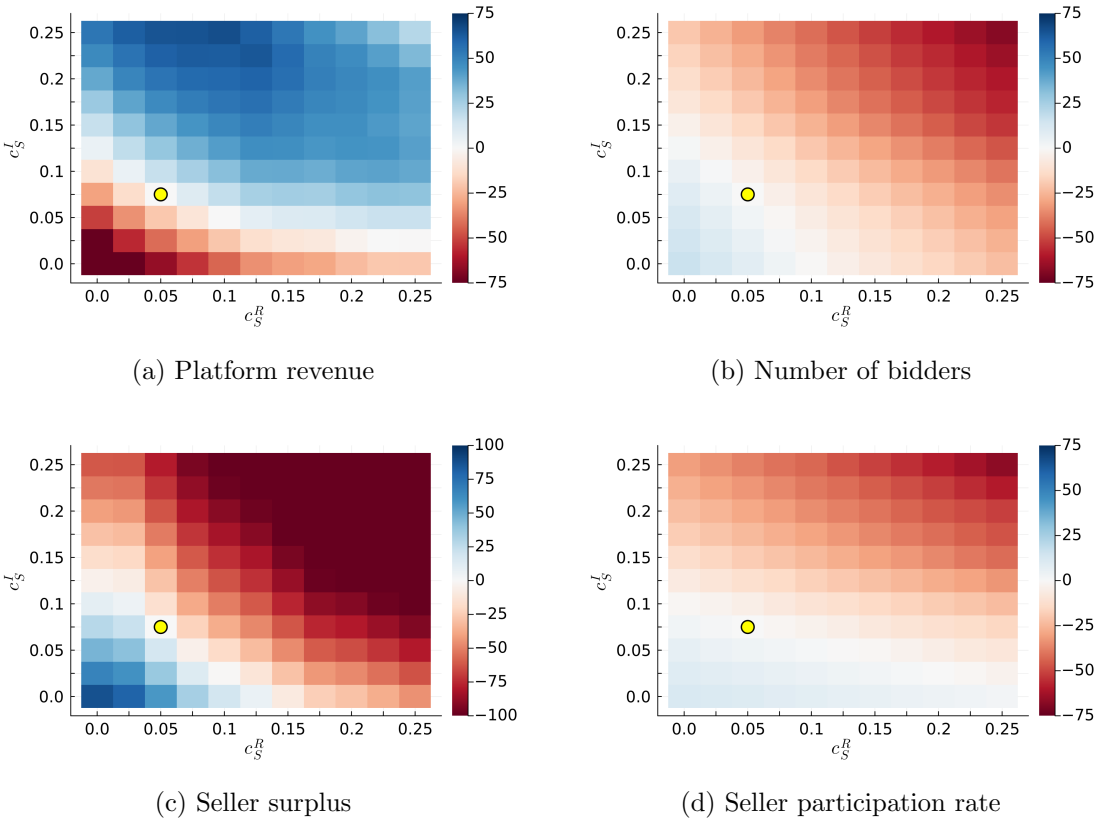
6.2 Optimal fee structures without information provision

As a benchmark, I examine how platform outcomes change under alternative fee structures in the absence of information provision. The average seller fees on the platform when the data were

collected are approximately $c_S^R = 0.05$ and $c_S^I = 0.075$. Since about half of all items in the data are listed by experienced sellers, I use $\omega = 0.5$ as the baseline fraction of new sellers, and I set $T = 15$ as the number of items that may be listed by each seller.

Changing insertion and revenue fees alone can significantly impact the platform's profits and bidder and seller outcomes. Figure 10 shows percent changes in platform revenue, the number of bidders, seller surplus, and the seller participation rate relative to the baseline fee structure. In general, increasing fees benefits the platform and reduces bidder and seller participation, as well as seller welfare. However, platform revenue in panel (a) is non-monotonic in the revenue fee: the reason for this can be seen in the severe entry effects for both bidders and sellers in panels (b) and (d). High fees reduce the probability that sellers list items on the platform and drive up the average reserve prices, which contributes to a lower number of bidders.

Figure 10: Percent changes in outcomes under alternative fee structures, $a = 0$ and $\omega = 0.5$



(a) Platform revenue

(b) Number of bidders

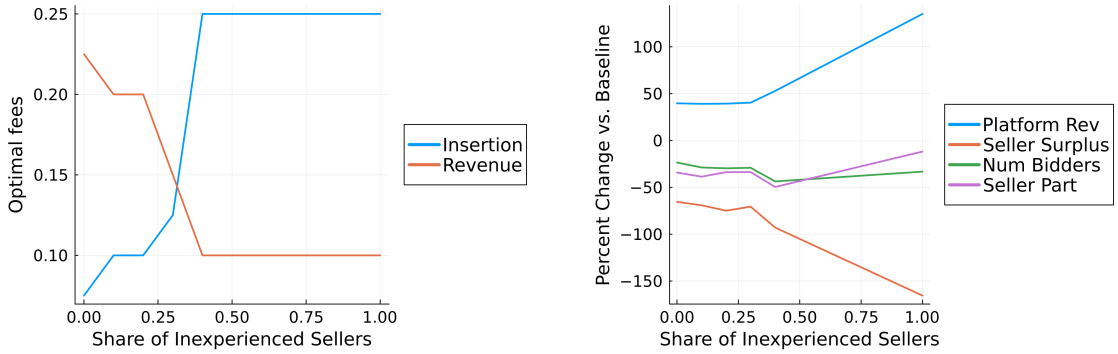
(c) Seller surplus

(d) Seller participation rate

Notes: Each scenario was evaluated from a grid of possible fees (in increments of 0.025 from 0 to 0.25), simulating 250 sellers for each combination of parameter values. Blue (red) squares represent a percentage increase (decrease) in a given outcome relative to the outcome under the baseline values of c_S^R and c_S^I. The yellow dot indicates the baseline fee structure in the dataset.

The effect of fees on platform outcomes varies with the composition of sellers—whether inexperienced or experienced—on the platform. Figure 11 plots the optimal seller fees under varying fractions ω of new sellers on the platform, as well as changes in various platform outcomes (under the optimal fee structure) relative to the baseline fee structure at $\omega = 0.5$. I keep all other parameters the same as in Figure 10. The optimal fee structure changes as ω increases: when most sellers are experienced, it is optimal to set a low insertion fee and instead charge a higher revenue fee. However, when the platform faces less knowledgeable sellers it becomes optimal to set a higher insertion fee and a lower revenue fee. This is because new sellers are more optimistic and therefore willing to pay a higher insertion fee, though this leads to lower realized surplus. At the same time, lower revenue fees decrease the average reserve price and allow for more successful transactions. Note that as the fraction ω of new sellers increases, seller surplus goes down while the overall entry rate increases, since the greater portion of inexperienced sellers are behaving suboptimally.

Figure 11: Optimal fees and outcomes for varying ω with no information provision ($a = 0$)



(a) Optimal fees (b) Platform outcomes

Notes: Each scenario was evaluated from a grid of possible fees (in increments of 0.025 from 0 to 0.25), simulating 250 sellers for each combination of parameter values. The platform outcomes of interest are platform revenue, seller surplus, number of bidders participating on the platform, and the seller participation rate. I simulate 250 sellers for each combination of parameter values.

Before proceeding, I briefly address possible reasons why the estimated optimal platform fees differs from the chosen platform fees in the data. First, many platforms pursue a growth-oriented strategy in their early years, during which they may forego short-term profits to develop a platform and gain additional users; for tractability I do not model these more complex dynamic incentives in addition to the detailed seller learning problem. I also treat the platform as a monopolist and ignore both existing competitors and the threat of competition (which limit the extent to which the platform may wish to increase fees) for the purposes of these counterfactuals. However, to address

unmodeled constraints by the platform during this period, I will also consider the case where the platform can only implement Pareto-improving fees for any information structure.[15] Regardless of why eBay did not implement higher fees during the sample period, this policy has since changed: as of November 2024, eBay charges relatively low (or zero) insertion fees for most moderately priced items, with revenue fees of 13-15% across many item categories.[16]

6.3 Optimal fees with information provision

I now quantify the platform's joint problem of choosing fees and information provision. To do so, I simulate multiple auctions for inexperienced and experienced sellers under both alternative fee structures and alternative amounts a of information provided to new sellers. I assume this information (in the form of "sample auctions" provided from the platform to the sellers) is learned costlessly by all new sellers, and I maintain the share of new sellers at $\omega = 0.5$.

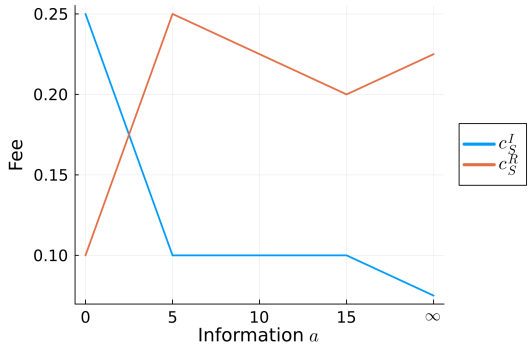
Figure 12 compares outcomes under different information regimes. I consider two cases: where the platform is choosing a fee and information structure that is optimal for its own profit maximization problem, and a second case where the platform constrains itself to search for only Pareto-improving choices. When the platform is unconstrained, it benefits most by *not* providing new sellers with any information and instead charging high insertion fees to naive sellers. Panel (a) of Figure 12 shows that the optimal insertion fee generally declines with the amount of information provided to new sellers, while the optimal revenue fee increases relative to zero information provision. This leads to more platform revenue, though with lower seller and bidder entry onto the platform. In particular, when the platform is a monopolist that is not constrained in any way (e.g. by regulation or threat of entry), the platform may act in a way that significantly decreases sellers' expected surplus. The difference in optimal fees across information structures highlights how understanding sellers' beliefs is important for optimal platform design.

In contrast, when the platform is constrained to make only Pareto-improving changes, the

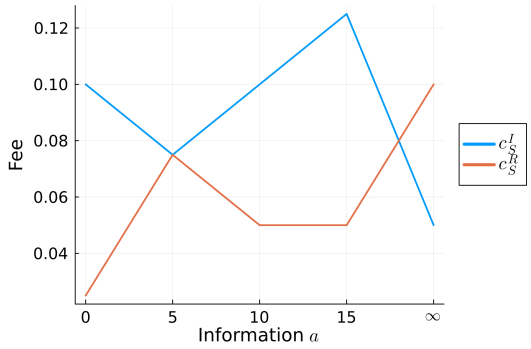
[15] Since bidders have a zero profit condition that determines entry, in expectation no bidder's welfare will be changed by these regimes. To compare bidder outcomes to those of sellers and the platform, I measure the average number of bidders that enter an auction on the platform. Any "Pareto" improvement in this exercise must maintain at least the same number of bidders as under the baseline fee structure—i.e., ensuring at least the same amount of possible bidder surplus is maintained—as well as weakly increasing both platform and seller surplus (with one strict improvement).

