Research Report

Department of Statistics

No. 2008:1

BAYESIAN INFERENCE AND VARIABLE SELECTION IN STRUCTURAL SECOND PRICE COMMON VALUE AUCTIONS

Bertil Wegmann
Mattias Villani

Department of Statistics, Stockholm University, SE-106 91 Stockholm, Sweden

BAYESIAN INFERENCE AND VARIABLE SELECTION IN STRUCTURAL SECOND PRICE COMMON VALUE AUCTIONS

BERTIL WEGMANN AND MATTIAS VILLANI

Abstract. Over the last decade structural econometric models of auction data have gained wide popularity, especially in the field of Internet auctions. To explore the determinants of bidder and seller behaviour we model eBay auctions as independent second price common value auctions, and assume a similar hierarchical Gaussian valuation structure as in Bajari and Hortacsu (2003). We use an efficient Bayesian variable selection algorithm that simultaneously samples the posterior distribution of the model parameters and does inference on the choice of covariates. We show that the bid function approximation in Wegmann (2007) gives fast, efficient and numerically stable inferences, and that the results are nearly identical to those from using the exact bid function. This allows us to use the approximate bid function and get fast explicit solutions of the equilibrium bid functions that can be inverted analytically, which in turn gives much faster and numerically more stable evaluations of the likelihood function. We apply the methodology to simulated data and to a carefully collected dataset of 1000 coin auctions at eBay. The structural estimates are reasonable, both in sign and magnitude, and the model fits the data well. Finally, we document good out-of-sample predictions from the estimated model.
Keywords: eBay, Markov Chain Monte Carlo, Online auctions.

1. Introduction

Structural econometric models of auction data have become increasingly common in the vast literature of auction theory. In a seminal piece of work, Laffont and Vuong (1996) emphasize that auction models are particularly suited for structural estimation since many datasets are available and well-defined game forms exists. Early examples of structural econometric auction models are Paarsch (1992), who develops an empirical framework to test between a private or a common value element in second price auctions of tree planting contract, and Elyakime et. al. (1997) that estimate the distributions of the buyers' and sellers' private values by nonlinear least squares from first price auction data on standing timber. More recently, economists have developed structural estimation methods in the field of Internet auctions where high-quality datasets can be collected for estimating auction models. Lucking-Reiley (2000) surveys 142 online auction sites and estimates the volumes of transactions at each individual auction site. By far, eBay was estimated to be the world's largest online auction site. At eBay, millions of items are sold every day in thousands of categories from which rich datasets become available to buyers and sellers through completed auction listings.

To explore the determinants of bidder and seller behaviour, Bajari and Hortacsu (2003) examine a dataset of coin auctions from eBay. They argue that these coin auctions possess a common value component and simplify their eBay auction model by assuming independency across auctions. In second price Internet auctions with a fixed end time, Ockenfels and Roth

Wegmann: *Department of Statistics, Stockholm University, SE-106 91 Stockholm. E-mail: bertil.wegmann@stat.su.se.* Villani: *Research Division, Sveriges Riksbank, SE-103 37 Stockholm, Sweden* and *Department of Statistics, Stockholm University. E-mail: mattias.villani@riksbank.se.* The authors thank Camilla Löw for carefully collecting the eBay data.

(2006) show that late bidding, i.e. that bidders tend to place bids very late in the auction, can occur in auctions with both private and common values. Similarly, in the spirit of Wilson (1977), Bajari and Hortacsu (2003) prove that late bidding is a Nash equilibrium in their independent symmetric common value model of eBay auctions. As a consequence, they estimate eBay auctions as independent second price common value auctions. In this environment each bidder is assumed to place only one bid in the very last minute of the auction, so that other bidders have no time to revise their bids. Wilcox (2000) and Schindler (2003) document similar findings for last-minute bidding of Internet auctions with fixed end time.

Bajari and Hortacsu (2003) develop an interesting model for eBay auctions. In their model, the common values are modelled as a heteroscedastic Gaussian regression, and entry into the auctions is stochastically determined by a Poisson regression. The use of auction specific covariates makes it possible to analyze aspects such as the effect of the seller's minimum bid on participation.

Following Bajari and Hortacsu (2003), we model eBay auctions as independent second price common value auctions and assume a similar hierarchical Gaussian valuation structure as in their model. Bajari and Hortacsu (2003) specify rather ad hoc which covariates to include in the model. One of the contributions of our paper is the use of Bayesian inference methodology that lets the data make this decision. We use an efficient Bayesian variable selection algorithm that simultaneously samples the posterior distribution of the model parameters and does inference on the choice of covariates. We conduct a small simulation study to check for the sensitivity with respect to the priors and evaluate the general performance of the variable selection procedure.

By exploiting a linearization property Bajari and Hortacsu (2003) reduce the computational complexity significantly in their model, but the inverse bid function in the very complicated likelihood function still needs to be evaluated by time-consuming numerical methods. A main contribution of our paper is that we show that the approximate bid function in Wegmann (2007) gives nearly identical inferences as when the exact bid function is used. This is an important result since the approximate bid function yields fast explicit solutions of the equilibrium bid functions that can be inverted analytically, which in turn gives much faster and numerically more stable evaluations of the likelihood function.

We apply our Bayesian algorithm on a carefully collected dataset consisting of bids and auction specific covariates from 1000 coin auctions at eBay. We perform reduced form analysis as well as estimation of the full structural model, in both cases using Bayesian inference with variable selection. The estimation results agree with intuition, model evaluations are excellent and even out-of-sample predictions perform well.

Section 2 presents the general model for second price common value auctions and discuss the equivalence to eBay auctions. Two censoring problems of the complicated likelihood function is discussed and the bid function approximation in Wegmann (2007) is stated. Section 3 describes Bayesian inference, specifies the prior distribution and describes the posterior variable selection algorithm in detail. In section 4 we examine the performance of the variable selection on a simulated data and compare the estimation results from the exact and the approximate bid function. Our collected dataset of 1000 coin auctions are documented and analyzed in Section 5. The fit of the estimated model is evaluated by comparing the observed data to data simulated from the model. Our finding is that estimated model gives an excellent fit. Section 5 concludes with an evaluation of the out-of-sample prediction ability of the model on a test sample from 50 additional auctions. Section 6 concludes.

2. Model for second price common value auctions

In an influential article Bajari and Hortacsu (2003, henceforth BH) model an eBay auction as a second price common value auction with stochastic entry. In their model, the seller sets a publicly announced *minimum bid (public reserve price)*, $r \geq 0$, and risk-neutral bidders compete for a single object using the same bidding strategy. The value of the object, v, is unknown and the same for each bidder at the time of bidding, but a prior distribution for v is shared by the bidders. To estimate v, each bidder receives a private signal x from the same distribution $x|v$. Thus, the auction involves symmetric bidders and a symmetric equilibrium why we can focus on a single bidder without loss of generality.

Let $f_v(v)$ denote the probability density function of v, $f_{x|v}(x|v)$ the conditional probability density function of $x|v$, and $F_{x|v}(x|v)$ the conditional cumulative distribution function of $x|v$. The bid function can then be written as[1]

$$b(x,\lambda) = \begin{cases} \frac{\sum_{n=2}^{\infty}(n-1)\cdot p_n(\lambda)\cdot\int_{-\infty}^{\infty} v\cdot f_{x|v}^2(x|v)\cdot F_{x|v}^{n-2}(x|v)\cdot f_v(v)\ dv}{\sum_{n=2}^{\infty}(n-1)\cdot p_n(\lambda)\cdot\int_{-\infty}^{\infty} f_{x|v}^2(x|v)\cdot F_{x|v}^{n-2}(x|v)\cdot f_v(v)\ dv}, & \text{if } x \geq x^\star \\ 0, \text{ if } x < x^\star \text{ or } b(x,\lambda) < 0, \end{cases} \tag{2.1}$$

where $p_n(\lambda)$ is the Poisson probability of $(n-1)$ bidders in the auction with λ as the expected value of the Poisson entry process. Bidders participate with a positive bid if their signal, x, is above the cut-off signal level $x^\star$. Given an arbitrary bidder with signal x, let y be the maximum signal of the other $(n-1)$ bidders. Then, the cut-off signal level is given in implicit form as (Milgrom and Weber, 1982)

$$x^\star(r,\lambda) = \inf_x \left(E_n E[v|X=x, Y<x, n] \geq r\right),$$

which gives

$$r(x^\star,\lambda) = \sum_{n=2}^{\infty} p_n(\lambda) \cdot \frac{\int_{-\infty}^{\infty} v \cdot f_{x|v}(x^\star|v)\cdot F_{x|v}^{n-1}(x^\star|v)\cdot f_v(v)\ dv}{\int_{-\infty}^{\infty} f_{x|v}(x^\star|v)\cdot F_{x|v}^{n-1}(x^\star|v)\cdot f_v(v)\ dv}, \tag{2.2}$$

where r is the minimum bid set by the seller. Hence, the minimum bid $r(x^\star,\lambda)$ is an explicit function of $x^\star$ and λ.

Let v_j denote the common value in auction j, where $j = 1,...,m$, and let x_{ij} denote the signal of the ith bidder in auction j, where $i = 1,...,n_j$, and n_j is the number of bidders that bid zero or place a positive bid in auction j. Similar to BH, we define an hierarchical normal model for valuations as

$$\begin{aligned} v_j &\overset{iid}{\sim} N(\mu_j,\sigma_j^2), \;\; j=1,...,m, \\ x_{ij}|v_j &\overset{iid}{\sim} N(v_j,\kappa\sigma_j^2), \;\; i=1,...,n_j, \\ \mu_j &= z'_{\mu j}\beta_\mu \\ \sigma_j^2 &= \exp\left(z'_{\sigma j}\beta_\sigma\right) \\ \lambda_j &= \exp\left(z'_{\lambda j}\beta_\lambda\right), \end{aligned}$$

where $z_j = (z'_{\mu j}, z'_{\sigma j}, z'_{\lambda j})'$ are auction specific covariates in the regression models for $(\mu_j,\sigma_j^2,\lambda_j)$ in auction j.

