

Identification and Estimation of a Bidding Model for Electronic Auctions*

Brent R. Hickman[†]

Department of Economics, University of Chicago

Timothy P. Hubbard

Department of Economics, Colby College

Harry J. Paarsch

Department of Economics, University of Central Florida

Abstract

Because of discrete bid increments, bidders at electronic auctions engage in shading, instead of revealing their valuations, which would occur under the commonly assumed second-price rule. We demonstrate that mis-specifying the pricing rule can lead to biased estimates of the latent valuation distribution and then explore identification and estimation of a model with a correctly-specified pricing rule. A further challenge to econometricians is that only a lower bound on the number of participants at each auction is observed. From this bound, however, we establish nonparametric identification of the arrival process of bidders—the process that matches potential buyers to auction listings—which then allows us to identify the latent valuation distribution without imposing functional-form assumptions. We propose a computationally tractable, sieve-type estimator of the latent valuation distribution based on B-splines, and then compare two parametric models of bidder participation, finding that a generalized Poisson model cannot be rejected by the empirical distribution of observables. Our structural estimates enable us to explore information rents and optimal reserve prices on eBay.

Keywords: eBay; electronic auctions; bid increments; pricing rule.

JEL Classification: D44, C72.

*Previous drafts of this paper circulated under the title, “Investigating the Economic Importance of Pricing-Rule Mis-Specification in Empirical Models of Electronic Auctions.” Additional material can be found in the Online Appendix provided on the journal’s website. We thank Dan Quint, Ken Judd, and Stéphane Bonhomme as well as the anonymous referees for feedback that led to significant improvements in both the substance and the presentation of the paper. We also thank Steven Mohr, to whom we are deeply indebted for the in-depth technical support that made the complex computations in this project possible. For useful feedback, we thank Jorge Balat, Michael Dinerstein, Joachim Freyberger, Srihari Govindan, Tong Li, Stephen Ryan, and Matthew Allen; seminar participants at the College of William and Mary, the University of Chicago, the University of Oklahoma, Wake Forest University, and Vanderbilt University; and conference participants at the eleventh annual International Industrial Organization Conference and the second annual Conference on Auctions, Competition, Regulation, and Public Policy. Hickman thanks the Department of Economics at the University of Wisconsin–Madison for hospitality during the period when much of this project was in development. During our research, we benefited immensely from the advice and insights of our friend and colleague Che-Lin Su, who passed away too soon on 31 July 2015. We dedicate the paper to his memory.

[†]Corresponding author: 1126 E 59th Street; Chicago, IL 60637.

1 Introduction

During the past two decades, electronic auctions (EAs) have become important market mechanisms in the world economy—reducing frictions by bringing together large numbers of buyers and sellers of an incredibly diverse range of goods. Between the first quarter of 2011 and the fourth quarter of 2013, the leading host of EAs, eBay, averaged \$4.95 billion in sales per quarter from auctions alone.¹ EAs are of interest to economists not just because of these sales volumes, but also because the sites are data-rich environments within which fundamental economic phenomena can be studied—for example, market design, market frictions, private information rents, and price discovery. Typically modeled along the lines of traditional auction formats, EAs also display other features that derive from their online implementations and present unique challenges for empirical researchers. An especially important feature is eBay’s proxy bidding procedure: a bidder reports a number to the server—kept private until the bidder is outbid—that represents the maximal amount she authorizes the server to bid, and in turn, the server acts on the bidder’s behalf to increase her standing offer up to her reported maximum. For this reason, researchers have typically modeled EAs as some variant of a second-price auction (SPA) because the winner (usually) pays a price linked to the second-highest proxy bid for the object on sale. Important qualifications apply to this procedure: when the EA software overtakes one bidder’s maximum bid on behalf of another, it forces the new lead bidder, whenever possible, to surpass the former lead bidder by a commonly known, discrete amount Δ , referred to as the *bid increment*.² Because the EA software forces the lead bidder to surpass the second highest bid by Δ , a necessary exception occurs when the top two proxy bids are within Δ of one another, because a jump in the full amount Δ would then surpass the high bidder’s maximum authorized bid. In this eventuality, the price is set at the value of the high bid, that is, the winner pays her bid as in the case of a first-price auction (FPA). Thus, transacted prices on eBay (and many other electronic bidding sites) are actually determined by a hybrid second-price/first-price rule. In other words, a positive probability always exists that the highest bid is the sale price, specifically, the winner pays her bid.

Because EAs constitute one of the largest applications of auctions in history, understanding the impact of this nonstandard pricing rule is an important and useful endeavor. In fact, several other real-world situations exist where minimum bid increments have been noted, but ignored, when the pricing rule is modeled. For example, McAfee and McMillan (1996) noted that the FCC spectrum auctions have minimum bid increments that are typically five to ten percent of the current price of a given license. We use the simpler, single-unit EA as an opportunity to account formally for the complications minimum bid increments introduce, and to highlight the consequences of ignoring this feature of the pricing rule.

Hickman (2010) derived the theoretical implications of the EA pricing rule and found that dominant strategies cease to exist in private-valued auction games, unlike at SPAs. In equilibrium, participants engage in bid shading, similar to that at FPAs, because with positive probability the winner’s bid determines the transaction price. Hickman also presented empirical evidence that even seemingly small bid increments can play a nontrivial role in affecting bidder behavior. Using data from a sample of laptop auctions held on eBay, where the average sale price was roughly \$300 and Δ was \$5, Hickman found that the final sale price was determined by the winner’s bid nearly one-quarter of the time.

Because the structural econometric analysis of data from auctions hinges critically on the form of the theoretical equilibrium, bid shading at EAs implies that the traditional SPA approach leads to an incorrectly specified bidding model, which can induce significant bias in the estimated latent valuation distribution,

¹These data were downloaded on 25 August 2015 from <http://www.statista.com/statistics/242267/ebays-quarterly-gross-merchandise-volume-by-sales-format>.

²Originally, bid increments may have been implemented as a security measure against cyber attacks. For example, setting an increment of $\Delta = \$2.50$ would increase considerably the cost of automated, high-frequency bid submission relative to the case where the price is allowed to adjust by as little as a penny. Bid increments are fixed by online auction houses, and openly advertised to market participants prior to bidding, in other words, common knowledge.

which in turn can bias inference concerning various other questions relating to bidders' private values, including information rents, optimal reservation price, or predicted revenue changes from additional bidder at an auction.³

In this paper, based on the theory developed by Hickman (2010), we develop a complete structural model of EA bidding with a correctly specified pricing rule. The first contribution of our paper is to demonstrate the empirical relevance of the pricing rule mis-specification using a simple Monte Carlo exercise calibrated to resemble a host of realistic scenarios at EAs.

Our second contribution is to develop a fully structural econometric model designed to be implemented using actual eBay data with possible bid shading as part of the equilibrium. We demonstrate that the model is nonparametrically identified by commonly available observables at EAs. Along the way, we also overcome another, independent challenge inherent in data from EAs: In many real-world settings, it is notoriously difficult to measure precisely participation at an auction, that is, to get an accurate measure of how many bidders competed for the object on sale; eBay auctions (and EAs in general) are no different. Within the structural econometrics of auctions literature, considerable research exists addressing this problem, but we take a novel approach to identification by modeling explicitly the process through which some bidder identities are revealed to the econometrician and others are not.⁴ Our model also takes into account the fact that on electronic platforms bidders and econometricians share the same informational disadvantage: neither observes the total number of potential bidders watching at home an object for sale with intent to bid on it.

In short, we model the number of competitors N as a random variable from the bidders' perspectives and demonstrate that its distribution is nonparametrically identified by the observable lower bounds at each auction. Even if the EA pricing rule were simply the second-price rule, this problem of imperfectly observed participation would still represent a significant roadblock to empirical work. Previous researchers, such as Song (2004a,b), developed methods to side-step the need to observe N when identifying the latent valuation distribution. However, without knowing N (or at least its distribution), many counterfactual simulations based on private-values estimates remain impossible. Ours is the first paper to solve this problem: we establish nonparametric identification of a bidding model with stochastic N under the available observables, which in principle makes counterfactual simulations possible. For tractability, our estimator relies on a parametric, generalized Poisson model of the bidder arrival process. We find in our eBay data strong evidence that there is little additional gain from a fully nonparametric estimator.

Our third contribution is to propose a novel sieve-type estimator of the latent valuation distribution, which incorporates B-splines for flexibility and Galerkin methods for tractability. Our choice of B-splines permits us to match a crucial prediction of the theory: at points where bid increments change discretely from one value to another, the equilibrium bid function contains an abrupt jog. Other popular estimation methods, such as kernel smoothing or orthogonal polynomials, either cannot accommodate the needs of the identification strategy or require an inordinate amount of numerical complexity to do so. B-splines provide a remarkable degree of functional flexibility and numerical stability, which we demonstrate below.

³One might also wonder about the possible revenue implications of the bid increment Δ itself. If, however, bidder valuations are independent and identically distributed; bidders are risk neutral; and the expected payment to a bidder with a value of zero is zero; then in a symmetric, increasing equilibrium of an auction game under a variety of auction formats and pricing rules, the seller will earn the same expected revenue—the well-known revenue equivalence proposition, first demonstrated by Vickrey (1961) for uniformly-distributed valuations, and later proven in general by Myerson (1981) as well as Riley and Samuelson (1981). Thus, even though the EA pricing rule changes bidder behavior, it will not change expected revenue when the above assumptions hold. Whether these assumptions are empirically relevant is beyond the scope of this paper. Instead, we focus on developing and implementing an empirical bidding model that takes into account the pressure to shade bids induced by positive increments at EAs.

⁴See Hickman, Hubbard, and Sağlam (2015) for a complete survey of previous methods for coping with an imperfectly observed number of bidders.

Our estimator also incorporates Galerkin methods—commonly used in physical sciences applications—to solve the differential equation defined by a representative bidder’s first-order condition. In ours and many other contexts, structural estimation involves reverse-engineering the distribution of a latent random variable (*e.g.*, private values) from an observable distribution (*e.g.*, bids), with an equilibrium mapping linking the two together (*e.g.*, Bayes-Nash bidding strategies). This situation often necessitates a *nested fixed-point* approach to estimation, where the equilibrium equations are re-evaluated each time model parameters are updated by the solver that optimizes the empirical criterion function. Our use of Galerkin methods reduces substantially the computational costs of implementing the estimator in a way that is similar in spirit to the *MPEC* approach pioneered by Su and Judd (2012). Rather than sequentially solving a new differential equation every time the empirical objective function is evaluated—which may occur thousands of times—our proposed method essentially requires solving a bidder’s first-order conditions *only once* by simultaneously optimizing the parameters of both the equilibrium mapping and the underlying structural primitives. Although we tailor our estimator to the particular differential equations governing equilibrium conditions for bidding in electronic auctions, our general method combining B-splines and Galerkin methods may be applicable in broader structural contexts outside of auctions as well.

We use our structural estimates to explore counterfactual simulations of the model that inform us about two aspects of online market design: we quantify the distribution of information rents to winners and explore optimal reserve prices for sellers on eBay. First, we find that despite the large numbers of buyers participating in the market winners on eBay garner significant information rents. Second, we resolve the puzzle of why the majority of sellers on eBay set reserve prices at zero: although the expected revenue curve (as a function of reserve price) has an interior global optimum, the difference between expected revenue at a reserve price of zero and the optimal expected revenue involves mere pennies.

The remainder of the paper has the following structure: in order to place our research in context, in the following subsection we briefly review the related literature, while in Section 2, we present the basic theory of bidding, identification, and estimation for a simplified version of the model—one with a fixed, known N . This simplification permits a parsimonious Monte Carlo exercise to explore the impact of pricing rule mis-specification. The results of this mis-specification analysis are presented in Section 3, while in Section 4, we extend the basic bidding model, identification results, and estimation method to handle more realistic aspects of real-world EAs, such as randomness in N , limited data availability, and nonconstant Δ . In Section 5, we implement our estimator using a sample of data from eBay laptop auctions and discuss our empirical results, while in Section 6, we present a summary of and the conclusion from our research. In the Online Appendix, we provide Monte Carlo results, additional figures, results from robustness checks, and a brief primer on B-spline.

1.1 Related Literature

Lucking-Reiley (2000) provided a guide to EAs for economists. In surveying Internet auctions in 1998, he found that 121 of 142 sites used the ascending-price format. Lucking-Reiley recognized the use of proxy bidding, but did not distinguish between SPAs and EAs. Similarly, Bajari and Hortaçsu (2004) discussed the prevalence of the ascending-price format at online auctions. Although they noted that bid increments existed, and discussed the eBay case explicitly, Bajari and Hortaçsu implicitly assumed that EAs were SPAs. Other empirical researchers (such as Roth and Ockenfels (2002), Adams (2007) as well as Zeithammer and Adams (2010)) recognized that bid increments exist, but nevertheless assumed that EAs are SPAs. Given that the ascending-price format is the predominant one at EAs, a natural first step in analyzing these auctions is to examine how such auctions are typically modeled.

Haile and Tamer (2003) presented what remains to be the most flexible open-auction format method available. Two important points should be kept in mind when comparing their research to ours: First, in

our case, the EA rules are clear enough that we can use the structure to identify the model fully; Haile and Tamer focused on partial identification, in part because their assumptions are quite flexible and tailored to fitting a host of settings that may deviate from the canonical “clock” model of an ascending-price auction. Second, Haile and Tamer considered an environment within which bidders are only permitted to tender bids from a discrete grid—for example, when the auctioneer raises the price in discrete jumps and asks for audience volunteers to pay the proposed price, or bidders cry out discrete bid jumps. At EAs, however, bidders may choose their own bids from a continuum (or at least a very fine grid), and the fixed increment instead governs the probability that they will pay their own bid on winning.

We investigate the economic and statistical importance of mis-specifying the EA pricing rule as a second-price rule. In the next section, we present a simple game-theoretic model of bidding at EAs where the number of competitors is known and constant across different auctions. We begin with the most austere model possible—to focus attention on the effects of neglecting a minimum bid increment when considering the pricing rule at EAs. Hortaçsu and Nielsen (2010) emphasized how crucial using the correct mapping between valuations and bids is to identification. We demonstrate that within the EA model this mapping is nonparametrically identified, and present a sensitivity analysis, describing the effects of mis-specification. We find that it can lead to significantly biased point estimates and policy prescriptions, for example, concerning optimal auction design. We then consider model identification and estimation under more empirically realistic assumptions when the true number of competitors is random and unknown to the econometrician.

2 Baseline Model: Fixed Participation

Consider a seller who seeks to divest a single object at the highest price. There are N potential buyers, each of whom has a privately known value for the object for sale. Each potential buyer views the private value of each of her competitors as a random variable V , which represents an independent draw from the cumulative distribution function (CDF) $F_V(v)$ that is twice continuously differentiable, and has a strictly positive probability density function (PDF) $f_V(v)$ on a compact support $[\underline{v}, \bar{v}]$, where \underline{v} weakly exceeds zero.⁵ This information is common knowledge to all bidders. For now, we also assume that the number of potential buyers N is fixed and known. In Section 4, we present a more realistic model of N when we establish identification and estimation using actual eBay data. For the current purposes, however, it will help to isolate the specific implications of pricing rule mis-specification if we begin with this simple model of auction participation. This environment is often referred to as the symmetric *independent private-values paradigm* (IPVP).

To reduce clutter, it will be useful later to let Z denote $V_{(1:N-1)}$, the highest valuation of $(N - 1)$ draws from the urn $F_V(v)$ —in symbols, $Z = \max \{V_1, V_2, \dots, V_{N-1}\}$. The random variable Z represents the highest private value of a bidder’s $(N - 1)$ opponents at the auction. Given that valuations are distributed independently and identically, the CDF and PDF of Z are $F_Z(z) = F_V(z)^{N-1}$ and $f_Z(z) = (N - 1)F_V(z)^{N-2}f_V(z)$, respectively.

The highest bidder is declared winner, but the price is determined by a special hybrid pricing rule, where the winner pays the smaller of either her bid or the second-highest bid plus a commonly-known bid increment Δ . Below, we characterize the symmetric equilibrium bid function and investigate the importance of the error that obtains from assuming the EA is an SPA.

⁵We use uppercase Roman letters to denote random variables and lowercase Roman letters to denote realizations of these random variables.

2.1 Deriving the EA Equilibrium Bid Function

Hickman (2010) showed that a unique, monotone, pure-strategy, Bayes–Nash equilibrium exists within the IPVP. Denote by $\beta(v)$ the symmetric bid function at an EA with bid increment Δ . At an EA two scenarios can determine the sale price: in the first, the highest losing bid is within Δ of the winner’s bid and she therefore pays her own submitted bid (the first-price rule is used); in the second, the two top bids are further apart, and the winner pays the highest losing bid plus Δ (a second-price rule is used). Note that the actual bid increment only enters the pricing equation directly in the event that a second-price rule is triggered, in which case a bidder’s own bid does not affect the sale price. Therefore, when optimizing bids on the margin, players only consider how the bid increment controls the threshold (below their own bid) at which a first-price rule is triggered, where their own bid determines sale price. With this in mind, we define a *threshold function*, $\tau(b)$, as follows:

$$\tau(b) = \begin{cases} \underline{v}, & b \leq \underline{v} + \Delta_1, \\ b - \Delta_1, & \underline{v} + \Delta_1 \leq b. \end{cases} \quad (1)$$

Suppose that a bid of b is the highest, and therefore the winner. Intuitively, whenever b is less than $\underline{v} + \Delta_1$ the first-price rule is triggered with certainty; otherwise, this only occurs when the maximum competing bid is within Δ_1 of b .

