

Structural Estimation of Auction Data through Equilibrium Learning and Optimal Transport

Research Paper

Markus Ewert¹ and Martin Bichler^{1,2}

¹ Technical University of Munich, School of Computation, Information and Technology, Munich, Germany

{markus.ewert,m.bichler}@tum.de

² Munich Data Science Institute, Munich, Germany

Abstract. We propose an extension to the two-step structural estimation framework for auction data introduced by Guerre et al. (2000). Traditionally, this technique relies on kernel estimators to approximate the distribution and density of observed bids, which are then used to infer the latent prior distribution over bidders' types. However, kernel-based methods suffer from statistical limitations, often leading to biased estimations in real-world applications. To address these challenges, we reformulate the estimation problem as an optimal transport task between the observed bids and a proxy equilibrium model. This approach enables analytical evaluation of the bid distribution and density, improving estimation accuracy. Through numerical experiments on experimental auction data, our estimator consistently outperforms established techniques in terms of Wasserstein distance, demonstrating its robustness and applicability.

Keywords: Structural Estimation, Auctions, Equilibrium Learning

1 Introduction

In many real-world settings, the allocation of goods, resources, and services from sellers to buyers is facilitated through auctions. These mechanisms are widely used across diverse economic sectors, from art and flower auctions to multi-billion dollar markets such as online advertisement auctions or crowdsourcing contests. A fundamental question in auction theory is how to maximize the auctioneer's expected revenue or minimize their expected costs in procurement settings. The design of these markets relies on information about bidders, which is often unobservable or only partially known.

For instance, Liu et al. (2014) showed that in crowdsourcing platforms, where users compete for prizes by submitting answers to questions, the quality of submissions strongly depends on selecting the optimal prize structure.¹ However, determining this structure requires knowledge of the contestants' ability levels. Similarly, search engines allocate slots for sponsored links based on bids for specific keywords in sponsored search auctions. A field experiment by Ostrovsky & Schwarz (2023) demonstrated that setting optimal reserve prices can substantially increase revenue in these markets. However,

¹ A crowdsourcing contest can be described as a multi-unit all-pay auction.

choosing these prices requires knowledge of bidders’ valuations to obtain a slot, which is typically unknown to the auctioneer.

In auction theory, private information is modeled as a prior distribution over bidders’ types, such as their valuations, costs, or ability levels. The Wilson doctrine posits that practical applications should minimize reliance on common knowledge assumptions (Wilson 1985). In the context of auctions, the prior distribution is assumed to be common knowledge among participants but inaccessible to the designers of such markets. Thus, research on structural estimation in auctions seeks to develop methods for recovering this latent information from the observed bidding behavior (Hickman et al. 2012).

Existing methods rely on strong assumptions and can be categorized as parametric or non-parametric methods. Traditional parametric approaches require knowledge of the analytical equilibrium strategies, which are only available for a small number of standard single-object auctions (Donald & Paarsch 1996). Moreover, they assume that the parametric form of the latent distribution is known ex-ante, which is rarely the case with field data. Conversely, deriving a non-parametric estimator is challenging (Shakhgildyan 2022), and resulting techniques are restricted to a specific auction mechanism.

Guerre et al. (2000) (GPV) introduced a two-step estimation framework to recover the prior distribution in *First-Price Sealed Bid* (FPSB) auctions, which serves as the foundation for subsequent non-parametric estimators in this domain. Their key insight is that the valuations can be expressed as a function of the number of bidders, the observed bids, their distribution, and density. However, since auctioneers only observe bids and the number of participants, the density and distribution must be estimated. Kernel estimators are commonly used for this purpose, but they suffer from several statistical limitations (Härdle 2004). These methods require large datasets to ensure estimation accuracy, which are often unavailable in real-world applications. Moreover, their predictive performance is highly sensitive to the choice of kernel type and bandwidth parameter, and they often exhibit bias near the boundaries of the support.

We propose an extension to the GPV framework that eliminates reliance on kernel estimators, making it more flexible and practical for real-world applications. Our approach reformulates the estimation problem as an optimal transport task between observed bids and a proxy equilibrium model, drawing from the optimal transport and equilibrium learning literature. This allows for the analytical evaluation of the density and distribution via the optimal transportation map between the proxy model and observed bids. To achieve this, we employ equilibrium learning techniques (Bichler et al. 2021) to solve the equilibrium problem in the proxy model in settings where an analytical equilibrium is unavailable.

We evaluate our approach using two experimental datasets and compare its performance against the GPV framework and the extension of Hickman & Hubbard (2015) that impose boundary correction methods on the kernel estimator. Our analysis shows that the proposed estimator consistently outperforms the kernel-based techniques in Wasserstein distance across all datasets. This demonstrates its robustness and applicability in real-world scenarios.