The model applies generally to second price common value auctions, but can also be customized to eBay auctions. On eBay, similar single objects are being auctioned simultaneously in parallel auctions. Bidders can thus choose to place bids in different auctions at the same time, and strategic behaviour across auctions is therefore possible. Following BH, however,

[1] See Wegmann (2007) for a complete derivation of the implicit solution in BH.

we assume that the auctions are independent conditional on a set of covariates. For a given auction, Ockenfels and Roth (2006) show that last-minute bidding can occur at equilibrium in auctions both with private and with common values, and BH show that last-minute bidding on eBay is a symmetric Nash equilibrium. BH therefore model their eBay auctions as independent second price common value auctions.

In Section 4 and 5 we assume this simplification of eBay auction models and use Bayesian techniques to estimate the model parameters $\beta = (\kappa, \beta_\mu, \beta_\sigma, \beta_\lambda)$. The likelihood function is complicated because of two censoring problems: *i*) some bidders may draw a signal $x < x^\star$, in which case their bids become unobserved, and *ii*) the highest bid is not observed because of eBay's proxy bidding system, see BH for a detailed description[2].

We will now present the likelihood function for a single auction. Let $b\left(x|\beta, r, z\right)$ be the equilibrium bid function in a given auction conditional on the parameter coefficients β, the minimum bid r, and the covariates z. Since the bid function is strictly increasing in signals, an inverse bid function exists, which we denote as $\phi\left(b|\beta, r, z\right).$ Let $f_b\left(b|\beta, r, z, v\right)$ be the probability density function of the bids conditional on $\left(\beta, r, z, v\right).$ Then, by variable substitution from $x|v$ to b, we get

$$f_b\left(b|\beta, r, z, v\right) = f_{x|v}\left[\phi\left(b|\beta, r, z\right)|v, \kappa, \sigma\right]\phi^{'}\left(b|\beta, r, z\right) \tag{2.3}$$

Let n be the number of bidders who submit a positive bid in a given auction and let $\mathbf{b} = (b_2, b_3, \ldots, b_n)$ be the vector of observed bids[3]. Then, the likelihood function for that auction is given by

$$\begin{aligned} & f_{\mathbf{b}}\left(b_2, b_3, \ldots, b_n|\mu, \sigma, \lambda, \beta, r, z\right) \\ & = \sum_{i=n_j+1}^{\bar{N}} p(i|\lambda)\cdot\int_{-\infty}^{\infty} F_{x|v}\left(x^\star|v, \kappa, \sigma\right)^{i-n}\cdot\left\{1 - f_{x|v}\left[\phi\left(b_2|\beta, r, z\right)|v, \kappa, \sigma\right]\right\}^{I(n\geq 1)} \\ & \quad \times\prod_{i=2}^{n} f_b\left(b_i|\beta, r, z, v\right)\cdot f_v(v|\mu, \sigma)dv, \end{aligned} \tag{2.4}$$

where $p(i|\lambda)$ is the Poisson probability of i bidders in the auction with λ as the expected value, $I\left(n \geq 1\right)$ is an indicator variable for at least one bidder in the auction, and $\bar{N}$ is an upper bound for the total number of potential bidders. For the sake of tractability, let $\bar{N} = 30$, as in BH.

A single evaluation of the likelihood function requires numerical integration to compute $b\left(x|\beta, r, z\right)$ in (2.1), followed by additional numerical work to invert and differentiate $b\left(x|\beta, r, z\right)$. The same applies to compute $x^\star$. This costly procedure needs to be repeated for each of the auctions in the data set. BH cleverly exploits a linearity property of the bid function that confines a large portion of the numerical work to a single auction which is then extrapolated linearly to the other auctions. Nevertheless, the likelihood evaluation suggested by BH is not fast enough to be routinely used for inference. Instead, we take advantage of the linear approximation of the bid function derived in Wegmann (2007). Wegmann's approximation is of an easily interpretable and computationally convenient form, defined as

$$b(x, \lambda) \approx c + \omega\mu + (1-\omega)x,$$

[2]In some auctions the highest proxy bid is less than the second highest proxy bid plus one bid increment. As a result, in those auctions the highest bid is actually observed and equal to the price the highest bidder pays.

[3]The highest bid b_1 is only observed if it is not higher than the second highest bid plus one bid increment, which happens quite seldom. In our collected datamaterial of 1000 eBay auctions the highest bid was only observed in 109 auctions. The likelihood function in (2.4) for such auctions is then modified accordingly.

where $c = -\frac{\sqrt{\kappa}\sigma\gamma\theta(\lambda-2)}{\gamma(\lambda-2)+1+\frac{\kappa}{2}}$, $\omega = \frac{\frac{\kappa}{2}}{\gamma(\lambda-2)+1+\frac{\kappa}{2}}$, $\theta = 1.96$ and $\gamma = 0.1938$. In addition, Wegmann (2007) shows that the cutoff signal can be similarly approximated as

$$x^{\star}(r,\lambda) \approx \frac{r - \sum_{n=2}^{\infty} p_n(\lambda)(c_r + \omega_r \mu)}{\sum_{n=2}^{\infty} p_n(\lambda)(1-\omega_r)},$$

where $c_r = -\frac{\sqrt{\kappa}\sigma\gamma\theta(\lambda-1)}{\gamma(\lambda-1)+\frac{1}{2}+\frac{\kappa}{2}}$, and $\omega_r = \frac{\frac{\kappa}{2}}{\gamma(\lambda-1)+\frac{1}{2}+\frac{\kappa}{2}}$. Moreover, the distribution of the bids in (2.3) simplifies to $b|v \in N[c+\mu,(1-\omega)^2\kappa\sigma^2]$, which speeds up the likelihood evaluation even more. Wegmann (2007) shows that his approximations are very accurate. We will later demonstrate that the approximation can be safely used without distorting inferences.

3. Bayesian inference and variable selection

3.1. **Prior.** Bayesian inference combines the likelihood function in (2.4) with a prior distribution on the unknown model parameters. The numerical algorithms that we use for sampling from the joint posterior distribution (see next section) can be used with any prior. This section proposes a particular prior that can be used with very limited input from the user. Our prior for β_μ and β_σ is motivated by the fact that the common values $v = (v_1, ..., v_n)'$ are modelled as a heteroscedastic regression

$$v = Z_\mu \beta_\mu + \varepsilon,\ \varepsilon_i \sim N(0, \sigma_i^2), \tag{3.1}$$

where $\sigma^2 = (\sigma_1^2, ..., \sigma_n^2)' = \exp(Z_\sigma \beta_\sigma)$. Premultiplying both sides of (3.1) by the $n \times n$ diagonal matrix $D^{1/2} = \mathrm{Diag}[\exp(-z_{\sigma 1}'\beta_\sigma/2), ..., \exp(-z_{\sigma n}'\beta_\sigma/2)]$ we obtain the homoscedastic regression

$$\tilde{v} = \tilde{Z}_\mu \beta_\mu + \tilde{\varepsilon},\ \tilde{\varepsilon}_i \overset{iid}{\sim} N(0,1), \tag{3.2}$$

where $\tilde{v} = D^{1/2}v$, $\tilde{Z}_\mu = D^{1/2}Z_\mu$ and $\tilde{\varepsilon} = D^{1/2}\varepsilon$. We can now specify a g-prior (Zellner, 1986) for β_μ, conditional on β_σ, as

$$\beta_\mu|\beta_\sigma \sim N[0, c_\mu(\tilde{Z}_\mu'\tilde{Z}_\mu)^{-1}] = N[0, c_\mu(Z_\mu' D Z_\mu)^{-1}].$$

where $c_\mu > 0$ is a scaling factor that determines the tightness of the prior. Setting $c_\mu = n$, where n is the number of auctions in the sample, makes the information in the prior equivalent to the information in a single auction (conditional on β_σ), which is a useful becnhmark. The marginal prior for β_σ is also taken to be a g-prior

$$\beta_\sigma \sim N[0, c_\sigma(Z_\sigma' Z_\sigma)^{-1}].$$

Turning to the Poisson entry model, we use the following g-prior for β_λ

$$\beta_\lambda \sim N[0, c_\lambda(Z_\lambda' Z_\lambda)^{-1}].$$

We use an inverse Gamma prior for κ, $\kappa \sim IG(\bar{\kappa}, g)$, where $\bar{\kappa}$ is the prior mean of κ and g are the degrees of freedom. The user thus needs to specify the five hyper parameters c_μ, c_σ, c_λ, $\bar{\kappa}$ and g. We document in Section 5 that the posterior distribution and the variable selection inference are not sensitive to the exact choice of these prior hyperparameters.

Finally, in the variable selection we also need a prior for the covariate subsets. We will here use the simple prior that includes covariates independently of each other with probability π, where π is referred to as the prior inclusion probability.

3.2. **A Metropolis-Hastings algorithm for variable selection.** It is clear from (2.4) that the likelihood function for second price common value auctions is highly non-standard, so the posterior distribution of the model parameters cannot be analyzed by analytical methods. The most commonly used algorithm for simulating from posterior distributions is the Metropolis-Hastings algorithm, which belongs to Markov Chain Monte Carlo (MCMC) family of algorithms, see e.g. Gelman et al. (2004) for an introduction. At a given step of the algorithm, a draw β_p is simulated from the proposal density $f(\beta_p|\beta_c)$, where β_c is the current draw of the parameters (i.e. the most recently accepted draw). The proposal draw β_p is then accepted into the posterior sample with probability

$$a(\beta_c \to \beta_p) = \min\left[1, \frac{p(\beta_p|y)/f(\beta_p|\beta_c)}{p(\beta_c|y)/f(\beta_c|\beta_p)}\right],$$

where $p(\beta|y)$ denotes the posterior density. If β_p is rejected, then β_c is included in the posterior sample. This sampling scheme produces (autocorrelated) draws that converge in distribution to $p(\beta|y)$. The $f(\beta_p|\beta_c)$ can in principle be any density, but should for efficiency reasons be a fairly good approximation to the posterior density. One possibility is to let $f(\beta_p|\beta_c)$ be the multivariate-t density $t(\hat{\beta}, -H^{-1}, h)$, where $\hat{\beta}$ is the posterior mode of $p(\beta|y)$, H is the Hessian matrix at the mode and h is the degrees of freedom. The multivariate-t density is here defined in terms of its mean and covariance matrix. The posterior mode and Hessian matrix can be easily obtained using a standard Newton-Raphson algorithm with BFGS update of the Hessian matrix (Fletcher, 1987).