To derive the optimal bidding strategy, let $\mathbf{B} = (B_1, B_2, \dots, B_N)$ denote the vector of all bids, and let B_{-n}^{max} denote the maximum bid among player n ’s opponents. Bidder n ’s optimization problem can be expressed as

$$\max_{b \in \mathbb{R}} \left\{ (v_n - b) \Pr[\tau(b) < B_{-n}^{max} \leq b] + [v_n - \mathbb{E}(B_{-n}^{max} | B_{-n}^{max} \leq \tau(b)) - \Delta] \Pr[B_{-n}^{max} \leq \tau(b)] \right\}. \quad (2)$$

The first term in the sum corresponds to winning the EA under a first-price rule, that is, when bidder n ’s bid and the second-highest bid are within Δ of each other. The second term corresponds to winning the EA under a second-price rule, that is, b_n exceeds the second-highest bid by at least Δ , and bidder n pays Δ more than the second-highest bid. We can express the probability of winning under the first-price rule given equilibrium bid function $\beta(v)$ with inverse $\beta^{-1}(b)$ as $\Pr[\tau(b) < B_{-n}^{max} \leq b] = \Pr[\beta^{-1}[\tau(b)] < Z \leq \beta^{-1}(b)] = F_Z[\beta^{-1}(b)] - F_Z[\beta^{-1}[\tau(b)]]$, and under the second-price rule as $\Pr[B_{-n}^{max} \leq \tau(b)] = \Pr[Z \leq \beta^{-1}[\tau(b)]] = F_Z[\beta^{-1}[\tau(b)]]$. The conditional expectation in the second term of expression (2) is

$$\mathbb{E}[B_{-n}^{max} | B_{-n}^{max} \leq \tau(b)] = \frac{\int_{\underline{v}}^{\beta^{-1}[\tau(b)]} \beta(u) f_Z(u) du}{F_Z[\beta^{-1}[\tau(b)]]}.$$

Maximizing a bidder’s expected surplus on the interval $[v_\Delta, \bar{v}]$ yields a boundary value problem which defines the function $\beta(v)$:

$$\beta'(v) = \frac{[v - \beta(v)] f_Z(v)}{F_Z(v) - F_Z[\beta^{-1}(\tau[\beta(v)])]}, \quad \text{and} \quad \beta(\underline{v}) = \underline{v}. \quad (3)$$

Note that for bids $\beta(v) \leq \underline{v} + \Delta$, (3) reduces to the standard first-price equilibrium differential equation, since $F_Z(\underline{v}) = 0$. This is for good reason since anyone planning to bid below that point will realize that for them the first-price rule will be triggered with certainty in the event that they win. Note also that the derivative in (3) exists everywhere since $\tau(\cdot)$ is a continuous function. This formulation, which was developed by Hickman (2010), is not particularly helpful to empirical researchers because the differential equation does not admit a closed-form solution. Fortunately though, after applying the inverse function

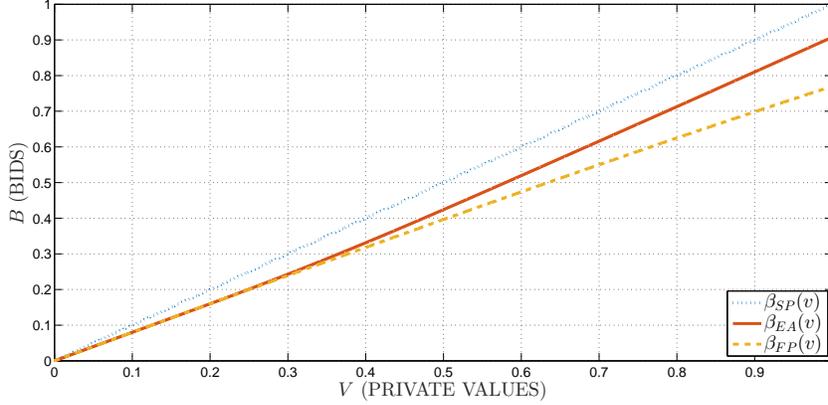


Figure 1: SPA, EA, and FPA Equilibria

theorem this differential equation can be re-cast solely in terms of the equilibrium inverse-bid function, so equation (3) can be rewritten as

$$\frac{d\beta^{-1}(b)}{db} = \frac{F_Z[\beta^{-1}(b)] - F_Z[\beta^{-1}(b - \underline{v} - \Delta)]}{[\beta^{-1}(b) - b] f_Z[\beta^{-1}(b)]} \quad (4)$$

where we have used the fact that $\beta(v)$ equals b and $\beta^{-1}(b)$ equals v . Thus, the EA equilibrium is characterized by a piecewise differential equation for optimal bidding on $[0, v_\Delta)$ and $[v_\Delta, \bar{v}]$.⁶

As an example of how equilibrium behavior at an EA differs from the standard SPA and FPA, we depict in Figure 1 the bid functions under the three pricing rules.⁷ In this figure, and below, we refer to the equilibrium bid function at an EA as $\beta_{EA}(v)$, the equilibrium bid function at an SPA as $\beta_{SP}(v) = v$, and the equilibrium bid function at an FPA as $\beta_{FP}(v) = \mathbb{E}[Z|Z < v]$. The important point to note is that the EA equilibrium bid function lay weakly between the SPA and the FPA bid functions. In fact, Hickman (2010, Theorem 3.6) showed that the EA model nests the SPA and FPA as special cases in the sense that, as $\Delta \rightarrow 0$, we get $\beta_{EA}(v) \rightarrow \beta_{SP}(v)$ uniformly, and as $\Delta \rightarrow \mathbb{E}[Z]$, we get $\beta_{EA}(v) \rightarrow \beta_{FP}(v)$ uniformly.

2.2 Nonparametric Identification and Estimation with Fixed N

The EA model is identified nonparametrically and can be estimated using a simple modification to the approach proposed by Guerre et al. (2000, GPV) for models of first-price, sealed-bid auctions. To see why, let G_B denote the cumulative distribution function of equilibrium bids, and note that $G_B(b) = F_V(v) = F_V[\beta^{-1}(b)]$, and therefore, $g_B(b) = \frac{f_V(v)}{\beta'(v)} = f_V[\beta^{-1}(b)] \frac{d\beta^{-1}(b)}{db}$ is the corresponding probability density function. Substituting these terms into equation (4) yields

$$v = b + \frac{G_B(b)^{N-1} - G_B[\tau(b)]^{N-1}}{(N-1)g_B(b)G_B(b)^{N-2}}. \quad (5)$$

This formulation of the equilibrium shows that each bidder's private value is point identified from a sample of bids. This is particularly useful because, when mis-specifying the pricing rule as SPA, the researcher implicitly assumes that bids *are* private valuations, whereas, given an estimate of the bid distribution and

⁶Despite the piecewise definition, $\beta(v)$ turns out to be continuously differentiable at the point v_Δ , since the derivative in equation (4) approaches the first price derivative as $v \rightarrow v_\Delta$ from above; see Hickman (2010).

⁷This example was constructed using $N = 3$, $V \sim \text{Rayleigh}(1)$, truncated to the interval $[0, 1]$, and $\Delta = 0.1$.

density, equation (5) allows for a simple error correction to adjust for demand shading, or the difference between private values v bids b .

In the following section, using simulated data, we present a sensitivity analysis to compare empirical results under the (incorrect) SPA assumption, which merely takes bids as private values and estimates their PDF nonparametrically, and the (correct) EA assumption. For the latter case, we construct a two-stage nonparametric estimator in the spirit of GPV: in the first stage, we construct an empirical analog of (5) using a kernel density estimator of g_B to get a sample of private value estimates, $\{\hat{v}\}$; in a second stage we then kernel smooth the density of \hat{v} to obtain an estimate of f_V . To avoid problems of sample trimming, we use the boundary-corrected GPV (BCGPV) estimator proposed by Hickman and Hubbard (2015).⁸ In order to provide a basis of comparison between the EA and SPA estimates, we also directly estimate the density f_V by kernel smoothing the sample of simulated private values as well. Although such a strategy would never be available to practitioners working with real data, it provides a baseline estimate by which to judge the closeness of the other two estimators to the true distribution, given a finite sample of data with which to work.

As a precursor to our sensitivity analysis, in Figure 2, we plot three estimated PDFs. We considered a simple case in which a researcher observes 1,000 auctions each having five bidders who draw valuations from a Weibull(0.5,4.0) defined on $[0, 1]$. In panels (a), (b), and (c) of the figure, we display representative results under a bid increment of 0.02, 0.05, and 0.1, respectively. In each panel the solid line is the kernel-smoothed nonparametric estimate of $f_V(v)$ based on the random valuations we generated; this is the best one could hope for from a nonparametric estimator; we denote it $\hat{f}(V)$. The other two estimators take as input the equilibrium bids from an EA that correspond with the randomly-drawn valuations. The dash-dotted line represents $\hat{f}_{EA}(V)$ —the nonparametric estimate under the correctly-specified EA model. The dashed line represents $\hat{f}_{SP}(V)$ —the nonparametric estimate under the mis-specified SPA model. Note that $\hat{f}_{EA}(V)$ is not visually distinct in the figure because it coincides almost exactly with $\hat{f}(V)$. The mis-specified nonparametric estimate attains a higher peak and is shifted to the left of the optimal estimate (and the EA nonparametric estimate), with the effect becoming more pronounced the larger is Δ . This occurs because, for a given valuation, the SPA model predicts a higher bid (bidders’ weakly dominant strategy is to bid their valuation) than the EA model, which involves shading bids in the hope that the item at auction is awarded under a first-price rule. As such, the valuation implied by a given bid is lower under a SPA-assumed pricing rule. It is also worth noting that \hat{f}_{EA} is naturally handicapped relative to \hat{f}_{SP} because the former is a two-step nonparametric estimator and has a slower convergence rate than the latter, a one-step nonparametric estimator. In the next section, our sensitivity analysis will test varying sample sizes, but the figure suggests the statistical challenges a two-step nonparametric estimator faces can be less than the cost of pricing rule mis-specification, something we formally pursue in the next section.

3 Mis-Specification Analysis

In this section we adopt three alternative specifications for the private value distribution—Exponential, with CDF $F(v) = 1 - \exp(-v/\theta)$; Power, with CDF $F(v) = v^\theta$; and Rayleigh, with CDF $F(v) = 1 - \exp(-v^2/2\theta^2)$ —in order to explore the implications of the non-standard pricing rule based on bid

⁸Traditional kernel density estimators are known to be inconsistent and biased at the boundary of the support. GPV proposed a solution to this problem that involved discarding data within a neighborhood of the sample extremes in order to preserve consistency within the interior of the support. In finite samples this creates several problems for inference. Hickman and Hubbard (2015) proposed an alternative approach based on boundary-corrected kernel density estimators which are uniformly consistent on the closure of the support. They described a number of attractive features of the boundary-corrected GPV estimation strategy, but for this paper the most important benefit is perhaps preserving the entire sample of data, making the one-step SPA estimator and the two-step EA estimators more comparable.

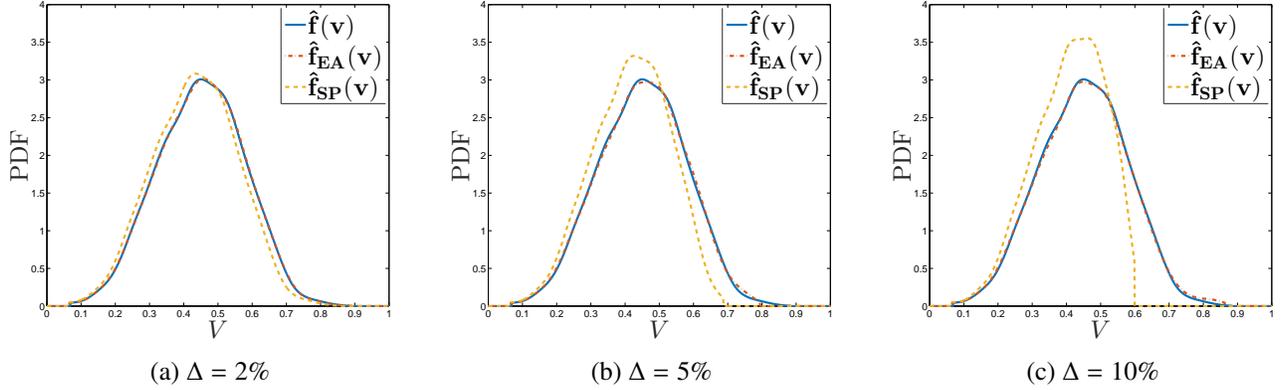


Figure 2: Comparison of Probability Density Function Estimates

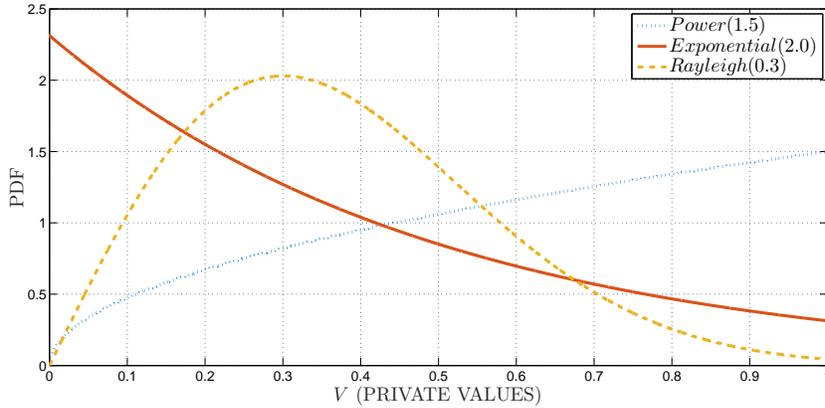


Figure 3: Truncated Probability Density Functions

increments. For each one we assumed that $[\underline{v}, \bar{v}]$ equals $[0, 1]$, and all distributions are truncated and re-normalized so that the corresponding densities integrate to one. For simplicity, in what follows we simply use the names of the untruncated distributions when referring to the truncated ones. The three PDFs are depicted in Figure 3. We chose these three in particular for our Monte Carlo exercise in order to evaluate the effect of the PDF having a mode at \underline{v} , an interior mode, and a mode at \bar{v} . Unless explicitly stated, we assumed the bid increment Δ is two percent of the highest valuation which is consistent with a major portion of eBay bid increments.⁹ Before presenting our Monte Carlo experiments, we perform two exploratory analyses to probe the realism of our three test distributions.

3.0.1 Error in the Equilibrium Bid Function

To quantify the effect of mis-specifying the EA pricing rule, we computed a relative measure of error in the implied bidding function. Define the relative error from modeling and solving for the equilibrium at an EA by assuming an SPA by $\varepsilon(v) \equiv \frac{\beta_{SP}(v) - \beta_{EA}(v)}{v} = \frac{v - \beta_{EA}(v)}{v} = 1 - \frac{\beta_{EA}(v)}{v}$, which can be interpreted as the percentage error in the predicted bid for a given valuation as $\beta_{SP}(v)$ equals v .

⁹Bid increments at eBay auctions are discussed at <http://pages.ebay.com/help/buy/bid-increments.html> which we accessed on 8/25/2015. Bid increments ranged from five percent for low-valued items (under \$5.00) to one percent for items that are selling for between \$2,500.00 and \$4,999.99. As an example, items with prices between \$5.00 and \$24.99 have a bid increment of \$0.50 which is ten percent of \$5 and two percent of \$25.

Table 1: Expected Relative Error in Bid Function under SPA Assumption

Distribution	$N = 3$	$N = 5$	$N = 10$
Exponential(2.0)	0.11584	0.08945	0.06245
Rayleigh(0.3)	0.06685	0.05607	0.04162
Power(1.5)	0.04540	0.03950	0.03119

In Table 1, we summarize the expected relative error involved in assuming an SPA—which we computed as $\mathbb{E}[\varepsilon(V)] = \int_{\underline{v}}^{\bar{v}} \varepsilon(v) f_V(v) dv$ —for each of the three distributions. The expected relative errors reported in Table 1 are all greater than three percent and can be as high as eleven percent. For each distribution we see that as the number of bidders increases, the expected error decreases. For an EA, the sign of this effect is not obvious: with more competition bidders behave more aggressively, but with more participants at an auction, the probability that the top two bids are within Δ of each other also increases. The numbers in the table suggest that the former competitive effect dominates the latter probabilistic effect.

3.0.2 Bid Increments and the Frequency of a First-Price Rule

Hickman (2010) found, within a sample of 1,128 eBay auctions for laptop computers, that 23.05% of final sale prices were generated by a first-price rule being triggered (because the top two bids were close together), as opposed to the often assumed second-price rule. For our test distributions, we simulated EA auctions involving three, five, and ten bidders, with bid increments of 0.01, 0.02, and 0.05.¹⁰ In each case we simulated one million auctions and computed the fraction where the first-price rule determined the transaction price.

In Table 2, we present these frequencies for the simulated scenarios. The results illustrate that, for a given distribution and a fixed number of players at auction, increasing the bid increment increases the share of transaction prices determined by the first-price rule. Likewise, for a given distribution and a fixed bid increment, increasing the number of players at auction increases the share of transaction prices determined by the first-price rule. The Power distribution dominates the Exponential one, but there is no clear ranking between these two and the Rayleigh. For a given (N, Δ) pair, the Power distribution involves a higher share of bid profiles triggering the first-price rule than the Exponential. The distributions we consider, with realistic bid increments, are capable of generating phenomena consistent with what is observed in actual data. Moreover, assuming an EA is an SPA involves potentially serious consequences.

To evaluate the effects of mis-specification, we employed a nonparametric empirical model under the correct EA assumption as well as the mis-specified SPA assumption. Next, we report the results of simulation experiments in which we used these distributions to demonstrate that mis-specification can lead to significantly different estimates of the latent valuation distribution. In subsection 3.2, we consider the implications that model mis-specification can have, estimating optimal auctions and quantifying the economic importance of the biased policy prescriptions and predictions deriving from the SPA assumption.

3.1 Simulation Experiments

We conducted a series of simulation experiments in which we varied model components, including $F_V \in \{\text{Exponential}(2.0), \text{Rayleigh}(0.3), \text{Power}(1.5)\}$, $N \in \{3, 5, 10\}$, and $\Delta \in \{0.02, 0.05, 0.10\}$. We also varied the sample size of auctions $T \in \{100, 300\}$. We simulated each instance $S = 1,000$ times and allowed the

¹⁰By calibrating $\bar{v} = 1$, the bid increments correspond to a percentage of the highest possible valuation.

Table 2: Frequency of First-Price Rule at EA Auctions

		Distribution		
Δ		Exponential(2.0)	Rayleigh(0.3)	Power(1.5)
$N = 3$	0.01	2.90%	4.04%	3.80%
	0.02	5.80%	8.04%	7.47%
	0.05	14.65%	20.23%	18.04%
$N = 5$	0.01	3.53%	5.21%	6.72%
	0.02	7.00%	10.27%	13.07%
	0.05	17.04%	25.35%	30.07%
$N = 10$	0.01	5.04%	6.54%	13.52%
	0.02	9.90%	12.85%	25.37%
	0.05	23.71%	31.48%	52.88%

econometrician to observe the bids from all potential bidders. For each simulation s , we did the following:

1. generated T N -tuples of valuations from a given distribution;
2. used the true EA bid function to map these valuations into bids which we assume the researcher actually observes;
3. assumed the bids came from an SPA and estimated the model via the one-step nonparametric estimator described in the previous section;
4. assumed the bids came from an EA auction and estimated the EA model via the two-step nonparametric estimator described in the previous section.

