Optimizing Reserve Prices in Display Advertising Auctions

Hana Choi* Carl F. Mela†

January 2, 2023

Abstract

This paper considers how a publisher should set reserve prices for real-time bidding (RTB) auctions when selling display advertising impressions through ad exchanges, a $115 billion market and growing. Conducting a series of field experiments to induce exogenous variation in reserve prices at a major publisher, we find an increase in publisher's revenues by 35% and find evidence that advertisers use a minimum impression constraint to ensure advertising reach.

Based on this insight, we construct a structural model of advertiser bidding model that accommodates impression constraints to infer the overall demand for advertising as a function of reserve prices. Using this demand model, we solve the publisher pricing problem. Consideration of minimum impression constraints in setting reserve prices generates a predicted increase in profits of nine percentage points over a solution that does not incorporate the constraint. In a final series of field experiments, we validate our model's predictions by showing ad revenues tend to be highest across exogenously varied reserve price levels closest to the imputed optimal reserve prices.

Keywords: Display Advertising, Pricing, Auctions, Reserve Prices, Optimization, Real-Time Buying (RTB), Programmatic Buying, Real-Time Bidding, Impression Constraints, Field Experiments, Fluid Mean-Field Equilibrium

JEL Classification Codes: D4, L1, L2, M3

---

*Assistant Professor of Marketing at the Simon School of Business, University of Rochester (email: hana.choi@rochester.edu, phone: 585-275-0790)

†The T. Austin Finch Foundation Professor of Business Administration at the Fuqua School of Business, Duke University (email: mela@duke.edu, phone: 919-660-7767).

The authors thank Santiago Balseiro, Bryan Bollinger, Mingyu (Max) Joo, Giuseppe Lopomo, Kenneth Wilbur, Hema Yoganarasimhan, and seminar participants at the 2019 Conference on Digital Experimentation, 2020 UT Dallas Frank M. Bass FORMS Conference, Arizona State University Statistics, Boston College, Boston University, City University of Hong Kong, the University College London, the University of Colorado Boulder, Cornell University, Duke-UNC Brownbag, Facebook Econ / Algorithms Meeting, the University of Georgia, Harvard University, HEC

1 Introduction

1.1 Overview

This paper considers how publishers (e.g., CNN, Wall Street Journal, Facebook, Youtube) should set reserve prices for auctions when selling display advertising impressions in real-time through ad exchanges (e.g., DoubleClick Ad Exchange, OpenX, Rubicon).[1] Display ad markets are estimated to be \$123 billion in the US in 2022, growing considerably from \$106 billion in 2021.[2] Drivers behind this 16% growth rate include an upswing in mobile activities, a proliferation in online video ad formats, and technological advancements in real-time buying (RTB) through ad exchange auctions. Display ad spending have surpassed even search ad spending and is forecasted to continue its rapid revenue ascent.[3] One reason RTB is an ascendant approach to selling display advertising is that it offers advertisers a litany of information regarding demographics, interests, and past behaviors of the person seeing the impression on the publisher's page, thereby affording advertisers considerable latitude when targeting customers. Because each impression is sold individually and in real time, the advertiser has substantial flexibility over both how many exposures to buy (a reach constraint) and when and how much to spend (a budget constraint). In spite of this rapid market growth, empirical research into publisher's strategies in display markets is limited.[4] This paper seeks to fill this void by developing an approach to help publishers to maximize their revenues by optimizing their reserve pricing in display ad markets.

To achieve this aim, we first consider how advertisers value ad impressions, and how advertisers' bidding behaviors can be affected by (i) own valuations, (ii) competing valuations, and (iii) the potential reach or budget constraint they face. Subsequently, we examine how the publisher should set its reserve price in response to advertisers' demand and bidding behaviors. To address these

Paris, HKUST, the University of Illinois at Urbana-Champaign, Johns Hopkins University, Lingnan University, the University of Miami, MSI State of Marketing Science Summit, Northwestern University, the University of Notre Dame, Rice University, the University of Rochester, Southern Methodist University, Stanford University, the University of Texas at Austin, Tilburg University, U.S. Department of Justice, the University of Vienna, and Yonsei University for comments and suggestions.

[1] User characteristics together with ad unit characteristics (e.g., time of ad delivery, site where ad is placed, ad location on a page etc.) define an ad impression - the unit of sale in display ad markets. If a site visitor sees two ads on the site, for example one located above-the-fold and the other located below-the-fold, two *independent* auctions are conducted for these two ad impressions.

[2] https://www.insiderintelligence.com/insights/programmatic-digital-display-ad-spending/

[3] Search advertising revenue totaled \$61.7 billion in 2020. (https://www.emarketer.com/content/search-advertising-resilient-thanks-ecommerce-channel),.

[4] Please see Choi et al. 2020 for operational details of display ad markets.

two questions, we collect a novel dataset on advertiser bidding behaviors from an ad exchange, and run a series of reserve pricing experiments in conjunction with a major publisher. Using this exogenous variation in reserve pricing, we develop a structural model of advertiser bidding behavior in competitive markets, and derive the optimal pricing policy on part of the publisher selling its inventory into the exchange market.

Several novel findings emerge. First, large revenue increases result by using a better reserve price. Using a naive assumption of no constraints to set reserve prices in our experiments increases publisher revenues an average of 35% across 12 studies. Second, the exogenous experimental variation enables us to subsequently test this naive assumption of no advertiser budget or impression constraints. Initial findings reveal that bids vary with reserves in a manner consistent with advertisers' use of minimum impression constraints. It becomes possible to further improve publisher profits by about 9 percentage points when accounting for these constraints. Finally, using an additional series of validation experiments, we find that revenues are generally optimized at the optimal reserve levels set by the model.

1.2 Relevant Research

Much of the display advertising literature focuses on measuring advertising effectiveness and the attendant consequences for advertisers' ad buying and targeting decisions (e.g., Balseiro and Gur 2019, Barajas et al. 2016, Cai et al. 2017, Celis et al. 2014, Ghosh et al. 2009, Goldfarb and Tucker 2011, Hoban and Bucklin 2015, Iyer et al. 2014, Johnson 2013, Johnson, Lewis, and Nubbemeyer 2016, Johnson, Lewis, and Reiley 2016, Lee et al. 2013, Lewis and Rao 2015, Rafieian and Yoganarasimhan 2021, Sahni 2015, Sayedi 2018, Tucker 2014, Tunuguntla and Hoban 2021, Xu et al. 2015, Yuan et al. 2013). In contrast, our focus is on publishers' pricing strategies in exchange markets.

