Empirical Analysis of Dynamic Bidding on eBay *

Denis Nekipelov†

August 2008

Abstract

Despite their relatively simple structure, internet auctions demonstrate a puzzling discrepancy between the predictions of standard auction theory and the actual behavior of bidders. In particular, despite the second-price structure of eBay auctions, eBay bidders frequently bid multiple times over the course of a single auction and cluster their bids at the very end. Optimal behavior in a private-value second-price auction implies that bidders submit a single bid. In this paper, I show that the coexistence of multiple, contemporaneous auctions for similar items coupled with uncertainty regarding rival entry can explain these features. I construct a continuous-time auction model with endogenous entry, establish non-parametric identification, and provide a method for full structural estimation. I use the method of indirect inference for construction of the semi-parametric estimation procedure. Empirical estimates using eBay auctions of pop-music CDs confirm that entry indeed depends on price. I then test the model against alternative explanations of observed bidding behavior using a detailed field experiment. The methods I develop for continuous-time analysis, identification, and estimation are applicable to many other settings with large dynamic markets.

*Denis Nekipelov, Department of Economics, University of California, Berkeley, e-mail: nekipelov@econ.berkeley.edu

†I am deeply indebted to my advisor Han Hong as well as Pat Bajari, Chris Timmins, Paul Ellickson, and Curt Taylor for their help and continuous encouragement. I would like to thank Peter Arcidiacono, Pat Bayer, Arie Beresteanu, Jesus Fernandez-Villaverde, Marjorie McElroy, and Juan Rubio-Ramirez for their suggestions and fruitful discussions, Drew Beauchamp and Ralph Boleslavsky for their comments on the earlier draft of this paper, and all participants of the Microeconometrics and the Microeconomic Theory lunches at Duke as well as seminars at Princeton, MSR, Berkeley, Stanford GSB, UPenn, Carnegie Mellon, Chicago GSB, Rice, Penn State, Georgia State, Minnesota, Washington University, NYU, Columbia, UCLA, and MIT. I am grateful to MIRC for financial support, and to Jason Blevins and Anna Kormilitsyna for tremendous help with the field experiment.

1 Introduction

Since it went online in 1995, eBay has both provided consumers an efficient market for the exchange of goods and researchers a fertile ground for testing academic theories. Although empirical evidence from this rich environment has validated many aspects of standard auction theory, there are some important features of bidding behavior that are inconsistent with the existing literature. In particular, it has been widely observed that the bidders on eBay tend to both bid multiple times for the same object and cluster their bids in the final seconds of a given auction[1]. This behavior is inconsistent with the theoretical predictions for a standard static second price auction, the mechanism which is used on eBay[2]. Under this mechanism, bidders should simply bid their reservation price.

I argue that the observed behavior follows directly from a defining feature of the eBay environment: the existence of many simultaneous auctions. Because similar items are frequently auctioned in close temporal proximity, the number of entrants into a particular auction will depend on price. Inter-dependence between price and entry pair with a specific structure of uncertainty in the market generates a non-linear stochastic price dynamics.

The goal of this paper is not to create a new auction model. I argue that consideration of the problem in continuous time and applying the results developed in the theoretical literature regarding the dynamic *public perfect equilibria* provides useful insights for the market structure. Moreover, considering the eBay environment from this perspective I find that it is possible to implement an estimator for the model with a considerable out-of-sample predictive power.

Employing a novel estimation method based on a form of indirect inference, I estimate the structural parameters of the model using data from a single auction category (pop music CDs) and then perform a controlled experiment aimed at identifying the impact of entry deterrence on optimal bidding behavior. The idea underlying the experiment is based on the fact that the incentive to deter entry should depend on the thickness of the market. In large markets (with the same number of bidders), entry into individual auctions should be significantly more sensitive to price than in small markets. Therefore, if we exogenously increase the number of objects being auctioned, we should observe a significant increase in early bidding. Consistent with the predictions of my model, I find that early bidding is indeed significantly more common when the number of contemporaneous auctions is increased.

This paper contributes to both the auction literature and a growing field of research aimed at

[1] See, for example, Bajari and Hortacsu (2003).

[2] eBay.com auctions utilize a second price format: the winning bidder is the bidder with the highest bid, while the price of the object is equal to the second highest bid. In a static auction model with independent valuations, the optimal strategy is to submit a single bid equal to the bidder's valuation.

estimating dynamic games. From a theoretical perspective, this paper emphasizes a fundamental difference between stand-alone auctions and those that take place in markets. In the latter case, the strategic behavior of individual bidders depends on the characteristics of every available auction, as well as the structure of information regarding potential rivals. In particular, this interdependence means that entry will depend on price. Of course, information regarding potential rivals is limited: bidders must infer the number of current entrants and their valuations. These features yield a rich and very complex strategic game.

Although solving the full strategic game is intractable given the dimension of the state space, imposing a few restrictions simplifies the problem considerably. In particular, I have chosen to represent entry using a reduced form Poisson process with price-dependent frequency. In particular, I approximate the full auction game with a competitive-type mechanism in which bidders directly influence only the price, influencing the payoffs of the other bidders indirectly. I capture the role of uncertainty by allowing for unobserved heterogeneity across auctions using a one-dimensional "visibility" parameter. Bidders have different a priori information regarding the visibility of the auction, but they can learn about it during the auction by observing the price path.

Although there is no closed-from solution for the bidding function, by formulating the model in continuous time, I am able to exploit a convenient linear algorithm for computing optimal strategies, thereby simplifying the computational burden significantly. This method offers an advantage over traditional value function iterations because the number of steps is fixed. More importantly, it provides a tractable avenue for estimation based on a form of indirect inference.

On the empirical front, this paper provides a novel technique for the full structural estimation of a continuous-time auction that can easily be extended to other dynamic settings. In particular, in Nekipelov (2006) I use a similar approach to tackle the dynamic principal-agent model of Holmstrom and Milgrom (1987). In both settings, the estimation method is based on a partial equilibrium approach. In the case of auctions, if bidders respond optimally to the movement of price, their bids should replicate the price process over time. Specifically, if we observe the distribution of price in an auction at any given time, and we can compute the corresponding distribution of optimal bids, at the true structural parameters these distributions should coincide. This idea is empirically implemented by minimizing the distance between these two distributions. I prove this technique produces consistent estimates of the structural parameter of the model and establish non-parametric identification. I then estimate the model using data from pop-music CD auctions on eBay. The estimation results are consistent with economic intuition. In particular, I observe that the entry rate is a decreasing function of price, while the average entry rate is consistent with the average number of bidders per auction. In addition, the estimates show that the price movement becomes more incremental towards the end of the auction, which also coincides with the empirical observations.

Having recovered the structural parameters, I am then free to perform counterfactual analyses, I focus my attention on a controlled experiment.

There have, of course, been other attempts at explaining the observed features of eBay bidding behavior. For instance, irrationality of bidders, tacit collusion, technical problems of processing bids in the last moment, and common values[3] have been suggested as alternative explanations. My model is unique in emphasizing the importance of endogenous entry in the bidding decision. Specifically, in my framework, any event that shifts the rate of entry into a group of auctions should also change the optimal bidding strategies. In particular, the incentive to deter entry should depend on the thickness of the market (other things being equal). In thick markets (holding the number of bidders constant) entry should be significantly more sensitive to price, making entry deterrence (and the early bidding associated with it) more common in larger markets. The alternative models listed above do not have such implications. Therefore, I can directly test the validity of my hypothesis against these alternative models via a controlled experiment. In particular, I artificially increase the size of the market for a specific CD by listing additional auctions. My empirical results reveal that an expansion in the size of the market indeed increases the frequency of early bidding, confirming the predictions of my model[4].

The structure of the paper is as follows. Section 2 describes the data used in the structural estimation, while Section 3 describes the model and derives the optimal bidding strategy. Section 4 outlines the identification strategy. Section 5 explains the estimation procedure and contains the main structural estimates of the model obtained using the data from eBay. Section 6 presents the results of the field experiment and a broader discussion. Section 7 concludes.

2 Bidding for pop-music CDs on eBay

In my theoretical framework, the competing strategic forces that drive bidding behavior depend on the process by which individual bidders learn about the attractiveness of *competing auctions,* as opposed to the value of a particular object in a *single auction.* For this reason, I will focus on private value auctions, requiring me to identify a class of eBay auctions for which this assumption is

[3]Existing attempts to explain the eBay data include the tacit collusion proposed in Ockenfels and Roth (2002) technical features of eBay (lack of information about the proxy bidding mechanism, offered in Ockenfels and Roth (2002), multiplicity of listings in Jank and Shmueli (2002) and Peters and Severinov (2004)) and behavioral reasons, such as bidder's irrationality or uncertainty about own private value as in Rasmusen (2003) and Hossain (2004), and the possibility of common values in Bajari and Hortacsu (2003).

[4]There have other field experiments recently conducted on eBay. Yin (2003), for instance, studies bidding for computers on eBay. Many studies of eBay discuss bidding for collectible items which also possess common value components. Relevant papers include Bajari and Hortacsu (2003) and Ockenfels and Roth (2002) and some other papers. To my knowledge, my experiment is the first one to focus on private-value auctions.

appropriate.[5] The auction category I have chosen is music CDs from the recording artist Madonna. The logic behind this choice is threefold. First, a CD is a relatively inexpensive item, so its purchase should have little or no wealth effects. Second, variation in the market price of CDs sold outside of eBay is likely to be small, implying that most bidders are consuming CDs for their intrinsic value, rather than buying them for speculative reasons. Finally, focusing on a single recording artist should further limit heterogeneity in the bidder population, further mitigating the influence of other forms of information asymmetry.

Data from eBay have been used by economists in several empirical studies, mainly in the context of common values[6]. A unique feature of my dataset is that it was constructed to analyze a continuous-time auction and to facilitate a related field experiment. The dataset contains information about both the auctioned objects (including information about the seller) as well as the bidding process itself. The data was collected using a *Perl* "spider" program, which gathered information about both items and bidding histories[7]. The information on bidding behavior characterizes the bid profile for each item, including the bidding histories (including an item identifier) and bid amounts (along with a bidder identifier). The sample statistics (timing of bids and their quantities) for this dataset are contained in Table 1[8]. Note that the duration of all auctions is normalized to 1, both for illustrative purposes and to facilitate subsequent estimation. (The reason for this normalization will be explained in a subsequent section). Along with these numerical variables, the dataset also contains categorical variables identifying the objects and the bidders who have submitted each bid. Table 1 reveals that the average bid is close to $14 and the average bidding time is very close to the end of the auction.

Along with the bidding profile, I collected additional information about each auction in the dataset. The variables, summarized in Table 1, include the duration of the auction, the buy-it-now price (if offered), the seller's proportion of positive feedback and score, a dummy variable indicating

[5]A private value structure requires that valuations of different bidders be uncorrelated. In practice, I will assume that a private value setting is appropriate so long as these correlations are sufficiently small compared to the variance of valuations.

[6]Yin (2003) examines the extent of the winners' curse in online auctions for computers, a category in which product authenticity is a major issue. Several other authors have analyzed bidding for collectibles, which also feature common value components. Relevant papers include Bajari and Hortacsu (2004) and Ockenfels and Roth (2002) and some other papers.

[7]The program was designed to work as follows. First it submitted a search query for music CDs sold through auctions (no eBay stores), and sorted the items according to the auction ending time. The program selects up to 20 pages of items so that the auctions for those items end in at most 4 hours. The program browses individual pages for those items and saves the exact time when the auction ends, the rating and the nickname of the seller, the characteristics of the auction for a CD. Then the program sleeps for 4 hours and goes to the page where eBay saves the bidding history. The process repeats the necessary number of times.

[8]The tables in this and further sections can be found in Appendix F

that the item's page includes a photo of the CD cover, the shipping cost, and the condition of the CD (used or new). Note that most sellers in the dataset have very high, positive feedback levels and scores. The average duration of the auction is 6.5 days, reflecting the fact that sellers can avoid additional restrictions by choosing an auction lasting less than 7 days. The buy-it-now price is very close to the average bid and the buy-it-now option is offered approximately half of the time. The condition variable reveals that half of the CDs on the market are used, while the picture dummy indicates that most auctions include a photo of the CD. Approximately 83% of the sellers in my dataset were located in the United States, while other sellers were primarily located in Argentina, Australia, Europe and Japan. I should also note that the four best-selling Madonna albums (according to Amazon.com) accounted for 16% of the eBay market. A "top 4" indicator variable will therefore be used to control for additional heterogeneity across auctions.

The dataset allows me to highlight a few basic features of bidding behavior, providing some initial support for my model. To describe these features, I constructed two additional variables: a variable corresponding to the arrival rate of bids (i.e. the number of bids in the auction per unit of time) and another variable equal to the size of the price jump (in the beginning of the auction, it is simply the first bid). Table 2 contains the results of some basic regression analysis aimed at illustrating how these new variables (as well as the number of bids in the auction) are affected by the characteristics of the auctioned objects, as well as time and price. This basic analysis is carried out using simple linear regressions of the characteristics of bidding on the characteristics of the objects. The first column contains the output from a fixed effect regression that explains the observed size of bid increments. The estimates reveal that bids become more incremental (i.e. smaller) toward the end of the auction. Moreover, the price tends to jump higher when the price level is high.

Column 2 focuses on the number of bidders who are active over the course of the auction. The estimates from this regression reveal that the entry of bidders (per unit of time) is lower in auctions with higher starting prices: price is a significant determinant of entry. This implies that we simply cannot ignore the dependence of bidding behavior on price in eBay auctions. Employing a static private value auction model will lead to incorrect conclusions

Apart from providing further motivation for my model of a continuous-time auction, this initial reduced-form regression analysis provides information important in constructing a sensible structural model. In particular, we can establish the relevance of various auction characteristics (potential instruments) by looking at the determinants of cross-auction variation in the number of bids and the average price jumps due to bidding. Column 3 examines the intensity of bidding in the auctions auctions. Here, we can see that the number of bids is lower if starting price is high and higher for the most recent albums, while a higher feedback score (of the seller) increases the number of bids. The results for number of bidders (Column 4) are quite similar to the results for the number of bids (Column 3). Therefore, we can conclude that the observed characteristics of cross-auction variation are significant and have sensible coefficients in these simple regressions. For this reason, I will use them as shifters in my structural model.

2.1 Continuous-Time Model

In this section I analyze a continuous-time model using the notion of the *public perfect equilibrium* extensively studied in Sannikov (2007). The structure of the repeated game with imperfect monitoring in continuous time considered in Sannikov (2007) suggests that players in the game have discrete action spaces, but they do not observe each other's actions directly. Instead, they observe the realizations of a diffusion process with a drift, where the drift is determined by actions of the players. The strategy in this case would be defined as a stochastic process measurable with respect to the filtration generated by the publicly observable process. This framework can be a good approximation of the environment of the online auctions. In the basic model for the auction on the Internet I choose to treat the price process as public information. The actions of individual bidders represented by their bids are assumed to be private. The difference between this environment and the model in Sannikov (2007) is that actions in the online auction are continuous and the number of rivals is unknown. I address this complication by introducing the beliefs of the bidders regarding the price movement. These beliefs are determined by the parametric family of probability distributions of the price, where the parameter is a private information of the player. The strategy of the bidder is determined as a stochastic process measurable with respect to the filtration generated by the beliefs rather than the publicly observable price process. In the subsequent discussion I formalize this structure.

The continuous-time auction model considered here describes a price process that starts at time 0 and ends at a time T. Consider the strategic behavior of a single bidder i who competes against her rivals by submitting multiple bids $b_i \in \mathbb{R}_+$ at any moment of time $t \in [0, T]$.

The auction has a second-price structure, so that the price of the object - $p \in \mathbb{R}_+$ - at a given time t is equal to the second highest bid among the bids submitted before time t. A bidder observes neither the identities nor the number of her rivals, and I denote the number of bidders at time t by N_t. She knows only that entry into the auction is characterized by a Poisson process with frequency $\lambda(t, p, \theta_0)$, a function of time, price and the visibility parameter θ_0[9], while its functional form is a common knowledge. This function is the instantaneous demand described above, whose dependence on price is motivated by equilibrium search considerations[10]

[9] Baldi, Frasconi, and Smyth (2003) observed that the number of hits at popular Internet sites can be described by a Poisson process

[10] Existing models of simultaneous search lead to the conclusion that the optimal stopping rule for the search

The visibility parameter θ_0 is exogenous and fixed for a specific auction. The set of possible visibility values Θ is convex, compact, and known by the bidders. Although θ_0 is not observed, bidders have initial beliefs regarding its value. These beliefs take the form of normal priors, truncated to the set Θ, with means μ_θ and variances σ_θ^2 drawn from the distributions $G_\mu(\cdot)$ and $G_\sigma(\cdot)$ respectively. Bidders with small mean-squared errors of the initial belief $(\mu_\theta - \theta_0)^2 + \sigma_\theta^2$ are considered "more informed", while bidders with large mean-squared errors are "less informed".

In addition to being uncertain about the visibility of the auction the bidders also have only imperfect observations of the price process. Due to the inability to continuously monitor the price in the auction (late at night, for example), a price change occurring at time t will be observed by the bidder at time $t + \epsilon$ where ϵ is independent and drawn from the same (non-negative) distribution across the bidders. Bidders are assumed to observe their own ϵ and it is constant throughout the auction for the given bidder. The instantaneous demand for an individual bidder with an observation delay ϵ will be denoted $\lambda_\epsilon(t,\, p,\, \theta)$. For example, one of possible forms for the transformation of $\lambda(\cdot)$ into $\lambda_\epsilon(\cdot)$ is $\lambda_\epsilon(t, p, \theta) = \alpha e^{a\epsilon} \lambda(t, p, \theta)$ for fixed a and α.

The bidder's valuation for the object is v and she is risk-neutral: hence if she obtains the object for price p, then her utility is $v - p$. The valuations of the bidders are independently drawn from the distribution $F(v)$. The bidder maximizes the expected utility[11] from winning the auction $E_0 \{(v - p_T)\, \mathbf{1}\, [b_T > p_T]\}$. This expected utility is positive only if the bid of the bidder under consideration at the end of the auction is the highest (and, thus, is higher than the price of the object equal to the second highest bid). The strategy of the representative bidder is characterized by a bidding function $b_{v\,\epsilon}(t,\, p,\, \mu,\, \sigma)$, which gives the optimal bid value at time t for a bidder with valuation v and observation delay ϵ if the price of the object is p and the beliefs of the bidder about the visibility have mean μ and variance σ. To facilitate derivation, I suggest the following decomposition of the bidding function:

$$b_{v\epsilon}(t,\, p,\, \mu,\, \sigma) = p + \eta(t,\, p,\, \mu,\, \sigma)\,. \tag{1}$$

The function $\eta(\cdot)$ will be referred to as the bid increment. In general $\eta(\cdot)$ depends on ϵ and v, but to facilitate the notation we will dropping these indices wherever it clear which bidder is described. As noted above, the bidder wins the auction if her bid increment is positive at the end of the auction.

Definition 1 *A pure strategy of the bidder i η_i is a stochastic process with sample paths in $\mathbb{R}_+$ which*

process will be price dependent. Such results are known for labor markets and matching markets more generally, e.g. Pissarides (1990), McAfee (1993), Peters (1991), Moen (1997).

[11] It can be argued that in some cases individuals might bid on eBay for reasons other than maximizing the expected surplus from winning when, for instance, they like gambling. A variety of possible other targets are discussed in Kagel (1995). In this paper, I focus only on maximization of expected surplus.

is progressively measurable with respect to the filter generated by the price process corresponding to a particular visibility θ_0.

The price is always equal to the second highest bid. From the perspective of a particular bidder who submits at least a minimum required bid higher than the current price, the price is equal to the minimum of her own bid and the highest bid of her rivals. In this way, if she submits a bid and is the current highest bidder, then the current price is determined by the highest bid of the remaining bidders. A new bidder determines the price if her bid is between the current highest bid and the second highest bid. Any particular bid of a bidder can determine the price only once per auction.