To evaluate the statistical performance of the estimators, we constructed tests of the null hypothesis that the sample of estimated pseudo-values recovered under the SPA and EA assumptions, respectively, came from the same distribution as the actual sample of simulated valuations. Specifically, we used a two-sample Kolmogorov–Smirnov test as well as an Anderson–Darling test based on Scholz and Stephens (1987). Results are presented in Tables A.6 and A.7 in the Online Appendix, respectively. Specifically, we present results from this simulation exercise by reporting the number of null hypotheses rejected (that the two distributions are the same, at the 5% level) as well as the median p -value of the Anderson–Darling test statistic for the instances involving the various distributions, bid increments, number of bidders at auction, and sample sizes.

The most robust result is that the two-step EA estimator always outperforms the one-step, albeit misspecified, SPA estimator. For the exponential and Rayleigh distribution cases, a null hypothesis is never rejected under the EA estimation. For a given bid increment, as either T or N increases, the number of times the SPA-based null hypothesis is rejected increases. Moreover, once Δ reaches five percent, nearly every simulation involving the SPA assumption allows the null hypothesis to be rejected. Regardless of the distribution, we never reject the null hypothesis for the two percent bid increment cases involving the EA-based estimates. The power distribution case, however, is notably difficult for both estimators if Δ or $L \equiv NT$ is sufficiently large, though the correctly-specified EA model is superior. The EA model gets rejected because the Kolmogorov–Smirnov test statistic was designed to be used on direct observations from a given CDF, whereas in the EA case, the pseudo values are estimated. Moreover, the Kolmogorov–Smirnov test statistic converges at rate $1/\sqrt{L}$ which is faster than the optimal nonparametric convergence rate for a two-step estimator $(\log L/L)^{1/5}$ derived by GPV (Proposition 1). The power distribution has a right-hand mode, where kernel-based estimators are known to have difficulty. For the power distribution, in addition to the full sample, we present a 95% sample (we drop the top 5% of each dataset) and show

that this issue is related to the upper boundary—the EA model is never rejected in these subsamples while the SPA is almost always for large enough L or Δ .

Under the SPA assumption, the private value distribution is the same as the bid distribution. Thus, structural estimation methods which employ the second-price rule will uncover the population bid distribution $G_B(b)$ as the sample size gets large. That is, if we denote the estimated valuation distribution under an SPA assumption given a sample of T auctions by $\widehat{F}_{\text{SP}}^T(V)$, then as the number of auctions in the sample increases, we have $\text{plim}_{T \rightarrow \infty} \widehat{F}_{\text{SP}}^T(V) = G_B(b) = F_V[\beta^{-1}(b)]$. Since $\beta(v)$ does not equal v when Δ is positive, it is clear that $\widehat{F}_{\text{SP}}(V)$ will fail to converge in probability to $F_V(v)$. As such, the estimated demand functions under the SPA assumption will always lie to the left of the estimated demand function under the EA assumption.

3.2 Economic Importance of Mis-Specification

To investigate the economic importance of the bias that obtains when a researcher estimates an EA under the SPA assumption, we considered two exercises which an econometrician might be asked to conduct: recommending a reserve price and predicting anticipated revenues if another bidder were to enter the auction. First, we computed the implied optimal auctions (involving optimally-chosen reserve prices) corresponding to each estimated distribution (based off the EA and the SPA assumptions), for each simulation s for each instance of our experiment. Denote by Ω the set of all auctions at which: (1) any bidder can submit a bid as long as it is greater than some value r^* ; (2) the buyer submitting the highest bid above r^* is awarded the object; (3) auction rules are anonymous in that each bidder is treated in the same way; and (4) there exists a monotone, symmetric, pure-strategy, Bayes–Nash equilibrium. At any auction satisfying these four conditions, the optimal reserve price r^* must satisfy

$$r^* = v_0 + \frac{[1 - F_V(r^*)]}{f_V(r^*)}$$

where v_0 is the seller’s valuation for the item at auction; see, Riley and Samuelson (1981) as well as Myerson (1981). In our simulation experiments, we assumed v_0 is zero and computed the optimal reserve price implied by the EA and SPA estimates for each simulation experiment.

In Table A.8 in the Online Appendix, we present the mean reservation price \hat{r}_{EA} and \hat{r}_{SP} implied by the estimates of the latent valuation distribution and density under each assumption, EA and SPA, respectively, along with their standard errors $\hat{\sigma}_{r_{\text{EA}}}$ and $\hat{\sigma}_{r_{\text{SP}}}$. The true optimal reserve price for the Exponential(2.0), Rayleigh(0.3), and Power(1.5) cases are 0.36077, 0.29905, and 0.54288, respectively.¹¹ The mean reservation price under the EA auction is always within a standard deviation of the true optimal reservation price. In contrast, under the SPA assumption, the mean reservation price is regularly at least two standard deviations away from the truth for large enough sample size and especially for the larger bid increments, suggesting that, from a policy perspective, mis-specification has important effects.

In our second exercise, we used the estimated latent value distributions to estimate what revenues would be were another bidder to participate at auction. This consideration is motivated by researchers who have pointed out that adding another bidder to the auction is often far more valuable than getting the reservation price exactly right. For example, Bulow and Klemperer (1996) showed that, in a world with a fixed number of participants, the optimal auction with a specific number of bidders provides less revenue than an auction with no reserve price, but one additional bidder. We consider an econometrician who might be asked to predict the expected revenues were another bidder to show up at auction. To do this,

¹¹Recall that the optimal reserve price does not depend on the number of bidders at auction, although we present estimates in Table A.8 for which the number of bidders at auction varied according to the cases we considered in our simulation experiments.

we appeal to the revenue equivalence theorem and note that we need only draw with replacement from each estimated valuation distributions and compute the average values of the second highest valuation a sufficiently high number of times. In this way, the root cause of any discrepancy is entirely attributed to the error deriving from the original mis-specification in recovering the valuation distribution from the observed bids.

In Table A.9 in the Online Appendix, we present the expected revenue an econometrician would predict were another bidder to enter the auction. For the reasons documented earlier, the SPA-estimated specification always underpredicts expected revenue. Note that, since our simultaneous EA model satisfies the assumptions of the revenue equivalence theorem, a practitioner should predict the same expected revenue regardless of the minimum bid increment Δ corresponding to the data from which the underlying distribution was estimated from. The results in the table illustrate the substantial variation in predictions from the SPA-based estimations for fixed values of N and T across the sections of the table in which Δ is changing. Specifically, as Δ increases, a practitioner assuming a SPA pricing rule would underpredict the expected revenue by larger amounts.

Our simulation results suggest that mis-specifying the pricing rule can result in significantly different estimates of the latent distribution of valuations. These differences carryover to policy recommendations because the SPA-estimated valuation distributions suggest significantly different policy estimates—optimal reserve prices. There are also consequences in applying estimates from the mis-specified model to predict expected revenues; for example, if another bidder were to participate at auction. In contrast, concerns about the two-step estimation process required of a correctly-specified EA model do not appear to cause issues in practice as the estimator performs quite well in these dimensions.

4 Identification and Estimation Under Stochastic Participation

We now extend our model to make it more compatible with real-world data from eBay. The principal challenges in empirical applications are three-fold: first, the econometrician does not observe all bids, but rather a selected sample of bids which may or may not be consistent with equilibrium play. Second, the total number of bidders participating in an auction is unobserved to the econometrician, and instead, only a lower bound on total participation can be gleaned from data. Third, real-world electronic auctions do not use constant bid increments, but bid increment schedules that are piecewise constant (that is, they discretely jump at specific, predetermined points in bid space). In this section, we provide tractable solutions to these three problems, and demonstrate that a model with a correctly-specified pricing rule is still non-parametrically identified under these more empirically realistic conditions. We then propose a sieve-type estimation strategy, based on B-splines, that we implement in the next section using data from eBay laptop auctions.

4.1 A Bidding Model With Stochastic N

One common characteristic of EAs is that the web-based interface makes it impossible to observe precisely the number of potential competitors within a given auction; that is, the number of users who are following an item with intent to bid on it. Difficulty in measuring N has long been a principal challenge within the empirical auctions literature, particularly when the bidders are able to observe N and adjust their bidding strategies with information unavailable to the econometrician. At an EA, however, the researcher and the bidders are on the same footing in that *neither* observes N . Therefore, we shall model participation from the perspective of a bidder as a stable stochastic process that exogenously allocates bidders to a given auction.

Specifically, let bidders view N as a random variable with probability mass function $\rho_N(n; \lambda) \equiv \Pr(N = n)$ indexed by a parameter vector λ and assume that they do not know the realization of N *ex ante*, when they compute their strategic bids.¹² In formulating the exogenous participation process this way, we allow for λ to be infinite dimensional so that the distribution of N may be fully nonparametric if, for example, $\lambda = \{\lambda_0, \lambda_1, \lambda_2, \dots\}$, where $\lambda_n = \Pr(N = n)$.

Once again, consider the auction from bidder 1's perspective, and let $M \equiv N - 1$ denote the number of opponents he faces. We also define

$$\rho_M(m; \lambda) \equiv \Pr(M = m | N \geq 2) = \Pr(N = m + 1 | N \geq 2) = \frac{\rho_N(m + 1)}{1 - \rho_N(0) - \rho_N(1)}, \quad m \in \{1, 2, 3, \dots\}$$

as the probability that bidder 1 faces exactly $m \geq 1$ opponents.¹³ Just as before when N was known *ex ante*, a bidder's strategic decision problem within an auction is how to respond optimally to his highest rival bid. We denote the highest rival valuation and bid as random variables V_M and B_M , respectively, and we denote their respective distributions as

$$F_M(V_M) = \sum_{m=2}^{\infty} \rho_M(m; \lambda) F_V(V_M)^m, \quad \text{and} \quad G_M(B_M) = \sum_{m=2}^{\infty} \rho_M(m; \lambda) G_B(B_M)^m. \quad (6)$$

As such, F_M and G_M are weighted sums of powers of their respective parent distributions, where the weights represent the probability of a given realization for the number of potential bidders.

With these adjustments to notation, the bidding model based on a hybrid pricing rule can be easily extended to handle stochastic exogenous participation. By inserting F_M into equation (3) above in place of F_Z , we get a new differential equation to redefine the equilibrium β for the stochastic participation case (with the same boundary condition as before); namely,

$$\beta'(v) = \frac{[v - \beta(v)] f_M(v)}{F_M(v) - F_M[\beta^{-1}(\tau[\beta(v)])]}, \quad \text{and} \quad \beta(\underline{v}) = \underline{v}. \quad (7)$$

Intuitively, whether N is known or stochastic, the highest rival bid is still a random variable to which each player best responds. The only difference now is that randomness comes from two sources: a given opponent's valuation is unknown, and the quantity of opponents is also unknown.

Going forward we denote the inverse bidding function by $\xi(b) \equiv \beta^{-1}(b) : [\underline{b}, \bar{b}] \rightarrow [\underline{v}, \bar{v}]$. As before, equation (7) can be transformed to express the inverse bidding relationship as

$$\xi(b) = v = b + \frac{G_M(b) - G_M[\tau(b)]}{g_M(b)}, \quad (8)$$

where $g_M(b)$ is the density corresponding to the distribution of the highest rival bid. These two equations establish that the form of the inverse bid function can be inferred, so long as λ and the parent distribution of bids G_B can be identified. If ξ and G_B are known, then F_V is also known since $F_V(v) = G_B[\xi(b)]$. Thus, structural identification now hinges crucially on recovering the distribution of N from data. This we address in the next subsection.

¹²Song (2004b) was the first to propose a bidding model of first-price auctions where bidders view unknown N as a random variable. She then showed that assuming N is distributed Poisson makes it possible to identify the distribution of private valuations without knowing the exogenous arrival rate of bidders. However, knowledge of λ is needed for counterfactuals and revenue projections using structural estimates. We take a different approach here which aims to identify both the private value distribution *and* the distribution of N from observable data. The advantages of our approach are that parametric assumptions are unnecessary and structural counterfactual simulations are well defined.

¹³Conditioning on the event that $N \geq 2$ comes from two facts. First, if bidder 1 exists, then he knows N is at least 1. Second, if $N = 1$ so that bidder 1 faces no opponent, then he wins the object at a price of Δ . But since his own bid has no bearing on the likelihood of this possibility, the $N = 1$ scenario does not enter his decision making.

4.2 Identifying Exogenous Participation Rates

The set of observables available from an electronic platform like eBay, however, presents several challenges. We assume the econometrician does not observe all bids, from which the parent distribution G_B could be easily estimated. Rather, he observes a selected subset, being only the second order statistic from a sample of stochastic size. Note that on eBay, the maximal bid is only observed at auctions where the first-price rule was triggered, roughly one in every five in our data. There may also be reasons to doubt whether the third-highest and other observed bid submissions are generated by equation (8), as we discuss below in Section 5.1. If the researcher has access to a broader subset of bids than what we describe here, then the estimator resulting from our identification strategy will have more statistical power. But in order to be conservative on the capabilities of our proposed method, we use only the highest losing bid (the second order statistic).

Furthermore, the econometrician does not observe N , the total number of bidders, directly, but rather, she sees only the number of bidders who submit tenders to the server, call it \tilde{N} . This number we argue is merely a lower bound: some bidders who watch an item with intent to bid may find that their planned bid was surpassed before they get around to submitting it. Thus, the list of actual participants is passed through a natural “filter process” which withholds some of them from view *before* the econometrician is allowed to see the list of observed participants.

4.2.1 The Filter Process

Underlying this idea is an assumption of simple intra-auction dynamics in the sense that ordering of bidders’ submission times is random.¹⁴ We assume that, prior to the auction, Nature generates a list of bidders, indexed $\{1, 2, \dots, n\}$, where n follows known distribution $\rho_N(n; \lambda)$, but each bidder is confined to an enclosed cubicle so that she cannot observe the realization of n . For each i Nature generates an iid private valuation v_i from F_V . Each bidder then formulates his strategic sealed bid, $\beta(v_i)$, and waits for Nature to come collect it from him. Nature visits each bidder in order of his index within the list, but if the highest two bids from previous tenders both exceed $\beta(v_i)$, then Nature skips bidder i ’s submission, discarding it as if it never happened. At the conclusion, Nature reports to the econometrician the number of *recorded* bidders.

Within this simple environment, for each bidder $i \geq 3$, whenever the second-highest bid from among $\{\beta(v_1), \dots, \beta(v_{i-1})\}$ exceeds $\beta(v_i)$, then i will not appear to have participated, even though she may have intended, *ex ante*, to submit a bid. Observing only a subset of potential bidders presents a challenge to the econometrician, but by explicitly modeling the filter process we can overcome it and still identify λ from observed lower bounds \tilde{n} . Moreover, if one is interested solely in modeling auction participation, then the filter process can be further simplified. Since equilibrium bidding is monotone and bidder visibility depends on the relation between rank ordering of bids and bid timing, we can re-cast the filter process in equivalent terms where Nature endows each bidder i with an iid quantile rank $Q_i \sim \text{Uniform}(0, 1)$, rather than a private valuation. Nature then walks through the (unordered) list $\mathbf{q} = \{q_1, q_2, \dots, q_n\}$ (for a given value of n) and reports the following number to the econometrician:

$$\tilde{n} = \begin{cases} n & \text{if } n \leq 2, \text{ and} \\ 2 + \sum_{i=3}^n \mathbf{1}(q_i^* < q_i) & \text{if } n \geq 3, \end{cases} \quad (9)$$

¹⁴Note that our assumption allows us to be agnostic concerning how agents individually decide on bidding early or late within the auction. We do rule out, however, the possibility of coordination on some observable aspect of the auction so that the *relative ordering* of their bid submission times is systematic, rather than random. See further discussion on our assumptions which simplify intra-auction dynamics in Section 5.2 below.

where q_i^* is the second-highest from among $\{q_1, q_2, \dots, q_{i-1}\}$, and $\mathbf{1}(\cdot)$ is an indicator function. This observation facilitates simulation of the filter process *without* knowing F_V *ex ante*, which in turn makes it possible to separate identification/estimation of λ and F_V .

From the above description, it is easy to see that the distribution of \tilde{N} for a given value of N is invariant to changes in λ . Therefore, the filter process can be repeatedly simulated for arbitrary hypothetical values of n , and we can treat the conditional probabilities $\Pr(\tilde{n}|n)$ as known quantities for arbitrary (\tilde{n}, n) pairs. Since \tilde{N} is observable, we can in turn treat its probability mass function, denoted $\tilde{\rho}_N(\tilde{N})$, as an observable since it can be directly estimated from data. Moreover, by the law of total probability, we have the following relationship which establishes identification of the exogenous participation process:

$$\tilde{\rho}_N(\tilde{n}) = \sum_{n=0}^{\infty} \Pr(\tilde{n}|n) \times \rho_N(n; \lambda). \quad (10)$$

Since the above argument does not rely on an assumption that λ is finite-dimensional, our identification result is, in fact, *nonparametric*. In other words, observed participation together with our model of the filter process are enough to identify the distribution of N on their own, without appealing to specific functional form assumptions on ρ_N . In practice, additional parametric assumptions (for example, specifying N as Poisson) may provide benefits such as statistical efficiency or numerical tractability, but they are not fundamentally necessary from an identification standpoint. Below in our empirical application, we shall see that the parametric generalized Poisson model (Consul and Jain (1973)), a two-parameter distribution, provides a remarkably tight fit to the data, leaving very little room (given our finite sample) for additional improvements to fit through functional form relaxations.

4.3 Identifying F_V

With the above result in hand, nonparametric identification of the remainder of the structural model is straightforward. Let $H(b)$ denote the *ex ante* distribution of the highest losing bid within an auction (the second order statistic from a sample of stochastic size N). Once again, since the highest losing bid is observable, we can treat H as observable since it can be directly estimated from data. It relates to the parent distribution of bids via the following mapping

$$H(b) = \sum_{n=2}^{\infty} \frac{\rho_N(n; \lambda)}{1 - \rho_N(0; \lambda) - \rho_N(1; \lambda)} \left(G_B(b)^n + n G_B(b)^{n-1} [1 - G_B(b)] \right). \quad (11)$$

For fixed λ this mapping is a bijection for each b in the bid support. Therefore knowing λ and H implies that the parent distribution of bids is identified.

With the above arguments in place, we can state formally our identification result. In order to fix notation, we define a model as a set of (potentially nonparametric) arrival probabilities $\rho_N(n; \lambda)$, $n = 0, 1, 2, \dots$ and a private value distribution F_V . Moreover, we assume that the observables available to the econometrician include $H(b)$, the distribution of the highest losing bid, and $\tilde{\rho}_N(\tilde{n})$, $\tilde{n} = 0, 1, 2, \dots$, the probability mass function for observed participation \tilde{N} (which is a lower bound on actual participation).