[16] See eBay's website for details on its current fee structure.

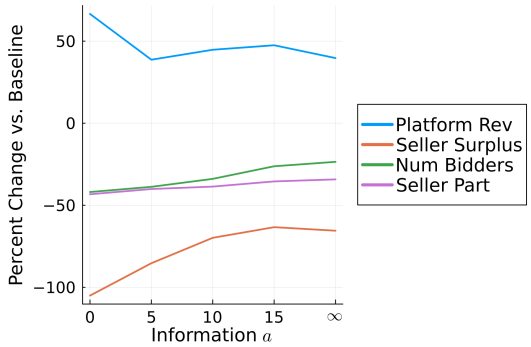
Figure 12: Optimal fees with information provision



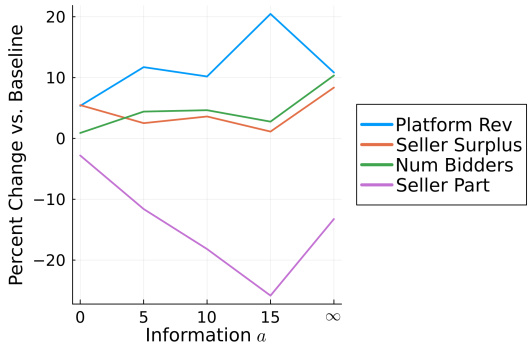
(a) Platform-optimal fees



(b) Pareto-optimal fees



(c) Outcomes under platform-optimal fees vs. baseline



(d) Outcomes under Pareto-optimal fees vs. baseline

Notes: Each scenario was evaluated from a grid of possible fees (in increments of 0.025 from 0 to 0.25) and information provision (for $a \in [0, 5, 10, 15, \infty]$, where ∞ represents perfect information provision i.e. the platform's maximum likelihood estimate being given to the sellers), simulating 250 sellers for each combination of parameter values. The platform outcomes of interest are platform revenue, seller surplus, number of bidders participating on the platform, and the seller participation rate. Panels (a) and (c) represent how optimal fees and platform outcomes change with the amount of information provision when the platform maximizes its own profits; panels (b) and (c) repeat this exercise with the constraint that the chosen fee structure must yield a Pareto improvement for the given level of information provision.

platform wishes to provide some information to sellers to help guide their behavior. Panel (b) of Figure 12 indicates that platform profits are maximized by giving $a = 15$ observations to new sellers and optimizing fees accordingly. This leads to a positive change in bidder surplus that is competed away by an increase in the number of bidders. Seller participation is lower under the alternative information structure than in the baseline setting because new sellers with higher private values v_0 are more selective—and profitable—in their entry. The platform still withholds some information in this case, since doing so allows it to charge somewhat higher fees from naive sellers who over-participate relative to the full-information setting, while still benefitting from shifting new sellers' behavior.

6.4 Seller willingness to pay for information

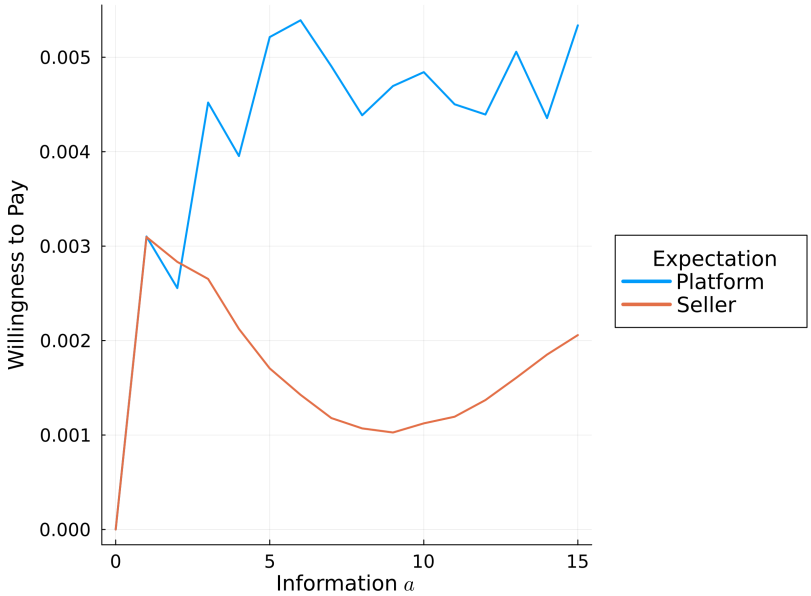
I now examine the difference between new sellers' *ex ante* expected value from receiving information from the platform and the true expected value from providing new sellers with information. As shown in the previous section, information provision can be used to improve platform profits by shaping new sellers' behavior. However, new sellers' biased beliefs mean they may not value their own learning as much as the platform. Throughout this exercise I assume seller fees are fixed at the baseline values of $c_S^R = 0.05$ and $c_S^I = 0.075$, and restrict attention only to new sellers ($\omega = 1$).

New sellers may not value information properly when their beliefs about the bidder arrival process are biased. A new seller's willingness to pay for information is the gap between expected surplus when receiving information from the platform and when learning unassisted. However, a new seller's expected evolution of their own beliefs depends on their current, biased beliefs about the true bidder arrival process. This means that they do not anticipate their beliefs or outcomes changing significantly from additional information. In contrast, the true expected returns from information are higher because new sellers' update their beliefs more significantly.

Figure 13 plots new sellers' willingness to pay for information under their initial, biased beliefs against the true expected gains from the provision of additional observations from the data-generating process. Note that the willingness to pay under sellers' expected path of future beliefs is lower than when conditioning on the platform's full information set.[17] Thus, new sellers' subjective expectation of the marginal value of information is lower than its true value. This difference helps explain why some platforms may offer seller tools for free rather than imposing additional fees. New sellers may undervalue these tools due to their biased beliefs, and the platform can benefit more from better-informed sellers' choices than from charging information fees to uninterested sellers, especially if there is some hassle cost for sellers in accessing and using this data.

[17] This holds with the exception of the first few units of information (observations) provided to sellers; the nonlinearity in expected surplus is due to taking expectations over profit which is itself nonlinear in the unknown parameters.

Figure 13: Estimated seller willingness to pay for information under different beliefs



Notes: This figure plots the willingness to pay (as a fraction of mean item value) for various amounts of information a when sellers know they will list 15 items for auction. These are simulated differences between sellers' expected willingness to pay for different amounts of information under both new sellers' expectation of the value of information and the platform's expectation of the value of new information. The blue line is the expected gain from information to sellers when conditioning on the platform's knowledge of the data-generating process and new sellers' initial prior, while the orange line is sellers' expected gain from information when conditioning on new sellers' biased beliefs of how their own beliefs will evolve. I use 2,500 simulated sellers for each amount of information and for both types of expectations.

7 Conclusions

This paper studies the problem of information provision by an auction platform where sellers face uncertainty about the bidder arrival process. I first present evidence that new sellers learn to set optimal reserve prices as they gain more experience. I pair a model of seller learning with a model of two-sided endogenous entry onto an auction platform to investigate how new seller behavior is driven by both selection and learning. The model implies a new reserve price formula that is designed to both attract bidders to the auction and extract surplus from them; failure to account for the negative effect of high reserve prices on bidder entry leads new sellers to set higher-than-optimal reserve prices. I show that sellers' beliefs about the bidder arrival process can be semiparametrically identified from reserve price data under certain conditions, estimate new sellers' beliefs, and evaluate their implications for the platform design problem.

These results highlight platforms' ability to influence user behavior outside of its well-known ability to charge different fees to different sides of the market. Information provision improves the

quality of the marketplace for new sellers, who are able to better optimize their entry decisions and pricing strategy. Sellers' improved profits increase the transaction volume on the platform, which increases platform profits while also inducing additional bidder entry through lower average prices on the platform. Thus, Pareto improvements can be achieved through more widespread access to important economic knowledge. However, this may not always align with platforms' incentives: withholding information may allow the platform to extract revenue from uninformed sellers (or possibly, in other settings, uninformed buyers). In particular, a large platform that faces little competition may not be incentivized to use additional information to improve seller or bidder outcomes.

More broadly, this paper speaks to the importance of information and its availability in the modern economy. Platforms are increasingly common in areas such as consumer products (eBay, Amazon), social media and advertising (Facebook/Meta, Twitter/X), financial services (NYSE, cryptocurrency exchanges), transportation (Uber, Lyft), and online betting (DraftKings, FanDuel). Users of each platform may lack crucial knowledge that affects their behavior on that platform, making them susceptible to the influence of information or misinformation in developing their beliefs. As the set of tools that can be used to shape information expands, helped in part by recent advances in generative AI, platforms face a more complex problem of how to shape their users' beliefs. Whether the platform benefits from doing so, and whether this comes at the expense of any of the platform's own users, is a central question in both the design and potential regulation of platform markets.

References

Aguirregabiria, Victor, and Arvind Magesan. 2020. "Identification and estimation of dynamic games when players' beliefs are not in equilibrium". *The Review of Economic Studies* 87 (2): 582–625.

Bajari, Patrick, et al. 2023. "Hedonic prices and quality adjusted price indices powered by AI". *arXiv preprint arXiv:2305.00044*.

Bronnenberg, Bart J, Jun B Kim, and Carl F Mela. 2016. "Zooming in on choice: How do consumers search for cameras online?" *Marketing science* 35 (5): 693–712.

Bulow, Jeremy, and Paul Klemperer. 1996. "Auctions versus Negotiations". *American Economic Review* 86 (1): 180–94.

Chernozhukov, Victor, et al. 2022. "Locally robust semiparametric estimation". *Econometrica* 90 (4): 1501–1535.

Cho, Sungjin, and John Rust. 2010. "The flat rental puzzle". *The Review of Economic Studies* 77 (2): 560–594.

Cogley, Timothy, Riccardo Colacito, and Thomas J Sargent. 2007. "Benefits from US monetary policy experimentation in the days of Samuelson and Solow and Lucas". *Journal of Money, Credit and Banking* 39:67–99.

Compiani, Giovanni, Ilya Morozov, and Stephan Seiler. 2023. "Demand estimation with text and image data".

Compiani, Giovanni, et al. 2022. *Online Search and Product Rankings: A Double Logit Approach*. Tech. rep. Working paper.

Davis, Andrew M, Elena Katok, and Anthony M Kwasnica. 2011. "Do auctioneers pick optimal reserve prices?" *Management Science* 57 (1): 177–192.

