

# Adjustment of Bidding Strategies After a Switch to First-Price Rules

Miguel Alcobendas  
Yahoo Research

Robert Zeithammer  
UCLA

October 29, 2021

**Abstract:** We document the response of bidders to a switch in auction pricing rules by a platform for selling online advertising impressions. The platform switched from a second-price auction to a first-price auction, so the same bidder bidding to show the same creative in the same location on the same webpage should bid less after the switch than before the switch. We show that bids indeed decline after the switch, but they do not decline enough given the actual competition each bidder is facing after the switch. To measure whether bids declined enough, we propose a nonparametric estimator of a lower bound on the bidder's valuation underlying each post-switch bid. Bids did not decline enough in that the estimated bounds substantially exceed the pre-switch valuations of showing the same creative. We find evidence of an incomplete and slow downward adjustment in bid magnitude over the period of months, whereby bids remain insufficiently shaded for about half the creatives we analyze even three months after the switch. Implications for analysis of bidding in first-price auctions and analysis of short-run A/B tests of different pricing rules are discussed.

**Contact:** Robert Zeithammer, [rzeitham@ucla.edu](mailto:rzeitham@ucla.edu)

## Introduction

Most platforms for real-time bidding on online display advertising impressions have recently shifted from a second-price to a first-price auction format (Despotakis et al 2019, Choi et al 2020). Compared to life under the old second-price rules, bidding in first-price auctions places a higher computational and informational burden on bidders in that they need to take the intensity of competition into account when shading their bids below their valuations. Most of the transitions happened abruptly, meaning that the entire platform switched to the new pricing rule at one time. Since the timing of the rule-switch is unrelated to the underlying advertiser's valuation of displaying a particular creative, the advertiser valuations of comparable impressions of the same creative do not change around the time of the switch. Because the timing of the rule-switch is also unrelated to customer visits to the websites that display the ads, the quality distribution of available impressions should also not change around the time of the switch. Therefore, bids of the same bidders on the same creative should fall relative to pre-switch levels. We analyze one transition from early 2019, and ask three questions: Did the bids by the same bidder bidding to show the same creative in the same location on the same webpage fall? If yes, did the bids fall enough? And if the bids fell enough, how long did the transition take?

We find that bids indeed declined after the switch for most of the long-running creatives we study, with a median decline of about 5 percent. But was this decline sufficient from the perspective of the individual bidders at the time? To answer this question about the magnitude of observed bid-shading, we propose a simple nonparametric estimator of the lower bound on advertiser valuations in first-price sealed-bid auctions. The estimator assumes that bidders have no detailed private information about the competition they are facing at any given moment, so we can interpret all bids of about the same magnitude at about the same time by the same bidder on the

same creative as facing the same competition (chance of winning). Following Guerre, Perigne, and Vuong (2000), we estimate this competition locally and non-parametrically, but we do not then construct a point estimate of the valuation as is customary in the econometrics literature about first-price auctions (e.g. Athey and Haile 2002). Instead of leveraging a global optimality assumption to get a point estimate of the valuation behind every observed bid, we only assume that the bidder prefers his actual shading of the bid to an alternative greater shading, and show that this weaker assumption implies a lower bound on a rationalizable valuation of the bidder a first-price auction. While weaker than the global optimality assumption, our assumption about bidder rationality is stronger than that in Haile and Tamer (2003), who analyze an English auction and only assume that “bidders neither bid more than their valuations nor let an opponent win at a price they would be willing to beat”. Because we study a sealed-bid auction, the second assumption of Haile and Tamer (2003) is not useful in our context, and we need our bidders to have probability beliefs about their chances of winning at different bid-levels to get a lower bound on valuation. Before the switch, the truth-revealing property of the second price auction allows us to directly equate valuations with bids. Comparing the post-switch valuation lower bounds to the pre-switch valuation magnitudes allows us to conservatively detect insufficient adjustment whenever the former exceeds the latter, on average. And we indeed find that the bid shading was insufficient for a vast majority of bidders and creatives: observed bids on the median creative imply that the bidder bid *as if* the switch from second-price to first-price rules increased his valuation of an impression at least 30 percent.

Our proposed measure provides information about bidder’s bidding strategy above and beyond the prices the bidder pays for impressions measure by CPM (cost per thousand). While CPM increased for all but one creative, such an increase can be attributed not only to insufficient

shading by the focal bidder, but also to an intensification of competition after the switch in the form of additional entry or insufficient shading by competing bidders. Our bounds estimator is designed to disentangle these two alternative explanations of the higher CPM by controlling for the actual competition within each post-switch auction.

Not only do we find insufficient shading right after the switch occurred, we are also unable to detect any downward trend in the lower bound on valuations expressed as a percentage of pre-switch valuation. In summary, we conclude that the bidders we study took more than three months to adjust to the new pricing rule if it ever adjusted at all. Every one of the three multi-creative bidders in our data has at least one creative with consistently insufficient shading throughout the data period, so we do not find evidence of heterogeneity in bidding sophistication across bidders.

We contribute to the growing literature on real-time bidding (RTB) on display advertising (see Choi et al 2020 for a recent review), which is the dominant form of digital advertising today, having surpassed search advertising in dollar terms in 2016. Most papers in the literature keep the market rules fixed and analyze the RTB market from the auctioneer’s perspective focusing on, for example, the problem of setting reserve prices (Choi and Mela 2018), the decision whether to use soft floors (Zeithammer 2019), or the strategy for incorporating dynamic ad sequencing (Rafieian 2019). In contrast to most of the literature, we focus on analyzing the bidders in this market (demand side platforms), and ask how good their bidding strategies seem to be. A closely related paper that also analyzes the switch from second-price to first-price rules is Despotakis, Ravi and Sayedi (2021). Unlike Despotakis, Ravi and Sayedi (2021), we do not attempt to model why the switch occurred, but merely measure the adjustment of bidding strategies to it.

The algorithmic bidders we study do not shade their first-price auction bids enough (as revealed though their implied valuations abruptly increasing after the switch to first-price rules),

just like their human counterparts in laboratory experiments following the seminal study by Cox, Smith and Walker (1988). The experimental literature documents that human bidders bid substantially more than the canonical risk-neutral model predicts, and bidder risk-aversion is the most common explanation of the excessive bidding (Bajari and Hortacısu 2005). We can only speculate about the underlying causes of the insufficient shading we document, but propose that it is unlikely that our bidders are risk-averse, i.e. experience diminishing marginal utility of monetary surplus. Figuring out whether the underlying reason for insufficient shading is one of the key research questions our results bring to light.

Regarding the implications of our results for the broader market for digital display advertising, it is important to note that our sample selection to long-running creatives with lots of bids throughout the period makes our conclusions limited to the bidders and creatives we study, and should not be taken as a general characterization of the entire market. Instead, we view our main contributions to be the development of our simple nonparametric estimator of the lower bound on advertiser valuations in first-price sealed-bid auctions, and an exposition of the difficulty of bidding-strategy adjustment in the early days of selling online advertising by first-price auctions.