1.1 Contribution

Our study advances the literature on structural estimation techniques in auctions. The contributions of this paper are as follows: First, we introduce our extension of the GPV framework that eliminates the need for kernel-based estimation procedures, making it significantly more flexible than existing techniques. Second, we conduct a numerical analysis comparing our estimator to established methods and find that our approach consistently outperforms them. The following paragraphs elaborate on these contributions in more detail.

The proposed estimator Established structural estimation techniques face several challenges. Parametric approaches require the auctioneer to impose a parametric assumption on the latent structure, which is often infeasible in real-world settings. Additionally, these techniques rely on knowledge of an analytical equilibrium strategy, which is often unavailable in closed form. On the other hand, non-parametric estimators are difficult to derive and specifically tailored to a particular auction setting. Additionally, existing approaches rely on complex non-parametric models to represent the distributions, which face statistical challenges like bias at the boundaries of the support, noisy observations, and limited data availability. Our estimator directly addresses these issues, making it more versatile in realistic environments.

Our key observation is that in single-item auctions, where the prior distribution is univariate, the density and distribution of the observed bids can be described using a proxy equilibrium model that is independent of the equilibrium problem bidders face in the data-generating process. The optimal transport map between observed bids and those induced in the equilibrium problem is then given by a monotonic function. Using its inverse, we can assess the distribution of the observed bids directly via the known distribution in the proxy model. This approach allows us to compute pseudo-valuations in the two-step estimation framework by Guerre et al. (2000). We can then construct a transport map between the prior distribution in the proxy model and the pseudo-valuations, enabling us to directly evaluate the density and distribution of the latent prior distribution to compute comparative statics.

In the proxy model, we solve an equilibrium problem. Analytical derivations of closed-form bidding strategies are only available for simple models (Krishna 2009). We employ equilibrium learning techniques to approximate equilibria in games instead of deriving them analytically (see Section 2.2).

Numerical analysis We conduct a numerical analysis on two different experimental datasets, where the prior distribution is known, to compare the performance of our framework against the existing methods. Our results demonstrate that our technique consistently outperforms the kernel-based techniques across all datasets in terms of the Wasserstein-2 distance. This superior performance translates into better-informed managerial decisions based on the estimation results. To assess the impact on auction design, we simulated how optimal reserve prices derived from each prior model influence the auctioneer’s expected revenue. In both datasets, our method yields reserve prices that generate the highest revenue gains.

1.2 Organization

This paper is structured as follows. We begin with a brief overview of the related literature on structural estimation and equilibrium learning. Next, we introduce the notation and present our estimation framework. We then describe the numerical analysis used to compare the performance of our approach in recovering the latent prior distribution against the kernel-based estimators from the literature. Finally, we conclude the article by summarizing our findings and discussing their managerial implications.

2 Related work

Our study contributes to the literature on structural estimation in auctions. The following section provides a brief overview. Additionally, since we leverage equilibrium learning techniques to approximate equilibria in these games, we also review relevant research in this area.

2.1 Structural Estimation in Auctions

Research on structural estimation focuses on developing methods to recover the latent structures of an econometric model from a set of observables (Manski et al. 1981). Since auction rules inherently define the data-generating process, these mechanisms are well-suited for structural estimation techniques. Considerable research was dedicated to developing methods to estimate the primitives of auctions to explain observed bidding behavior. These methods span a spectrum, with parametric approaches at one end and non-parametric estimators at the other. Most studies focus thereby on recovering the prior value distribution describing the bidders' types.

Parametric approaches assume that the auctioneer knows the shape of the latent structure but not its actual realization. By imposing a parametric model on this structure, the estimation task reduces to recovering the data-generating parameters. One of the earliest examples in this domain is the method proposed by Donald & Paarsch (1996), who assume that in a FPSB auction, the family of the prior distribution is known, but its parametrization is not. They develop a maximum likelihood estimator to infer these parameters using the observed bids and the number of participants. However, this approach requires knowledge of the analytical equilibrium strategy and its inverse. Computing the analytical equilibrium for arbitrary prior distributions is, even under the assumption of symmetric, risk-neutral bidders, often infeasible (Cai & Papadimitriou 2014). This computational burden, coupled with the challenge of correctly specifying the underlying distributional family in real-world settings, significantly limits the applicability of parametric approaches.

In contrast, non-parametric approaches neither rely on a parametric assumption nor require knowledge of the analytical equilibrium. Instead, they estimate the latent structure - specifically, the prior distribution - using flexible models such as kernel-density estimators. In their seminal work, Guerre et al. (2000) proposed a two-step estimation framework for non-parametrically recovering the data-generating prior distribution. Their key idea is to leverage the first-order condition of the expected utility to express the

valuations as a function of bids, their distribution, and the number of bidders. The first step involves estimating the bid distribution non-parametrically using kernel-density estimators. This estimate is then used to derive pseudo-valuations, which ultimately allow for the non-parametric estimation of the value distribution.