Doraszelski, Ulrich, Gregory Lewis, and Ariel Pakes. 2018. "Just starting out: Learning and equilibrium in a new market". *American Economic Review* 108 (3): 565–615.

Engelbrecht-Wiggans, Richard. 1987. "On optimal reservation prices in auctions". *Management Science* 33 (6): 763–770.

Erdem, Tülin, and Michael P Keane. 1996. "Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets". *Marketing science* 15 (1): 1–20.

Farrell, Max H, Tengyuan Liang, and Sanjog Misra. 2020. "Deep learning for individual heterogeneity: An automatic inference framework". *arXiv preprint arXiv:2010.14694*.

Fong, Jessica, Puneet Manchanda, and Yu Song. 2023. "How Effective is Suggested Pricing?: Experimental Evidence from an E-Commerce Platform". *Experimental Evidence from an E-Commerce Platform (November 10, 2023)*.

Foroughifar, Mohsen. 2023. "The challenges of deploying an algorithmic pricing tool: Evidence from airbnb". PhD thesis, University of Toronto (Canada).

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2016. "The slow growth of new plants: Learning about demand?" *Economica* 83 (329): 91–129.

Fox, Jeremy T, et al. 2012. "The random coefficients logit model is identified". *Journal of Econometrics* 166 (2): 204–212.

Freyberger, Joachim, and Bradley J Larsen. 2022. "Identification in ascending auctions, with an application to digital rights management". *Quantitative Economics* 13 (2): 505–543.

Gomes, Renato. 2014. "Optimal auction design in two-sided markets". *The RAND Journal of Economics* 45 (2): 248–272.

Haggag, Kareem, Brian McManus, and Giovanni Paci. 2017. "Learning by driving: Productivity improvements by new york city taxi drivers". *American Economic Journal: Applied Economics* 9 (1): 70–95.

Hasker, Kevin, and Robin Sickles. 2010. "eBay in the economic literature: Analysis of an auction marketplace". *Review of Industrial Organization* 37 (1): 3–42.

Hitsch, Günter J. 2006. "An empirical model of optimal dynamic product launch and exit under demand uncertainty". *Marketing Science* 25 (1): 25–50.

Hodgson, Charles, and Gregory Lewis. 2023. *You can lead a horse to water: Spatial learning and path dependence in consumer search*. Tech. rep. National Bureau of Economic Research.

Huang, Yufeng, Paul B Ellickson, and Mitchell J Lovett. 2020. "Learning to set prices". *Available at SSRN 3267701*.

Ichimura, Hidehiko, and Whitney K Newey. 2022. "The influence function of semiparametric estimators". *Quantitative Economics* 13 (1): 29–61.

Jullien, Bruno, Alessandro Pavan, and Marc Rysman. 2021. "Two-sided markets, pricing, and network effects". In *Handbook of Industrial Organization*, 4:485–592. 1. Elsevier.

Katkar, Rama, and David H Reiley. 2007. "Public versus secret reserve prices in eBay auctions: Results from a pokémon field experiment". *The BE Journal of Economic Analysis & Policy* 6 (2): 000010220215380637144 2.

Keller, Godfrey, and Sven Rady. 1999. "Optimal experimentation in a changing environment". *The review of economic studies* 66 (3): 475–507.

Kim, Yewon. 2020. "Customer retention under imperfect information". PhD thesis, The University of Chicago.

Klein, Benjamin, et al. 2005. "Competition in two-sided markets: The antitrust economics of payment card interchange fees". *Antitrust LJ* 73:571.

Levin, Dan, and James L Smith. 1996. "Optimal reservation prices in auctions". *The Economic Journal* 106 (438): 1271–1283.

Li, Xiaohu. 2005. "A note on expected rent in auction theory". *Operations Research Letters* 33 (5): 531–534.

List, John A. 2003. "Does market experience eliminate market anomalies?" *The quarterly journal of economics* 118 (1): 41–71.

— . 2011. "Does market experience eliminate market anomalies? The case of exogenous market experience". *American Economic Review* 101 (3): 313–317.

— . 2004. "Neoclassical theory versus prospect theory: Evidence from the marketplace". *Econometrica* 72 (2): 615–625.

Lu, Jay. 2019. "Bayesian identification: a theory for state-dependent utilities". *American Economic Review* 109 (9): 3192–3228.

Marra, Marleen. 2019. *Pricing and fees in auction platforms with two-sided entry*. Tech. rep. Sciences Po Departement of Economics.

Mela, Carl F, Jason MT Roos, and Tulio Sousa. 2023. "Advertiser Learning in Direct Advertising Markets". *arXiv preprint arXiv:2307.07015*.

Myerson, Roger B. 1981. "Optimal auction design". *Mathematics of operations research* 6 (1): 58–73.

Netzer, Oded, et al. 2012. "Mine your own business: Market-structure surveillance through text mining". *Marketing Science* 31 (3): 521–543.

Ostrovsky, Michael, and Michael Schwarz. 2016. *Reserve prices in Internet advertising auctions: a field experiment. Typescript*.

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

Quint, Daniel. 2017. "Common values and low reserve prices". *The Journal of Industrial Economics* 65 (2): 363–396.

Resnick, Paul, and Richard Zeckhauser. 2002. "Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system". In *The Economics of the Internet and E-commerce*. Emerald Group Publishing Limited.

Rochet, Jean-Charles, and Jean Tirole. 2003. "Platform competition in two-sided markets". *Journal of the european economic association* 1 (4): 990–1029.

Rothschild, Michael. 1974. "A two-armed bandit theory of market pricing". *Journal of Economic Theory* 9 (2): 185–202.

Simonsohn, Uri. 2010. "eBay's crowded evenings: Competition neglect in market entry decisions". *Management science* 56 (7): 1060–1073.

Strulov-Shlain, Avner. 2021. "More than a Penny's Worth: Left-Digit Bias and Firm Pricing". *Chicago Booth Research Paper* 19–22.

Tadelis, Steven, et al. 2023. *Learning, Sophistication, and the Returns to Advertising: Implications for Differences in Firm Performance*. Tech. rep. National Bureau of Economic Research.

Tong, Lester CP, et al. 2016. "Trading experience modulates anterior insula to reduce the endowment effect". *Proceedings of the National Academy of Sciences* 113 (33): 9238–9243.

Vickrey, William. 1961. "Counterspeculation, auctions, and competitive sealed tenders". *The Journal of finance* 16 (1): 8–37.

Wang, Zhide, and Nathan Yang. 2024. "Identification of Structural Learning Models". *Available at SSRN*.

Wang, Zhide, et al. 2024. "Retail Investment under Aggregate Fluctuations". *Available at SSRN*.

Wu, Ruiqi, Yufeng Huang, and Nan Li. 2023. "Platform Information Design and Competitive Price Targeting". *Available at SSRN*.

Xu, Boya, Yiting Deng, and Carl Mela. 2022. "A scalable recommendation engine for new users and items". *arXiv preprint arXiv:2209.06128*.

A Appendix

A1 Summary statistics and additional descriptive evidence

Table A1.1 shows selected summary statistics for the sample used in the data. To limit the effect of prediction error in estimated item values (described in more detail in 5.1), I drop all items with a standardized reserve price and standardized revenue greater than the 99th quantile of the respective variables.

Table A1.1: Summary statistics

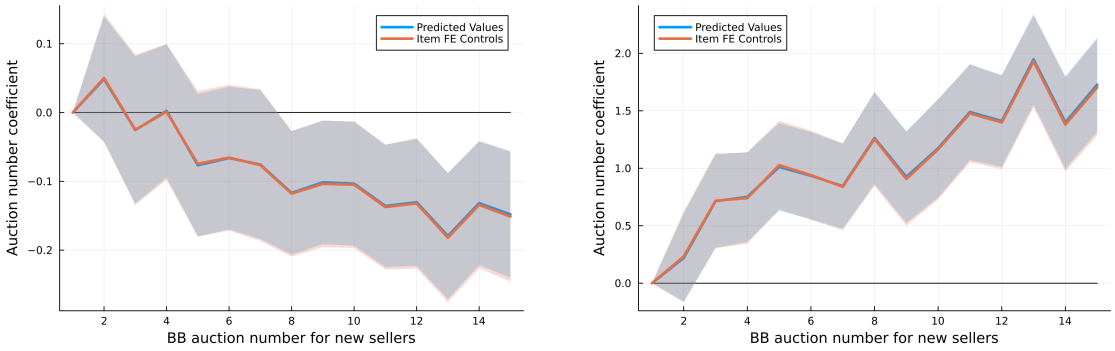
Variable	Mean	Std. Dev.	Minimum	Maximum
Minimum Bid	13.85	48.43	0.01	10,000
Reserve Price	15.39	139.4	0	68,000
Revenue	15.66	83.5	0	68,000
# Bidders	2.62	2.93	0	36
Sell	0.55	0.5	0	1
Fees	1.05	1.97	0	867.12

Notes: These statistics are from the 1,038,383 items included in the analysis data. The top 1% of the items (by standardized effective reserve price) have been removed from the analysis data.

Figure A1.1 estimates the same equation as Figure 2, but the dependent variables are the inverse hyperbolic sine (IHS) of effective reserve price and net revenue. I also restrict the sample to items where the seller has listed at least one other item with the same description, and run the regression with predicted item values and item fixed effects to compare the resulting estimates. The trends are quite similar whether using predicted item values or fixed effects, which suggests the predicted item values capture economically meaningful information. They are also similar to the trends in Figure 2, though with the caveat that the greater magnitude of the coefficients in panel (b) may be in part driven by re-listed items that were not sold the first time.[18]

[18] One limitation of the dataset is that I cannot observe seller inventories, so I cannot see how many of the identical items are true duplicates as opposed to relisting.

Figure A1.1: Time trends of variables of interest with auction experience



(a) IHS(Effective reserve) (b) IHS(Net revenue)

Notes: These regressions pool 1,639 new sellers' first 15 auctions with all auctions by 5,165 experienced sellers (defined as those with ≥47 auctions at the start of the data, which is the 75th percentile of initial experience). The sample is limited to sellers with at least 15 auctions in the data. The results are similar when using different values of T_{New}.

Figure A1.2 shows the words that most increased and decreased in their usage by new sellers in their first 15 auctions. While the trends are largely small, the words that most increased in frequency include "nr" (short for "no reserve") and "no" "reserve". This is consistent with sellers becoming more aware of the possible effect of their pricing decisions on bidder entry.

