Mass of Bids at Reserve*

Miguel Alcobendas
Yahoo Research

Robert Zeithammer
UCLA Anderson School of Management

December 12, 2025

Abstract

We document partial pooling of bids at the reserve-price level in a large market for digital display advertising that employs a first-price auction mechanism to sell individual impressions. Such partial pooling is unexpected because standard theory predicts that equilibrium bidding strategies in first-price auctions should not involve a substantial mass of bids at the reserve.

To explain the mechanism underlying the partial pooling, we analyze a unique feature of our data - a direct observation of the actual bidding strategy of one of the bidders. The bidder whose strategy we observe employs a heuristic bidding strategy that mechanically produces the mass of bids at reserve as follows: before bidding begins, the bidder pre-calculates a candidate bidding strategy based on only his valuation of the impression but without considering the reserve price. When an opportunity to buy an impression arises, the bidder censors the bid at exactly the reserve price whenever the valuation of the opportunity exceeds the reserve but the pre-calculated bidding falls below it. While we do not directly observe the bidding strategies of the remaining bidders in our data, we test for a data pattern consistent with the Heuristic model: when the reserve shifts downward for an exogenous reason, the tail of the bidding distribution above the higher of the two reserve levels should be the same. We find that a majority of the bidders do indeed exhibit this pattern.

To understand the economic consequences of the Heuristic Bidding strategy, we use a counterfactual analysis to compare the publisher's revenue and the bidders' surplus under the heuristic bidding strategy and the more theoretically grounded Bayes-Nash equilibrium bidding strategy. We find that the counterfactual Bayes-Nash equilibrium model involves generally higher bids than the

*We thank Jan Brueckner, Jason Choi, Shunto Kobayashi, Maciej Kotowski, and seminar participants at UCIrvine Paul Merage School of Business and Yahoo; as well as participants at International Industrial Organization Conference (2026) and ISMS Marketing Science Conference (2026) for their helpful comments.

Heuristic model, increasing the average auction revenue by 9% and decreasing bidders' surplus by 5%- both economically meaningful amounts. Widespread use of simplified bidding heuristics in online advertising can thus lead to sub-optimal outcomes for publishers.

Keywords: digital advertising, auctions,reserve prices, bidding strategies, ad tech

1 Introduction

Equilibrium bidding strategies in simple first-price sealed-bid auctions do not involve any pooling of bids at the reserve price level. Yet, we find that partial pooling at reserve is common in a prominent market for digital display advertising that uses the first-price sealed-bid auction to sell individual impressions. Specifically, we analyze the distributions of bids submitted by twelve large bidders who participated in the Yahoo ad exchange in 2023 and together account for 80 percent of revenues, and show that most of their bid distributions have a substantial mass of bids at reserve. In addition to documenting this surprising phenomenon, we explore potential mechanisms behind it using a unique auction-level dataset, and show that it can be economically meaningful in that it can lower the auction revenue by 9 percent relative to counterfactual revenue in a standard Bayes Nash equilibrium model that does not involve any pooling of bids.

To document the partial pooling phenomenon, we rely on a unique way the exchange that provided our data set reserve prices during the time of our data collection. Each bidder randomly faced one of only two different levels of the reserve price: the optimized "production" level and a much lower "monitoring" level the exchange used to learn about bids below the production reserve level. We find that most bidders bid exactly the reserve price with a large positive probability when the reserve is high (i.e. set at the production level), but this mass-point in the bid distribution disappeared when the reserve shifted to the monitoring level. The partial bid-pooling at the reserve price we find is thus not due to the distribution of valuations having a mass point just where the reserve price happens to be.

Yahoo's own bidding platform, the Yahoo Demand Side Platform (Yahoo DSP), provides a clear example of the "disappering" mass at the production-level reserve: When Yahoo DSP faces the production-level reserve price of \$0.71, it submits bids exactly equal to this amount with a probability of about 12 percent. But when the exchange randomly uses the much lower monitoring reserve price of \$0.15[1], the probability the same DSP bids exactly \$0.71, conditional on its bid being at or above \$0.71, falls to less than 1 percent. The partial bid-pooling at the production reserve price

[1] All \$ amounts are expressed in dollars per thousand impressions (CPM), unless stated otherwise.

we find is thus not due to the Yahoo DSP valuing 12 percent of impressions exactly where the production reserve price happens to be. This pattern of a disappearing mass point characterizes the behavior of a vast majority (11 of 12, i.e. 91 percent) of the large bidders we observe.[2]

The underlying mechanism most consistent with our data is that about half of the bidders use a simplified Heuristic Bidding model, which first solves for a univariate bidding strategy as a function of the valuation without considering the reserve price, and then censors the actual bids at exactly the reserve level whenever the valuation of the opportunity exceeds the reserve but the univariate strategy does not. This is the first paper documenting such heuristic behavior by bidding algorithms in online advertising. We offer both direct evidence for this simplifying heuristic by observing the bidding strategy of one of the bidders, and indirect evidence by testing a distributional equivalence predicted by the model.

For the Yahoo DSP bidder, we have access to the actual valuations corresponding to each bid, so we effectively observe the bidder's bidding strategy. This allows us to observe the above-described heuristic bidding in action, and document how its censoring mechanism produces the observed mass point at the production reserve. To provide direct evidence of the Heuristic model's underlying univariate bidding strategy, we also show that for high-enough valuations such that the bids exceed the production reserve level, the Yahoo DSP employs an identical bidding strategy regardless of which reserve price is in effect.

We extend this analysis to other bidders for whom valuations are unobserved and find that the bidding distributions as a function of reserve prices for about half of the bidders are consistent with the Yahoo DSP's Heuristic Bidding model, suggesting they use similar heuristics.

To measure the economic significance of the partial pooling at reserve arising from the heuristic bidding we discovered, we simulate a market in which all bidders follow heuristic bidding, and compare seller revenues and bidder surpluses to those in an otherwise identical market that follows Bayes Nash equilibrium bidding strategies (that does not involve any mass points at reserve). This exercise is complicated by the fact that the exchange recognizes two types of bidders (Regular DSPs and Rebroadcasters) who face different production reserve prices. We thus need to draw on the recently developed theory of first-price sealed-bid auctions with asymmetric bidders facing different reserve prices by Kotowski (2018) to characterize the counterfactual equilibrium bidding. On the other hand, the exercise is also simplified by the fact

[2]For each of these 11 bidders, the mass point at the production reserve (conditional on bids weakly above it) is significantly larger when the reserve is at the production level than when it is at the monitoring level. Among these 11 bidders, the mass point at the production-level reserve "disappears" in the sense that the probability of that bid amount drops from an average of 31 percent of bids to about 2 percent of bids.

we observe the underlying valuations for one of the bidder types (the Regular DSP, observed via the Yahoo DSP), so we do not have to separately estimate the valuations of both bidder types.

The counterfactual equilibrium without a mass of bids at reserve involves generally higher bids than the Heuristic model, increasing the average auction revenue by 9% and decreasing bidders' surplus by 5%- both economically meaningful amounts. The decrease in bidder surplus is not simply shared proportionately between the two bidder types - switching from the Heuristic model to the equilibrium one would make Regular DSPs better off while making Rebroadcasters worse off. These findings highlight the significant and differential effects that strategic bidding behavior can have on various types of auction participants.

2 Related Literature

Previous literature discusses inconsistencies between observational data and predictions from auction theory. For instance, the existence of a mass of bids at the reserve price was documented by Reiley (2006) in his collectible trading cards experiment. Reiley (2006) used a field experiment to verify the predictions of the theoretical model of reserve prices in first-price auctions developed by Riley and Samuelson (1981). Conditional on having a valuation strictly above the reserve price, standard theory of private-value first-price auctions predicts that increasing reserve prices does not raise bids merely to the level of the reserve, but rather strictly above it. In other words, equilibrium bidding strategies in simple first-price sealed-bid auctions do not involve any pooling of bids at the reserve price level. Reiley (2006) found that, inconsistent with the strategic behavior predicted by the standard theory, some bidders submitted bids equal to the higher reserve price, and concluded that these bidders behaved naively. We extend Reiley's work by finding similar bidding patterns in an economically larger market facilitated by supposedly sophisticated algorithms. In addition, we also directly document the bidding heuristic that leads to the partial pooling, and we estimate its effects on economic outcomes using counterfactual simulations.

Existing auction theory does include equilibrium auction models with partial pooling at reserve, but the assumptions of those models do not hold in the present setting. For example, Haile (2000) finds such partial pooling in markets with resale, but digital advertising impressions cannot be resold. Lauermann and Speit (2023) find that mass-points can occur in bidding for common-value objects, but the detailed targeting of digital display ads based on private information individual bidders have about each opportunity makes the private-value model more relevant here.

Golrezaei et al. (2021) also discuss inconsistencies between observational studies and theoretical predictions. Using data from Google's display advertising exchange

in a second-price auction environment, Golrezaei et al. (2021) show that the behavior of some buyers is not consistent with the predictions of classical auction theory. In particular, they found that a significant set of buyers lower their bids in response to higher reserve prices and concluded that this behavior can be explained by buyers' return on investment constraints. This observation contrasts with standard results in second-price auctions, where bidders have no incentive to change their bids as a function of reserve prices. In contrast to Golrezaei et al. (2021), we study first-price auctions, in which equilibrium bids do depend on the level of reserve, and we study a different deviation from theoretical predictions, namely the partial pooling at reserve.

Our counterfactual exercise is based on the theoretical work by Kotowski (2018), who characterizes the equilibrium bidding function in first-price auctions with asymmetric reserve prices and asymmetric bidders. Kotowski (2018) builds on prior studies of asymmetric first-price auctions, particularly Lebrun (1999, 2006), Lizzeri and Persico (2000), and Maskin and Riley (2003). While the Kotowski (2018) equilibrium allows us to perform counterfactual simulations with two bidder types each having a different distribution of valuations and facing a different reserve price, it does restrict us to the assumption that all bidders of a given type have the same distribution of valuations. We use the observed distribution of Yahoo DSP's valuations as the distribution of Regular DSP valuations, and we estimate the distribution of Rebroadcaster valuations using the approach of Flambard and Perrigne (2006). Therefore, our empirical approach also connects with the structural empirical literature on auctions.

3 Institutional Background: Real-time auctions for digital display advertising with two bidder types

Our analysis focuses on auctions of digital display advertising. The advertisements are displayed in rectangular boxes between or on the side of the main content of the online publisher which sells the ad space. In addition to location on the webpage, the ads also differ in size (e.g. 300x250 or 600x250 pixels). The "seller" in the auction is the *publisher* whose web page with some ad space for sale is being browsed by a consumer at the moment. In our data, the seller is a popular website in the yahoo.com domain - one of the top ten domains in the U.S. by traffic.

The role of "auctioneer" is fulfilled by an *ad exchange*, which collects the bids and determines the winner of the auction based on pre-announced rules. In our case, the ad exchange is the Yahoo Ad Exchange, which has been using the *first-price sealed-bid* rules with a reserve price since about three years prior to the time of the our data collection (2023).[3] These auction rules were also the established industry standard

[3]The Yahoo Ad Exchange, along with most other exchanges facilitating real-time bidding on digital advertising impressions switched to the first-price rules in 2019 (Despotakis et al. (2021)).

at other ad exchanges at the time of our data collection.

The auction is triggered when the user opens the publisher's web page through her browser. The ad exchange sends an ad request with information about the user and the site, among other parameters, to the potential bidders along with a reserve price.

We now describe the bidders in this market. Given the large volume of auctions and the complexity and required speed of online bidding, advertisers do not participate in these auctions directly, but instead engage with bidding agents who process the available information about the user and other impression characteristics, determine the bid amount, and act as bidders in the auction. The Yahoo Ad Exchange distinguishes between two types of such bidders: regular Demand-Side Platforms (DSPs) and Rebroadcasters. This distinction is relevant because each bidder type faces a different reserve price. Regular DSPs are agents of the advertisers as described above. Rebroadcasters are themselves third-party exchanges that rebroadcast the bid request to their own DSPs and bid on their behalf. Importantly for our analysis, Rebroadcasters face a higher production reserve price than Regular DSPs on Yahoo Ad Exchange. Therefore, Rebroadcasters effectively only compete at higher bid-levels, and there is a range of bids between the regular and Rebroadcaster reserve prices in which only Regular DSPs can win.

