LEARNING ADVERTISER VALUES FROM STRATEGIC BIDS IN PLATFORM POSITION AUCTIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Digital platforms allocate scarce user attention through position auctions, yet the advertiser values needed for platform design are not directly observed. In generalized second-price (GSP) auctions, bids are strategic and therefore cannot be interpreted as values. This poses a structural estimation problem: how can a platform recover latent advertiser value distributions from observed bids when analytical bid-to-value inversion is mechanism-specific or altogether unavailable? We propose a simulation-based estimator for non-truthful platform auctions. The method represents advertiser value distributions with flexible Bernstein-polynomial CDFs, computes approximate equilibrium bidding strategies for candidate primitives, and selects the primitives whose simulated bids match the observed bids in Wasserstein distance. A Bayesian-optimization-assisted SMC-ABC procedure guides the search while limiting the number of costly equilibrium computations. We validate the approach on benchmark auction simulations and on GSP laboratory data, and use the recovered value distributions to evaluate counterfactual platform-design choices such as optimal reserve prices. The framework offers a practical route to data-driven market design in digital platform auctions where standard inversion-based estimators are difficult to apply.

1 INTRODUCTION

Digital platforms increasingly rely on auction mechanisms to allocate scarce attention, visibility, and access to users. Sponsored search advertising is a canonical example: advertisers compete for ranked positions on a search results page, where higher positions typically generate more clicks. The generalized second-price (GSP) auction has been one of the central mechanisms used in these markets (Edelman et al., 2007; Varian, 2007). In such auctions, the platform observes bids, positions, clicks, and payments, but not advertisers' private values. These values, however, are a key primitive for platform design. Without estimates of advertiser values, the platform cannot systematically evaluate reserve prices, ranking rules, slot design, or alternative auction formats.

The empirical challenge is that GSP auctions are not truthful. Unlike VCG mechanisms, GSP does not generally make truthful bidding a dominant strategy (Edelman et al., 2007; Varian, 2007). Observed bids, therefore, cannot be interpreted directly as advertiser values. Instead, bids are strategic choices that depend on the auction rules, the distribution of competing bids, position-specific click-through rates (CTRs), and possibly reserve prices or quality scores. Recovering private values from bid data is thus a structural estimation problem: the analyst observes equilibrium outcomes, but the primitives that generated those outcomes are latent.

Classical structural auction estimation has made substantial progress in settings such as first-price sealed-bid auctions. The influential two-step approach of Guerre et al. (2000) (henceforth GPV) first estimates the bid distribution and then uses the first-order condition of the equilibrium bidding problem to invert observed bids into pseudo-values. A key attraction of this approach is that it does not require solving for the equilibrium bid function. However, its implementation relies on an analytical inversion formula that is specific to the auction format. While this inversion is particularly tractable in first-price auctions, it is generally infeasible in platform auctions.

In GSP-style platform auctions, the allocation of slots is rarely determined by raw bids alone. Instead, platforms rank participants using complex, proprietary quality scores calibrated to specific bidders, keywords, and advertisements. As established by Laffont et al. (2020) in the context of

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

procurement auctions, when such scoring rules are generalized or unobserved, the underlying economic primitives become fundamentally unidentified. While analytical inversion of equilibrium bidding strategies remains possible under specific scoring rules in stylized GSP environments (Kim & Pal, 2025), such identification results can be highly sensitive to the underlying model specifications. Thus, these tightly tailored inversions frequently become mathematically intractable or computationally prohibitive when applied to the nuanced architectures observed in the field.

This paper develops a simulation-based structural estimator for non-truthful platform position auctions. Instead of deriving a new bid-to-value inversion for each mechanism variant, we use the auction rules directly. For a candidate private-value distribution, we compute an approximate Bayes–Nash equilibrium (BNE), simulate bids from that equilibrium, and compare the simulated bid distribution with the observed bid distribution. The estimated value distribution is the one whose equilibrium bids best match the data. This shifts the empirical task from analytical inversion to treating equilibrium computation as a forward model for simulation-based inference and distributional matching.

The proposed framework has three main components. First, we represent latent advertiser value distributions using flexible Bernstein polynomial specifications. These specifications impose the monotonicity and boundary conditions required of cumulative distribution functions (CDFs) while allowing rich shapes. Second, we use an equilibrium-learning oracle to approximate BNE bidding strategies for candidate value distributions. Recent work has shown that learning-based and online-optimization-based methods can compute equilibria in auction games that are difficult to solve analytically (Bichler et al., 2021; 2023). Third, we adopt a Bayesian approach to the estimation problem. This allows us not only to recover the data-generating distribution but also to quantify estimation uncertainty. Because the likelihood of observed bids is unavailable in closed form, we estimate the model using a likelihood-free approximate Bayesian computation (ABC) procedure. Candidate distributions are evaluated by the Wasserstein distance between simulated and observed bid distributions, and a Bayesian-optimization surrogate reduces the number of expensive equilibrium computations (Toni et al., 2009; Brochu et al., 2010; Bernton et al., 2019a).