The model analyzed here contains covariates in the mean and variance functions of the common value distribution, as well as a set of covariates for modelling the Poisson entry process. This opens up the issue of which covariates to include in the model and the importance of the included covariates. Here we follow a Bayesian approach and compute the posterior probability that a covariate belongs to the model, *i.e.* that it has a non-zero regression coefficient. Starting with George and McCulloch (1993) and Smith and Kohn (1996), there has been a number of algorithms that simultaneously draw the regression coefficients from the posterior and does variable selection, all in a single run of the sampler. In particular, Nott and Leonte (2004) propose an efficient algorithm for variable selection in generalized linear models (GLM). Nott and Leonte's algorithm requires that the gradient and hessian matrix are available in closed form (which is the case for GLMs). The algorithm presented below is of similar form, but can be applied to any problem as long as the likelihood and prior can be evaluated numerically.

We present the algorithm for a general setting where β contains all the r model parameters and D denotes the available data. Consider now setting a subset of the elements in β to zero (any other value is also possible). In a regression situation, this is clearly equivalent to selecting a subset of the covariates. Let $\mathcal{J} = (j_1, ..., j_r)$ be a vector of binary indicators such that $j_i = 0$ iff the ith element of β is zero. We can view these indicators as a set of new parameters. We shall here for simplicity assume that the elements of $\mathcal{J}$ are independent a priori with $\Pr(j_i) = \pi$ for all i, so that π is the prior probabilty of including the ith covariate in the model.[4] The following algorithm samples β and $\mathcal{J}$ simultaneously using an extended Metropolis-Hastings algorithm. The algorithm uses the following proposal density

$$f(\beta_p, \mathcal{J}_p|\beta_c, \mathcal{J}_c) = g(\beta_p|\mathcal{J}_p, \beta_c)h(\mathcal{J}_p|\beta_c, \mathcal{J}_c),$$

where β_p and $\mathcal{J}_p$ are the proposed values for β and $\mathcal{J}$, β_c and $\mathcal{J}_c$ are the current values for β and $\mathcal{J}$, h is the proposal distribution for $\mathcal{J}$ and g is the proposal density for β conditional on

[4] A perhaps better way to view this is that the prior on the coefficients is given by a two-component mixture density with one of the components degenerate at zero (Smith and Kohn, 1996).

$\mathcal{J}_p$. The Metropolis-Hasting acceptance probability then becomes

$$a[(\beta_c,\mathcal{J}_c)\to(\beta_p,\mathcal{J}_p)]=\min\left(1,\frac{p(D|\beta_p,\mathcal{J}_p)p(\beta_p|\mathcal{J}_p)p(\mathcal{J}_p)/g(\beta_p|\mathcal{J}_p,\beta_c)h(\mathcal{J}_p|\beta_c,\mathcal{J}_c)}{p(D|\beta_c,\mathcal{J}_c)p(\beta_c|\mathcal{J}_c)p(\mathcal{J}_c)/g(\beta_c|\mathcal{J}_c,\beta_p)h(\mathcal{J}_c|\beta_p,\mathcal{J}_p)}\right),$$

where $p(D|\beta_p,\mathcal{J}_p)$ is the likelihood of the observed data conditional on β with zeros given by $\mathcal{J}_p$, $p(\beta|\mathcal{J})$ is the prior of the non-zero elements of β, $p(\mathcal{J})$ is the prior probability of $\mathcal{J}$.

We will propose $\mathcal{J}_p$ using the two updating steps: i) randomly picking a small subset of $\mathcal{J}_p$ and then always propose a change of the selected indicators (metropolized move), and ii) randomly picking a pair of covariates (one currently in the model and the other currently not in the model) and propose a switch of their corresponding indicators (switch move). Our experience is that these two simple updating rules work well in practise. Note also that the acceptance probability for these updates simplifies to

$$a[(\beta_c,\mathcal{J}_c)\to(\beta_p,\mathcal{J}_p)]=\min\left(1,\frac{p(D|\beta_p,\mathcal{J}_p)p(\beta_p|\mathcal{J}_p)p(\mathcal{J}_p)/g(\beta_p|\mathcal{J}_p,\beta_c)}{p(D|\beta_c,\mathcal{J}_c)p(\beta_c|\mathcal{J}_c)p(\mathcal{J}_c)/g(\beta_c|\mathcal{J}_c,\beta_p)}\right). \tag{3.3}$$

More sophisticated ways to propose can be easily implemented, e.g. the adaptive scheme in Nott and Kohn (2005), where the history of $\mathcal{J}$-draws is used to adaptively build up a proposal for each indicator.

The proposal density $g(\beta_p|\mathcal{J}_p,\beta_c)$ is obtained as follows. First we approximate the posterior with all covariates included with the $t(\hat{\beta},-H^{-1},h)$ density as described above for the case without variable selection. We then propose $\beta_p|\mathcal{J}_p$ from this multivariate t distribution conditional on the zero restrictions dictated by $\mathcal{J}_p$. Assume for notational simplicity that the elements in β has been rearranged so that $\beta=(\beta_0',\beta_p')'$ where β_0 are the p_0 zero-restricted elements of β under $\mathcal{J}_p$, and β_p are the non-zero parameters. Decompose $\hat{\beta}$ and $P=-H^{-1}$ conformably with the decomposition of β as

$$\begin{aligned}\hat{\beta} &= (\hat{\beta}_0,\hat{\beta}_p)' \\ \hat{P} &= \begin{pmatrix} P_{00} & P_{0p} \\ P_{p0} & P_{pp}\end{pmatrix}.\end{aligned}$$

Using result from the conditional distributions of subsets of multivariate-t variables (see e.g. Bauwens, Lubrano and Richard (1999, Theorem A.16)), we can now propose β_p conditional on $\beta_0=0$ from

$$\beta_p|(\beta_0=0)\sim t[\hat{\beta}_p+P_{pp}^{-1}P_{p0}\hat{\beta}_0,P_{pp},c(\beta_0),v+p_0].$$

$$c(\beta_0)=1+\hat{\beta}_0'(\hat{P}_{00}-\hat{P}_{0p}\hat{P}_{pp}^{-1}\hat{P}_{p0})\hat{\beta}_0,$$

using the parametrization of the t distribution in Bauwens et al. (1999). A similar algorithm has recently been suggested by Giordani and Kohn (2007) in their adaptive sampling framework. They propose to use a mixture of multivariate normals as a proposal density rather than a multivariate t. They document good performance on simulated data.

We use the mean acceptance probability and the inefficiency factor (IF) to measure the efficiency of the Metropolis-Hastings algorithm. The inefficiency factor is defined as $1+2\sum_{k=1}^{K}\rho_k$, where ρ_k is the autocorrelation at the kth lag in the MCMC chain for a given parameter and K is an upper limit of the lag length such that $\rho_k\approx 0$ for all $k>K$. The inefficiency factor approximates the ratio of the numerical variance of the posterior mean from the MCMC chain to that from hypothetical iid draws. Put differently, the IF measures the number of draws needed to obtain the equivalent of a single independent draw. IFs close to unity is therefore an indication of a very efficient algorithm.

Table 1. Comparing the inferences from the exact and approx. bid functions

Parameter	Mean		Std Dev		Incl prob	
	Exact	**Approx**	**Exact**	**Approx**	**Exact**	**Approx**
$\ln \kappa$	−1.1312	−1.1215	0.1151	0.1117	1	1
β_1	0.9955	0.9946	0.0048	0.0047	1	1
β_2	−0.1336	−0.1360	0.0316	0.0352	0.9108	0.8921
β_3	0.9531	0.9485	0.0128	0.0126	1	1
β_4	−0.3001	−0.3039	0.0718	0.0686	0.9082	0.9113
β_5	1.2533	1.2407	0.1009	0.1047	1	1
β_6	0.2213	0.2242	0.0210	0.0212	1	1
β_7	−0.0082	−0.0051	0.0332	0.0288	0.1189	0.0858
β_8	−2.3584	−2.3689	0.1655	0.1658	1	1

Note: $c = c_\mu = c_\sigma = c_\lambda = n$, $\bar{\kappa} = 0.25$, $g = 4$, and $\pi = 0.2$.

4. Simulations

In this section we evaluate the performance of the variable selection procedure. We are particularly interested in exploring the sensitivity of the posterior inclusion probabilities with respect to changes in prior hyperparameters $c_\mu, c_\sigma, c_\lambda, \bar{\kappa}$, and g. We set the prior inclusion probability to $\pi = 0.2$ and $\bar{\kappa} = 0.25$ throughout the simulations. See Section 5 for a sensitivity analysis with respect to π. We simulated 50 full data sets, each with 407 auctions, from the eBay model in BH, using their posterior mean estimates as parameter values. The covariates in the model were simulated independently to mimic the summary statistics in Table 1 and 2 in BH. The eBay auction model in BH for auction j, without a secret reserve price, can be written as[5]

$$\mu_j = \beta_1 \mathsf{BookVal}_j + \beta_2 \mathsf{Blemish}_j \cdot \mathsf{BookVal}_j - 2.18$$

$$\sigma_j = \beta_3 \mathsf{BookVal}_j + \beta_4 \mathsf{Blemish}_j \cdot \mathsf{BookVal}_j$$

$$log\lambda_j = \beta_5 + \beta_6 \mathsf{LogBookVal}_j + \beta_7 \mathsf{Negative}_j + \beta_8 \frac{\mathsf{MinBid}_j}{\mathsf{BookVal}_j}.$$

To speed up the computations we use the approximate bid solution in (2.3). Wegmann (2007) show that this approximation is very accurate by comparing the exact and approximate bid function for several different sets of parameter values. In addition, we have simulated several data sets to compare the inferences based on the approximate and exact bid function, and the two always gave very similar results. As an example, Table 1 displays the results for a randomly chosen dataset. The posterior inferences are nearly identical. We will therefore use the approximate bid solution throughout the paper.