A developing game theoretical literature addresses this concern. Despotakis et al. 2021 solve for the optimal display ad reserves in a single-shot game. Balseiro et al. 2015 consider reserves where advertisers face budget constraints over the duration their campaigns. We build on this literature in two ways; (i) we assume an empirical approach to set optimal reserve prices by developing a structural model to be used in practice and (ii) we consider the possibility of advertiser impression constraints (that is, the advertiser seeks to win at least a specific number of ad impressions over the duration of a campaign).[5,6]

In recovering advertisers' valuations in this dynamic setting, we adopt the fluid mean-field equilibrium (FMFE) framework developed in Balseiro et al. 2015. Of note, the solution to the advertiser bidding game only depends on the steady-state distribution of rival bids (and not on the current single auction-specific state or the rivals' individual states). This equilibrium concept provides a computationally tractable way of modeling advertiser bidding behaviors and competition while capturing the dynamic nature of the advertiser decisions across auctions.[7] We extend this framework to accommodate a minimum impression constraint, and suggest estimation and identification strategies to link the FMFE theory to our empirical context.

Finally, in contrast with most of this prior research, we conduct field experiments to provide exogenous variation to test assumptions used in our model and demonstrate the potential of using theoretical insights to improve auction outcomes. On this dimension, our research is related to Ostrovsky and Schwarz 2011, who conduct a large field experiment in the context of search advertising. They find that setting appropriate reserve prices guided by theory leads to substantial increases in seller revenues. We similarly corroborate the importance of setting reserve prices, but in the context of display ad auctions (instead of sponsored search auctions). Different from Ostrovsky and Schwarz 2011, we additionally examine the causal effect of reserve prices on advertiser bidding behaviors and use a structural model to back out advertiser primitives.

1.3 Organization

This paper is organized as follows. Section 2 first characterizes our data, demonstrating considerable heterogeneity in advertiser bidding behaviors. Next in Section 3, we present the advertiser bidding model and the publisher pricing model. Section 4 provides experimental evidence of constraints.

[5] Institutional practice suggests that advertisers set auction 'win rates', i.e., the number of impressions won / the number of bids submitted (https://www.adtaxi.com/blog-roll/rtb-win-rates-bigger-always-better). When 'win rates' decrease, advertisers increase their bids to ensure a certain number of ad impressions are won for a campaign. To abet this aim, demand side platforms such as Google Ad Manager typically enable advertiser impression tracking against goal (https://support.google.com/admanager/answer/2913775?hl=en). It is possible that such a constraint could reflect advertisers' minimum sales thresholds (perhaps to amortize various marketing, development, or other investments). Though advertisers cannot set sales thresholds directly when bidding for ads, they can opt to set impression thresholds which have downstream consequences for sales.

[6] Coey et al. 2021 develops a scalable approach to compute optimal reserve prices without the need to recover bidder valuations. Unlike our display advertising setting, bidders are assumed to have one opportunity to bid and play the weakly dominant strategy of bidding their valuations.

[7] Backus and Lewis 2016 and Hendricks and Sorensen 2015 study bidders having a unit demand with multiple opportunities to bid in second-price sealed-bid auctions, and apply the framework to eBay's market. They also similarly use the belief formation and the mean-field equilibrium concept in Krusell and Smith 1998, Weintraub et al. 2008, and Iyer et al. 2014, but for a unit demand without a (reach or budget) constraint.

Section 5 discusses the estimation method and the identification argument in inferring advertiser valuations, and Section 6 presents the estimation results. Finally in Section 7, we compute the optimal reserve prices and the revenue gains.

2 Data

In this section, we describe the publisher and its data source to describe the research context and to show that the advertiser bids differ for different types of impressions. Heterogeneity in bids is instrumental in setting reserve prices; it suggests that advertisers differ in their valuations for advertising impressions and that they consider information about impressions such as the type of ad inventory in setting their bids.

2.1 The Publisher and its Data

The data for this study are collected from a large, premium publisher, ranked within the U.S. Top 10 by comScore.[8] The publisher has over 35 brands (sites). We focus our attention on display ads and exclude video in our analyses.[9]

The publisher's ad exchange partner provides a report on auction outcomes related to this publisher's ad impressions and this report is generated at the daily level and includes advertiser ID, Demand Side Platform (DSP herein) ID,[10] day, site where ad was placed, ad type (ad size, ad location on a page, device), number of bids submitted, number of impressions won, bidding amount, payment amount, and click responses. Thus the observational unit (i.e., dimension) is at the advertiser-DSP-day-site-ad type delivered, and the metrics provided for each observational unit are number of bids submitted, number of impressions won, bidding amount, payment amount, and number of clicks received. While number of impressions won, payment amount, and number of click responses are available for all advertisers (we call this the "payment data"), number of bids submitted and the bidding amount are only available for a subset of advertisers who opt-in to share their data with the publisher (we call this the "bidding data").

The data are collected in three waves as outlined in Figure 1. We collect observational data in Wave 1 to characterize heterogeneity in advertiser bidding behaviors and to inform the design

[8] https://www.comscore.com/Insights/Rankings

[9] We adopt this sampling criterion due to the publisher's current selling policy, where video ads are rarely available for sale in the ad exchange.

[10] Demand Side Platforms (e.g., MediaMath, DataXu, Turn, Rocket Fuel, Adobe Advertising Cloud) facilitate the advertisers' bidding process.

Figure 1: Timeline

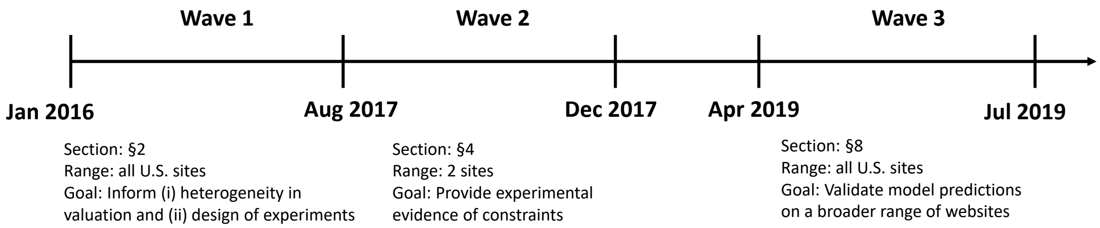

of reserve price field experiments in Wave 2. The reserve pricing field experiments in Wave 2 are used to discriminate between various theories regarding advertiser behavior as described in Section 4. These insights on advertiser behavior enable us to recommend an optimal reserve policy as described in Section 7. Wave 3 then conducts validation experiments to test this reserve price setting approach.

2.2 Summary Statistics

In this section, we provide summary statistics from the observational data (Wave 1) to show there exists considerable variation in advertiser bidding behaviors which can inform the design of reserve price field experiments.

A total of 15,745 advertisers participate in the auctions from January 2016 to August 2017. Summary statistics of advertiser buying behaviors are present in Table 1. Cost-per-mille (CPM) paid and bid CPM are scaled by a multiplicative constant for confidentiality. At the observational unit level (i.e., advertiser-DSP-day-site-ad type), advertisers on average win 14 impressions at \$0.93 CPM rate.[11] The minimum CPM payment is close to zero, because the publisher imposed no reserve prices for most of its inventory during this period.