In contrast to standard approaches based on inverting equilibrium conditions, which can be infeasible and require substantial analytical derivations, our method can be applied whenever the rules of the game can be implemented in code and equilibrium behavior can be computed. This makes it straightforward to accommodate changes in auction design without re-deriving a mechanism-specific inversion. Moreover, because our equilibrium computation is algorithmic, the same framework can be used to evaluate counterfactual market designs before deploying them in practice.

We validate the approach in GSP-style auction experiments. The experimental setting is useful because values are observed by the researcher but withheld from the estimator. This allows us to evaluate whether the method can recover the latent value distribution solely from bids. We then use the recovered value distributions to conduct platform-design counterfactuals, specifically reserve prices. These counterfactuals illustrate how the recovered economic primitives can be used for downstream platform-design decisions. Importantly, our objective is not to outperform specialized inversion-based estimators in settings where such estimators are available, but to provide a flexible structural estimation framework that remains applicable when changes in the auction mechanism make analytical inversion difficult or unavailable.

The paper makes four contributions. First, it frames the recovery of advertiser values in non-truthful GSP auctions as a structural estimation problem. Second, it provides a unified simulation-based alternative to mechanism-specific analytical inversion. Third, it uses equilibrium learning as a black-box forward model for likelihood-free structural inference, avoiding analytical inversion of the equilibrium mapping. Fourth, it shows how recovered primitives can inform counterfactual platform design. More broadly, equilibrium-learning-based structural estimation can support empirical market design when analytical inversion is impractical.

2 Related Literature

This paper contributes to four streams of literature: sponsored search and GSP auctions, structural estimation, equilibrium computation in auctions, and likelihood-free simulation-based estimation.

2.1 SPONSORED SEARCH AND GSP AUCTIONS

A large body of literature studies sponsored search auctions and the GSP mechanism. Edelman et al. (2007), Varian (2007), and Gomes & Sweeney (2014) provide foundational analyses of GSP position auctions and their equilibrium properties. A key insight from this literature is that GSP is not a truthful mechanism. Bidders generally have strategic incentives to shade or adjust bids, and equilibrium outcomes depend on the structure of position effects and payments. This makes GSP a natural environment for structural estimation: bids are observed, but private values are not.

Several papers study empirical and experimental behavior in sponsored search auctions. Börgers et al. (2013) analyze equilibrium bids in sponsored search auctions using theory and evidence. Experimental studies such as Che et al. (2017) and Bae & Kagel (2019) examine how subjects bid in GSP environments and how observed behavior relates to equilibrium predictions. These studies show that GSP is both practically important and behaviorally nontrivial. Our paper differs from this work by focusing on the inverse problem: given observed bids in a non-truthful position auction, how can a platform recover the latent distribution of advertiser values?

A related stream in information systems and machine learning studies sponsored search advertising and keyword auctions as digital platform mechanisms. Liu et al. (2010) analyze keyword auction design and show how weighting schemes and minimum-bid policies affect equilibrium matching and platform revenue. More recently, Kannan et al. (2023) studied GSP auction markets using Q-learning simulations and human-subject experiments, documenting how market frictions affect bidding behavior and allocative efficiency. Empirical research has also examined how sponsored search performance depends on ad rank, keyword characteristics, competing ads, organic search results, and consumer search intent (Ghose & Yang, 2009; Du et al., 2017; Agarwal et al., 2015; Agarwal & Mukhopadhyay, 2016; Bhattacharya et al., 2022). This literature establishes sponsored search as a central information-systems setting in which platform rules, ranking algorithms, advertiser bidding, and consumer attention jointly determine market outcomes. Our paper complements this work by addressing the inverse structural problem: recovering advertisers' latent value distributions from strategic bids in non-truthful platform auctions.

2.2 STRUCTURAL ESTIMATION OF AUCTIONS

Our work builds on the structural auction-estimation literature. The benchmark contribution is Guerre et al. (2000), who developed a two-step nonparametric estimator for first-price auctions. Their estimator uses the equilibrium first-order condition to transform observed bids into pseudo-values. This insight has been central to empirical auction work because it avoids repeatedly solving the auction game. Kim & Pal (2025) extend this result to a particular GSP auction, where the scoring rule is linear in the bids. General treatments of auction theory and auction design are provided by Krishna (2010) and Myerson (1981).

Kim and Pal also numerically compute equilibrium strategies for simulation and counterfactual analysis. The distinction is therefore not whether equilibrium computation is available, but how it is used. Their estimator relies on an analytically derived mapping from observed bids to latent values, whereas our estimator treats equilibrium computation as a forward model inside likelihood-free inference. Consequently, changing the auction mechanism in our framework requires modifying the simulator rather than re-deriving an identification/inversion formula.