### **Data: a price-rule switch on one platform in 2019**

The platform we study implemented the switch to first-price rules in the early part of 2019. We cannot disclose the identity of the platform or the exact date, so we measure time relative to the week in which the switch occurred. Our data consists of all bids on the platform for the most popular ad size and location on the same publisher's website for four months from one month before the switch to three months after the switch. Our goal is to analyze how the same bidder bids

on showing the same creative in the same location on the same webpage before and after the switch. In these auctions, the advertisers do not bid directly, but rather via "Demand Side Platforms" (DSPs for short). DSPs pick and choose among the millions of ad opportunities available on the internet on behalf of their advertiser clients, and formulate bids. Therefore, we focus on the top DSPs (hereafter also called "bidders") in terms of bid volume<sup>1</sup>, and select their creatives that received enough serious bids throughout the time period. Specifically, we look for creatives that received at least one thousand bids every week, with every week's median bid above the reserve price. Fortunately for our goal of analyzing the effect of the switch while holding as many things as possible constant, the reserve price on the exchange we study stayed around the same level throughout the observation period and did not respond to the switch in any discernible way. While idea of our purpose of studying the adjustment of bidding strategies over time, our selection of long-running creatives is clearly not representative of the average creative campaign which tends to be more short-lived.

The data selection procedure described above yields 12.3 million bids on 10 creatives by 4 different bidders – all major players in the industry. We cannot reveal the identities of the advertisers, only that the creatives we study advertise a range of consumer products and services, including consumer durables, home-goods retailers, and investment services. To protect the identity of the bidders, we label them with letters of the alphabet in no particular order. Each bidder's creative is then assigned a number in no particular order. Thus, our unit of observation is a particular bidder (e.g. C) bidding on a particular creative (e.g. 2), labeled "C2". Table 1 summarizes the bidding data by month along with the average cost per thousand impressions (CPM) defined as the amount of money paid divided by the number of thousands of auctions won.

---

<sup>1</sup> One of the bidders is associated with the platform at the corporate level. We exclude this bidder from our analysis because their incentives may be different from the simple surplus maximization we assume in our model.

Our sample is ideal for measuring the adjustment of bidding strategies over time, but it is in no way representative of the typical creative in the market – most last less time, and there are also many smaller bidders who bid too sporadically to be useful for our analysis. Our sample selection thus makes our conclusions limited to the bidders and creatives we study, and should not be taken as a general characterization of the entire market.

**Table 1: Data summary**

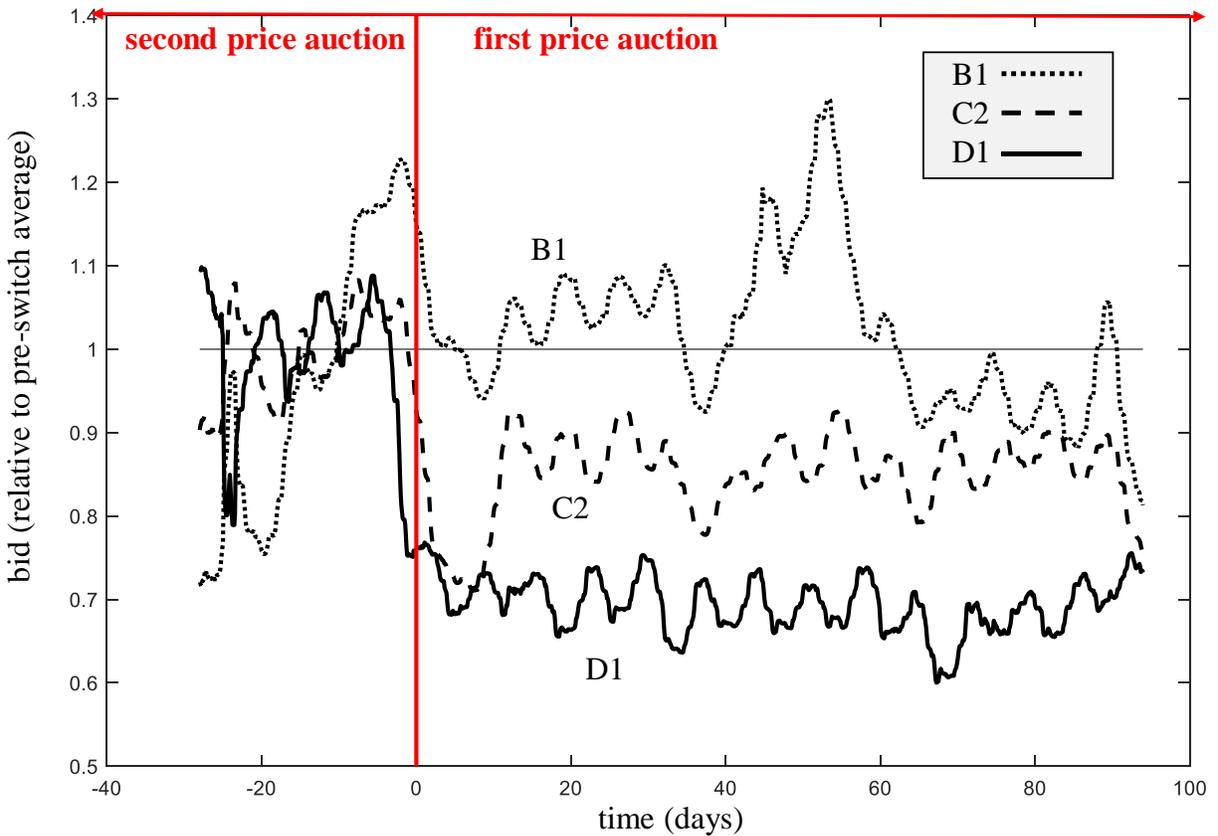
Creative	Number of bids	median bid by month (% of month before switch)			average CPM by month (% of month before switch)			Line type in Figure 1
		1 after	2 after	3 after	1 after	2 after	3 after	
A1	968,238	101%	97%	112%	157%	149%	164%	-
B1	778,412	107%	110%	99%	120%	124%	105%	dotted
B2	505,714	104%	89%	74%	124%	119%	101%	-
B3	123,187	91%	89%	94%	205%	203%	205%	-
C1	4,253,188	88%	85%	93%	139%	136%	137%	-
C2	3,254,961	81%	81%	83%	139%	137%	137%	dashed
D1	214,450	72%	70%	70%	94%	95%	92%	solid
D2	159,370	135%	110%	109%	165%	137%	130%	-
D3	1,612,881	55%	81%	94%	86%	126%	142%	-
D4	487,630	99%	132%	132%	114%	142%	151%	-
Average	1,235,803	93%	95%	96%	134%	137%	136%	

To motivate the development of our bound estimator, we now briefly interpret the data summary in Table 1 starting with bid amounts. The average decrease of 5 percent conceals a lot of variation: a month after the switch, bid decreased substantially (about 20 percent) for seven (about two thirds) of the creatives we analyze, and increased for the remaining creatives mostly to a lesser amount. We conclude that, consistent with theory, the bidders we study did generally shade their bids down in response to the switch in pricing rule. However, we can also rule out a pure bidder heterogeneity story as a potential explanation of differences in adjustment to the price-rule switch: the creatives of bidder D include both extremes of the post-switch difference (highest increase for D2, highest decrease for D1). Analogously, bidder B reduced bids on two creatives (B2 and B3)

while increasing them on a third one (B1). We were surprised not to find more systematic differences at the bidder level.