During the time of our data collection, the Yahoo ad exchange set reserve prices by type of bidder (Regular DSP vs. Rebroadcaster) and inventory (i.e. the upper-right banner space in Yahoo Mail vs. the top banner in Yahoo news). However, reserve prices did not generally vary from one impression to another during the day or across bidders within a type other than an occasionally and randomly switching to a much lower level to monitor the bidding landscape. The monitoring reserve level was the same for both types of bidders. We exploit this randomization in our main analysis in two ways: first, we use it to show that the mass of bids at the production reserve disappears when the reserve level shifts down to the monitoring level. Second, we use the Rebroadcaster bid observations with reserve at the monitoring level to estimate the Rebroadcaster valuations on the full support needed for our counterfactual simulations.

4 Theory

What does auction theory predict about the market we just described? In this section, we first establish the central theoretical claim that partial pooling at reserve is incompatible with equilibrium bidding in simple first-price auctions, and then present two theoretical models of the market we are studying: 1) a fully rational equilibrium model adapted from the literature, and 2) a Heuristic model we propose to explain the partial pooling at reserve we observe.

4.1 Mass points at the reserve price are inconsistent with equilibrium

To see why partial pooling at reserve is incompatible with equilibrium bidding, consider a focal bidder with valuation $v > r$ where r is the reserve price in the auction, and assume the bidder's competition involves a mass point of bids at r. The focal bidder's belief about the competition is summarized by his probability of winning $Pr(win|b)$ at different bid-levels b. The mass-point of competing bids at r means there exists some $M > 0$ such that for any $\epsilon > 0$, $Pr(win|r + \epsilon) \geq Pr(win|r) + M$.[4] Assuming a quasi-linear utility function, the bidder's best response to his belief about the probability of winning solves:

$$\max_{b \geq r} \Pr(win|b)(v - b) \tag{1}$$

It is immediate that for any $M > 0$ there exists an $\epsilon > 0$ such that bidding $r + \epsilon$ dominates bidding r. Formally, bidding r yields $Pr(win|r)(v - r)$ while bidding $r + \epsilon$ yields at least $[\Pr(win|r) + M](v - r - \epsilon)$, which exceeds the payoff from bidding exactly r for all $\epsilon < \frac{M(v-r)}{2(M+\Pr(B<r))}$. In words: bidding just over r avoids the possibility of a tie with the highest competitor at r at the cost of an arbitrarily small increase in price paid.

The only time such a small increase in price paid is not worth the increase in the chance of winning is the situation when the bidder values the good exactly r, so paying more would result in a loss.[5] So one possible explanation of mass points at the reserve is that the seller happens to set the reserve exactly at a point where the valuation distribution of at least some bidders has a mass point.

Note that the above discussion does not make any assumption about the rationality, number, or symmetry of the competitors the focal bidder is facing. As long as at least some of the competitors bid the reserve price r with positive probability, the above argument, and consistent with previous literature (e.g., Reiley (2006) and Lebrun (1999)), shows that a rational bidder would never respond by bidding r except in the event that his valuation equaled r exactly. In other words, an equilibrium bidding strategy that placed a positive mass of bids on r would unravel because rational bidders would respond by bidding slightly above r.

[4]The exact formula for M depends on the number of competitors and tie-breaking rules. For example, if there is only one competitor who bids B with a mass point of size m at r in that $Pr(B = r) = m$ and ties are broken at random, then $Pr(win|r + \epsilon) \geq Pr(B < r) + m = Pr(win|r) + m/2$ because $Pr(win|r) = Pr(B < r) + m/2$. So in this example, $M = m/2$.

[5]The right hand side of the inequality is positive as long as $v > r$, but it becomes zero when $v = r$.

4.2 Modeling Assumptions

We consider a model with two types of bidders: Regular DSPs indexed throughout by $i = 1$ subscript, and Rebroadcasters indexed throughout by $i = 2$ subscript. Regular DSP and Rebroadcaster types consist of N_1 and N_2 potential bidders, respectively, with $N_1 + N_2 = N$ and $N \geq 2$. Let $v_{i,n}$ for $n = 1, ..., N_i$ denote the valuations of N_i potential bidders of type i. Assume the valuations $v_{i,n}$ are drawn independently from continuous distributions $F_i[\cdot]$ with full support on the interval $[0, \bar{v}]$, and denote the corresponding densities $f_i[\cdot]$.

The auctioneer runs a first-price sealed-bid auction with different public reserve price for each bidder type r_i, where $r_1 \leq r_2 \leq \bar{v}$, and rejects any bid below the corresponding reserve price. The number of actual type i participants in the auction may differ from the number of potential bidders N_i as some bidders may have a valuation below the corresponding reserve price r_i. We assume that bidders know their own valuation as well as their chance of winning at different bid-levels. Having defined our modeling assumptions, we now introduce the Heuristic model that seems to fit the bidding behavior we observe.

4.3 Heuristic Bidding Strategy

Consider a single bidder of type $i \in \{1, 2\}$, who knows its average chance of winning at different bid-levels $\Pr(win|b)$ but does not know how the chance of winning changes as the reserve price changes. The following process describes a seemingly reasonable bidding strategy such a bidder can follow:

1. Given a set of valuations v (for example, a fine grid spanning the $[0, \bar{v}]$ interval), pre-compute the univariate bidding strategy $\tilde{b}_i(v)$ that maximizes the expected surplus $\Pr(win|b)(v - b)$ separately for each valuation in the set.

2. When a bid-request arrives with a specific reserve price r_i, find the valuation of the opportunity v, and bid $\tilde{b}_i(v)$ when it is high-enough that $\tilde{b}_i(v) > r_i$. When $\tilde{b}_i(v) < r_i < v$, simply bid the reserve r_i so as not to pass up on a potentially valuable trade.

Mathematically, let the heuristic bidding strategy $\tilde{\beta}_i(v|r_i)$ of type i bidder facing a reserve price r_i be:

$$\tilde{\beta}_i(v|r_i) = \begin{cases} 0 & if \quad v < r_i \\ r_i & if \quad v \in [r_i, \tilde{v}_i] \\ \tilde{b}_i(v) & if \quad v > \tilde{v}_i \end{cases} \tag{2}$$

In words, the bidder cannot profitably participate if its valuation of the impression opportunity v is below its reserve r_i. If the valuation is greater than the reserve and below some valuation threshold $\tilde{v}_i = \tilde{b}_i^{-1}(r_i)$, then the bidder submits a bid equal to r_i. For valuations above $\tilde{v}_i$, the bidder's response is determined by the univariate bidding function $\tilde{b}_i(v)$.

Note that the heuristic bidding strategy does not involve any notion of equilibrium. Instead, it is a simple response to a challenging environment the bidders find themselves in. In real-time bidding for online advertising, bidders face millions of ad requests every second from very heterogeneous publishers. For each request, bidders must decide which ad they are going to use in the response along with the corresponding bid amount. Employing a very complex bidding strategy optimized in real time may add latency in the bid response, increasing the chances that the response takes longer than the time limit imposed by exchanges. The heuristic defined above easily handles variation in the reserve price values enclosed in the bid request.

We can only speculate about the origins of this heuristic bidding in the market we study. One possible origin is in the market's initial embrace of the second-price auction in which bidding does not depend on reserve or the probability of winning. Once the market switched to the first-price auction around 2019, the bidders already had working models that calculated valuations of impression opportunities, and they needed to adjust these valuations to come up with bids to submit in first-price auctions. Several papers analyze this transition, and conclude it was difficult for the bidders (Alcobendas and Zeithammer (2021) and Goke et al. (2024)), so it is not surprising that the bidders were searching for simplified solutions to the bidding problem.

Using experiments to measure the average probability of winning at different bid levels $\Pr(win|b)$ is a tractable first pass at the problem of developing first-price auction bidding strategies. It is thus possible that the strategists developing the new bidding algorithms learned $\Pr(win|b)$, but then needed to somehow adjust for the reserves changing beyond their control. We propose that the above heuristic bidding strategy is a pragmatic solution to this problem. The heuristic nature of the strategy arises because bidders ignore the reserve in the first stage rather than conditioning on it, as in $\Pr(win|b, r)$. If they conditioned on reserves, they could eventually derive the Bayes-Nash equilibrium bidding strategies to which we now turn.

4.4 Counterfactual: Bayes-Nash Equilibrium Bidding Strategies

We now consider a counterfactual bidding strategy where both Regular DSPs and Rebroadcasters bid according to Bayes-Nash equilibrium strategies. Theorem 2a in Kotowski (2018) shows that if $r_1 < r_2$ and the following condition holds

$$\int_{r_1}^{\bar{v}} F_2[r_2]^{N_2} F_1[z]^{N_1-1} dz < \int_{r_2}^{\bar{v}} F_2[z]^{N_2-1} dz \tag{3}$$

there exists a Bayesian-Nash equilibrium where all bidders of type $i \in \{1, 2\}$ adopt the strategy

$$\beta_i(v|r_i) = \begin{cases} 0 & if \;\; v < r_i \\ b_i(v) & if \;\; v \in [r_i, \hat{v}_i] \\ \hat{b}_i(v) & if \;\; v > \hat{v}_i \end{cases} \tag{4}$$

where $\hat{v}_i$ is a specific valuation threshold that Kotowski characterizes, and:

(a) The function $b_i : [r_i, \hat{v}_i] \to [r_i, \bar{v}]$ is given by the symmetric Bayes Nash equilibrium that would arise if the bidder faced only competitors of its own type, characterized by Riley and Samuelson (1981)

$$b_i(v) = v - \int_{r_i}^{v} \left[\frac{F_i[z]}{F_i[v]}\right]^{N_i-1} dz \tag{5}$$

(b) The function $\hat{b}_i : [\hat{v}_i, \bar{v}] \to [r_2, \bar{v}]$ is defined as $\hat{b}_i = \phi_i^{-1}$, where (ϕ_1, ϕ_2) is the unique solution to the system of differential equations

$$[\phi_i(b_i) - b_i] \left[\frac{N_j f_j \left[\phi_j(b_i)\right] \phi_j'(b_i)}{F_j[\phi_j(b_i)]} + \frac{(N_i - 1) f_i \left[\phi_i(b_i)\right] \phi_i'(b_i)}{f_i[\phi_i(b_i)]}\right] = 1 \tag{6}$$

for $i \neq j$, satisfying the boundary conditions

- $\phi_1(\eta^*) = \phi_2(\eta^*) = \bar{v}$ for some $\eta^* \in (r_2, \bar{v})$
- $\phi_1(b^*) = \hat{v}_1$ and $\phi_2(b^*) = \hat{v}_2$ where $b^* = b_2(\hat{v}_2)$

To gain intuition for the system of Equations 6, note that each equation is a type-specific first-order condition of the optimization problem of a type i bidder who maximizes the expected surplus when bidding against opponents of both types:

$$(v - b) F_j[\phi_j(b)]^{N_j} F_i[\phi_i(b)]^{N_i-1} \tag{7}$$

whenever valuations are above the threshold $\hat{v}_i$.[6] Interestingly, the equilibrium bidding function $\beta_1(\cdot)$ of the lower type (in the present context, this is the Regular DSP) is discontinuous at $\hat{v}_1$, see Kotowski (2018)'s Figure 2, Case 1. While Kotowski (2018) also analyzes the equilibrium for cases where Equation 3 does not hold, our estimation results (presented later) confirm this condition is satisfied in our setting. Therefore, Equation 4 characterizes the Bayes-Nash bidding behavior for buyers here.

5 Data Description

We employ exchange bidding data from auctions selling display advertising space on a popular website in the *yahoo.com* domain. We consider a random sample of user impressions from the United States during one day in February 2023, stratified according to the two reserve-price levels (monitoring and production). For each auction in the sample, the exchange data contains all bids submitted by all bidders (Regular DSPs and Rebroadcasters).