Kernel-density estimators are prone to bias near the boundaries of the support, which can impact the estimation performance in recovering the latent distribution. Guerre et al. (2000) address this issue by trimming observations near the boundaries to improve kernel performance. However, trimming reduces the effective sample size in both estimation stages and undermines the performance of standard bandwidth selection methods. To overcome this, Hickman & Hubbard (2015) apply a boundary correction technique originally proposed by Zhang et al. (1999), which mitigates boundary bias without discarding data and allows for the use of conventional bandwidth rules. Their numerical analysis demonstrates that this leads to more robust estimations of the framework by Guerre et al. (2000). However, these estimators require large amounts of data to produce accurate estimates. This is often not the case in real-world scenarios, where the number of observations is restricted, or datasets are truncated (Hickman et al. 2012), indicating that the applicability of this technique in the field is limited.

The approach by Guerre et al. (2000) is restricted to FPSB auctions with symmetric, risk-neutral bidders, where the auctioneer knows the number of participants ex-ante. These assumptions are often too restrictive in practice, and thus, research developed extensions to mitigate this issue. Li et al. (2002) considered a model with affiliated price values rather than independent bidder types, while Flambard & Perrigne (2006) extended the framework to settings with bidders that are asymmetric in their types. A common practice to increase the auctioneer's expected revenue is to set reserve prices, which was covered by Li & Perrigne (2003). Laboratory experiments in auctions indicate that bidders frequently overbid their risk-neutral equilibrium strategy, where one explanation is that bidders are risk-averse (Cox et al. 1985, 1988). Although models that consider estimating the prior distribution and a risk-averse utility simultaneously are not identified (Campo et al. 2011), Campo et al. (2011) and Lu & Perrigne (2008), among others, developed strategies to achieve this with additional information about the observed behavior. Furthermore, Shakhgildyan (2022) extended the approach to multi-unit all-pay auctions, accommodating both risk-neutral and risk-averse bidders under independent and affiliated private values.

Unlike existing approaches, our estimator imposes fewer restrictive assumptions, making it more flexible. Since we focus on scenarios where the bidders' types and bids are univariate random variables, we can express the transportation map between the proxy equilibrium and the observed bids analytically. In contrast to kernel-based estimators, our approach eliminates the need for extensive hyperparameter tuning, which would otherwise introduce a bias in the estimation. This, in turn, enhances managerial decisions based on the approximated latent structure, as inaccurate estimates of the data-generating process could lead to economically detrimental decisions, such as a lower expected revenue of the auctioneer.

2.2 Equilibrium Learning

As mentioned previously, deriving an analytical equilibrium is generally infeasible for arbitrary auctions, and equilibrium strategies are only known for a few standard auctions. Advances in the field of equilibrium learning allow for the numerical approximation of game equilibria instead. Broadly speaking, these techniques find an approximate equilibrium through an iterative process of adaption by learning the dynamics of the game as opposed to solving the underlying system of differential equations (Fudenberg & Levine 1998). Most research in this area focuses on computing Nash equilibria in finite, complete-information games. Examples include the class of best or better response dynamics (Abreu & Rubinstein 1988, Fudenberg & Levine 1998), as well as regret minimization algorithms (Hart & Mas-Colell 2000, Bowling 2004).

Auctions are Bayesian games with continuous types and actions where these methods cannot be applied. However, there exist algorithms that intentionally were developed to find approximate equilibria in these types of games. Bichler, Fichtl & Oberlechner (2023) introduced a method that computes the Nash equilibrium in the approximate game, which is a discretized version of the original Bayesian game. They further show that if the procedure converges to an equilibrium in this approximate game, then it is an approximation of the equilibrium in the continuous game. Similarly, Bichler et al. (2021) proposed *Neural Pseudo Gradient Ascent* (NPGA) where they represent the strategies as neural networks and iteratively optimize their parameters by following the gradient dynamics of the game using simultaneous gradient ascent. Numerical experiments demonstrate that both algorithms approximate equilibria with small errors in settings where an analytical equilibrium is known. Moreover, they also found equilibria in settings where an analytical derivation was intractable (Bichler, Kohring & Heidekrüger 2023, Bichler, Kohring, Oberlechner & Pieroth 2023, Bichler, Gupta, Mathews & Oberlechner 2023).