To check that the posterior variable selection procedure assigns small posterior inclusion probabilities to insignificant covariates in the model, we include one superfluous covariate for each regression in $(\mu_j, \sigma_j, \lambda_j)$, where each covariate is drawn independently from the standard normal distribution. In Figure 1, inclusion probabilities for each of the covariates in 50 simulated datasets are shown as Boxplots for the priors with $g = 10$ and different values on $c = c_m = c_s = c_l$. Other priors than $g = 10$ and $c_m = c_s = c_l$ were also used and gave very similar results. As we can see in Figure 1, the inclusion probabilities for the most significant variables are all close to one and differ very little across the different priors. The inclusion probabilities for parameter β_4 and β_7 are low even if they were included with non-zero coefficients in the data-generating process. This is not surprising, as they are quite close to zero

[5]See Appendix A for a detailed description of covariates.

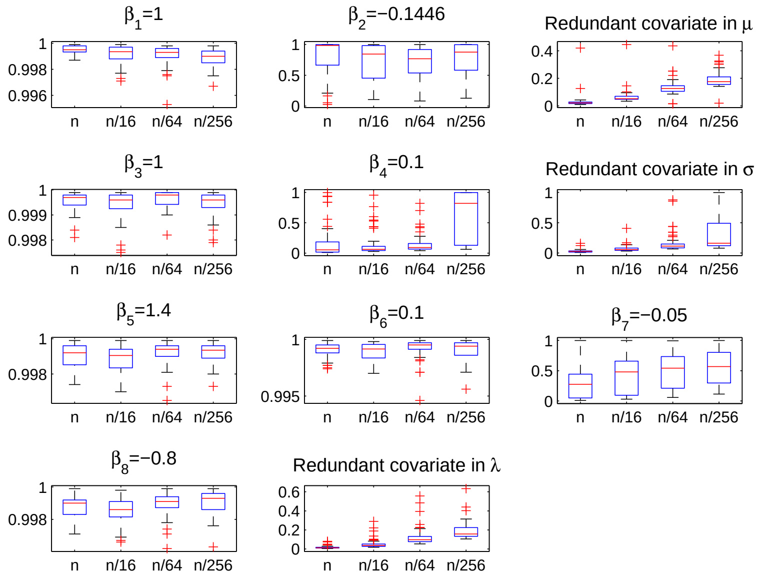

FIGURE 1. Boxplots of posterior inclusion probabilities across simulated data sets for four different values of the prior hyperparameter $c = c_\mu = c_\sigma = c_\lambda$. $\bar{\kappa} = 0.25$, $g = 4$ and $\pi = 0.2$ was used in all simulations.

in the generating model. As a consequence, the posterior variable selection process consider them more or less as unnecessary covariates and removes them from the model with large probability. Another finding from Figure 1 is the phenomena of higher inclusion probabilities for insignificant variables as the prior information becomes more informative, which corresponds to smaller scale factors c. Since the prior information is tighter around zero, the cost of including a covariate with a small coefficient is smaller which in turn increases the posterior inclusion probability.

It turns out that a simple way to characterize the posterior inclusion probabilities is to use *Bayesian t-ratios*. We define these ratios as

$$t_{Bayesian} = \frac{\hat{\theta}}{s(\hat{\theta})},$$

where $\hat{\theta}$ and $s(\hat{\theta})$ are the posterior mode and the approximate/asymptotic posterior deviation, respectively, from the optimization of the posterior function. In Figure 2, the inclusion probabilities increase sharply around a threshold whose value depends on the prior hyperparameters (the dashed line in Figure 2 marks out the 1.96 threshold used in classical t-tests at a 5% significance level). As c decreases (tighter prior) the thresholds move to the left and the curves flatten out. Posterior distributions with more informative prior distributions requires smaller t-ratio values for a certain inclusion probability.

Finally, most of the Metropolis-Hastings runs gave an acceptance probability in the range 0.6-0.8, only a few below 0.1, and none below 0.25 for the prior $c_m = c_s = c_l = n$.

5. Application to eBay auction data

At eBay, millions of items are listed into thousands of categories and subcategories. To search for an item, eBay provides a search engine which allows buyers to specify auction

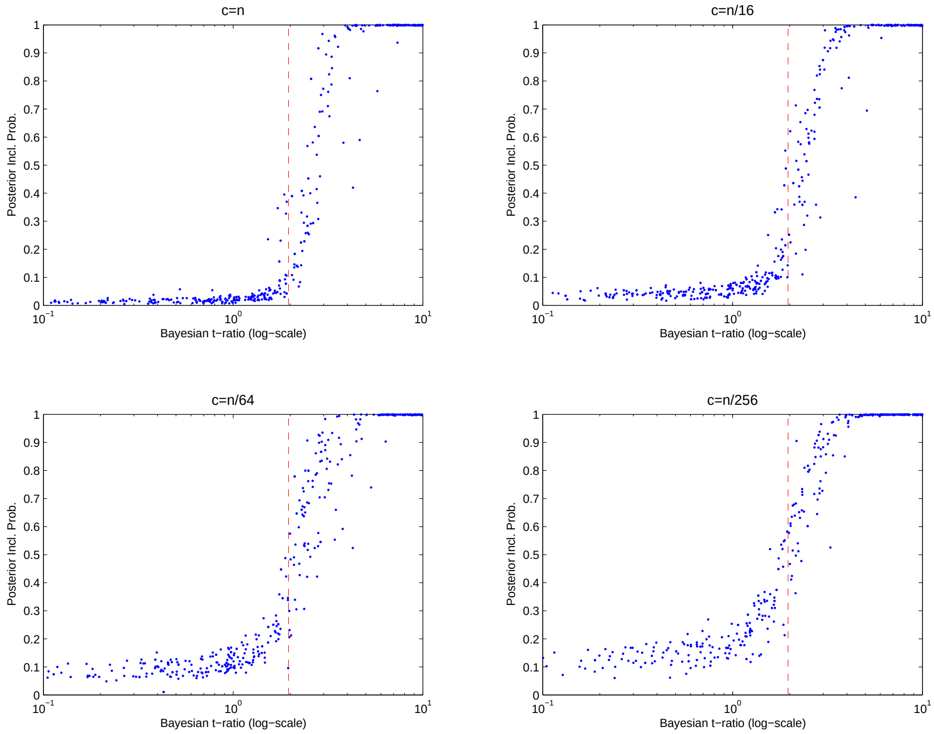

FIGURE 2. The relation between Bayesian t-ratios and the posterior inclusion probabilities for the simulated datasets. Each subplot corresponds to a different value of the prior hyperparameter $c = c_m = c_s = c_l$.

characteristics, e.g. by keywords, location, price range, or ending time. In addition, the search engine is also a useful tool to review completed auctions for recent transactions. In either case, the listings typically contain detailed description of the item, the quantity that is being sold, seller characteristics, reservation price policy, the highest bidder, and current bids as reported on "'the bid history page"'[6]. Hence, high-quality datasets with detailed auction characteristics and all but the highest bid can be collected at eBay and be used for estimating auction models.

In this section, we document and analyze our collected eBay auction data and examine the performance of the variable selection algorithm in Section 3.

5.1. **Description of the data.** The data used for estimation was collected between November 7 to December 19, 2007 and December 27, 2007 to January 22, 2008 from completed eBay auctions, and contains auction characteristics and bids from 1000 auctions of $U.S$ proof sets[7].

[6]Not all bids are observed. The bid of the highest bidder is only observed if his bid is below the next highest bid plus one bid increment.

[7]U.S. proof sets can be defined as the specially packaged set of Proof coins issued annually and sold by the U.S. Mint. The proof set is characterized by the different denominations of proof coins struck by a nation in a particular year.

We exclude multi-unit objects, auctions with a *Buy It Now* option [8], and Dutch auctions. We also collected data from 50 additional auctions between January, 23 to January, 29 2008 that are used in Section 5.4 to evaluate the out-of-sample predictions of auction prices. The bids recorded on eBay's Bid History Page are supposed to correspond to the final bids for each bidder. A careful inspection of the bids reveals, however, that some bids are only a tiny fraction of the object's book value, and cannot realistically represent serious final bids. We will therefore exclude the most extreme bids in our main analysis. A bid b is excluded if $b \leq \delta \cdot \min(\mathsf{BookValue},\mathsf{Price})$, where $\delta = 0.25$ in the benchmark estimations (this excludes 107 bids from the 3742 bids in the sample). We also present results for the case where all bids are used in the estimation ($\delta = 0$) and $\delta = 0.5$ (431 bids removed). As in BH, we believe that these coin auctions on eBay possess a common value component and for simplicity we model the eBay auctions as independent second price common value auctions[9].

In Table 2, summary statistics are shown for all collected covariates[10] and bids in all auctions, see Appendix A for a detailed description.

Table 2. Summary statistics for the eBay data

Variable	**Mean**	**Std Dev**	**Min**	**Max**	μ	σ	λ
Book Value($)[10]	32.30	37.26	7.50	399.50	x	x	x
Price/Book Value, no Blemish	0.76	0.30	0.16	3.63			
Price/Book Value, Blemish	0.61	0.24	0.05	1.20			
Minor Blemish	0.09				x	x	x
Major Blemish	0.03				x	x	x
Number of Bidders	3.74	2.70	0	15.00			
Minimum Bid/Book Value	0.35	0.34	0	1.34			x
Seller's Feedback Score	2427	2883	0	12568			
PowerSeller	0.53				x	x	x
ID Seller	0.06				x	x	x
Unopen	0.10				x	x	x
%Sold	0.86						

The last three columns specifies the covariates used in the models for μ, σ, and λ in the next sections.

5.2. **Preliminary results from reduced form regressions.** To specify suitable functional forms for the regressions in (μ, σ, λ), we created scatterplots to illustrate the linear relationship between the dependent variables and functions of non-dummy covariates. Since we do not observe μ, σ, or λ in the data we derive simple estimates for each auction as follows.

The approximate bid solution in (2.3) as a function of (μ, σ) can be written as

$$b(x) = h\sigma + \omega\mu + (1-\omega)x,$$

where

$$h = -\frac{\sqrt{\kappa}\gamma\theta(\lambda-2)}{\gamma(\lambda-2)+1+\frac{\kappa}{2}}, \text{ and } \omega = \frac{\frac{\kappa}{2}}{\gamma(\lambda-2)+1+\frac{\kappa}{2}}.$$

[8] In some auctions buyers have the option to purchase the item directly at a certain price. If so, the auction is ended and the buyer pays an amount equal to the Buy It Now price.

[9] See BH for a complete treatment of the equivalence between the eBay auction format and a second price common value auction.