Figure A1.2: Trends in the frequency of words in item descriptions

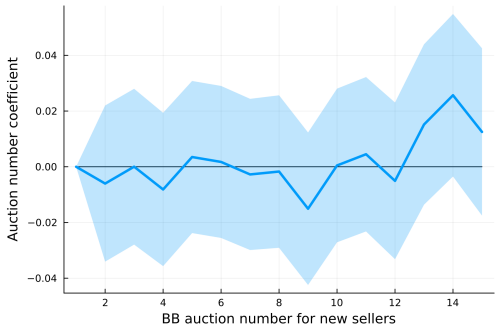


Notes: These are the 10 most positive and 10 most negative (by absolute value) coefficients when regressing $\mathbb{1}$[item contains word] on the inverse hyperbolic sine (IHS) of new sellers' auction number (among the first 15 auctions of sellers who have at least 15 auctions in the data or sellers with >75th percentile of experience at the start of the data), along with predicted item value, IHS(feedback count), feedback percentage, and seller and month fixed effects.

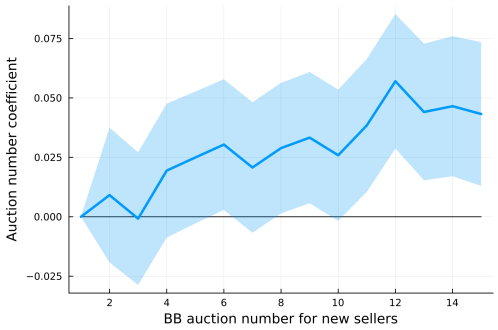
Figure A1.3 shows additional trends in non-price variables among sellers with at least 15 auctions. New sellers show some trends in the timing of an auction (panels (b) and (c)), where they favor

shorter auctions that end on weekends. As shown in panel (d), new sellers also become less likely to feature items. To ensure that estimation is computationally tractable, and since these trends are generally smaller in magnitude relative to the baseline averages of each variable, I focus attention on the choice of reserve prices.

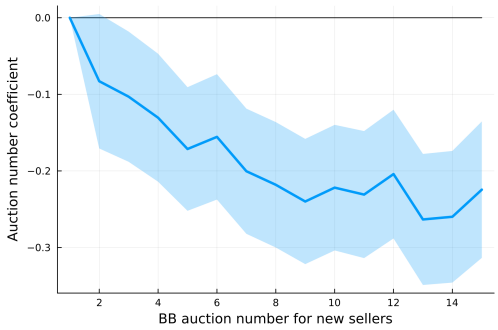
Figure A1.3: Trends in non-price variables among new sellers



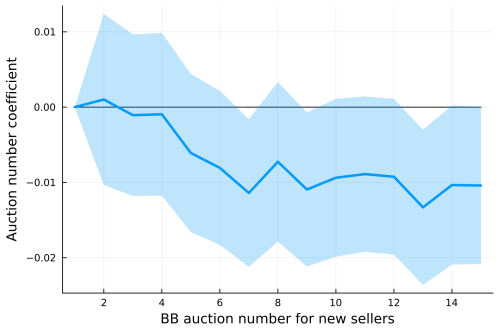
(a) End time in evening



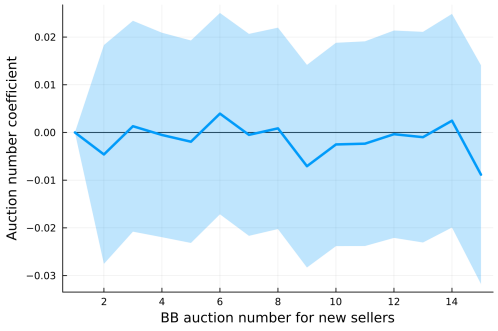
(b) End time on weekend



(c) Auction length (days)



(d) Featured item



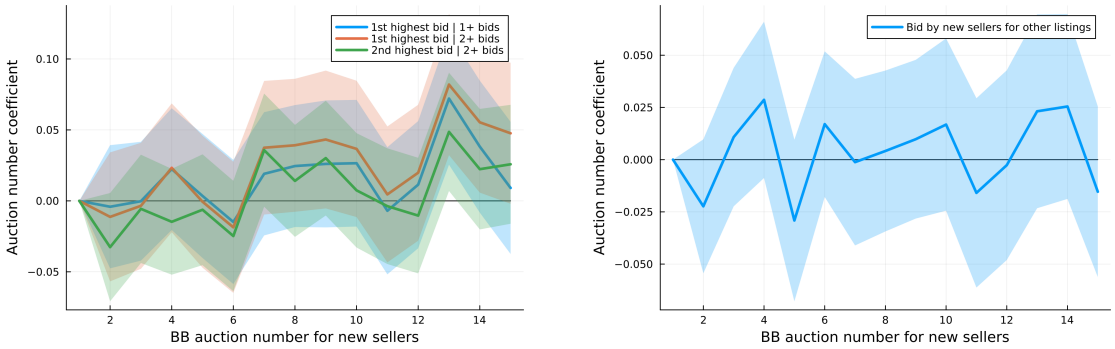
(e) Has secret reserve

Notes: These figures display coefficients of the first 15 auctions of new sellers in a regression of the various outcomes on predicted item values, seller feedback scores, inverse hyperbolic sine (IHS) of auction experience, and seller and month fixed effects. The regression is on new sellers with at least 15 auctions in the data and experienced sellers (defined as sellers with >75th percentile of experience at the start of the data).

Figure A1.4 estimates similar regressions for bids within auctions, though with the number of

observed bids in each auction additional control. Panel (a) shows the trends in first and second highest bids faced by new sellers, since these are the only bids known to reflect the first and second highest values in an ascending IPV setting. Panel (b) shows the trend in bids placed by new sellers on other listings on or before their kth listing. The trend lines for both figures are relatively noisy, with a small but statistically insignificant trend upward in panel (a). Note that panel (a) conditions on the number of observed bidders, which is a noisy measure of the true number of bidders that is unobserved due to eBay's increasing minimum bid rule.

Figure A1.4: Time trends of bids with auction experience



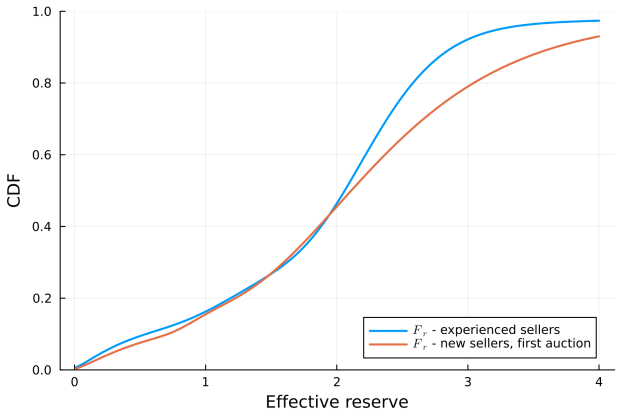
(a) Trends in bids faced by new sellers

(b) Trends in bids placed by new sellers

Notes: These regressions pool bids from 1,639 new sellers' first 15 auctions with bids from all auctions by 5,165 experienced sellers (defined as those with $\geq$47 auctions at the start of the data, which is the 75th percentile of initial experience). The sample is limited to sellers with at least 15 auctions in the data. Panel (a) restricts the sample to all auctions with at least 1 or 2 bids (as specified in the legend), and panel (b) examines bids placed by new sellers before they list their kth item. Both panels control for seller fixed effects as in section 3.2, as well as fixed effects for the number of other observed bids in each auction.

Figure A1.5 shows the estimated CDFs of the reserve prices among experienced vs. inexperienced sellers in their first auction, for all inexperienced sellers that list at least 15 items. The distributions largely align for approximately 50% of listings, but diverge for higher-reserve listings.

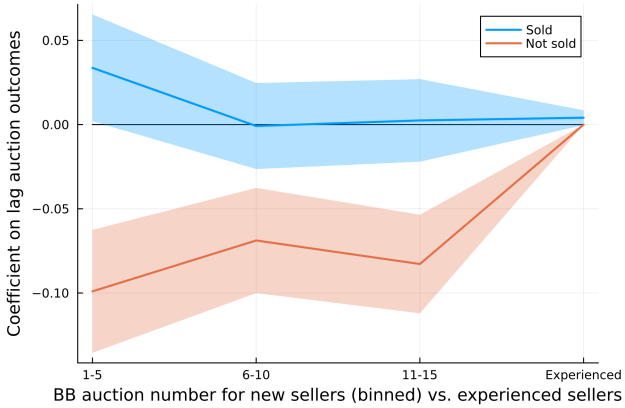
Figure A1.5: Distributions of reserve prices among new sellers in their first auction vs. experienced sellers



Notes: These are the estimated CDFs of the standardized reserve prices for new sellers in their first auction (among new sellers who have at least 15 auctions in the data) and experienced sellers (those with >75th percentile of experience at the start of the data). I use 5-component log-Normal mixture models to smooth the estimated CDFs in both cases.

Figure A1.6 plots coefficients from a regression of standardized effective reserve prices on interactions between seller experience indicators and lag auction outcome indicators. Note that only the first coefficient for the "lag sale" indicator is positive and statistically significant (though it is close), while new sellers failing to sell an item in the previous auction is more strongly (and negatively) correlated with effective reserve prices.

Figure A1.6: Coefficients for regressing lag auction sale outcomes on current effective reserve price, by experience bin



Notes: The plotted coefficients are from regressing standardized effective reserve prices on seller experience indicators interacted with lag auction sale indicators, along with 95% confidence intervals for the estimated coefficients. This regression controls for lag effective reserve price, bidder feedback, feedback score, and month fixed effects.

A2 Testing for common values

One potential concern in this setting is that Beanie Babies have a common value component that is unobservable to the bidders, opening the possibility of a winner's curse that affects bidding strategies. I regress log standardized bids on the number of bidders observed in the auction, which is a noisy but biased measure of the total number of bidders who have entered the auction. I examine both the total number of bidders and the number of serious bidders, or those whose bids exceed the reserve price, since some bids may be mechanically lower due to a low minimum bid allowing low-value bidders to submit bids that they would not ordinarily be able to submit. I also control for the standardized minimum bid and other covariates that may influence bidder arrival, namely the log item value, the inverse hyperbolic sine of the feedback count, and sellers' feedback percentage score, along with month fixed effects. These regressions are not causal, as the number of bidders is likely correlated with other drivers of bidder values.