The strength of GPV-style estimation is also its limitation for our setting. The bid-to-value inversion is derived from the first-order condition of a particular auction format. In GSP position auctions, the bidder's marginal incentive depends on movements across positions, differences in click-through rates, payments determined by neighboring bids or scores, reserve prices, and quality adjustments. As a result, analytical inversion becomes mechanism-specific. Our paper addresses this limitation by replacing closed-form inversion with an equilibrium-based simulation procedure. The estimator asks which candidate distribution, when combined with the auction rules and equilibrium behavior, produces a bid distribution that matches the data.

Our approach is related to simulation-based structural estimation methods such as simulated method of moments and indirect inference (McFadden, 1989; Gouriéroux et al., 1993). Like these, we compare observed data with data simulated from a structural model. The distinctive feature of our setting

is that each parameter evaluation requires solving a Bayesian game. This makes computational efficiency central and motivates the use of equilibrium learning and Bayesian optimization.

2.3 Equilibrium computation and learning in auction games

The proposed estimator relies on an equilibrium oracle to approximate the BNE bidding strategy for each candidate value distribution. Because analytical equilibria are rarely tractable outside of highly stylized environments, we leverage recent advancements in computational game theory to bridge this gap. Specifically, we utilize algorithmic approaches, such as neural-network-based learning (Bichler et al., 2021) and simultaneous online dual averaging (SODA) (Bichler et al., 2023), to efficiently approximate these equilibrium strategies. In our experiments, we leverage SODA as it converges to an approximate equilibrium within seconds.

We use this literature as a computational foundation for structural estimation. Equilibrium learning allows us to evaluate candidate primitives even when a closed-form equilibrium is unavailable. In this respect, our paper connects equilibrium computation with empirical market design: the equilibrium oracle is not the final object of interest, but a component inside an estimator that recovers latent economic primitives from observed platform data.

2.4 Likelihood-free inference and Wasserstein matching

Because the likelihood of observed bids is generally unavailable in closed form, we use likelihood-free inference. Specifically, ABC methods replace exact likelihood evaluation with simulation and comparison between observed and simulated data (Toni et al., 2009). We adapt this logic to auction estimation, where the intractability comes from the equilibrium mapping from values to bids. The discrepancy between observed and simulated bid distributions is measured using the Wasserstein distance. This metric is attractive because it compares entire distributions and is well-suited to simulation-based parameter estimation (Bernton et al., 2019a;b). We combine these ideas with Bayesian optimization for expensive black-box functions (Brochu et al., 2010). The Bayesian-optimization surrogate screens candidate parameters before invoking the computationally expensive equilibrium oracle. Finally, initial exploration of the parameter space uses space-filling designs (Mckay et al., 2000) and simplex transformations (Pillards & Cools, 2005).

3 GSP Auctions

We consider a platform that allocates K advertising positions to N bidders, where $K \leq N$. Positions differ in their CTRs, $\alpha_1 > \alpha_2 > \cdots > \alpha_K > 0$. Bidder i has a private value per click v_i, drawn independently from a distribution F_i, and submits a per-click bid b_i. In the simplest rank-by-bid version of the GSP auction, bidders are ordered by their bids. In many search implementations, bidders are instead ranked by quality-adjusted scores, $s_i = q_i b_i$, where q_i is a platform-assigned quality or relevance score. Such quality-adjusted ranking is central in sponsored search auctions because the platform cares about expected revenue and user relevance, not only about the advertiser's bid (Edelman et al., 2007; Varian, 2007). If bidder i receives position k, its expected payoff is $u_i = \alpha_k(v_i - p_i)$, where p_i is the per-click payment determined by the next-ranked bidder's bid or score. The exact payment rule depends on the institutional environment. The key property of GSP is that truthful bidding is not generally a dominant strategy. Therefore, recovering F_i from observed bids requires a model of bidding behavior.

This raises the question of why the analytical inversion underlying GPV-style structural estimation is difficult to extend to GSP auctions. In these settings, the marginal incentive to change a bid depends on the probability of moving between adjacent positions, the CTRs of those positions, the payment rule, quality scoring rules, and reserve prices. Thus, analytical inversion is highly mechanism-specific. While special stylized cases may admit an inversion, the derivation becomes cumbersome or unavailable in platform-relevant variants.

Since GSP auctions can admit multiple equilibria, we restrict attention throughout to the symmetric BNE in strictly increasing strategies. Under incomplete information, this monotone symmetric equilibrium is unique within its class (Gomes & Sweeney, 2014). This restriction is a maintained

assumption rather than a property we derive as for some primitives, such an equilibrium could not exist. In all experiments, the recovered strategies were consistent with this restriction.