[10] $\mathsf{BookValue}$ is replaced with $\ln(\mathsf{BookValue})$ for the regressions in σ and λ.

Table 3. Posterior inference for the reduced form regressions

Coeff	Covariate	Mean	Stdev	t-ratio	Incl Prob	IF
κ		3.042	–	–	–	–
μ	Const	24.191	0.293	88.377	1.000	5.1981
	BookStd	0.672	0.008	78.600	1.000	2.6350
	Book*Power	0.031	0.011	2.183	0.394	1.6808
	Book*ID	−0.008	0.024	−0.763	0.007	–
	Book*Unopen	0.235	0.016	14.187	1.000	3.1463
	Book*MinBlem	0.040	0.017	2.084	0.166	–
	Book*MajBlem	−0.243	0.021	−12.129	1.000	3.3420
	Book*LargNeg	−0.017	0.013	−0.666	0.019	–
	σ_ε	5.720	0.161	–	–	2.544
	R^2	94.5%				
σ	Const	2.247	0.059	29.568	1.000	2.066
	LBookStd	2.054	0.075	26.881	1.000	1.682
	LBook*Power	−0.025	0.035	−0.246	0.009	–
	LBook*ID	0.103	0.076	1.159	0.017	–
	LBook*Unopen	0.272	0.044	5.812	1.000	1.786
	LBook*MinBlem	0.038	0.058	0.405	0.015	–
	LBook*MajBlem	−0.171	0.080	−1.998	0.101	–
	LBook*LargNeg	0.096	0.050	2.009	0.073	–
	σ_ε	1.3092	0.036	–	–	1.993
	R^2	57.2%				
λ	Const	1.059	0.023	34.966	1.000	2.124
	Power	−0.028	0.038	−0.559	0.021	–
	ID	−0.402	0.092	−4.275	1.000	1.618
	Unopen	0.439	0.049	8.800	1.000	1.601
	MinBlem	−0.045	0.058	−0.871	0.048	–
	MajBlem	−0.232	0.091	−2.420	0.664	1.615
	LargeNeg	0.078	0.054	1.260	0.077	–
	LBookStd	−0.112	0.029	−4.173	0.992	1.858
	MinBidStd	−1.901	0.070	−26.679	1.000	2.024

Upper: linear regression with $\hat{\mu}$ as dependent variable.
Middle: linear regression with $\ln\hat{\sigma}^2$ as dependent variable.
Lower: Poisson regression with the number of bidders as dependent variable.
$\hat{\mu}$ and $\hat{\sigma}^2$ are the MLE conditional on the posterior mean of κ from the structural model.

Since $x|v \overset{iid}{\sim} N\left[\mu, \kappa\sigma^2\right]$, we have

$$b(x|v) \overset{iid}{\sim} N\left[\mu + h\sigma, (1-\omega)^2\kappa\sigma^2\right].$$

The likelihood function for μ and σ^2 is then obtained by integrating out v.

It is now straightforward to verify that the maximum likelihood estimates of μ and σ^2, conditional on κ, are given by

$$(\hat{\mu}, \hat{\sigma}^2) = \left(\bar{b} - \frac{hs_b}{(1-\omega)\sqrt{\kappa}}, \frac{s_b^2}{(1-\omega)^2\kappa}\right),$$

where we make the usual correction of the sample variance, defined as $s_b^2 = \frac{\sum_{i=1}^n (b_i - \bar{b})^2}{n-1}$.

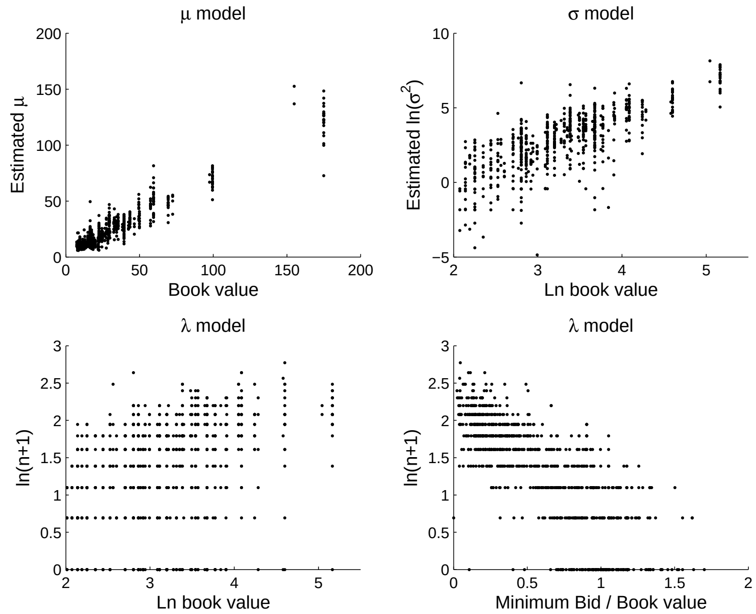

FIGURE 3. Scatterplots of $\hat{\mu}, \ln \hat{\sigma}^2$, and $\ln(n_j + 1)$ against the suitable transformations of the covariate BookValue, and $\ln(n_j + 1)$ against the non-dummy covariates in the entry model.

Note that these estimates can only be computed for auctions with at least two bidders (if the highest bid is observed, otherwise three) why the results should be treated with care. Figure 3 displays scatterplots of $\hat{\mu}$ and $\ln \hat{\sigma}^2$ against the suitable transformations of the covariate BookValue, and the relationship seems to be appropriately linear. Figure 3 also plots $\ln(n_j+1)$, where n_j is the number of bidders with non-zero bids in auction j, against the non-dummy covariates in the entry model. The relationships are again very close to linear.

As a precursor to the analysis of the full structural model in Section 5.3, we perform separate reduced form regressions with $\hat{\mu}$ and $\ln \hat{\sigma}^2$ as dependent variables, respectively. The estimates of μ and σ^2 are conditioned on $\kappa = 3.042$, the posterior mean in the benchmark structural model estimated in the next section. Other values for κ, e.g. the estimate in BH of $\kappa = 0.25$, did not have any major effects on the results. Table 3 (upper and middle parts) presents the results from a full Bayesian analysis with variable selection using the algorithm in Section 3.2. Note that the non-dummy covariates have been standardized (demeaned) and a constant is added to each model. As an example, BookStd is the covariate BookValue minus its mean. This reduces the correlation among the coefficients, which is beneficial for numerical stability. The same standardizations are also used in the structural model in Section 5.4.

The covariates BookStd and interactions of book values and Unopen and MajBlem clearly belong in the model for $\hat{\mu}$, and their estimated coefficients are large and of the expected sign. Similar results were also obtained from a Bayesian Tobit regression with auction prices as the dependent variable[11], see Table 4. The strongest predictors in the regression for $\ln \hat{\sigma}^2$ are LBookStd (standardized ln(BookValue)), and the interaction with Unopen.

[11]In auctions where the objects remain unsold, Tobit regressions use the information that the price is smaller than the auctions' minimum bids.

Table 4. Posterior inference for the Tobit price regression

Covariate	Mean	Stdev	t-ratio	Incl Prob	IF
Const	27.177	0.230	95.824	1.000	1.789
BookStd	0.757	0.007	86.353	1.000	1.796
Book*Power	−0.005	0.009	−0.557	0.039	–
Book*ID	0.032	0.021	1.676	0.118	–
Book*Unopen	0.321	0.018	17.322	1.000	2.063
Book*MinBlem	0.003	0.014	0.291	0.035	–
Book*MajBlem	−0.151	0.023	−6.673	1.000	1.601
Book*LargNeg	0.006	0.013	0.398	0.032	–
σ_ε	6.664	0.158	–	–	1.938

Tobit regressions with auction prices possibly censored at the minimum bid as dependent variable.

Finally, the last part of Table 3 gives the result for a Poisson regression with the number of bidders with a non-zero bid as the dependent variable. Here the covariates MinBidStd (standardized MinBid/BookValue), LBookStd, and the interactions LBook*ID, LBook*Unopen, LBook*MajBlem come out as the most relevant covariates. All significant coefficients are of the expected sign, except the negative coefficient on LBookStd and possibly also the result that eBay verified sellers (ID) seem to attract less bids, all else equal. The negative coefficient on LBookStd in the Poisson regression is somewhat unexpected, especially given positive relation in the scatterplot in Figure 3. These conflicting pieces of evidence can be reconciled, however, by noting that LBookStd is rather strongly negatively related to MinBidStd, the main driver of the entry process.

The Metropolis-Hastings acceptance probabilities for all reduced form regressions were very high (70-90%), the inefficiency factors are close to unity, and convergence was excellent.

5.3. **Estimation results.** Table 5 reports our parameter estimates for the structural model, specified in Table 2 in Section 5.1. With a few exceptions, the parameter estimates in Table 5 agree fairly well with the reduced form estimates. Just as in Section 5.2 the estimated coefficients are mostly of the expected sign and have reasonable magnitudes. The (significant) negative coefficient on LBookStd in the model for λ is puzzling. The same result was also obtained in the reduced form analysis, but here the this covariate is highly significant. The covariates BookStd and LBookStd play a very central role in the models for μ and σ, respectively. The large negative sign for MinBidStd is explained by the fact that a higher minimum bid implies a higher cut-off signal, which reduces the number of positive bids. Overall, the posterior inclusion probabilities are either close to 0 or 1, which gives conclusive evidence on which covariates that are important for explaining valuations and participation in eBay auctions for common value objects. It interesting to note that eBay's detailed seller information seems to be of little use to buyers: the covariates Power, ID and LargNeg (dummy for a large proportion of negative feedbacks from buyers) are almost invariably removed from the model by the algorithm. We experimented with other transformations of the negative feedback score and also transformations of the overall feedback score as substitutes for Power and ID, with unchanged results.