Table A2.2: Regression of log standardized bid values on number of bidders

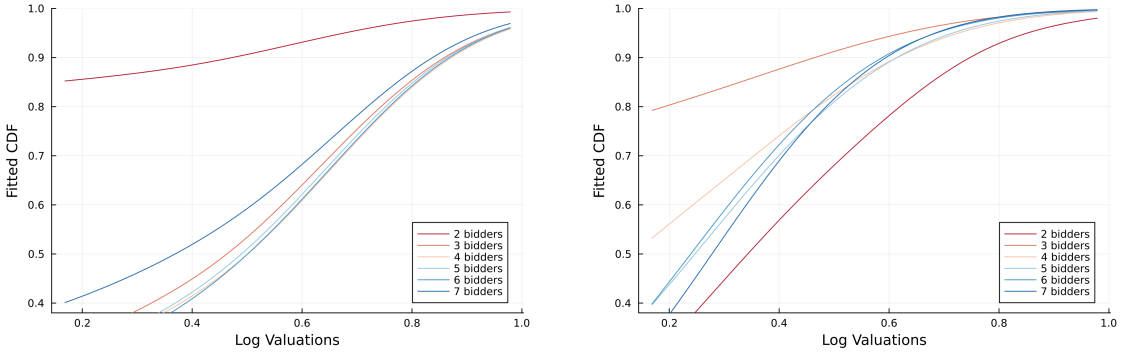
Variable	(a)	(b)
Number of bidders	-0.008	-
	(2e-4)	-
Number of serious bidders	-	0.004
	-	(1e-4)
Adjusted R^2	0.159	0.159

Notes: These regressions also control for IHS(feedback), feedback percentage, log mean item value, standardized minimum bid, month fixed effects, though these coefficients are omitted from the table. Specifications (a) and (b) differ by whether the number of bidders or the number of serious bidders (defined as bidders whose bid is equal to or greater than the reserve price) is used.

Table A2.2 shows that while there is some statistically significant correlation between the number of bidders and the log standardized bid values, its magnitude is small. I also use the two highest bids, binning by the number of either observed or serious bidders, and estimate the corresponding distributions; under common values, these distributions should be stochastically ordered so higher-bidder auctions have lower value distributions. Figure A2.7 presents the estimated distributions where I bin distributions for auctions by number of serious bidders (those with a bid above the secret reserve price, to discard unusually low bids where there is a low minimum bid) and the total number of observed bidders. In panel (a), the distributions are quite similar for all but the auctions with only 2 serious bidders. In panel (b), the distribution estimated on auctions with only 2 bidders is to the right of the remaining distributions, but the others have no clear stochastic ordering that

is indicative of common values. While not conclusive, this is suggestive evidence that independent private values is a useful modeling assumption in this setting.

Figure A2.7: Estimated distributions of bidder valuations by the number of observed bidders per auction



(a) Binned by number of serious bidders

(b) Binned by number of observed bidders

Notes: I use 5-component Gaussian mixture models to estimate the distributions for each binned number of bidders. Serious bidders are defined as those with bids above the reserve price. These distributions are plotted between the 5th and 95th percentile of the second-highest bids in the data.

A3 Proofs

Proof of Proposition 1

This closely follows the corollary in Marra (2019), though with the caveat that r is observed in this setting and thus directly impacts the Poisson mean Λ. First, denote

$$\pi_B(r \mid n, c) = \frac{1}{n} \cdot \mathbb{E}\Big[v_{n:n} - (1 + c_B^R) \max\{v_{(n-1):n}, r^*\} \;\Big|\; v_{n:n} \geq r^B\Big] \cdot (1 - F_B(r^B)^n)$$

Since F_B satisfies the strict monotone hazard rate property, Li (2005) implies that $\mathbb{E}[v_{(n+1):(n+1)} - v_{n:(n+1)}] < \mathbb{E}[v_{n:n} - v_{(n-1):n}]$; this holds when conditioning on r since r is set before the auction and does not vary with the number of bidders n that arrive. Additionally, $\frac{1}{n}(1 - F_B(r)^n) \geq \frac{1}{n+1}(1 - F_B(r)^{n+1})$,[19] so $\pi_B(r \mid n)$ is decreasing in n.

We temporarily abuse notation to write the probability mass at n given Λ as $p_n(\Lambda)$. Since arrival is Poisson, $p_n(\Lambda')$ first-order stochastically dominates $p_n(\Lambda)$ for $\Lambda' > \Lambda$. Increasing Λ therefore

[19] To see this, first note that $F_B(r) \in [0, 1]$, and the expression is equivalent to showing $(1 + \frac{1}{n})\frac{1-F_B(r)^n}{1-F_B(r)^{n+1}} \geq 1$. Let $x \in [0, 1]$, and note that $1 - x^n = (1 - x)b_n$, where $b_n \equiv \sum_{k=0}^{n-1} x^k$. Then $\frac{1-x^n}{1-x^{n+1}} = \frac{b_n}{b_n + x^n} = \frac{1}{1+\frac{x^n}{b_n}}$, and $\frac{1}{n} \geq \frac{x^n}{b_n}$, since $\sum_{k=0}^{n-1} x^k \geq \sum_{k=0}^{n-1} x^n$.

decreases $\sum_{n=1}^{\mathbf{N}_B-1} \pi_B(r \mid n,c)p_n(\Lambda)$ since $\pi_B(r \mid n)$ is monotonically decreasing in n. Thus, there exists a unique Λ that solves the zero profit condition.

Similarly, $\pi_B(r \mid n,c)$ is decreasing in r, since the measure of the set $\{v_{n:n} \geq r^B\}$, the probability of winning $1-F_B(r^B)^n$, and the winner's expected surplus are all all decreasing in r. Since higher r corresponds to strictly lower surplus, Λ is strictly decreasing in r.

Proof of Proposition 3

We differentiate the first-order condition with respect to v_0:

$$\frac{\partial^2 \Pi(v_0, r \mid b,c)}{\partial r \partial v_0} = (1-c_S^R)R_{rr}(r^*(v_0))\frac{\partial r^*(v_0)}{\partial v_0} + K_r(r^*(v_0) \mid b) + v_0 K_{rr}(r^*(v_0))\frac{\partial r^*(v_0)}{\partial v_0}$$

Rearranging, we have

$$\frac{\partial r^*(v_0)}{\partial v_0} = \frac{K_r(r^*(v_0) \mid b)}{-[(1-c_S^R)R_{rr}(r^*(v_0)) + v_0 K_{rr}(r^*(v_0))]}$$

The denominator is positive at the interior optimum because it is the negative second-order condition of the seller's profit maximization problem. Thus, $\frac{\partial r^*}{\partial v_0}$ when $K_r(\cdot \mid b) > 0$. Further, since the inverse of an increasing function is itself increasing, the virtual type function $\psi(\cdot \mid b,c)$ is monotonic increasing.

Lemma 1. Let $p_n(x \mid \delta) = \frac{1}{n!}\exp(-\Lambda(x \mid \delta))\Lambda(x \mid \delta)^n$, where $\Lambda(x \mid \delta) = \exp(\delta_1 + \delta_2 \rho(x))$ and and $\rho(x^*) = 0$. Then for each $k = 1, 2, ...,$ $\frac{\partial^k}{\partial x^k} p_n(x^* \mid \delta)$ has the form

$$\sum_{\ell=0}^{k} \delta_2^\ell \cdot h_{\ell,k}\left(n, \delta_1, \left\{\frac{\partial^t}{\partial x^t}\rho(x^*)\right\}_{t=1}^{k-\ell+1}\right)$$

where each $h_{\ell,k}$ is known.

Proof

We first show that for every $k = 1, 2, ...,$ $\frac{\partial^k}{\partial x^k} p_n(x \mid \delta)$ has the form

$$\sum_{\ell=0}^{k} \delta_2^\ell \cdot \tilde{h}_{\ell,k}\left(n, \Lambda(x \mid \delta), \left\{\frac{\partial^t}{\partial x^t}\rho(x)\right\}_{t=0}^{k-\ell+1}\right)$$

for known $\tilde{h}_{\ell,k}$. Beginning with $k = 1$, note that

$$\begin{aligned}\frac{\partial^1}{\partial x^1} p_n(x \mid \delta) &= \delta_2 \cdot p_n(x \mid \delta) \cdot [n - \Lambda(x \mid \delta)] \cdot \rho'(x) \\ &\equiv \delta_2 \cdot \tilde{h}_{1,1}(n, \Lambda(x \mid \delta), \rho'(x))\end{aligned}$$

since p_n is a known function of Λ. Now suppose $\frac{\partial^k}{\partial x^k} p_n(x \mid \delta)$ has the form above. Then using the shorthand $\tilde{h}_{j,\ell,k}$ to denote the first derivative with respect to the jth argument of $\tilde{h}_{\ell,k}$, note that by the chain rule,

$$\begin{aligned}&\frac{\partial}{\partial x} \sum_{\ell=0}^{k} \delta_2^\ell \cdot \tilde{h}_{\ell,k}\left(n, \Lambda(x \mid \delta), \left\{\frac{\partial^t}{\partial x^t}\rho(x)\right\}_{t=0}^{k-\ell+1}\right) \\ &= \sum_{\ell=0}^{k} \delta_2^\ell \cdot \left[\tilde{h}_{2,\ell,k}\left(n, \Lambda(x \mid \delta), \left\{\frac{\partial^t}{\partial x^t}\rho(x)\right\}_{t=0}^{k-\ell+1}\right) \cdot \Lambda(x \mid \delta) \cdot \rho'(x) \cdot \delta_2\right. \\ &\qquad \left. + \tilde{h}_{3,\ell,k}\left(n, \Lambda(x \mid \delta), \left\{\frac{\partial^t}{\partial x^t}\rho(x)\right\}_{t=0}^{k-\ell+1}\right)^\intercal \left\{\frac{\partial^{t+1}}{\partial x^{t+1}}\rho(x)\right\}_{t=1}^{k-\ell+2}\right] \\ &\equiv \sum_{\ell=0}^{k+1} \delta_2^\ell \cdot \tilde{h}_{\ell,k+1}\left(n, \Lambda(x \mid \delta), \left\{\frac{\partial^t}{\partial x^t}\rho(x)\right\}_{t=0}^{k-\ell+1}\right)\end{aligned}$$

Evaluating the expression above at x^*, where $\rho(x^* = 0)$ and therefore $\Lambda(x^* \mid \delta) = \exp(\delta_1)$, yields the desired result.

Proof of Proposition 5

By assumption (i), restricting attention to all sellers with the same history $\mathcal{H}$ is equivalent to restricting attention to sellers with identical beliefs b.

By Proposition 3, $\psi(\cdot \mid b, c)$ is increasing. Using the seller first-order condition, a change-of-variables can be applied to the reserve price distribution to write it in terms of the virtual type function and the known seller value distribution:

$$\mathbb{P}[r \leq x] = \mathbb{P}[v_0 \leq \psi(x \mid b, c)] = F_S(\psi(x \mid b, c))$$

Further, the selection rule for seller entry implies $\bar{v}(b, c) \geq v_0$. Taken together, the probability that a reserve price is less than or equal to x, conditional on beliefs b and cost vector c, is $F_S(\psi(x \mid$

$b, c))/F_S(\bar{v}(b,c))$. Since by assumption (ii) the entry threshold $\bar{v}(b,c)$ is known, inverting the empirical reserve price distribution $\phi(x \mid \mathcal{H})$ of sellers with history $\mathcal{H}$ yields the virtual type function:

$$\psi(x \mid b,c) = F_S^{-1}\Big[\phi(x \mid \mathcal{H}) \cdot F_S(\bar{v}(b,c))\Big]$$

for all x such that $\psi(x \mid b,c) < \bar{v}(b,c)$. By assumption (iv) this support includes a positive-measure interval including some value x^* for which $\rho(x^*) = 0$; in what follows we restrict attention to this interval.