4 Estimation Framework

Our approach replaces mechanism-specific analytical inversion with simulation-based estimation. Let F_θ denote a candidate distribution of advertisers' private values, parameterized by θ. For each F_θ, the GSP game described above induces an equilibrium bidding strategy profile β_θ and an associated equilibrium bid distribution G_θ. The structural estimation problem is to recover the latent value distribution from the observed bid distribution G_{obs}. The model therefore defines the mapping

$$F_\theta \longmapsto \beta_\theta \longmapsto G_\theta.$$

Rather than deriving a bid-to-value inversion formula, we search for the parameter θ whose model-implied bid distribution best matches the observed bid distribution.

The estimator relies on four theoretical observations. We state them here informally:

Result 1: Bids are induced by values and equilibrium strategies. If bidder i's equilibrium bidding strategy β_i is strictly increasing, then the bid distribution is induced by the value distribution through the equilibrium bidding strategy. In particular, at $b_i = \beta_i(v_i), G_i(b_i) = F_i(v_i)$. If β_i is differentiable, the corresponding bid density satisfies $g_i(b_i) = {}^{f_i(v_i)}\!/_{\beta_i'(v_i)}$. Thus, although bids are not equal to values in a non-truthful mechanism such as GSP, the observed bid distribution contains information about the latent value distribution through the equilibrium.

Result 2: Identification condition within the maintained family. The estimator searches over a parametric family of value distributions $\mathcal{F}_\Theta = \{F_\theta : \theta \in \Theta\}$. For each candidate parameter θ, the auction model Γ and equilibrium oracle imply a bid distribution $G_\Gamma(F_\theta)$. Identification requires that two different candidate value distributions do not generate the same equilibrium bid distribution. Formally, if $\theta_1 \neq \theta_2$, then we require

$$G_\Gamma(F_{\theta_1}) \neq G_\Gamma(F_{\theta_2}).$$

Under this condition, the population bid distribution uniquely pins down the parameter θ within the chosen family. This is a parametric, model-based identification statement: it says that θ can be recovered under the maintained auction model and parameterization. It does not imply that arbitrary value distributions can always be recovered from any GSP dataset, especially when bids are censored, participation is endogenous, quality scores are unobserved, or equilibrium selection is ambiguous.

Result 3: Bernstein CDFs provide a flexible value-distribution family. We represent F_θ using Bernstein-polynomial CDFs. For bidder i, the CDF is written as

$$F_{i,\theta}(v) = \sum_{j=0}^{m} \theta_{ij} \binom{m}{j} v^j (1-v)^{m-j}, \qquad v \in [0,1],$$

with coefficients satisfying

$$0 = \theta_{i0} \leq \theta_{i1} \leq \cdots \leq \theta_{im} = 1.$$

These restrictions ensure that $F_{i,\theta}$ is a valid CDF over the unit interval. Bernstein CDFs with monotone coefficients can approximate any continuous CDF on bounded support arbitrarily well as the degree m increases. This provides a flexible finite-dimensional representation of latent value distributions while maintaining the shape restrictions required of a distribution function.

Result 4: Population target and approximation error. The minimum-distance objective has a well-defined population target. Let G_0 denote the population bid distribution generated by the true value distribution F_{θ_0} under the maintained auction model. If the model is correctly specified and θ_0 is identified within $\mathcal{F}_\Theta$, then $W_p\left(G_{\theta_0}, G_0\right) = 0$ and, for any $\theta \neq \theta_0$, $W_p\left(G_\theta, G_0\right) > 0$.

Thus, in the population problem, the true parameter θ_0 is the unique minimizer of $Q(\theta) = W_p\left(G_\theta, G_0\right)$. In practice, the estimator replaces G_0 by the empirical bid distribution G_{obs}, and

replaces G_θ by a simulated bid distribution generated from an approximate equilibrium. The implemented criterion therefore approximates $Q(\theta) = W_p(G_\theta, G_0)$ with a finite-sample, simulation-based criterion. The resulting estimation error can be interpreted as coming from four sources: sampling error in G_{obs}, simulation error in the approximation of G_θ, equilibrium-oracle error in the computation of β_θ, and ABC tolerance error from accepting parameter values that are close to, but not exactly at, the minimum. The simulations and experiments below evaluate the magnitude of these errors in settings where the true values are known.

For each candidate θ, we construct F_θ, compute approximate equilibrium bidding behavior, simulate bids, and compare the simulated bid distribution G_θ to the observed bid distribution G_{obs}. The estimator solves the minimum-distance problem:

$$\hat{\theta} = \arg\min_{\theta \in \Theta} W_p(G_\theta, G_{\text{obs}}),$$

where W_p is a Wasserstein distance. For univariate marginal bid distributions, this distance can be computed from quantile functions:

$$W_p(G_\theta, G_{\text{obs}}) = \left(\int_0^1 \left| G_\theta^{-1}(t) - G_{\text{obs}}^{-1}(t) \right|^p dt \right)^{1/p}.$$

The discrepancy $W_p(G_\theta, G_{\text{obs}})$ guides the search over value distributions. Parameter values that generate bids close to the observed data receive higher weight and are used to propose new candidates, whereas parameter values that generate poor matches are less likely to be retained. In our implementation, the search is carried out using a Bayesian-optimization-assisted sequential Monte Carlo (BOSMC) ABC procedure. Sequential Monte Carlo concentrates probability mass on parameter regions that fit the observed bids well, while the Bayesian-optimization surrogate reduces the number of expensive equilibrium computations.