To check for the sensitivity of the prior hyperparameters, we repeated the estimations using several different priors. In Table 6 we use the estimated model in Table 5 as a benchmark and compare posterior means given various prior settings. Almost all parameter estimates are insensitive to changes in priors, especially in the model for λ. Notable changes only appear

Table 5. Posterior inference for the structural model

Parameter	Covariate	Mean	St Dev	t-ratio	Incl prob	IF
κ	-	3.042	0.155	-	-	27.668
μ	Const	28.334	0.213	-12.462	1.000	10.857
	BookStd	0.733	0.010	58.062	1.000	15.728
	Book*Pow	0.022	0.011	1.775	0.021	-
	Book*ID	0.116	0.030	3.420	0.967	31.986
	Book*Unopen	0.392	0.027	13.928	1.000	13.413
	Book*MinBlem	-0.028	0.019	-1.958	0.020	-
	Book*MajBlem	-0.252	0.028	-8.984	1.000	9.827
	Book*LargNeg	-0.006	0.020	0.682	0.008	-
$log(\sigma^2)$	Const	3.340	0.027	93.668	1.000	23.766
	LogBookStd	1.477	0.035	43.556	1.000	32.272
	LBook*Pow	0.008	0.015	0.996	0.007	-
	LBook*ID	0.048	0.033	1.582	0.010	-
	LBook*Unopen	0.283	0.020	12.249	1.000	15.164
	LBook*MinBlem	-0.007	0.020	-0.867	0.008	-
	LBook*MajBlem	0.017	0.031	0.897	0.007	-
	LBook*LargNeg	0.043	0.019	2.109	0.044	-
$log(\lambda)$	Const	1.448	0.022	60.094	1.000	38.941
	Pow	0.054	0.041	0.824	0.008	-
	ID	-0.076	0.058	-1.455	0.005	-
	Unopen	0.299	0.042	7.366	1.000	23.823
	MinBlem	-0.025	0.036	-1.475	0.004	-
	MajBlem	-0.145	0.075	-2.117	0.012	-
	LargNeg	0.027	0.030	1.343	0.008	-
	LogBookStd	-0.087	0.025	-3.562	0.865	25.164
	MinBidStd	-1.452	0.076	-25.307	1.000	65.725

Note: $c = n, \bar{\kappa} = 0.25, g = 4$, and $\pi = 0.2$. The last column displays the inefficiency factors.

in the model for σ, where the estimated value of the constant increases and the estimated value of the parameter for LBookStd decreases with more prior information. The tighter prior distribution around zero reduces the impact of LBookStd and to compensate for this the estimated constant is increased. The next to last column in Table 6 shows that varying the prior inclusion probabilty, π, does not affect the posterior mean estimates, as expected.

Finally, the last column of Table 6 gives the posterior mean estimates when the benchmark model is estimates on all bids ($\delta = 0$), see Section 5.1. The main difference in results is in κ and σ who both naturally increases as we face a wider dispersion in bids. Estimations with $\delta = 0.5$ (results not shown) reduced κ to 2.27. It should be noted however that the relation between κ and the unconditional variance of the bids is complicated by the bidders' strategic behavior. One way to see this is by looking at the unconditional variance of the bids $\text{var}(b) = (1-\omega)^2(\kappa+1)\sigma^2$ under the approximate bid function. var(b) obviously increases with κ via the factor $(\kappa + 1)$, but κ also affects var(b) via ω (the weight placed on the prior mean) in a non-linear fashion that depends on λ. Figure 4 displays Std(b) as a function of κ for $\sigma^2 = 1$, for different values on λ. It is seen from Figure 4 that Std(b) is fairly insensitive to κ over a large region. There are larger differences in Std(b) across λ for a given κ.

Table 6. Sensitivity analysis - Posterior means

Coeff	Covariate	Bench	$c = n/16$	$c = n/64$	$c = n/256$	$\pi = .5$	$\delta = 0$
κ		3.042	3.245	**3.676**	3.217	3.013	3.251
μ	Const	28.334	28.441	28.544	**28.974**	28.2385	28.536
	BookStd	0.733	0.728	**0.705**	**0.610**	0.730	0.743
	Book*Power	0.022	0.021	0.025	0.017	0.022	0.018
	Book*ID	0.116	0.108	0.093	0.068	0.116	0.107
	Book*Unopen	0.392	0.382	0.344	**0.275**	0.392	0.389
	Book*MinBlem	−0.028	−0.031	−0.024	−0.018	−0.026	−0.045
	Book*MajBlem	−0.252	−0.254	−0.246	−0.228	−0.255	−0.244
	Book*LargNeg	−0.006	−0.005	0.003	−0.001	0.003	−0.003
σ	Const	3.340	**3.415**	**3.522**	**3.612**	3.333	**3.441**
	LBookStd	1.477	**1.367**	**1.176**	**0.956**	1.473	**1.646**
	LBook*Power	0.008	0.013	0.028	0.029	0.007	0.010
	LBook*ID	0.048	0.033	0.029	−0.006	0.039	0.051
	LBook*Unopen	0.283	0.259	**0.218**	**0.146**	0.283	0.256
	LBook*MinBlem	−0.007	−0.008	0.014	**0.046**	−0.021	0.025
	LBook*MajBlem	0.017	0.030	0.057	0.072	0.018	**−0.066**
	LBook*LargNeg	0.043	0.040	0.046	0.050	0.038	0.028
λ	Const	1.448	1.441	1.428	**1.400**	1.446	1.470
	Power	0.054	0.007	0.026	0.014	0.021	0.025
	ID	−0.076	−0.072	−0.066	−0.086	−0.073	−0.054
	Unopen	0.299	0.302	0.281	0.317	0.305	0.279
	MinBlem	−0.025	−0.059	−0.055	−0.057	−0.031	−0.013
	MajBlem	−0.145	−0.131	−0.074	−0.139	−0.159	−0.173
	LargeNeg	0.027	0.044	0.007	0.035	0.034	0.013
	LBookStd	−0.087	−0.072	−0.058	−0.090	−0.091	−0.051
	MinBidShare	−1.452	−1.436	−1.437	−1.507	−1.472	−1.437

The columns display the posterior mean of each parameter conditional on inclusion.
The bench column displays the results for the benchmark model, the other columns are results for Variations of the benchmark model. Bold numbers indicate that the estimate lies outside the benchmark model's 95% posterior credibility interval.

In Table 7, we also check for the sensitivity of the prior hyperparameters by comparing posterior inclusion probabilities for various prior settings and using the estimated model in Table 5 as a benchmark. Overall, small inclusion probabilities for the benchmark model tend to increase as the prior information becomes more precise (c increases), which is exactly the same expected result as we discerned from the simulation study in Section 4. Again the results are not overly sensitive to c. Furthermore, increasing the prior inclusion probability from 0.2 to 0.5 does not overturn the previous results on the variable selection inference.

5.4. **Model evaluations and predictions.** We will now evaluate the fit of the model by comparing the observed data to simulated data from the estimated benchmark model. Given the observed auction specific covariates, we simulated 100 new complete data sets for each of a 100 systematically sampled posterior draws of the model parameters. This gives us 10.000 full data sets, each with bids from a 1000 auctions.

Following BH, we compare the observed and simulated data through two summary statistics: within-auction bid dispersion and cross-auction heterogeneity. As we can see, in Figures 5 and

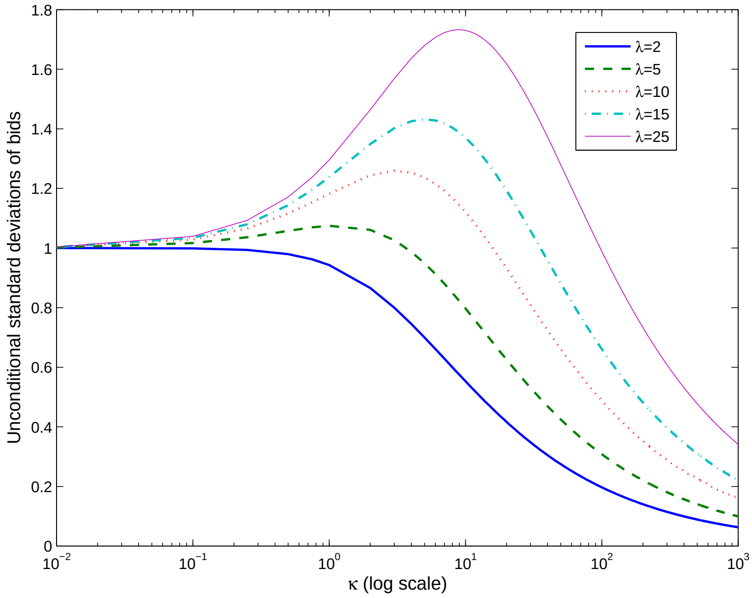

FIGURE 4. The unconditional standard deviation of bids as a function of the variance scale parameter κ.

6 the observed within-auction bid dispersion and the cross-auction heterogeneity are very well captured by the model. This is not the case in BH where they greatly underestimate within-auction bid dispersion and overestimate cross-auction bid dispersions.

A more severe test of the model is to evaluate its out-of-sample predictions. To test this, we use the estimated benchmark model in Section 5.3 to predict a fresh data set of 50 additional auctions of *U.S.* proof sets. Given the covariates from these auctions, we simulated price distributions for each auction in a similar way as above. The parameters are drawn from the posterior distribution computed using only the previously analyzed 1000 auctions. Figures 1 and 2 display the predictive distributions in the 50 out-of-sample auctions. Note these distributions have three components: i) a probability that the item is not sold (Pr(No)), ii) a point mass at the minimum bid (Pr(Min)), which is the price when there is a single bidder in the auction, and iii) a continuous density when there is more than one bidder. The predictive price distributions look reasonable and capture the observed prices very well in most cases. This fact together with the good fit of the bid dispersion indicates strongly that our estimated eBay auction model is quite accurate in explaining seller and bidder behaviour at eBay.