Bidders form beliefs regarding the distributions of the price process at each instant. If bidder i has full information regarding parameter θ_0 in the entry process, the beliefs of bidder i about the price movement can be characterized by the finite-dimensional distributions of the stochastic process:

$$dp_{t,\epsilon_i} = h\left(t,\, p_t,\, \eta_i\right)\, dJ_{\epsilon_i}\left(t,\, p_t,\, \theta_0\right), \tag{2}$$

where dp_t are changes in the price over infinitesimal time intervals dt, and $dJ_\epsilon\left(t,\, p_t,\, \theta_0\right)$ are increments of the Poisson process with frequency $\lambda_\epsilon\left(t,\, p_t,\, \theta_0\right)$.[12] This equation reveals that the price, as observed by an individual bidder, evolves in jumps. The size of the jump at time t is equal to $h\left(t,\, p_t,\, \eta_i\right)$ where the bid increment η is defined at (1) while the timing of the jumps is governed by the Poisson process $J_\epsilon(\cdot)$, so that $n_t^J = \int_0^t dJ_\epsilon\left(\tau,\, p_\tau,\, \theta_0\right)$ is the number of price jumps from the beginning of the auction up to time t[13]. Thus if n_t^J is the number of price jumps up to t and t_i are the times when jumps occur for $i = 1, \ldots, n_t^J$, then the price can equivalently be written as:

$$p_{t,\epsilon_i} = \sum_{i=1}^{n_t^J} h\left(t_i,\, p_{t_i},\, \eta_i\right). \tag{3}$$

This expression shows that the price of the object at time t is equal to the total sum of the price jumps up to time t.

[12] A more formal treatment of the conditions providing the existence of the function in the described model is given in technical companion for this paper. I will be using footnotes in this section to give short comments regarding the formal properties of the functions in the model. To describe the model formally I consider a process defined on a complete probability space $(\Omega,\, \mathbb{F},\, \mathbf{Q})$ and $\{\mathfrak{F}_t,\, t \in [0,\, T]\}$, a filter such that $\mathfrak{F}_t \subset \mathbb{F}$ and $J(t,\, x)$, and a Poisson measure for $t \in [0,\, T]$ adapted to filter $\{\mathfrak{F}_t\}$.

[13] Formally we need to make sure that the solution to this stochastic differential equation exists for given $h(\cdot)$ in terms of a stochastic Itô integral. For that purpose we impose two requirements on $h(\cdot)$: (i) $\int_0^T \left|h(t,\, x_t, \eta)\right|^2 \lambda_\epsilon(t,\, x_t)\, dt < +\infty$ with probability 1; (ii) The compensated process $\int_0^t h(t,\, x_t, \eta)\, \overline{J_\epsilon(dt,\, dx)}$ is a local square integrable martingale adapted to the filter $\{\mathfrak{F}_t\}$ with piecewise - continuous sample functions.

Second, the timing of the price jumps is determined by the rest of the bidders, while the observed timing is contaminated by observation delay error. In the absence of observation delays, bidders would coordinate their bidding such that all participants would bid at the same time when new information arrived.

Equation (2) contains the visibility parameter θ_0, which is unobserved by the bidders. Because a bidder forms her bidding strategy to optimally influence the auction price, it is important that she predict the visibility parameter precisely. This prediction incorporates prior information, which the bidders can have from previous bidding experience and from observing the behavior of the price in the auction. In this way, as the price evolves during the auction, the precision of the bidder's estimate of the visibility of the auction increases. This process can be the result of strategic learning. One way to describe the estimation of the visibility parameter by the bidder is to consider Bayesian updating of the prior beliefs given the stochastic movement of the price.[14]

To compute the dynamics of these beliefs I use the linear filtration method developed in Van Schuppen (1977). While an optimal linear filter (as shown in Lipster and Shiryaev (2001)) is infinite-dimensional, for convenience of computing the bidder's estimate of the visibility of the auction, I restrict the analysis to the first two moments[15]. Nekipelov (2007) shows that the posterior distribution for θ_0 can be described by a normal posterior distribution with mean μ_t, and variance parameter σ_t. The mean and variance parameter of the bidder's belief regarding the auction's visibility change over time according to the following system of stochastic differential equations:

$$\begin{aligned} d\mu_t &= \frac{\partial\lambda(t, p_t, \mu_t)}{\partial\theta} \frac{\sigma_t}{\psi(t, p_t, \mu_t)} \left\{dp_t - \psi(t,\, p_t,\, \mu_t)\, dt\right\}, \\ d\sigma_t &= -\frac{1}{\lambda(t, p_t, \mu_t)} \left(\frac{\partial\lambda(t, p_t, \mu_t)}{\partial\theta}\right)^2 \sigma_t^2\, dt. \end{aligned} \tag{4}$$

In this expression $\psi\left(t,\, p_t,\, \mu_t\right) = h\left(t,\, p_t,\, \eta_t\left(t, p_t, \mu_t, \sigma_t\right)\right) \lambda_\epsilon\left(t,\, p_t,\, \mu_t\right)$ characterizes the expected growth rate of price per unit of time.

[14]The specific problems of construction of Bayesian estimates for Poisson - type processes are described, for example in Kutoyants (1998) and Karr (1986). For the Gaussian processes the solution to nonlinear mean-square inference problem may be found, for example in Lipster and Shiryaev (2001) in the form of the system of stochastic differential equations. In a more general case of point-stochastic processes, a solution becomes complicated, relies heavily on martingale properties of the underlying stochastic processes and, more importantly, becomes a computationally intensive problem. A constructive way of building linear Bayesian forecasts for the parameters of a Poisson process with a variable frequency is shown in Grandell (1972). The generalization of Grandell's method follows from the integral representation of martingales and is discussed in detail, for example, in Lipster and Shiryaev (2001). Here I will discuss an easier and more intuitive approach which, despite being less rigorous, provides a straightforward way to construct linear Bayesian estimates for the unknown parameters.

[15]In other words if $\mathfrak{F}_{xt}$ are the σ-algebras generated by the sample trajectories of the price process, then my assumption is that $E\{(\theta_t - \theta^*)^3 | \mathfrak{F}_{xt}\} = \sigma_3$ for any $t \in [0,\, T]$.

2.2 Equilibrium

In reality, auction markets are complex and involve repeated strategic interactions between buyers, sellers and the auction company. The equilibrium that I examine in this paper considers the behavior of bidders in a single auction and takes as given the behavior of the sellers and the auction engine itself.

In my model, the bidder forms beliefs regarding the distribution of the price at each instant. In equilibrium, these beliefs should be consistent with the true price distribution generated by aggregate behavior of multiple bidders, while bidders have to maximize their expected payoffs from winning the auction. To formalize the notion of the equilibrium, I use the results in Sannikov (2007) and combine them with the Bayesian structure of the game that I consider in this paper. Following Sannikov (2007) I call the equilibrium here a *perfect public equilibrium.* The components of equilibrium are: (i) entry of new bidders, (i) types of new bidders and their beliefs regarding the price process, (iii) a profile of strategies of bidders who have entered into the auction up to time t, (iv) the price process corresponding to the second-highest bid at each instant.

Definition 2 *A profile of strategies* $\left\{(\eta_{i,\tau})_{i=1}^{N_\tau}\right\}_{\tau=0}^{T}$ *along with beliefs of bidders regarding the price process* $\mathbf{P}\left\{p_t < x,\, J\left(t, p_t, \mu_t^i\right) < y\right\}$ *and means and variances of the beliefs of the bidders regarding the visitbility* $\left\{\mu_t^i,\, \sigma_t^i\right\}_{i=1}^{N_t}$ *constitute a perfect public equilibrium if for each* $t \in [0,\, T]$*:*

- *The entry of the bidders is a Poisson process with the frequency* $\lambda(t,\, p,\, \theta_0)$ *for a given* $\theta_0 \in \Theta$ *so that the number of bidders who entered up to time* t *is* N_t *and is unobserved by the participating bidders*
- *Observable public histories consist of price paths up to time* t
- *Bidders' types include valuations, initial beliefs regarding the visitbility, and observation delays. The valuations of the bidders are drawn from the distribution* $v^i \sim F(\cdot)$*, the initial parameters of the beliefs are* $\mu_0^i \sim G_\mu(\cdot)$ *and* $\sigma_0^i \sim G_\sigma(\cdot)$*, and the observation delay errors are uniformly distributed* $\epsilon^i \sim U[0,\, 1]$ *for all bidders* $i = 1, \ldots, N_t$ *and for each* $t \in [0,\, T]$[16]
- *Each bidder* i *maximizes the expected payoff from winning the auction given her beliefs regarding the finite-dimensional distributions of the price process*

$$\mathbf{P}\left\{\int h\left(t,\, p_t,\, \eta_i\right)\, dJ_{\epsilon_i}\left(t,\, p_t,\, \mu_t^i\right) < x,\, J\left(t, p_t, \mu_t^i\right) < y\right\}$$

  and her type

[16] As I will show in the next section, the distribution of the observation error and the frequency of entry are not separately identified. For this reason, I assume here that the distribution of observation errors is uniform to estimate the frequency of entry.

- *A profile of actions* $\left\{ (\eta_{i,\tau})_{i=1}^{N_\tau} \right\}_{\tau=0}^{T}$ *along with the Poisson entry process* N_t *generate the price process such that its finite-dimensional distributions are consistent with beliefs of bidders regarding the price process:*

$$\mathbf{P}\left\{ \int_0^t \max_{i=1,\dots,N_t,\, i\neq J} \eta_{i,t} < x, J\left(t, p_t, \theta_0\right) < y \right\}$$
$$= \mathbf{P}\left\{ \int h\left(t,\, p_t,\, \eta_i\right) \, dJ_{\epsilon_i}\left(t,\, p_t,\, \mu_t^i\right) < x,\; J\left(t, p_t, \mu_t^i\right) < y \right\} \Big|_{\substack{\mu_{i,t}=\theta_0 \\ \epsilon_i = 0}}. \tag{5}$$

for $J = argmax_j \eta_{j,t}$.

It is important to note that in my model bidders can submit multiple bids, and equation (5) refers to the last bid of every bidder who entered into the auction up to time t. This equilibrium concept suggests that the price increase at time t is driven by the second order statistic of the bid increment function computed at price p_t. In this way, a competitive dynamic equilibrium is a solution of the stochastic differential equation (5).

Let me first consider the steps that establish the existence of equilibrium in the continuous-time auction model. It is straightforward to establish existence and uniqueness of the fixed point by considering the individual optimization problem of the bidder. An equilibrium in this setting corresponds to the solution of a collection of individual bidding problems of the participating bidders. This collection is a system of partial differential equations describing the law of motion of the value functions of the bidders in the auction.

To prove uniqueness of the equilibrium, we must account for entry of bidders into the auction. By assumption, this entry is driven by a Poisson process with variable frequency. If this frequency is finite for any value of the price, time, and visibility parameter, then the number of bidders entering into the auction per unit of time is finite with probability one. This means that the equilibrium in the continuous-time auction can be represented as a finite collection of optimal bidding problems of individual bidders with probability one. As a result, this collection has a unique fixed point that is the equilibrium in the continuous-time auction. A more detailed treatment of this result is presented in the technical companion to this paper. I will use this uniqueness result to describe the likelihood of the continuous-time auction model.

The following theorem formalizes the above intuition.

Theorem 1 *Suppose that the stochastic differential equations, representing the dynamics of state variables (the price and beliefs) have continuously differentiable coefficients and square - integrable solutions with probability one. In this case the equilibrium represented by Definition 3 exists and is unique.*

Proof:

Existence. In Appendix A we establish that the solution to the individual bidding problem exists. The individual bidding problem can be reduced to the partial differential equation with boundary condition. In this way, the equilibrium can be reduced to the system of partial differential equations. As the proof in Appendix A still holds for vector-valued equations, it leads to the existence of the solution to such system of equations. Thus, it proves the existence of the equilibrium.

Uniqueness. We will now prove that the equilibrium is unique for the entire problem using the fact that the number of strategic bidders follows a Poisson process, so it is quite unlikely to have a situation when the number of bidders is extremely large. We assume that the set of potential strategic bidders is countable. We can write down the problem (7) for each strategic bidder in the set of bidders $\mathbb{I}$.

Let us take finite subsamples of players from $\mathbb{I}$. For each fixed N - the size of taken subsamples we will obtain a multidimensional boundary problem which has a unique solution for the same reason as a one-dimensional problem. It is known that the set of all finite subsets of a countable set is a countable set[17]. Therefore we can index each subsample of size N by some natural number k and attribute to it an element of a positive sequence $\{x_n\}_{n=1}^{\infty}$ such that $\sum\limits_{n=1}^{\infty} x_n = 1$. Further note that as players are coming at a Poisson rate we can find some positive constants c_1 and c_2 that the probability of encountering the subsample sample of size N is at least $c_1 e^{-c_2 N}$. Then the probability of observing a subsample of size N indexed k is at least $x_k\, c_1 e^{-c_2 N}$. Total probability of observing a subsample of the size at most N is at least

$$\sum_{n=0}^{N} \sum_{k=1}^{\infty} x_k\, c_1 e^{-c_2 n} = \sum_{n=0}^{N} c_1 e^{-c_2 n} = 1 - e^{-(N+1)c_2}.$$

While drawing subsamples of size N from $\mathbb{I}$ we have the following properties:

(i) For each of the subsamples the equilibrium exists and is unique.

(ii) The total probability of the observing a subsample of a size at most N is approaching 1 as N goes to infinity.

If we now make N go to infinity, these two conclusions imply that with probability 1 the equilibrium exists and is unique.

By forming this construction not only we provided a method for calculating the equilibrium but with the aid of operator theory immediately obtained the result about the existence and uniqueness of the equilibrium.

This proof shows also that it is quite complicated to prove the uniqueness in case when parameter set is infinite-dimensional (that is the relevant functions are defined nonparametrically). In fact

[17] See Hrbacek and Jech (1999)

to prove the uniqueness in that case we would have to add some mechanism that exponentially increases explanatory power of the non-parametric estimate as the size of the grid increases.
Q.E.D.

2.3 Individual Bidding Behavior with Consistent Beliefs

In the previous subsection, I characterized the public perfect equilibrium with consistent beliefs which generates the dynamic behavior in the considered environment. I assumed that bidders maximize their expected utilities from winning the auction by submitting optimal bids. In equilibrium beliefs of individual bidders regarding the price distribution should be consistent with the true distribution of the price. In the following two subsections I will describe individual best responses. In this subsection I will provide a method for finding individual best responses of the bidders provided that they have consistent beliefs regarding the price distribution. In case of consistent beliefs, bidding strategies are described by the bid increment functions $\eta(t, p, \mu, \sigma) \in \mathbb{R}_+$ defined as the difference between the submitted bid of the bidder and the current price. Expected utility maximization is constrained by the dynamics of the price, described by the stochastic differential equation (2). The visibility of the auction θ_0 is unobserved by the bidders, but they try to infer its value from observing the actual movement of price in the auction. I assume that the bidders infer the visibility of the auction by forming priors and update their beliefs as the price in the auction changes. This allows me to write the problem of the bidder as:

$$
\begin{aligned}
&\max_{\eta(\cdot)} E_0\{(v - p_T)\,\mathbf{1}\,[\eta_T > 0]\} \\
&dp_{t,\epsilon} = h\,(t,\, p_t,\, \eta)\; dJ_\epsilon(t,\, p_t,\, \mu_t), \\
&d\mu_t = \tfrac{\partial\lambda(t,p_t,\mu_t)}{\partial\theta}\tfrac{\sigma_t}{\psi(t,p_t,\mu_t)}\left\{dp_t - \psi(t,\, p_t,\, \mu_t)\,dt\right\}, \\
&d\sigma_t = -\tfrac{1}{\lambda(t,p_t,\mu_t)}\left(\tfrac{\partial\lambda(t,p_t,\mu_t)}{\partial\theta}\right)^2 \sigma_t^2\, dt. \\
&\theta_t|_{t=0} \sim N\,(\mu_0,\, \sigma_0)\,, \quad p_t|_{t=0} = p_0.
\end{aligned}
\tag{6}
$$

The first equation of this dynamic optimization problem represents the objective of the bidder, which is the expected utility of winning the auction. This objective reflects the fact that the bidder obtains a positive utility from the auction only if she wins it (so that her bid is the highest, implying that the bid increment at the end of the auction is strictly positive $\eta_T > 0$). The expectation $E_t[\cdot]$ is taken over the information set up to time t which includes the paths $\{p_t,\, \mu_t,\, \sigma_t\}_{t=0}^T$.

The second equation demonstrates the price process corresponding to beliefs of an individual bidder regarding the price movement. It describes the dynamics of the price as a jump process with Poisson-driven jumps. The frequency of price jumps is $\lambda_\epsilon(t,\, p_t,\, \mu_t)$. The bidders observe the price jumps with a random delay ϵ. The mean of the belief of the bidder about the visibility of the auction is μ_t.

The third and the fourth equations represent the evolution of the mean and variance of the bidder's beliefs regarding visibility. These beliefs are driven by the price changes so that the mean of the distribution shifts when the auction price jumps. Moreover, if the expected price growth rate is elastic with respect to visibility, then the variance of the bidder's beliefs will decrease over time.

The last line in the dynamic optimization problem (6) sets the initial conditions. Bidders' initial beliefs about the visibility parameter , θ_0, take the form of a normal distribution with mean μ_0 and variance σ_0. The price at the beginning of the auction is the reserve price for the object set by the seller. It is assumed to be observable by the bidders and equal to p_0.

2.4 Individual Best Responses with Consistent Beliefs

My approach to solving the problem (6) uses the Bellman equation formulation, which has been studied extensively in the literature on stochastic dynamic optimal control. Here, we can define the value function of the bidder as:

$$V\left(t,p,\mu,\sigma\right)=E_t\left\{\left(v-p_T\right)\mathbf{1}\left[\eta_T>0\right]\right\}.$$

The function $V(\cdot)$ specifies the expected surplus of the bidder from winning the auction given that, at time $t\in[0,\,T]$, the price in the auction is p and the belief about the visibility of the auction has mean μ and variance σ. The value function at the end of the auction is the utility from winning the auction (to win the auction the bidder must have the highest bid, but the price is determined by the second highest bid, so $\eta_T=b-p_T>0$). The vector of state variables $(p_t,\,\mu_t,\,\sigma_t)'$ forms a sufficient statistic for the dynamics information embedded in the entire price path up to time t. The dynamic evolution for the state variables $(p_t,\,\mu_t,\,\sigma_t)'$ over time is a Markov process.

The optimal behavior of the representative bidder is derived by applying the Itô calculus to the value function. Intuitively, the derivation uses the fact (due to the Bellman principle) that the expected surplus of the bidder at time t is equal to the maximum over all possible bid increments of the expected surplus at time $t+dt$. The expectation is taken over the distribution of all possible price changes in the interval of time dt. The bidders optimally influence the price process to maximize their expected payoff. As a result, we can represent the law of motion of the value function of the bidder as:

$$\begin{aligned}&\tfrac{\partial V(t,p,\mu,\sigma)}{\partial t}+\sup_{\eta\in\Xi}\Bigg[-\tfrac{\sigma}{h}\tfrac{\partial\lambda}{\partial\theta}\tfrac{\partial V}{\partial\mu}-\tfrac{\sigma^2}{\lambda}\left(\tfrac{\partial\lambda}{\partial\theta}\right)^2\tfrac{\partial V}{\partial\sigma}+\\&+V\left(t,\,p+h\left(t,p,\eta\right),\,\mu+\tfrac{\partial\lambda}{\partial\theta}\tfrac{\sigma}{\lambda},\,\sigma\right)\lambda_\epsilon\left(p+h\left(t,p,\eta\right),\,\mu+\tfrac{\partial\lambda}{\partial\theta}\tfrac{\sigma}{\lambda},\,t\right)-\\&-V(t,\,p,\,\mu,\,\sigma)\,\lambda_\epsilon\left(p,\,\mu,\,t\right)\Bigg]=0,\end{aligned}\tag{7}$$

$$V(T,\,p,\,\mu,\,\sigma)=\sup_{\eta\in\Xi}\left\{\left(v-p\right)\mathbf{1}\left[\eta>0\right]\right\}.\tag{8}$$

In this equation, the space of the control functions Ξ limits my analysis to the bounded and piecewise-continuous bid increment functions, simplifying further derivations[18]. Equation (7) is a partial differential equation for the value function of the bidder $V(\cdot)$ with boundary condition (8) which implies that the value function of the bidder at the end of the auction has to be equal to her utility from winning the auctioned object[19]. Note that the boundary condition in (7) implies that the optimal bid increment at the end of the auction can take a range of values $\eta_{v\epsilon}(T, p, \mu, \sigma) \in (0, +\infty)$. If we analyze the optimal strategy in the interval $[T - \tau, T]$ then the first equation in the boundary problem (7) suggests that as $\tau \to 0$ then $v - p$ is optimal for $\eta_{v\epsilon}(T, p, \mu, \sigma)$ if $v > p$. Therefore, it is optimal for the bidder to submit a bid equal to her true valuation at the last moment of the auction, analogous to the behavior in a static second-price auction.

2.5 Discussion of the Model

A collection of individual bidding problems generates a dynamic *public perfect equilibrium*. This individual bidding problem is represented by the boundary problem (7) and (8). I will first describe the general properties of this problem and then discuss the interpretation of the individual components of the decision problem.