Figure 1 plots the 4-hour moving average of bids on three different creatives by three different bidders (see last column of Table 1), and it shows that C2 and D1 definitely adjusted to a fairly constant lower level after the switch while B1 did not do so. It is worth noting that the weekly fluctuation in bid amount seems synchronized between B1 and C2, but runs opposite for D1. In other words, B1 and C2 bid more on weekends while D1 bids less.

**Figure 1: Bids on three selected creatives**



While it is clear from Table 1 and Figure 1 that several creatives shaded bids in the same direction as theory would predict, a question remains whether they shaded their bids enough given their valuations of the impressions and the competition they faced (we will define sufficient bid

shading more formally in the next section). The cost per thousand impressions (CPM) data in Table 1 suggests that the adjustment was mostly insufficient: the effective price of an impression increased for all but one creative (D1). One possible explanation for the increase in CPM is insufficient shading by the bidders. Alternatively, competition may have intensified due to additional entry or insufficient shading by competing bidders. Our bounds estimator is designed to disentangle these two alternative explanations of the higher CPM by calculating a lower bound on valuations the post-switch bids imply, and comparing this bound to the pre-switch levels of valuations (revealed directly as bids thanks to the second-price rules) of the same creative by the same bidder.

### **Theory: Deriving a lower bound on the bidder’s valuation from an observed bid in a first-price sealed-bid auction**

Consider a single first-price sealed-bid auction (FPSB) with a reserve price of  $R$ , and let  $b$  be a bid submitted in the auction by one participating bidder. The bidder believes his probability of winning the auction is captured by a cumulative distribution function  $G(b)$  on  $[R, \infty)$ , and values winning the auction some amount  $x$ . Note that if this auction is actually one of many auctions selling similar things, the valuation  $x$  reflects both the option value of other opportunities to win the same object and the number of identical objects the bidder wants to win.<sup>2</sup>

The bidder solves the optimization problem  $b^* = \arg \max_b G(b)(x - b)$ . Note that we are assuming the bidder plays his best response to  $G$ , but we are not assuming any equilibrium. This

---

<sup>2</sup> We thus abstract from dynamic auction issues, assuming the bidder has solved the relevant dynamic program to compute his net valuation of a single opportunity when facing a stream of opportunities. See the appendix for a formalized example of this argument and additional references. We also assume this net valuation is stable throughout the time period.

is a strength of our approach since we are analyzing a market in transition. Obviously, the optimal bid involves some amount of “shading” in the sense that  $b^* < x$ . The fact that the bidder does not shade further by some additional amount  $s$  yields a lower bound on  $x$  as a percentage of  $b^*$ :

$$G(b^*)(x - b^*) > G(b^* - s)(x - b^* + s) \Leftrightarrow \frac{x}{b^*} > 1 + \left(\frac{s}{b^*}\right) \left(\frac{G(b^* - s)}{G(b^*) - G(b^* - s)}\right) \quad (1)$$

We now apply the above inequality to derive a lower bound on valuation  $x$  from the observed bid. Let  $b$  be the observed bid facing  $G$ . Since the bidder does not choose any counterfactual bid on  $[R, b)$ , the bidder’s valuation  $x$  must satisfy:

$$x > LBV(b) \equiv b + b \max_{s \in [0, b-R]} \left(\frac{s}{b}\right) \left(\frac{G(b-s)}{G(b) - G(b-s)}\right) \quad (2)$$

Nekipelov, Syrgkanis, and Tardos (2015) show that under mild assumptions, an analogous bound can apply even when the bidder is merely using no-regret learning as opposed to solving the optimization problem we posit.

Consider a weaker assumption inspired by the “zero-intelligence” bidder idea of Gode and Sunder (1993). Suppose the bidder prefers the observed bid to a random lower bid instead of preferring it to every lower bid pointwise. Then, the analogue to equation 1 becomes:

$$G(b^*)(x - b^*) > E_{s \in [0, b-R]} \left[ G(b^* - s)(x - b^* + s) \right] \Leftrightarrow x > b^* + \frac{E[sG(b^* - s)]}{G(b^*) - EG(b^* - s)} \quad (3)$$

and the analogue of equation 2 becomes:

$$x > LBV_0(b) \equiv b + \frac{E_{s \in [0, b-R]} [sG(b-s)]}{G(b) - E_{s \in [0, b-R]} G(b-s)} \quad (4)$$

We now turn to our empirical strategy for using equations 2 and 4 with our data.

## Empirical identification of winning probability under affiliated private values

Consider the problem of inference about the valuation  $x$  of one focal bidder  $i$  participating in a set of interchangeable auctions. Suppose the analyst observes  $K$  auctions indexed by  $k=1,2,\dots, K$  in which the focal bidder submitted a bid. Specifically, the analyst observes the focal bidder's bid  $b_k$  and the highest competing bid  $Y_k$ . All the auctions are identical in terms of the observables, but differ in terms of some quality unobservable to the analyst yet privately observable to the bidders in the sense that their private values are positively correlated (APV). For example, the focal bidder is a particular DSP (Demand Side Platform, see Data section for institutional details) bidding to show one particular creative, the  $K$  auctions all sell a single impression on the same property at approximately the same time, and the bidders have private information from cookies about the individuals making the different impressions. Such information is naturally positively correlated, e.g. most retailers would like to reach more affluent consumers, *ceteris paribus*. The private information about impression quality gives rise to the different private value  $x_{k,i}$  the focal bidder  $i$  has of each impression  $k$  of the same creative.

The analyst would like to compute  $LBV(b_{k,i})$  using equation 2, but he does not know  $G$ , which clearly depends on  $k$  because of the unobserved (to the analyst) quality of the impression, as well as on  $i$  in the sense of bidder perspective (we are not assuming symmetry or equilibrium). In the absence of the unobserved auction heterogeneity in quality (which causes the affiliation of the underlying private values) or bidder asymmetries, the analyst could simply compute the empirical cumulative distribution function of  $Y_k$  across all auctions in which the focal bidder tried to show the focal creative, and use it in place of  $G$  as in Guerre, Perigne, and Vuong (2000). In such a world, a high value of  $b_{k,i}$  could be attributed solely to a high draw of  $x_{k,i}$  instead of to a combination of high  $x_{k,i}$  and stiffer competition arising from higher  $Y_k$  distributed according to