Table 7. Sensitivity analysis - Posterior inclusion probabilities

Coeff	Covariate	Bench	$c = n/16$	$c = n/64$	$c = n/256$	$\pi = .5$	$\delta = 0$
κ		–	–	–	–	–	–
μ	Const	1.000	1.000	1.000	1.000	1.000	1.000
	BookStd	1.000	1.000	1.000	1.000	1.000	1.000
	Book*Power	0.021	0.111	**0.223**	**0.213**	**0.144**	0.037
	Book*ID	0.967	0.977	0.946	0.838	0.982	**0.762**
	Book*Unopen	1.000	1.000	1.000	1.000	1.000	1.000
	Book*MinBlem	0.020	0.066	0.089	**0.134**	0.050	0.061
	Book*MajBlem	1.000	1.000	1.000	1.000	1.000	1.000
	Book*LargNeg	0.008	0.031	0.064	0.072	0.016	0.007
σ	Const	1.000	1.000	1.000	1.000	1.000	1.000
	LBookStd	1.000	1.000	1.000	1.000	1.000	1.000
	LBook*Power	0.007	0.028	**0.295**	**0.772**	0.038	0.021
	LBook*ID	0.010	0.037	0.056	0.071	0.096	0.015
	LBook*Unopen	1.000	1.000	1.000	1.000	1.000	1.000
	LBook*MinBlem	0.008	0.020	0.048	**0.631**	0.019	0.008
	LBook*MajBlem	0.007	0.044	**0.133**	**0.786**	0.019	0.069
	LBook*LargNeg	0.044	**0.203**	**0.547**	**0.915**	**0.235**	0.010
λ	Const	1.000	1.000	1.000	1.000	1.000	1.000
	Power	0.008	0.021	0.031	0.069	0.016	0.001
	ID	0.005	0.019	0.058	**0.136**	0.025	0.006
	Unopen	1.000	1.000	1.000	1.000	1.000	1.000
	MinBlem	0.004	0.027	0.055	**0.105**	0.010	0.001
	MajBlem	0.012	0.053	0.037	**0.246**	0.090	**0.159**
	LargeNeg	0.008	0.025	0.017	0.063	0.019	0.003
	LBookStd	0.865	**0.701**	**0.472**	**0.979**	**0.986**	**0.070**
	MinBidShare	1.000	1.000	1.000	1.000	1.000	1.000

The columns display the posterior inclusion probability of each parameter.
The bench column displays the results for the benchmark model, the other columns are results for variations of the benchmark model. Bold numbers indicate a difference from bench by more than 0.1.

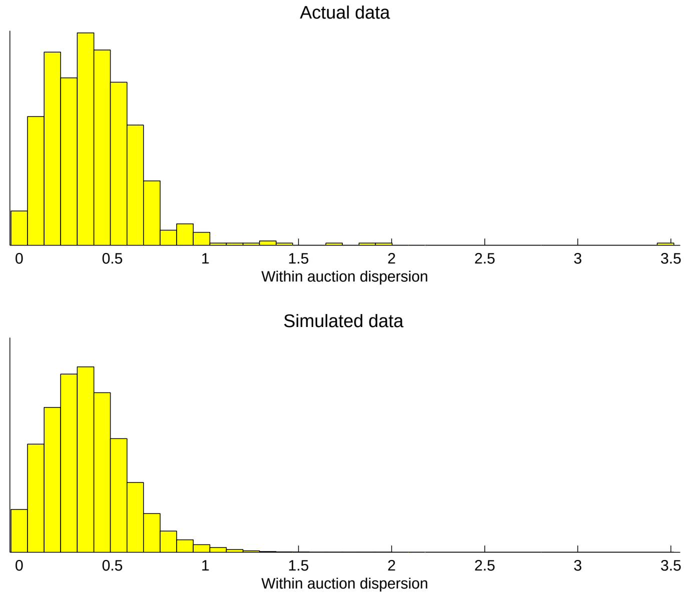

FIGURE 5. Posterior predictive check of the within-auction dispersion. The within-auction dispersion is defined as the difference between the highest observed bid and the lowest bid divided by the auctioned item's book value.

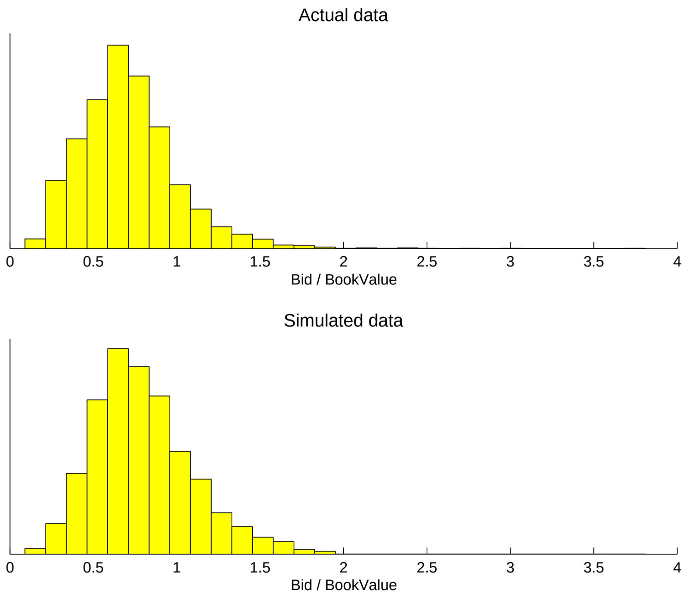

FIGURE 6. Posterior predictive check of the cross-auction heterogeneity. The figure displays histograms of the bids divided by the corresponding book value in that auction.

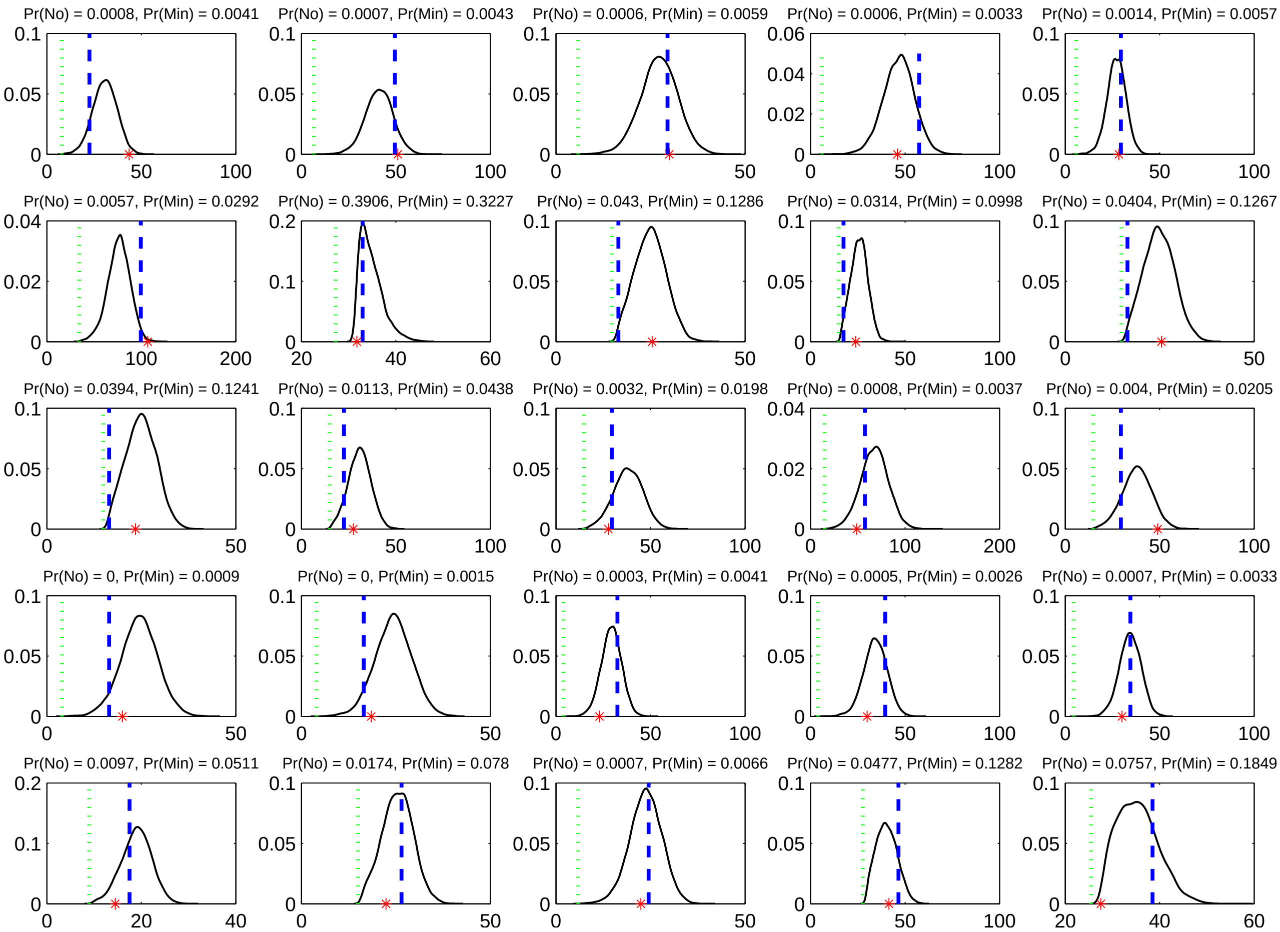

FIGURE 7. Out-of-sample prediction for the first 25 auctions in the evaluation sample. Each subplot displays the realized price marked out by a star (if the item is sold), the minimum bid (vertical dotted line), and the book value (vertical dashed line). Pr(No) and Pr(Min) are the predictive probabilities of no bids and a single bid (in which case the winner pays the seller's posted minimum bid), respectively. The solid curve is the predictive density conditional on at least two bids.