The problem (7), (8) has a two-component structure. The component in the *supremum* describes optimal bidding at time t. This component allows me to compute the bidding strategy for a fixed instant as it produces the optimal response of the bidder at time t to the current price, and the mean and variance of the bidder's beliefs regarding auction visibility. After the optimal bidding strategy is computed at time t, I can compute the time derivative of the value function of the bidder and, as a result, obtain the value function for the previous instant of time[20]. Having discussed the general

[18] In fact I require that $\eta \in \Xi$ are continuous and differentiable almost everywhere functions defined on $[0, T] \times [0, \overline{X}] \times \Theta \times \Sigma$. I assume that these functions are bounded by some $\overline{X} < \infty$ and that $h(\cdot, \eta(\cdot))$ are measurable functions of finite variance with respect to the Poisson measure.

[19] Technically speaking, I have obtained a boundary problem for the expected surplus only. This determines the behavior of the function on one of the boundaries but does not provide the information about the derivatives. Such a boundary problem is often classified as a Dirichlet problem.

[20] This two component structure of the problem naturally enables me to use the standard Euler algorithm to compute the optimal behavior of the bidder. The algorithm would begin from the boundary condition which states that at the end of the auction (T) the value function of the bidder is equal to her utility from winning the auction. Then I make a step backwards in time τ and recompute the value function at time $T - \tau$. The step size τ determines the precision of the computed value function. For this value function I compute the optimal bid increment of the bidder and evaluate the time derivative of her value function. This process repeats until it reaches the beginning time 0. This approach has a significant computational advantage over the value function iteration approach which is used to compute the optimal behavior in an infinite horizon discrete time optimization problem. As compared to the value function iterations, there is no need to prove the convergence of the iteration method and wait until the convergence is reached. In my case the number of steps is fixed and is equal to the number of grid points of time

structure of the bidder's problem, let us now proceed with the interpretation of the individual components describing bidding behavior at a specific point of time in the auction.

The last two terms in equation (7) describe the main strategic tradeoffs faced by each bidder. These strategic tradeoffs highlight the role of visibility in the bidder's behavior: higher visibility will in general imply that the bidders will bid more aggressively. The information about the visibility influences the bidding frequency and allows me to distinguish the more informed bidders from the less informed ones. Let me consider these two terms by components. The component

$$V\left(t,p+h\left(t,p,\eta\right),\mu,\sigma\right)\lambda_{\epsilon}\left(t,p+h\left(t,p,\eta\right),\mu\right)-V(t,p,\mu,\sigma)\lambda_{\epsilon}\left(t,p,\mu\right)$$

describes the expected change in the value function due to the price jump. The tradeoff of the bidder is between bidding early and bidding late. Bidding early deters entry, decreasing the probability of price jumps and resulting in a higher value function at the end of the auction. On the other hand, the direct effect of bidding early is to raise the price throughout the auction, which decreases the value function of the bidder at the end of the auction. The optimal time to bid occurs when these two effects offset each other. The movement of the value function is determined by the product of the current value function and the rate of entry into the auction $\lambda_{\epsilon}(\cdot)$ (instantaneous demand). As a result, the bidder maximizes the total value of the entry rate (value times demand) and extracts a surplus from bidding by choosing the optimal point on the demand schedule by means of bidding early. This behavior constitutes entry deterrence.

The component

$$V\left(t,p,\mu+\frac{\partial\lambda}{\partial\theta}\frac{\sigma}{\psi},\sigma\right)\lambda_{\epsilon}\left(t,p,\mu+\frac{\partial\lambda}{\partial\theta}\frac{\sigma}{\psi}\right)-V\left(t,p,\mu,\sigma\right)\lambda_{\epsilon}\left(t,p,\mu\right)$$

reveals the role of information in the strategic decision of the bidders. Note that the size of this component depends on the variance of a bidder's belief regarding the visibility of the auction, so that the effect will be different for bidders with different variances of prior beliefs. If the variance of the prior is large, the bidder does not have much information about the visibility of a specific auction. As a result, the value function of the bidder will significantly depend on the behavior of price. As the variance of the bidder's belief diminishes (towards the end of the auction), the value function will become less sensitive to the new information and the strategy of this bidder will converge to the strategy of the fully informed bidder. Therefore bidders with large prior variances will prefer rapid changes in price because this allows them to reduce the variance of their beliefs faster and extract a higher surplus from bidding. The other type of bidders includes those who initially have low variances of beliefs. In the limit, if the variance of the prior is zero, these bidders

while the convergence of Euler's steps for the differential equations is well known. Moreover, such an approach allows me to use a continuous state space which is not always feasible in the discrete time case.

will not change their strategies as price changes. Specifically, if this type of bidder deters the entry of the other potential bidders at the beginning of the auction (or the instantaneous demand is initially low), she will wait until the end of the auction to submit her final bid since there are no incentives for her to bid more frequently. In this way, she will not allow her information about the visibility of the auction to be absorbed by the price. This type of behavior constitutes learning prevention.

The component $-\left\{\frac{\sigma}{\lambda}\frac{\partial\lambda}{\partial\theta}\right\}\frac{\partial V(t,p,\mu,\sigma)}{\partial\mu}$ can be interpreted as the effect of the mean of the bidder's beliefs on the value function. Specifically, this effect depends on the variance of beliefs and the sensitivity of the entry rate $\lambda(\cdot)$ with respect to the changes in the visibility. If the variance of the belief distribution σ is small, then the effect of a change in the mean of bidder's beliefs is small. Similarly, if the instantaneous demand λ is very sensitive to changes in visibility, then the value function of the bidder will be more sensitive to changes in the mean of bidder's beliefs.

The component $-\frac{\sigma^2}{\lambda}\left(\frac{\partial\lambda}{\partial\theta}\right)^2\frac{\partial V(t,p,\mu,\sigma)}{\partial\sigma}$ can be interpreted the effect of the variance of a bidder's beliefs on her value function. This indicates that as bidders obtain more information about entry into the auction over time, their beliefs regarding visibility become closer to the truth. As a result, an increase in the variance of the bidder's belief will be followed by a decrease in the bidder's value function. Intuitively, this implies that the bidder values more precise information from the auction.

It is important to mention that my model is designed to analyze a partial equilibrium in a single auction within a context of a large auction market. In general equilibrium in the market visibility is endogenous characteristic: entry rate into particular auctions will depend on the total number of bidders in the market, the number of sellers, and characteristics of an auction (such as a position in the search contents, spelling errors in the name, number of cross-listed items).

2.6 Identification

The model of the continuous-time auction considered in this paper has a complex structure. In the theoretical section,I represented this structure in an algorithmic form which allowed me to describe the behavior of the bidders as a reaction to a stochastically changing price. The reaction of the bidders is summarized by the second highest bid which, in equilibrium, should be equal to the price in the auction. In order to fit the model to the data, one makes a "guess" about the structural functions of the model (the instantaneous demand, the size of the price jumps, the distribution of valuations, and the distribution of bidders' beliefs), simulates the optimal behavior of the bidders, computes the second highest bid at each instant, and matches the simulated second highest bid to the price that is actually observed. In order to have meaningful estimates of the structural functions, I will need to assure that the result of the specified matching procedure establishes a one-to-one correspondence between the simulated distribution and structural parameters, that is, to verify that the model is identified.

Definition 3 *We will say that the model of the continuous-time auction is identified if there is only one joint probability distribution of price jump sizes and the jump times* $\mathbf{P}\left\{\int h\left(t,p_t,\eta,\theta,\gamma\right)\,dJ\left(t,p_t,\theta,\gamma\right)<\pi,\ \int dJ\left(t,p,\theta,\gamma\right)<\tau\right\}$ *corresponding to a specific collection of characteristics of the theoretical model: instantaneous demand function* $\lambda\left(t,\,p_t\right)$*, price jump size function* $h\left(t,\,p_t,\,\eta\right)$*, the distribution of valuations* $F\left(v,\,\theta\right)$*, and the distributions of initial beliefs* $G_\mu\left(\mu_0\right)$ *and* $G_\sigma\left(\sigma_0\right)$

This definition suggests that if one simulates the equilibrium behavior of the bidders, this will also determine the characteristics of the theoretical model that produce the same simulated distribution of prices as that which is actually observed. According to this definition, identification is achieved if no other set of characteristics of the model produces the same distribution of simulated prices.

My strategy now will be to provide a set of conditions that allow me to non-parametrically identify the model using Definition 3. As a result, the potential outcome of the estimation procedure will be the complete functional form of the instantaneous demand function, the price jump function, and the distributions of bidders' valuations and beliefs. From the point of view of Definition 3, the model will be identified if two conditions hold. First, we should be able to estimate the joint distribution of the price jumps and their timing from the data. Second, we should be able to evaluate the structural functions of the model by using some method of inverting the distribution of price jumps. As a result, identification of the model is partially assured by the structure of the data and partially assured by the functional form assumptions that allow me to invert the estimated distribution. The data collected from eBay contain the complete price paths for multiple auctions. I will assume that my data collection methodology provides a uniform unbiased sample of auctions. Namely, the collected sample of auctions will be considered free of selection bias (i.e. I did not pre-select the auctions with certain visibility). Moreover, I will assume that the data are collected for very uniform and similar items, guaranteeing that the functions driving the continuous-time auction (the instantaneous demand, the size of the price jumps, and the distribution functions) are the same across auctions, conditional on individual-specific covariates. These data collection assumptions assure that the dataset possesses a certain degree of ergodicity, so that observing repetitions of the auction with certain parameters is equivalent to observing a cross-section of simultaneous similar auctions. As a result, my data have two dimensions: on the one hand, I have a complete record of individual bids corresponding to the continuous-time observations of the price process over time. On the other hand, for each instant I can see the distribution of prices across the auctions. In this way, under certain assumptions I will be able to disentangle the characteristics which are constant over time by looking at the cross sectional dimension, and then recover the remaining characteristics from the data over time.

I will now formulate a set of assumptions which will allow me to establish identification of the theoretical model.

- *Assumption 1.*

  The support of the joint distribution of prices and the price jumps

  $$\mathbf{P}\left\{\int_0^t h\left(t, p_t, \eta, \theta, \gamma\right)\, dJ\left(t, p_t, \theta, \gamma\right) < \pi,\ \int_0^t dJ\left(t, p, \theta, \gamma\right) < \tau\right\}$$

  is a convex compact set for any $t \in [0,\, T]$.

- *Assumption 2.*

  The instantaneous demand function observed by the bidders $\lambda_\epsilon\left(t, p, \theta\right)$ is decreasing with respect to the observation delay ϵ.

- *Assumption 3.*

  The instantaneous demand function observed by the bidder $\lambda_\epsilon\left(t, p, \theta\right)$ is strictly increasing in the visibility of the auction.

- *Assumption 4.*

  The bid increment function $\eta\left(t, p_t, v, \mu, \sigma\right)$ is non-decreasing in the bidder's valuation v and the mean of the prior belief of the bidder μ about the visibility of the auction, while bidder beliefs are independent from their valuations. Moreover, for each v there is a price $\overline{p}$ such that for all prices p higher than $\overline{p}$ we have $\eta\left(t, p, v, \mu, \sigma\right) = 0$.

- *Assumption 5.*

  The bid increment function $\eta\left(t, p_t, v, \mu, \sigma\right)$ is less sensitive to changes in the price for bidders with smaller initial variance of beliefs about the visibility of the auction. That is for $\sigma' > \sigma$ and $\Delta p > 0$:

  $$\left|\eta\left(t, p + \Delta p, v, \mu, \sigma'\right) - \eta\left(t, p, v, \mu, \sigma'\right)\right| \leq \left|\eta\left(t, p + \Delta p, v, \mu, \sigma\right) - \eta\left(t, p, v, \mu, \sigma\right)\right|$$

Before beginning the proof of identification, let me discuss the role of each assumption. The first assumption suggests that there are no "holes" in the joint distribution of the number of price jumps and the price. In this way, if one this distribution is observed at any moment of time, one can invert it to get the corresponding number of price jumps and the value of the price.

The second assumption allows me to specify the direction in which the observation delay influences the observed instantaneous demand. This assumption suggests that the bidders with higher observation delays are observing the price jumps less frequently.

The third assumption formalizes the notion of the visibility of the auction as a measure of "attractiveness" of a specific auction to the bidders. This enables me to estimate the visibility parameter from multiple auctions based on ranking them by the number of active bidders.

The fourth assumption allows me to distinguish aggressive bidding by bidders with high valuations from aggressive bidding by bidders with high beliefs regarding the visibility. This separation allows me to estimate the distributions of valuations and the initial beliefs on different parts of the support of the bid increment function.

The fifth assumption allows me to identify bidder types based on their different reactions to the price jumps.

The assumptions allow me to obtain a unique set of structural functions for which a set of simulations from my model replicates the observed data. I will now outline the identification strategy and then proceed to a formal proof of identification.

For some given instantaneous demand function and the price jump size function, we can compute the optimal bidding function for individual bidders. The bidding function reflects the optimal bid value for the bidder given her valuation, beliefs, and the current price of the object for each moment of the auction. The argument for identification of the distribution of valuations is similar to that in the static empirical auction literature (e.e., Guerre, Perrigne, and Vuong (2000), Campo, Guerre, Perrigne, and Vuong (2003), and Athey and Haile (2002)). If one assumes that the bidding function is monotone with respect to the valuation at each moment of time and that the beliefs are independent from valuations, then the distribution of the number of active bidders across auctions given the price will reflect the distribution of valuations. As a result, a set of cross-sectional observations will identify the distribution of valuations if we can compute the optimal bidding function.

Once the distribution of valuations is available, it becomes possible to identify the distribution of beliefs. Two facts are important for the identification of this distribution. First, if we assume that strategic learning occurs faster for bidders whose beliefs are close to the true visibility, then we will be able to identify the mean beliefs of the bidders by the distance between their observed bidding patterns and the optimal bidding pattern (computed for an auction with given visibility and given structural functions). Second, as the model predicts that bidders with more diffuse priors bid more frequently, we will be able to sort the bidders according to the relative number of their bids and, in this way, separate bidders by size of variance of initial beliefs. In this way, I will identify the distribution of the bidders' beliefs from the observations across time by measuring the relative frequencies of bidding of different bidders and the distance of their bidding patterns from the optimal one.

Finally, given the distributions of valuations and beliefs, for a given instantaneous demand and price jump size function, I will be able to simulate the behavior of the second highest bid. If we match the distribution of the simulated second highest bid and the distribution of actually observed prices over time, we reduce the estimation of the model to the estimation of parameters of the compound Poisson process. Such estimation is described in detail in the literature (for instance, in Møller and Waagepetersen (2004), Lipster and Shiryaev (2001), and Karr (1986)) and provides a unique outcome under conditions which are assumed to be satisfied here.

I now use the stated assumptions 1-5 to prove identification of my theoretical model.

Theorem 2 *Suppose that we observe distributions of the form*

$$\mathbf{P}\left\{p_t,\, J_t\right\}, \quad \textit{for each}\; t \in [0,\, T]\,,$$

where p_t is the price in the auction at time t and J_t is the number of price jumps from the beginning to time t. Then the following results are valid for identification from this collection of distributions:

- *The distribution of observation delay errors ϵ is not identified*
- *The price jump size and the instantaneous demand function are identified*
- *The distribution of valuations of the bidders is identified*
- *The distributions of the beliefs of the bidders are identified up to the scale provided by the visibility of the auction: $G_\mu\left(\frac{\mu}{\theta_0}\right)$ and $G_\sigma\left(\frac{\sigma}{[\theta_0]^2}\right)$*

Proof: The first statement follows immediately from the structure of the model. Suppose that $\lambda_\epsilon\left(t,\, p_t\right)$ is the instantaneous demand. We observe the realizations of the mean $\int \lambda_\epsilon\left(t,\, p_t\right)\, dF(\epsilon)$ because ϵ was assumed to be independent across the auctions. For any differentiable monotone transformation $\varphi(\cdot)$ we can find a constant C and a function $\phi(\cdot)$ such that:

$$\int \lambda_\epsilon\left(t,\, p_t\right)\, dF(\epsilon) = \int \varphi\left(\lambda_\epsilon\left(t,\, p_t\right)\right) \frac{1}{C}\phi\left(\epsilon\right)\, dF(\epsilon),$$

and $\frac{1}{C}\int \phi\left(\epsilon\right)\, dF(\epsilon) = 1$. Thus, the distribution of the observation delay errors cannot be identified simultaneously with the instantaneous demand function. For this reason I set this distribution to be uniform and impose the normalizations $\lambda_\epsilon\left(t, x\right)\big|_{\epsilon=0} \lambda\left(t, x\right)$ and $\lambda_{\epsilon'}\left(t, x\right) < \lambda_\epsilon\left(t, x\right)$ if $\epsilon' > \epsilon$. This normalization implies that the instantaneous demand observed without any delay is equal to the actual instantaneous demand.

The argument for the second statement is the same as that for the general Poisson processes. It known that the parameters of a Poisson process of the this kind are identified.

Now consider the last two statements. By assumption from the proof of the first statement of the theorem, the distribution of the observation delay errors is uniform on $[0,\, 1]$. The instantaneous demand is assumed to be monotone with respect to the observation delay errors. If $\overline{N}$ refers to the maximum number of bidders observed in the auction, the conditional distribution of the second highest bid increment in the auction can be written as:

$$\mathbf{P}_\epsilon\left(\max_{(2)} \eta\left(t, p, v, \mu_0, \sigma_0, \epsilon\right) < \pi\right) = \sum_{k=2}^{\overline{N}} \binom{\overline{N}}{k} \left[\frac{\pi}{\eta\left(\cdot, 1\right) - \eta\left(\cdot, 0\right)}\right]^k \left[1 - \frac{\pi}{\eta\left(\cdot, 1\right) - \eta\left(\cdot, 0\right)}\right]^{\overline{N}-k}$$

The bid increment function is non-decreasing in the valuation of the bidder. This formula shows that the probability $\mathbf{P}_\epsilon\left(\max_{(2)} \eta(t,p,v,\mu_0,\sigma_0,\epsilon) < \pi\right)$ is a monotone one-to one map from the bid increment function to the the conditional distribution of the second highest bid. As a result, I can conclude that the conditional probability of the bid increment function is increasing in the valuations of the bidders. Then for two subsets $\mathbb{V}_1$ and $\mathbb{V}_2$ of $\mathbb{R}_+$ such that $\mathbb{V}_1 \leq \mathbb{V}_2$, the probability distributions of the second highest bid given the initial belief are also ordered:

$$\mathbf{P}_\epsilon\left(\sum_{\tau<t} \max_{(2)} \eta(t,p,v_1,\mu_0,\sigma_0,\epsilon) < \pi\right) \leq \mathbf{P}_\epsilon\left(\sum_{\tau<t} \max_{(2)} \eta(t,p,v_2,\mu_0,\sigma_0,\epsilon) < \pi\right),$$

for all $v_1 \in \mathbb{V}_1$ and $v_2 \in \mathbb{V}_2$, and the summation $\sum\limits_{\tau<t}$ is over all discrete moments of price jumps up to time t. This means that we can partition the distribution of valuations into very small regions. If the unconditional distribution of bid increments in two distinct small regions is the same, the density of valuations in the region with higher valuations is lower.

Now let me consider whether I can provide a similar ordering for the conditional distribution of the second highest bid with respect to the parameters of the bidder's beliefs. By assumption 4, a higher mean of the belief distribution implies that the bidder expects a higher demand for the object. As it follows from the model, the bid increment function is non-decreasing with the instantaneous demand, implying that the bid increment function is non-decreasing in the mean of the bidder's belief. In this case, for two subsets of Θ such that $\mathbb{T}_1 \leq \mathbb{T}_2$ we should have that:

$$\mathbf{P}_\epsilon\left(\sum_{\tau<t} \max_{(2)} \eta(t,p,v,\mu_1,\sigma_0,\epsilon) < \pi\right) \leq \mathbf{P}_\epsilon\left(\sum_{\tau<t} \max_{(2)} \eta(t,p,v,\mu_2,\sigma_0,\epsilon) < \pi\right),$$

such that $\mu_1 \in \mathbb{T}_1$ and $\mu_2 \in \mathbb{T}_2$.