Proposition 4.1. *Under the assumptions of Section 4.2.1, the bidding model $(\{\rho_N(n; \lambda)\}_{n=0}^{\infty}, F_V)$ is nonparametrically identified from the observables $(H(b), \{\tilde{\rho}_N(\tilde{n})\}_{\tilde{n}=0}^{\infty})$.*

Proof: Equation (10) establishes identification of the nonparametric bidder arrival probabilities $\{\rho_N(n; \lambda)\}_{n=0}^{\infty}$ from the distribution of observed lower bounds under the model of the filter process described in Section 4.2.1. Given known arrival probabilities, the bijectivity of the mapping (11) establishes that the parent distribution of bids G_B is nonparametrically identified from the observables.

This in turn means that we can now construct G_M from equation (6) using the parent bid distribution and the bidder arrival probabilities.¹⁵ Moreover, if G_M is known then we can re-construct the inverse bidding function $\xi(b) = \beta^{-1}(v)$ from equation (8). Finally, this also implies that the private value distribution F_V is nonparametrically identified through the relationship $F_V(v) = G_B[\beta(v)]$. ■

4.4 Model Extension: Nonconstant Δ

Having presented the basic identification strategy, we now extend our model to handle a final challenge. Above, we assumed that the bid increment Δ is constant, but at most real-world EAs Δ changes at pre-determined transition points. For simplicity, consider a single transition point, as the extension generalizes straightforwardly for two or more transition points.

Suppose we have a transition point, denoted by b^* , and two bid increments, denoted $\Delta_1 < \Delta_2$: Δ_1 applies when the second-highest bid is strictly less than b^* , and Δ_2 applies when it is weakly above.¹⁶ Recall that when optimizing bids on the margin, players only consider how the bid increment controls the threshold (below their own bid) at which a first-price rule is triggered, so that their own bid determines sale price. Thus, the important detail to keep in mind is that when one's own bid b passes the transition point b^* , the threshold at which the first-price rule is triggered changes, and moves further away from b , since $\Delta_1 < \Delta_2$. With this in mind, we re-define an adjusted threshold function as

$$\tau(b) = \begin{cases} \underline{v}, & b \leq \underline{v} + \Delta_1, \\ b - \Delta_1, & \underline{v} + \Delta_1 \leq b < b^* + \Delta_1, \\ b^*, & b^* + \Delta_1 \leq b < b^* + \Delta_2, \text{ and} \\ b - \Delta_2, & b^* + \Delta_2 \leq b. \end{cases} \quad (12)$$

Intuitively, whenever one's own bid is less than $b^* + \Delta_1$, then the first-price rule will only be triggered by B_M within Δ_1 of b ; and likewise, whenever one's bid is weakly above $b^* + \Delta_2$, the first-price rule will only be triggered by B_M within Δ_2 of b . Between these two points, the first-price threshold remains constant at b^* : if one's own b is in the interval $[b^* + \Delta_1, b^* + \Delta_2)$ then B_M within Δ_1 of b implies $B_M > b^*$, but at the same time, Δ_2 cannot be involved in a second-price outcome until $b \geq b^* + \Delta_2$. Note, however, that the difference $b - \tau(b)$ steadily increases from Δ_1 to a value of Δ_2 , from which it follows that $\tau(b)$ is a continuous function.

Given this fact, existing results by Hickman (2010) establish that the equilibrium bid function with a transition point is still continuous, with right- and left-hand derivatives that are the same everywhere.¹⁷ Therefore, inserting the expanded version of the threshold function above into equations (7)–(8) suffices to characterize the equilibrium and establish nonparametric identification of the bidding model with transition points as well. The principal challenge that transition points will pose is on the implementation of an estimator, which we discuss in the following section.

¹⁵Note that $H(b)$, the distribution of the highest of $(N - 1)$ draws, and $G_M(b)$, the distribution of the second highest of N draws, are *not* the same distribution. To see why, consider the task of repeatedly simulating order statistics based on fixed N . For each simulation, N iid realizations are generated from some distribution and stored in an unordered list. To compute the value of the second highest from the list of N , we find the maximum, discard it, and take the maximum of the remaining $(N - 1)$ draws; that is, in this case, the highest draw from the original list of N is discarded with certainty. To compute the highest draw from a sample of size $(N - 1)$, we merely discard the first observation of the unordered list and then find the maximum of the remaining $(N - 1)$ draws. But this procedure only discards the maximum of the original N draws with probability $(1/N)$. Therefore, the two random variables cannot have the same distribution.

¹⁶We shall also assume for simplicity that $\underline{v} + \Delta_1 < b^*$.

¹⁷See Lemma 3.3.1 and Proposition 3.3 of Hickman (2010).

4.5 Two-Stage Estimator

In this section, we propose a simple estimator for the bidder arrival process, as well as a sieve-type estimator of the private value distribution F_V based on B-splines. Our choice of B-splines is motivated partly by their ability to accommodate elements of the empirical model flexibly, such as the kink in the bid function at b^* , which alternative methods such as kernel smoothing or global polynomials cannot easily do.

To fix notation, let $\{y_t, \tilde{n}_t\}_{t=1}^T$ denote a sample of auctions; for each we observe the highest losing bid, y_t , and the number of observed participants, \tilde{n}_t . Although we have demonstrated nonparametric identification of the bidder arrival process, we shall focus our discussion on estimation of a model in which λ is finite-dimensional. In our empirical application, we shall specify the distribution of N as generalized Poisson, or

$$\rho_n(n; \lambda) = \frac{\lambda_1(\lambda_1 + n\lambda_2)^{n-1}}{n!} e^{-\lambda_1 - n\lambda_2}, \quad 0 < \lambda_1, \quad 0 \leq \lambda_2 < 1.$$

Later, we show that this two-parameter model leaves little room for further improvements to data fitting through more flexible functional forms: the generalized Poisson model is able to generate a distribution over the observables $\tilde{\rho}_N(\tilde{n}; \lambda)$ that lay within the nonparametric 95% confidence bounds of the empirical distribution of \tilde{N}_t . For the present discussion, however, it suffices to consider any known parametric family $\rho_N(n; \lambda)$ which is indexed by a finite-dimensional parameter vector, λ . Where appropriate, we shall discuss further concerns and complications that would arise if the parametric assumptions are relaxed.

4.5.1 First Stage: Estimating λ

We begin by constructing a simulation routine that mimics the filter process and allows us to estimate $\Pr(\tilde{n}|n)$. Fix a finite upper bound \bar{N} and for each $n \in \{0, 1, 2, \dots, \bar{N}\}$ simulate $s = 1, 2, \dots, S$ auctions wherein a list of n_s independent (unordered) quantile ranks $\mathbf{q}_{n_s} = \{q_{1s}, \dots, q_{n_s}\}$ are drawn from the Uniform(0, 1) distribution. For each such list, we then compute \tilde{n}_s according to the definition in equation (9). For each n , the simulated conditional frequencies are then computed as

$$\hat{\Pr}(\tilde{n}|n) = \frac{1}{S} \sum_{s=1}^S \mathbf{1}(\tilde{n}_s = \tilde{n}).$$

Note that the simulation frequencies are zero whenever $n < \tilde{n}$.

In a slight change of notation, we now redefine the model-generated frequencies of \tilde{N} as

$$\tilde{\rho}_N(\tilde{n}; \lambda) = \sum_{n=0}^{\bar{N}} \hat{\Pr}(\tilde{n}|n) \times \rho_N(n; \lambda), \quad (13)$$

and define the empirical frequencies as $\hat{\rho}_N(k) \equiv \frac{1}{T} \sum_{t=1}^T \mathbf{1}(\tilde{n}_t = k)$. Finally, letting $\tilde{\mathbf{n}} = \{k_1, k_2, \dots, k_L\}$ denote the complete set of unique observed values of \tilde{n} in the data, we define a nonlinear least squares (NLS) estimator as the optimizer of the following objective function:

$$\hat{\lambda} = \underset{\lambda \in \mathbb{R}^k}{\operatorname{argmin}} \left\{ \sum_{l=1}^L \left[\tilde{\rho}_N(k_l; \lambda) - \hat{\rho}_N(k_l) \right]^2 \right\} \quad (14)$$

In words, the estimate $\hat{\lambda}$ is chosen to make the model-generated frequencies of observed bidders as close to the empirical frequencies as possible.¹⁸

¹⁸An analogous nonparametric estimator could be similar, but with additional complications. Reverting back to the case

As a practical matter, specifying \bar{N} and S involves a trade-off between computational cost and numerical accuracy. For the former, we judged $\bar{N} = 100$ to be a sensible choice for several reasons. First, note that the Poisson probability $\rho_N(100; 40) \approx 7.315 \times 10^{-16}$, so the terms truncated out of the infinite sum in (13) will be on or below the order of machine precision whenever N is roughly Poisson with a parameter that is weakly less than 40. Second, while eBay auctions are known for high participation rates, 40 is still quite a large number. For example, in our empirical application with laptop computer data we observed a maximum of 11 participants in any given auction. Thus, our choice of $\bar{N} = 100$ ensures that the finite truncation which we must impose on equation (9) will have no discernible effect for a wide array of eBay datasets. In turn, a relatively low truncation point allowed us to simulate a large number of auctions, or $S = 10^{10}$, which delivers at least 5 (and up to 6) digits of accuracy in each cell of the matrix $\hat{\text{Pr}}(\tilde{n}|n)$. In other words, if the conditional probabilities in equation (13) above are expressed as percentages, then our simulation estimates are accurate to within one one-thousandth of a percentage point.¹⁹ One advantage of our approach is that these need only be simulated once and, given our choice of \bar{N} and S , can be re-used with any dataset for which the average participation rate is less than 40. A copy of the matrix of simulated conditional probabilities $\hat{\text{Pr}}(\tilde{n}|n)$, $n = 1, \dots, 100$, as well as MATLAB code implementation of the estimator, is available from the authors upon request.²⁰

Before moving on, one caveat is worth discussion. In the above proposal we have implicitly maintained a *scale invariance assumption* (SIA)—that the filter process may be simulated using quantile ranks without respect to the underlying equilibrium bid values—in order to make our estimator tractable by simplifying computation of the conditional probability matrix $\text{Pr}(\tilde{n}|n)$. While the resulting simulations are quite involved (see footnotes 19 and 20), the SIA allows for $\text{Pr}(\tilde{n}|n)$ to be computed only once and then re-used in many different settings, and more importantly, it can be re-used for each separate evaluation of the objective function in (14). Having to re-compute the conditional probability matrix repeatedly during runtime would render implementation infeasible. While the SIA is approximately true, small deviations from it exist due to the bid increment, Δ : the EA price evolution in practice may filter out an additional small

where $\lambda = \{\lambda_0, \lambda_1, \lambda_2, \dots\}$, $\lambda_n = \text{Pr}(N = n)$, is infinite dimensional. The main challenge now is that only finitely many elements of λ can be estimated with finitely much data. Therefore, for finite sample size T , we begin by choosing an upper bound, $\bar{N}_T < \infty$, after which we restrict $\lambda_n = 0$ whenever $(n - 1) > \bar{N}$ and define the following:

$$\begin{aligned} \{\hat{\rho}_N(n; \lambda)\}_{n=0}^{\bar{N}_T} &= \underset{\lambda \in \mathbb{R}^{\bar{N}_T}}{\text{argmin}} \left\{ \sum_{l=1}^L \left[\bar{\rho}_N(k_l; \lambda) - \hat{\rho}_N(k_l) \right]^2 \right\} \\ &\text{subject to } \sum_{n=0}^{\bar{N}} \hat{\rho}_N(n; \lambda) = 1. \end{aligned} \tag{15}$$

Choice of \bar{N}_T involves the usual variance-bias tradeoff. For larger \bar{N}_T , less bias arises from setting high-order elements of λ to zero, but as (\bar{N}_T/T) gets large the variance of the estimator will increase as well. A second challenge involves specifying the rate at which \bar{N}_T should optimally grow as $T \rightarrow \infty$. However, our empirical application suggests strongly that solving these problems would produce little benefit above available finite-dimensional parametric methods, so we do not address them here.

¹⁹The probabilities $\hat{\text{Pr}}(\tilde{n}|n)$ are computed as the sample mean of a Bernoulli random variable $\mathbf{1}(\tilde{N} = \tilde{n}|N = n)$. Since the sample mean is known to converge at rate \sqrt{S} , our simulation error is on the order of $1/\sqrt{10^{10}} = 10^{-5}$, but may be even less. Simulation was performed in 100 blocks of 10^8 simulations. As a check on accuracy, these can be used to compute 100 different estimates of the conditional probability matrix. Taking standard deviations across all 100 estimates for each (\tilde{n}, n) pair (and excluding pairs that trivially render a conditional probability of zero), we get mean and maximum standard deviations of 1.47×10^{-5} and 5.55×10^{-5} , respectively. Of course, averaging across these 100 estimates (as our final conditional probability matrix does) should further improve the precision for each (\tilde{n}, n) pair, reducing the mean and maximum standard deviations further to roughly 1.47×10^{-6} and 5.55×10^{-6} , respectively.

²⁰In all, we simulated the filter process for $10^{10} \times 98$ separate auctions (10^{10} simulations for each $N \in \{3, \dots, 100\}$). Computation was performed in parallel using a cluster of 150 MATLAB workers for 310 hours. We periodically reset the seed so as to avoid surpassing the periodicity of the random number generator.

number of bidders from being observed, even though their planned bids would exceed the second-highest from previous bid submissions. Specifically, at a given point in time, with positive probability the next bidder to arrive may wish to submit a bid that is less than Δ above the second-highest previous bid, and in that case the posted price will have already updated to a level slightly above her planned bid.²¹ Because of this deviation from the SIA induced by the bid increment, our data will tend to somewhat undercount the number of observed bidders, relative to what would be the case if the SIA were never violated.

To address this concern, we propose a simple, data-driven correction which allows us to still use our conditional probability matrix $\Pr(\tilde{n}|n)$, whose computation relied on the SIA. Note that under the complication introduced by Δ , there is a positive fraction of the time that k bidders are observed in the data, but if the SIA were perfectly true there would actually have been at least $(k + 1)$ observed bidders. Fortunately, our data contain some additional information which provides clues as to how this process played out. Within each eBay auction, one can observe the complete price path during the life of the auction, from which one can also deduce whether the final sale price involved the triggering of a first-price rule—this happened whenever the value of the final price adjustment was strictly less than Δ . Thus, we add to the set of observables an additional variable \mathcal{F}_t , being 1 if a first-price rule was triggered within auction t , and 0 otherwise. Note that, by definition, this variable informs us on the equilibrium frequency with which the top two order statistics were within Δ of one another, which is also connected to deviations from the SIA. We use this information to adjust our estimator as follows. First, for each value k_l contained in the vector $\tilde{\mathbf{n}}$, compute

$$P_{k_l} \equiv \frac{\sum_{t=1}^T \mathbf{1}(\mathcal{F}_t = 1 \cap \tilde{n}_t = k_l)}{\sum_{t=1}^T \mathbf{1}(\tilde{n}_t = k_l)}, \quad l = 1, \dots, L$$

which estimates the conditional Bernoulli probability that a first-price rule was triggered, given $\tilde{N} = k_l$. Second, we compute an adjusted empirical mass function for \tilde{N} , call it $\hat{\rho}_N^{adj}$, by re-sampling from the data with replacement $S > T$ times, and building an adjusted sample $\{\tilde{n}_s^{adj}\}_{s=1}^S$ where whenever $\tilde{n}_t = k_l$ is sampled in the s^{th} draw, we assign a value

$$\tilde{n}_s^{adj} = \begin{cases} k_l + 1, & \text{with probability } P_{k_l}, \text{ and} \\ k_l & \text{otherwise.} \end{cases} \quad (16)$$

The adjusted empirical mass function then becomes $\hat{\rho}_N^{adj}(k_l) = \sum_{s=1}^S \mathbf{1}(\tilde{n}_s^{adj} = k_l)/S$, for each $k_l \in (\tilde{\mathbf{n}} \cup (k_L + 1))$. Note that if $P_{k_l} = 0 \forall k_l$, then the adjusted mass function will be the same as the original for large S . Third, rather than using the raw empirical mass function $\hat{\rho}_N$ in the objective function (14), we substitute $\hat{\rho}_N^{adj}$ instead. For practical purposes, since the re-sampling step happens only once during runtime, it is not terribly costly to choose S fairly large. For our implementation, we chose $S = 10^7$, which means that the simulation error will have little or no effect on the first four digits of $\hat{\rho}_N^{adj}$.

Intuitively, this adjustment recognizes that the raw empirical distribution of \tilde{N} is stochastically dominated by the one we would observe if the SIA were never violated, and it tends to push the former toward the latter. However, it only partially offsets the over-elimination of bidders in reality, since the conditional Bernoulli probabilities P_{k_l} only inform us on the tendency for the top order statistics to be close together. Therefore, we also propose and execute a robustness check (see Figure A.17 in the Online Appendix) in order to assess the magnitude of the remaining problem. Briefly, we generated a sample of data from our point estimates in which for each simulated auction we knew the true n , valuations, and equilibrium bids for all n potential bidders, as well as the \tilde{n} that would result under the SIA and the \tilde{n}^{raw} which accounts for the over-elimination due to Δ . We then performed estimation using three hypothetical scenarios: (1) with

²¹Note that it is the value of Δ itself which directly controls the degree of deviation, with no other indirect effects since monotone bidding strategies do not by themselves violate the SIA.

an ideal dataset in which the true n was known for each auction and can therefore be directly estimated (for a baseline comparison); (2) an estimator that maintains the SIA, even though the generated \tilde{n}^{raw} data are subject to over-elimination; and, (3) an estimator that employs the same set of generated \tilde{n}^{raw} data but also incorporates a corresponding simulated \mathcal{F} variable to perform our proposed correction. We also provide, for comparison, a plot of the private value CDF that results from mis-specifying EAs as simple, second-price auctions. As is clear from Figure A.17, the overall bias in our baseline estimator of F_V is small, relative to the effect of pricing-rule mis-specification. Moreover, our proposed adjustment eliminates nearly all of the bias due to deviations from the SIA. We take this as evidence that our estimator does not unduly oversimplify the actual EA price dynamics, and achieves an appropriate balance of statistical accuracy and computational tractability.