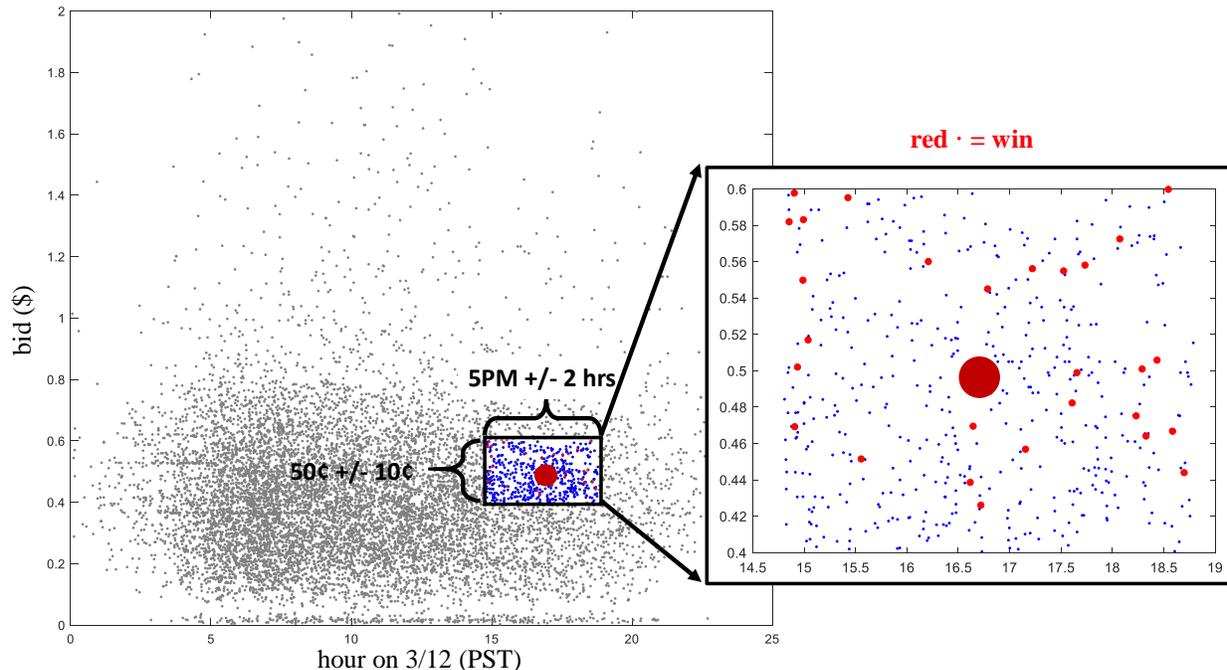
some  $G_{k,i}$ . We propose to account for bidder asymmetry and value affiliation by estimating  $G_{k,i}$  separately from the perspective of each bidder  $i$ , and also estimating it locally both in time and in bid magnitude: Specifically, for each  $k$ , we estimate  $G_{k,i}$  as the empirical distribution of  $Y_j$  faced by bidder  $i$  in the subset of auctions in which bidder  $i$  wanted to show the same creative as in the focal auction, and for which  $|b_{k,i} - b_{j,i}| < \varepsilon$  for some small  $\varepsilon$  and the time of auction  $j$  is close to the time of the focal auction  $k$ .

Figure 2 illustrates our approach on a particular focal auction held at 5 PM on 3/12/2019, in which the bidder offered 50 cents (all bids are on the CPM basis, i.e. per thousand impressions). That day, that bidder made 10,757 bids to show the same creative as in the focal auction, shown as grey dots in Figure 2. To estimate the competition faced in the focal auction, we focus on only the 476 auctions “near” the focal auction in the sense of occurring within two hours and in the sense of the focal bidder bidding within 10 cents of the focal amount. Figure 2 zooms in on these 476 “nearby” auctions, and shows in red the 28 bids that exceeded their respective highest competing bid  $Y_j$ , and thus resulted in a win. The probability of winning the focal auction when bidding 50 cents was thus about 6 percent. The key idea behind our estimator is that we can analyze the cumulative distribution function of the 476 bids *competing* with the nearby auctions, and derive the probability of winning at a lower counterfactual amount. For example, the probability of winning when bidding 45 cents would have been only 4 percent, and not bidding that low thus implies a lower bound of 60 cents on valuation  $x$  as shown in equation 1:

$$0.06 \underset{G(0.50)}{(x-0.50)} > 0.04 \underset{G(0.45)}{(x-0.45)} \Leftrightarrow x > 0.60 . \text{ Equation 2 defines the } LBV(b_{k,i}) \text{ in this auction as}$$

the highest such lower bound among all counterfactual bids between 30 cents (the reserve price) and 50 cents (the observed bid amount).

**Figure 2: Illustration of how we define local competition**



This approach is most similar to Li, Perrigne and Vuong (2002), who studied the equilibrium bids in a symmetric APV FPSB. In contrast, we study only the best response of one bidder, so we have to re-compute the  $G_{k,i}$  separately from each bidder's perspective for each auction. Thanks to this local focus, we can accommodate asymmetry in valuation distributions and we do not require any equilibrium to hold – a better assumption for a market in transition.

We now carefully describe the assumptions that make our approach work. Formally, decompose each bidder's valuation into quality of the impression commonly known among bidders but unobserved by the analyst  $q_k$ , and the private-information residual  $v_{k,i}$ :  $x_{k,i} = q_k + v_{k,i}$ . Our assumption is that our focal bidder  $i$  cannot perform this decomposition, and instead constructs beliefs about competition based on the level of his own overall valuation  $x_{k,i}$ . Specifically, we assume that if  $x_{k,i} = x_{j,i} = z$ , then bidder  $i$  believes  $G_{k,i}(b) = G_{j,i}(b) \equiv G(b; z)$  for every  $b$ , and so

the bid level indexes competition beliefs. Such a bidder then naturally submits the same bid  $b_{k,i} = b_{j,i} = \arg \max_b G(b; z)(z - b)$  in the two auctions. For each observed bid  $b_{k,i}$ , the subset of auctions  $j$  for which  $|b_{k,i} - b_{j,i}| < \varepsilon$  for some small  $\varepsilon$  thus isolates the auctions in which the bidder held approximately the same belief. Some sort of continuity of beliefs in valuations is clearly required here, but we do not think additional definitional formalism is necessary to get our idea across.

It is important to highlight the main limitation implied by our key assumption that that beliefs about competition depend only on the bidder's own overall valuation (in addition to time and other observables). If bidders receive an auction-specific signal about the intensity of competition and this signal is not perfectly correlated with their overall valuation, our approach may not work well. Since both stiffer competition and a higher valuation drive bids up, two identical bids by bidder  $i$  in auctions  $j$  and  $k$  would then not necessarily correspond to the same beliefs about competition, and our identification of the relevant competition from “nearby” auctions would not work.

### **Estimation of Lower Bound on Valuation (LBV): practical considerations**

Having discussed the assumptions under which our approach empirically identifies the competition relevant to a given bid, we now turn to practical issues that arise in an application of our method. There are two issues: first, our *LBV* computations cannot be applied to bids below reserve because such bids already have the lowest (zero) chance of winning. We cannot somehow ignore such bids since they occur both before and after the switch, and so we conservatively let  $LBV(b) = LBV_0(b) = b$  whenever  $b < reserve$ .

The second practical issue is the tradeoff is between “localness” of the competition and precision of the estimate of  $G$  because there aren’t infinitely many auctions for the same creative occurring exactly at the same time as every focal auction. We solve this tradeoff by requiring at least 100 observations of relevant “nearby” competition for a given focal auction before we compute  $LBV$  for the focal bid using equations (2) and (4), and otherwise we again conservatively let  $LBV(b) = LBV_0(b) = b$ . We also make the definition of “nearby” auctions adaptive to the local density of auctions, as we describe next.