Algorithm 1 BOSMC: Likelihood-Free Structural Estimation for GSP Auctions

Input: observed bids D, auction rules Γ, Bernstein degree m, prior over θ, tolerances $\{\epsilon_\ell\}$
for generation $\ell = 1, \ldots, L$ **do**
    Propose candidate value-distribution parameters θ
    Use the surrogate model to screen candidates
    for each retained candidate **do**
        Construct F_θ
        Compute approximate equilibrium bidding strategy β_θ
        Simulate the bid distribution G_θ
        Compute $d_\theta = W_p(G_\theta, G_{\text{obs}})$
        if $d_\theta \leq \epsilon_\ell$ **then**
            Accept θ into the particle population
        end if
    end for
    Update weights, proposal distribution, tolerance, and surrogate
end for
Output: estimated value distribution $F_{\hat{\theta}}$ and accepted particles

Algorithm 1 provides a simulation-based approximation to the population minimum-distance estimator. It should not be interpreted as providing an unconditional convergence guarantee. Such a guarantee would require additional assumptions: the value-distribution parameter must be identified, the empirical bid distribution must converge to its population counterpart, the equilibrium oracle must be exact or asymptotically accurate, the number of simulated bids must grow, the ABC tolerance must shrink, and the proposal distribution must continue to place mass near the population minimizer. In the present paper, we therefore treat BOSMC-ABC as a practical estimator and evaluate its reliability through benchmark simulations, GSP laboratory data, and counterfactual validation.

5 EXPERIMENTAL EVALUATION

We evaluate the estimator in three steps. First, we conduct benchmark simulations in GSP environments where the true data-generating value distribution is known and an established structural estimator is available for comparison. These simulations serve as a rigorous baseline for validating the estimation procedure. Second, we apply the method to a controlled GSP laboratory experiment. This laboratory setting is particularly advantageous because the researcher observes bidders' true private values but intentionally withholds them from the estimator, allowing us to directly assess whether the method successfully recovers latent values from bids generated by human subjects exhibiting potentially noisy equilibrium strategic behavior. Third, we apply the estimator to a synthetic dataset calibrated to mirror the field experiment by Ostrovsky & Schwarz (2023). Using the recovered distributions from this environment, we perform a counterfactual market design analysis to determine optimal reserve prices and compare the resulting expected platform revenues.

5.1 BENCHMARK SIMULATIONS

In the first step, we validate the estimator in controlled auction environments where the true data-generating value distribution is known. We consider a GSP auction with a linear scoring rule ($s_i = q_i b_i$) and vary the number of bidders $N \in \{3, \dots, 7\}$ across uniform and beta distributions. For each configuration, we generate 100 random draws of the underlying distributional parameters with 2000 bid samples per bidder; notably, for the beta distributions, we also randomize the scale factor to evaluate the estimators across varying supports. Subsequently, we utilize the distributions recovered by both our proposed approach (BOSMC) and the benchmark estimator (KimPal) to compute and compare the revenue-maximizing reserve prices for each auction setting.

Table 1 summarizes the average Wasserstein-2 (W_2) distances and their standard deviations for the individual prior distributions. While both approaches yield low Wasserstein distances, implying that both can recover the data-generating prior well, our approach achieves, on average, a lower Wasserstein distance than the benchmark model by Kim & Pal (2025) across both prior models. A likely explanation for this gap is that GPV-like estimators rely on kernel density estimators, which are known to be prone to errors near the boundaries of distributions with compact support (Hickman & Hubbard, 2015).

Table 1: Estimation performance in benchmark simulations. Average normalized W_2 distances between estimated and true value distributions; standard deviations are reported in parentheses.

Prior	BOSMC W_2	KimPal W_2
Beta	0.067 (0.093)	0.104 (0.154)
Uniform	0.030 (0.011)	0.050 (0.033)

In addition to estimation performance, we compare the expected revenue induced by the optimal reserve prices under the estimated prior distributions with the revenue achieved by the optimal reserve price under the true data-generating distribution (DGP). Table 2 shows the average expected revenues for both experiments. Similar to the estimation quality, the expected revenues are close to the true optimal value across both estimators.

Table 2: Expected platform revenues in benchmark simulations using reserve prices optimized from the estimated value distributions, compared with the theoretical optimum under the true DGP.