Despite their empirical success, it is yet unknown whether and under which circumstances these techniques converge to an equilibrium. For the special case of first-price and all-pay auctions, Ahunbay & Bichler (2025) established ex-ante guarantees under which no-regret learners converge to an equilibrium. The technique by Bichler, Fichtl & Oberlechner (2023) satisfies the criteria of such a learner. Nonetheless, by placing a grid over the type and action spaces and assessing whether deviating from the current strategy improves the expected payoff of a bidder, both methods can numerically verify that they found an approximate equilibrium. The error is determined by the *utility loss*, which we explain in detail in Section 3.1. We will leverage NPGA as the equilibrium learning technique to represent the equilibrium oracle in our numerical analysis.

3 Estimation Framework

The following section first provides an overview of the notation and definitions used. It then introduces the identification strategy for recovering latent distributions from observed bidding data in single-item auctions.

3.1 Preliminaries

An auction is a Bayesian game $\Gamma = (N, F, \mathcal{V}, \mathcal{A}, u)$. $N \in \mathbb{N}$ is the number of risk neutral bidders, and $[N] = \{n \in \mathbb{N} | n \leq N\}$ is the set of players. We use $i \in [N]$ to index a specific bidder and $-i$ to denote all bidders except i . Prior to the auction, the players have common knowledge of a joint prior distribution F over the set of type profiles $\mathcal{V} = \times_{i \in [N]} \mathcal{V}_i$, with density f . We assume *independent-private-values*. This study considers the type as the valuation for a specific item. Meanwhile, our framework also applies to other games, such as procurement auctions or contests, where the type describes the costs of exerting an effort. After observing their private types by drawing from this prior distribution, each bidder i determines their bid $b_i \in A_i$, with $A = \times_{i \in [N]} A_i$, using a bid function $\beta_i : \mathcal{V}_i \rightarrow A_i$. As common in the literature on auctions, we assume that β is strictly monotonically increasing in the type. We denote the inverse bid function of player i as $\beta_i^{-1} : A_i \rightarrow \mathcal{V}_i$. After the bidders submit their bids to the auctioneer, they observe their utility $u_i : \mathcal{V}_i \times \mathcal{A} \rightarrow \mathbb{R}$ given the current bid profile $b = (b_1, \dots, b_N)$ and their type v_i . In this article, we focus on the ϵ -Bayes Nash Equilibrium (ϵ -BNE) as the solution concept for this game.

Definition 1. An *ex-ntem* ϵ -BNE of a game Γ is a strategy profile $\beta^* = (\beta_1^*, \dots, \beta_N^*)$, in which none of the players can unilaterally deviate from their current strategy to improve their expected utility by more than ϵ :

$$\begin{aligned} & \mathbb{E}_{v_{-i} \sim F}(u_i(v_i, \beta_i(v_i), \beta_{-i}^*(v_{-i}))) - \mathbb{E}_{v_{-i} \sim F}(u_i(v_i, \beta_i^*(v_i), \beta_{-i}^*(v_{-i}))) \\ & \leq \epsilon \quad \forall \beta_i, \forall i \in [N] \end{aligned} \quad (1)$$

As mentioned in the previous section, we can numerically verify that NPGA has found an ϵ -BNE if it converged. The following metric is the relative ex-ante utility loss, which quantifies the relative error of the approximated BNE to the best response. We found an exact BNE if the best response corresponds to the found strategy profile.

Definition 2. (Bichler et al. 2021) For a batch of type samples v_i with size n_{batch} , the *relative ex-ante utility loss* $\hat{\mathcal{L}}_i(\beta_i)$ for player i is given by

$$\hat{\mathcal{L}}_i(\beta_i) = 1 - \frac{\hat{u}_i(\beta)}{\hat{u}_i(\beta) + \hat{l}_i(\beta)}, \quad (2)$$

where $\hat{u}_i(\beta)$ is the approximated utility, given the current strategy profile β , and $\hat{l}_i(\beta)$ is an estimator for $\epsilon = \max_i \lambda_i$ and given by

$$\hat{l}_i(\beta) = \frac{1}{n_{batch}} \sum_{v_i} \max_w \lambda_i(v_i, b_w, \beta). \quad (3)$$

λ_i is the expected gain in utility of bidder i if they deviate from their current strategy β_i to the alternative action b' , where the valuations of the remaining player are conditionally dependent on v_i :

$$\lambda_i(v_i, b', \beta) = \frac{1}{n_{batch}} \sum_{v_{-i} | v_i} (u_i(v_i, b', \beta_{-i}(v_{-i})) - u_i(v_i, \beta_i(v_i), \beta_{-i}(v_{-i}))) \quad (4)$$

Identificaton Structural estimation techniques in auctions seek to recover the primitives that lead to observed bidding behavior. Thus, we distinguish between the set of observables O , which includes all information that is directly available, and the latent structure S , which encompasses all information inaccessible to the auctioneer. For example, an auctioneer has access to a dataset $D = \{(b_{i,t}), [N]\}_{t=1}^T$ that records the bids submitted by the player i at time t and the set of players $[N]$; this dataset is part of the observables O . In contrast, the corresponding valuations, expressed by the prior distribution, constitute the latent structure $S = (F)$. Whether such a structure can be inferred solely from the observables using a structural estimation technique depends on its *identifiability*.