The above example in Figure 2 defined auctions “close to the time of the focal auction  $k$ ” in terms of occurring within two hours. Such a fixed-window definition of proximity identifies thousands of nearby auctions for peak-time auctions of high-volume creatives while not finding enough (at least 100, see previous paragraph) nearby auction for, for example, nighttime auctions of lower-volume creatives. We address this seasonality issue by allowing the time-window that defines temporal proximity to vary across creatives and time as follows: For any given focal auction, a candidate auction by the same bidder on the same creative is considered *nearby in terms of time* when both of the following conditions are satisfied:

- 1) it occurred within 6 hours of the focal auction and
- 2) there occurred fewer than 500 auctions by the same bidder on the same creative between the time of the focal auction and the time of the candidate auction.

This proximity rule automatically expands the time-window during low-volume times while tightening it in busy periods, guaranteeing between 100 and 1000 nearby auctions every time we actually compute the  $LBV$ .

For each creative, Table 2 documents the proportion of observations for which we computed our  $LBV$  measures (in percents of pre-switch average valuation), and gives the average

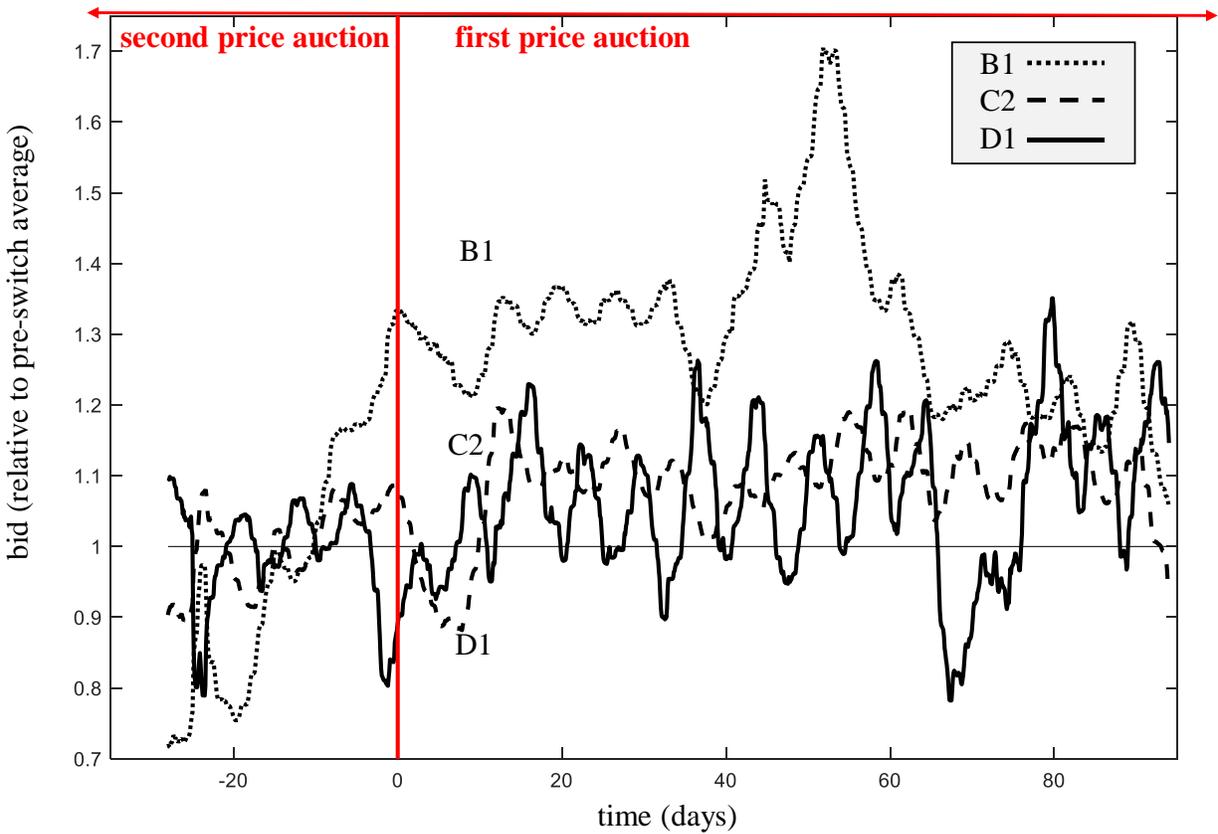
sample sizes used in those computations. While there is obvious variation across the creatives, the table shows that on average, our method computes an LBV for about two thirds of the bids, and bases the estimate of  $G$  on about 325 local auctions involving the same creative that occurred within about 94 minutes (average window with is 3.14 hours = 188 minutes) of the focal auction. For the remaining one third of the observations, our method conservatively assigns an  $LBV=100\%$  (i.e. valuations of those impressions are interpreted to be at least the average valuation of an impression before the switch), so any average LBV (for example, we report a 4-hour moving average LBV below) is also conservative. A median LBV, on the other hand, likely suffers from less bias, and we resort to the median in the full analysis in the next section.

**Table 2: Implementation details of the LBV computation**

Creative	Number of bids	Classification of observed bids according to practical issues in LBV computation			Properties of observations used in computing the $LBV \geq 1$	
		bids above reserve with at least 100 nearby auctions $\rightarrow LBV \geq 1$	bids above reserve with fewer than 100 nearby auctions $\rightarrow LBV = 1$	bids below reserve $\rightarrow LBV = 1$	average number of nearby auctions	average size of "nearby" window (hours)
A1	968,238	59%	10%	30%	276.8	1.29
B1	778,412	67%	9%	23%	265.9	2.65
B2	505,714	57%	12%	31%	298.1	3.55
B3	123,187	55%	45%	1%	241.1	5.17
C1	4,253,188	34%	21%	45%	201.9	0.72
C2	3,254,961	41%	22%	37%	200.8	0.89
D1	214,450	78%	9%	13%	364.2	5.44
D2	159,370	68%	30%	1%	262.6	6.43
D3	1,612,881	63%	15%	22%	286.4	1.39
D4	487,630	98%	1%	1%	857.6	3.90
<b>Average</b>	<b>1,235,803</b>	<b>62%</b>	<b>17%</b>	<b>20%</b>	<b>325.5</b>	<b>3.14</b>