The last parameter of interest is the variance of the bidders' beliefs at the beginning of the auction. By Assumption 5, the squared bid increment of bidders with more diffuse priors is more sensitive to the price than that of bidders with less diffuse priors. By monotonicity, the distribution of the second highest bid is more concentrated if the variance of the initial belief is smaller. In this way, the more informed bidders (those with lower variance of beliefs regarding visibility) will be using bidding strategies that are close to the optimal one. As a result, we can write that for two sets $\Sigma_1 \leq \Sigma_2$:

$$\int \mathbf{P}_\epsilon\left(\sum_{\tau<t} \max_{(2)} \eta(t,p,v,\mu_0,\sigma_1,\epsilon) < \pi\right) \mathbb{J}_\gamma(dt,dp,\theta)$$
$$\leq \int \mathbf{P}_\epsilon\left(\sum_{\tau<t} \max_{(2)} \eta(t,p,v,\mu_0,\sigma_2,\epsilon) < \pi\right) \mathbb{J}_\gamma(dt,dp,\theta),$$

if $\sigma_1 \in \Sigma_1$ and $\sigma_2 \in \Sigma_2$. We can denote the integral

$$\int \mathbf{P}_\epsilon \left(\sum_{\tau<t} \max_{(2)} \eta\left(t, p, v, \mu, \sigma, \epsilon\right) < \pi \right) \mathbb{J}_\gamma \left(dt, dp, \theta\right) = K\left(v, \mu, \sigma, \pi\right).$$

I have provided a set of assumptions assuring that the kernel function $K(\cdot)$ will be monotone with respect to its arguments. Moreover, given the specific Poisson measure $\mathbb{J}_\gamma\left(t, p, \theta\right)$ we can compute this kernel function. As a result, we reduce the problem of finding the unknown distribution to the problem of solving the integral equation:

$$\mathbf{P}\left\{p_t < \pi\right\} = \int \int \int K\left(v, \mu, \sigma, \pi\right) dF\left(v, \theta\right) dG_\mu\left(\mu\right) dG_\sigma\left(\sigma\right)$$

with a monotone kernel. This is a Volterra-type integral equation for the unknown distribution functions, which is known to have a unique solution. This suggests that given the price process and the visibility parameter we can reproduce the distribution of the valuations and the beliefs of the bidders. To find the unknown distributions, I find the subsets of valuations and the parameters of the bidders' corresponding to the same cumulative probability of the auction price. Then, if in one region the kernel is larger than in another, we can conclude that the probability density in the first region is smaller. In this way I rank the probability masses in all points of the distributions under consideration. As the probability masses must sum to one given the rank for each point, I find the actual probability mass in this point.

Finally, note that given the monotonicity of the instantaneous demand with respect to the visibility of the auction, we can shift the visibility of all auctions by the same constant. This will result in an equivalent shift of the bidder's beliefs. In this case, if the visibility parameter is not fixed, we cannot separately identify the distribution of beliefs and the visibility parameter. For this reason, I restrict the support of the visibility of the auction to the segment $[0, 1]$. In this case, we can estimate the visibility by ranking the auctions according to the level of instantaneous demand when price and time are the same.

If the instantaneous demand function and the size of price jumps are estimated from the collection of price trajectories, then they will represent the characteristics of the price process averaged over characteristics of the bidders: valuations and observation errors. Therefore, they characterize the conditional moments of the price process - average jump magnitude given time and price and an average frequency. Nonparametric identification of such a system of moments can be formally justified using arguments from the identification theory for nonseparable models (such as in Chesher (2003) and Matzkin (2003)), as well as the analysis of nonseparable systems with an independent instrument (as in Imbens and Newey (2002)). As I have shown above, the structure of the simulations from the model is quite similar to that in nonseparable triangular systems of moments considered in Chesher (2003), providing an additional argument for identification.

3 Estimation

Differential equations (7) rarely have closed form solutions. Specifically, if the model is formulated non-parametrically, then the solution of the optimal bidding problem can only be obtained numerically. To describe the complete equilibrium model, it would be necessary to solve a large set of individual bidding problems, making calculations of the likelihood intractable.

The existing methods of empirical analysis of complex structural models suggest using simulation to estimate the structural parameters. Examples of these methods are simulated method of moments (including indirect inference as a special case[21]) and two-stage inference (including the hedonic approach[22]). Unfortunately, these methods are not suited for the analysis of continuous-time models[23].

I provide a new method for estimation of a continuous-time model as that described above. The data allow me to observe all the actual bids for a specific auction retrospectively. As a result, I have continuous-time observations for each auction, and the estimation method suggested in this paper uses this feature. The idea of the method is the following. From the data we observe the actual price in the auction, which is equal to the second highest bid at any instant. From the model, I can compute the optimal response of a number of bidders to the price movement. This response describes the bids of bidders given the price. If we simulate the entry of bidders, we can compute the second highest bid for the specified group of bidders, which I will call the response of the structural model to the price. In equilibrium, this response should coincide with the price. As the process of entry is stochastic, the distribution of this response should coincide with the distribution of the price. By minimizing the distance between the distribution of the price and the distribution of the response of the structural model, we can estimate the "true" values of the parameters of the structural model.

Let me now describe the estimation algorithm in more detail. A defining feature of this paper is that I consider the entire price path in the auction as the unit of observation, rather than individual bids (the latter is the case in most existing empirical research on auctions). This captures

[21]This method usually referred to as indirect inference (II) was analyzed in Gourieroux, Monfort, and Renault (1993) and elaborated later to its efficient version adapted to the score estimation method with the auxiliary model generated by the semiparametric estimate of the likelihood of the data, described for example in Gallant and Tauchen (2002).

[22]For instance the two-stage method proved to be useful for estimation of models of imperfect competition as in Bajari, Benkard, and Levin (2007) and Pesendorfer and Schmidt-Dengler (2003) of dynamic games as in Hotz and Miller (1991) and, auctions as in Guerre, Perrigne, and Vuong (2000), differentiated product markets as in Berry, Levinsohn, and Pakes (1995) and Bajari and Benkard (2002).

[23]The simulated methods of moments relies on the specific choice moment equations, which is unclear in the model under consideration. The two-stage methods have not been developed for continuous-time models and they can have extremely large standard errors.

the continuous-time structure of the model, where the equilibrium is described by the stochastic behavior of price. Let p_t be the observed price in the auction at time t, N_t the number of bidders in the auction at time t, and suppose that the auction proceeds from time 0 to time T. Let $f(p_t,\, N_t,\, t,\, \gamma_0)$ denote the joint distribution of the object price and the time when the price jumps[24]. This distribution is characterized by a structural parameter $\gamma_0 \in \Gamma$. Since it is possible to observe prices of multiple auctions, it is possible to record several price trajectories and estimate the marginal distribution of price p_t and jump time t using a kernel estimator:

$$\widehat{f}\,(p,\, t) = \frac{1}{n}\sum_{k=1}^{n}\frac{1}{h_p\, h_t}\sum_{i=1}^{I_k}\kappa\left(\frac{p_i^{(k)}-p}{h_p}\right)\kappa\left(\frac{t_i^{(k)}-t}{h_t}\right), \tag{9}$$

where n is the number of observed auctions, k is the index of an auction, I_k is the number of price jumps on the auction k, $\kappa(\cdot)$ is a kernel function, and h_p and h_t are bandwidth parameters.

In estimation, I impose several restrictions on the kernel function and bandwidth parameters. These conditions are needed to provide consistency and asymptotic normality of the estimates[25].

A problem for estimation in a continuous-time auction is that the movement of the price depends on the entry of bidders, which is considered latent. If there is a large number of overlapping auctions (so that entry of bidders to the auction can be considered independent across auctions), the kernel estimator will still yield the correct estimate of the marginal distribution of prices and the moments of price jumps. Formally this implies that, under the aforementioned restrictions:

$$\widehat{f}\,(p,\, t) \overset{n\to\infty}{\longrightarrow} \int f(p,\, N,\, t,\, \gamma_0)\, dN,$$

where N denotes the latent number of bidders who have entered into the auction. A proof of this fact is provided in Nekipelov (2007). The listed restrictions on the kernel functions and bandwidth parameters are needed to establish the asymptotic normality of the resulting estimates. This suggests that for each price p and jump time t, the estimate of the density $\widehat{f}\,(\cdot)$ is asymptotically normal:

$$\sqrt{nh_th_p}\left(\widehat{f}(p,t) - f(p,t)\right) \overset{d}{\longrightarrow} N\left(0,\, f(p,t)\left\{\int_0^{\infty}\kappa^2\left(\psi\right)\, d\psi\right\}^2\right). \tag{10}$$

This is an important result establishing that even in the continuous-time case, the density of the estimate possesses the property of asymptotic normality.

[24] I assume that the joint distribution of realizations of stochastic processes y_t and N_t exists. Formal conditions can be found in Gihman and Skorohod (1979).

[25] These restrictions basically specify that the density that I am trying to estimate actually exists, the kernel function is very smooth and decreases at a fast rate to suppress the influence of "outliers", and that bandwidth parameters go to zero as the sample size increases so that there is no asymptotic bias in the estimates.

I first simulate the entry of bidders N_t given parameter vector γ for every instant t. Given the structural parameter γ and the optimal bidding problem for each price, I can then calculate the second highest bid in the auction at any given instant. This determines the "response" of the structural model to the data:

$$\widehat{p}_i^{(k)}(\gamma) = \varphi_{\widehat{N},\gamma}\left(p_i^{(k)}\right).$$

I have already assumed that entry of the bidders is independent across auctions. Therefore, similar to the previous arguments, we can estimate the marginal density of the second highest bid $\widehat{p}(\gamma)$ for every t given parameter vector γ by the kernel estimator as:

$$\widehat{f}_\gamma(p,\, t) = \frac{1}{n}\sum_{k=1}^{n}\frac{1}{h_t h_p}\sum_{i=1}^{I_k}\kappa\left(\frac{\widehat{p}_i^{(k)}(\gamma)-p}{h_p}\right)\kappa\left(\frac{t_i^{(k)}-t}{h_t}\right),$$

By construction, the number of bidders who have entered the auction $\widehat{N}$ is independent across auctions. This suggests that:

$$\widehat{f}_\gamma(p,\, t) \xrightarrow{p} \int f(p_t,\, N,\, t,\, \gamma)dN,$$

where convergence is justified by the same arguments as in (15).

At this point, given that the data generating process is induced by the structural model under analysis, we can check whether the observed data were generated from the model with the structural parameter γ. A problem that arises here is that the structural models $\varphi_{N,\gamma}$ for different vectors γ are not nested, so that a standard specification test will not apply in this case. Since my task is analogous to a model selection problem, it is convenient to use the Kullback - Leibler Information Criterion (KLIC), which has proven to be useful even for a choice between nonparametrically specified models[26]. The idea now will be to compare the empirical model with the structural model (based on the simulated response of bidders). At the point when the parameter of the structural model is $\gamma = \gamma_0$, both models should work equally well.

Let I_k be the number of price jumps in the auction k and n be the total number of auctions. The Kullback - Leibler Information Criterion takes the form:

$$\widehat{\text{KLIC}} = \frac{1}{n}\sum_{k=1}^{n}\sum_{i=1}^{I_k}\log\left[\frac{\widehat{f}(p_i^{(k)},\, t_i^{(k)})}{\widehat{f}_\gamma(p_i^{(k)},\, t_i^{(k)})}\right] \tag{11}$$

For a consistent estimator of the density, $\widehat{f}(\cdot)$ goes to $f(\cdot)$ pointwise. In Nekipelov (2007), it is shown that the estimate for the structural parameter γ obtained from minimization of the KLIC function will be asymptotically normal, so that:

$$\sqrt{nh_t h_p}\,(\widehat{\gamma}-\gamma_0) \xrightarrow{d} N\left(0, Q^{-1}\,\Omega\, Q^{-1}\right), \tag{12}$$

[26]Vuong (1989) and Chen, Hong, and Shum (2005) show that the KLIC is a powerful tool for model selection.

where

$$Q = E\left\{\int_0^T \frac{\partial f(p_\tau,\tau,\gamma_0)/\partial\gamma}{f(p_\tau,\tau,\gamma_0)}\, dJ(\tau,p_\tau)\right\} \quad \text{and} \quad \Omega = 2\left(\int_0^\infty \kappa^2(\psi)\, d\psi\right)^2 E\left\{\int_0^T \frac{dJ(\tau,p_\tau)}{f(p_\tau,\tau,\gamma_0)}\right\}.$$

A significant problem that might arise in this context is that optimization of (11) requires simulating the model response $\widehat{p}$ for each parameter value γ. For especially large models, this simulation can be extremely time consuming. However, the computational burden can be reduced substantially by using Bayesian methods for simulation. In fact, unlike a deterministic optimization procedure, a Bayesian estimation procedure utilizes all of the intermediate output (such as function and parameter values) to form the posterior distribution[27].

Let me first describe the procedure for estimation through MCMC and then characterize the asymptotic properties of the obtained estimates. First, we obtain a non-parametric estimate of the density of the data $\left\{\left(p_i^{(k)}, t_i^{(k)}\right)_{i=1}^{I_k}\right\}_{k=1}^{n}$. Then we construct a random walk sampler from the density $\varphi(\gamma) = \zeta \exp\left\{-\widehat{KLIC}\,(\gamma)\right\}$ for some normalizing constant ζ, where the hat refers to the fact that the criterion uses an estimate of the density. The random walk sampler is paired with the Metropolis-Hastings method for sampling from an arbitrary density. The method is performed as follows.

The algorithm is initialized at some parameter value γ_0. One then proposes a parameter vector γ drawn from the standard normal distribution (when sampling, it may be known that some of the parameters are confined in some area, in which case the draws are made from the truncated normal distribution). If the KLIC function decreases at the new parameter value, the draw is "accepted" (i.e. we add it to the database of the draws from $\varphi(\gamma)$). If the KLIC increases, then we make a draw z from the uniform distribution on [0, 1] and if $\exp\left\{\widehat{KLIC}(\gamma_0) - \widehat{KLIC}(\gamma)\right\} < z$, then the draw is "accepted", otherwise it is "rejected" and the current value of γ remains at the level γ_0. This process is repeated sufficiently many times to achieve convergence (meaning that the Markov chain of γ becomes stable)[28].

Note that these expressions are valid only if the generated Markov chain is stationary. Stationar-

[27] Chernozhukov and Hong (2004) have shown that for specific classes of loss functions in minimum distance - type procedures, the first and second moments of posterior distribution formed from the minimum distance criterion have the same asymptotic properties as the maximizer to the minimum distance criterion and its variance.

[28] The fact that the resulting values of γ will represent the distribution with the density $\varphi(\cdot)$ is possible to establish in the following way. Let $\gamma_0 \sim \varphi(\gamma)$. Then acceptance of γ is justified by the fact that $\exp\left\{\widehat{KLIC}(\gamma_0) - \widehat{KLIC}(\gamma)\right\} < z$. The probability of this event:

$$\mathbf{P}\left\{\gamma \text{ is accepted}|\gamma_0\right\} = \int_0^1 \mathbf{1}\left\{e^{\left\{\widehat{KLIC}(\gamma_0)-\widehat{KLIC}(\gamma)\right\}} < z\right\} dz = \exp\left\{\widehat{KLIC}(\gamma_0) - \widehat{KLIC}(\gamma)\right\}$$

The unconditional density is thus expressed as: $f(\gamma) = \mathbf{P}\left\{\gamma \text{ is accepted}|\gamma_0\right\}\varphi(\gamma_0) = \zeta\, e^{\left\{-\widehat{KLIC}(\gamma)\right\}}$, which is exactly the density $\varphi(\gamma)$. This implies that γ is a draw from this density.

ity conditions are described in Robert and Casella (2006), and are satisfied for posteriors generated by minimum distance estimators with a single minimum and a distance measure increasing at a high enough rate.

Given a significant flexibility of the estimation method, it does not achieve the $\sqrt{n}$-convergence. However, the model is highly non-linear, and even to achieve a non-parametric convergence rate, we need to impose a list of restrictions on the data.

Assumptions

A 1 Realizations $\left\{p_\tau^{(k)},\, t_\tau^{(k)},\, N_\tau^{(k)}\right\}_{\tau\in[0,\,T]}$ are independent for each $\tau \in [0, T]$ and $k = 1, \dots, n$, have joint and marginal distributions with finite second moments.

A 2 The kernel is a Borel-measurable function such that

- $\int \kappa(\psi) d\psi = 1$,
- $\int |\kappa(\psi)| d\psi < \infty$,
- $\kappa(\cdot)$ is twice continuously differentiable a.s. and approaches zero with its first and second derivatives as $|\psi| \to \infty$.
- $\sup_{\psi\in\mathbb{R}} |\kappa(\psi)| < \infty$

A 3 The marginal distribution function of the process $\left\{p_\tau^{(k)}, t_\tau^{(k)}\right\}_{\tau\in[0,\,T]}$ has a continuous Hessian which is bounded in some neighborhood of each point of the domain.

A 4 $h_t,\ h_p \to 0$ as $n \to \infty$.

A 5 $n\, h_t\, h_p \to \infty$ as $n \to \infty$.

These assumptions ensure the appropriate convergence properties of the non-parametric density estimates. Moreover, these properties are helpful for the proof of asymptotic normality of the structural estimates. The following theorem establishes the convergence properties of the estimated densities and the properties of the structural estimates.

Theorem 3 *Suppose that Assumptions 1-5 hold. Then the density of the distribution of price jumps can be consistently estimated at step 1.*

Assume that in addition we require that

$$E\left\{\int_0^T \frac{\lambda(\tau, p_\tau)}{f^2(p_\tau, \tau, \gamma_0)}\, d\tau\right\} < \infty,$$

then the structural estimates obtained by minimization at step 7 are asymptotically normal:

$$\sqrt{n h_t h_p}\, (\widehat{\gamma} - \gamma_0) \xrightarrow{d} N\left(0, Q^{-1}\, \Omega\, Q^{-1}\right),$$

where

$$Q = E\left\{\int_0^T \frac{\partial f(p_\tau,\tau,\gamma_0)/\partial\gamma}{f(p_\tau,\tau,\gamma_0)}\, dJ(\tau, p_\tau)\right\} \quad \textit{and} \quad \Omega = 2\left(\int_0^\infty \kappa^2(\psi)\, d\psi\right)^2 E\left\{\int_0^T \frac{dJ(\tau,p_\tau)}{f(p_\tau,\tau,\gamma_0)}\right\}.$$

The proof of this theorem is given in Appendix D.

Although Theorem 3 looks very similar to the standard asymptotic results in the non-linear settings, it is very different in the nature from the results derived for i.i.d. observations. In our case we have to employ the fact that different price paths are generated by the same stochastic process. Surprisingly, summation over observations as in 11 fulfills this task, because we can now look at the function of the entire price path. These functions should be independent and identically distributed across different paths. Theorem 3 further projects these properties onto the regularity properties of the structural estimates.

Chernozhukov and Hong (2004) show that the asymptotic behavior of the posterior mean and variance is the same as the asymptotic behavior of the estimates implied by the minimum distance function. In Nekipelov (2007) it is shown that the asymptotic behavior of the MCMC estimate of γ under a fixed, non-parametric estimation of the data density is:

$$\sqrt{nh_t h_p}\,(\widehat{\gamma}_{MCMC} - \gamma_0) \xrightarrow{d} N\left(0, \frac{1}{2}Q^{-1}\,\Omega\,Q^{-1}\right) \tag{13}$$

As compared to the asymptotic variance, the variance of the MCMC estimate has an additional factor of $\frac{1}{2}$. This occurs because the variance of the non-parametric estimate of the data density is not taken into account in the MCMC procedure. To obtain the correct variance estimate, we need to double the MCMC variance estimate.

4 Results of structural estimation

The theoretical section developed a continuous-time model of eBay auctions in which the optimal bidding problem is expressed as a partial differential equation characterizing the law of motion of the expected payoff from winning the auction. However, I also argued that the computation of the likelihood of the complete equilibrium model of the auction is infeasible. To resolve this problem, I proposed an estimation method based on the assumption that, in equilibrium, the "response" of the bidders as reflected in the optimal second highest bid should be equal to the price in the auction.

In the previous part I proved that the model is non-parametrically identified. Although this analysis is valid for general non-parametric inference, my estimation procedure is based on a Markov Chain Monte Carlo approach. This means, in particular, that we will need to estimate the posterior distribution of the structural parameters using a sequence of simulations. However, the computational burden of the simulation increases with the dimension of the parameter vector, requiring more evaluations of the distance function. For this reason, I will need to specify the structural model parsimoniously enough to allow the chain to explore the posterior parameter distribution in an efficient manner. These parametric assumptions, hence, are made solely for computational convenience.