Table 1: Summary Statistics of Advertiser Behaviors

Per Observation Unit		Mean	Median	Std Dev	Min	Max
All	#Impressions Won	14.32	1	760.74	0	$2.99M$
	CPM paid (\$)	0.93	0.59	1.94	0.00	4646.95
Opt-In	#Impressions Won	6.98	0	720.88	0	$2.29M$
	CPM paid (\$)	1.18	0.71	2.34	0.00	1459.91
	# Bids Submitted	230.15	4	8306.13	1	$14.44M$
	Bid CPM (\$)	1.48	0.66	12.25	0.00	37500.00

Note: Summary CPM paid statistics are calculated conditioned on #impressions won > 0. CPMs are scaled to preserve publisher confidentiality.

Bids are observed from the 81% of the advertisers who opted-in (default setting) to share their

[11] The number of observations is not reported in Table 1 and Table 2 to conceal the daily total number of ad impressions for the publisher firm.

bidding information. Opt-in advertisers pay a higher CPM ($1.18) than the full sample average ($0.93), but buy a much smaller number of impressions, constituting about 20% of the total revenues. Hence, there may be selection concerns arising from using the bid CPMs for inferring advertiser valuations in estimation. Fortunately, data are available on the CPM paid from all advertisers. Accordingly, we rely on the payment data in our estimation and only use the number of bids submitted and the bid CPM data when reporting data patterns (sub-section 2.3) and experimental findings (sub-section 4.2.2).

2.3 Heterogeneity in Valuations

In order to understand how to set reserve prices, it is necessary for the publisher to determine the distribution of advertiser valuations. The results below suggest substantial heterogeneity in bids with respect to the observed characteristics, and thus the potential to use reserve prices to price discriminate.

Using the bid data from January 2016, Figure 2 depicts the distribution of bid CPM z-scores.[12] The bid CPM z-score is calculated as (bid CPM - weighted mean) / (weighted stdev), where the number of bids for a given observational unit (advertiser-DSP-day-site-ad type) is used as the weight.[13] The x-axis represents these bid CPM z-scores, and the y-axis represents the percentage of number of bids. The figure evidences considerable heterogeneity in bidding with the minimum of -0.22 and the maximum value being 19506.13.

Figure 2: Bid CPM Z-Scores: Heterogeneity

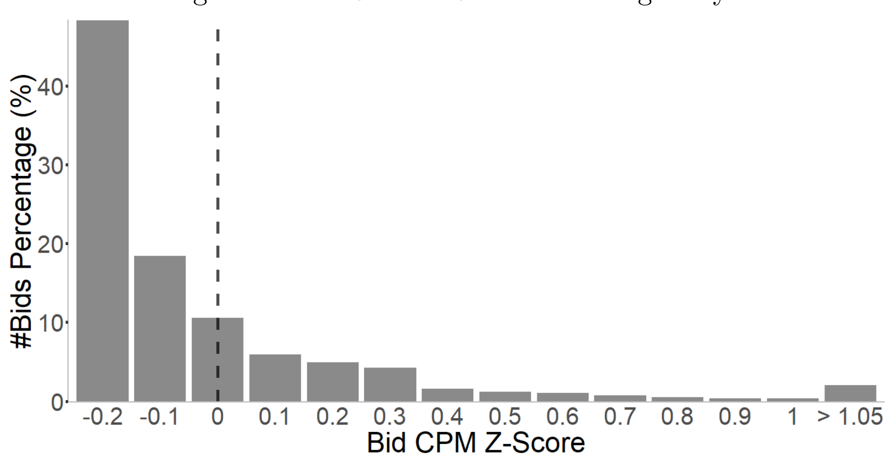

[12]The results from other months are qualitatively similar.

[13]Bid CPM observations are daily averages (i.e., total bidding amount/#bids) for given observational units (advertiser-DSP-day-site-ad type) so we use the number of bids as the weight in constructing the bid CPM z-score.

Table 2: Bid CPM Z-Score: Weighted Least Square Regression

DV: Bid CPM Z-Score	Estimates	SE
Desktop vs. Tablets	**0.11*****	(0.02)
Mobile vs. Tablets	**0.05*****	(0.02)
Web vs. App	0.02	(0.05)
Above the Fold vs. No Info	**0.04*****	(0.01)
Mid vs. No Info	0.01	(0.02)
Below the Fold vs. No Info	0.01	(0.02)
Size	**0.02*****	(0.00)
Day FE	Yes	
Site FE	Yes	
Advertiser FE	Yes	
DSP FE	Yes	
R-squared	0.14	

Note: An observation unit is (advertiser-DSP-day-site-ad type). Standard errors in parentheses are clustered at site level. ***denotes 1% significance.