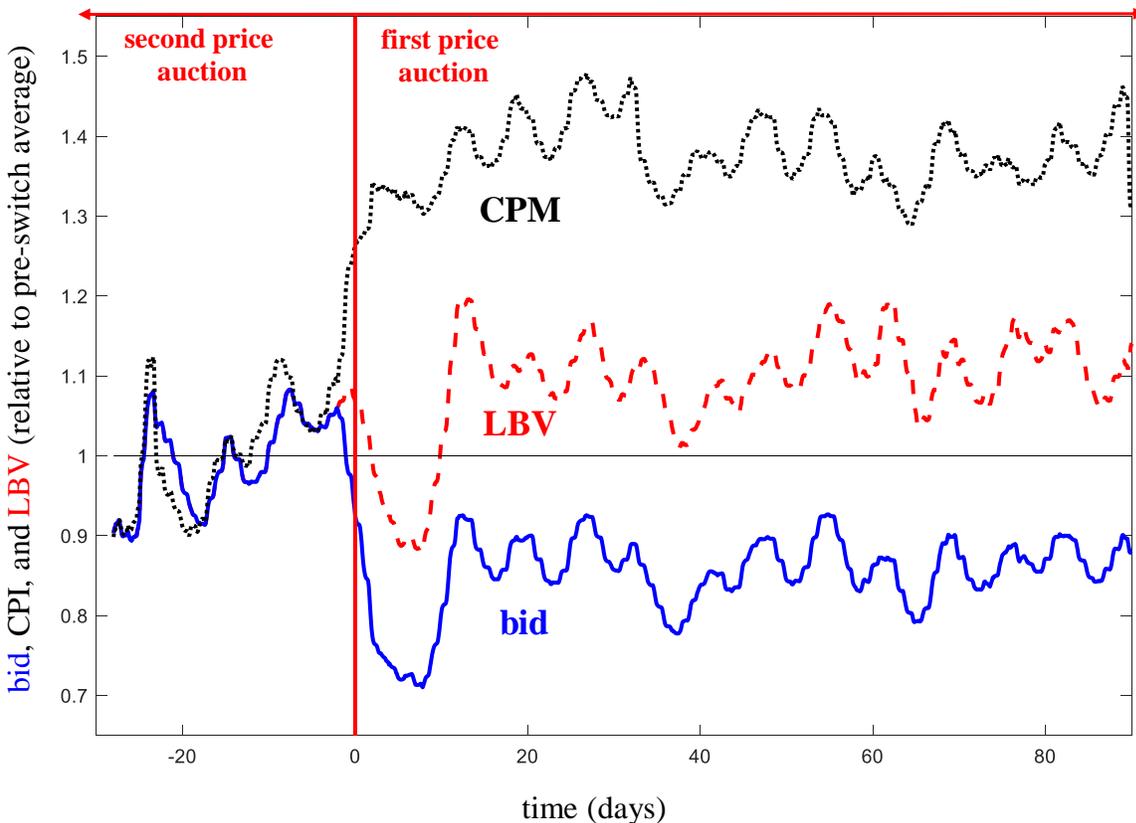
Before summarizing our analysis of all 11 creatives in our data, we return to the three creatives shown in Figure 1, and illustrate the application of our LBV calculation to them in Figure 3. By comparing the LBV to the pre-switch levels of valuations of the same creative by the same bidder, we can detect insufficient shading whenever the post-switch LBV systematically exceeds the pre-switch valuation. From this comparison shown in Figure 3, it is clear that B1 did not shade bids down sufficiently until perhaps the very end of the time period, and C2 did not shade bids sufficiently after about the first two weeks. Since the LBV of D1 oscillates around the same magnitude as the pre-switch valuation levels, we cannot reject the hypothesis that D1 shaded sufficiently.

**Figure 3: Lower bounds on valuations the creatives in Figure 1**



As explained above, the LBV measure isolates insufficient shading given the competition at the time, as opposed to the CPM measure which derives simultaneously from both influences. To illustrate the difference in practice, consider Figure 4, which shows the two measures along with the bid amount for creative C2. Note that in the tumultuous two weeks after the switch, the LBV and CPM measures are telling different stories: the LBV is consistent with adequate shading while the elevated CPM indicates the competition has increased enough to raise prices compared to pre-switch levels. We conclude that our proposed measure provides information about bidder's bidding strategy above and beyond the prices the bidder pays for impressions.

**Figure 4: Cost per impression (CPM) vs. Lower Bound on Valuation for one creative**



Having described our proposed bound estimator and illustrated its use on three creatives, we now turn to the analysis of the entire dataset.

## Summary of empirical results for all creatives

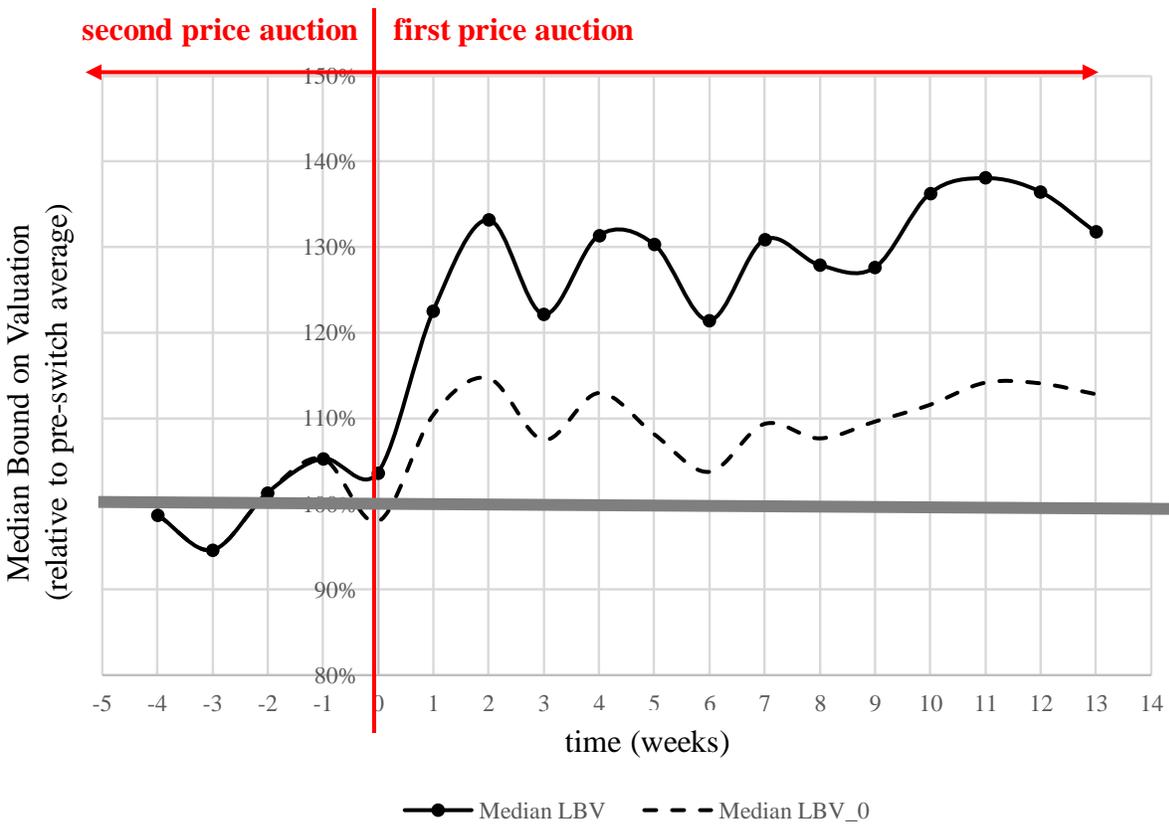
For each creative in our data, Table 3 shows the median weekly LBV expressed relative to the pre-switch median bid of the same creative. Week zero during which the switch occurred is omitted because it involved both rules. The highlighted cells show creative-weeks in which the median LBV exceeds the average pre-switch valuation by at least 10 percent. It is immediately evident that the D1 creative is in a minority, joined only by B3 as an example of a creative, for which our method does not suggest insufficient shading. Looking back at Table 1, B3 is also another example of a creative with CPM (which rises dramatically after the switch) telling a different story from LBV (which does not rise, and hence does not explain the increase in CPM). All other creatives exhibit insufficient shading for the entire three months after the switch. The median LBV remains about 30 percent above the pre-switch average valuation throughout.