In addition to the exchange bidding data, we also observe contemporaneous valuation information coming from Yahoo's own DSP. Specifically, we observe Yahoo DSP's valuations of *all* the ad opportunities in our sample, including both the valuations behind every bid the Yahoo DSP submitted and valuations of opportunities on which the Yahoo DSP did not bid. This information allows us to reconstruct Yahoo DSP's bidding strategy.

Table 1 presents summary statistics of key variables in the dataset, with standard deviations in parentheses. The column marked as *All* reports aggregate information, while columns *Monitoring Reserve* and *Production Reserve* slice the data according to the reserve price value enclosed in the bid request. Recall that the *Monitoring* reserve price value is smaller than the *Production* one and equal to all bidders. On the other hand, the reserve price used in production depends on the bidder type: Regular DSP and Rebroadcaster. Recall also that our sample is stratified according to the reserve-price level, so the approximately equal number of auctions in each stratum of Table 1 arises from our sampling strategy, not from actual prevalence of each reserve level in the field (where only 5 percent of auctions have the monitoring reserve).

The top panel of Table 1 describes our sample of exchange bidding data, which consists of over 1.55 million bids from 225,379 auctions, split roughly evenly into the two reserve-price strata. As expected, the auctions with the lower *Monitoring* reserve receive more bids of lower magnitude. More interestingly, they also clear for a lower price, indicated by *Winning Bid*. The top panel also shows the first evidence of the

[6]Note that the terms in Equation 6 have been rearranged with respect to the Equation 7 appearing in Theorem 2a in Kotowski (2018) to align with the definitions used here.

Table 1: Summary Statistics

	All	Monitoring Reserve	Production Reserve
Auction - Yahoo Exchange:			
No. Bids	1,553,592	1,197,821	355,771
No. Auctions	225,379	115,791	109,588
Bid (× \$1000)	0.75	0.62	1.21
	(0.49)	(0.44)	(0.38)
No. Bidders	6.89	10.34	3.25
	(5.95)	(5.78)	(3.37)
Winning_Bid (× \$1000)	1.31	1.18	1.48
	(0.56)	(0.59)	(0.46)
%_Bids_at_Reserve	12.54%	8.37%	26.59%
Yahoo DSP:			
No. Auctions	225,379	115,791	109,588
No. Bids	53,773	34,176	19,597
Bid (× \$1000)	1.18	0.94	1.43
	(0.78)	(0.76)	(0.71)
Winning_Bid (× \$1000)	1.72	1.59	1.89
	(0.86)	(0.91)	(0.77)
Winning_Frequency	8.74%	9.43%	8.02%
%_Bids_at_Reserve	6.17%	2.98%	11.74%
Shaded_Bid (%)	45.34%	44.61%	46.62%
	(17.65%)	(19.12%)	(14.67%)
Valuation (× \$1000)	2.35	2.36	2.34
	(39.72)	(41.23)	(38.05)

Notes: Standard deviations of variables are shown in parentheses.
The variables Bid, $Winning_Bid$, and $Valuation$ are in thousands of dollars.

main finding of this paper: 26.59% of the bids in the *Production* column are at the reserve price.

The bottom panel of Table 1 describes the additional data from the Yahoo DSP, which bids the *Production* reserve 11.74% of the time in the *Production* bucket. Under the Heuristic Bidding model described in the previous section, the Yahoo DSP enters auctions more often when the reserve is lower, submits bids exactly at reserve less often, and shades the bid down from valuation by a similar percentage amount. The bottom row of the panel can be interpreted as a randomization check: the auctions assigned the monitoring reserve do not differ from the auctions assigned the production reserve in terms of how much the Yahoo DSP values the opportunity ($2.36 versus $2.34). Comparing the two panels of Table 1, we see that the Yahoo DSP tends to bid more than the average bidder, and wins about 9 percent of the auctions it enters.

Table 2 summarizes our exchange bidding data by bidder type and reserve level. While both bidders' type face the same monitoring reserve price ($r_0 = \$0.15$), Rebroadcasters have a higher reserve price in production: $r_2 = \$1.15$ ¿ $r_1 = \$0.71$. Our dataset contains 874K valid bids submitted by Regular DSPs and 678K from Rebroadcasters. As expected, the lower reserve price value explains the larger number of submitted bids in the monitoring bucket independently of the bidder type. Interestingly, the overall average bid submitted by Regular DSPs is significantly lower than the average Rebroadcaster bid, even in the monitoring data stratum where both types face the same reserve price. We take this as preliminary evidence that the Rebroadcasters have generally higher valuations than the Regular DSPs, and the ad exchange is thus well advised to impose higher reserves on them. One reason why rebroadcasters have higher valuations is that they each represent the highest of several regular DSP bidders.

Rebroadcasters also bid at reserve much more often than Regular DSPs: Table 2 shows that when the reserve price is at the production level, 42.42% of Rebroadcaster bids equal reserve vs. only 14.79% of Regular-DSP bids.

6 Evidence for the Heuristic Bidding model

6.1 Direct evidence: the Yahoo DSP bidding function

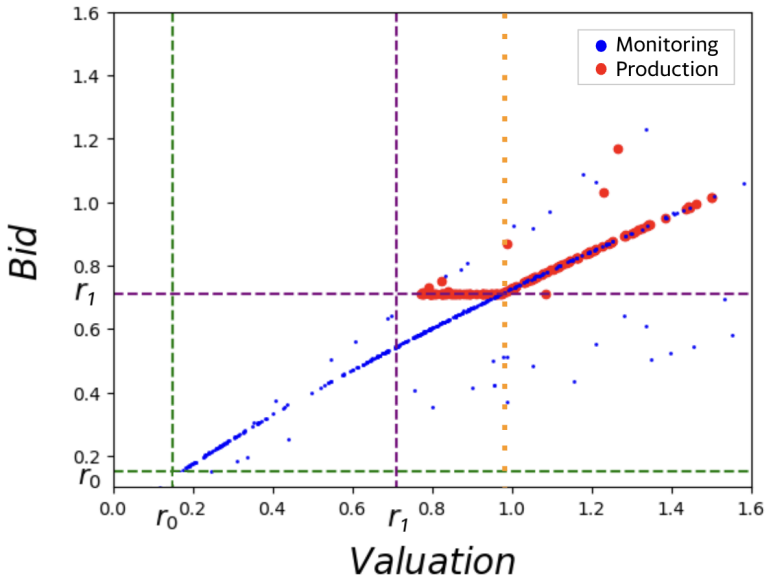
Yahoo DSP data provides information about its submitted bids and corresponding valuations for each opportunity and ad campaign, allowing us to infer Yahoo DSP bidding function. As an illustration, Figure 1 displays the observed valuations of a randomly chosen ad campaign with significant valuation support, and corresponding bids for monitoring and production reserve prices (blue and red dots respectively).

Table 2: Summary Statistics by Bidder Type

	All	Monitoring Reserve	Production Reserve
Regular DSPs		$r_0 = \$0.15$	$r_1 = \$0.71$
No. Bids	874,887	670,944	203,943
No. Auctions	180,325	108,699	71,626
Bid (× \$1,000)	0.66	0.53	1.10
	(0.50)	(0.42)	(0.44)
%_Bids_at_Reserve	13.00%	12.45%	14.79%
No. Bidders	3.89	5.80	1.86
	(3.29)	(3.08)	(2.05)
Winning_Bid (× \$1,000)	1.32	1.20	1.47
	(0.82)	(0.84)	(0.75)
Winning_Frequency %	41.75%	46.24%	37.02%
Rebroadcasters		$r_0 = \$0.15$	$r_2 = \$1.15$
No. Bids	678,705	526,877	151,828
No. Auctions	157,059	96,903	60,156
Bid (× \$1,000)	0.86	0.73	1.33
	(0.44)	(0.39)	(0.24)
%_Bids_at_Reserve	11.95%	3.17%	42.42%
No. Bidders	3.01	4.56	1.38
	(3.12)	(3.32)	(1.80)
Winning_Bid (× \$1,000)	1.27	1.13	1.46
	(0.46)	(0.45)	(0.42)
Winning_Frequency %	43.42%	49.28%	37.22%
No_Winners_Frequency	14.82%	4.47%	25.74%

Notes: The variables *Bid*, and *Winning_Bid* are in thousands of dollars.

Figure 1: Observed Yahoo DSP Bid Responses for a Particular Ad-Campaign



Note: Bid responses for *Monitoring* and *Production* buckets

Consistent with our proposed heuristic mechanism, for bid-levels strictly above the reserve price set in production ($r_1 = \$0.71$), the function mapping Yahoo DSP valuations to bids is the same independently on the reserve price enclosed in the bid request.

Differences arise for valuations between the production reserve level r_1 and the 98 cents valuations (vertical dotted line in orange) for which Yahoo DSP bids equal the reservation value r_1 when the exchange sends the production reserve price in the bid request. On the other hand, for requests with the monitoring reserve price r_0 and valuations within the same interval, Yahoo DSP bidding decision for the ad campaign is strictly increasing. This pattern is a direct evidence of our proposed mechanism whereby the DSP effectively censors the bids at the reserve price whenever the underlying candidate bid based on valuation alone does not clear at the reserve but the valuation of the opportunity does. Our data also shows that, in the production bucket the number of auctions where Yahoo DSP censors bids at the reserve is not negligible: 38% of the total number of responses in the ad campaign illustrated in Figure 1.

While Figure 1 shows only one campaign the pattern holds for other Yahoo DSP advertising campaigns more generally. Across all the Yahoo ad campaigns, 11.74% of bid responses facing production reserve prices are censored to the reserve price value, and bidding strategies in the two reserve-price buckets broadly coincide above the production reserve level.

6.2 Indirect Evidence: Disappearing Mass Point at Reserve and Right-Tail Comparisons

Having documented that the Heuristic model fits the behavior of the Yahoo DSP well, we now provide evidence that it is also consistent with the behavior of at least half of the other DSPs (based on tail-comparison tests of the bid distributions between the two reserve levels), if not a vast majority of them (based on disappearing mass point at the production reserve level). We begin with the partial pooling evidence.

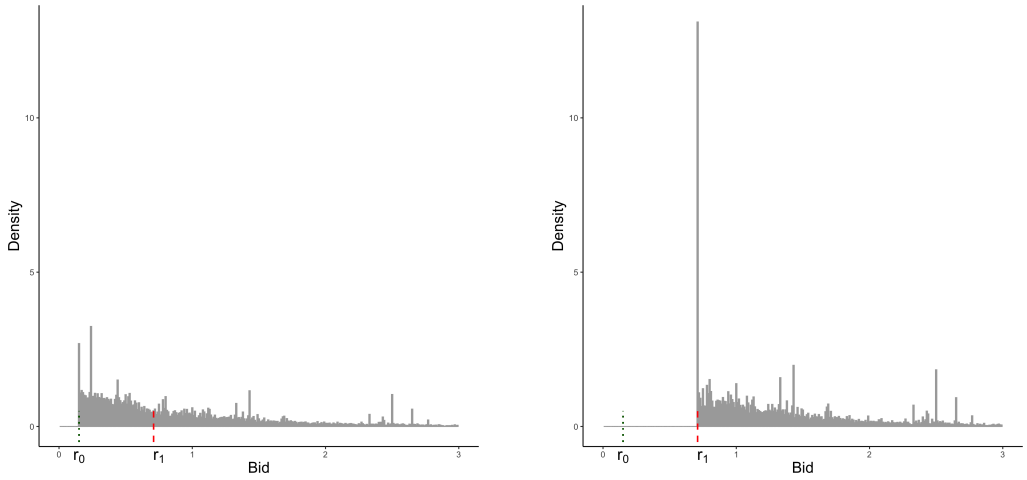
The data summaries in Tables 1 and 2 already include aggregate evidence that both types of bidders engage in partial pooling at the reserve price when the reserve price is set at the production level. In this section, we analyze the bidding behavior more thoroughly at the bidder level by examining the empirical distributions of bids submitted by different bidders, starting with the Yahoo DSP as a representative example.