Prior	BOSMC Rev.	KimPal Rev.	DGP Rev.
Beta	0.861	0.874	0.875
Uniform	0.748	0.748	0.748

However, an interesting divergence occurs for the Beta distribution: while BOSMC achieves a superior global fit (W_2), the KimPal estimator yields an expected revenue slightly closer to the DGP baseline. This counterintuitive result stems from the fundamental difference between global distributional approximation and localized mechanism design objectives. Because W_2 evaluates the entire

support, BOSMC effectively minimizes global discrepancies and avoids the severe boundary errors that plague kernel-based estimators. In contrast, optimal reserve prices are not dictated by the global distribution, but rather by the local geometry of the inverse hazard rate in the interior of the support, where the virtual valuation crosses zero (Gomes & Sweeney, 2014). Thus, while KimPal's kernel density estimation produces larger overall errors due to boundary bias, its local smoothing coincidentally provides a highly accurate fit in the specific interior region where the reserve price binds.

These results also clarify the intended role of our estimator. In auction environments for which a specialized inversion-based estimator is available, BOSMC is not expected to uniformly dominate such an estimator for every downstream objective. Rather, the benchmark demonstrates that the simulation-based approach can recover economically relevant primitives with comparable downstream performance while avoiding reliance on a mechanism-specific inversion formula. Its principal advantage therefore lies in portability across auction environments rather than in uniformly superior reserve-price optimization within the benchmark setting.

Overall, the benchmark simulations serve two purposes. First, they verify that the simulation-based estimator performs well in environments where standard structural estimators are available for comparison. Second, they evaluate whether errors in distributional recovery translate into economically meaningful losses in reserve-price counterfactuals.

5.2 GSP Laboratory Experiment

Our main validation exercise uses a GSP laboratory experiment conducted by Bae & Kagel (2019). In this setting, participants act as advertisers competing for multiple ad positions. Each round, three bidders receive private values drawn uniformly from the interval $[0, 100]$ and submit per-click bids. The platform assigns positions using the GSP rule, setting per-click payments based on the next-highest bid or a quality-adjusted score. The experiment has two treatments with different position-specific CTRs: the first with two slots at 11% and 10% (Treatment 1), and the second with two slots at 11% and 3% (Treatment 2).

The experiment records bids, values, assigned positions, payments, and auction-level parameters. For estimation, we use only the information typically available to a platform econometrician: bids, positions, payments, and the auction rules. Crucially, the true experimental values are held out and used solely for validation. This design allows us to rigorously evaluate whether the estimator can recover the latent value distribution from strategic bid data alone.

Table 3 presents the estimation results for the experimental data, comparing the Wasserstein-2 (W_2) distances between the estimated and empirical value distributions. In the first treatment, our proposed BOSMC approach substantially outperforms the KimPal benchmark. This performance advantage persists in the second treatment. Furthermore, the BOSMC estimator exhibits lower standard deviations across both experimental environments, indicating more stable recovery. Ultimately, these results demonstrate that the simulation-based framework is highly robust to variations in platform design and reliably recovers the underlying fundamentals regardless of slot-quality decay.

Table 3: W_2 distances between estimated and true empirical value distributions across the two CTR treatments; standard deviations are reported in parentheses.

Treatment	BOSMC (W_2)	KimPal (W_2)
11-10	152.862 (9.603)	216.940 (16.475)
11-3	13.305 (2.675)	15.281 (3.311)

5.3 Counterfactual Platform Design

Search engines and advertising platforms treat granular bid records and auction parameters as strictly proprietary due to their immense strategic and commercial importance. To demonstrate the empirical applicability of our estimator despite this limitation, we conduct a counterfactual analysis of platform design. Specifically, we recover the underlying advertiser value distributions from observed bids to subsequently evaluate the revenue implications of updating platform reserve prices. Because internal platform logs remain inaccessible, we leverage a synthetic dataset published by

Ostrovsky & Schwarz (2023). This dataset is carefully calibrated to capture the structural complexities of real-world search advertising markets, providing a highly realistic environment for evaluating our counterfactual reserve-price policies.

The dataset comprises multiple keywords, each defined by the number of active bidders and the parameters of a log-normal prior value distribution.[1] From this dataset, we randomly sample 100 keywords and simulate equilibrium bids based on their structural parameters. We then apply both our estimator and the benchmark to these generated bids to recover the latent value distributions. In each auction, there was an existing reserve price assumed of $\$0.10$. Finally, we derive the optimal reserve prices under these estimated priors and compute the platform's expected revenue.

Overall, we find that the average Wasserstein distance is lower for the benchmark estimator ($W_2 = 0.03$) compared to our approach ($W_2 = 0.07$).[2] This discrepancy arises because the data-generating prior is unbounded on the upper tail, whereas our estimator assumes a prior with compact support. Consequently, our approach recovers the correct overall shape of the prior but fails to precisely capture the extreme upper tail, leading to the gap in the Wasserstein distance. The expected revenues, however, are almost identical at 0.54 for both estimators, closely matching the theoretical revenue under the optimal reserve price for the true DGP. This represents a 4% increase in expected revenue across all keywords, compared to the prevailing reserve price. Given the scale of advertising platforms, even modest revenue gains can be economically significant.