Definition 3. A latent structure S can be **identified** by an observable O if there does not exist a different structure S' that induces the same observables under the rules defined in the underlying game Γ :

$$S \xrightarrow{\Gamma} O_S \neq O_{S'} \xleftarrow{\Gamma} S' \quad \forall S' \neq S \quad (5)$$

We can tailor this general definition to the example mentioned above. If we are only interested in recovering the prior distribution F , the set of observables is given by the empirical distribution of the submitted bids G with density g and the number of participants N .² Thus, identification is given if there does not exist a different distribution F' with density f' leading to the same bid distribution. Since we assume monotonicity of the equilibrium bid functions, we can express the density of the bids g in terms of the value density f and the bid function β using the change-of-variables formula. This leads to the following definition of identification:

$$g(b) = f(\beta_f^{-1}(b)) \left| \frac{d}{db} \beta_f^{-1}(b) \right| \neq f'(\beta_{f'}^{-1}(b)) \left| \frac{d}{db} \beta_{f'}^{-1}(b) \right| = g'(b) \quad \forall f' \neq f \quad (6)$$

Assuming that identification holds, this definition implies that it is sufficient to find a prior model f that induces an equilibrium bid distribution that matches the observed bids under the given auction rules. However, it remains necessary to validate this assumption for specific settings.

In the next section, we propose a new estimation framework that recovers latent structures from observables, given that the identification assumption is met. We quantify the similarity between two probability measures using the Wasserstein- p distance.

Definition 4. (Villani et al. 2009) For two probability measures μ and ν , the **Wasserstein- p distance** W_p on a Polish space (\mathcal{X}, d) with $p \in [1, \infty)$ is given by

$$W_p(\mu, \nu) = \left(\inf_{J \in \mathcal{J}(\mu, \nu)} \int_{\mathcal{X}} d(x, y)^p dJ(x, y) \right)^{\frac{1}{p}} \quad (7)$$

Intuitively speaking, this metric determines the minimal costs to solve the underlying optimal transport problem between the two measures, i.e., it finds the optimal coupling

² We assume that agents are risk-neutral. If this assumption is not applicable, then the utility parameters must also be considered either in the set of observables, if they are known, or in the latent structure. The support of the type and bid distributions implicitly defines the type and action spaces.

J of the set of all possible couplings \mathcal{J} , which determines how much mass from μ has to be moved to recover ν . There exists a closed form solution to this problem if μ and ν are univariate probability distributions P and Q , and d is given by the p -norm (Vallender 1974):

$$W_p(P, Q) = \left(\int_0^1 |P^{-1}(z) - Q^{-1}(z)|^p dz \right)^{\frac{1}{p}} \quad (8)$$

P^{-1} and Q^{-1} are the quantile functions of the distributions P and Q . In our numerical experiments, we leverage the Wasserstein-2 distance.

3.2 Recovering Latent Distributions from Bidding Data

This section introduces the estimation framework for recovering a latent structure S from the set of observables O in single-item auctions. In an auction, bidders face the following equilibrium problem. Given a prior distribution F^* , bidders select their bids according to the equilibrium bid function β^* , inducing a bid distribution G^* with density g^* . While the realization of F^* is common knowledge among the bidders, it is inaccessible to the auctioneer. Thus, we assume that O encompasses all information characterizing the bidding behavior, except for the prior distribution F^* , which constitutes the latent structure ($S = (F)$). Furthermore, we assume that bidders are ex-ante symmetric. Thus, the observables are defined as $O = (N, M, D)$, N is the number of bidders, M is the underlying mechanism, and D is a dataset that contains the observed bids from all participants.

Since the prior distribution F^* is unknown, the corresponding bid function β^* is also inaccessible to the auctioneer. Consequently, it is infeasible to directly compute the value distribution from the bid distribution. Instead, Guerre et al. (2000) leveraged the first-order condition of the expected utility to express valuations as a function ζ , which depends on the observed bids b , number of bidders N , the density g^* , and distribution G^* of bids in the FPSB auction:

$$v = \zeta(b, N, g^*, G^*) = b + \frac{1}{N-1} \frac{G^*(b)}{g^*(b)} \quad (9)$$

The dataset only contains a finite number of bids, and thus, it is necessary to estimate G^* and g^* . Instead of using kernel methods, we propose using a proxy equilibrium model to estimate those.