The baseline structural functions describing the dynamics of the price are the instantaneous demand and the size of price jumps. The specification of the instantaneous demand as observed by the individual bidder (given the observation error ϵ) was assumed to have a logistic form:

$$\lambda_\epsilon(t,\, x,\, \theta) = \alpha\, \epsilon \exp\left(\theta\right) \frac{exp\left(\mathcal{P}(t,x)\right)}{1+exp\left(\mathcal{P}(t,x)\right)} = \alpha_\lambda\, \epsilon \mathcal{L}\left(\mathcal{P}(t,x)\right),$$

where $\mathcal{P}(t,x)$ is a polynomial in time and price, $\mathcal{L}(\cdot)$ is the logistic function, and ϵ is uniformly distributed on [0, 1]. The logistic structure of the frequency rules out cases where the frequency is very large, making the optimization procedure unstable[29]. Flexibility in the structure is captured by a polynomial form in the argument of the logistic function.

The magnitude of the price jumps is generated by the logistic form, similar to that for the instantaneous demand function

$$h(t,x,\eta) = \alpha_h \mathcal{L}\left(\mathcal{Q}(t,x,\eta)\right),$$

where $\mathcal{Q}(\cdot)$ is a polynomial function of time, price and the individual bid increment. This polynomial structure allows a semi-parametric representation of both the instantaneous demand and the size of the price jumps.

The non-parametric flexibility of the structural model increases with the power of the polynomials $\mathcal{P}$ and $\mathcal{Q}$. However, this also increases the number of unknown coefficients in the polynomial representations, leading to a decrease in computation speed. For this reason the degrees of polynomials $\mathcal{P}$ and $\mathcal{Q}$ were restricted to 2. The polynomial $\mathcal{P}\left(\cdot\right)$ in the instantaneous demand function is thus quadratic in time t and current price x, taking the form:

$$\mathcal{P}\left(t,\, x\right) = a_0 + a_1\, x + a_2\, x\, t + a_3\, x^2$$

The price jump size function h is an exponential of the quadratic function of time t, price x, and the bid increment η:

$$\mathcal{Q}\left(t,\, x,\, \eta\right) = b_0 + b_1\, x + b_2\, t + b_3\, \eta\, x + b_4\, \eta + b_5\, \eta^2.$$

The parameters characterizing individual bidders in the continuous-time auction model are their valuations and the parameters of their prior beliefs regarding the visibility of the auction. The valuation of each player was assumed to take the form $v = \underline{v} + \sigma_v\, \xi^2$, where ξ is a standard normal random variable. The parameters $\underline{v}$ and σ_v are estimated. Since this implies that the

[29]For large frequencies the solution to the individual bidding problem becomes unstable

normal variables ξ are independent across bidders, the bidders have private valuations as in the theoretical model.

To simplify the computations of the individual bidding problem, I assumed that there are only three types of bidders (the intermediate type includes all bidders except for the bidders with degenerate beliefs), and each entering bidder is exogenously assigned a type with certain probability. The first type of bidders - the non-strategic bidders - bid their valuations immediately after entry. The probability of the first type is denoted Δ_u. The second type of bidders - the learning bidders - have imprecise information about the visibility of the auction. I estimate the parameters of the distribution of the initial beliefs of these bidders - the mean μ_θ and variance σ_θ. The distributions of initial beliefs across bidders were taken to be normal, truncated to be positive. The probability of the second type is denoted Δ_i. The third type of bidders - the experienced bidders - are assumed to know the visibility of the auction exactly and behave according to the optimal strategy. The probability of the third type is the complement $1 - \Delta_i - \Delta_u$.

I used my MCMC estimation approach to recover the parameters of bidding for pop music CDs on eBay. The estimation was based on data for 1441 eBay auctions. The estimation procedure was organized in the following way. For a sample of 50 potential bidders, I generated the observation errors from a uniform distribution on $[0, 1]$ and the private values of the bidders from the distribution of valuations. Then for the proposed values of the structural parameters α, a_i $i = 1, \ldots, 3$, b_i $i = 1, \ldots, 5$, Δ_u and Δ_i, I computed the optimal bid for each bidder and identified the second highest bid. On the basis of this bid, I computed the value of the KLIC. This calculation was incorporated into a Metropolis-Hastings procedure, drawing 15,000 MCMC observations from the quasi-posterior once the chain became visually stable. The estimation results are represented below.

$$
\begin{array}{lllllllllll}
\mathcal{P}(t,\,x) = & 8.7611 & - & 10.025x & + & 15.2702t & + & 23.4179tx & - & 29.8975x^2 \\
 & (.011) & & (4.313) & & (10.9870) & & (13.3762) & & (12.2918) \\
\lambda(t,\,x,\,\theta) = & 4.9091 & \epsilon e^{\theta}\mathcal{L}(\mathcal{P}) & & & & & & & \\
 & (2.2520) & & & & & & & & \\
\mathcal{Q}(t,\,x,\,\eta) = & 38.2823 & + & 38.1648x & - & 55.0350t & + & 44.8797\eta & - & 68.9198\eta^2 \\
 & (21.9060) & & (4.313) & & (27.0920) & & (18.5953) & & (35.0370) \\
h(t,\,x,\,\eta) = & 24.7271 & \mathcal{L}(\mathcal{Q}) & & & & & & & \\
 & (15.3709) & & & & & & & &
\end{array}
$$

$E[\mu_\theta] =$	0.5073	$\text{var}[\mu_\theta] =$	0.6164
	(0.2823)		(0.3991)
$E[\sigma_\theta] =$	0.8781	$\text{var}[\sigma_\theta] =$	0.8408
	(0.3642)		(0.4226)
$\Delta_i =$	0.1697	$\Delta_u =$	0.0982
	(.0445)		(0.876)
$\underline{v} =$	0.6636	$\sigma_v =$	9.1738
	(0.7723)		(4.4822)

There are several features which can be noted from these structural estimates. The first equation shows the estimated parameters for the polynomial arguments of the instantaneous demand. The results suggest that the instantaneous demand is downward sloping in price. It also becomes larger over time. Moreover, the quadratic term for the price variable is significant and negative, implying that the instantaneous demand is less price-elastic when price is high. This suggests that it is optimal for the bidders to raise the price when it is relatively low in order to have a greater impact on entry by other bidders. One can also see that this can explain the observed pattern of price jumps in the auctions for pop music CDs. In fact, I have observed that the size of price jumps tends to fall towards the end of the auction, consistent with my observation that the impact of the individual bids on the price is smaller when the price is high.

The second equation describes the instantaneous demand function in terms of the ”scale” parameter α. This parameter can be interpreted as the average expected number of bidders who enter the auction. The estimated parameter is equal to 4.9, consistent with the observed bidding patterns since the average number of bidders in an auction is 3.82 in the data.

The third equation represents the polynomial in the price jump size function. The coefficient estimates suggest that the price jumps become larger as we approach the end of the auction. The influence of individual bids decreases with the price of the object (as the logistic function approaches its horizontal asymptote). When the price of the object is high, individual bids have a smaller effect on the size of price jumps in the auction. If the price is fixed, then the influence of the individual bids increases with the size of the individual bid increments. This is reflected in the positive coefficient on the quadratic term of the individual bid increment.

The last four values are the parameters of the distribution of valuations and the shares of the first two types of bidders. The parameter for the distribution of valuations suggests that the lower bound of the support of valuations is statistically indistinguishable from zero. This also implies that the expected valuation of the bidders is equal to \$18.2 (from the mean of the χ^2 distribution). Another key parameter is the proportion of non-strategic bidders Δ_u. The value of the estimated proportion and the corresponding standard error suggest that I cannot reject the hypothesis that the actual number of non-strategic bidders is zero. From parameter Δ_i the proportion of the "learning" bidders is 16.9% and is statistically significant. Given these results, I can claim that the estimates of the structural model suggest that the majority of bidders are in fact behaving strategically.

4.1 Counterfactual simulations of optimal behavior

Previous analysis shows that my continuous-time auction model is identified from the data. In particular, this implies that if I obtain consistent estimates of the parameters of the model, simulations can be used to demonstrate the properties of the true model. Therefore, it is possible to illustrate how bidding behavior changes in response to exogenous changes in the determinants of the rate entry into the auction and bidder beliefs (i.e., comparative statics). The results of the analysis in this section quantify the effect of the parameters of the model (specifically, the parameters of the entry rate and the bidder beliefs) on the timing and sizes of bids.

Previously I have described two kinds of aggressive bidding behavior that arise in the model. The first is entry deterrence, when bidders raise the price in the auction early to prevent the entry of other bidders. The second is learning prevention, when bidders who are more informed about the unobserved visibility of the auction tend to bid late to hide their information. In general, we can see that these types of behavior are reflected by the timing of bids. For this reason, I will be focusing my analysis of the model at the timing. In addition, as aggressive behavior might lead to relatively large bids of the bidders, I will also consider the average size of bids in the auctions (relative to the final price).

Counterfactual simulations for the given parameters are much easier to implement than simulations for estimation of the model. To estimate the model, I needed to evaluate the behavior of many bidders and compare the second highest bid among these bidders to the actual price in the auction. To make counterfactual simulations, it is sufficient to look at the behavior of a single bidder (because the behavior of other bidders in the auction is fully described by the Poisson price jumps). During the simulations I record the fraction of early bids submitted by a bidder as well as the size of her bids (as proportion of the final price)[30]. If we are able to simulate a "representative" sample of bidders (in terms of their valuations and beliefs), then will consistently describe the

[30]For convenience of exposition I define an early bid as a bid submitted during the first 60% of the auction duration. Additional simulations, however, show that the results are not qualitatively different if the timing of early bidding moves up to 95% of the auction duration.

characteristics of the model by looking at the features of their bidding[31].

In the subsequent discussion I chose to analyze the model by changing the appropriate coefficients proportional to their values at the benchmark (estimated parameter values). The first feature of interest is the motive for learning prevention. Theoretical considerations suggest that a definitive role in the incentive to prevent learning of other bidders is played by the variance of the bidder's beliefs. The behavior of the bidder with a large variance of beliefs about the visibility was earlier referred to as "strategic learning". Such a bidder should frequently update her bids in response to price jumps. The behavior of bidders with small variance of beliefs is associated with learning prevention. To analyze the influence of beliefs on the size and timing of bids I moved the average variance of initial beliefs of the bidders. Table 3 shows changes in early bidding behavior and sizes of bids if the average variance of bidders' beliefs is set to 0.5, 1.5 and 2.0 of the benchmark value. The fraction of early bids tends to increase if the average variance of the bidders beliefs increases. Bid sizes tend to diminish with an increase in the average variance of beliefs. This implies that if the variance of beliefs of a "representative" bidder is high, then we can observe the bidding pattern attributed earlier to strategic learning. Specifically, bids become more frequent and more incremental. On the other hand, if the beliefs of a representative bidder are precise, then early and frequent bidding become rare due to learning prevention.

To study entry deterrence, I analyzed the response of the model to the changes in the instantaneous demand. In the discussion of the theoretical model, I noted that the sensitivity of entry into the auction to price determines entry deterrence behavior. In my model entry into the auction is characterized by the frequency of the Poisson entry process $\lambda(\cdot)$. In the estimation this frequency was chosen to be a logistic function a quadratic polynomial of time and price, multiplied by a scale factor. To change the sensitivity of entry with respect to price, I moved the scale factor α and the coefficient a_1 for the linear price term in the quadratic polynomial. For both coefficients an increase in the coefficient leads to an increase in the marginal entry rate. Table 3 shows the results of simulations of changes in the bidding behavior in response to changes in α and a_1. One can see that, in general, increases in both coefficients lead to increases in early bidding and the average

[31] A concrete implementation of the counterfactual simulations will be organized in the following way. (i) choose specific structural parameters which identify the model. The benchmark model is the model with the parameters obtained from the structural estimation. (ii) given the structural parameters draw the characteristics of the bidder: valuation, mean and variance of the initial beliefs about the visibility of the auction, and observation delay error. The distributions of these characteristics are defined by the structural parameters. (iii) for fixed characteristics of the bidder compute the optimal bidding function. The optimal bidding function reflects the value of the optimal bid given time, price and current beliefs about the visibility of the auction. (iv) simulate the price behavior taking into account the optimal bidding function. This allows me to record the timing and the values of bids of the bidder of interest. I repeat steps one to four sufficiently many times (in this analysis - 1000 times) to build a representative sample of the bidders.

sizes of bids. The numbers, however, show that the response to changes in the scale factor α is larger. This reflects the logistic structure of the Poisson rate $\lambda(\cdot)$ such that large increases in its argument do not lead to large changes in the function values. The results of simulations, however, confirm the theoretical prediction that if the sensitivity of entry with respect to price increases, this leads to more frequent early bidding to prevent entry of other bidders.

Simulation analysis of entry deterrence produces one important testable implication of the model. Specifically, if we observe two markets with similar overall number of bidders, but in one market entry is more price-sensitive, we should expect more early bidding in this market. In general, the entry rate reflects the search equilibrium in the auction market and it is determined by bidder's choices of auctions. One of determining factors for the probability of auction choice is the number of available choices. If the number of choices grows, while the number of potential bidders stays the same, then the probability of entry into a specific auction decreases. In addition, a high price in the auction negatively affects the surplus of the entering bidders. Therefore the probability of entry into the auction with a high price decreases if the number of choices grows. As a result, the sensitivity of entry with respect to price decreases if the the number of available auctions increases and the average number of bidders in the market remains the same. One way of empirically testing this pattern is to analyze the variation in the early bidding behavior across different categories of items on eBay with different thickness of markets, controlling for the number of bidders. *Ceteris paribus* the frequency of early bidding should be positively correlated with the number of listed auctions. Another way of testing this implication for early bidding is to run a field experiment where the number of auctions in the market is increased for a period of time. The simulation results imply that we should observe an increase in early bidding behavior in these circumstances. The results of such experiment are discussed in the next section.

5 Bidding for musical CDs on eBay: A Field Experiment

5.1 Methodology

The results from the structural estimation show that the data from bidding for pop music CDs on eBay is consistent with expectations. Specifically, I found that the rate of entry into the auction decreases as the price grows. I also found that a significant portion (about 16%) of bidders can be labeled as the "learning" type and they update their beliefs by observing the price in the auction. These findings suggest that if the bidders are behaving optimally from the point of view of my theoretical model, then we should be able to observe two predicted features of bidding.

The first feature, entry deterrence by bidding early, arises due to the dependence of entry into the auction on price. Specifically, if it possible to change the structure of this dependence, we should be able to observe a change in the early bidding behavior. For instance, we should expect the rate of entry to change with the size of the market. If the market would become larger (i.e. the number of listed items increases) then the probability of entry into auctions with high prices would decrease, since the chance of winning the item in an auction with no other bidders is higher. As a result, we should expect that early bidding would become more beneficial in these circumstances.

The second feature, learning prevention by more experienced bidders, connects the late bidding behavior with past bidding experience on eBay. Past bidding experience allows the bidders to predict the entry rate into the auction by forecasting the visibility of the auction. This feature has different implications for bidders with different prior experience. Less experienced bidders will want to "experiment" with the auction to improve their information regarding the visibility. They will have an incentive to submit bids only at the end of the auction to conceal their information. For this reason, bidders with different prior experience will respond differently to the exogenous changes in the market, such as an increase in market size. As I argued before, an increase in market size increases the incentive to bid early. If the bidder is experienced, but does not bid early in the market of original size, she may start bidding early if the market expands. As a result, an increase in market size should increase the probability of early bidding by more experienced bidders.

As previously noted, these two features should be observed in bidding behavior given the parameter estimates I obtained from the data and the optimal bidding behavior suggested by the model. However, such effects should be directly observable if the market size increases within the framework of a controlled field experiment. Field experiments on eBay have been used by several applied researchers to study bidding in Internet auctions. Examples include numerous studies analyzing various static models of eBay bidding behavior, such as the effects of seller's reputation (Melnik, Alm, 2002), revenue equivalence (Lucking-Reiley, 1999), and reserve prices (Reiley, 2006). However, to my knowledge, this is the first paper to study the implications of entry and experience on dynamic bidding behavior.

For my experiment, I chose a specific small market on eBay and increased the size of the market by listing additional auctions. I formed a control sample of auctions by observing the unperturbed market. To form a treatment sample of auctions, I listed enough additional auctions to double the size of the market relative to the treatment group. According to my model, if the treatment increases the sensitivity of entry to price (the elasticity of instantaneous demand increases) then we should observe the two features of bidding behavior discussed above. Specifically, the probability of early bidding across all types of bidders should increase, but it should grow more for more experienced bidders.

I chose a market of music CDs on eBay for my experiment. These items are relatively inexpensive and uniform so that it can be assumed that the bidders have private valuations. Furthermore, if the analysis focuses on a specific CD, then the individual market is quite small, meaning that a researcher can significantly influence the supply by listing one or two additional items at a time.

In my experiment, I look at bidding on eBay for the Robbie Williams' CD "The Greatest Hits", which was released in 2004. The supply for this CD is quite stable on eBay with an average of 24 items listed every day. This number of listed items allows me to shift the supply significantly by listing additional items. On the other hand, the number of items is sufficiently large to yield a dataset of about two hundred items over a period of two months. I form the dataset by first observing the unperturbed market. Then I double the size of the market and look at the change in bidding behavior.

I next analyze the frequencies of early jump bidding in the treatment (i.e., post-experiment) and control (pre-experiment) groups. According to the theoretical model, the frequency of early jump bidding should increase with the size of the market. I construct a dummy variable for each auction equal to 1 if there is an early jump bid and zero otherwise. Conditional on the auction characteristics (such as seller's feedback score and the location of the seller) we should expect a positive relationship between this dummy variable and the treatment group dummy.

The second relationship studies the dependence between multiple bidding and the bidder's experience. The model predicts that less experienced bidders should bid more frequently than more experienced bidders. Moreover, if the instantaneous demand is low, then experienced bidders will bid only at the last moment of the auction. If the instantaneous demand grows, then the experienced bidders should start submitting early jump bids to prevent entry by potential rivals. The incentive to prevent entry of other bidders offsets the incentive to prevent learning of other bidders. As a result, in the regression across bidders of the number of early bids submitted by a specific bidder on the treatment dummy, experience of the bidder and the interaction of the treatment dummy with experience indicator (conditional on the auction characteristics), there should be a positive coefficient for the interaction term, and a negative coefficient on experience.

5.2 Control dataset

The control dataset in my study is a set of auctions for the same CDs when the market was not inflated. In total, I collected data for 151 auctions with the earliest auction starting on August 23, 2006 and the last auction ending on October 4, 2006. 15 auctions had items located in North America, 75 in Europe (predominantly the UK), 48 in Asia (mainly China and Taiwan), and 13 in Australia. The shipping cost can be substantial if the item is located very far from the winning bidder. In these circumstances, the shipping costs might play an important role in bidding behavior and I will use regional indicators to capture this effect in the data.

I collected the following auction-specific characteristics: buy-it-now price, starting bid, picture dummy, duration (in days), a dummy equal to 1 if the item is available only in the country of origin (regional auction dummy), shipping cost, the percentage of positive feedback for the seller, seller's feedback score, and a dummy equal to 1 if the seller has a store. One minor issue with the prices and costs is that the auctions in the UK and in Australia use, correspondingly, British pounds and Australian dollars. All prices and shipping costs were converted to US dollars on the basis of daily average FOREX exchange rate on the day of transaction. Since there were no significant fluctuations of the US dollar exchange rate to these currencies, this conversion should not significantly impact the results. Table 4[32] presents summary statistics for the variables in the control sample. Note that the highest buy-it-now prices and the starting bids correspond to rare autographed CDs which could be considered collectibles. These were still included in the dataset because they possess the basic properties of the standard CDs. Most of the items in the sample have a picture. It is worth noting that in most cases this is a picture that is offered by eBay automatically and corresponds to the cover of the album [33]. The average duration of the auction is 7 days, with with a range of 3 to 10 days[34]. The variable "Regional auction" denotes a dummy variable equal to 1 if the item is available only to the bidders in the country where the item is located. There are 19 such auctions in the sample, some of which were located in Australia and the rest in the UK. In the sample, 70% of CDs are new. Most of the used CDs were auctioned from the UK (40 in total) and none of them were auctioned from Asia. The average shipping cost (consisting of the cost of postage and a handling fee) is close to the average sale price of the CD. This makes the total highest prices of the item on eBay close to the lowest prices available at online retail stores (such as Amazon.com). The largest shipping costs are for the largest distance between the seller and the buyer. The shipping cost shown in the table corresponds to the shipping cost to the United States, while the empirical analysis uses the bidder-specific shipping cost. The feedback variable does not seem to be very informative as a majority of sellers in the sample have a 100% positive feedback, probably due to the existing eBay culture of leaving negative feedback only in the case of a very bad experience. A more informative measure is the feedback score, which indicates the experience of the eBay participant and is proportional to the number of transactions made by the user on eBay. Most of the sellers in the sample have very high feedback scores, and, as the mean of the store dummy indicates, almost 60% of the sellers have stores. This suggests that most sellers in the sample are very experienced.