**Table 3: Relative LBV by week**

week	A1	B1	B2	B3	C1	C2	D1	D2	D3	D4	median
-4	161%	85%	83%	73%	96%	102%	99%	106%	102%	99%	99%
-3	95%	89%	86%	100%	84%	91%	101%	94%	98%	99%	95%
-2	84%	111%	114%	100%	104%	98%	101%	94%	123%	101%	101%
-1	105%	133%	120%	100%	113%	106%	99%	123%	91%	101%	105%
1	230%	147%	142%	109%	121%	117%	107%	190%	60%	124%	123%
2	177%	151%	141%	100%	143%	126%	101%	149%	74%	117%	133%
3	118%	152%	144%	100%	140%	126%	105%	178%	95%	118%	122%
4	129%	152%	141%	101%	133%	123%	98%	169%	143%	126%	131%
5	153%	145%	135%	142%	77%	116%	109%	142%	108%	126%	130%
6	157%	163%	114%	109%	118%	125%	107%	133%	109%	198%	121%
7	135%	187%	117%	103%	120%	127%	109%	160%	144%	195%	131%
8	137%	153%	109%	103%	129%	127%	107%	115%	173%	203%	128%
9	168%	139%	92%	96%	129%	120%	93%	126%	153%	196%	128%
10	198%	145%	98%	103%	133%	122%	92%	164%	140%	193%	136%
11	144%	139%	89%	103%	143%	131%	115%	166%	137%	193%	138%
12	146%	136%	86%	103%	137%	121%	107%	149%	145%	197%	136%
13	121%	126%	188%	111%	137%	107%	126%	137%	157%	181%	132%

Several of the creatives in Table 3 exhibit a large variance in pre-switch median weekly bids. If we conservatively take the highest pre-switch median weekly bid as the estimate of pre-switch median valuation of the impression, seven of the ten creatives continue to exhibit insufficient shading in that the post-switch median LBV exceeds the pre-switch valuation estimate consistently over a long period.

**Figure 4: Lower Bound on Valuations over time, median across creatives**



In addition to suggesting widespread insufficient bid shading even months after the switch, there is no detectable downward trend in the median LBV over time. Figure 4 plots the median LBV by week shown in the last column of Table 3 along with the analogue for  $LBV_0$  reported in detail in Table 4 in the Appendix. The latter metric indicates that even if we assume that each bidder preferred its bid to a random lower bid (as opposed to a “cherry-picked” lower bid that

maximizes our bound), we still find evidence of insufficient shading throughout the three months after the switch. Unlike in the case of LBV that suggests only D1 shaded sufficiently, the  $LBV_0$  metric suggests that about half of the creatives adjusted bids enough to the new pricing rule. For both metrics, the median lines in Figure 4 represent a wide range of behaviors of the different creatives shown in Table 3.

In summary, we conclude that the bidders we study took more than three months to adjust to the new pricing rule if they ever adjusted at all. Thanks to the insufficient bid-shading, advertisers paid higher prices for impressions during all those months. At least part of the median and average 39 percent increase in CPM evident from Table 1 can thus be attributed to insufficiently shaded bidding strategies of the bidders associated with the creatives, and not merely to increased competition or insufficient shading by other bidders participating in the auctions we study. Every one of the three multi-creative bidders in our data has at least one creative with consistently insufficient shading throughout the data period, so we do not find evidence of heterogeneity in bidding sophistication across bidders.

## **Discussion**

We analyzed the response of bidders (demand-side platforms) to a switch in auction pricing rules from second-price to first-price on one online advertising auction platform in 2019. Bidding in a first-price sealed bid auction is difficult, and we find evidence that the bidders we studied struggled to adjust to the switch on their long-running creatives that enable our analysis. Specifically, the bidders bid on the creatives we analyzed *as if* the switch in pricing rules suddenly increased their valuations of an impression by at least 30 percent for at least three months after the switch. In other

words, the bidders we observe actively bidding to show the same creative in the same space throughout the study period did not shade their bids enough in response to the switch.

We can reach the above substantive “as if” conclusion thanks to a new bound estimator we developed specifically for the purpose of measuring insufficient bid shading. Building on classic work in the econometrics of first-price auctions, our approach derives a lower bound on valuation in each post-switch auction from bid magnitude and the intensity of local competition. By considering local competition not only in the sense of time but also in the sense of bid magnitude, our bound estimator accounts for affiliation in valuations of impressions – a natural assumption in the online advertising context. Armed with an estimate of the lower bound on valuation in each post-switch auction, we then compare the distribution of the bounds with the distribution of pre-switch valuations revealed directly as bids (thanks to second-price rules). When the lower bound on post-switch valuations exceeds the pre-switch valuations of showing the same creative in the same location on the publisher’s website, we conclude that the bidder did not shade the first-price bid sufficiently. We reach this conclusion in 8 of the 10 creatives we study. Every one of the three multi-creative bidders in our data has at least one creative with consistently insufficient shading throughout the three months after the switch covered by our data, so we do not find evidence of heterogeneity in bidding sophistication across bidders.

The insufficient shading has profound implications on the cost of advertising to advertisers. For the median creative, costs per impression rose by about 36 percent after the switch and remained elevated for months. Note that this does not immediately imply that the revenue of the bidding platform rose – an analysis of revenues generated by each of our creatives (not reported here in detail) suggests no systematic pattern, presumably because advertiser budgets varied over time. Moreover, recall that our sample of creatives is selected to address our question of interest,

not to be representative of the market as a whole. However, the CPM increase does mean that our bidders had to pay more to show their creatives after the switch. The bound estimator discussed above allows us to attribute this price increase to insufficient shading by frequently participating bidders as opposed to an increase in competition after the switch or another external factor beyond each bidder's control. It is, however, not clear from our results whether a single rational bidder can somehow recover his pre-switch costs of advertising by an improved bidding strategy. Similar to the argument in Deltas and Engelbrecht-Wiggans (2001), it may be that the seeming irrationality of the competing bidders increases costs even for a bidder who shades bids correctly.

When we look at the lower bound on valuations over time, we do not find a detectable downward trend in the three months after the switch. We conclude that the bidders took more than three months to adjust to the new pricing rule if they ever adjusted at all. The bid-adjustment process, while heterogeneous, was thus surprisingly slow given the technological capabilities the industry.

Since the switch we study occurred early in the calendar year, it is possible that the elevated LBV we find is a result of a seasonal increase in impression valuations over time. However, such an increase would be gradual, not abrupt right after the switch. Nevertheless, as the switch-date recedes further and further into the past, seasonality and other shifts make our assumption of a constant valuation more and more suspect. Our conclusion that the initial strategy adjustment to first-price rules was insufficient thus stands on firmer ground than our conclusion that bid-shading remained insufficient for the entire duration of our data sample.