Figure 3: Advertiser Bid CPM: Advertiser and Site Fixed Effects

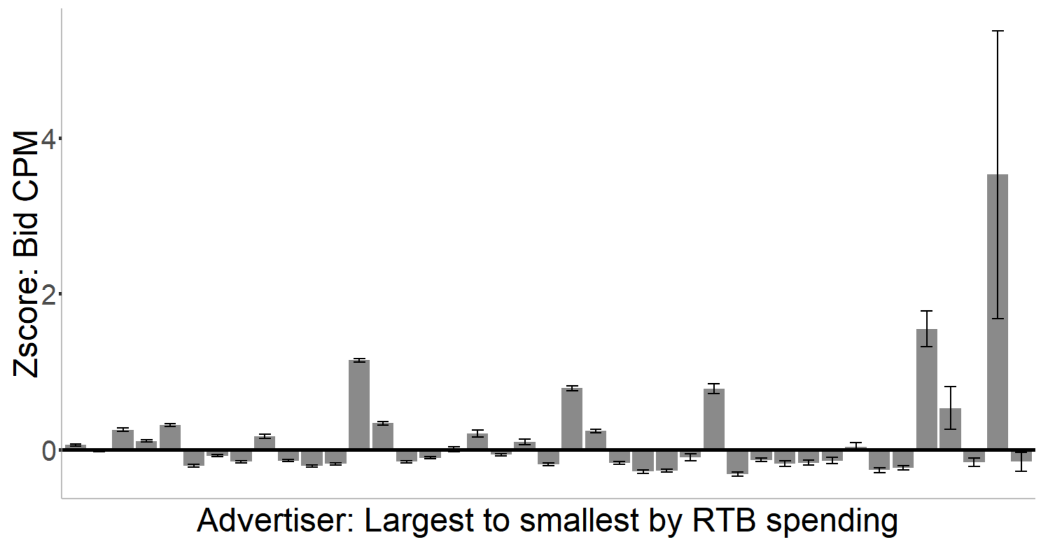

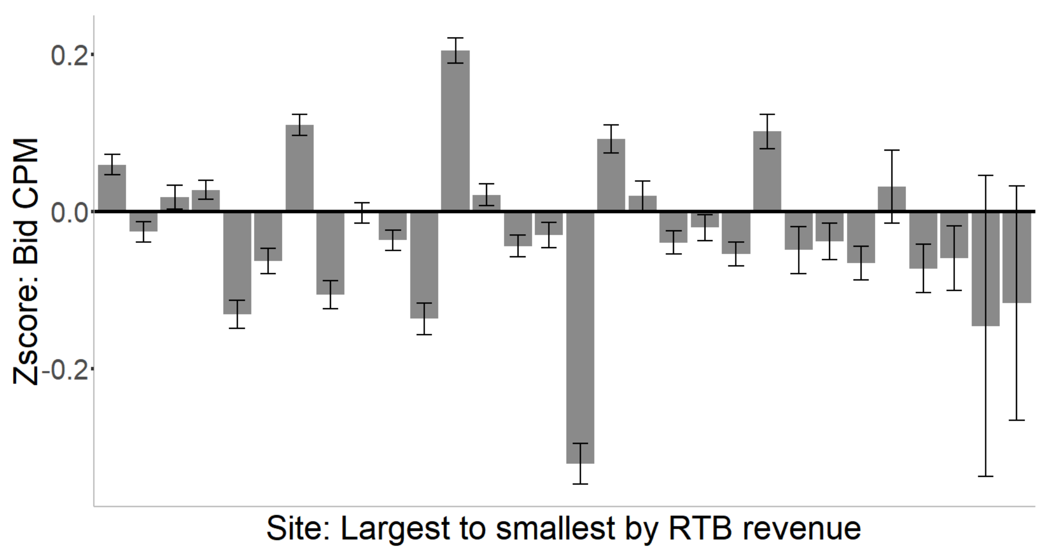

Note: Advertisers and sites are ordered by the spending/revenues in the real-time-buying (RTB) ad auctions.

To explain the variation in the advertisers' bid CPMs, we estimate a weighted least square regression of bid CPM z-scores on the ad type (device, web vs. app, ad location, ad size), day, site, DSP, and advertiser. The weight used in the least square regression is the number of bids submitted for each observational unit.[14] Table 2 indicates that advertisers bid more for the desktop ads (relative to mobile, tablets), above-the-fold ads (relative to mid, below-the-fold, or no info), and bigger size ads. These results imply that optimal reserve prices can be set differentially based on these ad types. Figure 3 plots the advertiser fixed effects and the site fixed effects estimated from the weighted least square regression of bid CPM zscores. Both advertiser and site explain considerable variation in bid CPMs, suggesting they can also be used to price discriminate when setting reserve prices.[15]

3 Model

This section develops a model of advertiser ad buying as a function of advertisers' valuations and reserve prices. Subsequently the publisher's reserve price decision is modeled taking advertisers' ad buying decisions into account. Our approach builds upon the theoretical work of Balseiro et al. 2015, who use a fluid mean-field equilibrium (FMFE) to approximate the dynamic optimal bidding problem of budget-constrained advertisers buying multiple impressions for a campaign. In this section, we extend these results to the context of a minimum impression constraint. Analogous to a budget constraint, a minimum impression constraint can induce inter-temporal, dynamic interactions over multiple second-price auctions that influence how reserve prices affect advertiser ad buying and publisher revenues. As we shall show, this constraint induces advertisers to bid higher than their true valuations (subject to a participation constraint). We conclude this section by contrasting the two types of constraints and how they affect advertiser bidding behaviors in Section 3.5.

3.1 Model Overview

The advertising game proceeds as follows:

1. Publisher: The publisher moves first and selects the reserve price, r, in the second-price auction. The publisher's objective is to maximize its expected revenues from the auctions.

[14] The ad size metric is calculated in (pixel^2/10000) unit. For example, we calculate the ad size for a (300×250) ad to be 7.5 ($= 300 \times 250 \div 10000$).

[15] The results from using bid CPMs ($) in lieu of bid CPM zscores are qualitatively similar. We report bid CPM z-scores instead of levels to preserve the confidentiality of the bid CPM levels.

The publisher knows the distribution from which the advertiser's valuation is drawn, but does not know the advertiser's realized true value for each particular available impression.

2. Advertisers: The advertisers move in the second stage. Each advertiser decides how much to bid b, based on the realized valuation v for each given impression, the reserve price r, and its expectation of the distribution of bids the advertiser will compete against. The advertiser's objective is to maximize its expected utility (= valuation − cost) across auctions over the duration of the advertiser's ad campaign and given its minimum impression level, y, for that ad campaign.

Using backward induction, we first solve the advertiser's ad buying problem, and then characterize the publisher's reserve pricing problem.

3.2 Assumptions

This sub-section outlines the model assumptions used to solve the game.

3.2.1 Auction Rule

The ad exchange conducts a second-price, sealed-bid auction for each available impression. An impression is delivered to the winner who bids the highest above the reserve price (i.e., the winner's ad is displayed to the consumer). If no advertiser bids above the reserve price, the impression remains unsold and is perishable. The winner pays the second highest bid, or the reserve price if the second highest bid falls below the reserve price. Hence, we consider a second-price auction mechanism.

Generally speaking, walled gardens (e.g., Amazon or Facebook where the exchange and publisher are integrated) use first price auctions (FPA) while independent exchanges used second price auctions (SPA) (Despotakis et al. 2021). Generalizing our approach across auction mechanisms is conceptually straightforward. When the estimation model uses a second price auction, assessing the effect of a FPA requires a counterfactual analysis. In contexts where data are generated using FPAs, the estimation model would be adapted to bidder behavior under FPAs, while the effect of an SPA could be conducted using a counterfactual model.

3.2.2 Symmetric Independent Private Values

Following the symmetric independent private value assumption commonly adopted in prior work (e.g., Ostrovsky and Schwarz 2011, Balseiro et al. 2015), advertiser valuations are assumed to be drawn independently and identically from the conditional distribution $F_V(v|z)$, where Z are auction specific observed covariates. After controlling for these observed covariates, advertiser valuations are assumed to be symmetric, independent, and do not depend on other bidders' private information.[16]

3.2.3 Reserve Price

The reserve price is assumed to be known to all potential bidders (advertisers). In practice, the reserve price is not announced to the advertisers, but they can infer it from their repeated auction experience (for example, via experimentation, automated bidding, and machine learning algorithms).[17]

3.2.4 Utility Function

Advertisers are assumed to have a quasi-linear utility function, where utility is defined as the sum of the advertiser's valuations from the impressions won less the total payment for the impressions. Details on the advertiser utility function will be presented in equation 1.

3.2.5 Fluid Mean-Field Equilibrium (FMFE)

We adopt the fluid mean-field equilibrium (FMFE) concept as it employs a couple of advantageous approximations (Balseiro et al. 2015). First, it considers a *mean-field* approximation (Weintraub et al. 2008, Iyer et al. 2014) to relax the informational requirements of agents, requiring only that they know some aggregate and stationary representation of the competitors' bids, as opposed to tracking the specific action of each and every of the potentially 1000's of other bidders. Second, this FMFE equilibrium concept considers a *stochastic fluid* approximation from the revenue management literature that reduces the computational burden in addressing the complex inter-temporal dynamics of the advertisers' bidding problem.