The virtual type function $\psi(x \mid b,c)$ is proportional to the ratio of $R_r(x \mid b)$ and $K_r(x \mid b)$, and its derivatives are

$$\frac{\partial^k}{\partial x^k}\psi(x \mid b,c) = -(1-c_S^R)\frac{\partial^k}{\partial x^k}\frac{R_r(x \mid b)}{K_r(x \mid b)} = -(1-c_S^R)\sum_{\ell=0}^{k}\binom{k}{\ell}\Big(\frac{\partial^{k-\ell}}{\partial x^{k-\ell}}R_r(x \mid b)\Big)\cdot\Big(\frac{\partial^\ell}{\partial x^\ell}\big(K_r(x \mid b)\big)^{-1}\Big)$$

where by Faà di Bruno's formula

$$\frac{\partial^\ell}{\partial x^\ell}\big(K_r(x \mid b)\big)^{-1} = \sum_{t=0}^{\ell}\frac{(-1)^t \cdot t!}{(K_r(x \mid b)^{t+1})}B_{\ell,t}\Big(\frac{\partial}{\partial x}K_r(x \mid b),\ \ldots\ ,\ \frac{\partial^{\ell-t+1}}{\partial x^{\ell-t+1}}K_r(x \mid b)\Big)$$

in which $B_{\ell,t}$ are Bell polynomials. In turn, for both functions $G \in \{R, K\}$, we have

$$\frac{\partial^k}{\partial x^k}G(x \mid b) = \sum_{\ell=0}^{k}\binom{k}{\ell}\left(\sum_{n=0}^{\infty}\Big(\frac{\partial^{k-\ell}}{\partial x^{k-\ell}}G_n(x)\Big)\cdot\Big(\frac{\partial^\ell}{\partial x^\ell}p_n(x \mid b)\Big)\right)$$

Thus, the kth derivative of the virtual type function is a known function of the 0,...,kth derivatives of R_n, K_n, and p_n.

We now turn our attention to the Poisson mass function and its derivatives, which are the only arguments of ξ_k that depend on beliefs b. Expanding $\frac{\partial^k}{\partial x^k}p_n(x \mid b)$ and imposing independence between the marginal beliefs about the two parameters yields

$$\frac{\partial^k}{\partial x^k}p_n(x \mid b) = \int\int\left[\frac{\partial^k}{\partial x^k}p_n(x \mid \delta)\right]b_{\delta_1}(\delta_1)b_{\delta_2}(\delta_2)d\delta_1 d\delta_2$$

Evaluating this at x^* and applying Lemma 1, this can be expanded to yield

$$\sum_{\ell=0}^{k} \mathbb{E}_{b_{\delta_2}}[\delta_2^\ell] \cdot \hat{h}_{\ell,k}(n)$$

where $\hat{h}_{\ell,k}(n) \equiv \int h_{\ell,k}(n, \delta_1, \{\frac{\partial^t}{\partial x^t}\rho(x^*)\}_{t=1}^{k-\ell+1}) b_{\delta_1}(\delta_1) d\delta_1$ is known under the assumption that b_{δ_1} is known. Note that $\mathbb{E}_{b_{\delta_2}}[\delta_2^\ell]$ does not depend on n, so this term can be pulled out of all the infinite sums in which it appears, i.e. for $G \in \{R, K\}$

$$\begin{aligned}
\frac{\partial^k}{\partial x^k} G(x^* \mid b) &= \sum_{\ell=0}^{k} \binom{k}{\ell} \left(\sum_{n=0}^{\infty} \left(\frac{\partial^{k-\ell}}{\partial x^{k-\ell}} G_n(x^*) \right) \cdot \left(\sum_{t=0}^{\ell} \mathbb{E}_{b_{\delta_2}}[\delta_2^t] \cdot \hat{h}_{t,\ell}(n) \right) \right) \\
&= \mathbb{E}_{b_{\delta_2}}[\delta_2^k] \left(\sum_{n=0}^{\infty} G_n(x^*) \hat{h}_{k,k}(n) \right) + \sum_{\ell=0}^{k-1} \binom{k}{\ell} \left(\sum_{n=0}^{\infty} \left(\frac{\partial^{k-\ell}}{\partial x^{k-\ell}} G_n(x^*) \right) \cdot \left(\sum_{t=0}^{\ell} \mathbb{E}_{b_{\delta_2}}[\delta_2^t] \cdot \hat{h}_{t,\ell}(n) \right) \right)
\end{aligned}$$

Note the $k-1$th derivative of the virtual type function is a function of the kth raw moment of b_{δ_2}. By assumption, the $k-1$th derivative of $\psi(x^* \mid b, c)$ is invertible in the coefficient of this first term, yielding identification of the kth raw moment from the $k-1$th derivative of $\psi(x^* \mid b, c)$ and knowledge of lower-order moments. Since b_{δ_2} satisfies the Carleman condition, its moments uniquely characterize the distribution, and b_{δ_2} is identified up to the $\bar{k}$th moment.

A4 Demand estimation details

I use text data from item descriptions to estimate the average value for each item. Since item descriptions are seller-provided, there is significant variation in how words are spelled, which poses a challenge for tractably estimating item values. To address this, I manually created a crosswalk of individual words to their apparent intended word to decrease the dimensionality of the space item descriptions (e.g., replacing "beaneis" and "babys" with "beanies" and "babies"). I then constructed a dictionary of the 5,368 words that appear at least 10 times in the cleaned item descriptions. I also include indicators for each month in the dataset.

I tested multiple neural network architectures for γ via out-of-sample validation and with built-in dropout layers to find the architecture that achieved the lowest out-of-sample loss using the likelihood derived in A6. In particular, I use 80% of the sample to train the model, 10% of the sample to test out-of-sample loss during training, and the remaining 10% of the sample for out-of-

sample validation after training. I used an early stopping rule to determine the number of epochs with which to train the full model: I use the smallest number of epochs after which the testing loss fails to improve for 10 consecutive epochs. I then select the architecture with the lowest validation loss. The resulting architecture has 9,111,233 parameters; additional information on the various architectures is presented in Table A4.1. This table shows that the nonparametric specifications outperform the parametric model in the first line of the table. Larger numbers of parameters generally improve validation loss, though there are diminishing returns to increased complexity. The empirical results of the paper are similar using different architectures.

Table A4.1: Neural network architectures and performance

Model	# 1st-Layer Nodes	# Parameters	Train Loss	Test Loss	Val Loss	Epochs	R^2
1	-	5,374	2.6013	2.6165	2.6104	46	0.691
2	512	2,900,321	1.0885	1.2647	1.2680	34	0.984
2	1,024	5,782,881	0.9965	1.2264	1.2374	38	0.992
2	1,536	8,665,441	1.0006	1.2316	1.2752	30	0.989
3	512	3,083,969	1.0509	1.2638	1.2791	47	0.987
3	1,024	6,097,601	0.9144	1.1934	1.2007	43	0.993
3	1,536	9,111,233	0.8590	1.1774	1.1933	46	-

Notes: Model 1 is fully parametric, with a Gaussian distribution for log-values and no hidden nodes (i.e., log-values are modeled as a linear combination of word-specific fixed effects). Models 2 and 3 both use 5-component Gaussian mixture models, with a varying number of nodes in the first hidden layer; both have 4 hidden layers with corresponding dropout of 50%, 40%, 30%, and 20% for each. The second through fourth hidden layers contain 256, 64, and 16 nodes for model 2 and 512, 128, and 32 nodes for model 3. The number of parameters is the total number of trained parameters in each specification. The train, test, and validation loss columns denote the loss of each of the 80%, 10% and 10% samples used in comparing each of the models. The number of epochs is chosen via early stopping, since subsequent training after the listed number of epochs yields no test loss improvement for at least 10 epochs. The R^2 is taken from regressing the fitted item values from each architecture on the architecture with the lowest validation loss.

A5 Orthogonalization of the likelihood function

This section derives a method to estimate the true parameter ϑ_0 without bias due to estimation error for the nonparametric component γ_0. I denote the log-likelihood as ℓ; its derivation is shown in the following section. All relevant data for this demand-side likelihood is abbreviated as $\mathbf{D}_d$ to differentiate it from the data $\mathbf{D}$ that is used by sellers in updating their beliefs.

Denote the score function for the structural parameters ϑ_0 as

$$g(\vartheta \mid \gamma, \mathbf{D}_d) = \frac{\partial \ell(\vartheta, \gamma \mid \mathbf{D}_d)}{\partial \vartheta}$$

and note that $\mathbb{E}[g(\vartheta_0 \mid \gamma_0, \mathbf{D}_d)] = 0$. To derive a Neyman orthogonal score $g^*(\vartheta \mid \gamma, \mathbf{D}_d)$ for the average score $\mathbb{E}[g(\vartheta_0 \mid \gamma_0, \mathbf{D}_d)]$, Ichimura and Newey (2022) provide a method for finding a candidate first-stage influence function that will be added to the original score. I follow the steps in their Proposition 1 to show how this applies to a setting with both low and high dimensional parameters, where we orthogonalize with respect to the high dimensional parameter.

By way of notation, γ_0 as the true high-dimensional parameter under the true distribution function, and γ_τ is the perturbation in the direction of some alternative $\tilde{\gamma}$ (i.e., $\gamma_\tau = (1-\tau)\gamma_0 + \tau\tilde{\gamma}$). The Gateaux derivative $\frac{\partial}{\partial \tau}$ is the derivative with respect to τ from above evaluated at zero ($\tau \downarrow 0$). I assume that $\mathbb{E}[\frac{\partial}{\partial a}\frac{\partial}{\partial a}\ell(\vartheta, \gamma(X) + a \mid \mathbf{D}_d) \mid X = x] = 0$, which implies $\mathbb{E}[b(X)\ell(\vartheta, \gamma(X) + a \mid \mathbf{D}_d) \mid X = x] = 0$ for all b.