One way to interpret the lack of adjustment we document is that first-price auctions were still relatively new to the DSPs at the time of the switch we study (early 2019), so coming up with optimal bidding strategies required more effort than it does today. Anecdotal evidence from the

time period suggests that most of the bidders we study did not even actively take the reserve price into account when formulating their bids – something all DSPs do routinely today. Taking the reserve price into account is a clear acknowledgement of first-price rules because it is necessary for bid optimization in first-price auctions, but not in the legacy second-price auctions that have dominant strategies. We have no doubt that bidding strategies have gotten a lot more sophisticated in the few years since our sample, so our paper should be viewed as a study of adjustment to a pricing mechanism with which bidders are not too familiar, not as a critique of today’s industry participants.

The insufficient shading we document means that valuation estimates using standard econometrics of first-price auctions (e.g. Athey and Haile 2002) on current data from real-time bidding (RTB) markets for display advertising may be biased upwards. Unlike in the case of human bidders, where risk-aversion seems to fix this bias (Bajari and Hortacısu 2005), it is not clear why RTB algorithms exhibit the same bias. More research is needed to figure out whether the adjustment we looked for eventually happened, and what was the underlying cause of the insufficient adjustment.

Assuming that the strategy adjustment eventually happened, its slow speed we document implies that analysts of real-time bidding on online display advertising should not rely on short-run A/B tests when evaluating the profitability of different auction pricing rules. Davies (2019) describes a prominent example of such a test when she says “Google has spent the last few months testing the outcome of running first-price auctions across 10% of its Google Ad Manager inventory.” Our results imply that such a test of FPSB vs SPSB is likely to wildly overestimate the long-run profitability of FPSB because it does not allow sufficient time for the surprisingly slow adjustment in bidding strategies to occur. A/B experiments are clearly here to stay as a permanent

part of real-world mechanism design toolkit, but more work on bidder learning and other adjustment is needed to correctly extrapolate short-run tests into long-run predictions.

## **Appendix:**

### **Example of how using a valuation net of future opportunities reduces the dynamic problem to a static one**

Suppose the bidder in fact faces a rapid (no time discounting) infinite sequence of opportunities, and only wants to obtain a single impression (“unit demand” - a simple and tractable example of a budget constraint). Such a bidder solves the following dynamic program for the NPV of having a unit-demand valuation of  $v$ :  $V(v) = \max_b G(b)(v-b) + [1-G(b)]V(v)$ .

It is clear that  $V$  solves:  $0 = \max_b G(b)[v-V(v)-b]$ , so the bidder bids as if he were in a single auction and his valuation were net of the option value of losing:  $x \equiv v - V(v)$ . See Milgrom and Weber (2000) for additional analysis of this issue.

**Table 4: ALBV: Lower Bound on Valuation averaged over counterfactual smaller bids**

week	A1	B1	B2	B3	C1	C2	D1	D2	D3	D4	median
-4	145%	81%	83%	78%	97%	103%	99%	101%	99%	99%	99%
-3	85%	85%	85%	107%	85%	92%	101%	90%	95%	99%	91%
-2	75%	106%	113%	107%	104%	99%	101%	90%	119%	101%	103%
-1	94%	127%	119%	107%	114%	107%	99%	118%	88%	101%	107%
1	158%	119%	120%	117%	111%	96%	90%	149%	55%	111%	114%
2	123%	123%	121%	95%	120%	100%	82%	125%	63%	110%	115%
3	95%	122%	121%	91%	118%	99%	85%	131%	76%	109%	104%
4	99%	123%	119%	98%	115%	99%	80%	130%	109%	113%	111%
5	113%	116%	113%	110%	78%	95%	87%	110%	85%	116%	110%
6	116%	128%	101%	107%	106%	99%	86%	105%	85%	149%	106%
7	104%	143%	102%	107%	107%	99%	85%	109%	109%	155%	107%
8	108%	119%	98%	102%	112%	99%	90%	99%	133%	162%	105%
9	119%	113%	89%	103%	112%	98%	81%	104%	116%	161%	108%
10	135%	116%	93%	110%	114%	98%	82%	116%	107%	157%	112%
11	109%	114%	86%	102%	120%	104%	96%	117%	106%	158%	108%
12	111%	112%	84%	110%	116%	98%	92%	113%	109%	154%	110%
13	97%	103%	144%	119%	115%	91%	99%	115%	115%	150%	115%

## References

- Athey, Susan and Philip Haile (2002) Identification of Standard Auction Models. *Econometrica* 70(6): 2107-2140.
- Bajari, Patrick and Ali Hortaçsu (2005) Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data. *Journal of Political Economy* 113 (4): 703-41.
- Choi, Hana and Carl. F. Mela (2019). Display advertising pricing in exchange markets. Working paper.
- Choi, Hana and Carl F. Mela and Santiago R. Balseiro and Adam Leary (2020) Online Display Advertising Markets: A Literature Review and Future Directions. *Information Systems Research* 31 (2): 556-575.
- Cox, James C., Vernon L. Smith, and James M. Walker (1988) Theory and Individual Behavior of First-Price Auctions. *Journal of Risk and Uncertainty* 1 (1), 61-99.
- Davies, Jessica (2019). “What to know about Google’s implementation of first-price ad auctions”, posted on Digiday.com on September 6, 2019. <https://digiday.com/media/buyers-welcome-auction-standardization-as-google-finally-goes-all-in-on-first-price/>
- Deltas, George & Engelbrecht-Wiggans, Richard (2001) Auctions with an inexpert bidder. *Economics Letters* 73(1): 35-42.
- Despotakis, Stylianos and R Ravi and Amin Sayedi (2021) First-price auctions in online display advertising. *Journal of Marketing Research*. Forthcoming.
- Gode, D. K., & Sunder, S. (1993). Allocative efficiency of markets with zero intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy* 101: 119–137.
- Guerre, Emmanuel and Isabelle Perrigne and Quang Vuong (2000). Optimal Nonparametric Estimation of First-Price Auctions. *Econometrica* 68(3): 525-574.
- Haile, Philip and Elie Tamer (2003) Inference with an Incomplete Model of English Auctions *Journal of Political Economy* 2003 111(1): 1-51.
- Li, Tong, Isabelle Perrigne, and Quang Vuong (2002) Structural Estimation of the Affiliated Private Value Auction Model. *The RAND Journal of Economics* 33(2): 171-93.

Milgrom, Paul R. and Robert Weber (2000). A theory of auctions and competitive bidding, II P.D. Klemperer (Ed.), *The Economic Theory of Auctions*, vol. 1, Edward Edgar, Cambridge, UK.

Nekipelov, Denis and Vasilis Syrgkanis, and Eva Tardos (2015) *Econometrics for Learning Agents*. Proceedings of the Sixteenth ACM Conference on Economics and Computation. June 2015, Pages 1–18.

Rafieian, Omid (2019). Revenue-optimal dynamic auctions for adaptive ad sequencing. Working paper.

Zeithammer, Robert (2019) Soft Floors in Auctions. *Management Science* 65(9): 3949-4450.