The two approximations in FMFE and their benefits come with several assumptions. First, the mean-field approximation assumes that the distribution of competitors' bids across auctions

[16] This specification allows for values that are correlated via the observed attributes Z. For example, Z may contain site dummies to capture that site1 is valued more than site2 among the advertisers. Z can also contain month dummies to control for seasonality.

[17] https://www.aarki.com/blog/understanding-hard-and-soft-price-floors-in-programmatic-media-buying

is stationary with the number of advertisers in the market being large (Balseiro et al. 2015, Iyer et al. 2014). Because only a small fraction of competitors participate in any given auction and that set varies from auction to auction, the distribution their bids "averages out" and appear stationary over the many auctions an advertiser faces during the duration of its campaign (which often extends for weeks or months). In a large market where any individual player has little effect on the subsequent play of others, a bidding strategy that relies on the aggregate and stationary representation of the competitors' bids approximates well the rational behavior of advertisers. This approximation is well suited for our display advertising context where thousands of advertisers compete over a duration of a campaign and where there is little value in tracking and forecasting the exact actions of all individual bidders. Institutionally, tracking the specific actions of all agents is impractical and implausible because advertisers generally do not have access to complete information about bid history. Rather, advertisers often rely on "bid landscape" information provided by the ad exchanges (Ghosh et al. 2009, Iyer et al. 2014, Balseiro et al. 2015). Bid landscapes are stationary representations of bidding distributions over a recent time horizon, which is exactly the relaxed information requirement in FMFE.

Second, the stochastic fluid approximation assumes that bidders' minimum impression constraints need only be satisfied ex-ante, in expectation when bidders solve for their optimal bidding strategies. Here, the rationale is that impressions arrive with sufficient frequency over a campaign duration that the minimum impression constraints can be readily met in expectation. That is, an advertiser has a large number of bidding opportunities over the campaign length (Gallego and Van Ryzin 1994). With this assumption and the stochastic fluid approximation, one can yield a simple and behaviorally appealing characterization of the optimal bidding function: an advertiser needs only to adjust its own valuation by a *constant scale* (see Proposition 1). This constant factor guarantees that the advertiser's minimum impression constraint is met in expectation at the end of the ad campaign, thus it captures the dynamic nature of the advertiser decisions well without much computational burden.

3.3 Advertiser Bidding Model

The goal of the advertiser bidding model is to determine the optimal bidding policy for advertisers. Subsequently, we will use the advertiser bidding model developed in this section to match the optimality condition to the moments in the data (such as advertisers' payments) in order to infer the distribution of advertiser valuations. Once these valuations are known, it becomes possible to explore the effect of reserve prices on advertiser bidding and the revenue implications when advertisers face practical constraints – the ultimate goal of this paper.

The amount of advertising inventory varies over time. Following Balseiro et al. 2015, we assume that the arrival process of available advertising impressions ($\simeq$ users) at any given point in time follows a Poisson distribution with intensity η. We allow for advertiser k to have a minimum impression level y_k for the campaign length s_k (i.e., the duration over which an advertiser is using its bidding rule). Then, ηs_k (the expected arrivals times the duration of the bidding interval) indicates the total number of impressions arriving during the campaign period.

For a given impression i that arrives in this bidding interval (campaign length), we denote advertiser k's value as v_{ik}, which is assumed to be drawn independently and identically from a continuous cumulative distribution $F_V(\cdot|Z)$. Z are observed characteristics (e.g., ad type). Further, let D be the steady-state maximum of the competitors' bids, where the publisher is also considered as one competitor that submits a bid equal to r. The distribution of D will endogenously be determined in equilibrium, which we denote as F_D.

The advertiser maximizes its expected utility (= valuation − cost) from the ad auctions, given the minimum impression constraint and the participation constraint. With the assumptions made in sub-section 3.2, we focus on the bidding strategy $\beta_\theta^F(v_i|F_D,\theta)$ for an advertiser type $\theta = (s, y)$ (where advertiser type is defined by the campaign length and the minimum impression level) to be a function of the advertiser's own valuation v_i. The advertiser faces the optimization problem given by

$$J_\theta^F(F_D) = \max_b \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{b(V) \geq D\right\}(V-D)\right] \tag{1}$$

$$\text{s.t. } y_\theta \leq \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{b(V) \geq D\right\}\right]$$

$$0 \leq \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{b(V) \geq D\right\}(V-D)\right]$$

where the expectation is taken over both F_V (the distribution of valuations) and F_D (the distribution of the maximum competing bid). To define a well-behaved optimization problem, we assume $y_k<\eta s_k$, that is the minimum impression level is lower than the total available impressions.[18]

[18] If the advertiser is assumed to maximize its expected utility given a budget constraint (Balseiro et al. 2015), the constraint in the model will be specified as $B_\theta \geq \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{b(V) \geq D\right\}\right]$ where B_θ is the maximum budget level.

In the first line, ηs_θ indicates the total number of impressions arriving during the campaign period. $\mathbf{1}\{b(V) \geq D\}$ indicates the probability of winning the auction on a given arrival, where the advertiser's bid is higher than the maximum of the competitors' bids. Lastly, $(V - D)$ indicates valuation minus payment, where the payment is consistent with the second-price rule.

In the second line, the right hand side is the expected number of impressions won at the end of the campaign period. The inequality constraint assures that the expected number of impressions won is greater than the minimum impression level y_θ. This inequality can also be written as $\frac{y_\theta}{\eta s_\theta} \leq E_{V,D}[\mathbf{1}\{b(V) \geq D\}]$, implying that the advertiser tries to attain a minimum auction winning rate of $\frac{y_\theta}{\eta s_\theta}$.

The third line captures the advertiser's participation constraint that its expected utility in bidding in the exchange channel is greater than zero. Below we characterize advertisers' optimal bidding strategies, assuming that the participation constraints hold (do not bind) in equilibrium. In sub-section 3.4, we discuss how the participation constraint is imposed in calculating the optimal reserve price.