In this setting, Assumptions 1 and 2 of Ichimura and Newey (2022) are that there exist $\alpha_1(\vartheta \mid x)$ and $\alpha_2(\vartheta \mid x)$ with finite variance (and where α_2 is bounded away from zero) such that

$$\begin{aligned}
\frac{\partial}{\partial \tau}\mathbb{E}\left[\frac{\partial}{\partial \vartheta}\ell(\mathbf{D}_d, \vartheta, \gamma_\tau(X))\right] &= \frac{\partial}{\partial \tau}\mathbb{E}\left[\alpha_1(\vartheta \mid X)\gamma_\tau(X)\right] \\
\frac{\partial}{\partial \tau}\mathbb{E}\left[b(X)\frac{\partial}{\partial \gamma}\ell(\mathbf{D}_d, \vartheta, \gamma_\tau(X))\right] &= \frac{\partial}{\partial \tau}\mathbb{E}\left[b(X)\alpha_2(\vartheta \mid X)\gamma_\tau(X)\right]
\end{aligned}$$

By the chain rule and iterated expectations on the score above, we have

$$\begin{aligned}
\alpha_1(\vartheta \mid x) &= \mathbb{E}\left[\frac{\partial g(\vartheta \mid a, \mathbf{D})}{\partial a}\Big|_{a=\gamma(X)} \middle| X = x\right] \\
\alpha_2(\vartheta \mid x) &= \mathbb{E}\left[\frac{\partial^2 \ell(\vartheta \mid a, \mathbf{D})}{\partial a^2}\Big|_{a=\gamma(X)} \middle| X = x\right]
\end{aligned}$$

Writing the derivative of the likelihood with respect to the scalar output of γ as

$$\tilde{g}(\vartheta \mid \gamma, \mathbf{D}) = \frac{\partial \ell(\vartheta \mid a, \mathbf{D})}{\partial a}\Big|_{a=\gamma(X)}$$

we can combine these terms to form the orthogonal score

$$g^*(\vartheta \mid \gamma, \mathbf{D}_d) = g(\vartheta \mid \gamma, \mathbf{D}_d) - \alpha_1(\vartheta \mid x) \cdot \alpha_2(\vartheta \mid x)^{-1} \cdot \tilde{g}(\vartheta \mid \gamma, \mathbf{D}) \tag{12}$$

This orthogonal score may then be used to estimate θ while removing bias due to the plug-in estimator γ_0.

The nuisance parameters α_1 and α_2 are projections of second derivatives of ℓ onto the space of covariates X that enter γ. Unlike regression settings, each depends on the structural parameters ϑ; this is similar to Example 3 of Chernozhukov et al. (2022). As in that setting, initial estimators $\hat{\gamma}$ and $\hat{\theta}$ can be constructed using sample splitting, and then plugged into α_1 and α_2 to get predicted values and estimate the conditional expectations $\hat{\alpha}$ using nonparametric regression on X. These "plugin" estimators form the nuisance parameter $\hat{\alpha}(x) = \hat{\alpha}_1(\hat{\vartheta} \mid x) \cdot \hat{\alpha}_2(\hat{\vartheta} \mid x)^{-1}$ that yields a version of equation (12) that will be used in estimation (I omit the multiple indices used in sample splitting for ease of exposition):

$$g^*(\vartheta \mid \gamma, \mathbf{D}_d) = g(\vartheta \mid \gamma, \mathbf{D}_d) - \hat{\alpha}(x) \cdot \tilde{g}(\vartheta \mid \gamma, \mathbf{D}_d)$$

This orthogonal moment can be used as in standard GMM both to estimate ϑ without bias and construct the asymptotic variance matrix of the structural parameters. I follow the steps in Chernozhukov et al. (2022) with threefold sample splitting.

A6 Likelihood derivation: demand side

I now derive the demand side likelihood as presented in equation (10). Subscripts for j and t will be omitted where possible (since this focuses on the likelihood contribution for any given auction), as well as the dependence on parameters ϑ_d, to streamline notation.

The likelihood that no bids are observed is simply the likelihood that the highest bid (integrated over the arrival distribution) is smaller than the minimum bid. Thus, the likelihood contribution

for observing no bids is

$$\begin{aligned}\sum_{n=0}^{\infty}\frac{\Lambda^n e^{-\Lambda}}{n!}\underbrace{F_B(m)^n}_{\mathbb{P}[\text{all bids below } m]} &= \\ &= e^{-\Lambda[1-F_B(m)]}\sum_{n=0}^{\infty}\frac{[\Lambda F_B(m)]^n e^{-[\Lambda F_B(m)]}}{n!} \\ &= e^{-\Lambda[1-F_B(m)]}\end{aligned}$$

where the third equality holds since the sum is the integral of a Poisson density with mean $\Lambda F_B(m)$.

The likelihood contribution from auctions with one observed bidder uses the fact that one bid is not censored, but the other $n-1$ are. For any n bidders that arrive,

$$\begin{aligned}\sum_{n=1}^{\infty}\frac{\Lambda^n e^{-\Lambda}}{n!}\underbrace{nF_B(m)^{n-1}f_B(v^{(1)})}_{\mathbb{P}[\text{only 1 bid above } m]} &= \\ &= \frac{f_B(v^{(1)})}{F_B(m)}e^{-\Lambda[1-F_B(m)]}\sum_{n=0}^{\infty}\frac{[\Lambda F_B(m)]^n e^{-[\Lambda F_B(m)]}}{n!}n \\ &= f_B(v^{(1)})\Lambda e^{-\Lambda[1-F_B(m)]}\end{aligned}$$

where the second equality holds since the summand is 0 for $n=0$, and the last equality holds since the sum is the first moment of the Poisson distribution with mean $\Lambda F_B(m)$.

The last case, where $N_j \geq 2$, combines all possible arrival orders with at least 2 bidders. The precise values of other bidders are not necessary to construct the partial likelihood; the two highest

bids provide enough information about the arrival process.

$$\begin{aligned}
&\sum_{n=2}^{\infty} \frac{\Lambda^n e^{-\Lambda}}{n!} \underbrace{f_B(v^{(1)} \mid v^{(2)}) f_B(v^{(2)} \mid n \text{ bids})}_{\mathbb{P}[2 \text{ highest bids}]} = \\
&\quad = f_B(v^{(1)} \mid v^{(2)}) \sum_{n=2}^{\infty} \frac{\Lambda^n e^{-\Lambda}}{n!} n(n-1)(1 - F_B(v^{(2)})) F_B(v^{(2)})^{n-2} f_B(v^{(2)}) \\
&\quad = f_B(v^{(1)} \mid v^{(2)}) \frac{(1 - F_B(v^{(2)}))}{F_B(v^{(2)})} \frac{f_B(v^{(2)})}{F_B(v^{(2)})} e^{-\Lambda[1 - F_B(v^{(2)})]} \\
&\qquad\qquad \cdot \sum_{n=0}^{\infty} \frac{[\Lambda F_B(v^{(2)})]^n e^{-[\Lambda F_B(v^{(2)})]}}{n!} (n^2 - n) \\
&\quad = f_B(v^{(1)}) f_B(v^{(2)}) \Lambda^2 e^{-\Lambda[1 - F_B(v^{(2)})]}
\end{aligned}$$

where the third equality follows since $n^2 - n = 0$ for $n = 0, 1$, the fourth equality comes from the difference of the first and second raw moments of the Poisson distribution and further simplification. Combining the likelihood component of each case ($N = 0$, $N = 1$, or $N \geq 2$) with the density of the reserve price, we obtain equation (10).

Test with simulated data

For the simulations, I use a modified version of the data-generating process various architectures for the item value index γ. In the first architecture, I assume item j's log-value γ_j is known (i.e. $\gamma(\gamma_j) = \gamma_j$). In the following, I assume item values are a function of 50 indicator variables, each of which is randomly generated with average probability 0.1. I allow item values to be generated from a dense neural network (mapping from 50 indicator variables to two hidden layers of 10 nodes each before outputting to a scalar).

The number of bidders is Poisson distributed with mean $\Lambda_j = \exp(\delta_{0,1} + \delta_{0,2}\gamma_j)$, so in these simulations bidder arrival does not depend on the reserve price. I also set $m = r$ in the simulations, so the minimum bid and reserve price are the same. I model $r_j \sim \mathcal{N}(\mu_r, \sigma_r^2)$ and $v_{ij} \sim \mathcal{N}(0, \sigma_B^2)$. I report the results for this set of simulations in Table A6.1.

Table A6.1: Simulations for maximum likelihood estimation (demand)

	Regression: γ_0 on $\hat{\gamma}$		Structural Parameters				
	Intercept	Coef	σ_B^2	μ_r	σ_r^2	$\delta_{0,1}$	$\delta_{0,2}$
True values	0.0	1.0	0.5	0.25	0.5	1.0	0.5
Known γ_0							
$N = 2{,}000$	-	-	0.498	0.251	0.499	1.019	0.469
	-	-	(0.009)	(0.012)	(0.008)	(0.031)	(0.047)
$N = 10{,}000$	-	-	0.5	0.25	0.5	1.022	0.473
	-	-	(0.004)	(0.005)	(0.004)	(0.012)	(0.02)
Estimated γ_0 (uncorrected)							
$N = 2{,}000$	0.002	0.976	0.501	0.261	0.491	1.095	0.748
	(0.112)	(0.091)	(0.06)	(0.112)	(0.02)	(0.155)	(0.164)
$N = 10{,}000$	0.031	1.099	0.55	0.283	0.528	1.049	0.749
	(0.055)	(0.04)	(0.05)	(0.087)	(0.014)	(0.088)	(0.108)
Estimated γ_0 (orthogonalized)							
$N = 2{,}000$	-0.022	0.91	0.661	0.235	0.625	0.809	0.163
	(0.134)	(0.119)	(0.098)	(0.136)	(0.042)	(0.243)	(0.296)
$N = 10{,}000$	0.035	1.069	0.59	0.273	0.556	0.985	0.554
	(0.07)	(0.054)	(0.062)	(0.101)	(0.021)	(0.105)	(0.153)

Notes: Average (standard deviation) parameters are from 100 simulations for each case. Starting values were chosen randomly using Julia's Flux package initializations.

A7 Likelihood approach: supply side belief estimation

Several functions (e.g. R, K, and their derivatives with respect to r, and expectations with respect to belief densities or bidder values) involve multiple integrals and/or summations, and are therefore infeasible to compute repeatedly for all possible parameters ϑ_s (the parameters of the seller value and cost distributions) and ϑ_b (the parameters of the seller prior beliefs). I use dense neural networks to approximate several functions used in the estimation procedure.

Each neural network maps from $\mathbb{R}^M$ to $\mathbb{R}$, and each is composed of one input layer, 9 hidden layers, and one output layer. The activation function for each hidden layer is leakyrelu, and the number of nodes from input layer to output layer for each neural network is: M, 50, 100, 100, 200, 300, 3000, 300, 200, 100, 100, 50, 1. The activation function for the output layer is listed with the associated function below.