Overall, this analysis demonstrates that our approach can be successfully applied in an environment calibrated to bidding behavior on real-world platforms. Furthermore, it shows how the framework can be used to evaluate market-design changes and their revenue implications. Nonetheless, this synthetic environment remains an abstraction of the real world, and future work should aim to conduct similar analyses using actual field data.

6 Conclusions

To address the challenge of data-driven market design in digital environments, this paper develops a simulation-based structural estimator for non-truthful GSP platform auctions. The method is explicitly designed for complex platform architectures where analytical bid-to-value inversion is computationally intractable or highly sensitive to proprietary institutional details. By integrating flexible value-distribution parameterizations, equilibrium-learning algorithms, and likelihood-free inference, the proposed framework empowers platform operators to reliably recover latent economic primitives from operational bid data and rigorously evaluate counterfactual design choices.

The primary advantage of this framework is its flexibility: it is uniquely suited for complex platforms where auction rules can be simulated, yet standard econometric inversions remain analytically prohibitive. However, this approach has boundaries. Methodologically, it relies on approximate equilibrium behavior and well-defined equilibrium selection; consequently, systematic behavioral deviations may be inadvertently absorbed into the estimated value distributions. Computationally, the equilibrium oracle incurs non-trivial estimation costs. Empirically, robust recovery requires a rich dataset of observable bids, positions, and payments. Finally, as with any structural model, counterfactual predictions remain conditional on the maintained assumptions and may be sensitive to the chosen discrepancy metric and support specifications.

Despite these limitations, the practical implications are substantial. The framework's core operational advantage is decoupling structural estimation from analytical tractability. Modern platforms deploy complex mechanisms, such as proprietary quality scores or dynamic slots, that can render traditional analytical inversions difficult or unavailable. Because our approach requires only a simulable auction mechanism and a suitable equilibrium oracle, it can accommodate changes in auction design without requiring a new closed-form econometric solution. Engineers can instead update the simulation environment and use the resulting model to estimate latent advertiser values and evaluate counterfactual market designs. As illustrated by the 4% expected revenue improvement in our field-calibrated synthetic application, recovering these primitives can have economically meaningful implications for mechanism optimization.

[1] Notably, the log-normal distribution is an assumption made by Ostrovsky & Schwarz (2023).

[2] Increasing SMC particles per generation would better approximate the upper tail and narrow this gap, but would not significantly alter the optimal reserve price.

AI USE STATEMENT

In this work, we used generative AI tools for language editing of the manuscript, including grammar correction and limited rephrasing to improve clarity and readability. We have not used generative AI tools for the development of the theoretical results, the generation or modification of code, or the generation of data, and the remaining required disclosure tasks are not applicable to this work. Additionally, we used generative AI tools for minor stylistic revisions of the manuscript. We have reviewed all AI-assisted work. All AI-assisted language edits were manually reviewed to ensure that they did not alter the substantive content, theoretical arguments, or scientific claims of the manuscript. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REFERENCES

Ashish Agarwal and Tridas Mukhopadhyay. The impact of competing ads on click performance in sponsored search. *Information Systems Research*, 27(3):538–557, 2016. doi: 10.1287/isre.2016.0637.

Ashish Agarwal, Kartik Hosanagar, and Michael D. Smith. Do organic results help or hurt sponsored search performance? *Information Systems Research*, 26(4):695–713, 2015. doi: 10.1287/isre.2015.0593.

Jinsoo Bae and John H. Kagel. A dataset collected in a lab experiment for the generalized second price auction. *Data in Brief*, 26:104469, 2019. doi: 10.1016/j.dib.2019.104469.

Espen Bernton, Pierre E. Jacob, Mathieu Gerber, and Christian P. Robert. Approximate bayesian computation with the wasserstein distance. *Journal of the Royal Statistical Society: Series B*, 81 (2):235–269, 2019a. doi: 10.1111/rssb.12312.

Espen Bernton, Pierre E. Jacob, Mathieu Gerber, and Christian P. Robert. On parameter estimation with the wasserstein distance. *Information and Inference: A Journal of the IMA*, 8(4):657–676, 2019b. doi: 10.1093/imaiai/iaz003.

Siddharth Bhattacharya, Jing Gong, and Sunil Wattal. Competitive poaching in search advertising: Two randomized field experiments. *Information Systems Research*, 33(2):599–619, 2022. doi: 10.1287/isre.2021.1073.