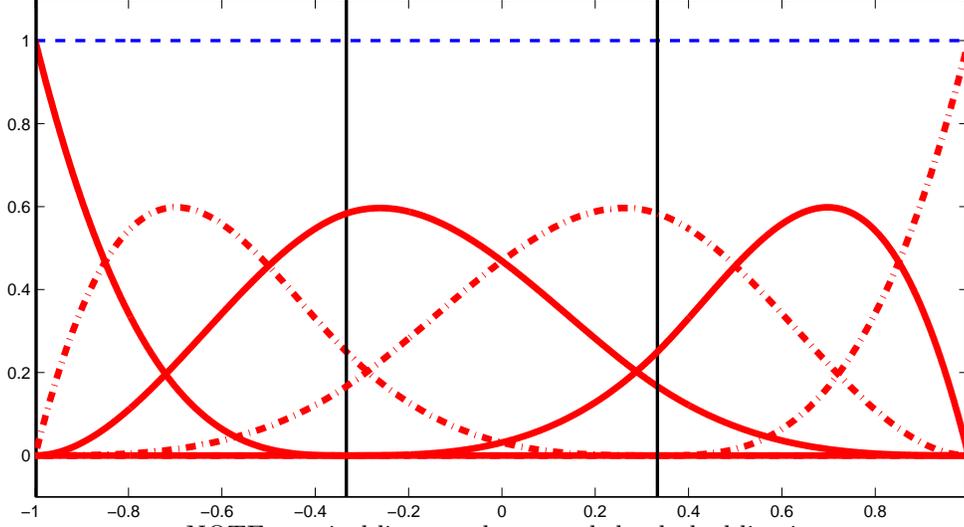
4.5.2 Second Stage: B-spline Estimator of F_V

Our identification argument above states that knowledge of G_M is sufficient to recover the private value distribution. A seemingly natural way to proceed then would be to estimate $G_B(b)$ directly from the observables and then reverse-engineer F_V . Implementing this approach can be difficult, as \hat{G}_B must also ensure that the mappings implied by equation (7) and/or (8) are monotone in order to be consistent with equilibrium bidding. These requirements rule out kernel density estimators, in favor of sieve estimation where a finite-dimensional parametric restriction is imposed and then gradually relaxed as the sample size increases. Even then, parameterizing G_B and then optimizing some empirical criterion function subject to the constraints mentioned above, along with enforcing monotonicity and boundary conditions of G_B itself, poses a formidable numerical challenge. The resulting constraint set becomes very complicated and highly nonconvex, making it difficult to compute admissible initial guesses and converge to global optima afterward.

We propose an alternative approach whereby we directly parameterize the private value distribution F_V as a flexible B-spline function. Still, since $G_B(b) = F_V[\xi(b)]$, building and optimizing an empirical criterion function of the observables—order statistics of bids—requires finding the solution to a differential equation based on F_V given in (7) (or equivalently, (8)). To solve this piece of the puzzle, we employ the *Galerkin method* which is commonly used to solve differential equations numerically in physical sciences applications.²² This approach involves parameterizing the inverse bid function ξ as a B-spline as well, and afterward enforcing its adherence to the conditions of the boundary value problem defined by the equilibrium FOCs. This is done by defining a grid of points on the domain (where the number of grid points is at least as large as the number of free parameters in the B-spline function), and then augmenting the estimator objective function with extra terms that penalize it for deviations from equilibrium. Our approach of augmenting an extremum estimator with the Galerkin method has the added benefit of being relatively inexpensive to compute: rather than repeatedly updating parameter values for F_V and then solving a differential equation in sequence, modeling ξ as a B-spline allows us to fit F_V to the data while *simultaneously* adjusting ξ to conform to the equilibrium conditions required by theory. We explain our approach concretely below, but first a brief word on B-splines is in order.

B-splines have many attractive properties that are well-adapted to our application. First, they behave identically to piecewise splines, so by the Stone–Weierstrass theorem they can fit a broad class of non-parametric curves to arbitrary precision, given a fine enough partition of the domain. They also mimic attractive properties of piecewise splines: being locally low-dimensional they are numerically stable, and adjusting parameters to improve model fit at one point will have little or no effect on behavior of the

²²See Zienkiewicz and Morgan (2006) for details on Galerkin method approximation. Hubbard and Paarsch (2014) discuss and compare various ways in which researchers have solved the differential equation(s) characterizing equilibrium behavior at auctions.



NOTE: vertical lines are knots and the dashed line is the sum of all basis functions.

Figure 4: Fourth-Order (Cubic) B-Spline Basis Functions

B-spline function at points outside of a relatively small neighborhood. This is in contrast to global polynomials, where adjusting parameters to improve fit at one point may have drastic consequences for the behavior of the function at points far away.

On the other hand, like global polynomials, B-spline basis functions are globally defined—in fact, the name “B-spline” is short for *basis spline* because of this—making their functional values, derivatives, and integrals less cumbersome to compute than piecewise splines. In addition, B-splines are more adaptable in applications where the researcher may have *a priori* information about regions of the domain where the function is likely to display a high degree of curvature or complexity. For global polynomials, if the researcher wishes to infuse extra flexibility at a single point, he must increase flexibility on the *entire* functional domain. In contrast, B-splines allow for surgical targeting of an arbitrary degree of curvature within a small neighborhood of a particular point, up to and *including* infinite curvature at that point if kinks or discontinuities are known to occur. This property will be particularly useful in dealing with transition points where the bid increment discretely shifts and the bid function must therefore adjust itself to account for the abrupt change in the strategic environment just above b^* . A brief primer on construction of our B-spline functions is included in the Online Appendix.

To outline our estimator, let $\mathbf{j}_J = \{j_1 < j_2 < \dots < j_J < j_{J+1}\}$ denote a set of unique “knots” which partition the private value support into J subintervals, with $j_1 = \underline{v}$ and $j_{J+1} = \bar{v}$, and let $\mathbf{k}_K = \{k_1 < k_2 < \dots < k_K < k_{K+1}\}$ denote a partition of the bid support into K subintervals, with $k_1 = \underline{b} = \underline{v}$ and $k_{K+1} = \bar{b}$. These knot vectors uniquely define a set of $(J + 3)$ and $(K + 3)$ 4th-order (cubic) B-spline basis functions $\{\mathcal{V}_j(\cdot)\}_{j=1}^{J+3}$ and $\{\mathcal{B}_k(\cdot)\}_{k=1}^{K+3}$, where $\mathcal{V}_j : [\underline{v}, \bar{v}] \rightarrow \mathbb{R}$ for each j and $\mathcal{B}_k : [\underline{b}, \bar{b}] \rightarrow \mathbb{R}$ for each k .

The basis functions are uniquely defined from their respective knot vector through the Cox–de Boor recursion relation formula with concurrent boundary knots (see the Online Appendix for details). The resulting basis behaves the same as a set of piecewise cubic splines that are constrained to be C^2 at the endpoints of the subintervals defined by the knot vector. Each basis function is C^2 everywhere on the global domain, with the C^2 conditions at the knots being built into the recursion formula. Moreover, each one is nonzero on at most four of the K subintervals (though some are nonzero on fewer than four), and exactly four of the cubic basis functions are nonzero on each of the K subintervals, which explains the name “4th-order B-spline,” and also why there are $(K + 3)$ basis functions in total. Figure 4 depicts a Basis

on the interval $[-1, 1]$ partitioned by a uniform knot vector with $K = 3$.

Letting $\boldsymbol{\mu} = \{\mu_1, \dots, \mu_{J+3}\} \subset \mathbb{R}^{J+3}$ and $\boldsymbol{\alpha} = \{\alpha_1, \dots, \alpha_{K+3}\} \subset \mathbb{R}^{K+3}$ denote sets of weights, we can now parameterize F_V and ξ as

$$\hat{F}_V(v; \boldsymbol{\mu}) = \sum_{j=1}^{J+3} \mu_j \mathcal{V}_j(b), \quad \text{and}$$

$$\hat{\xi}(b; \boldsymbol{\alpha}) = \sum_{k=1}^{K+3} \alpha_k \mathcal{B}_k(b).$$

Given this parameterization, the bid distribution is $\hat{G}_B(b; \boldsymbol{\mu}, \boldsymbol{\alpha}) = \hat{F}_V(\hat{\xi}(b; \boldsymbol{\alpha}); \boldsymbol{\mu})$, and, in a slight change of notation, we can redefine the model-generated distribution of the highest losing bid as

$$H(b; \boldsymbol{\mu}, \boldsymbol{\alpha}, \hat{\lambda}) = \sum_{n=2}^{\bar{N}} \rho_N(n|N \geq 2; \hat{\lambda}) \left(\hat{G}_B(b; \boldsymbol{\mu}, \boldsymbol{\alpha})^n + n \hat{G}_B(b; \boldsymbol{\mu}, \boldsymbol{\alpha})^{n-1} [1 - \hat{G}_B(b; \boldsymbol{\mu}, \boldsymbol{\alpha})] \right). \quad (17)$$

This function can now be directly compared to its empirical analog, $\hat{H}(b) \equiv \frac{1}{T} \sum_{t=1}^T \mathbf{1}(y_t \leq b)$, in order to build a criterion function that will form the basis of an estimator. While optimizing the empirical criterion function, we must also impose the equilibrium conditions to ensure that \hat{F}_V and $\hat{\xi}$ will be jointly consistent with the theory model. In other words, our final estimate for $\hat{\xi}$ must constitute a valid solution to the boundary value problem in (7). In order to accomplish this, we specify a uniform grid of checkpoints $\{b_l\}_{l=1}^L \subset [\underline{b}, \bar{b}]$, with $L \geq K + 3$, and we introduce four additional parameters, ε_0 , ε_1 , ε_2 , and ε_3 , and a residual function based on equations (12) and (8):

$$R(b_l; \boldsymbol{\mu}, \boldsymbol{\alpha}, \hat{\lambda}) = \left(\hat{\xi}(b_l; \boldsymbol{\alpha}) - b_l \right) \hat{g}_M(b_l; \boldsymbol{\mu}, \boldsymbol{\alpha}, \hat{\lambda}) - \hat{G}_M(b_l; \boldsymbol{\mu}, \boldsymbol{\alpha}, \hat{\lambda}) + \hat{G}_M(\tau(b_l); \boldsymbol{\mu}, \boldsymbol{\alpha}, \hat{\lambda}) \quad (18)$$

where $\hat{G}_M(b; \boldsymbol{\mu}, \boldsymbol{\alpha}, \hat{\lambda}) = \sum_{n=2}^{\bar{N}} \rho_N(n|N \geq 2; \hat{\lambda}) \hat{G}_B(b; \boldsymbol{\mu}, \boldsymbol{\alpha})^{(n-1)}$. With that, we can now define a constrained NLS estimator as

$$(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\alpha}}) = \underset{\left(\begin{smallmatrix} \boldsymbol{\mu}, \boldsymbol{\alpha} \\ \varepsilon_0, \varepsilon_1, \varepsilon_2 \end{smallmatrix} \right) \in \mathbb{R}^{J+K+6} \times \mathbb{R}_+^3}{\operatorname{argmin}} \left\{ \sum_{t=1}^T \left[H(x_t; \boldsymbol{\mu}, \boldsymbol{\alpha}, \hat{\lambda}) - \hat{H}(x_t) \right]^2 + P_0 \varepsilon_0 + P_1 \varepsilon_1 + P_2 \varepsilon_2 + P_3 \varepsilon_3 \right\}$$

subject to:

$$\begin{aligned} \hat{F}_V(\underline{v}; \boldsymbol{\mu}) &= 0, & \hat{F}_V(\bar{v}; \boldsymbol{\mu}) &= 1, \\ \hat{f}_V(v; \boldsymbol{\mu}) &> 0, & v &\in [\underline{v}, \bar{v}], \\ \hat{\xi}(\underline{b}; \boldsymbol{\alpha}) &= \underline{v}, \\ \hat{\xi}'(b; \boldsymbol{\alpha}) &> 0, & b &\in [\underline{v}, \bar{b}], \\ \mathcal{N} \left[\{R(b_l; \boldsymbol{\mu}, \boldsymbol{\alpha}, \hat{\lambda})\}_{l=1}^L \right] &\leq \varepsilon_0, & b_l &\leq \underline{v} + \Delta_1, \\ \mathcal{N} \left[\{R(b_l; \boldsymbol{\mu}, \boldsymbol{\alpha}, \hat{\lambda})\}_{l=1}^L \right] &\leq \varepsilon_1, & \underline{v} + \Delta_1 &< b_l < b^* + \Delta_1, \\ \mathcal{N} \left[\{R(b_l; \boldsymbol{\mu}, \boldsymbol{\alpha}, \hat{\lambda})\}_{l=1}^L \right] &\leq \varepsilon_2, & b^* + \Delta_1 &\leq b_l < b^* + \Delta_2, \\ \mathcal{N} \left[\{R(b_l; \boldsymbol{\mu}, \boldsymbol{\alpha}, \hat{\lambda})\}_{l=1}^L \right] &\leq \varepsilon_3, & b^* + \Delta_2 &\leq b_l, \\ \varepsilon_0 \geq 0, \varepsilon_1 \geq 0, \varepsilon_2 \geq 0, \varepsilon_3 \geq 0, \end{aligned} \quad (19)$$

where $\mathcal{N}[\cdot]$ is a norm function such as the L_∞ -norm (the sup-norm), $\sup_{l \in \{1, 2, \dots, L\}} |R(b_l; \boldsymbol{\mu}, \boldsymbol{\alpha}, \hat{\lambda})|$, or the L_2 -norm, $\left(\sum_{l=1}^L R(b_l; \boldsymbol{\mu}, \boldsymbol{\alpha}, \hat{\lambda})^2 \right)^{1/2}$, and (P_0, P_1, P_2, P_3) are pre-specified penalty parameters.²³

²³Note that the sup-norm variant of the estimator may be implemented without violating differentiability of the Lagrangean

Intuitively, the residual function equals zero when the private value distribution \hat{F}_V together with the numerical bid function $\hat{\xi}$ exactly conform to the first-order conditions of a bidder's decision problem, and the constraint that $\hat{\xi}(\underline{b}; \alpha) = \underline{v}$ enforces the boundary condition. Thus, the vector $(\varepsilon_0, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ controls the degree of numerical error in the approximated solution to the piecewise differential equation. Moreover, by fitting the parameterized distribution F_V to the data while simultaneously searching over parameter values for $\hat{\xi}$ that conform to the first order conditions, given F_V , we avoid the computational cost involved in repeatedly solving differential equations *each time* the objective function is evaluated, which can easily number into the thousands. Instead the essentially only solve the equilibrium *once*, and our NLS estimator with Galerkin ODE solution is in the spirit of the MPEC (Mathematical Programming with Equilibrium Constraints) method pioneered by Su and Judd (2012): we choose the parameters of the parent distribution G_B so that the model-generated order statistic quantiles match the empirical order statistic quantiles as closely as possible, while penalizing the objective function for deviations from the equilibrium conditions in (12) and (8). Although we tailor our estimator here to the equilibrium of electronic auctions, our approach may be applicable to broader structural contexts outside of auctions where the distribution of a latent, structural random variable is linked to an observable distribution through an equilibrium mapping which must otherwise be repeatedly solved during estimation runtime.

The main difference between our proposal here and the way Galerkin methods are typically implemented stems from how the equilibrium conditions are enforced. The most common implementation is to define the residuals $R(b_l)$ on a grid of checkpoints $\{b_1, \dots, b_L\}$ with $L = K + 2$. This, plus a boundary condition, produces a square system of $K + 3$ equations in $K + 3$ unknowns, which can be solved exactly on the grid of checkpoints (see, *e.g.*, Hulme (1972) for further discussion). However, following this standard approach during estimator runtime would once again necessitate solving the differential equation once for every objective function evaluation, which can number well into the thousands. This is why we opt for the objective function penalization approach described above (with over-fitting, or $L \geq K + 2$), which allows for the parameters of \hat{F}_V and $\hat{\xi}$ to move independently during runtime. This essentially means that the differential equation need only be solved once. As a check on output though, having obtained the point estimate $\hat{F}_V(v; \hat{\mu})$, the researcher can easily verify the solution for $\hat{\xi}(b; \hat{\alpha})$ by re-solving the equilibrium boundary value problem in (12) and (8) using the standard Galerkin solution method based on a square system of nonlinear equations, *holding $\hat{\mu}$ fixed*, and even with a finer knot vector than \mathbf{j}_J , if desired.

4.5.3 Enforcing Boundary Conditions and Shape Restrictions

Before moving on, a final attractive property of B-splines is their ability to facilitate known shape restrictions on the latent function they parameterize. Note from Figure 4 that, by construction, the extremal basis functions are the only ones to attain nonzero values at the boundaries, and they both equal one exactly at their respective endpoints. This makes enforcement of boundary conditions very easy: $\hat{F}_V(\underline{v}; \hat{\mu}) = 0 \Leftrightarrow \mu_1 = 0$, $\hat{F}_V(\bar{v}; \hat{\mu}) = 1 \Leftrightarrow \mu_{J+3} = 1$, and $\hat{\xi}(\underline{b}; \hat{\alpha}) = \underline{v} \Leftrightarrow \alpha_1 = \underline{v}$, which also reduces the number of free parameters by three. Finally, B-splines also have some other remarkable properties that make enforcement of monotonicity conditions (and hence, concavity/convexity also, if needed) quite simple: it turns out that for any B-spline function $\hat{F}_V(v; \mu) = \sum_{j=1}^{J+3} \mu_j \mathcal{V}_j(b)$, we have that $\hat{F}'_V(v) \geq (>)0 \Leftrightarrow \alpha_{j+1} \geq (>)\alpha_j$, for each $j = 1, \dots, J + 2$ (see de Boor (2001, p. 115)). Therefore, enforcing the shape and boundary conditions required by the theory reduces to adding a set of equality constraints and linear inequality constraints directly on the parameter values themselves. This presents

objective function: one can replace the absolute value operator with a two-part, non-linear constraint where $R(b_l; \mu, \alpha, \hat{\lambda}) < \varepsilon_i$ and $-R(b_l; \mu, \alpha, \hat{\lambda}) < \varepsilon_i$ must both be satisfied for each relevant (l, i) pair, for $l = 1, 2, \dots, L$ and $i = 0, 1, 2, 3$. In our experience we have found the L^∞ (sup) norm to be much more stable and accurate than the L^2 (Euclidean) norm. See Figure A.19 in the Online Supplement for a comparison.