Figure 2 displays the bidding distribution of the Yahoo DSP for a particular inventory as a histogram with bins one cent wide. As discussed earlier, each bidder in our sample randomly faces two reserve prices: a low monitoring level and a higher production level. Figure 2a captures Yahoo DSP bid responses when facing the *Monitoring* reserve ($r_0 = \$0.15$), while Figure 2b illustrates responses when the reserve is set at the *Production* level ($r_1 = \$0.71$). In both figures, we observe a mass of bid responses at the applied reserve price. But in Figure 2a, we do not observe a mass of bids at the *Production* reserve price value (vertical dashed red line at r_1). In the language of the introduction, this mass point seems to have "disappeared" when the reserve switched to the monitoring level. Note also that Figure 2a does not exhibit any mass points nearly as high (13%) as the r_1 mass-point in Figure 2b, so we can reject the possibility that the disappearing mass point has simply shifted to a different bid level. In summary, Figure 2 shows that Yahoo DSP considers the reserve price value when deciding how much to bid for the advertising opportunity, and it submits bids equal to the reserve price much more frequently than bids just above the reserve when the reserve is set at the production level.

By comparing the right tail of the bid distribution strictly above the highest of the two reserve prices in Figure 2a to the equivalently defined right tail of the bid distribution in Figure 2b (i.e., bid distributions strictly above r_1), we can test whether the bidding strategy depends on the reserve price beyond the disappearing mass point. We use the Kolmogorov-Smirnov test to compare the two distribution tails, and find that we cannot reject the hypothesis that both distributions are equal (see first row Table 3). Figure 3 shows the two cumulative distribution functions, and illustrates they are nearly identical to each other.

The data patterns described above are consistent with the Heuristic Bidding model, which employs the same univariate bidding strategy for high-enough valuations, and

Figure 2: Yahoo DSP Bid Distributions for Monitoring/Production Reserve Prices



(a) *Monitoring* Reserve Price (b) *Production* Reserve Price

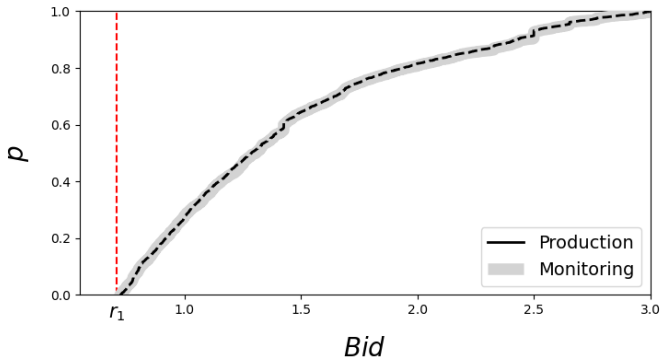
merely censors the bids at the reserve price whenever the valuation exceeds the reserve but the univariate candidate bid does not.

In addition to providing evidence for the Heuristic model, we also note that having a strategy that does not depend on reserve prices is inconsistent with previous theoretical work on first-price auctions. As Reiley (2006) and Lebrun (1999), among others, point out, the equilibrium bidding response should be strictly increasing with reserve prices. This inconsistency motivates our counterfactual exercise, where we explore differences between the Heuristic Bidding model that is consistent with the observational data and the expected behavior in the Bayes-Nash equilibrium.

The partial pooling bid pattern evidenced by Figure 2 is not unique to Yahoo DSP. Other bidders participating in the Yahoo exchange show similar behavior. Figures A.1, A.2, and A.3 in the Appendix show the bid distributions for *Monitoring* and *Production* reserve prices of several Regular DSPs and Rebroadcasters, which, together with Yahoo DSP, account for more than 80% of the total spend on the analyzed display inventory. For each bidder in the figures, we display the bid distributions for *Monitoring* and *Production* reserve prices in a manner analogous to Figure 2, and the corresponding cumulative distribution functions of bids strictly above the production reserve price in a manner analogous to Figure 3. The figures show that a mass of bids at the reserve price level is a common phenomenon among bidders, and in several cases, the bid distributions above the production reserve value do not depend on the applied reserve price.

Table 3 provides further evidence of bid pooling and the heuristic mechanism employed by bidders to determine bids. The table compares the *Monitoring* and

Figure 3: Bid Distributions Above Production Reserve: Production vs Monitoring Buckets



Notes: Cumulative Distribution Functions of bids in the two experimental buckets conditional on bids exceeding the Production reserve level plus one cent.

Production buckets of twelve bidders who together account for more than 80% of total ad spend on the analyzed inventory. For each *Bidder_id*, the table shows the bidder's *Type*, and the conditional percentages of submitted bids at the *Production* reserve price value in the two experimental buckets defined as follows:

- $P_0 = Pr\left(bid \in [r_{prod}, r_{prod} + \$0.01) | reserve = r_0, bid \geq r_{prod}\right)$
- $P_{prod} = Pr\left(bid \in [r_{prod}, r_{prod} + \$0.01) | reserve = r_{prod}, bid \geq r_{prod}\right)$

where *prod* stands for the production level of reserve, i.e. r_1 for regular DSPs and r_2 for Rebroadcasters. For example, Yahoo DSP's percentage of bids at or just above the *Production* reserve price ($r_{prod} = r_1 = \$0.71$) is 0.96% in the Monitoring bucket with reserve set to r_0, significantly lower than the 11.74% percentage in the *Production* bucket. The column denoted as $H_0 : P_0 = P_{prod}$ indicates whether we reject the hypothesis that both proportions are equal at a 5% significance level. Results suggest that in all cases, we do reject this hypothesis, and in all but one of the twelve cases, P_{prod} is thus statistically significantly larger than P_0. In other words, over 90 percent of the bidders we analyze exhibit the disappearing mass point at the production reserve level in their bid distributions.

Table 3 also provides evidence from the distributions of bids above the production reserve level (plus one cent). Column $KS(D)$ reports the two-sample Kolmogorov–Smirnov statistic, and the column $H_0 : KS(D) = 0$ indicates if we reject the hypothesis that both distributions are the same at the 5% significance level based on the p-value reported next to it.

If bidders do not adjust their high-value bids in response to reserve prices, we should not reject the null hypothesis that the right tails of the bidding distribution

are the same. For instance, the KS statistics for Yahoo DSP cannot reject this hypothesis. Considering all the bidders, we find that this tail-comparison test cannot reject their null hypotheses for six bidders. In summary, the tail-comparison tests summarized on the right side of Table 3 indicate that the bidding distributions of about half the bidders we analyze are consistent with the Heuristic model.

Table 3: Bidding Distributions Comparison between Reserve Price Levels

Bidder_Id	Type	P_0	P_{prod}	$H_0: P_0 = P_{prod}$	$KS(D)$	$pvalue$	$H_0: KS(D) = 0$
YahooDSP	RegularDSP	0.96%	11.74%	Reject	2.07E-02	0.44	No Reject
A	Rebroadcaster	2.61%	17.29%	Reject	2.52E-01	3.99E-95	Reject
B	Rebroadcaster	4.95%	39.26%	Reject	1.11E-01	2.44E-19	Reject
C	RegularDSP	0.76%	2.02%	Reject	1.66E-02	0.15	No Reject
D	RegularDSP	2.87%	5.44%	Reject	8.43E-02	1.59E-16	Reject
E	RegularDSP	1.75%	17.21%	Reject	1.08E-02	0.57	No Reject
F	RegularDSP	1.95%	38.51%	Reject	2.72E-02	0.58	No Reject
G	Rebroadcaster	2.22%	35.78%	Reject	1.92E-01	3.12E-47	Reject
H	Rebroadcaster	2.13%	79.97%	Reject	5.57E-02	0.21	No Reject
I	Rebroadcaster	0.98%	61.93%	Reject	3.25E-01	3.34E-88	Reject
J	Rebroadcaster	4.58%	30.66%	Reject	3.59E-02	0.07	No Reject
K	Rebroadcaster	1.91%	1.58%	Reject	4.08E-01	9.73E-80	Reject

Note: The Heuristic model implies a "Reject" of $H_0 : P_0 = P_{prod}$ and "No Reject" of $H_0 : KS(D) = 0$.
All statistical tests are performed at 5% significance level.

7 Estimation of Rebroadcaster Valuations and Heuristic Bidding Strategies of Both Bidder Types

The available tractable Bayes Nash equilibrium model of Kotowski (2018) only admits two bidder types, with all bidders of the same type having the same distribution - an assumption we thus need to enforce in our counterfactual exercise. To perform our counterfactual analysis, we therefore first need to estimate the distribution of valuations for both types of bidders. We assume that Regular DSPs all have the same valuation distribution as Yahoo DSP, which is observed in our dataset. For Rebroadcasters, we use a structural model to recover their valuations. Specifically, we assume all Rebroadcasters follow the Heuristic Bidding model, and we use bidding data from the *Monitoring* bucket with low reserve price r_0 to infer their valuations. This section first explains the estimation procedure, and then summarizes the estimates.

7.1 Rebroadcaster Valuation Distribution: Estimation Procedure

We use tools developed in previous literature to recover the distribution of Rebroadcaster valuations $F_2[\cdot]$. Since both bidder types face the same reserve price r_0 in the *Monitoring* bucket, we can follow the two step nonparametric procedure suggested by Flambard and Perrigne (2006) to recover $F_2[\cdot]$. In this section, we outline the details of our estimation procedure. Flambard and Perrigne (2006) analyze a lowest-bid procurement auction, so we need to adapt their procedure to our highest-bid auction as follows: Let $\tilde{\phi}_i = \tilde{b}_i^{-1}$ be the inverse of the univariate bidding function of type i in Equation 2. The Heuristic Bidding model then proposes that Rebroadcasters derive their univariate bidding function $\tilde{b}_2(v)$ by maximizing their expected surplus in an auction without a reserve:

$$(v-b)F_1[\tilde{\phi}_1(b)]^{N_1}F_2[\tilde{\phi}_2(b)]^{N_2-1} \tag{8}$$

Recall that in our dataset, we observe the distribution of valuations F_1 of the Yahoo DSP, as well as its univariate bidding strategy $\tilde{b}_1(v)$. In this estimation procedure, we assume all the other Regular DSPs (bidders of type $i = 1$) have the same distribution of valuations and use the same univariate bidding functions when formulating their heuristic bids. Given the known inverse bid function of Regular DSPs $\tilde{\phi}_1$, the function $\tilde{\phi}_2$ is thus the unique solution to the corresponding first-order condition[7]:

$$\left[\tilde{\phi}_2(b_2) - b_2\right]\left[\frac{N_1 f_1\left[\tilde{\phi}_1(b_2)\right]\tilde{\phi}_1'(b_2)}{F_1[\tilde{\phi}_1(b_2)]} + \frac{(N_2-1)f_2\left[\tilde{\phi}_2(b_2)\right]\tilde{\phi}_2'(b_2)}{F_2[\tilde{\phi}_2(b_2)]}\right] = 1 \tag{9}$$

Satisfying the boundary conditions

- $\tilde{\phi}_1(r_1) = \tilde{v}_1$ and $\tilde{\phi}_2(r_2) = \tilde{v}_2$
- $\tilde{\phi}_1(\gamma_1^*) = \tilde{\phi}_2(\gamma_2^*) = \bar{v}$ for some $\gamma_1^* \in (r_1, \bar{v})$ and $\gamma_2^* \in (r_2, \bar{v})$

Flambard and Perrigne (2006) show that the differential Equation 9 can be rewritten in terms of bid distributions. The binding reserve introduces a complication because we do not observe observe Rebroadcaster bids below r_0. Therefore, we can only recover the truncated distribution of valuations above this reserve. Let $G_i^*(b) = (F_i[\phi_i(b)] - F_i[r_0])/(1 - F_i[r_0])$ be the distributions for each bidder type

[7]This first-order condition is the analogue of Equation 3 in Flambard and Perrigne (2006) in our auction setting.

$i \in \{1,2\}$ truncated below at r_0, and let g_i^* be the corresponding densities $g_i^*(b) = (f_i[\phi_i(b)]\phi_i'(b))/(1 - F_i[r_0])$. Then, following Flambard and Perrigne (2006), express all elements in Equation 9 in terms of the truncated bid distributions, and invert the first-order condition as follows:

$$v_2 = b_2 + \frac{1}{\frac{N_1 g_1^*(b_2)(1-F_1[r_0])}{(1-F_1[r_0])G_1^*(b_2)+F_1[r_0]} + \frac{(N_2-1)g_2^*(b_2)(1-F_2[r_0])}{(1-F_2[r_0])G_2^*(b_2)+F_2[r_0]}} \tag{10}$$

Proposition 1 in Flambard and Perrigne (2006) shows that if certain conditions hold, Equation 10 nonparametrically identifies $F_2[\cdot]$ on $[r_0, \bar{v}]$ from the distributions of observed bids G_i^*.