Definition 5. *Let F be an atomless distribution with density f and let β denote the BNE strategy under F , which maps to a bid distribution G with density g . Notably, this **proxy equilibrium** is entirely independent of the data-generating model.*

Under the monotonicity assumption of the bid function, the BNE strategy acts as a push-forward measure from the value distribution to the bid distribution, which we denote as $G = \beta_{\#F}$ in the following. This formulation allows us to evaluate G and g using the inverse bid function β^{-1} and the prior distribution F in this proxy model.

Given the distribution of the bids in the proxy model, we can compute the optimal transport map T , which defines a push-forward measure from the proxy model to the observed bid distribution, i.e., $G^* = T_{\#G}$. Since we focus on univariate distributions,

one transportation plan is given by $T(b) = (G^{*-1}(G(b)))$ (Villani et al. 2009). As before, this formulation allows us to express the density g^* and distribution G^* of observed bids in terms of the proxy model’s density g and distribution G . Furthermore, we combine both transport maps into a single transformation, enabling us to directly evaluate g^* and G^* from f and F :

$$G^* = T_{\#G} = T_{\#\beta_{\#F}} \quad (10)$$

We leverage this relation to compute pseudo-values for each bid in the datasets using Equations (9). This, in turn, allows us to define another transport map T_v between the prior in the proxy model and the data-generating value distributions expressed by the pseudo-values, i.e., $T_v(v) = (F^{*-1}(F(v)))$. This mapping enables the computation of comparative statistics, like the optimal reserve price in an auction. In our numerical experiments, we employ neural networks to smoothly interpolate all transportation maps.³

4 Numerical Analysis

We apply our estimation framework to laboratory datasets, where the prior distribution is defined by the underlying experimental design. This setup allows us to evaluate the performance of our approach (OTGPV) and compare it to the GPV framework and its boundary corrected variant introduced by Hickman & Hubbard (2015) (BCGPV). Performance is assessed along two dimensions. First, we compare the Wasserstein distance between the experimental prior distribution and the distributions produced by the estimation methods. Second, knowledge of the prior distribution enables the auctioneer to compute the optimal reserve price to maximize their expected revenue. Accordingly, we evaluate how the estimation affects the calculation of the reserve price and the resulting revenue of the auctioneer. The settings consist of FPSB auctions, in which the prior distribution is non-parametrically identified from the bid distribution (Guerre et al. 2000).

We base our numerical analysis on two datasets kindly provided by the original authors, which we briefly describe below. In all experiments, the prior distribution was a Uniform distribution. For consistency, we scaled the support of these distributions, along with their corresponding bids, to the interval $[0, 1]$ without loss of generality.

In our proxy model, we assume a Uniform distribution over $[0, 0.5]$. This ensures numerical stability even with small sample sizes. We tested other models like Beta and Gaussian distributions, but due to the sparse observations near their support boundaries, they yielded slightly poorer results. Increasing the sample size would improve performance.

Data by Ockenfels & Selten (2005) The first dataset contains the two-player FPSB auction data generated by Ockenfels & Selten (2005) (Ock). The experiment included two treatments: a *No-Feedback* (NF) treatment, where bidders only received information about their own bids, and a *Feedback* (F) treatment, where all bids were made public. There were eight groups of six people per treatment, with each participant taking part in 140 auctions under random matching within the groups. Bidding was restricted to prevent participants from submitting bids exceeding their valuations.

³ We tested various network structures, and the approximation of these functions are robust. This indicates that optimizing the networks is less extensive than for kernel estimators.

Table 1. The Wasserstein-2 distance between the estimated and true value distributions for the different approaches. The rows under the OTGPV column indicate whether we used the analytical equilibrium or NPGA.

Dataset	GPV	BCGPV	OTGPV	
			Analytical	NPGA
Kag	0.08	0.05	0.04	0.06
Ock (F)	2.74	1.69	0.10	0.18
Ock (NF)	2.60	2.37	0.13	0.17

Data by Kagel & Levin (1993) The second dataset describes the bidding data of the *cross-over* treatment in the experiment conducted by Kagel & Levin (1993) (Kag). In this experiment, two groups of five bidders first participated in two five-player FPSB auctions, followed by a single ten-player auction. Our dataset includes only the five-player auction data, so we focus exclusively on this setting. The bidding behavior was restricted by an initial budget provided to the participants, which they could use throughout all auctions. All bid-value pairs were announced after each auction round. Unfortunately, the dataset does not specify how many subjects participated in this experiment.

4.1 Results

We present the results of our numerical analysis in the following subsection. Thereby, we first provide an overview of the identification results. Subsequently, we discuss the implications for setting the optimal reserve price based on these estimates.