Proposition 1. *Suppose that $E[D] < \infty$. An optimal bidding strategy that solves (1) is given by*

$$\beta_\theta^F(v|F_D) = v + \mu^*$$

where μ^* is the optimal solution of the dual problem

$$inf_{\mu \geq 0}\, \eta s_\theta E_{V,D}[\mathbf{1}\{V \geq D - \mu\}(V - D + \mu)] - \mu y_\theta$$

That is the advertiser bids higher than its own valuation by a constant factor μ^*, which is the optimal dual (Lagrangian) multiplier of the minimum impression constraint. Intuitively, this means that the advertiser foregoes some utility to satisfy the constraint. This constant factor μ^* guarantees that the advertiser meets the minimum impression constraint at the end of the campaign period. Of note, the dynamic nature of the repeated auctions is captured by this constant factor μ^*, and the bidding strategy becomes static in the sense that μ^* does not depend on the current single auction-specific state.[19]

Proposition 2. *If the participation constraints do not bind in equilibrium, the equilibrium can be*

[19] In the budget constraint case (Balseiro et al. 2015), advertisers will shade down theirs bids by a constant factor to account for the future bidding opportunities (future value). In our minimum impression constraint case, advertisers shade up by this constant factor μ^* to satisfy the constraint at the end of the campaign period.

characterized as follows:

$$\beta_\theta^F\left(v|F_D\right) = v + \mu^*$$

where μ^* *is*

$$\begin{cases} \mu^* = 0 & \text{if } y_\theta < \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{V \geq D\right\}\right] \\ y_\theta - \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{V + \mu^* \geq D\right\}\right] = 0 & \text{if } y_\theta \geq \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{V \geq D\right\}\right] \end{cases}$$

The proposition states that if the minimum impression constraint is not binding, then in equilibrium advertisers will bid truthfully ($\mu^* = 0$). On the other hand, if the minimum impression constraint does bind, then advertisers will bid higher than the true valuations, where μ^* solves the implicit function $y_\theta - \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{V + \mu^* \geq D\right\}\right] = 0$. Of note, $y_\theta - \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{V + \mu^* \geq D\right\}\right]$ is the number of expected impressions shy of the minimum impression level at the end of the campaign, when the optimal bid function is employed.

Based on this proposition, the cost of the minimum impression constraint, μ^*, increases with the minimum impression level (y_θ), and decreases with the number of impressions and length of campaign (ηs_θ). Perhaps more importantly for our purposes, an increase in the second highest payment, D, lowers $E_{V,D}\left[\mathbf{1}\left\{V + \mu^* \geq D\right\}\right]$ and thus increases μ^*, the bid premium. Because an increase in reserves can increase D when the reserve binds (it binds when the second highest valuation is lower than the reserve), reserves can lead to higher bids, consistent with our experimental data. The proofs for Proposition 1 and Proposition 2 are in online Appendix A.1.

The primitives to be estimated are (F_V, F_D, $\boldsymbol{\mu}$) given the reserve price r.

3.4 Optimal Publisher Ad Auction Reserve Price

The publisher's objective is to maximize the long-run expected revenues from the auctions given advertiser valuations. Within the independent private value (IPV) paradigm, the publisher can maximize the revenues from the RTB auctions by choosing the reserve price optimally.

3.4.1 Publisher's Optimization Problem

We denote $G_\theta\left(\boldsymbol{\mu}, r\right) = E_{V,D}\left[\mathbf{1}\left\{V + \mu_\theta \geq D\right\} D\right]$ to be the expected payment of a θ-type advertiser when advertisers bid according to the profile $\boldsymbol{\mu}$, and the publisher sets a reserve price r (G_θ term is the product of the payment times the likelihood the impression is won). We define $I(\boldsymbol{\mu}, r) = F_D\left(r|\boldsymbol{\mu}\right)$ as the probability that the impression is not won in the exchange.

The publisher's problem can then be written as

$$\max_r \eta \left[\sum_\theta \{ p_\theta s_\theta G_\theta (\boldsymbol{\mu}, r) \} + cI(\boldsymbol{\mu}, r) \right] \tag{2}$$

$$\text{s.t. } \mu_\theta \geq 0 \perp y_\theta \leq \eta s_\theta E_{V,D} [\mathbf{1}\{V + \mu_\theta \geq D\}] \ \forall \theta \in \Theta$$

where p_θ is the probability that an arriving advertiser is of type θ and $c > 0$ is the publisher's valuation (i.e., the outside option value if the impression is not won by some advertiser in the exchange).[20]

The first term in the sum in the first line indicates the average expenditure of the advertisers (i.e., the second highest bids when the auctions are won by some advertisers). The second term in the first line indicates the publisher's outside option value when the impression is not won by any advertiser, as it reflects the product of the scrap value of the impression and the probability the impression is not sold. The constraints in the second line reflect the conditions for the Lagrangian multipliers in ensuring the minimum impression constraints (either the constraint binds or it does not).

3.4.2 Advertiser Participation Constraints

In the counterfactual, where we increase the reserve price to the optimal level, some advertisers will start to face binding participation constraints as the reserve price increases.[21] To incorporate the effect of the participation constraints in the policy simulation, we ascertain whether the bidding profile $\boldsymbol{\mu}$ satisfies the participation constraint $(0 \leq \eta s_\theta E_{V,D} [\mathbf{1}\{V + \mu_\theta \geq D\}(V - D)])$ at the considered reserve price level r for $\forall \theta \in \Theta$.

In the case the participation constraint does not hold for some advertiser types θ, we calculate the maximum μ_θ that satisfies the participation constraint such that

$$\bar{\mu}_\theta = \max_{\mu_\theta \geq 0} [0 \leq \eta s_\theta E_{V,D} [\mathbf{1}\{V + \mu_\theta \geq D\}(V - D)]]$$

This is, the advertiser increases its bid to $v + \bar{\mu}_\theta$ to bid as closely as possible to its minimum impression level while satisfying the participation constraint.[22]

[20] Note that, by definition, $I(\boldsymbol{\mu}, r) = F_D(r|\boldsymbol{\mu}) = 1 - \sum_\theta \{p_\theta s_\theta E_{V,D} [\mathbf{1}\{V + \mu_\theta \geq D\}]\}$. Thus, Equation (2) can be interpreted as the publisher's value over selling and not selling an impression.

[21] In the extreme case where all advertisers face binding impressions constraints, advertisers' bidding strategies will reach ∞ as the reserve price increases to ∞ without the participation constraints.

[22] Because of the participation constraints, the advertiser may not achieve the minimum impression level under the counterfactual. In this case, we are assuming that advertisers purchase as many impressions as possible toward the minimum impression level while satisfying the participation constraints. Alternatively, we can specify the advertisers to drop-out all together from the exchange channel when the participation constraints do not meet, but we think the

Given the publisher's outside option value c, the primitives to be recovered from the policy simulation are the optimal reserve price r^* and the corresponding outcomes (F_V, F_D, $\boldsymbol{\mu}$) at r^*.

3.5 The Effect of Constraints on Advertiser Bidding

The theoretical prediction under the standard second-price, sealed-bid auction is that bidders bid truthfully, meaning that it is a weakly dominant action for the advertiser to bid the true valuation for an ad impression (e.g., the value they place on displaying an ad to the consumer). The truth-telling strategy is tractable and often assumed in models of display markets (e.g., Celis et al. 2014, Sayedi 2018) because advertisers' unobserved valuations for the ad impressions can directly be inferred from the observed bids. However, there exists some empirical evidence that advertisers face practical constraints when bidding in the ad exchange auctions (Balseiro et al. 2015; Balseiro and Gur 2019, Ghosh et al. 2009), in which case advertisers may deviate from bidding their true valuations. Therefore, we consider how advertiser facing these practical constraints would change their bidding behaviors when reserve prices are considered. Specifically, we consider the effect of reserve prices when advertisers face (i) a maximum budget constraint (Balseiro et al. 2015; Balseiro and Gur 2019), (ii) a minimum impression constraint (Ghosh et al. 2009), and (iii) no binding constraints.