[32]See Appendix

[33]This picture and the list of songs can be generated automatically on eBay website if the seller uses the search option allowing to track the album information from the bar code on the cover of the disc. For this reason,it is possibly an unimportant factor for the bidders.

[34]This can be explained by the fact that there is no additional listing fee on eBay for a 7-day listing (and if the item was not sold in can also be relisted without extra charge), while a 10 day listing costs an extra 40 cents, and a 3 day listing requires a certain level of feedback of the seller.

Turning to the characteristics of the bidders in the control sample, Table 5 summarizes the main parameters of the bidders. First of all, we can see that the feedback scores of the bidders are significantly smaller than the feedback score of the sellers. This implies that on average the sellers in my sample on average have more experience. The percentage of positive feedback, similar to that of the sellers, is equal to 100% for most of the bidders. The last two variables are the dummy "experience" indicators constructed from the observations of past purchase histories of the bidders. The first is equal to one if a specific bidder has won an auction for any music CD on eBay during the last three months. The second variable is equal to 1 if a bidder has won any auction on eBay within the last three months. The CD "experience" variable shows that, in my sample, 45% of the bidders have won a CD on eBay, while 94% of bidders have won an auction on eBay during the last three months. These variables are informative from the point of view of my theoretical model since they indicate that the bidders are familiar with the price behavior on eBay during the auction and, thus, can use the information they have to improve their bidding. The last characteristic reported in the table is the dummy equal to 1 if the auctioned item and the bidder are located in the same region. As one can see, 62% of bidders preferred to bid on items on the same continent.

Let's now consider the characteristics of bidding on eBay auctions in the sample. Table 6 has only 99 observations which reflects that only 99 auctions in the sample ended with a sale.

The table shows that the maximum number of bids in the sample was equal to 10 while the maximum number of bidders was equal to 8. On average, 2.3 bids per bidder were observed in the auctions where bidders were active. The average sale price of the CD was $5.16 which is significantly lower than the price offered by online US retailers (over $20 by Amazon.com). However, this auction price also reflects the fact that some of the CDs in the sample are used. The average bid across auctions in the sample is slightly less than the winning bid, with a maximum value of $25.04.

5.3 Treatment dataset

To construct the treatment dataset I inflated the market size by listing additional items. I used 5 different sellers' accounts and listed from 5 to 10 CDs on each account. This allowed me to double the average number of listed CDs per day in the auction. During the experiment, I sold 60 music CDs. Some of these CDs had to be relisted because there were no bids in the auctions. The actual dates of the experiment were from October 4, 2006 to November 10, 2006. To avoid the "transition effects" in the beginning, and taking into account that I did not have enough CDs at the end to double the market, I only used the data from 20 days of the experiment. I used the data for 156 auctions with the earliest auction starting on October 8, 2006 and the last auction ending on October 28, 2006. In total, 70 had items located in North America (these were predominantly the items listed from my accounts), 56 in Europe, 25 in Asia, and 5 in Australia.

I collected the same variables for the auctions in the treatment dataset as in the control dataset. Table 4 presents the statistics for the auction characteristics in the treatment sample. The characteristics of the items in the treatment sample are quite similar to those in the control sample. The duration of the auctions in the treatment sample is close to 7 days. The percentage of new CDs is higher in the treatment dataset, which can be partly explained by the fact that all the CDs that I listed were new. There are also fewer auctions which restrict the shipping to the country of item location. This could be explained by the fact that I listed my items without such restriction, making the CDs available to overseas bidders.

Let us now look at the characteristics of the bidders in the treatment sample (Table 8). The average experience characteristics, the dummies for CD purchases and all purchases of bidders on eBay as well as their feedback scores, are very similar in the treatment and control samples. The number of bidders in the treatment sample is smaller because the treatment sample covers a shorter period of time (20 days as compared to 31 days in the control sample). The only visible difference between the treatment and the control sample is that the fraction of bidders who are located in the same region as the item is larger in the treatment sample. This can be explained by the fact that, before my experiment, many bidders from the U.S. were bidding for items in Europe, predominantly in the U.K.

Table 9 has only 46 observations of winning bids and average bids, reflecting that only 46 auctions (out of 156) in the sample ended with a sale. There were almost no sales for the auctions located in Asia because the sum of the winning bid and the shipping cost was lower than for items outside Asia. The table shows that the maximum number of bids in the sample was equal to 10 while the maximum number of bidders was equal to 12. On average .692 bids per bidder were observed in the auctions. The average sale price of the CD was $ 3.67, which is lower than the control sample.

5.4 Analysis of early bidding

Using the data from my experiment, I can analyze the correspondence between the predictions of my theoretical model and the results of the experiment. The model predicts that bidding early becomes more attractive to bidders due to the increased size of the market. To capture this effect, I constructed an early jump bid dummy variable. This dummy variable is equal to one if, during the first 80% of the duration of the auction[35], there was a bid which increased the price by a certain threshold. The thresholds of price increases determining the jump bid were chosen to be 10%, 15%, and 20% of the final price in the auction respectively. For jump bid dummy variables constructed with these thresholds, I estimated probit models across the auctions for the combined treatment

[35] I also used other thresholds such as 30% and 50% of the duration but the results were quite similar

and control groups. The dependent variable in these probit models is the jump bid dummy, while the independent variables include the treatment dummy and a set of auction-specific regressors including availability of picture, duration, availability outside the country of item location, seller's feedback and feedback score, and an indicator that the seller has a store. The results of the probit regressions are presented in Table 10.

The estimates in Table 10 suggest that the treatment dummy is significant for all jump bid thresholds. This means that the results of the experiment support the model's prediction that an increase in supply leads to increased early bidding. This effect is quite significant in absolute terms: the probability of early bidding in auctions with doubled supply is 11% higher than in auctions without this treatment.

The coefficients for the characteristics of the auctions demonstrate their effect on early bidding. Specifically, one can see that including a picture of the CD increases early bidding. On the other hand, longer auctions tend to have early jump bids less frequently than short auctions. The increase in the shipping cost decreases overall bidding and, in this way, decreases early bidding as well. Finally, one can see that early bidding is more likely in auctions where the seller has a store. This last result coincides with my interpretation of the visibility of the auction. Specifically, if eBay stores have many cross-listed items, then the auctions of eBay stores are more "visible". As a result, we should expect that the instantaneous demand in the auctions of eBay stores is higher and more price-elastic, meaning that there is a greater incentive to bid early in these auctions. This is confirmed by the data from my field experiment.

The information about individual bidders, specifically the experience dummies and the feedback scores, allows me to analyze the effect of bidding experience on early bidding. My model predicts that bidders have an incentive to deter entry and prevent learning. These incentives have the opposite effect on experienced bidders. The entry deterrence incentive forces bidders to bid early to discourage other bidders from entering the auction. The learning prevention incentive leads bidders to bid late to prevent learning by other bidders. The relative effect of these two incentives should change during my experiment. We know that the increase in supply makes the instantaneous demand more price - elastic. In this case, for more experienced bidders, the incentive to deter entry should dominate the learning prevention incentive. As a result, we should expect an increase in early bidding by experienced bidders. To analyze the effect of increased supply on bidding, I constructed a variable equal to the number of bids that a single bidder submitted in the first 80% of the auction duration. To measure experience, I used the experience dummy variables described above. In addition, as I also used the feedback score of the bidder, which is approximately equal to the total number of transactions the bidder had had on eBay. I ran OLS regressions of the number of early bids on the experience indicators and the interactions between the treatment dummy and the experience indicator (controlling for the bidder location). My model predicts that more experienced bidders should bid more frequently when the market size doubles, implying that we should observe a positive coefficient on the interaction between the treatment dummy and the experience variable.

Table 11[36] contains the coefficient estimates for the models with different bidding experience indicators. In all three cases, I find a positive and significant coefficient on the interaction term, confirming the predictions of the theoretical model. Moreover, we can see that for the bidding experience indicators, the coefficient is negative, confirming that more experienced bidders indeed submit early bids less frequently due to the dominance of the learning prevention incentive. The location dummies are insignificant. The analysis of different specifications of these models (controlling for more characteristics of the auctions) does not significantly alter the effect on the interaction between the treatment dummy and the experience indicator.

6 Conclusion

Although internet auctions are designed as standard second-price auctions, the co-existence of multiple simultaneous auctions for similar items together with the information heterogeneity inherent in Internet auctions significantly affects bidding behavior. In this paper, I show that in such an environment two types of aggressive behavior by bidders can occur: entry deterrence and learning prevention. To deter rival entry, bidders submit early aggressive bids to raise the price in the auction, thereby influencing the rivals to enter elsewhere. To prevent their rivals from learning the true "visibility" of the auction, informed bidders will bid late. The observed bidding behavior arises from the interplay between these two forces.

In this paper, I develop a methodology for modeling auctions in continuous-time, allowing for both endogenous entry and uncertainty, two central characteristics of eBay auctions. I estimate the model using data from auctions for pop-music CDs, validating my theoretical model with empirical data. As an independent test of the model, I conduct a field experiment on eBay. The experiment tests the prediction of my model that, in thicker markets (with the same number of bidders), entry deterrence should be more frequent. I increased the number of listed items in the market for a particular music CD and found a significant increase in early bidding. This result both validates my dynamic model of endogenous entry and uncertainty, and also distinguishes it from alternative models that lack similar predictions.

[36]Feedback score variable was normalized by 10,000 to obtain reasonable coefficient values

References

Andrews, D. (1994): "Empirical Process Methods in Econometrics," in *Handbook of Econometrics*, ed. by R. F. Engle, and D. McFadden, vol. 4. Elsevier Science.

Athey, S., and P. Haile (2002): "Identification of Standard Auction Models," *Econometrica*, 70(6), 2107–2140.

Bajari, P., and L. Benkard (2002): "Estimation with Heterogenous Consumers and Unobserved Product Characteristics: A Hedonic Approach.," *NBER Working Paper*.

Bajari, P., L. Benkard, and J. Levin (2007): "Estimatong Dynamic Models of Imperfect Competition," *Econometrica*, 75(5), 1331–1370.

Bajari, P., and A. Hortacsu (2003): "The winner's curse, reserve prices, and endogenous entry: empirical insights from eBay auctions," *RAND Journal of economics*, 34(2), 329–355.

——— (2004): "Economic Insights from Internet Auctions," *Journal of Economic Literature*, XLII, 457–486.

Baldi, P., P. Frasconi, and P. Smyth (2003): *Modeling the Internet and the Web: Probabilistic Methods and Algorithms*. John Wiley & Sons.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Price in Market Equilibrium," *Econometrica*, 63.

Campo, S., E. Guerre, I. Perrigne, and Q. Vuong (2003): "Semiparametric Estimation of First-Price Auctions with Risk Averse Bidders," *working manuscript, University of Southern California*.

Chen, X., H. Hong, and M. Shum (2005): "Nonparametric Likelihood Ratio Model Selection Tests between Parametric and Moment Condition Model.," *Working paper*.

Chernozhukov, V., and H. Hong (2004): "An MCMC Approach to Classical Estimation," *Journal of Econometrics*, 115(2), 293–346.

Chesher, A. (2003): "Identification of Nonseparable Models," *Econometrica*, 71(5), 1405–1441.

Dunford, N., and J. T. Schwarz (1958): *Linear Operators. Part I: General Theory*. Wiley.

Gallant, R., and G. Tauchen (2002): "Simulated Score Methods and Indirect Inference for Continuous - Time Models," *Working paper, Duke University*.

Gihman, I. I., and A. V. Skorohod (1979): *Controlled Stochastic Processes*. Springer-Verlag.

Gourieroux, C., A. Monfort, and E. Renault (1993): "Indirect Inference.," *Journal of Applied Econometrics*, 8, 85–118.

Grandell, J. (1972): "Statistical inference for doubly stochastic Poisson process," in *Stochastic Point Processes: Statistical Analysis, Theory and Applications*, ed. by P. A. Lewis. Wiley.

Guerre, E., I. Perrigne, and Q. Vuong (2000): "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, 68(3), 525–574.

Hossain, T. (2004): "Learning by Bidding," *Working paper, Hong Kong University of Science and Technology.*

Hotz, V., and R. Miller (1991): "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, 60, 397–429.

Hrbacek, K., and T. Jech (1999): *Introduction to Set Theory*. Marcel Decker, Inc.

Imbens, G., and W. Newey (2002): "Identification and Estimation of Triangular Simulteneous Equations Models Without Additivity," *Working paper.*

Jank, W., and G. Shmueli (2002): "Dynamic Profiling of Online Auctions Using Curve Clustering," *Working paper.*

Kagel, J. H. (1995): "Auctions: A Survey of Experimental Research," in *Handbook of Experimental Economics*. Princeton University Press.

Karr, A. F. (1986): *Point Processes and Their Statistical Inference*. Marcel Decker, Inc.

Kosorok, M. (2008): *Introduction to Empirical Processes and Semiparametric Inference*. Springer Verlag.

Kutoyants, Y. (1998): *Statistical Inference for Spatial Poisson Processes*. Springer-Verlag.

Lipster, R., and A. Shiryaev (2001): *Statistics of Random Processes II*. Springer-Verlag.

Matzkin, R. (2003): "Nonseparable Estimation of Nonadditive Random Functions," *Econometrica*, 71(5), 1339–1375.

McAfee, R. (1993): "Mechanism Design by Competing Sellers," *Econometrica*, 61(6), 1281–1312.

Moen, E. (1997): "Competitive Search Equilibrium," *Journal of Political Economy*, 105(2), 385–411.

Møller, J., and R. P. Waagepetersen (2004): *Statistical Inference and Simulation for Spatial Point Processes*. Chapman & Hall / CRC.

Nekipelov, D. (2007): "Empirical Content of Continuous-Time Principal-Agent Models," *Working paper.*

Ockenfels, A., and A. Roth (2002): "Late Bidding in Second Price Internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction.," *Working Paper, Harvard University.*

Pesendorfer, M., and P. Schmidt-Dengler (2003): "Identificication and Estimation of Dynamic Games," *NBER Working Paper*, (9726).

Peters, M. (1991): "Ex Ante Pricing in Matching Games: Non Steady States," *Econometrica*, 59(5).

Peters, M., and S. Severinov (2004): "Internet Auctions with Many Traders," *Working paper*.

Pissarides, C. A. (1990): *Equilibrium Unemployment Theory*. Oxford: Blackwell.

Rasmusen, E. (2003): "Getting Carried Away in Auctions as Imperfect Value Discovery.," *Working paper, Indiana University*.

Robert, C. P., and G. Casella (2006): *Monte Carlo Statistical Methods*. Springer - Verlag.

Sannikov, Y. (2007): "Games with Imperfectly Observable Actions in Continuous Time.," *Econometrica*, 75(5), 1285–1329.

Silverman, B. W. (1986): *Density Estimation for Statistics and Data Analysis*. Chapman and Hall.

Van Schuppen, J. H. (1977): "Filtering, Prediction and Smoothing for Counting Process Observations, a Martingale Approach," *SIAM Journal of Mathematics*, 32(3), 552–570.

Vuong, Q. (1989): "Likelihood - Ratio Tests for Model Selection and Non - Nested Hypotheses.," *Econometrica*, 57, 307–333.

Yin, P.-L. (2003): "Information dispersion and auction prices," *Working paper, Stanford University*.

A Uniqueness of the Solution to (7).

Consider an operator Γ such that for a bounded function $g(x, y)$ defined on $[0, T]\times[0, \overline{Y}]$ we have $\Gamma\circ g(x, y) = x\,y/T - \int_x^T \lambda(t, y)\,(g(t, y + h(t, y)) - g(t, y))\,dt$. Suppose that functions $\lambda(x, y)$ and $h(x, y)$ are continuous and bounded on $[0, T] \times [0, \overline{Y}]$.

Let $\|u(x, y) - y\| < b$ for $y \in [0, \overline{Y}]$ and $|x - T| < a$ for some positive constants a and b. Let these conditions hold for some functions $u_1(x, y)$ and $u_2(x, y)$.

As the functions forming the operator are continuous and bounded, we can find a constant c determined by their Lipschitz constants[37] such that

$$\|\Gamma\, u_1(x, y) - \Gamma\, u_2(x, y)\| \le c\|u_1 - u_2\|\,\|x - T\|$$

The same procedure can be applied to the operator Γ^2 (as Γ can be applied to the function $\Gamma\circ g(x, y)$) which gives the expression:

$$\|\Gamma^2\, u_1(x, y) - \Gamma^2\, u_2(x, y)\| \le c^2\|u_1 - u_2\|\,\|x - T\|^2/2$$

Iterating this procedure we obtain for Γ^p:

$$\|\Gamma^p\, u_1(x, y) - \Gamma^p\, u_2(x, y)\| \le c^p\|u_1 - u_2\|\,\|x - T\|^p/(p!)$$

Choosing big enough p we will have that $\frac{c^p\,\sup\|x-T\|^p}{p!} < \alpha < 1$ and

$$\|\Gamma^p\, u_1(x, y) - \Gamma^p\, u_2(x, y)\| \le \alpha\|u_1 - u_2\|$$

Therefore Γ^p is a contraction mapping for the considered class of bounded functions. By Dunford and Schwarz (1958) the operator Γ has a unique fixed point if Γ^p is a contraction mapping for some $p \in \mathbb{N}$. This implies that the equation defined as $g = \Gamma \circ g$ has a fixed point which is unique. Such fixed point is the solution to the equation (7). This proves the necessary result.

B Derivation of the Optimal Filter for the Learning Problem

The derivation of the optimal filter here will follow that in the paper Van Schuppen (1977), where the expression for the general filtering of count processes was derived.

Let us look for the solution to the filtering problem to be of the form

$$d\theta_t = \xi_t\,\frac{1}{h(t, x_t, \theta_t)}\,\{dx_t - \psi(t, x_t, \theta_t)\,dt\}\,.$$

Denote $dP(\theta) = dx_t - \psi(t, x_t, \theta_t)\,dt$. Therefore one can write:

$$dx_t = \psi(t, x_t, \theta^*)dt + h(t, x_t, \theta^*)dP(\theta^*),$$

[37] If we impose the restriction that $y + h(x, y) \in [0, \overline{Y}]$ for all $x \in [0, T]$, then we will only need that both functions λ and h are bounded to have this result. Lipschitz property will be required to perform consequent steps.

where θ^* is the true parameter value. Denote $e_t = \theta_t - \theta^*$ and $e_t^\psi = \psi(t,\, x_t,\, \theta_t) - \psi(t,\, x_t,\, \theta^*)$, $\psi^* = \psi(t,\, x_t,\, \theta^*)$, $\psi_t = \psi(t,\, x_t,\, \theta_t)$, $h^* = h(t,\, x_t,\, \theta^*)$, $h_t = h(t,\, x_t,\, \theta_t)$. We will also use the notation $[\cdot,\, \cdot]_t$ for quadratic covariation between the processes.

Using the introduced notations we can write:

$$dP(\theta_t) = -\frac{1}{h_t} e_t^\psi + \frac{h^*}{h_t} dP(\theta^*)$$

Using the rules of stochastic calculus for count processes we can write the following sequence of calculations:

$$e_t P(\theta_t) = e_s P(\theta_s) + \int\limits_s^t e_{\tau-}\, dP(\theta_\tau) + \int\limits_s^t P(\theta_{\tau-})\, de_\tau + \int\limits_s^t d[e,\, P(\theta_t)]_t =$$
$$= e_s P(\theta_s) - \int\limits_s^t \tfrac{e_{\tau-}\, e_\tau^\psi}{h_\tau}\, d\tau + \int\limits_s^t \tfrac{h^*}{h_\tau} e_\tau\, dP(\theta^*) + \int\limits_s^t P(\theta_\tau)\xi_\tau\, dP(\theta_\tau) + \int\limits_s^t \xi_\tau \left\{ \tfrac{\psi_\tau}{h_\tau}\, d\tau + dP(\theta_\tau) \right\}$$

Due to Van Schuppen (1977) $E\{e_t\, P(\theta_t) | \mathfrak{F}_{ns}\} = 0$ for sample-generated σ-algebras $\mathfrak{F}_{ns}$. Also one can conclude that $E\{e_t\, P(\theta_t) | \mathfrak{F}_{ns}\} = E\{\int\limits_s^t \left\{ -\frac{e_\tau e_\tau^\psi}{h_\tau} + \xi_\tau \frac{\psi_\tau}{h_\tau} \right\} d\tau | \mathfrak{F}_{ns}\}$.

From this expression it follows that $\xi_t = \frac{E\{e_\tau e_\tau^\psi | \mathfrak{F}_{ns}\}}{E\{\psi_\tau | \mathfrak{F}_{ns}\}}$.