Table 1 compares the Wasserstein-2 distances of the estimated distributions to the data-generating distribution. While all methods achieve low Wasserstein distances in the Kag dataset, the kernel-based estimators yield significantly higher distances in the Ock dataset, implying that kernel based methods not only suffer from problems at the supports boundaries but are also prone to biased estimates due to the high variance. This suggests that our approach remains robust even in scenarios where participants exhibit higher variance in their bidding behavior. The table further indicates no significant difference in using the analytical bid function and NPGA to determine the equilibrium bid function. The small discrepancies are likely due to numerical instabilities in NPGA, which could be mitigated by increasing the sample size and the number of learning iterations. Notably, the utility losses were low in all experiments (smaller than 0.02).

Table 2 reports the resulting relative gain in expected revenue for the auctioneer when we compute the optimal reserve price using the approximated distribution. The results reflect the accuracy of the recovered distributions: While in the Kag dataset all approaches yield a reserve that does not increase the expected revenue significantly, the kernel-based estimators fail to determine an effective price that increases the expected revenue in the Ock datasets. In fact, it even leads to a reserve price that refrains all bidders from participating in the auction, leading to an expected revenue of 0.00. Conversely, our estimator yields a reserve price that significantly increases the auctioneers' expected revenue.

Table 2. Relative expected revenue gains from setting the optimal reserve price based on the estimated distributions. The rows under the OTGPV column indicate whether we used the analytical equilibrium or NPGA.

Dataset	GPV	BCGPV	OTGPV	
			Analytical	NPGA
Kag	1.01	1.01	1.00	1.00
Ock (F)	0.00	0.00	1.18	1.18
Ock (NF)	0.00	0.00	1.20	1.20

5 Conclusion

In this study, we proposed an extension to the two-step estimation framework by Guerre et al. (2000) for identifying the prior distribution from observables in auctions under the identification assumption. The core idea is to eliminate reliance on kernel estimators to express the density and distribution of the bids. Instead, we determine the optimal transport maps between the observed bid distribution and a proxy equilibrium model, which enables us to analytically evaluate their density and distribution. We employ equilibrium learning techniques to approximate BNE strategies in settings where it is infeasible to derive the equilibrium analytically. Our numerical analysis demonstrates that these design choices significantly improve the applicability compared to kernel-based approaches.

Despite these promising results, our study has two key limitations. First, our numerical analysis is based on two experimental datasets that describe FPSB auctions, where the prior distribution is uniform. Since we were not able to acquire more realistic datasets, future research should investigate the quality of our estimation framework in other datasets based on field data rather than only laboratory bidding behavior. Second, although laboratory experiments in auctions suggest that bidders tend to overbid the predicted risk-neutral equilibrium strategy, we assumed that bidders are risk-neutral. This could explain the higher variance in bidding behavior observed in the dataset by Ockenfels & Selten (2005). Future research should investigate the effects of setting a reserve price by considering behavioral motives like risk aversion. However, this would require more information about the risk preferences and the utilities of the bidders.

As the number of studies assessing structural estimation techniques in laboratory environments is very limited, our findings offer valuable insights for auction designers and policymakers. By eliminating the reliance on kernel estimators, our approach provides a more robust and flexible framework for estimating bidders’ private values. This has direct implications for auction design. For instance, improved estimation of the prior distribution allows auctioneers to set more accurate reserve prices, reducing the risk of inefficient allocations or lost revenue due to mispricing.

6 Acknowledgements

Bichler and Ewert gratefully acknowledge the financial support from the Deutsche Forschungsgemeinschaft (DFG) (grant no. BI 1057/1-10). Bichler is member of the Munich Data Science Institute (MDSI).