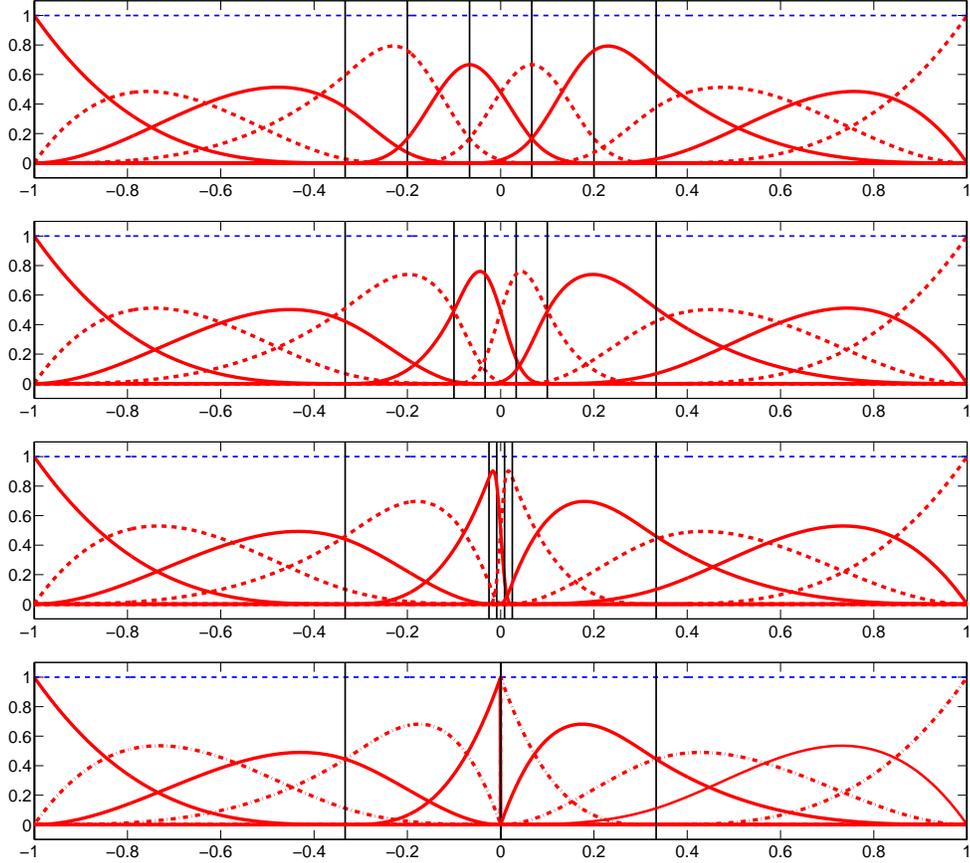


Figure 5: Fourth-Degree (Cubic) B-Spline Basis Functions with Additional Knots

another substantial numerical advantage over global polynomials, where enforcing shape conditions can require imposition of complicated, non-linear constraints on the functional values, which may lead to non-convex constraint sets and increase the tendency of a solver to get stuck at local optima.

4.5.4 Accommodating Transition Points

Given the above derivations, we know that the primitives of the strategic environment remain fixed for bids below $b^* + \Delta_1$ and above $b^* + \Delta_2$, but between these two points there is an abrupt (though continuous) increase in the probability that a first-price rule will be triggered, which means that the degree of demand shading is also expected to abruptly increase on this segment of the domain. Fortunately, the Cox–de Boor recursion formula (with concurrent boundary knots, see the Online Appendix), which we used to compute our B-spline basis functions, was developed with such a contingency in mind. Specifically, it allows for arbitrary curvature at a particular point by adding additional knots that are closely spaced together.

Recall from above that we defined knot vector $\mathbf{k}_K = \{k_1 < \dots < k_{K+1}\}$ in bid space, which in turn uniquely defines $(K + 3)$, C^2 basis functions. For simplicity of discussion, let $K > 2$ and suppose that $(b^* + \Delta_1, b^* + \Delta_2) \subset [k_2, k_3]$. We now modify the knot vector by inserting four additional knots between k_2 and k_3 so that we now get a knot vector with $(K + 5)$ elements

$$\mathbf{k}'_K = \{k_1 < k_2 < \kappa_1^* \leq \kappa_2^* \leq \kappa_3^* \leq \kappa_4^* < k_3 < \dots < k_{K+1}\}.$$

This insertion will have several effects on the recursion formula, which we briefly outline here; the interested reader is once again directed to the Online Appendix for additional details. First, we increase

the number of B-spline basis functions by four, and in particular, exactly eight basis functions will be nonzero on the subinterval $[k_2, k_3]$ now. Second, each set of basis functions—uniquely determined by a particular configuration of the knot vector—places an implicit bound on the second derivative of the B-spline function, but this bound is relaxed within a certain neighborhood as the additional knots are spaced more closely together. Figure 5 illustrates this phenomenon: each of the four panes depict the knot vector from figure 4 with four additional knots inserted between k_2 and k_3 and placed at varying distances from one another. As these get closer together, the amount of curvature displayed nearby in the four new basis functions becomes ever more extreme, and in the limit as the four knots are stacked on top of one another, two of the basis functions become kinked, and two are even discontinuous. Although the region with increased flexibility collapses as the bound on the second derivative is relaxed, the researcher may spread the increased flexibility over a wider area by including more knots, if desired. This is useful because usually the conditions which require extra flexibility within a particular neighborhood are not meant to profoundly impact the behavior of the B-spline function at points far away. Once again, the intuitive properties of B-splines come to bear; recall that each basis function is non-zero on at most 4 subintervals defined by the knot vector. Suppose then that the researcher wishes to locate 4 new knots within a small neighborhood of the middle subinterval, so that $k_2 < \kappa_1^* \leq \kappa_2^* \leq \kappa_3^* \leq \kappa_4^* < k_3$, but he does not wish the influence of the new corresponding basis functions to extend outside of $[k_2, k_3]$ either. In this case, inclusion of 4 additional knots in $[k_2, \kappa_1^*)$ and 4 additional knots in $(\kappa_4^*, k_3]$ (for a total of 12 new knots) would suffice.

A final change induced by the inclusion is that the adjacent basis functions also change shape as well, relative to Figure 4. This highlights an important difference between B-splines and global polynomials. With the latter, inclusion of new terms leads to a nested model structure, wherein new basis functions are simply being added to the pre-existing set. For B-splines, however, as we pack the functional domain with an increasing number K of knots, we get a sequence of *non-nested* models for the underlying function we wish to approximate, since adding one additional knot will cause the pre-existing basis functions with non-zero domain nearby that knot to change form. Numerically, this is of little consequence, but econometrically, this fact will have bearing on interpretation of varying model fit as $K \rightarrow \infty$.

5 Empirical Application

5.1 Data

We now consider a sample of laptop auctions collected between April and June of 2008 from the eBay website, which provides extensive information on item characteristics and bid histories. This application highlights various challenges present in real-world data, including those addressed in the previous section. The largest seller during the data collection period was a firm by the name of CSR Technologies (henceforth CSRT), which purchased large quantities of second-hand laptops for resale on eBay. CSRT’s product line was mostly made up of Dell Latitude laptops, which come in several different configurations. Because laptop and auction characteristics can be important determinants of the price a computer will receive, we attempt to homogenize our sample as much as possible.

CSRT’s most common laptop configuration, comprising 733 total auction listings, included an Intel Pentium 4M processor with a clock speed of 1.4GHz, 512MB of RAM, 30GB of hard drive, a DVD-ROM optical drive, a 14.1-inch screen, and with the Windows XP Professional operating system installed. All laptops in this sample were described by the seller as either “refurbished” or “used,” and all corresponding auctions lasted for 24 hours. We restrict our sample to only laptops sold by CSRT, which ensures constant seller reputation, exchange and upgrade policy, flat shipping rate (\$36), auction setup and so on. Moreover, CSRT used a template format for the display of each of its auction listings on eBay, making appearance

during the auction uniform as well. After these restrictions, the only auction characteristics which vary are the bidding data themselves. This unique and homogeneous sample of auctions will allow us to abstract away from further complications (such as unobserved heterogeneity) as we develop an empirical methodology to tackle the already formidable challenges inherent in ideal eBay data.

5.2 Sample Paring

For each auction, we have detailed bidding information, including the timing and amount of each bid submitted to the eBay server, as well as the identity of the bidder who submitted it. Looking into the bid history allows us to calculate the number of observed bid submissions and the number of observed unique bidders. The first empirical challenge we encounter is that bidders may play different, potentially non-equilibrium actions at various points in the auction; for example, submitting low, cheap-talk bids early on, and then later bidding based on best-response calculations resembling those in a sealed-bid auction. Empirically, a significant fraction of observed bid amounts, particularly those submitted early on in the life of the auction, fall too far below realistic transaction prices to be taken seriously.

On the other hand, in the vast majority of cases, the top two bids arrive within the final 30 minutes of the auction. Table 3 shows that the average time to end when the winning bid arrives is 24.04 minutes, and the median time to end is 1.43 minutes. For the highest losing bid, the mean and median time remaining are 24.29 and 3.10 minutes, respectively. Previous empirical work on eBay has established these phenomena as empirical regularities and our data are no different.

5.2.1 Intra-Auction Dynamics

This discussion hints at a need to deal with the issue of intra-auction dynamics on eBay in some way. We adopt an approach similar to that of Bajari and Hortaçsu (2003) by partitioning the auction run time into two stages.²⁴ Taking the total auction time to be T , we treat the first stage as an open-exit ascending auction played until $T - \varepsilon$; the final (terminal) stage, of length $\varepsilon < T$, is treated as a sealed-bid auction.²⁵ During the first stage of the auction bidders submit cheap-talk bids which convey little, if any, information on the likely sale price which will result from the auction.

As in Bajari and Hortaçsu (2003), bidders formulate a strategic sealed bid for submission during the terminal stage. We assume that this strategic bid ignores the outcome of the initial cheap-talk stage and the terminal stage is like a sealed-bid auction in the sense that bidders do not re-optimize their strategic bids in response to observed price-path dynamics during the terminal stage. Bidders may, however, submit these bids directly to the server to take advantage of eBay's automated proxy bidding capability, or they may choose to incrementally increase their bid submissions up to the level of their strategic bid on their own. In other words, we assume bidders formulate strategic bids based on the distribution of private values and their expectation of N at the beginning of the terminal period, and afterward they stick to their planned strategic bid through the end of the auction. Finally, consistent with the previous section, price-path dynamics within the terminal period are assumed to be simple in the sense that ordering of bidders' submission times is random rather than coordinated.

²⁴Bajari and Hortaçsu (2003) were motivated by the fact that eBay does not fit within Milgrom and Weber's [1982] "open-exit" ascending auction format where bidders' exit decisions are observable to their competitors. On the contrary, eBay bidders may re-join the auction at any time after an initial proxy bid is surpassed.

²⁵Some other work within the related literature, including Nekipelov (2007), has attempted to develop an explicit empirical model for determining both levels *and* timing of bids within an auction. One challenge to such an undertaking is that fully formed equilibrium theories of bid timing within an auction are rare. Such an exercise would introduce considerable complexity and is beyond the scope of our current purposes. Thus, we follow the standard approach of adopting assumptions which simplify intra-auction dynamics for tractability.

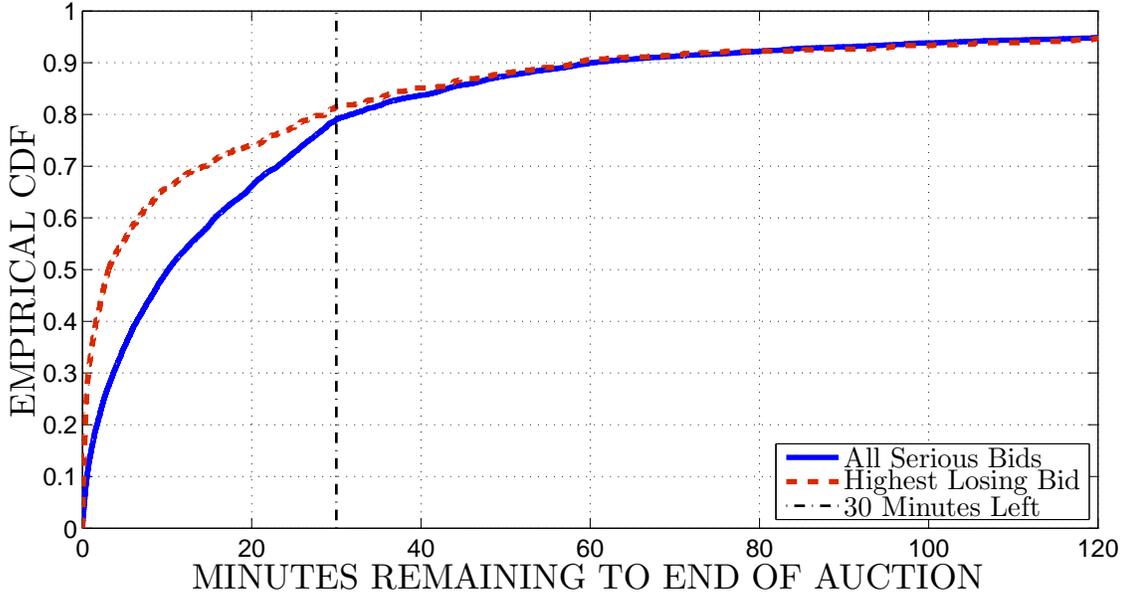


Figure 6: Empirical Distributions: Time Remaining when Bids are Submitted

5.2.2 Serious Bidders and Strategic Bids

This partitioning of the auction into an initial cheap-talk period followed by a terminal sealed-bid auction leads to the following definition: a *serious bid* is one which affects the price path within the terminal period; likewise, a *serious bidder* is one who is observed to submit at least one serious bid. This distinction allows for the possibility that some observed bidders early on in the life of the auction were merely “fishing for a steal” or casually dabbling, rather than seriously vying to win a laptop like other bidders who remain active close to the end. Of course, the possibility always exists that some bidders who are determined to be non-serious by the above criterion had serious intent to compete for the item, but were priced out before submitting a serious bid. This, however, is just part of the problem that our proposed estimator for λ based on our explicit model for the filter process is meant to solve: recall that observed participation is merely treated as a lower bound on actual participation when identifying the arrival rate of bidders within an auction.

In our empirical application, we specify the terminal period as the last 30 minutes of an auction. During this period, we see an average of 4.69 observed serious bidders. Figure 6 provides a justification for this choice. It shows two empirical distribution functions: one for time remaining when the highest losing bid is submitted, and one for time remaining across all serious bids. Note the highest losing bid arrives with fewer than 30 minutes left in over 80% of the auctions in our sample. Note also that 20.98% of all serious bids are submitted *prior* to the final 30 minutes of the auction. This possibility is naturally built into our serious bidder criterion, so as to avoid making too sharp a distinction between the two stages of the auction. Whenever there are at least two bidders over the life of an auction, the two highest bids during the cheap talk phase are the ones which set the initial price for the final 30 minutes. Therefore, the individuals who submitted them will be logged as serious bidders, even if they are not observed to actively bid during the final 30 minutes.

The majority of serious bidders only submit a single proxy bid, but just under one third of serious bidders are observed to raise their own bid levels incrementally with multiple submissions during the terminal stage. This discrepancy motivates our final data restriction: in order to be conservative we assume only that the highest losing submission (that is, the second-highest overall bid) is reflective of a strategic

Table 3: Descriptive Statistics

Variable	Mean	Median	St. Dev.	Min	Max	# Obs
Time Remaining (minutes)						
Winning Bid Submission:	24.04	1.43	97.50	0.00	1,392.65	733
High Loser Bid Submission:	24.29	3.10	61.76	0.00	633.40	733
Observed Participation						
\tilde{N} (serious bidders only):	4.69	4	1.70	2	11	733
Monetary Outcomes						
Sale Price (w/o shipping):	\$299.63	\$300	\$31.31	\$202.50	\$405	733
Highest Losing Bid:	\$295.42	\$299	\$31.20	\$200	\$400	733
First-Price Frequency:	22.24%	–	–	–	–	–

equilibrium bid consistent with the model from Section 4.1 above. Within the context of our simplified model of intra-auction dynamics, if some bidders log multiple bid submissions, we interpret such behavior as raising bid levels manually up to the point of their planned strategic bid (which we assume is not updated during auction runtime), rather than availing themselves of the automated proxy bidding system.²⁶ Thus, many submissions by serious bidders may still represent only lower bounds on their equilibrium strategy. In defense of our implicit trust in the highest losing bid, we appeal to the Haile and Tamer (2003) assumption that bidders will never allow an opponent to win at a price they would be willing to pay, where willingness to pay is interpreted here as a planned strategic bid, rather than a private valuation. By this logic one can at least be confident that the highest losing submission is reflective of a full strategic bid, and it will always be available since it is always recorded by Nature as the filter process runs its course.²⁷ Of course, if the researcher is able to incorporate more bidding data per auction in a reliable way, then it would improve the statistical precision of resulting estimates.

In total, our various sample restrictions leave us with 733 observed highest loser bids (that is, all auctions in our data had at least two serious bidders).²⁸ Table 3 displays descriptive statistics on bid timing, observed participation, sale prices, and highest losing bids. With probability 0.2224 the sale price takes on a value of exactly the highest bid for our sample. Note, too, that bids and sale prices do not include the flat shipping fee of \$36.

²⁶Our assumption that bidders choose not to update their planned strategic bids during auction runtime is based in the idea that electronic auctions exist in larger market places where a loser today may return tomorrow to bid again. Bodoh-Creed et al. (2016) find empirical evidence in support of this view; see Section 6.1 for a brief discussion on intra-auction dynamics.

²⁷Nature also always records the attendance of the winner, but the eBay bidding system only reports exact amounts for bids that were surpassed by the next lowest bid plus Δ . Thus, the winning bid itself is only observable to the econometrician for auctions where price is generated by a first-price rule, which occurs in roughly 22.24% of the auctions in our sample. In principle, inclusion of these bids may improve the statistical efficiency of the estimator, but it would come with added complexity and computational cost. For simplicity sake, we ignore them here. Despite this additional data loss, the confidence bounds we get on our estimates are still remarkably tight.

²⁸In addition, we used one final sample restriction in order to avoid numerical instability issues in the upper tail of the private value distribution. The original dataset contained 736 observations fitting the description above, but for three of these, the highest losing bid was four standard deviations or more above the mean. We drop these three auctions from the sample, leaving us with 733 total auctions.

5.2.3 Practical Matters

As a final practical matter, we had to choose two knot vectors \mathbf{j}_J and \mathbf{k}_K in order for the B-spline basis functions \mathcal{V}_j , $j = 1, \dots, J + 3$ and \mathcal{B}_k , $k = 1, \dots, K + 3$ to be well-defined. The first challenge is to specify the endpoints. Since the sample extrema are superconvergent estimators of the support endpoints, and since theory requires that $\beta(\underline{v}) = \underline{v}$, we set $j_1 = k_1 = \min_t\{y_t\}$ and $k_{K+1} = \max_t\{y_t\}$. Because the bid function does not touch the 45°-line at the upper end, we can only bound \bar{v} from below and must, therefore, make a guess at an appropriate value. We chose $j_{J+1} = \bar{b} + 2\Delta_2$ as an approximate value for the upper bound of the private value support. Our point estimates, discussed below, suggest that this was a reasonable choice which does not seem to drive results in any meaningful way.