We handle N_i and $F_i[r_0]$, the remaining parameters in Equation 10, as follows: for N_i, we use the maximum number of observed participants $N_{i,t}^*$ of each type i across all observed auctions T as an estimate of the number of potential bidders N_i , which are assumed to be constant across all auctions. That is, for N_i we set the estimator $\hat{N}_i = max_{t=1,...,T} N_{i,t}^*$ in Equation 10. As mentioned above, $F_1[\cdot]$ is observed in our data on all support $[0, \bar{v}]$, and, consequently, $F_1[r_0]$ is also observed. For $F_2[r_0]$, we rely on the fact that the expected number of observed participants $E(N_{2,t}^*) = [1 - F_2(r_0)]N_2$. Solving for the unknown probability, we thus let $F_2[r_0]$ equal one minus the ratio between the average observed number of participants and the number of potential bidders as estimator: $\hat{F}_2[r_0] = 1 - \sum_{t=1}^{T} N_{2,t}^*/(T\hat{N}_2)$.[8]

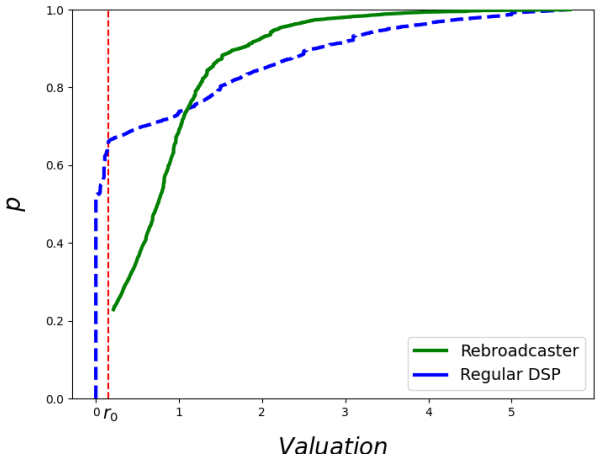
7.2 Rebroadcaster Valuation Distribution: Estimation Results

To capture changes in impressions opportunities patterns during the day, we perform our analysis hourly. As discussed in the previous section, we assume that the valuation distribution of regular bidders is the same as the observed one from Yahoo DSP. For Rebroadcasters, we use Equation 10 to recover the valuations, and use nonparemetric kernel techniques to estimate the corresponding distribution.

As an illustration, Figure 4 shows the valuation cumulative distribution functions of the Regular DSPs (dashed blue line) and Rebroadcasters (solid green line) during the 13th hour of the day (UTC time). Figures A.4 and A.5 in the Appendix display the distributions of bidders for each hour of the day, and they all look qualitatively

[8]Note that in the estimation of F_2, we interpret the small mass of rebroadcaster bids at the monitoring reserve as arising from valuations equal to r_0. This simplifying assumption does not impact our counterfactual simulations because those simulations care about the bidding behavior at much higher valuation levels, namely above the production reserve prices (r_1 and r_2), which are both much larger than r_0.

Figure 4: Valuation CDFs: Regular DSPs and Rebroadcasters (Hour 13th)



Notes: Plots of the valuation CDF of Regular DSPs and the truncated valuation CDF of Rebroadcasters.

similar. The Rebroadcasters' cumulative distribution function is truncated for valuations below r_0 (dashed red line).[9] This is not the case for Regular DSPs for which we do observe Yahoo DSP valuations independently of the submitted bids. Note that not all opportunities are valuable for Regular DSPs. In particular, Figure 4 shows that around 50% of opportunities have no value for Regular DSPs.

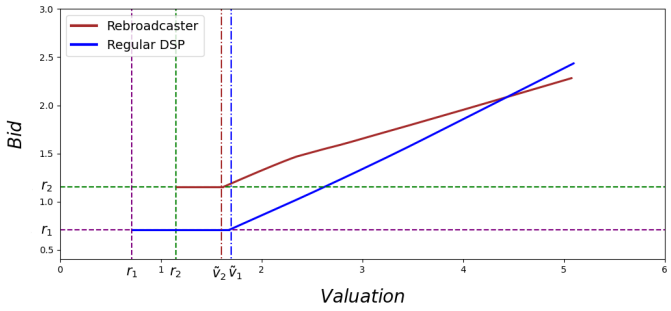
One way to interpret Figure 4 and its 23 analogous figures in the Appendix corresponding to the remainig hours of the day is that the Rebroadcaster valuations are higher than the valuations of Regular DSPs. This makes sense given the fact that Rebroadcasters run their own auctions with regular DSP bidders. The figure also shows that the Rebroadcaster valuations are more concentrated, with a thinner right tail.

7.3 Estimated Heuristic Bidding Strategies

Given the known valuation distribution $F_1[\cdot]$ of the Regular DSPs and the estimated distribution $F_2[\cdot]$, we can infer the heuristic bidding strategy of the Rebroadcasters $\tilde{b}_2(\cdot)$ characterized in Equation 2. To obtain the corresponding Regular DSP bid function $\tilde{b}_1(\cdot)$ in Equation 2, we use OLS and regress the Yahoo DSP bids on valuations and the square of valuations, using observations with bids strictly above r_0.

[9]Kernel methods struggle to accurately estimate bidding densities at the boundaries. This explains why the Rebroadcasters' cumulative distribution function does not reach the reserve price r_0 in Figure 4. To address this problem, we employ a mirroring technique to improve bidding density estimation.

Figure 5: Current Heuristic Bidding Functions (Hour 13th)



Notes: Plots of the current heuristic bid functions $\tilde{\beta}_1$ and $\tilde{\beta}_2$ of Regular DSPs and Rebroadcasters respectively

Given $\tilde{b}_i(\cdot)$, building $\tilde{\beta}_i(\cdot|r_i)$ is straightforward for any reserve price value $r_i \geq r_0$. This is specially useful in our counterfactual exercise where we compare the publisher's revenue and bidders' surplus of the heuristic bidding strategy currently used in production with the counterfactual Bayes-Nash equilibrium strategy.

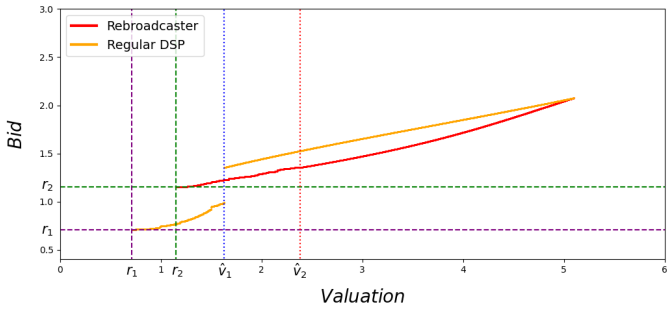
Figure 5 displays the observed heuristic bidding function $\tilde{\beta}_1(\cdot|r_1)$ of the Regular DSPs (blue line) and the estimated heuristic bidding function $\tilde{\beta}_2(\cdot|r_2)$ of the Rebroadcasters (purple line) during the 13th hour of the day when facing the asymmetric reserve prices used in production (r_1 and r_2), as characterized in Equation 2. As previously discussed, each bidder type does not participate in the auction if valuations are below the assigned reservation value. Bidders submit a bid that is equal to the reserve price when valuations are above the reserve price and below the corresponding thresholds $\tilde{v}_1$ and $\tilde{v}_2$, which are equal to \$1.84 and \$1.67 respectively. Finally, bidding functions are strictly increasing for valuations above $\tilde{v}_1$ and $\tilde{v}_2$.

8 Counterfactual Simulation: Bayes-Nash Equilibrium Model

Suppose all bidders use the Heuristic Bidding model. How would the auction revenue change if they all switched to Bayes Nash equilibrium strategies? And what if the auctioneer also adjusted the reserve prices to maximize revenue given the new equilibrium strategies? In this section, we answer these two counter-factual questions by comparing the expected revenues under the Heuristic model of Section 4.3 to revenues under the Bayes-Nash Equilibrium model of Section 4.4.

Given valuations and corresponding distribution estimates, we can simulate the impact of bidders adopting a Bayes Nash equilibrium bidding strategy ($\beta_1(\cdot|r_1)$ and $\beta_2(\cdot|r_2)$) characterized in Equation 4. The portion of the function captured by the

Figure 6: Bidding Functions: Bayes-Nash Equilibrium (Hour 13th)



Notes: Plots of the Bayes-Nash bid functions β_1 and β_2 of Regular DSPs and Rebroadcasters respectively

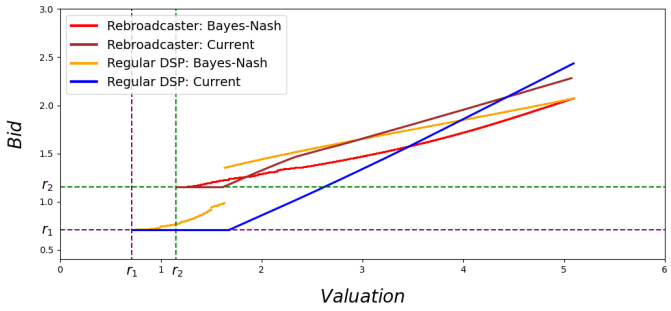
function $b_i(\cdot)$ can be nonparametrically estimated using Equation 5. For $\hat{b}_i(v)$, we follow Bajari (2001) and Hubbard et al. (2013), and use a polynomial approximation of the functions $\phi_1(\cdot)$ and $\phi_2(\cdot)$ (see Section A.1 in Appendix for details). Note that the Bayes-Nash equilibrium characterization in Equation 4 is valid if inequality in Equation 3 holds. We confirmed this condition holds for all hours (see Table A.1 in the Appendix).

As in the previous section, figures below illustrate the results for the 13th hour of the day. Figure 6 shows the Bayes-Nash equilibrium bidding strategy. The orange lines correspond to the Bayes-Nash bidding function employed by Regular DSPs ($\beta_1(\cdot|r_1)$), and the red line captures Rebroadcasters bidding function ($\beta_2(\cdot|r_2)$), evaluated at the production reserve prices r_1 and r_2 respectively. The functions are the result of the equilibrium conditions characterized in Equation 4. Both bidding functions are strictly increasing starting at the corresponding reserve price: r_1 for Regular DSPs, and r_2 for Rebroadcasters. As discussed in Kotowski (2018), there is a discontinuity in the Regular DSP bidding function at the valuation $\hat{v}_1$ that is equal to \$1.68. At $\hat{v}_1$, Regular DSPs are indifferent between bidding as if Rebroadcasters did not participate in the auction (Equation 5), and fully internalizing Rebroadcasters' competition (differential Equation 6).

Figure 7 overlaps the current heuristic and the Bayes-Nash bidding strategies. It shows that under the Heuristic model, Rebroadcasters tend to bid more than in the Bayes-Nash equilibrium while Regular DSPs tend to bid less. This is a common pattern across all hours of the day (see Figures A.6 and A.7 in the Appendix). We therefore expect that switching to the Bayers-Nash would make Regular DSPs win more auctions. But to understand the impact of the switch on bidder surpluses and publisher revenue, we need to perform counterfactual simulations to which we now turn.

The impact of different bidding strategies on metrics like publisher revenue and bidder surplus depends on the valuation distributions. To estimate the counterfactual,

Figure 7: Bidding Functions: Bayes-Nash vs Current Heuristic (Hour 13th)



Notes: Plots of the current and Bayes-Nash bid functions of Regular DSPs and Rebroadcasters

we proceed as follows: First, we recover valuations by sampling auction data from the monitoring bucket with reserve price r_0. Each sample contains information about the number of participants of each type and their corresponding bids, allowing us to recover valuations using the respective estimated inverse heuristic bidding functions. We then compute bid responses for both the heuristic and the Bayes-Nash bidding strategies when reserve prices are set to production levels (r_1 and r_2), enabling us to determine the average bid, bidder surplus, and corresponding publisher revenue for each strategy. We simulate 48,000 auctions and report the results in Tables 4 and 5, with standard deviations in parentheses.