Table 3: Theory Predictions

Mechanism	Reserve Level		Effect of Imposing Reserve Prices		
	$r = 0$	$r = r^*_{nc} > 0$	Bid CPM	#Impressions Won	Total Payment
Max Budget Constraint	Not Bind	Not Bind	No Change	$-$	$+$
	Not Bind	Bind	$-$	$-$	$+$
	Bind	Bind	$+/-$	$+/-$	No Change
Min Impression Constraint	Not Bind	Not Bind	No Change	$-$	$+$
	Not Bind	Bind	$+$	$-$	$+/-$
	Bind	Bind	$+$	No Change	$+$

Table 3 presents the theory predictions on bid CPM, the number of impressions won, and the total payment, when the reserve price is changed from $r = 0$ (i.e., no reserve price) to $r^*_{nc} > 0$. r^*_{nc} is the optimal reserve calculated under the no constraint, single-shot standard (truth-telling) model.[23] As the reserve price increases from ($r = 0$) to ($r^*_{nc} > 0$), advertisers face tighter constraints, and each row represents a possible scenario of the underlying state: (not bind, not bind), (not bind, bind), (bind, bind) where bind implies the respective constraint binds. Each constraint is considered in isolation. That is, when we consider the budget constraint, we assume that the minimum impression constraint does not bind in both ($r = 0$) and ($r^*_{nc} > 0$). Note that the case of no constraint model is

former is more realistic in our context.

[23] Online Appendix A.2 details the rationale for these predictions.

nested within the non-binding scenario (not bind, not bind).

Table 3 suggests that, by varying the reserve price and observing how the metrics change, we can determine which of these explanations is more consistent with the advertiser behavior given the bids observed in our data. Such a comparison is necessary for developing a pricing model concordant with advertiser behavior and of theoretical interest in its own right. Hence, before estimating a model for the purpose of setting reserve prices, the next section describes an experiment intended to assess which of these constraints is predominant, if any.

4 Experimental Evidence of Constraints

As shown in the preceding section, advertisers' bidding behaviors are affected by the practical constraints advertisers face regarding reach (minimum number of impressions to buy) or budget. Therefore, a set of field experiments were designed to vary the reserve prices in order to understand their impact on the publisher's revenues as well as on the advertisers' bidding behaviors.

In this section, we begin by characterizing the design of these experiments. Next, we detail the results. We find that i) advertisers' bidding behaviors are not consistent with truth-telling but appear to reflect a minimum impression constraint, and ii) our use of theory-based reserve prices leads to a 35% increase in revenues for the publisher, providing concrete evidence there is substantial room to enhance its pricing outcomes.

4.1 Experimental Setting

4.1.1 Paired Randomized Experiments

Twelve paired randomized experiments (Athey and Imbens 2017) were concurrently conducted on two selected websites. The two sites and the twelve pairs were chosen to be closest in terms of ad characteristics, contents, user demographics, revenues, and number of impressions (user visits). Two units in a pair were randomized into the treatment and the control group. For example, (site1, desktop, below-the-fold, 300x250) and (site1, desktop, above-the-fold, 300x250) were paired for the first experiment, and the randomly chosen (site1, desktop, below-the-fold, 300x250) was assigned to the treatment group, whereas (site1, desktop, above-the-fold, 300x250) was assigned to the control group. Table 9 in the online Appendix includes the details of the experimental pairs and the assignment of treatment/control groups. Each experimental pair is referred to by an Experiment Pair ID (i.e., Experiment 1 to 12). The unit in a pair (i.e., a treatment or a control) is henceforth referred to as an experimental cell.

The sample of experimental observations collected between 10/18/2017 - 12/03/2017 constitute the post-treatment period when the reserve prices were changed for the treatment group. The same duration sample from 08/30/2017 - 10/15/2017 constitutes the pre-treatment period. Both the pre- and post- periods are used for the difference-in-difference (DiD) analyses. Overall, we observe $5,382,829$ payment data observations from all advertisers and $3,635,899$ bid data observations from the opt-in advertisers at the (advertiser-DSP-day-site-ad type) level.

The table of balance reported in online Appendix (Table 10) shows that none of the observables are statistically different between the treatment and control groups in the pre-treatment period.[24] Nevertheless, to control for potential level differences between the treatment and control groups, our identification strategy is to compare the treatment and the control within a pair over time by conducting a DiD analysis. Appendix B.1.3 discusses the pre-trend assumption required for our identification.

4.1.2 Setting Reserve Prices

As advertisers vary bids across observed ad characteristics such as site, device, ad location, and size (Table 2), the publisher can price discriminate to enhance its revenues by setting different reserve prices across these observed ad characteristics. Hence, we design the experimental cells to set reserve prices at the site-ad type (device, ad location on a page, size) level.[25]

The experimental reserve prices in the treatment condition were calculated using a naive theoretical model presuming that advertisers do not face any binding constraints and that advertisers play a single-shot game for each available ad impression. All that is required to test the theoretical predictions of a binding budget or impression constraint is to ascertain if and how bids change with reserve prices. That is, the level of the treatment reserve prices need not be at the optimal level for testing, merely that the reserve prices vary exogenously. The suggested experimental reserve prices based on the naive, standard model provide good initial points and are likely to increase the publisher's revenues toward the global maximum relative to the condition of no reserve

[24]Exact matching pairs of cells is infeasible. For example, matching (Site1 above-the-fold, Desktop, 300x600) with (Site2, above-the-fold, Desktop, 300x600) could yield different average bid CPM levels in the pre-treatment period, because Site 1 and Site 2 though similar are not identical.

[25]Although possible to further vary reserve prices across days, advertisers, and consumers, we do not because the publisher's concern for fairness across advertisers and its preference toward a simpler pricing scheme prohibit this practice. In this regard, the effect we report on the publisher's revenue when setting reserve prices at the (site-ad type) level can be considered as a lower bound to the gains possible in a more flexible system.

prices. Further improvements on the publisher's revenue are possible when practical constraints are considered in setting the reserve prices, as we shall discuss later.

During the experiments, the reserve prices were set to these calculated levels for the treatment group, while they were kept at the historical levels for the control group (i.e., no reserve prices).[26] The details of calculating the experimental reserve prices can be found in online Appendix B.1.4.

4.2 Experimental Results

We begin by reporting experimental results regarding the publishers' revenues to gauge the effectiveness of setting reserve prices in auctions. Next, we discuss the effect of reserve prices on advertisers bidding behaviors to asses which of the constraints (reach or budget) is predominant, if any.