References

- Abreu, D. & Rubinstein, A. (1988), 'The structure of nash equilibrium in repeated games with finite automata', *Econometrica: Journal of the Econometric Society* pp. 1259–1281.
- Ahunbay, M. Ş. & Bichler, M. (2025), On the uniqueness of bayesian coarse correlated equilibria in standard first-price and all-pay auctions, in 'Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)', SIAM, pp. 2491–2537.
- Bichler, M., Fichtl, M., Heidekrüger, S., Kohring, N. & Sutterer, P. (2021), 'Learning equilibria in symmetric auction games using artificial neural networks', *Nature machine intelligence* **3**(8), 687–695.
- Bichler, M., Fichtl, M. & Oberlechner, M. (2023), 'Computing bayes–nash equilibrium strategies in auction games via simultaneous online dual averaging', *Operations Research*.
- Bichler, M., Gupta, A., Mathews, L. & Oberlechner, M. (2023), 'Low revenue in display ad auctions: Algorithmic collusion vs. non-quasilinear preferences', *arXiv preprint arXiv:2312.00243*.
- Bichler, M., Kohring, N. & Heidekrüger, S. (2023), 'Learning equilibria in asymmetric auction games', *INFORMS Journal on Computing* **35**(3), 523–542.
- Bichler, M., Kohring, N., Oberlechner, M. & Pieroth, F. R. (2023), 'Learning equilibrium in bilateral bargaining games', *European Journal of Operational Research* **311**(2), 660–678.
- Bowling, M. (2004), 'Convergence and no-regret in multiagent learning', *Advances in neural information processing systems* **17**.
- Cai, Y. & Papadimitriou, C. (2014), Simultaneous bayesian auctions and computational complexity, in 'Proceedings of the Fifteenth ACM Conference on Economics and Computation', ACM, Palo Alto California USA, pp. 895–910.
- Campo, S., Guerre, E., Perrigne, I. & Vuong, Q. (2011), 'Semiparametric estimation of first-price auctions with risk-averse bidders', *The Review of Economic Studies* **78**(1), 112–147.
- Cox, J. C., Smith, V. L. & Walker, J. M. (1985), 'Experimental development of sealed-bid auction theory; calibrating controls for risk aversion', *The American Economic Review* **75**(2), 160–165.
- Cox, J. C., Smith, V. L. & Walker, J. M. (1988), 'Theory and individual behavior of first-price auctions', *Journal of Risk and uncertainty* **1**, 61–99.
- Donald, S. G. & Paarsch, H. J. (1996), 'Identification, Estimation, and Testing in Parametric Empirical Models of Auctions within the Independent Private Values Paradigm', *Econometric Theory* **12**(3), 517–567.
- Flambard, V. & Perrigne, I. (2006), 'Asymmetry in procurement auctions: Evidence from snow removal contracts', *The Economic Journal* **116**(514), 1014–1036.
- Fudenberg, D. & Levine, D. K. (1998), *The theory of learning in games*, Vol. 2, MIT press.

- Guerre, E., Perrigne, I. & Vuong, Q. (2000), 'Optimal nonparametric estimation of first-price auctions', *Econometrica : journal of the Econometric Society* **68**(3), 525–574.
- Härdle, W. (2004), *Nonparametric and semiparametric models*, Springer Science & Business Media.
- Hart, S. & Mas-Colell, A. (2000), 'A simple adaptive procedure leading to correlated equilibrium', *Econometrica* **68**(5), 1127–1150.
- Hickman, B. R. & Hubbard, T. P. (2015), 'Replacing sample trimming with boundary correction in nonparametric estimation of first-price auctions', *Journal of Applied Econometrics* **30**(5), 739–762.
- Hickman, B. R., Hubbard, T. P. & Sağlam, Y. (2012), 'Structural econometric methods in auctions: A guide to the literature', *Journal of Econometric Methods* **1**(1), 67–106.
- Kagel, J. H. & Levin, D. (1993), 'Independent Private Value Auctions: Bidder Behaviour in First-, Second- and Third-Price Auctions with Varying Numbers of Bidders', *The Economic Journal* **103**(419), 868.
- Krishna, V. (2009), *Auction theory*, Academic press.
- Li, T. & Perrigne, I. (2003), 'Timber sale auctions with random reserve prices', *Review of Economics and Statistics* **85**(1), 189–200.
- Li, T., Perrigne, I. & Vuong, Q. (2002), 'Structural estimation of the affiliated private value auction model', *RAND Journal of Economics* pp. 171–193.
- Liu, T. X., Yang, J., Adamic, L. A. & Chen, Y. (2014), 'Crowdsourcing with all-pay auctions: A field experiment on taskcn', *Management Science* **60**(8), 2020–2037.
- Lu, J. & Perrigne, I. (2008), 'Estimating risk aversion from ascending and sealed-bid auctions: The case of timber auction data', *Journal of Applied Econometrics* **23**(7), 871–896.
- Manski, C. F., McFadden, D. et al. (1981), *Structural analysis of discrete data with econometric applications*, MIT press Cambridge, MA.
- Ockenfels, A. & Selten, R. (2005), 'Impulse balance equilibrium and feedback in first price auctions', *Games and Economic Behavior* **51**(1), 155–170.
- Ostrovsky, M. & Schwarz, M. (2023), 'Reserve prices in internet advertising auctions: A field experiment', *Journal of Political Economy* **131**(12), 3352–3376.
- Shakhgildyan, K. (2022), 'Nonparametric identification and estimation of all-pay auction and contest models', *Review of Economic Design* pp. 1–39.
- Vallender, S. (1974), 'Calculation of the wasserstein distance between probability distributions on the line', *Theory of Probability & Its Applications* **18**(4), 784–786.
- Villani, C. et al. (2009), *Optimal transport: old and new*, Vol. 338, Springer.
- Wilson, R. (1985), *Game-theoretic analyses of trading processes*, Institute for Mathematical Studies in the Social Sciences, Stanford University.
- Zhang, S., Karunamuni, R. J. & Jones, M. C. (1999), 'An improved estimator of the density function at the boundary', *Journal of the American Statistical Association* **94**(448), 1231–1240.