Table 4 compares the overall impact of switching from the current heuristic to the Bayes-Nash bidding strategy under two scenarios: one with current reserve prices and another where the publisher adjusts reserve prices to maximize its revenue[10]. The change would benefit publishers, increasing their expected auction revenue by 7 cents in the Bayes-Nash equilibrium with current reserve prices, and 12 cents when the publisher adjusts reserve prices (a 5.14% and 8.83% increase respectively). However, this comes at a cost to bidders, who would experience a reduction in surplus per auction. In particular in the Bayes-Nash equilibrium with optimized reserve prices (a 6-cent reduction, representing a 5.08% decrease). Current and the Bayes-Nash equilibrium without reserve prices adjustments strategies maintain a similar proportion of auctions (around 4.83%) with no bids above the reserve price, as bidders in both cases submit bids whenever their valuations exceed the reserve prices. This is no longer the case when the publisher adjusts reserve prices. Since the new reserve prices are higher, the number of non-cleared auctions increases (from 4.84% to 6.76%).

[10]In the Bayes-Nash equilibrium with optimized reserve prices, we employ a grid search to identify reserve prices r_1 and r_2 that maximize the publisher's expected revenue. For each candidate pair (r_1, r_2), we calculate the expected revenue under the resulting Bayes-Nash equilibrium and select the pair yielding the highest revenue. In our analysis the optimal reserve prices are \$0.91 for r_1 and \$1.25 for r_2.

Table 4: Counterfactual simulation of Bayes-Nash Equilibrium

	Current	Bayes-Nash	Lift ($)	Lift (%)	Bayes-Nash Reserve Price	Lift ($)	Lift (%)
Publisher revenue per request	1.36	1.43	0.07	5.14	1.48	0.12	8.83
	(0.44)	(0.42)	(0.11)		(0.38)	(0.23)	
Bidder surplus per request	1.18	1.17	-0.01	-0.8	1.12	-0.06	-5.08
	(0.83)	(0.81)	(0.30)		(0.98)	(0.32)	
% Non-Cleared Auctions	4.84%	4.83%		-0.21	6.76%		39.67

Notes: Simulated results are based on 48,000 auctions randomly drawn from the data.

Table 5: Counterfactual simulation of Bayes-Nash Equilibrium by Bidder Type

	Current	Bayes-Nash	Lift ($)	Lift (%)	Bayes-Nash Reserve Price	Lift ($)	Lift (%)
Winning Frequency							
Regular DSPs	39.95%	57.96%		18.01%	57.81%		17.86%
Rebroadcaster	56.39%	38.20%		-18.19%	36.43%		-19.96%
Surplus per request							
Regular DSPs	0.73	0.83	0.10	13.70	0.79	0.06	8.22
	(0.91)	(0.88)	(0.47)		(1.08)	(0.48)	
Rebroadcaster	0.44	0.34	-0.1	-22.73	0.33	-0.11	-25.00
	(0.60)	(0.61)	(0.30)		(0.59)	(0.32)	
% All Bids Below Reserve							
Regular DSPs	21.95%	21.95%		0.00%	24.30%		2.35%
Rebroadcaster	28.00%	28.00%		0.00%	33.50%		5.50%

Notes: Simulated results are based on 48,000 auctions randomly drawn from the data.

Table 5 breaks down these effects by bidder type, revealing substantial heterogeneity. Regardless of reserve-price optimization, Regular DSPs win much more often (57% of the auctions vs. 39% currently - an increase of about 46 percent) at the expense of a flip-side reduction in the winning chances of Rebroadcasters. The increased chance of winning naturally gives Regular DSPs more surplus, with the reserve price optimization clawing some of it back: at the status-quo reserve prices, Regular DSPs would be about 14 percent better off under the Bayes-Nash Equilibrium model, but this gain is reduced to about 8 percent once the auctioneer is allowed to re-optimize the reserves. On the losing side of the market, the Rebroadcasters would lose almost a quarter of their surplus either way. Finally, Table 5 also shows that reserve-price re-optimization in response to the switch from heuristic to equilibrium bidding would also slightly reduce market efficiency, increasing the percentage of auctions with no bids from each bidder type above their respective reserve prices (by between 2 and 5 percentage points).

9 Discussion

We document a simplified bidding heuristic that leads to partial pooling at the reserve price in auction markets for digital display advertising. Our evidence is both direct - in the form of the actual bidding strategy observed for one of the bidders - and indirect in the form of two patterns in the distribution of bids: 1) a mass point at reserve that disappears whenever the reserve randomly shifts to a much lower level, and 2) a right tail of the bid distribution that does not change when the the reserve randomly shifts to a much lower level.

To understand the economic consequences of this non-equilibrium heuristic behavior and its associated partial pooling of bids at reserve, we use a counterfactual analysis to compare the publisher's revenue and the bidders' surplus under the heuristic and the Bayes-Nash Equilibrium bidding strategies. We find that the bidding heuristic with partial pooling can lead to suboptimal outcomes for publishers relative to Bayes-Nash equilibrium strategies. Specifically, we find that all bidders switching to a Bayes-Nash equilibrium bidding strategies could increase publisher revenue by about 5 percent, and allowing the publisher to then re-optimize its reserve prices could increase its revenue by 9 percent. These estimates suggest that the heuristic bidding strategies we document benefit the advertisers by generally keeping bids artificially low compared to what they would be if bidders were fully rational and did not face any computational constraints.

The general increase in bids due to the switch from heuristic to equilibrium bidding would also result in a general increase in reserve prices compared to the status quo levels. Specifically, we find that reserve prices of both bidder types would rise, and the chance none of the bids exceed the reserves would increase by about two percentage points from the current 4.8% level. Therefore, the potential benefit to publishers from more strategic bidding would likely come with a decrease in overall market efficiency.

We find substantial heterogeneity in the effects the switch in bidding strategies would have on the two types of bidders we analyze. The type of bidder who would increase its bids most under the equilibrium model is the Regular DSP, who would also end up winning more auctions and receive higher expected surplus. On the flip-side, Rebroadcasters would experience a large decrease (25%) in surplus under Bayes-Nash equilibrium bidding, and a dramatic reduction (19.96%) in the chance of winning.

Our study underscores the importance of understanding real-world bidding behavior in online ad auctions. Classical auction theory, while providing a useful framework, does not always accurately predict how bidders act in practice. The prevalence of simpler, heuristic-based strategies, like the one documented and analyzed in this paper, highlights the need for more empirical research to inform the development of more effective auction mechanisms.

References

Miguel Alcobendas and Robert Zeithammer. 2021. Slim Shading in Ad Auctions: Adjustment of Bidding Strategies to First-Price Rules. *SSRN Working Paper* (March 2021). https://ssrn.com/abstract=403600 Available at SSRN.

Patrick Bajari. 2001. Comparing Competition and Collusion: A Numerical Approach. *Economic Theory* 18, 1 (2001), 187–205. http://www.jstor.org/stable/25055415

Stylianos Despotakis, R. Ravi, and Amin Sayedi. 2021. First-Price Auctions in Online Display Advertising. *Journal of Marketing Research* 58, 5 (2021), 888–907. https://doi.org/10.1177/00222437211030201 arXiv:https://doi.org/10.1177/00222437211030201

Véronique Flambard and Isabelle Perrigne. 2006. Asymmetry in Procurement Auctions: Evidence from Snow Removal Contracts. *Economic Journal* 116, 514 (2006), 1014–1036. https://EconPapers.repec.org/RePEc:ecj:econjl:v:116:y:2006:i:514:p:1014-1036

Shumpei Goke, Gabriel Y. Weintraub, Ralph Mastromonaco, and Sam Seljan. 2024. Bidders' Responses to Auction Format Change in Internet Display Advertising Auctions. *SSRN Working Paper* (June 2024). https://ssrn.com/abstract=4021221 Available at SSRN.

Negin Golrezaei, Ilan Lobel, and Renato Paes Leme. 2021. Auction Design for ROI-Constrained Buyers. In *Proceedings of the Web Conference 2021* (Ljubljana, Slovenia) *(WWW '21)*. Association for Computing Machinery, New York, NY, USA, 3941–3952. https://doi.org/10.1145/3442381.3449841

Philip A Haile. 2000. Partial Pooling at the Reserve Price in Auctions with Resale Opportunities. *Games and Economic Behavior* 33, 2 (2000), 231–248. https://doi.org/10.1006/game.1999.0781

Timothy P. Hubbard, René Kirkegaard, and Harry J. Paarsch. 2013. Using Economic Theory to Guide Numerical Analysis: Solving for Equilibria in Models of Asymmetric First-Price Auctions. *Computational Economics* 42 (2013), 241–266.

Maciej H. Kotowski. 2018. On asymmetric reserve prices. *Theoretical Economics* 13, 1 (2018), 205–237. https://doi.org/10.3982/TE1824 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.3982/TE1824

Stephan Lauermann and Andre Speit. 2023. Bidding in Common-Value Auctions With an Unknown Number of Competitors. *Econometrica* 91, 2 (2023), 493–527. https://doi.org/10.3982/ECTA17793 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA17793

Bernard Lebrun. 1999. First Price Auctions in the Asymmetric N Bidder Case. *International Economic Review* 40, 1 (1999), 125–142. http://www.jstor.org/stable/2648842

Bernard Lebrun. 2006. Uniqueness of the equilibrium in first-price auctions. *Games and Economic Behavior* 55, 1 (2006), 131–151. https://doi.org/10.1016/j.geb.2005.01.006

Alessandro Lizzeri and Nicola Persico. 2000. Uniqueness and Existence of Equilibrium in

Auctions with a Reserve Price. *Games and Economic Behavior* 30, 1 (2000), 83–114. https://doi.org/10.1006/game.1998.0704

Eric Maskin and John Riley. 2003. Uniqueness of equilibrium in sealed high-bid auctions. *Games and Economic Behavior* 45, 2 (2003), 395–409. https://doi.org/10.1016/S0899-8256(03)00150-7 Special Issue in Honor of Robert W. Rosenthal.

David H. Reiley. 2006. Field Experiments on the Effects of Reserve Prices in Auctions: More Magic on the Internet. *The RAND Journal of Economics* 37, 1 (2006), 195–211. http://www.jstor.org/stable/25046235

J. G. Riley and W. F. Samuelson. 1981. Optimal Auctions. *American Economic Review* 71, 3 (1981), 381–392.

A Appendix

A.1 Details of Polynomial Approximation of $\hat{b}_i(v)$

For $\hat{b}_i(v)$ in Equation 4, we follow Bajari (2001) and Hubbard et al. (2013), and use the following polynomial to approximate the functions ϕ_1 and ϕ_2 :

$$\phi_i(b) = \eta^* + \sum_{m=0}^{3} \alpha_{i,m}(b - \eta^*)^m \tag{11}$$

for $i \in \{1, 2\}$. η^* and coefficients $\alpha_{i,m}$ for $m \in \{0, 1, 2, 3\}$ are parameters to be estimated along with the valuation thresholds $\hat{v}_1$ and $\hat{v}_2$. To recover the parameters we derive the function $H_{i,t}(\cdot)$ from Equation 4 as follows

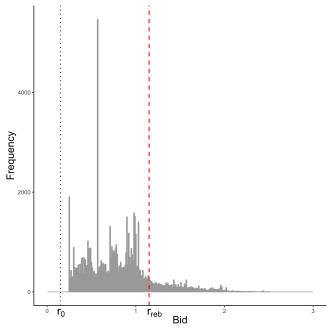
$$H_{i,t}(b, \alpha_i, \alpha_j, \hat{v}_i, \eta^*) = \begin{cases} 0 & if\ \phi_i(b) < r_i \\ b_i(\phi_i(b)) - \phi_i(b) + \int_{r_i}^{\phi_i(b)} \left[\frac{\hat{F}_i(z)}{\hat{F}_i(\phi_i(b))}\right]^{N_i - 1} dz & if\ \phi_i(b) \in [r_i, \hat{v}_i] \\ 1 - [\phi_i(b) - b]\left[\frac{N_j \hat{f}_j[\phi_j(b)]\phi_j'(b)}{\hat{F}_j[\phi_j(b)]} + \frac{(N_i - 1)\hat{f}_i[\phi_i(b)]\phi_i'(b)}{\hat{F}_i[\phi_i(b)]}\right] & if\ \phi_i(b) > \hat{v}_i \end{cases} \tag{12}$$

for $i \neq j$ and $i, j \in \{1, 2\}$. $\hat{F}_i[\cdot]$, $\hat{F}_j[\cdot]$, $\hat{F}_i[\cdot]$ and $\hat{f}_j[\cdot]$ correspond to the cumulative distribution function and probability distribution function estimates of the valuations. $\alpha_i = (\alpha_{i,0}, \alpha_{i,1}, \alpha_{i,2}, \alpha_{i,3})$ and $\alpha_j = (\alpha_{j,0}, \alpha_{j,1}, \alpha_{j,2}, \alpha_{j,3})$ correspond to the set of parameters defining the inverse bidding functions in Equation 11.