Note that in the first approximation $e_t^\psi = \frac{\partial \psi(t, x_t, \theta_t)}{\partial t}(\theta_t - \theta^*)$, therefore we get the following formula for the linear approximation of the filtered process: $\xi_t = \frac{\partial \psi_t}{\partial \theta} \frac{\sigma_t(\theta_t)}{\psi_t}$, where $\sigma_t(\theta_t) = E\{(\theta_t - \theta^*)^2 | \mathfrak{F}_{ns}\}$.

Now we will find the dynamics of the conditional variance of the estimate given the third moment σ_3. Consider a stochastic differential for function $\psi(\cdot)$:

$$d\psi(t, x_t, \theta_t) = \left\{ \frac{\partial \psi}{\partial t} + \frac{\partial \psi}{\partial x}\psi - \frac{\partial \psi}{\partial \theta} \frac{\xi_t}{h_t} e^\psi \right\} dt + \{\psi(t, x_t + h_t, \theta_t + \frac{h^*}{h_t}\xi_t) - \psi(t, x_t, \theta_t)\}\, dP(\theta^*)$$

Denote $A = \frac{\partial}{\partial \theta}\left\{ \frac{\partial \psi}{\partial t} + \frac{\partial \psi}{\partial x}\psi \right\}$ and $B = \frac{\partial \psi(t, x_t + h_t, \theta_t}{\partial \theta} - \frac{\partial \psi(t, x_t, \theta_t)}{\partial \theta}$.

Write:

$$de_t^\psi = \left\{ A\, e_t - \frac{\partial \psi}{\partial \theta} \frac{\xi_t}{h_t} e_t^\psi \right\} dt + B\, e_t\, dP(\theta^*).$$

Then for the differential of the quadratic covariation:

$$d[e,\, e^\psi] = B\, e_t \xi_t \frac{\psi_t}{h_t}\, dt + B\, e_t \xi_t\, dP(\theta^*).$$

Therefore:

$$d(e_t\, e_t^\psi) = e_{t-} de_t^\psi + e_{t-}^\psi de_t + d[e,\, e^\psi] = \left\{ e_t (A\, e_t - \tfrac{\partial \psi}{\partial \theta} \tfrac{\xi_t}{h_t} e_t^\psi) - e_t^\psi \xi_t \tfrac{e^\psi}{h_t} + B\, e_t \xi_t \tfrac{\psi_t}{h_t} \right\} dt +$$
$$+ r_t^e\, dP(\theta^*) = \left\{ A\, e_t^2 - \tfrac{\partial \psi}{\partial \theta} \xi_t \tfrac{e_t^\psi e_t}{h_t} - \xi_t \tfrac{(e_t^\psi)^2}{h_t} + B\, e_t \xi_t \tfrac{\psi_t}{h_t} \right\} dt + r_t^e\, dP(\theta^*)$$

Applying the same tools as for deriving the main equation, we obtain that for $r_t = E\{r_t^e | \mathfrak{F}_{ns}\}$ the expression in terms of expectations looks like $r_t = \frac{E\{e_\tau (e_\tau^\psi)^2 | \mathfrak{F}_{ns}\}}{E\{\psi_\tau / h_\tau | \mathfrak{F}_{ns}\}}$.

Using the first-order expansion we obtain that: $r_t = \left(\frac{\partial \psi}{\partial \theta} \right)^2 \frac{h_t\, \sigma_3}{\psi_t}$.

Therefore, the equation for the dynamics of $\sigma_t(\theta_t)$ can be calculated by using the expression for r_t and conditioning of the drift for the process $e_t\, e_t^\psi$ and using the expansion for e^ψ:

$$d\sigma_t(\theta_t) = \left\{ A \left(\frac{\partial \psi}{\partial \theta} \right)^{-1} - \frac{1}{h_t \psi_t} \left(\frac{\partial \psi}{\partial \theta} \right)^2 \right\} \sigma_t^2(\theta_t)\, dt + \frac{\partial \psi}{\partial \theta} \frac{\sigma_3}{\psi_t} \left\{ dx_t - \psi(t,\, x_t,\, \theta_t)\, dt \right\}.$$

By letting $\sigma_3 \equiv 0$, this concludes the derivation of (4).

C Numerical Example.

The problem (7) can be solved by a standard methods used to solve partial differential equations. However, its specific structure allows one to construct a simple Newton-type algorithm without using multiple-grid methods which are usually used to solve partial differential equations because the equation (7) does not contain derivatives with respect to the state variable. Let us consider the following example for the case when $h(t, x, \eta) = \max\{h(t, x), \eta\}$.

Example. Here we will first present the algorithm and then discuss its speed and stability. As many other boundary problems, this problem will be solved backwards, using boundary condition as a starting point. Let $Z_{t\,i}$ be array with the values of Z with index $t = 1, \ldots, T$ referring to time labels and index $i = 1, \ldots, N$ referring to the state variable labels. Let set $I_{t\,i}$ refer to the set of $\inf_{\eta>0} Z(t, x{+}\max\{h(x,\, t),\, \eta\})$ for each time label t and state variable label i. Let the points x_i form the grid for the state variable and τ be the time step.

Step 0.

Initialize $Z_{T\,i} = x_i$. Initialize $I_{T\,i} = x_i + h(x_i,\, T)$.

Step 1.

Make a step in time for every i

$$Z_{t-1\,i} = Z_{t\,i} + \tau\, \lambda_\epsilon(x_i,\, t)\, (I_{t\,i} - Z_{t\,i})$$

Step 2. Find $\min_{j \geq i}\{Z_{t-1\,j}\}$ for every i. If $argmin_{j\geq i}\{Z_{t-1\,j}\} \neq i$, then $I_{t-1\,i} = Z_{t-1\,j}$ otherwise take the grid point x_j closest to the point $x_i + h(x_i)$ and set $I_{t-1\,i} = Z_{t-1\,j}$.

The steps 1, 2 should be repeated until $t = 0$.

From the practical viewpoint, an important implementation problem is the scaling of the parameters of Poisson process. In fact, too rapid growth of the frequency can lead to the explosive behavior of the solution to the analyzed equation. Intuitively it will imply that if the frequency growth over time is higher than certain level then the outcome becomes unpredictable for the buyer.

We illustrate the influence of time on the expected final object price for the process with the following functions. First we consider a time-stationary problem with the functions $\lambda(x,\, t) = \alpha_1 e^{-\beta_1 x^2}$ and $h(x,\, t) = \alpha_2 e^{-\beta_2 x^2}$. The parameters of the functions were $\beta_1 = \beta_2 = 0.01$, $\alpha_2 = 20$ and $\alpha_1 = 0.1$. The figure (1) below shows the surface corresponding to the solution of the boundary problem (7) with these parameters.

The graph shows the values of expected final price given current price for every moment of sale duration under the optimal control from the buyer.

The obtained picture is consistent with the simulation results, obtained earlier. In fact for the moments of time close to the beginning of the sale, the expectation for the final price given optimal control is a function with a set minima for a region of relatively small initial prices while for higher prices there is a one-to-one correspondence between the initial price and the expected current price. Similarly, as predicted by the simulations, the dependence tends to a one-to-one correspondence for most of the price levels when the sale is close to the end.

Adding a time dimension to the problem changes the structure of solution. The characteristics of the Poisson process now look like $\lambda(x,\, t) = \alpha_1 e^{-\beta_1\, x^2}\, e^{\gamma_1\, t}$ and $h(x,\, t) = \alpha_2 e^{-\beta_2\, x^2}\, e^{\gamma_2\, t}$. The added parameters

Figure 1: Computation result for the time-stationary problem.

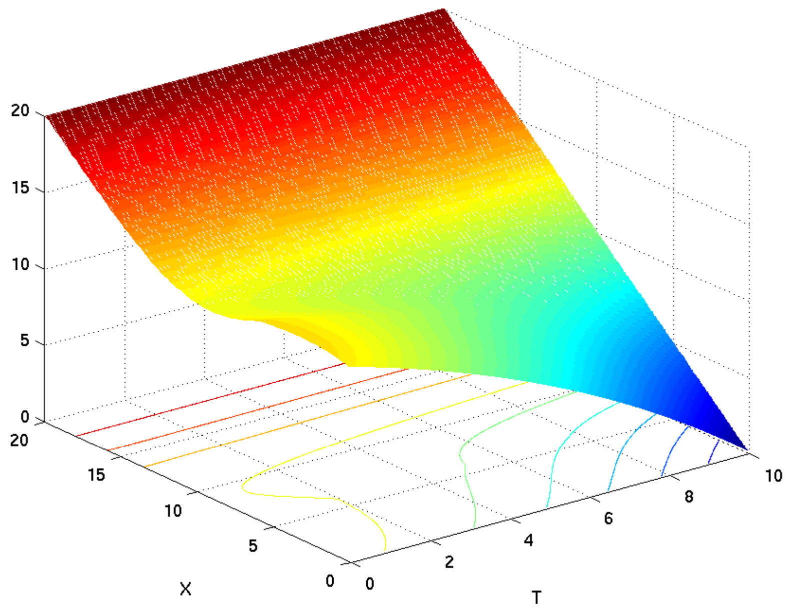

γ_1 and γ_2 were set equal to 0.1. The expected price at the end of the sale process generated under such parameter values is shown on the figure (2).

Figure 2: Computation result for the model with exponentially growing parameters.

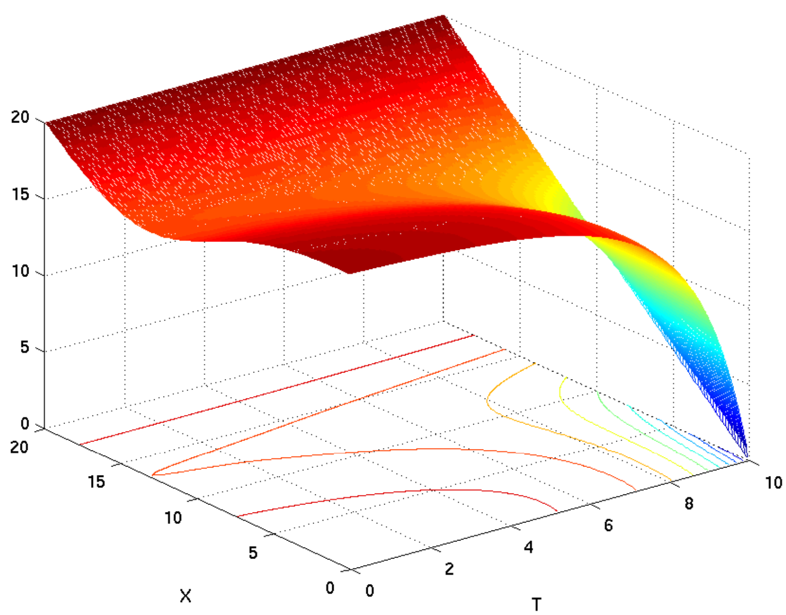

The graph shows that adding an exponentially growing time component to the parameters of the process inflates the expectation of the final price at the beginning of the sale. The comparison of the graphs for time-dependent and time-independent processes demonstrates that for a wide range of initial values the expected final price is higher for the time-dependent process. However, the general structure of the contour sets of the solution remains the same with a minimum and convergence to a one-to-one correspondence for high initial values of the process.

The existence and uniqueness of the solution of problem (7) can be proved in the same way it is proved for the the boundary problems for the partial differential equations. The proof is based on the fact that the boundary problem (7) can be rewritten as an integral equation. So that the solution to the problem (7)is also a solution to operator equation with the integral operator. Due to the result from operator theory [38], a solution to operator equation exists and is unique if some power of the generating operator is a contraction operator. The complete proof is given in the Appendix A.

The same technique can be used to prove the convergence of the chosen algorithm. In this case the operator can be rewritten in terms of the finite differences, while the result remains the same.

This example demonstrated that the problem of the agent can be solved by a simple linear algorithm which does not require value function iterations. Moreover, as the problem does not contain the derivatives of order higher than one, the requirements for the calculation grid are quite mild.

D Asymptotic properties

D.1 Consistency of the density estimate

Our goal is to prove convergence of the smoothed density estimates of the observable variables to their marginal distributions. The observed outcome is an equilibrium path of the dynamic game, represented by a collection of variables $\left\{p^{(k)},\, N^{(k)}\right\}_{k=1}^{n}$, where $p^{(k)}$ is the price path in auction k and $N^{(k)}$ is number of bidders. We will use a kernel estimator to estimate the marginal density of the observed price and the instants of price jumps. We assume that there exists a joint distribution of the number of bidders and the price at the instant t $f\,(p,\, N,\, t)$. Given assumptions **A1 - A5** let us show that the marginal density of observed bids, evaluated as a kernel density estimate of the joint distribution of prices and instants of price jump will be pointwise converging in probability to the marginal distribution, that is:

$$\widehat{f}\,(p,\, t) \stackrel{n\to\infty}{\longrightarrow} \int f(p,\, N,\, t)\, dN,$$

where the last integral should be treated as an integral over a counting measure.

For a single realization (price path) k:

$$\xi_k = \frac{1}{h_p\, h_t} \sum_{i=1}^{I_k} \kappa\left(\frac{p_i^{(k)} - p}{h_p}\right) \kappa\left(\frac{t_i^{(k)} - t}{h_t}\right) = \frac{1}{h_p\, h_t} \int_0^T \kappa\left(\frac{p_\tau^{(k)} - y}{h_y}\right) \kappa\left(\frac{\tau - t}{h_t}\right) dJ(\tau, p_\tau^{(k)}),$$

[38] Dunford and Schwarz (1958)

by definition where $J(\cdot)$ is the Poisson measure generating the stochastic process $\left\{p^{(k)},\, t^{(k)},\, N^{(k)}\right\}$. This integral has finite expectation and variance due to assumptions A1 and A2 and generalized Cauchy-Schwartz inequality.

Note also that

$$f(p,t) = \lim_{\substack{\Delta_t\to 0\\ \Delta_p\to 0}} \frac{1}{\Delta_p\,\Delta_t} E\left\{\int_p^{p+\Delta_p}\int_t^{t+\Delta_t} dJ(\tau,p_\tau)\right\} = \lim_{\substack{\Delta_t\to 0\\ \Delta_p\to 0}} E\left\{\eta(p,t,\Delta_p,\Delta_t)\big|p,\, t\right\} =$$
$$= \lim_{\substack{\Delta_t\to 0\\ \Delta_p\to 0}} \frac{1}{\Delta_p\,\Delta_t} E\left\{\int_0^T \mathbf{1}\{|p_\tau - p| < \Delta_p\}\,\mathbf{1}\{|\tau - t| < \Delta_t\} dJ(\tau,p_\tau)\right\}.$$

Then

$$E\left\{\sup|\xi_k - \eta(p,t,h_p,h_t)|^2\,\middle|p,\, t\right\} = E\left\{\sup\left|\int_0^T \left[\frac{1}{h_p\,h_t}\kappa\left(\frac{p_\tau-p}{h_p}\right)\kappa\left(\frac{\tau-t}{h_t}\right) -\right.\right.\right.$$
$$\left.\left.\left.-\frac{1}{h_p\,h_t}\mathbf{1}\{|p_\tau - p| < h_p\}\,\mathbf{1}\{|\tau - t| < h_t\}\right] dJ(\tau,p_\tau)\right|^2\right\} \leq$$
$$\leq 4\,E\left\{\int_0^T\int_0^\infty \left[\frac{1}{h_p\,h_t}\kappa\left(\frac{p_\tau-p}{h_p}\right)\kappa\left(\frac{\tau-t}{h_t}\right) - \frac{1}{h_p\,h_t}\mathbf{1}\{|p_\tau - p| < h_p\}\,\mathbf{1}\{|\tau - t| < h_t\}\right]^2\right.$$
$$\left.\lambda(\tau,p_\tau)\,dp_\tau\,d\tau\right\} \to 0,$$

as h_t and h_p approach 0 due to Doob's inequality, Cauchy-Schwartz inequality and the fact that the kernel function approaches the step function in $\mathbf{L}^2$ norm. This implies that $E\{\xi_k \,|\, p,t\} \to f(p,t)$ as h_p and $h_t \to 0$.

Due to Gihman and Skorohod (1979) we can express the variance of $\xi(y,t)$ as:

$$var\left(\xi_k \mid p,t\right) = \frac{1}{h_t^2 h_p^2}\int_0^T\int_0^\infty \kappa^2\left(\frac{p_\tau - p}{h_p}\right)\kappa^2\left(\frac{\tau-t}{h_t}\right) E\left\{\lambda(\tau,p_\tau)\right\}\,d\tau\,dp_\tau$$

Defining $\lambda(\cdot) \equiv 0$ when $\tau > T$, we can rewrite:

$$\begin{aligned} var\left(\xi_k \mid p,t\right) &= \frac{1}{h_t^2 h_p^2}\int_0^\infty\int_0^\infty \kappa^2\left(\frac{p_\tau-p}{h_p}\right)\kappa^2\left(\frac{\tau-t}{h_t}\right) E\left\{\lambda(\tau,p_\tau)\right\}\,d\tau\,dp_\tau = \\ &= \left(h_t\,h_p\right)^{-1} f(p,t)\left\{\int_0^\infty \kappa^2\left(\psi\right)\,d\psi\right\}^2 + o\left(\left(h_t\,h_p\right)^{-1}\right) \end{aligned} \tag{14}$$

Note then that the kernel density estimate represents a U-statistic for ξ_k, and thus standard projection results will apply to its representation. Therefore, for $\overline{\xi}(p,t) = \frac{1}{n}\sum_{k=1}^{n}\xi_k$ we can apply the projection theorem to prove consistency. Specifically, for any $\delta > 1/2$ (for instance, by **?**):

$$\overline{\xi}(p,t) - f\left(p,t\right) = \frac{1}{n}\sum_{k=1}^{n}\left[E\{\xi_k \mid p,t\} - f\left(p,t\right)\right] + o_p\left(n^{-1}\left(\log\, n\right)^{\delta}\right).$$

Therefore, $\overline{\xi}(p,t) = f\left(p,t\right) + O_p\left(\left(nh_ph_t\right)^{-1}\right)$ uniformly in p and t, which proves consistency of the estimate.

Applying the expression for the asymptotic distribution for the non-degenerate U-statistic we obtain that:

$$\sqrt{nh_th_p}\left(\overline{\xi}(p,t) - \xi(p,t)\right) \xrightarrow{d} N\left(0,\; f(p,t)\left\{\int_0^\infty \kappa^2\left(\psi\right)\,d\psi\right\}^2\right) \tag{15}$$

D.2 Asymptotic properties of the KLIC estimator

Consider a single realization of the stochastic process $\{p_\tau^{(k)}, t_\tau^{(k)}\}_{\tau\in[0,T]}$. For this realization one can write:

$$\sum_{i=1}^{I_k}\log\left[\frac{\widehat{f}_{\theta_0}(p_i^{(k)},\, t_i^{(k)})}{\widehat{f}_{\theta}(p_i^{(k)},\, t_i^{(k)})}\right] = \int_0^T \log\left[\frac{\widehat{f}_{\theta_0}(p_\tau^{(k)}, \tau)}{\widehat{f}_{\theta}(p_\tau^{(k)}, \tau)}\right]\,dJ\left(\tau,\, p_\tau^{(k)}\right),$$

where $J(\cdot)$ is a Poisson measure, generating stochastic process $\{p_\tau^{(k)}, t_\tau^{(k)}\}_{\tau\in[0,T]}$. Let ξ_1 and ξ_2 are nuisance parameters in non-parametric estimates of $f_{\theta_0}(\cdot)$ and $f_\theta(\cdot)$. Let us denote the latter stochastic integral by $I^{(k)}(\theta,\, \xi_1,\, \xi_2)$.

Note now that the KLIC can be written as

$$I_n\left(\theta,\, \widehat{\xi}_1,\, \widehat{\xi}_2\right) = \frac{1}{n}\sum_{k=1}^{n} I^{(k)}\left(\theta,\, \widehat{\xi}_1,\, \widehat{\xi}_2\right).$$

This is an M-estimator for the finite-dimensional parameter θ (it could also be sub-classified as Z-estimator because in the optimum the KLIC should be equal to zero). Under correct specification the estimator for θ_0 corresponds to the population moment equation:

$$E_{\theta_0}\left\{I^{(k)}\left(\theta^*,\, \xi_1,\, \xi_2\right)\right\} = 0$$

almost everywhere in the compact parameter space Θ. We assume that the model is identified and this equation has a unique solution in the interior of Θ.