The next step is to select the number and placement of the knots. The primary concern in selecting K is to minimize numerical error in the approximate solution to the equilibrium differential equation. We chose a grid of 48 knots to be placed uniformly over the set $[\underline{b}, b^* + \Delta_1 - 4\epsilon] \cup [b^* + \Delta_1 + 8\epsilon, \bar{b}]$, with $\epsilon = (\Delta_2 - \Delta_1)/4$. We also inserted an intermediate knot grid $\{b^* + \Delta_1 - 3\epsilon, b^* + \Delta_1 - 2\epsilon, \dots, b^* + \Delta_1 + 7\epsilon\}$, with 12 additional subintervals, to provide extra flexibility for the transition region where the bid increment changes from Δ_1 to Δ_2 , but without projecting the influence of this region to points far away. This gives us $K = 60$ total knots and 63 total basis functions for the inverse bid function approximant $\hat{\xi}$. The vector of points where equilibrium conditions were enforced was a uniform grid of 500 points on $[\underline{b}, \bar{b}]$. From experimentation, this configuration of knots and checkpoints seemed to deliver a reasonable trade-off between computational cost and numerical accuracy: larger values of either K or L do not improve model fit or numerical accuracy in any meaningful way.

As for the principal knot vector \mathbf{j}_J , the primary concern is goodness of fit to the distribution of the observables. One challenge is first to reduce the dimensionality of the decision problem, in a data-driven way, if possible. Our proposal is to specify a uniform grid of J quantile ranks $\{q_1, \dots, q_{J-1}, q_{J+1}\}$ spanning $[0, 1]$, and then to map them into bid space using the empirical quantile function $\hat{H}^{-1}(q)$. We then replace the uppermost knot j_{J+1} with a value of $(\bar{b} + 2\Delta_2)$, as mentioned above, in order to account for the fact that these knots will govern the behavior of \hat{F}_V in private value space. Finally, because our highest loser bid distribution is skewed to the right with a long, curved upper tail, we insert an additional knot $j_J = (j_{J-1} + j_{J+1})/2$ to provide some added flexibility in this region. The advantage of this approach is that it reduces the $(J + 1)$ -dimensional decision of knot choice to a single decision of J , the number of knots, while letting the data determine the placement of the knots. On the other hand, this method concentrates knots more densely within higher quantiles of \hat{F}_V , since we are choosing knots at the quantiles of highest losing bids. However, this is a natural problem inherent in estimating a parent distribution from any dataset consisting of order statistics: observations selected from the within-auction sample extremes will always be more informative of the higher quantiles of the parent distribution. Thus, in finite samples, knot choice involves the familiar bias-variance tradeoff. We propose the above method because it allows the observables to determine knot location. Statistically optimal knot choice, while an interesting problem, is beyond the scope of the current exercise, and is therefore left to future work.²⁹

Finally, when enforcing the equilibrium conditions, we found the L_∞ -norm most effective. The problem with the other L_p norms (with finite p ; for example, the L_2 -norm) is that they allow for a small number of drastic deviations from the FOCs, so long as the solution to $\hat{\xi}$ is well-behaved at most domain points. This can lead to poor performance of the solution method. Although both methods produce a spline estimate with the same overall trend, the L_∞ -norm, by directly disciplining the worst-behaved segments of the spline, is able to achieve a better overall fit, while avoiding wild oscillations around the bidder optimality

²⁹By experimenting with alternatives where knots are placed uniformly in V space instead, we found that the shape of the point estimate for \hat{f}_V changed somewhat near the lower end, but variance of the estimator increased significantly. The simple intuition is that with uniform knots in V space there are several model parameters whose values are being determined by sparse data near the lower extreme of the sample, which is likely leading to higher mean squared error.

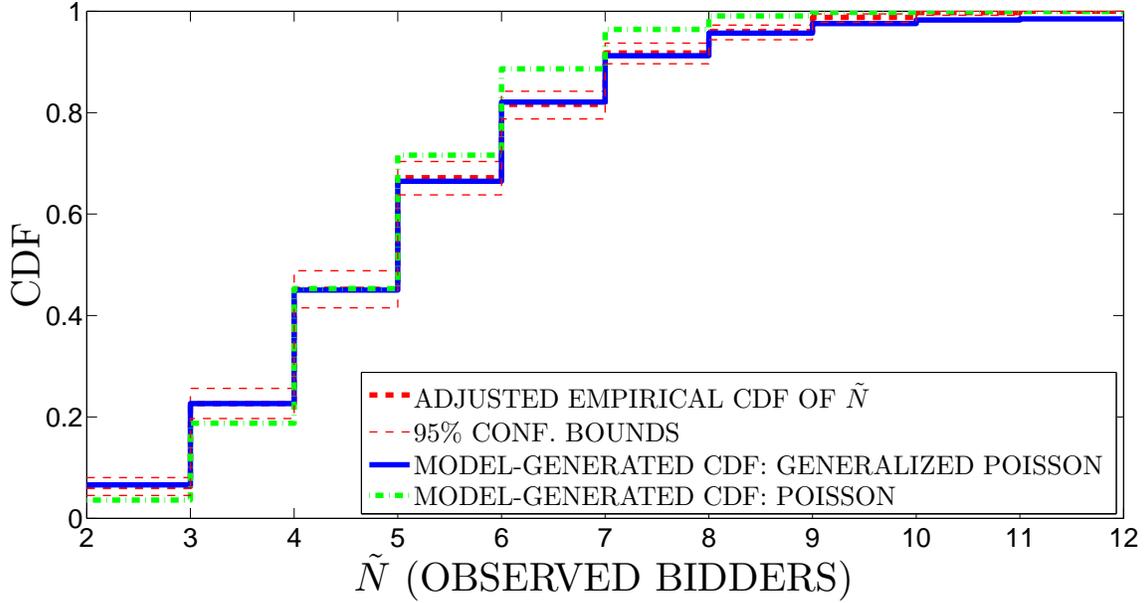


Figure 7: MODEL FIT: Empirical CDF of \tilde{N} vs. Parametric Model-Generated CDFs

condition. As an illustration of this point, Figure A.19 in the Online Appendix plots the relative approximation error of the B-spline approximation, $R(b)/\hat{\xi}(b; \hat{\alpha})$ under the L_∞ and L_2 norms. The plot shows that the approximation error under the L_∞ norm is several orders of magnitude smaller along much of the functional domain. Moreover, few remedies exist to improve this comparison without drastically increasing computational load; for example, the picture remains virtually unchanged if we double the number of domain checkpoints. Thus, we recommend the L_∞ norm for practical use. We chose values of $P_0 = 10$ and $P_1 = P_2 = P_3 = 200$ as this allowed for a good balance of least squares fit and small numerical error.

5.3 Empirical Results

We first separately estimated two parametric models of the bidder arrival process—a generalized Poisson model indexed by $(\lambda_1^{gP}, \lambda_2^{gP})$ and a standard Poisson model indexed by λ^P (where the second parameter in the generalized model is restricted to be zero). Under the generalized Poisson model we get both a higher mean, 11.569, and standard deviation, 6.832, of the random variable N as compared to the Poisson model with 10.023 and 3.166, respectively. Figure 7 depicts a comparison of the empirical distribution of \tilde{N} versus the two parametric models. Both fit the data fairly well, but the generalized Poisson model emerges as the clear winner: on the majority of the sample domain the model-generated distribution of observables under the generalized Poisson model is within the nonparametric 95% confidence bounds of the adjusted empirical distribution of observed bidders. This implies that, although there still may be asymptotic gains from a more flexible functional form, the current sample size precludes further improvement without more data.

For the private value distribution we settled on a value of $J = 11$, meaning 12 knots in V space and 14 total parameters, with 12 of them being free parameters after we enforce the boundary conditions on \hat{F}_V . This choice seemed to provide a high degree of flexibility for fitting patterns in the data, while producing a high degree of numerical accuracy. Table 4 displays our chosen knots $\{j_1, \dots, j_{12}\}$, parameter point estimates, and bootstrapped standard errors in parentheses. Figure 8 conveys how the model specifications $J = 9, 10, \dots, 13$ all fit the data remarkably well. In selecting $J = 11$ we conducted a bootstrap exercise

Table 4: Parameter Estimates and Standard Errors

$\hat{\lambda}_1$	$\hat{\lambda}_2$	j_1	j_2	j_3	j_4	j_5	j_6	$\hat{\mu}_1 = 0$	$\hat{\mu}_{14} = 1$
5.760	0.502	200.00	255.08	267.86	278.18	287.19	294.99	$\hat{\mu}_2$	$\hat{\mu}_3$
(0.424)	(0.049)	j_7	j_8	j_9	j_{10}	j_{11}	j_{12}	0.0291	0.3443
		301.87	309.10	318.90	337.78	373.89	410.00	(0.0513)	(0.0547)
$\hat{\mu}_4$	$\hat{\mu}_5$	$\hat{\mu}_6$	$\hat{\mu}_7$	$\hat{\mu}_8$	$\hat{\mu}_9$	$\hat{\mu}_{10}$	$\hat{\mu}_{11}$	$\hat{\mu}_{12}$	$\hat{\mu}_{13}$
0.6345	0.7300	0.7775	0.8206	0.8621	0.8994	0.9374	0.9593	0.9897	0.9989
(0.0201)	(0.0152)	(0.0126)	(0.0108)	(0.0079)	(0.0064)	(0.0052)	(0.0049)	(0.0050)	(0.0016)
ε_0	ε_1	ε_2	ε_3	$SSR _{J=11}$					
9.8×10^{-6}	3.4×10^{-6}	8.6×10^{-5}	9.9×10^{-5}	0.1194					

to compare the mean integrated squared error across these specifications.³⁰ All of them implied virtually identical pointwise mean values for the CDF, meaning that there should be no meaningful difference in bias. Therefore, the integrated mean squared error comparison reduces to one of integrated variance, which is lowest for specification $J = 11$. Figure A.18 in the Online Appendix also shows that the relative numerical errors under the different model specifications are all quite small and roughly the same.

Figure 8 illustrates the fit between the empirical CDF $\hat{H}(b)$ (thick dashed line), and the model-generated version under varying choices for J . The thick solid line corresponds to our preferred specification with $J = 11$. Figures 9 and 10 depict the estimated private value distribution and inverse bid function with bootstrapped 95% confidence bounds.³¹ In each figure, the thin vertical lines indicate knots. The inverse bid function plot is zoomed in below the 90th percentile of the private value distribution so that the features of the function can be seen more clearly. Note the abrupt transition in the degree of demand shading above the point $b^* = \$250$. The confidence bounds on the inverse bid function are remarkably tight because within the relevant ranges for Δ , private value distributions that significantly differ tend to still largely agree on the appropriate degree of demand shading. For completeness, Figure A.18 in the Online Appendix presents the relative approximation error, or the ratio of the residual function $R(b)$ to the units of $\hat{\xi}(b; \hat{\alpha})$. As suggested in Section 4.5.2 above, we also compute a final check on our solution for $\hat{\xi}(b; \hat{\alpha})$ by holding the structural parameters $\hat{\lambda}$ and $\hat{\mu}$ fixed, and computing a standard Galerkin approximant, call it $\tilde{\xi}(b; \tilde{\alpha})$, by solving a square system of non-linear equations on an appropriately chosen grid of checkpoints. This we do with a knot vector $\tilde{\mathbf{j}}_{\tilde{J}}$ where $\tilde{J} = 5J$ for improved accuracy (under the fixed point estimates for distributional parameters), and we then compare the two solutions; see Figure A.20 in the Online Appendix. The knot vector we selected followed the same rule as the one used in estimation, just with more knots; the grid of checkpoints included the midpoint of each subinterval, the endpoints of the knot vector, and a final point chosen half way between the lowermost knot and the lowermost midpoint. The sup-norm distance between our runtime ODE solution $\hat{\xi}(b; \hat{\alpha})$ and the ex-post, standard Galerkin implementation $\tilde{\xi}(b; \tilde{\alpha})$ (using a square system of equations) turns out to be twenty cents, but over the majority of the

³⁰A broader question concerns the optimal rate at which model complexity should increase as the sample size grows. Given our rule for knot placement (uniform in quantile rank space), could one construct a consistent estimator which includes a rule for choice of $J \rightarrow \infty$ as a function of sample size T ? Given that the Stone–Weierstrass Theorem applies to B-splines as it does to global polynomials, the answer is likely yes. However, it is beyond the scope of this work, as is the question of optimal rate of increase for J , and is left to future research.

³¹To compute our bounds we estimated the model parameters λ , μ , and α on 1,000 bootstrapped samples, holding the knot vectors \mathbf{k}_K and \mathbf{j}_J fixed.

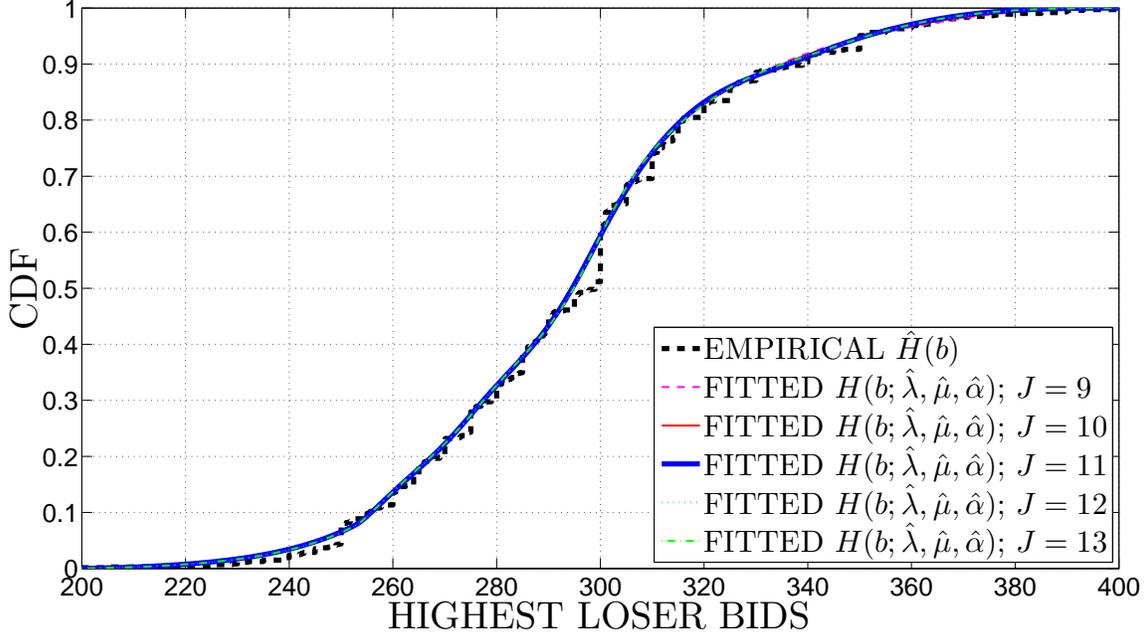


Figure 8: MODEL FIT: Empirical $H(b)$ and Estimated $H(b)$ with $J = 9 - 13$

domain, the two differ by less than five cents. Figure A.20 displays the runtime solution with a thick, dark line, and the ex-post solution with a thin, light line so that the two can be visually distinguishable. This comparison indicates that our point estimates for the private value distribution are based on a differential equation solution with an acceptable level of numerical error.

It is interesting to note that both the point estimate $\hat{\xi}(b; \hat{\alpha})$ and the lower confidence bound lay significantly above the 45°-line, which represents the hypothetical inverse bid function under a second-price equilibrium. To put the picture into context, the upper panel of Table 5 summarizes the degree of estimated demand shading for bids occurring at various quantiles of the observable distribution $H(b)$. Within the 90-10 range, the difference between the estimated EA bid function and the second-price bid function ranges from \$3.41 to \$5.29. Note that the information on bid shading in the table implicitly represents the magnitude of the mis-specification problem from ignoring the hybrid EA pricing rule as well: under a second-price assumption, the econometrician would estimate the parent distribution of bids using our model of the filter process, and then incorrectly interpret G_B as being the same as F_V since bids and private values are the same under a second-price rule. Our approach is to incorporate equations (12) and (8) into the estimation step as a correction factor. The numbers from Tables 3 and 5 indicate that failing to do so would result in mis-specifying point estimates of private valuations by between 11% and 17% of a standard deviation of the observable distribution H . It is also worth mentioning that if we compute an alternative B-spline estimator, call it \tilde{F}_V , using the same knot vector as our preferred specification (with $J = 11$) and the generalized Poisson point estimate $\hat{\lambda}^{sp}$, but under the second-price assumption (where we just interpret observed bids as private values), then we get a CDF estimate which is strongly rejected by the EA bidding model. Specifically, the estimator \tilde{F}_V is shifted to the left and parts of it lay outside of the bootstrapped 99% confidence bounds of $\hat{F}_V(v; \hat{\mu})$.³²

³²See Figure A.17 in the Online Appendix for a robustness check in which we compare the bias induced by pricing-rule mis-specification with the bias resulting from deviations in the scale invariance assumption.

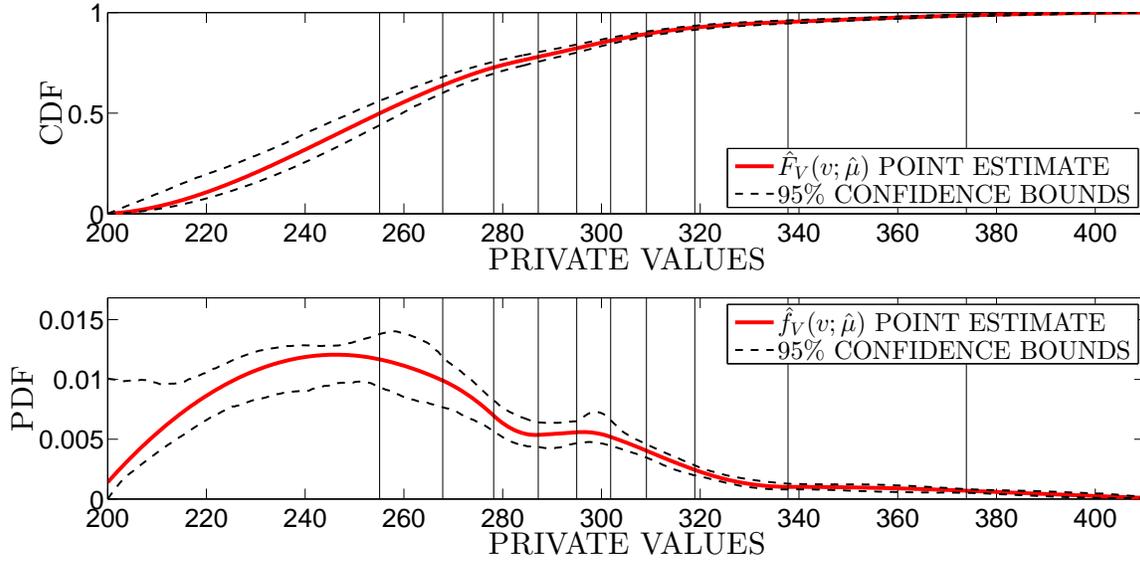


Figure 9: Estimated Valuation Distribution and Density with Knots

Table 5: Absolute and Relative Differences Between Private Values and Bids

	10 th	25 th	Median	75 th	90 th
	Percentile	Percentile	Percentile	Percentile	Percentile
High Loser Bid:	\$255.50	\$275.00	\$299.00	\$311.86	\$339.00
Bid Shading: $V - \beta(V)$	\$3.41	\$4.61	\$4.54	\$5.29	\$4.61
	10 th	25 th	Median	75 th	90 th
Buyer Information Rents	Percentile	Percentile	Percentile	Percentile	Percentile
$\mathcal{I} = V_{(1:N)} - \min\{B_{(2:N)} + \Delta, B_{(1:N)}\}$	\$5.01	\$9.19	\$22.51	\$44.46	\$69.36

5.4 Model Simulations and Counterfactual Analysis

We begin by exploring auction winners' market rents arising from private information on their willingness to pay. We define information rents as the following random variable

$$\mathcal{I} \equiv V_{(1:N)} - P = V_{(1:N)} - \min\{B_{(2:N)} + \Delta, B_{(1:N)}\},$$

or the winner's private valuation minus the price she pays. Note that the randomness in \mathcal{I} comes from it being a function of three separate random variables: N , $V_{(1:N)}$, and $B_{(2:N)}$. In our case, where we only occasionally observe the highest bid, and where we never directly observe N , our structural model estimates are needed in order to simulate the distribution of \mathcal{I} .