To construct each approximation, I generate datasets on which to train each neural network for various parameter values. The bounds of each variable used in the approximations are chosen to cover the empirical support of the corresponding variables where they are observed (e.g., $\delta_{j,0,1}$)

and sufficiently large support where they are unobserved (e.g. seller prior parameters). I use 99% of each dataset for training and 1% for holdout validation. I train each neural network on the respective training datasets in batches of 50 for 100 epochs before training the network on the full training dataset for 50 epochs; I exit training early if the mean square error of the holdout sample is less than 1e-5. The approximations (in bold) are constructed in the following order, with additional details listed for each approximation and the construction of the associated datasets. Each function is fit by minimizing mean square prediction error, though in some cases transformations are applied to improve relative accuracy for some parameter values.

1. *Functions with δ known.*
   (a) *Evaluate revenue and keep probabilities.* Using the estimated bidder value distribution and arrival parameters, I first evaluate R_n and K_n for $n = 0, 1, ..., 150$. I then construct 316 Chebyshev nodes in each dimension for $r \in [0.01, 6.25]$ and $\Lambda \in [-1.5, \ln(150)]$ and evaluate R and K, respectively, by taking their dot product with $\{p_n(r, \Lambda)\}_{n=0}^{150}$ evaluated at each node (I chose 316 because Floor($100{,}000^{0.5}$) = 316).
      i. **Expected revenue** R (softplus activation). Inputs: r and $\delta_{j,0,1}$.
      ii. **Keep probability** K (sigmoid activation). Inputs: r and $\delta_{j,0,1}$.
   (b) *Search for optimal reserve price.* I construct 316 Chebyshev nodes in each dimension for $v_0 \in [-1.25, 6.25]$ and $\delta_{j,0,1} \in [-1.5, \ln(150)]$ and search for the optimal reserve price r^* in 0.01, 0.02, ..., 6.25 along with the expected profit and seller surplus (profit minus outside option value) at the optimum.
      i. **Virtual type** ψ (identity activation). Inputs: r^* and $\delta_{j,0,1}$.
      ii. **Optimal reserve price** ψ^{-1} (identity activation). Inputs: v_0 and $\delta_{j,0,1}$.
      iii. **Expected surplus** $\Pi^* - v_0$ (exponential activation) Inputs: v_0 and $\delta_{j,0,1}$. Since expected surplus is positive when entry costs are zero, I minimize the mean square prediction error of the *log* expected surplus. This increases the relative accuracy of predicted expected surplus where it is small, which is important for precisely approximating the entry threshold in the next step.
   (c) *Entry threshold.* I then construct 316 Chebyshev nodes in each dimension for $c_E \in [0, 0.5]$ and $\delta_{j,0,1} \in [-5, \ln(150)]$ and search for the maximum $v_0 \in [0, 6.25]$ such that

expected surplus is weakly positive. Since expected surplus is monotonic in v_0, I use a binary search algorithm (i.e., evaluating expected surplus at the midpoint of $[0, 6.25]$, determining whether $\bar{v}$ lies above or below the midpoint, and iterating with additional intervals) until the difference in successive iterations is less than 0.01.

i. **Entry threshold** $\bar{v}$ (identity activation). Inputs: c_E and $\delta_{j,0,1}$.

2. *Functions with δ unknown.* I approximate the following functions for the full model with 2-dimensional unknown parameter δ_0 and the arrival coefficient model where only $\delta_{0,2}$ is unknown. The prior parameters are the mean $\{\mu_{0,1}, \mu_{0,2}\}$, standard deviations $\sigma_{0,1}$ and $\sigma_{0,2}$, and correlation $\tilde{\rho}_0$. Due to the higher dimensionality due to the belief parameters, I sample 100,000 input values for each step rather than constructing a grid of Chebyshev nodes. Unless otherwise specified, each input is sampled uniformly on the stated support.

   (a) *Search for optimal reserve price.* I sample $v_0 \sim F_S$ (bounded on $[0.0, 6.25]$), $\delta_{j,0,1} \sim U[-1.5, \ln(150)]$, $\mu_{0,2} \sim U[-0.75, 0.75]$, $\sigma_{0,i} \sim U[0.01, 1.0]$ for $i = 1, 2$, and $\tilde{\rho}_0 \sim U[-0.9, 0.9]$ ($\sigma_{0,2}$ and $\tilde{\rho}_0$ are only sampled for the full model). I then search for the optimal reserve price r^* in 0.01, 0.02, ..., 6.25 along with the expected profit and seller surplus (profit minus outside option value) at the optimum, integrating over sellers' belief distributions to do so using Gauss-Hermite quadrature with 5 points in each dimension.[20]

      i. **Virtual type** ψ (identity activation). Inputs: r^*, $\delta_{j,0,1}$, $\mu_{0,2}$, $\sigma_{0,1}$, $\sigma_{0,2}$, $\tilde{\rho}_0$.

      ii. **Optimal reserve price** ψ^{-1} (identity activation). Inputs: v_0, $\delta_{j,0,1}$, $\mu_{0,2}$, $\sigma_{0,1}$, $\sigma_{0,2}$, $\tilde{\rho}_0$.

      iii. **Expected surplus** $\Pi - v_0$ (exponential activation). Inputs: v_0, $\delta_{j,0,1}$, $\mu_{0,2}$, $\sigma_{0,1}$, $\sigma_{0,2}$, $\tilde{\rho}_0$. As with step 2(c), I minimize mean square prediction error using the log expected surplus.

   (b) *Entry threshold.* I sample $c_E \sim U[0, 1]$ (bounded on $[1.25, 6.25]$), $\delta_{j,0,1} \sim U[-1.5, \ln(150)]$, $\mu_{0,2} \sim U[-0.75, 0.75]$, $\sigma_{0,i} \sim U[0.01, \sqrt{0.5}]$ for $i = 1, 2$, and $\tilde{\rho}_0 \sim U[-0.95, 0.95]$ ($\sigma_{0,2}$ and $\tilde{\rho}_0$ are only sampled for the full model). I use the same binary search algorithm as in 1(c) to find $\bar{v}$.

      i. **Entry threshold** $\bar{v}$ (identity activation). Inputs: c_E, $\delta_{j,0,1}$, $\mu_{0,2}$, $\sigma_{0,1}$, $\sigma_{0,2}$, $\tilde{\rho}_0$.

[20] This procedure yields candidate quadrature nodes $\ln(\Lambda)$ at which I evaluate the expected revenue and keep probaiblities. For $\ln(\Lambda) \in [-5, \ln(150)]$ I use the evaluated R_n and K_n as in step 1(a) above, and for $\ln(\Lambda) \notin [-5, \ln(150)]$ I extrapolate using the fitted values of R and K.

(c) *Updating process.* I sample 100,000 draws of $v_0 \sim F_S$ (truncated on $[0, 6.25]$), $Z'\lambda \sim \mathcal{N}(2, 1)$ (truncated on $[-1.5, \ln(150)]$), $\mu_{0,i} \sim U[-0.75, 0.75]$ for $i = 1, 2$, $\sigma_{0,i} \sim U[0.01, 1.0]$ for $i = 1, 2$, and $\tilde{\rho}_0 \sim U[-0.9, 0.9]$ ($\sigma_{0,2}$ and $\tilde{\rho}_0$ are only sampled for the full model). I sample $r^* \sim 0.2\mathcal{N}(0.25, 0.5^2) + 0.7\mathcal{N}(1.25, 0.5^2) + 0.1\mathcal{N}(3.0, 1.5^2)$ (a Gaussian mixture), truncated on $[0, 6.25]$.[21] I also simulate profit signals ϵ as in equation (6) from $\mathcal{N}(0, \sigma^2_\Pi(r))$, where $\sigma^2_\Pi(r) = \exp(\tilde{\sigma}_{0,\Pi} + \tilde{\sigma}_{1,\Pi} r)^2$ is estimated using the empirical difference between profit signals and expected profit among experienced sellers.

I compute the updated beliefs via a Laplace approximation to the posterior for each observation. I first use a resilient backpropogation algorithm to search for the new maximum a posteriori estimates $\delta^*_{j,0,1}$ and $\mu^*_{0,2}$ given prior parameters, r^*, and ϵ. I then evaluate the updated covariance parameters ($\sigma^*_{0,1}$, $\sigma^*_{0,2}$, and $\tilde{\rho}^*_0$) by evaluating the Hessian of the posterior evaluated at the maximum a posteriori estimate. I drop all evaluations for which this algorithm returns either posterior parameters outside the simulation bounds or an invalid covariance matrix, and evaluate the neural networks using the resulting posterior parameters.

i. **Updating means $\mu_{0,1}$ and $\mu_{0,2}$** (identity activation). Inputs: ϵ, $\delta_{j,0,1}$, $\rho(r^*)$, $\mu_{0,2} \cdot \rho(r^*)$, $\sigma_{0,1}$, $\sigma_{0,2}$, $\tilde{\rho}_0$. Instead of approximating each updated parameter directly, I minimize mean square prediction error for the standardized difference $(\mu^*_{0,i} - \mu_{0,i})/\sigma_{0,i}$ for $i = 1, 2$; this ensures more accurate update steps for the sellers' mean parameters when beliefs are more highly concentrated.

ii. **Updating standard deviations $\sigma_{0,1}$ and $\sigma_{0,2}$** (identity activation). Inputs: ϵ, $\delta_{j,0,1}$, $\rho(r^*)$, $\mu_{0,2} \cdot \rho(r^*)$, $\sigma_{0,1}$, $\sigma_{0,2}$, $\tilde{\rho}_0$. Instead of approximating each updated parameter directly, I minimize mean square prediction error for the ratio $(\sigma^*_{0,1})/\sigma_{0,i}$ for $i = 1, 2$; this ensures more accurate update steps for the sellers' covariance parameters when beliefs are more highly concentrated.

iii. **Updating posterior correlation $\tilde{\rho}$** (identity activation). Inputs: ϵ, $\delta_{j,0,1}$, $\rho(r^*)$, $\mu_{0,2} \cdot \rho(r^*)$, $\sigma_{0,1}$, $\sigma_{0,2}$, $\tilde{\rho}_0$. Instead of approximating the updated parameter directly, I minimize mean square prediction error for the difference $\tilde{\rho}^*_0 - \tilde{\rho}_0$ for $i = 1, 2$.

[21] I chose this after experimenting with various distributions from which to sample; this density over-samples for relatively low draws of the reserve price where R can vary significantly with r^* depending on the value of $\delta_{0,2}$.

I chose both $\rho(r^*)$ and $\mu_{0,2} \cdot \rho(r^*)$ as inputs after experimenting with various architectures.