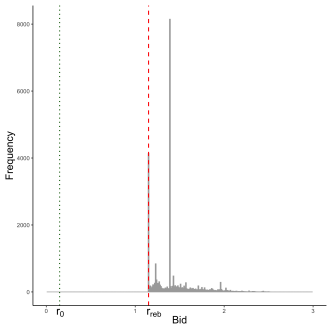
We solve the following constrained optimization problem to obtain α_1, α_2, $\hat{v}_1$, $\hat{v}_2$ and η^*

$$\begin{aligned} \min_{\alpha_1, \alpha_2, \hat{v}_1, \hat{v}_2, \eta^*} & \sum_i \sum_t [H_{i,t}(b, \alpha_i, \alpha_j, \hat{v}_i, \eta^*)]^2 \\ \text{s.t.} \quad & \phi_1(\eta^*) = \phi_2(\eta^*) = \bar{v} \\ & \phi_1(b_2(\hat{v}_2)) = \hat{v}_1 \\ & \phi_2(b_2(\hat{v}_2)) = \hat{v}_2 \end{aligned} \tag{13}$$

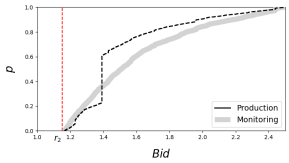
Figure A.1: Bid Distributions for Low/High Reserve Prices Bidders



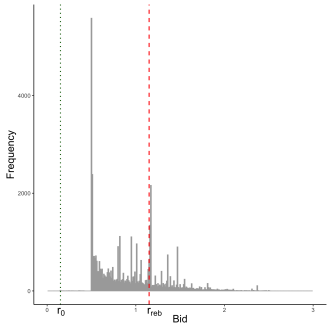
(a) Low Reserve Rebroadcaster A



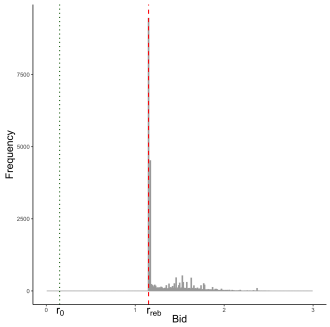
(b) High Reserve Rebroadcaster A



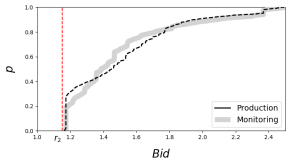
(c) CDFs Rebroadcaster A



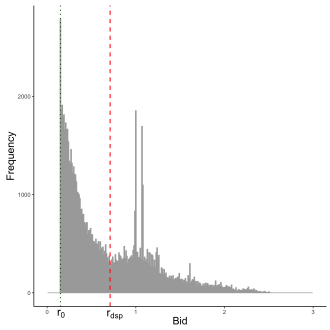
(d) Low Reserve Rebroadcaster B



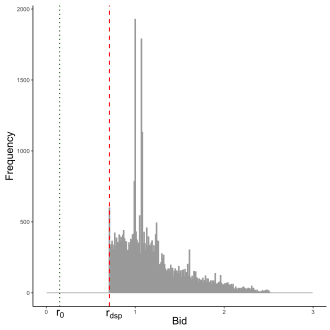
(e) High Reserve Rebroadcaster B



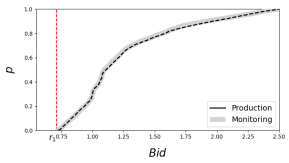
(f) CDFs Rebroadcaster B



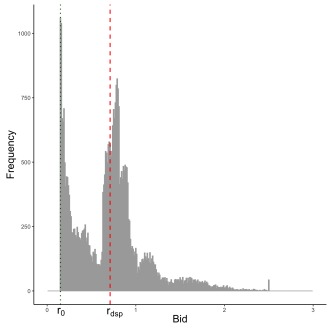
(g) Low Reserve Regular DSP C



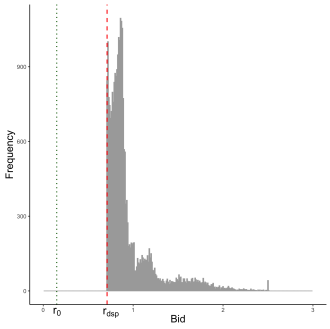
(h) High Reserve Regular DSP C



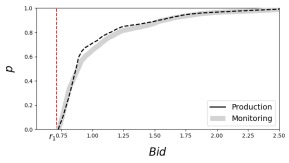
(i) CDFs Rebroadcaster C



(j) Low Reserve Regular DSP D

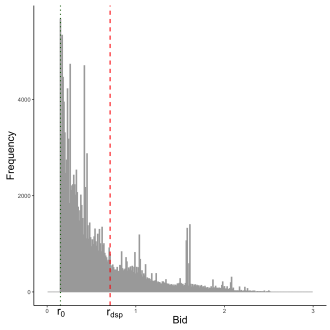


(k) High Reserve Regular DSP D

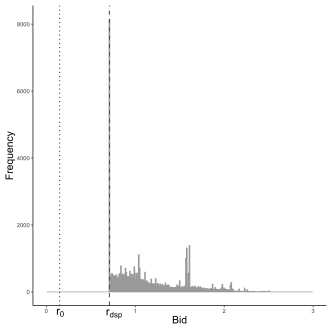


(l) CDFs Rebroadcaster D

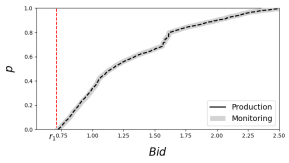
Figure A.2: Bid Distributions for Low/High Reserve Prices Bidders



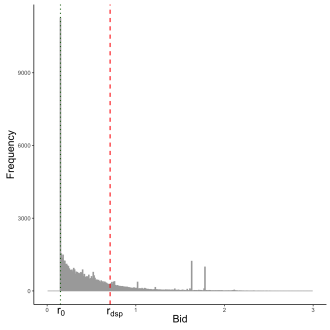
(a) Low Reserve Regular DSP E



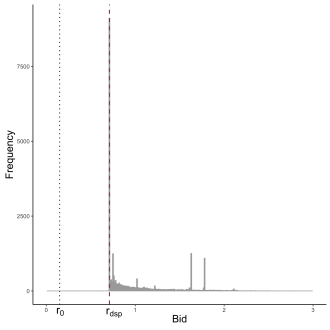
(b) High Reserve Regular DSP E



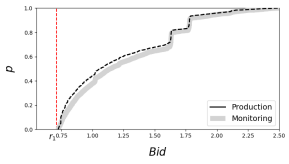
(c) CDFs Rebroadcaster E



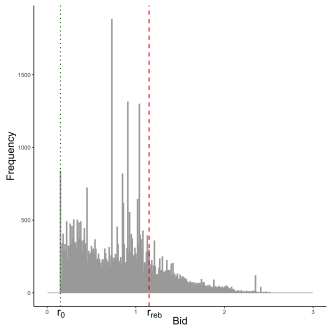
(d) Low Reserve Regular DSP F



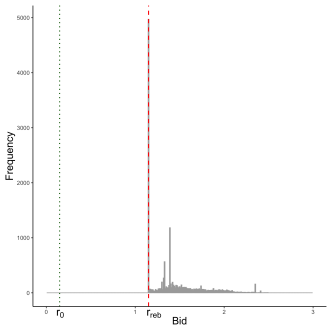
(e) High Reserve Regular DSP F



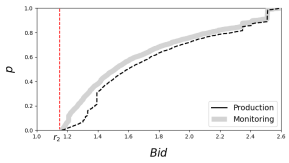
(f) CDFs Rebroadcaster F



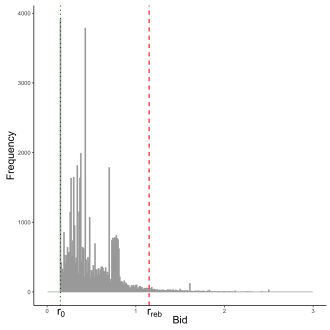
(g) Low Reserve Rebroadcaster G



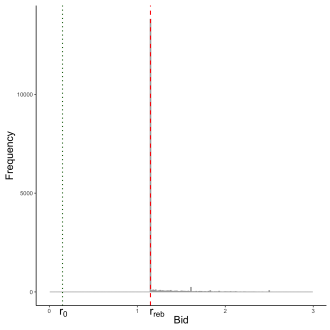
(h) High Reserve Rebroadcaster G



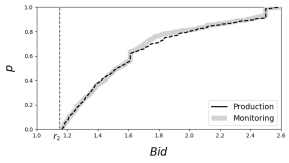
(i) CDFs Rebroadcaster G



(j) Low Reserve Rebroadcaster H

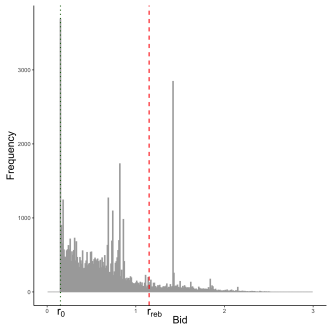


(k) High Reserve Rebroadcaster H

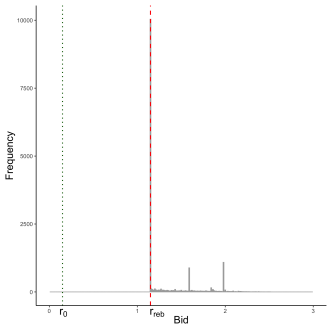


(l) CDFs Rebroadcaster H

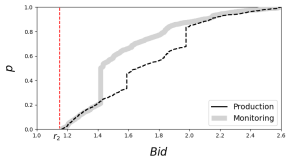
Figure A.3: Bid Distributions for Low/High Reserve Prices Bidders