To do so, we proceed in three steps: first, we generate a random value of N from the generalized Poisson distribution with our point estimate $\hat{\lambda}^{gp}$; second, we generate N random draws from the private value distribution $\hat{F}_V(v; \hat{\mu})$ and map them into bid space using a cubic interpolant of the inverse of $\hat{\xi}(b; \hat{\alpha})$; finally, we collect the highest private valuation and the two highest bids to compute a simulated value for \mathcal{I} . We followed this process 1,000,000 times to get a large enough sample to reliably represent the moments of the distribution of \mathcal{I} . The lower panel of Table 5 summarizes information rents at several

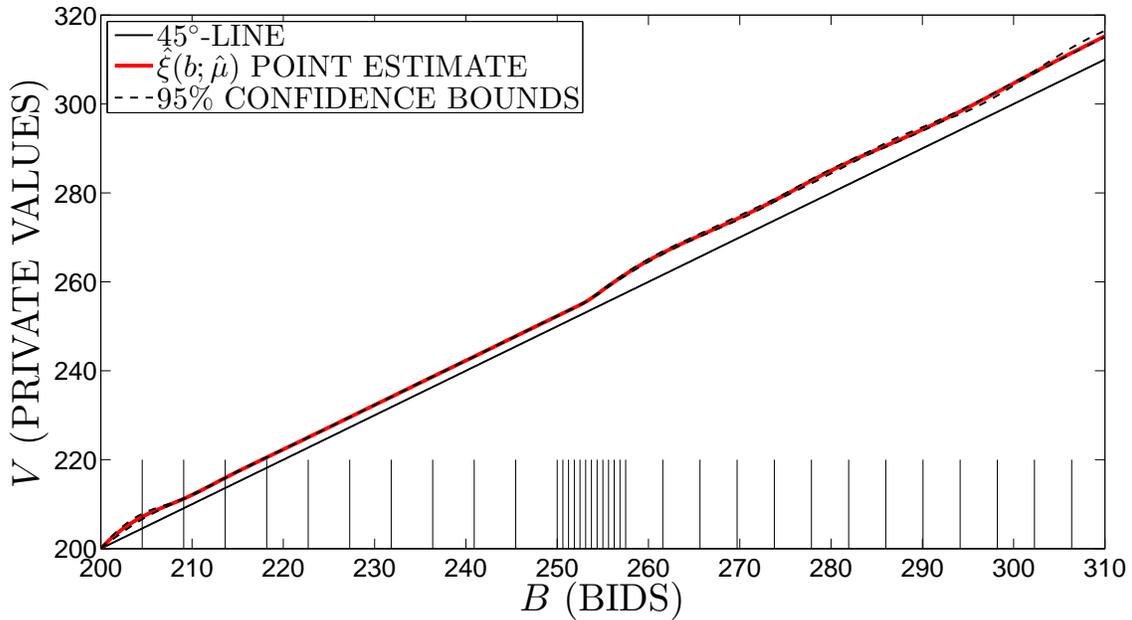


Figure 10: Estimated Equilibrium Inverse Bidding Function

quantiles of the winner distribution. On average, the winner in an auction from our sample retained an information rent of \$30.57, but the distribution is highly skewed to the right with a large standard deviation of \$26.36. From a seller’s perspective, one of the main concerns is how to maximize their own revenues by extracting as much of these information rents as possible through auction design. We now conclude our discussion with a counterfactual exploration of optimal reserve prices.

One puzzling regularity within eBay data is that the vast majority of sellers set reserve prices at or near zero. In the 733 auctions from our empirical application, a reserve price is never used, but this could represent default seller policy (given that all of these auctions come from the same seller, CSRT). Nevertheless, in a broader set of used laptop data from the same time period, with 13,193 separate auctions, only 724 of them, or 5.5% overall, involved a reserve price that exceeded \$1; 359, or 2.6%, had a reserve price that exceeded \$10; and 156, or 1.1%, had a reserve price that exceeded \$100. The sparse use of reserve prices might be surprising given the literature on optimal mechanisms in which a positive reserve price is suggested even for a seller who values the item at \$0. Equipped with our estimates of the private value distribution and the parameters of the generalized Poisson distribution, we can consider counterfactual experiments to understand why.

First though, following McAfee and McMillan (1987) and Harstad et al. (1990), observe that because we have a symmetric, independent private-values model with an unknown number of risk-neutral bidders, standard auctions will be revenue equivalent so long as bidders have the same beliefs about the number of potential bidders. Our EA model has the same equilibrium allocation rule as the canonical first- and second-price formats—the bidder with the highest valuation always wins the object (efficiency)—and the expected payoff of the lowest possible type \underline{v} is zero; therefore, the revenue equivalence principle applies. Second, optimal auction design does not depend on the distribution of N ; see McAfee and McMillan (1987). Taken together, these facts imply that the optimal auction can be implemented using a first- or second-price format, or via an EA.

This simplifies computation of counterfactuals: we can generate the correct expected revenues by considering a second-price auction, which avoids the need for solving the EA bid function using our estimated

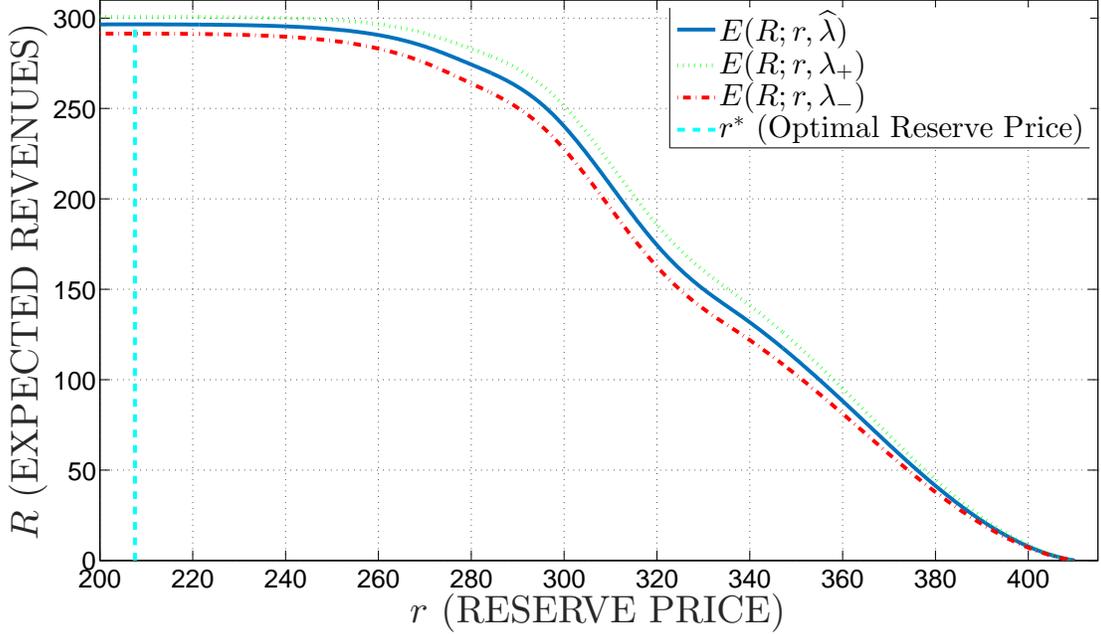


Figure 11: Counterfactual Revenue Curve Comparisons

valuation distribution for every possible choice of reserve or value of the participation parameters. Let R denote revenue, and r denote reserve price. Expected revenues are given by

$$\begin{aligned}
 E(R|r) = r[1 - F_V(r)] \sum_{n=1}^{\infty} \rho_N(n; \lambda) n F_V(r)^{n-1} \\
 + \int_r^{\bar{v}} v[1 - F_V(v)] f_V(v) \sum_{n=1}^{\infty} \rho_N(n; \lambda) n(n-1) F_V(v)^{n-2} dv,
 \end{aligned} \tag{20}$$

where the integral is solved using an adaptive recursive Simpson's rule. The maximizer, r^* , of this expression satisfies the equation

$$r^* = v_0 + \frac{1 - F_V(r^*)}{f_V(r^*)}, \tag{21}$$

first derived by Myerson (1981), where v_0 is the seller's valuation. For the purpose of this discussion, we shall assume throughout that $v_0 = 0$.

Figure 11 depicts three expected revenue curves as a function of reserve price r : one deriving from model point estimates (solid line), and two hypothetical curves where the generalized Poisson parameters are altered. For both of these we compute an alternative λ vector so that $\text{Var}[N]$ is held constant while for one, call it λ_+ , $E[N]$ increases by one relative to the point estimates, and for the other, call it λ_- , $E[N]$ decreases by one. A well-known result concerning reserve prices in auctions by Bulow and Klemperer (1996) is that the benefit from attracting an additional bidder to an auction exceeds any possible gain by optimizing the reserve price. This result plays out strongly in the figure: the impact on expected revenues from changing $E[N]$ by one is hundreds of times greater than the impact of moving from $r = 0$ to $r = r^*$. However, our empirical results say something even stronger about optimal reserve prices on eBay: they are, in fact, almost entirely irrelevant to begin with! Although the expected revenue curves in the figure have an interior global optimum, under our model estimates, the maximum benefit a seller may reap by optimizing the reserve price is estimated to be \$0.0288. To put this number in perspective, if we assume

that an eBay seller values her time at a conservative \$10/hour, then if it takes her longer than a mere 11 seconds to decide what the optimal reserve price should be, she would be better off by simply setting $r = 0$ and allocating her time toward some other, more profitable use. This resolves the puzzle of why the vast majority of eBay sellers choose non-binding reserve prices: the online auction house has done its job of attracting buyers to the market well enough so there is no longer any need to worry about this aspect of auction design.

6 Discussion and Conclusion

6.1 Model Extensions

Two possible extensions are worth comment before we conclude. Our estimation framework presented here assumes away two common aspects of online auction markets: binding reserve prices and inter-auction, dynamic incentives. Specifically, we have assumed here that sellers set reserve prices at a value of zero, as this was consistent with our empirical application to laptop data. We have also shown in our counterfactual analysis that if sellers have reservation values of zero, then they have little to gain by optimizing this particular aspect of auction design. However, in many datasets involving online auctions, nontrivial fractions of sellers do choose binding reserve prices. From an estimation standpoint, this addition introduces another source of sample selection: some actual bidders who would otherwise have been observed are withheld from the econometrician’s view because they happen to have low private valuations. Incorporating binding reserve prices into the filter process involves significant complications of both the model and estimator—it is no longer possible to estimate λ separately from G_B since both objects influence the link between the distributions of N and \tilde{N} —and is explored in ongoing work by Bodoh-Creed et al. (2016). Another important assumption underlying our framework presented here is the idea that individual auctions may be treated in isolation, as a series of static, one-shot games. In reality though, online auction markets for many consumer products are more complicated environments where bidders’ opportunity costs of losing are not the full amount of their private valuations, since they can always return in a future period if they lose an auction today. This positive continuation value gives rise to an additional source of demand shading from intertemporal incentives. Still, much can still be learned from a framework like ours which characterizes only purely static strategic demand shading. Bodoh-Creed et al. (2016) develop a model of dynamic platform markets like eBay. They characterize equilibrium bidding behavior and show that in the dynamic platform setting there is an intuitive layering of the intertemporal and static demand shading incentives. Their results imply that our identification and estimation strategy, based on a static, one-shot bidding model, can be applied in broader settings where bidding is influenced by the future outside option.

6.2 Conclusion

EAs are important market mechanisms in the world today, with eBay in 2010 accounting for total sales of \$25 billion from auction listings alone.³³ In this paper, we have made four contributions to research on EAs: First, we documented how failing to account for the nonstandard EA pricing rule can bias estimates of the latent valuation distribution. Second, we demonstrated that a realistic EA bidding model is nonparametrically identified by observables readily available. As part of this exercise, we solved another, independent problem by proposing a new identification strategy for inferring the distribution of N

³³See the eBay Annual Report 2010 available at <http://investor.ebayinc.com/annuals.cfm>. Our calculation is based on the gross merchandise value (GMV) multiplied by the reported share of GMV obtained under an auction-style format.

using observable lower bounds on bidder participation. Third, we proposed an estimation strategy to recover model parameters from actual eBay data in a flexible yet computationally tractable way. Although conceptually complex, the method we employed is easily implementable using a MATLAB toolbox that is available from the authors on request. Thus, this research fills two important methodological gaps in the literature: (1) researchers no longer need to assume (incorrectly) that eBay bids are private valuations; (2) having estimates of both the latent valuation distribution and the *ex ante* participation rate permits model simulations, where previous work only managed to identify the latent valuation distribution. Finally, we used our framework to shed light on empirically relevant questions in online auction design. For example, our model estimates explain why reserve prices are commonly not used by sellers: given the number of bidders that typically participate, they play very little role in determining revenues. On the other hand, varying the expected number of bidders that participate can play a significant role under observed market conditions. This suggests that the most relevant aspects of online market design take place at the level of the auction house itself, where policy levers exist to produce movements on the margin of market-wide, buyer-seller mix. Future research to investigate this question further will surely yield fruitful new insights, and will benefit from the methodological advances we developed here.

References

- Christopher P. Adams. Estimating demand from eBay prices. *International Journal of Industrial Organization*, 25:1213–1232, 2007.
- Patrick Bajari and Ali Hortaçsu. The winner’s curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions. *RAND Journal of Economics*, 34:329–55, 2003.
- Patrick Bajari and Ali Hortaçsu. Economic insights from Internet auctions. *Journal of Economic Literature*, 42:457–486, June 2004.
- Aaron Bodoh-Creed, Joern Boehnke, and Brent Hickman. How efficient are decentralized auction platforms? *Becker Friedman Institute for Research in Economics Working Paper No. 2016-22*, 2016.
- Jeremy Bulow and Paul Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1): 180–194, 1996.
- P. C. Consul and G. C. Jain. A generalization of the Poisson distribution. *Technometrics*, 15(4):791–799, 1973.
- Carl de Boor. *A Practical Guide to Splines Revised Edition*. New York: Springer-Verlag, 2001.
- Emmanuel Guerre, Isabelle Perrigne, and Quang H. Vuong. Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68:525–574, 2000.
- Philip A. Haile and Elie Tamer. Inference with an incomplete model of English auctions. *Journal of Political Economy*, 111:1–51, 2003.
- Ronald M. Harstad, John H. Kagel, and Dan Levin. Equilibrium bid functions for auctions with an uncertain number of bidders. *Economics Letters*, 33(1):35–40, 1990.
- Brent R. Hickman. On the pricing rule in electronic auctions. *International Journal of Industrial Organization*, 28:423–433, 2010.

- Brent R. Hickman and Timothy P. Hubbard. Replacing sample trimming with boundary correction in nonparametric estimation of first-price auctions. *Journal of Applied Econometrics*, 35:739–762, 2015.
- Brent R. Hickman, Timothy P. Hubbard, and Harry J. Paarsch. Investigating the economic importance of pricing rule mis-specification in electronic auctions, typescript, Department of Economics, University of Chicago, 2015.
- Ali Hortaçsu and Eric R. Nielsen. Commentary: Do bids equal values on eBay? *Marketing Science*, 29: 994–997, 2010.
- Timothy P. Hubbard and Harry J. Paarsch. On the numerical solution of equilibria in auction models with asymmetries within the private-values paradigm. *Handbook of Computational Economics*, volume 3, edited by Karl Schmedders and Kenneth L. Judd, 37–115, 2014.
- Bernie L. Hulme. One-step piecewise polynomial galerkin methods for initial value problems. *Mathematics of Computation*, 26(118):415–426, 1972.
- David Lucking-Reiley. Auctions on the Internet: What’s being auctioned, and how? *Journal of Industrial Economics*, 48:227–52, 2000.
- R. Preston McAfee and John McMillan. Auctions with a stochastic number of bidders. *Journal of Economic Theory*, 43(1):1–19, 1987.
- R. Preston McAfee and John McMillan. Analyzing the airwaves auction. *Journal of Economic Perspectives*, 10:159–175, 1996.
- Paul R. Milgrom and Robert J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50: 1089–1122, 1982.
- Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:55–73, 1981.
- Denis Nekipelov. Entry deterrence and learning prevention on eBay. *mimeo*, 2007.
- John G. Riley and William F. Samuelson. Optimal auctions. *American Economic Review*, 71:381–392, 1981.
- Alvin E. Roth and Axel Ockenfels. Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the Internet. *American Economic Review*, 92:1093–1103, 2002.
- Fritz W. Scholz and Michael A. Stephens. K -Sample Anderson–Darling Tests. *Journal of the American Statistical Association*, 82:918–924, 1987.
- Unjy Song. Nonparametric estimation of an eBay auction model with an unknown number of bidders, typescript, Department of Economics, University of British Columbia, 2004a.
- Unjy Song. Nonparametric identification and estimation of a first-price auction model with an uncertain number of bidders, typescript, Department of Economics, University of British Columbia, 2004b.
- Che-Lin Su and Kenneth L. Judd. Constrained optimization approaches to estimation of structural models. *Econometrica*, 80(5):2213–2230, 2012.

William S. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

Robert Zeithammer and Christopher P. Adams. The sealed-bid abstraction in online auctions. *Marketing Science*, 29:964–987, 2010.

Olgierd C. Zienkiewicz and Kenneth Morgan. *Finite Elements and Approximation*. Mineola, NY: Dover Publications, 2006.