Martin Bichler, Max Fichtl, Stefan Heidekrüger, Nils Kohring, and Paul Sutterer. Learning Equilibria in Symmetric Auction Games using Artificial Neural Networks. *Nature Machine Intelligence*, 3:687–695, August 2021. doi: 10.1038/s42256-021-00365-4.

Martin Bichler, Maximilian Fichtl, and Matthias Oberlechner. Computing Bayes-Nash Equilibrium Strategies in Auction Games via Simultaneous Online Dual Averaging. *Operations Research*, 73(2):1102-1127, 2023. doi: 10.1287/opre.2022.0287.

Tilman Börgers, Ingemar Cox, Martin Pesendorfer, and Vaclav Petricek. Equilibrium bids in sponsored search auctions: Theory and evidence. *American Economic Journal: Microeconomics*, 5 (4):163–187, 2013. doi: 10.1257/mic.5.4.163.

Eric Brochu, Vlad M. Cora, and Nando de Freitas. A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning, 2010.

Yeon-Koo Che, Syngjoo Choi, and Jinwoo Kim. An experimental study of sponsored-search auctions. *Games and Economic Behavior*, 102:20–43, 2017. doi: 10.1016/j.geb.2016.10.008.

Xiaomeng Du, Meng Su, Xiaoquan Michael Zhang, and Xiaona Zheng. Bidding for multiple keywords in sponsored search advertising: Keyword categories and match types. *Information Systems Research*, 28(4):711–722, 2017. doi: 10.1287/isre.2017.0724.

Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007. doi: 10.1257/aer.97.1.242.

Anindya Ghose and Sha Yang. An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Science*, 55(10):1605–1622, 2009. doi: 10.1287/mnsc.1090.1054.

Renato Gomes and Kane Sweeney. Bayes–nash equilibria of the generalized second-price auction. *Games and economic behavior*, 86:421–437, 2014.

Christian Gouriéroux, Alain Monfort, and Eric Renault. Indirect inference. *Journal of Applied Econometrics*, 8(S1):S85–S118, 1993. doi: 10.1002/jae.3950080507.

Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong. Optimal Nonparametric Estimation of First-price Auctions. *Econometrica : journal of the Econometric Society*, 68(3):525–574, May 2000. ISSN 0012-9682, 1468-0262. doi: 10.1111/1468-0262.00123.

Brent R Hickman and Timothy P Hubbard. Replacing sample trimming with boundary correction in nonparametric estimation of first-price auctions. *Journal of Applied Econometrics*, 30(5):739–762, 2015.

Karthik Kannan, Vandith Pamuru, and Yaroslav Rosokha. Analyzing frictions in generalized second-price auction markets. *Information Systems Research*, 34(4):1437–1454, 2023. doi: 10.1287/isre.2022.1187.

Dongwoo Kim and Pallavi Pal. Nonparametric estimation of sponsored search auctions and impact of ad quality on search revenue. *Management Science*, 71(12):10047–10066, 2025. doi: 10.1287/mnsc.2023.02052.

Vijay Krishna. *Auction Theory*. Academic Press, Amsterdam Boston Heidelberg London New York Oxford Paris San Diego San Francisco Singapore Sydney Tokyo, second edition, 2010. ISBN 978-0-12-374507-1.

Jean-Jacques Laffont, Isabelle Perrigne, Michel Simioni, and Quang Vuong. Econometrics of scoring auctions. 2020.

De Liu, Jianqing Chen, and Andrew B. Whinston. Ex ante information and the design of keyword auctions. *Information Systems Research*, 21(1):133–153, 2010. doi: 10.1287/isre.1080.0225.

Daniel McFadden. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica : journal of the Econometric Society*, 57(5):995–1026, 1989. doi: 10.2307/1913621.

M. D. Mckay, R. J. Beckman, and W. J. Conover. A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output From a Computer Code. *Technometrics : a journal of statistics for the physical, chemical, and engineering sciences*, 42(1):55–61, February 2000. ISSN 0040-1706, 1537-2723. doi: 10.1080/00401706.2000.10485979.

Roger B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, February 1981. ISSN 0364-765X, 1526-5471. doi: 10.1287/moor.6.1.58.

Michael Ostrovsky and Michael Schwarz. Reserve prices in internet advertising auctions: A field experiment. *Journal of Political Economy*, 131(12):3352–3376, 2023.

Tim Pillards and Ronald Cools. Transforming low-discrepancy sequences from a cube to a simplex. *Journal of Computational and Applied Mathematics*, 174(1):29–42, February 2005. ISSN 03770427. doi: 10.1016/j.cam.2004.03.019.

Tina Toni, David Welch, Natalja Strelkowa, Andreas Ipsen, and Michael P.H Stumpf. Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of The Royal Society Interface*, 6(31):187–202, February 2009. ISSN 1742-5689, 1742-5662. doi: 10.1098/rsif.2008.0172.

Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007. doi: 10.1016/j.ijindorg.2006.10.002.