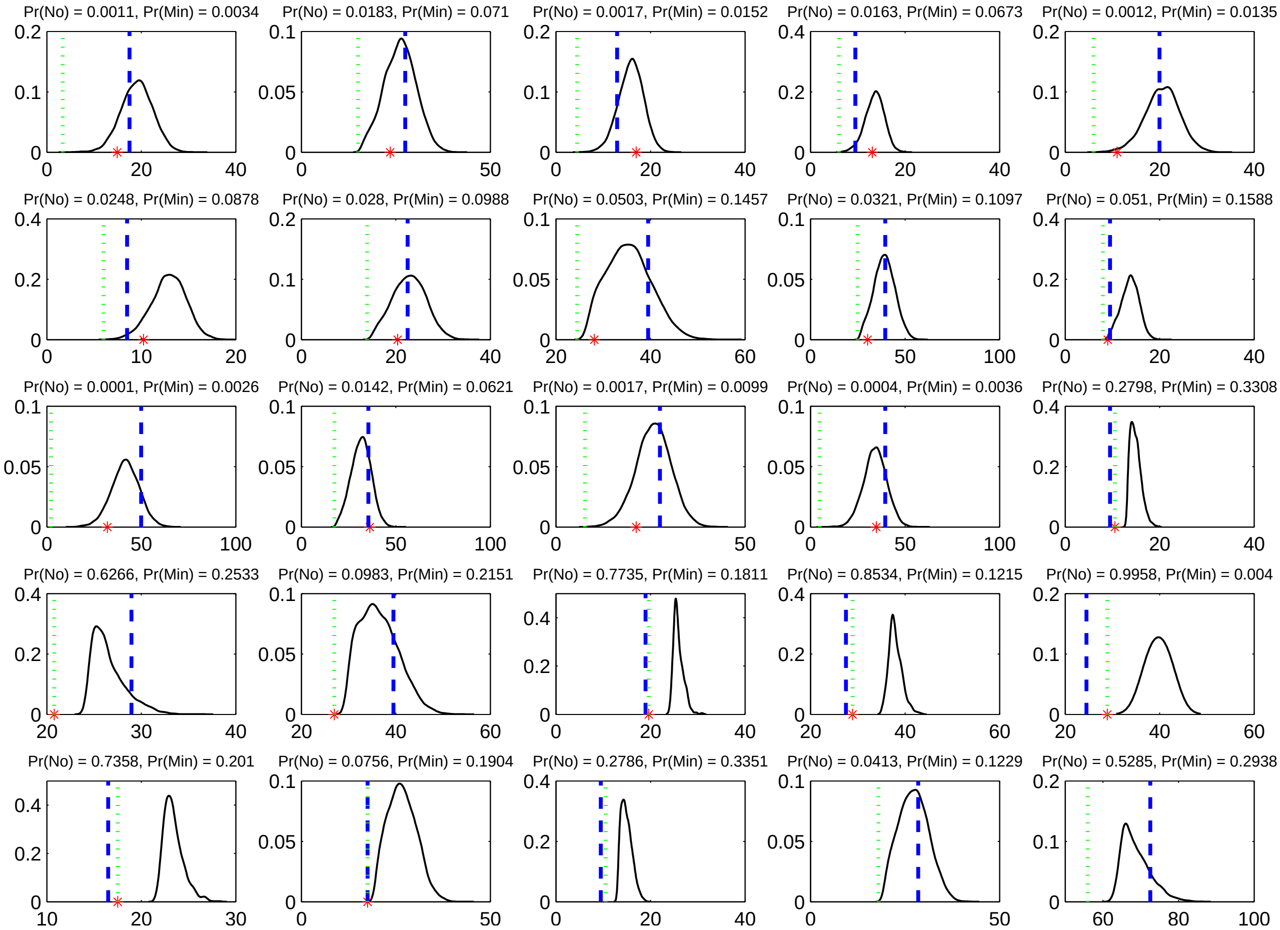

FIGURE 8. Out-of-sample prediction for the last 25 auctions in the evaluation sample. Each subplot displays the realized price marked out by a star (if the item is sold), the minimum bid (vertical dotted line), and the book value (vertical dashed line). Pr(No) and Pr(Min) are the predictive probabilities of no bids and a single bid (in which case the winner pays the seller's posted minimum bid), respectively. The solid curve is the predictive density conditional on at least two bids.

6. Conclusions

In contrast to private value auctions, the estimation of models within the common value framework is technically challenging. An important result of our paper is therefore the nearly identical inferences for the approximate bid function in Wegmann (2007) and the exact bid function. Thus, the approximate bid function is used to get fast, explicit and analytically invertible solutions of the equilibrium bid function, which in turn gives much faster and numerically more stable evaluations of the likelihood function.

We have shown how efficient Bayesian inference with variable selection can be performed in a general setting, including structural models for second price common value auctions. Simulated and real data was used to document the good performance of algorithm, and to investigate the sensitivity of the variable selection to changes in the prior hyperparameters. Bayesian t-ratios turned out to be a simple way to characterize the posterior inclusion probabilities.

We collected a high-quality dataset from coin auctions on eBay and analyzed it both with reduced form regressions and using a structural model for second price common value auctions. The results pointed strongly to book values as the most important predictor of common values. The minimum bid turned out to be the main determinant of the number of bids in an auction. Interestingly, the detailed seller information provided by eBay, and eBays feedback score system seemed to be of very little value to the buyers. No such covariate did consistenly get a low posterior inclusion probability in the estimations. Finally, unopened coin envelopes attracted high bidding activity and were sold at unusually large prices.

The estimated eBay model captured the within-auction bid dispersion and the cross-auction heterogeneity very well. A more severe test was to evaluate the out-of-sample predictions. The predictive price distributions looked reasonable and captured the observed prices very well in most cases. Overall, this fact together with the good fit of the bid dispersion indicates strongly that our estimated eBay auction model is quite accurate in explaining seller and bidder behaviour at eBay.

Finally, possible extensions for future research could be to take into account and model cross-auction heterogeneity. A careful inspection of the bids revealed that some bids are only a tiny fraction of the objects book value, and cannot realistically represent serious final bids. One possible explanation could stem from the presence of incremental bidding. Ockenfels and Roth (2006) argue that late bidding may also arise out of equilibrium as a best reply to incremental bidding. Another possibility could be the effect that bidders are searching the eBay marketplace for low-price auctions, see Sailer (2006) for non-parameterically identification of bidding costs in settings where the bidder searches with a "'reservation bid"' for low-price auctions. Other possible extensions include auctions with both a private and a common value element of the object, multiunit objects, or auctions with risk-averse bidders that we do not cover in this paper.

Appendix A: Description of eBay covariates

The following list describes each covariate for our collected data of 1000 eBay auctions:

Book Value: The price of the item as reported by the large Internet coin seller *Golden Eagle Coins* at http://www.goldeneaglecoin.com.

Minor Blemish: Dummy variable, coded as 1 if the proof set had a minor damage on the box or packaging according to careful subjective assessment of the item using the seller's detailed descriptions and pictures.

Major Blemish: Dummy variable, coded as 1 if at least one coin were missing or if other major imperfections were present.

Seller's Feedback Score: For each transaction, buyers can rate the seller in terms of reliability and timeliness in delivery. Next to the seller's ID the net positive responses is displayed.

PowerSeller: Dummy variable, coded to be 1 if the seller is a *PowerSeller*. A PowerSeller is a seller who ranks among the most successful sellers in terms of product sales and customer satisfaction on eBay. The icon next to a member's user ID means that the seller meets the criteria for being a PowerSeller.

ID Seller: Dummy variable, coded to be 1 if the seller's ID is verified.

UnOpened: Dummy variable, coded to be 1 if the proof set is delivered sealed and unopened in its original envelope. Most common for proof sets that origins from the fifties or sixties.

References

[1] Bajari, Patrick and Ali Hortacsu (2003), The Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions, *The Rand Journal of Economics*, **3(2)**, 329-355.

[2] —– (2004), Economic Insights from Internet Auctions, *Journal of Economic Literature*, **42**, 457-486.

[3] —– (2005), Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data, *Journal of Political Economy*, **113(4)**, 703-741.

[4] Bauwens, L. Lubrano, M., and Richard, J-F. (1999),*Bayesian Inference in Dynamic Econometric Models*. Oxford University Press: Oxford.

[5] Engelbrecht-Wiggans, R. and R. J. Weber (1979), On the Nonexistence of Multiplicative Equilibrium Bidding Strategies, *Cowles Foundation Discussion Paper No. 523*, Yale University.

[6] Elyakime, B., Laffont J.J., Loisel, O., and Vuong, Q. (1997) Auctioning and Bargaining: An Econometric Study of Timber Auctions With Secret Reservation Prices, *Journal of Business & Economic Statistics*, **15(2)**, 209-20.

[7] Fletcher, R. (1987), *Practical Methods of Optimization*, Wiley, New York.

[8] Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (2004). *Bayesian Data Analysis, 2nd edition*, Chapman & Hall, New York.

[9] George, E., and McCulloch, R. (1993), Variable selection via Gibbs sampling, *Journal of the American Statistical Association*, **88**, 881-889.

[10] Giordani, P. and Kohn, R. (2008), Adaptive MCMC, Forthcoming in Advances in Econometrics.

[11] Laffont, J.J. and Vuong, Q. (1996) Structural Analysis of Auction Data, *The American Economic Review*, **86(2)**, 414-20.

[12] Lucking-Reiley, D. (2000), Auctions on the Internet: What's Being Auctioned, and How?, *Journal of Industrial Economics*, **48(3)**, 227-252.

[13] Milgrom, Paul and Robert Weber (1982), A Theory of Auctions and Competitive Bidding, *Econometrica*, **50**, 1089-1022.

[14] Nott, D. J., and Kohn, R. (2005), Adaptive sampling for Bayesian variable selection, *Biometrika*, **92**, 747-763.

[15] Nott, D. J., and Leonte, D. (2004), Sampling schemes for Bayesian variables selection in generalized linear models, *Journal of Computational and Graphical Statistics*, **13**, 362-382.

[16] Ockenfels, A. and Roth, A. E. (2006), Late and multiple bidding in second price Internet auctions: Theory and evidence concerning different rules for ending an auction, *Games and Economic Behaviour*, **55**, 297-320.

[17] Paarsch, J. Harry (1992), Deciding between the Common and Private Value Paradigms in Empirical Models of Auctions, *Journal of Econometrics*, **51**, 191-215.

[18] Sailer, K. (2006), Searching the eBay Marketplace, mimeo, Department of Economics, University of Munich.

[19] Schindler, J. (2003), Late Bidding on the Internet, mimeo, Vienna U. Econ. and Bus. Admin.

[20] Smith, M., and Kohn, R. (1996), Nonparametric regression using Bayesian variable selection, *Journal of Econometrics*, **75**, 317-344.

[21] Wegmann, B. (2007), Bid function approximations of second price common value auctions, Research Report 2007:4, Department of Statistics, Stockholm University.

[22] Wilcox, R. (2000), Experts and Amateurs: The role of Experience in Internet auctions, *Marketing letters*, **11(4)**, 363-74.

[23] Wilson, R. (1977), A Bidding Model of Perfect Competition, *Review of Economic Studies*, **44**, 511-18.

[24] Zellner, A. (1986), On assessing prior distributions and Bayesian regression analysis with g-prior distributions. In *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti*, 233–243. North-Holland/Elsevier.