This moment condition is associated with the integral over the counting measure. For this reason, for asymptotic derivations we may use Donsker results derived for semiparametric estimators of counting models. To provide the local representation of the KLIC estimator for the sample we follow the reasoning in Andrews (1994). Sample moment condition can be represented as a mean-square expansion:

$$0 = I_n\left(\widehat{\theta},\, \widehat{\xi}_1,\, \widehat{\xi}_2\right) = I_n\left(\widehat{\theta}, \widehat{\xi}_1, \widehat{\xi}_2\right) - I_n\left(\widehat{\theta}, \xi_1, \xi_2\right) + \frac{\partial}{\partial\theta'} I_n\left(\theta^*, \xi_1, \xi_2\right)\left(\widehat{\theta} - \theta_0\right),$$

where the derivative $\frac{\partial}{\partial\theta'} I_n$ should be understood in the mean-square sense and θ^* is a point in the intersection of ϵ-balls around $\widehat{\theta}$ and θ_0.

Consider first $I_n\left(\theta_0, \xi_1, \xi_2\right)$. As an average of n i.i.d. mean zero variables

$$I_n\left(\theta_0, \xi_1, \xi_2\right) = O_p\left(n^{-1}\right).$$

The first part of the expansion belongs to Donsker class, as an integral of the totally bounded function with respect to the Poisson measure, for instance, due to Kosorok (2008). As a result it converges to a Gaussian process. Then we can see that $\widehat{\xi}_1$ and $\widehat{\xi}_1$ under θ_0 are estimates of density from two independent samples from the same distribution. Thus, assuming that Poisson measure does not depend on parameter θ:

$$\frac{\partial}{\partial\theta'} I_n\left(\theta^*,\, \widehat{\xi}_1,\, \widehat{\xi}_2\right) \xrightarrow{p} -Q = -E\left\{\int_0^T \frac{\partial f\left(p_\tau, \tau, \theta_0\right)/\partial\theta}{f\left(p_\tau, \tau, \theta_0\right)}\, dJ(\tau, p_\tau)\right\}$$

A sufficient condition for the finiteness of Q is, according to Gihman and Skorohod (1979), that the integral:

$$\int_0^T \int_0^\infty \left[\frac{\partial f\left(p, \tau, \theta_0\right)/\partial\theta}{f\left(p, \tau, \theta_0\right)}\right]^2 \lambda(\tau, p)\, d\tau dp < \infty, \tag{16}$$

is finite.

The expression for the parameter under consideration can be written as:

$$\begin{aligned}\sqrt{n}\left(\widehat{\theta} - \theta_0\right) &= Q^{-1}\sqrt{n} I_n\left(\theta_0,\, \widehat{\xi}_1,\, \widehat{\xi}_2\right) + o_p(1) = \\ &= Q^{-1}\left[\sqrt{n}\left(I_n\left(\theta_0,\, \widehat{\xi}_1,\, \widehat{\xi}_2\right) - I_n^*\left(\theta_0,\, \widehat{\xi}_1,\, \widehat{\xi}_2\right)\right) + \sqrt{n} I_n^*\left(\theta_0,\, \widehat{\xi}_1,\, \widehat{\xi}_2\right)\right],\end{aligned} \tag{17}$$

where $I_n^*(\theta,\,\xi_1,\,\xi_2)=\frac{1}{n}\sum\limits_{k=1}^{n}E\{J^{(k)}(\theta,\,\xi_1,\,\xi_2)\}$.

As ξ_1 and ξ_2 are nuisance parameters in the non-parametric density estimation, it has been shown in (15)that, similarly to standard kernel estimators as in Silverman (1986) such estimates are pointwise asymptotically normal. Using Fubbini theorem one can argue that

$$\mathrm{Var}_\xi\left[I^{(k)}\left(\theta_0,\,\widehat{\xi}_1,\,\widehat{\xi}_2\right)\right]=E\{\int_0^T\mathrm{Var}_\xi\left[\log\left[\frac{\widehat{f}(y_\tau,\,\tau)}{\widehat{f}_{\theta_0}(y_\tau,\,\tau)}\right]\right]dP(\tau,y_\tau)\}.$$

The existence of this variance is justified by the existence of finite variances of $\log\left(\widehat{f}(\cdot)\right)$ and $\log\left(\widehat{f}_\theta(\cdot)\right)$. For kernel density estimates the existence of such variances is justified by the fact that $f(\cdot)$ and $f_\theta(\cdot)$ are poitwise greater than 0 on their domain. In particular, if the simulated sample is independent from the actual sample of trajectories we can write the asymptotic expression for the variance given that the density estimates are obtained from the kernel smoother with a kernel function $\kappa(\cdot)$ given assumptions A4 and A5 as:

$$var_\xi\left[\sqrt{nh_th_y}I_n^*\left(\theta_0,\,\widehat{\xi}_1,\,\widehat{\xi}_2\right)\right]\overset{n\to\infty}{\longrightarrow}2\left(\int_0^\infty\kappa^2(\psi)\,d\psi\right)^2E\left\{\int_0^T\frac{1}{f(y_\tau,\tau,\theta_0)}\,dP(\tau,y_\tau)\right\}=\Omega_\xi$$

Note that this variance is only driven by the variance in the estimation of joint density but not by the point process *per se*. The reason for this is that if we could have a perfect estimate of the distribution of the process y and the timing of its jumps, then KLIC under true parameter values will be equal to zero and thus the variance of the corresponding stochastic integral would be equal to zero as well. The only source of variance in J_n^* is therefore the error in non-parametric density estimate.

According to Gihman and Skorohod (1979), a sufficient condition for the finiteness of this variance is, similarly to the above, that:

$$\int_0^T\int_0^\infty\frac{\lambda(\tau,p)}{\left(f(p,\tau,\theta_0)\right)^2}\,d\tau\,dp<\infty. \tag{18}$$

We can further write then that:

$$\sqrt{nh_th_p}I_n^*\left(\theta_0,\,\widehat{\xi}_1,\,\widehat{\xi}_2\right)\overset{d}{\longrightarrow}N\left(0,\Omega_\xi\right).$$

Denote now $v_n(\xi)=\sqrt{n}\left(I_n(\theta_0,\,\xi_1,\,\xi_2)-I_n^*(\theta_0,\,\xi_1,\,\xi_2)\right)$. Under true $\xi=\xi^0$ as $J^{(k)}(\cdot)$ are independent while $v_n(\xi^0)\equiv 0$.

Assuming that the integral condition in the statement of the Theorem holds with $\Lambda=\int_0^T\int_0^\infty\frac{\lambda(\tau,p)}{(f(p,\tau,\theta_0))^2}\,d\tau\,dp$, then:

$$E\left\{\sup_{\|\xi-\xi'\|\le\delta}\left|I^{(k)}(\theta_0,\xi')-I^{(k)}(\theta_0,\xi)\right|^2\right\}\le 4\Lambda\delta^2,$$

by Doob's inequality. Then given mean - square convergence of non-parametric estimates of density under $\mathbf{L}^2$ norm, according to Andrews (1994), $v_n(\cdot)$ is stochastically equicontinuous in ξ which implies that $v_n\left(\widehat{\xi}\right)-v_n\left(\xi^0\right)\overset{p}{\longrightarrow}0$.

Given these assumptions one can write the asymptotic distribution of (17) for the kernel density estimates:

$$\sqrt{nh_th_p}\left(\widehat{\theta}-\theta_0\right)\overset{d}{\longrightarrow}N\left(0,Q^{-1}\,\Omega_\xi\,Q^{-1}\right) \tag{19}$$

This establishes the fact that the obtained estimates of the vector of structural parameters θ are asymptotically normal. The finiteness of variance of such estimates is justified by the fact that as the state variable p increases, the frequency of the Poisson measure should decay faster than the density of the distribution of state variable for each moment of time.

D.3 Convergence of the MCMC estimator.

To prove the convergence of the MCMC estimator consider the difference:

$$\widehat{KLIC}(\theta) - \widehat{KLIC}(\theta_0) = \frac{1}{n}\sum_{k=1}^{n}\int_0^T \log\left\{\frac{\widehat{f}(y,t,\theta)}{f(y,t,\theta_0)}\right\} dP(t,y) = -J_n\left(\theta,\, \widehat{\xi}_1, \xi_2\right). \tag{20}$$

The asymptotic behavior of $J(\cdot)$ has been obtained above. The only difference is that above we discussed the variance of $J^{(k)}(\theta,\, \widehat{\xi}_1,\, \widehat{\xi}_2)$ while in (20) we have $J^{(k)}\left(\theta,\, \widehat{\xi}_1, \xi_2\right)$. By the same Fubbini theorem argument as above one can obtain $var_\xi\left[J^{(k)}\left(\theta_0,\, \widehat{\xi}_1,\, \xi_2\right)\right] = E\{\int_0^T var_\xi\left[\log\left[\frac{\widehat{f}_{\theta_0}(y_\tau,\tau)}{f_{\theta_0}(y_\tau,\tau)}\right]\right] dP(\tau, y_\tau)\}$.

Therefore:

$$\begin{aligned} var_\xi\left[\sqrt{nh_th_y}J_n^*\left(\theta_0,\, \widehat{\xi}_1,\, \xi_2\right)\right] \stackrel{n\to\infty}{\longrightarrow} \left(\int_0^\infty \kappa^2(\psi)\,d\psi\right)^2 E\left\{\int_0^T \frac{1}{f(y_\tau,\tau,\theta_0)}\, dP(\tau, y_\tau)\right\} = \\ = \tfrac{1}{2}\Omega_\xi \end{aligned}$$

This in turn implies that the asymptotic behavior of MCMC estimate of θ under fixed non-parametric estimate of the data density is:

$$\sqrt{nh_th_y}\left(\widehat{\theta}_{MCMC} - \theta_0\right) \stackrel{d}{\longrightarrow} N\left(0, \frac{1}{2}Q^{-1}\,\Omega_\xi\,Q^{-1}\right) \tag{21}$$

E Monte Carlo Chains of the Parameters

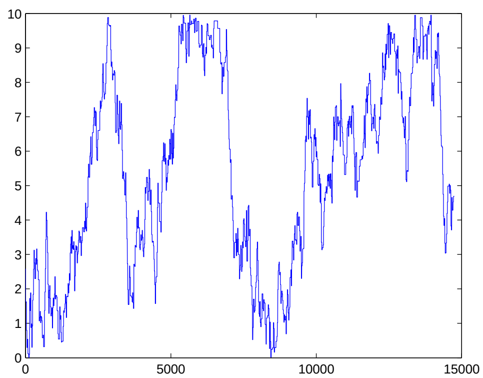

(a)

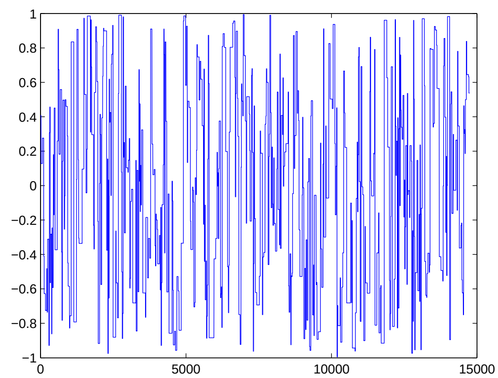

(b)

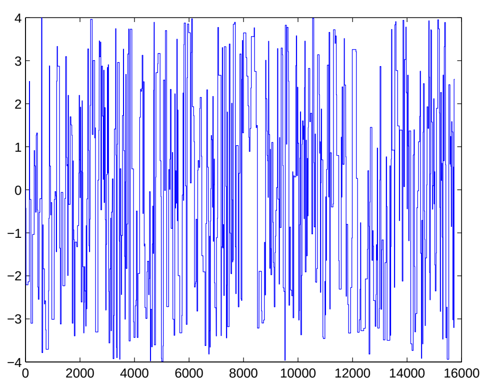

(c)

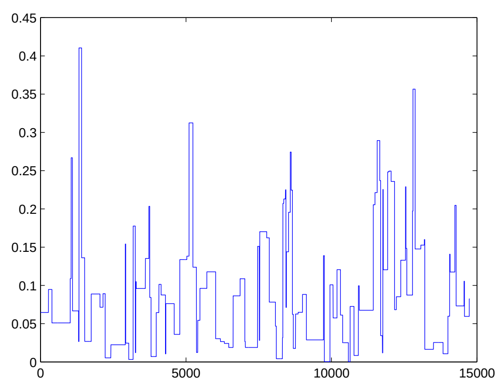

(d)

Figure 3: a) Parameter α; b) Parameter a_2; c) Parameter a_5; d) Probability λ_1

F Tables

Table 1: Summary statistics for auctions for Madonna's CDs

Variable	*N. obs*	*Mean*	*Std. deviation*	*Min*	*Max*
			Auctions		
Buy-it-now price	249	14.7676	46.7879	0.01	577
Starting bid	1132	5.0984	9.2041	0.01	175
Condition	1281	0.5097	0.5001	0	1
Picture	1281	0.9633	0.1880	0	1
Duration	1281	6.5011	1.3270	1	10
Store	1281	0.3575	0.4794	0	1
% seller's positive feedback	1281	98.6860	8.6711	10	100
Seller's feedback score	1281	9131.411	34726.98	0	325614
			Bidding		
time	1591	0.7775	0.3010	0.0015	1
bid	1591	13.8287	28.3903	0.01	699

Table 2: Results of regression analysis.

Dependent variable	*price jump*	*bids per time*	*number of bids*	*number of bidders*
Regressor	(1)	(2)	(3)	(4)
Time	-16.141	-	-	-
	(6.43)**			
Bid	1.152	-	-	-
	(-38.36)**			
Starting price	-	-0.091	-0.004	-0.005
		(2.17)**	**(2.11)***	**-1.86**
Confessions CD	-	10.27	0.513	1.075
		(3.10)**	**(2.80)****	**(2.45)***
Confessions Tour	-	32.263	1.197	1.642
CD+DVD		**(3.22)****	**(3.37)****	**(3.07)****
Duration	-	-6.224	-0.003	-0.033
		(4.15)**	**-0.09**	**-0.54**
Picture	-	2.644	0.154	0.158
		-0.92	**-0.87**	**-0.41**
Shipping cost	-	0.625	0.008	0.096
		-1.1	**-0.31**	**-1.6**
% seller's positive	-	-0.035	-0.003	-0.002
feedback		**-0.34**	**-0.54**	**-0.19**
Seller's feedback	-	1.232	0.072	0.118
score		**(5.17)****	**(6.10)****	**(4.27)****
Constant	-0.752	50.528	0.866	0.878
	-0.37	**(3.52)****	**-1.35**	**-0.81**
Observations	1591	1280	1280	1280
F-statistic	737.6	28.78	7.64	4.75

t-statistics are given in the parenthesis. Coefficients for which standard errors are marked by ** are significant at 1% confidence level marked by * are significant at 5% confidence level.

Table 3: Response to changes in beliefs and the instantaneous demand

Changes in average variance of beliefs $E[\sigma_\theta]$				
Factor	0.5	1.0	1.5	2.0
Fraction of early	4.2%	6.5%	10.1%	15.3%
bids	(2.9)	(2.2)	(8.2)	(7.0)
Average ratio of bid and	15.4%	13.8%	9.2%	9.1%
the final price	(5.4)	(7.3)	(7.4)	(6.9)
Changes in the coefficient α				
Factor	0.5	1.0	1.5	2.0
Fraction of early	4.2%	6.5%	12.1%	18.6%
bids	(3.1)	(2.2)	(4.9)	(7.4)
Average ratio of bid and	14.1%	13.8%	14.4%	21.0%
the final price	(8.7)	(7.3)	(8.9)	(12.0)
Changes in the coefficient a_1				
Factor	0.5	1.0	1.5	2.0
Fraction of early	5.7%	6.5%	8.2%	11.4%
bids	(2.5)	(2.2)	(3.4)	(5.3)
Average ratio of bid and	10.2%	13.8%	15.3%	14.7%
the final price	(6.9)	(7.3)	(9.0)	(9.3)

Standard errors are in the parentheses

Table 4: Characteristics of the auctions for "Greatest hits" CDs in the control group

variable	N obs.	mean	stdev	min	max
Buy it now price (US $)	80	14.97	17.60	1.89	63.72
Starting bid (US $)	117	3.49	5.34	.01	37.24
Picture (yes=1)	151	.894	.308	0	1
Duration	151	7.67	1.89	3	10
Regional auction	151	.125	.332	0	1
Condition (new=1)	151	.701	.458	0	1
Shipping cost (US $)	151	7.28	6.21	0	37.2
Seller's feedback	151	99.68%	.354%	96.8%	100%
Seller's feedback score	151	13028.66	21740.65	0	164546
Store (yes=1)	151	.589	.493	0	1

Table 5: Characteristics of the bidders for ”Greatest hits” CDs in the control group

variable	N obs.	mean	stdev	min	max
Feedback score	156	221.0	561.3	0	4611
Feedback	156	99.68%	1.01	92.6	100
Bough CD last 3 months (yes=1)	156	.455	.499	0	1
Bought from eBay (yes=1)	156	.942	.233	0	1
Same region (yes=1)	156	.621	.485	0	1

Table 6: Characteristics of bidding behavior in the control group

variable	N obs.	mean	stdev	min	max
Number of bids	151	.715	5.364	0	20
Winning bid (US $)	99	5.16	3.80	.01	26.56
Average bid (US $)	99	4.02	3.48	.01	25.04
Number of bidders	151	.656	2.011	0	8

Table 7: Characteristics of the auctions for ”Greatest hits” CDs in the treatment group

variable	N obs.	mean	stdev	min	max
Buy it now price (US $)	45	7.98	7.33	2.98	34.37
Starting bid (US $)	142	5.42	8.14	.18	54.28
Picture (yes=1)	156	.8205	.3849	0	1
Duration	156	7.685	2.234	3	10
Regional auction	156	.02564	.1585	0	1
Condition (new=1)	156	.9102	.2867	0	1
Shipping cost (US $)	156	4.92	2.89	1.59	18.1
Seller's feedback	156	99.70%	.90%	89.7%	100%
Seller's feedback score	156	34448	65604	0	176262
Store (yes=1)	156	.4615	.5001	0	1

Table 8: Characteristics of the bidders for "The Greatest hits" CDs in the treatment group

variable	N obs.	mean	stdev	min	max
Feedback score	81	311.20	790.03	0	4795
Feedback	81	99.58%	1.51%	87%	100%
Bought CD last 3 months (yes=1)	81	.4814	.5027	0	1
Bought from eBay (yes=1)	81	.9135	.2827	0	1
Same region	81	.8765	.3310	0	1

Table 9: Characteristics of bidding behavior in the treatment group

variable	N obs.	mean	stdev	min	max
Number of bids	156	.6923	1.754	0	12
Winning bid (US $)	46	3.67	3.45	.98	14.99
Average bid (US $)	46	3.39	3.27	.98	14.99
Number of bidders	156	.5192	1.127	0	7

Table 10: Effect of market expansion on early bidding: dependent variable - early bidding dummy

	Threshold		
Variable	10%	15%	20%
Treatment dummy	1.673	1.561	1.970
	(.242)***	**(.270)*****	**(.394)****
Picture	1.129	1.122	1.124
	(.513)**	**(.508)****	**(.525)****
Duration	-.118	-.0740	-.0673
	(.0314)***	**(.0309)****	**(.0320)****
Regional	.4335	.225	.373
	(.155)***	**(.156)**	**(.160)****
Condition	-.00941	-.0293	.0808
	(.130)	**(.134)**	**(.138)**
Shipping cost	-.0397	-.0432	-.0467
	(.0191)**	**(.0201)****	**(.0211)****
Seller's feedback score	1.25e-5	1.13e-5	1.28e-5
	(9.26e-6)	**(8.48e-6)**	**(9.54e-6)**
Seller's feedback	.811	.605	.576
	(.222)***	**(.206)*****	**(.221)*****
Store	.117	.484	.673
	(.172)	**(.169)*****	**(.177)*****
constant	-82.75	-62.75	-60.53
	(22.24)***	**(20.69)*****	**(22.11)*****
N. obs.	307	307	307
Pseudo-$\mathbf{R}^2$	.1869	.1849	.1995

Table 11: Multiple bidding as a function of experience

	Model		
Variable	(1)	(2)	(3)
Bought CD last 3 months	-.0710		
	(.0418)*		
Bought CD × Treatment	.637		
	(.173)***		
Bought from eBay		-.0841	
		(.0776)	
Bought from eBay × Treatment		1.136	
		(.266)***	
Feedback score			-.224
			(.123)*
Feedback score × Treatment			3.943
			(2.151)*
Europe dummy	-.00935	-.00826	-.00733
	(.0623)	**(.0631)**	**(.0426)**
Australia dummy	.236	.172	.234
	(.258)	**(.192)**	**(.248)**
Constant	.154	.156	.121
	(.0591)***	**(.108)**	**(.0399)*****
N. obs.	237	237	237
$\mathbf{R}^2$	.0694	.3967	.0793