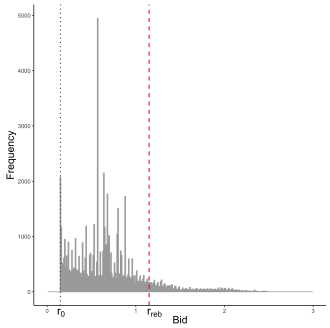
(a) Low Reserve Regular DSP I



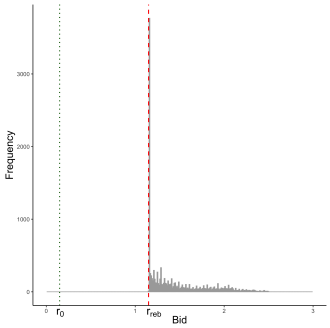
(b) High Reserve Regular DSP I



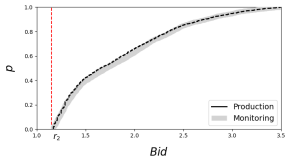
(c) CDFs Rebroadcaster I



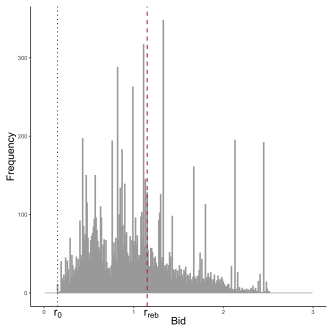
(d) Low Reserve Regular DSP J



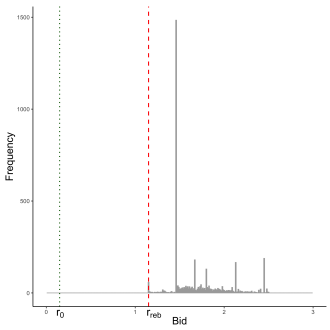
(e) High Reserve Regular DSP J



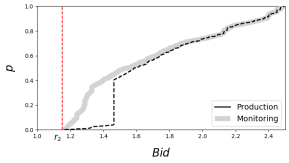
(f) CDFs Rebroadcaster J



(g) Low Reserve Rebroadcaster K

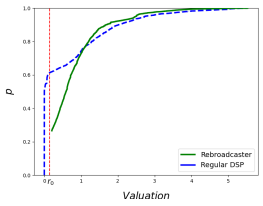


(h) High Reserve Rebroadcaster K

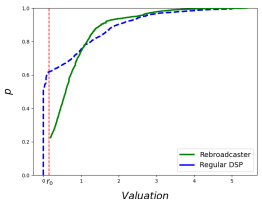


(i) CDFs Rebroadcaster K

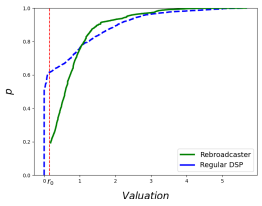
Figure A.4: Valuation CDFs per Hour 0-11: Regular DSPs and Rebroadcasters



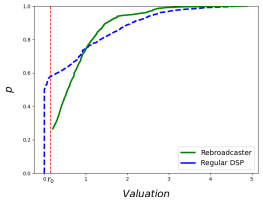
(a) CDFs Hour 0



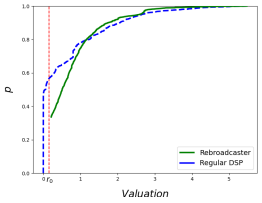
(b) CDFs Hour 1



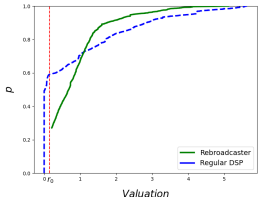
(c) CDFs Hour 2



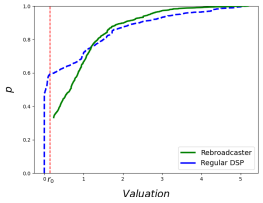
(d) CDFs Hour 3



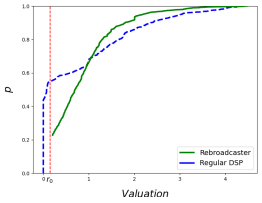
(e) CDFs Hour 4



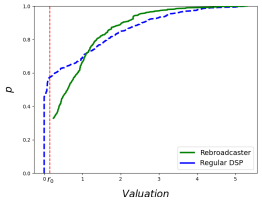
(f) CDFs Hour 5



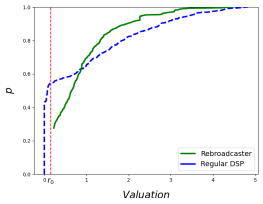
(g) CDFs Hour 6



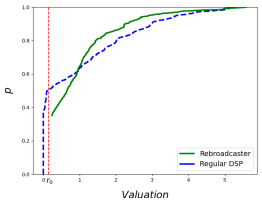
(h) CDFs Hour 7



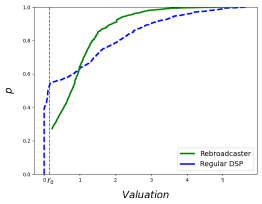
(i) CDFs Hour 8



(j) CDFs Hour 9

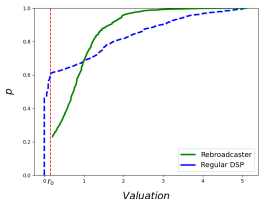


(k) CDFs Hour 10

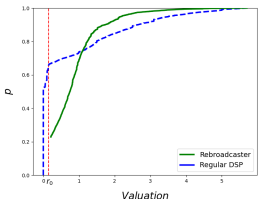


(l) CDFs Hour 11

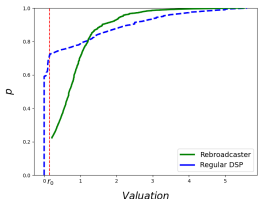
Figure A.5: Valuation CDFs per Hour 12-23: Regular DSPs and Rebroadcasters



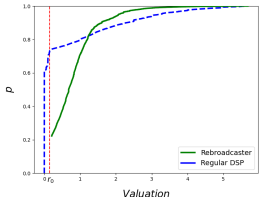
(a) CDFs Hour 12



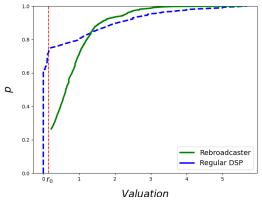
(b) CDFs Hour 13



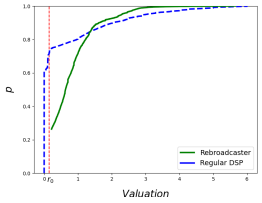
(c) CDFs Hour 14



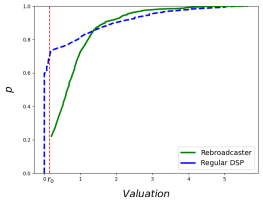
(d) CDFs Hour 15



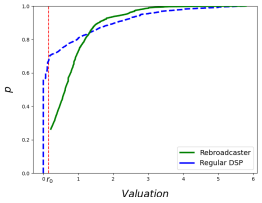
(e) CDFs Hour 16



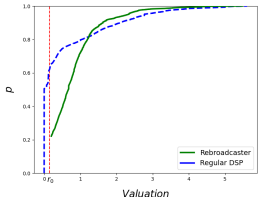
(f) CDFs Hour 17



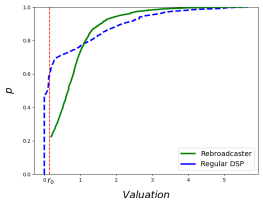
(g) CDFs Hour 18



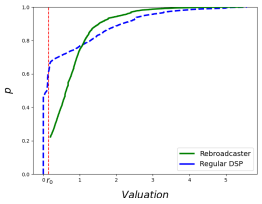
(h) CDFs Hour 19



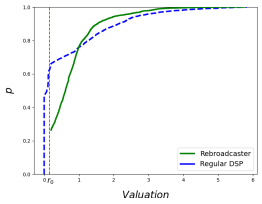
(i) CDFs Hour 20



(j) CDFs Hour 21

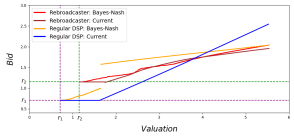


(k) CDFs Hour 22

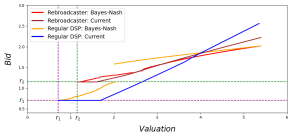


(l) CDFs Hour 23

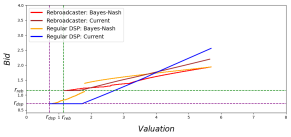
Figure A.6: Bidding Strategies per Hour 0-11: Regular DSPs and Rebroadcasters



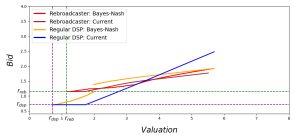
(a) Hour 0



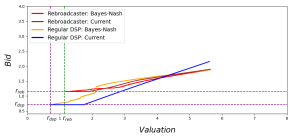
(b) Hour 1



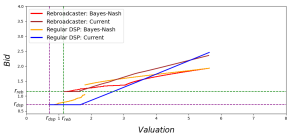
(c) Hour 2



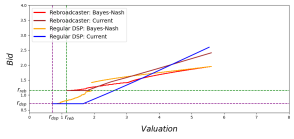
(d) Hour 3



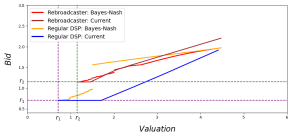
(e) Hour 4



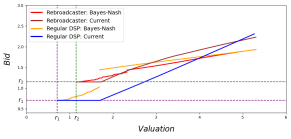
(f) Hour 5



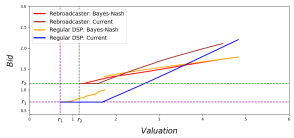
(g) Hour 6



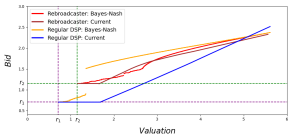
(h) Hour 7



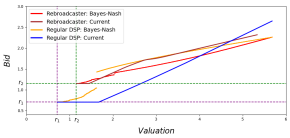
(i) Hour 8



(j) Hour 9

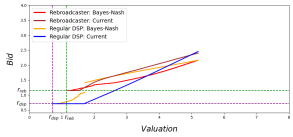


(k) Hour 10

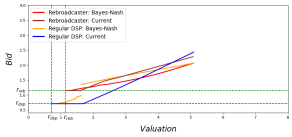


(l) Hour 11

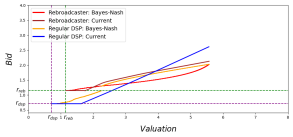
Figure A.7: Bidding Strategies per Hour 12-23: Regular DSPs and Rebroadcasters



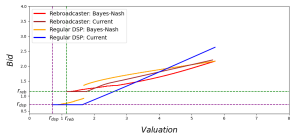
(a) Hour 12



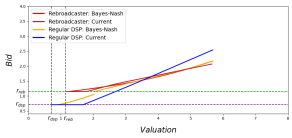
(b) Hour 13



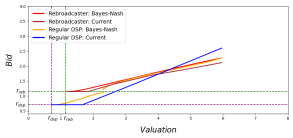
(c) Hour 14



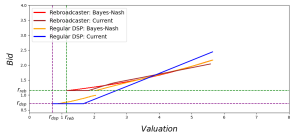
(d) Hour 15



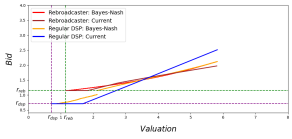
(e) Hour 16



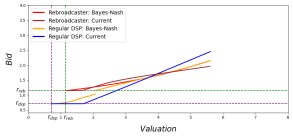
(f) Hour 17



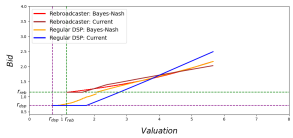
(g) Hour 18



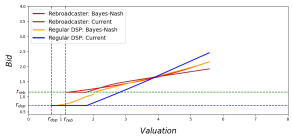
(h) Hour 19



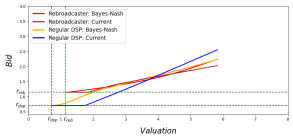
(i) Hour 20



(j) Hour 21



(k) Hour 22



(l) Hour 23

Table A.1: Bayes-Nash Equilibrium Bidding Strategies: Inequality in Equation 3

Hour	$\int_{r_1}^{\bar{v}} F_2[r_2]^{N_2} F_1[z]^{N_1-1} dz$	$\int_{r_2}^{\bar{v}} F_2[z]^{N_2-1} dz$	Inequality Hold
0	2.04	4.24	Yes
1	1.88	4.11	Yes
2	1.72	4.11	Yes
3	2.05	4.29	Yes
4	2.29	4.26	Yes
5	1.43	4.12	Yes
6	1.66	4.14	Yes
7	1.18	4.17	Yes
8	1.79	4.37	Yes
9	1.21	3.32	Yes
10	1.23	3.98	Yes
11	1.13	4.11	Yes
12	1.01	3.66	Yes
13	1.22	3.52	Yes
14	1.73	4.11	Yes
15	1.83	4.29	Yes
16	2.13	4.30	Yes
17	2.24	4.54	Yes
18	1.97	4.15	Yes
19	2.16	4.39	Yes
20	1.79	4.12	Yes
21	1.80	4.19	Yes
22	1.87	4.12	Yes
23	2.26	4.42	Yes

Notes: The table indicates if the inequality in Equation 3 holds for each hour of the day.