4.2.1 Effect on eCPM

The outcome measure considered is *eCPM* (effective CPM, industry vernacular), which yields a per supplied impression revenue.

$$eCPM = \frac{\text{Revenue}}{\text{\# Impressions Supplied to Ad Exchange (in thousand)}}$$

Table 4 shows the treatment effect on eCPM, where eCPM is multiplied by a common, multiplicative constant for confidentiality. The increase in revenues (holding the number of impressions supplied to the ad exchange the same) is 35%, thereby affirming the importance of setting the reserve prices in running the auctions.

$$\begin{aligned}\text{Increase in revenue} &= \frac{(0.49 - 0.37) - (0.36 - 0.37)}{0.37} \\ &= \frac{0.13}{0.37} \simeq 35\%\end{aligned}$$

Table 4: Treatment Effect on eCPM ($)

Group	Reserve	Pre (08/30/17 - 10/15/17)	Post (10/18/17 - 12/03/17)
Treatment	Yes	0.37	0.49
Control	No	0.37	0.36

Figure 4 plots the treatment effect by the experimental pair, and includes a dotted horizontal line

[26] For Experiments 8 and 10, positive reserve prices were in place in the pre-period. We provide additional discussions on these two cells in sub-section B.3 in online Appendix.

representing the overall percentage DiD change in eCPM across experiments, 35%. The treatment effects on eCPM from various DiD specifications (e.g., with various control variables) are qualitatively similar and reported in online Appendix B.2.1.[27]

Figure 4: Treatment Effect on eCPM by Experiment

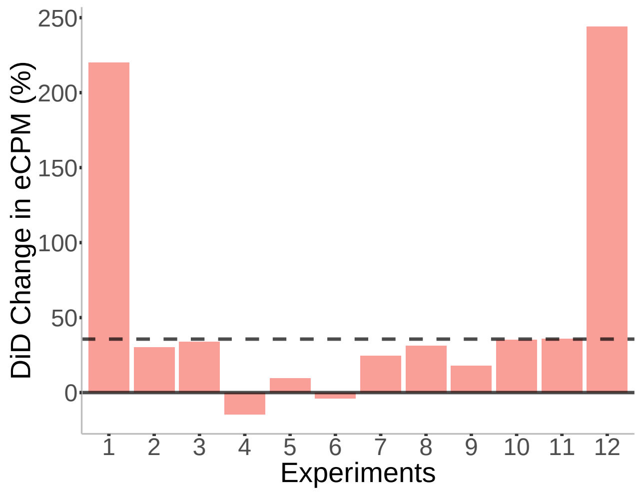

4.2.2 Effect on Bidding Behaviors

To assess how advertisers' bidding behaviors are affected by the exogenous increase in reserve prices, we consider their bid CPM data. If advertisers bid truthfully, the distribution of bids will remain invariant to the experimental manipulation of the reserve price. However, this is not the case. Figure 5 plots the distribution of changes in bid CPMs for those advertisers in the treatment group, bidding both in the pre- and post- periods for Experiment 1. The changes are computed as the advertiser's average bid CPM in the post- minus that in the pre-period, then these bid CPM changes are multiplied by a common, multiplicative constant to maintain data confidentiality. The probability density function's mass is greater in the positive region, suggesting advertisers increase their bids in the post-period in response to an increase in the reserve price. The p-value for the Kolmogorov-Smirnov test against a standard normal distribution is < 0.001.

While Figure 5 is suggestive that advertisers increase bid CPMs in response to the change in the reserve prices, there might exist confounding factors, like seasonality. Figure 6 plots the the average

[27] With only 12 experiments, the potential to test the underlying mechanism driving the heterogeneity in treatment effects across cells is limited. Auction theory suggests that the gains from using reserve prices would be large when auctions are thin thus the reserve prices often bind. Running the regression of (DiD % change in eCPM by experiment) on (DiD % change in ebids by experiment) where (ebids= #bids for opt-in / # impressions supplied to ad exchange' of the treatment group in the post-period) yields a negative coefficient, though it is not statistically significant possibly due to the small sample.

Figure 5: Changes in Bid CPM ($)

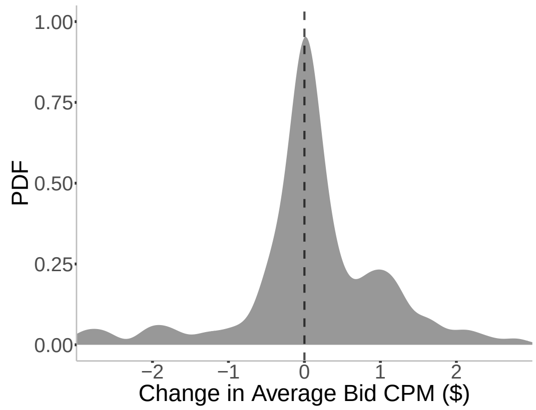

bid CPM trend line, where the vertical dotted line indicates the start date of the experiments. In the pre-period, the control group mean level is higher than the treatment group. But the mean level for the treatment group becomes higher as soon as reserve prices are imposed. Even though reserve prices are not published, advertisers respond to the changes in reserve prices within days by adjusting their bidding strategies, suggesting they can quickly infer the reserve price levels based on the outcomes of their many bids.

Figure 6: Bid CPM Trend

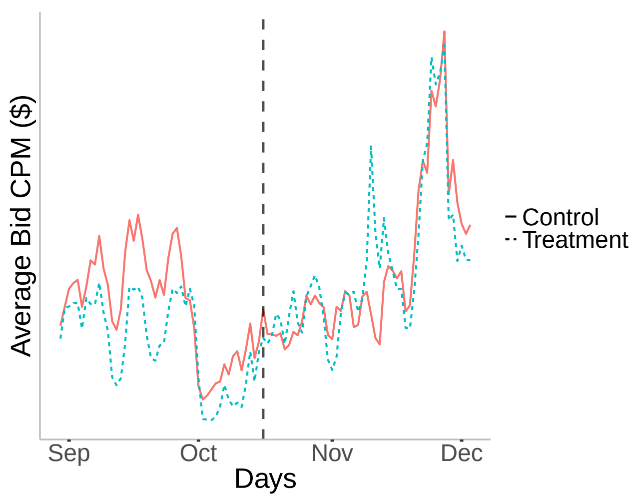

Note: Y-axis levels are not displayed for confidentiality.

The first column in Table 5 presents the result from a DiD regression of bid CPM. Day fixed effects are included to control for the time trend such as seasonality. Further, we include advertiser fixed effects as our emphasis is on within advertiser bid changes over time and the number of ad impressions to control for inventory. Each observation in the DiD regression is weighted by the number of bids submitted.[28] The increase in (scaled) bid CPM with respect to the exogenous increase in reserve prices is \$0.039. This 19% increase suggests that advertisers do not bid their true valuations.[29, 30]

Table 5: Treatment Effect on Bidding Behaviors

DV (Scaled)	Bid CPM (\$)	# Impressions Won	Total Payment (\$)
Treated × Post	$\mathbf{0.039}^{**}$	0.351	$\mathbf{0.005}^{**}$
P-value (Randomization Inference)	0.037	0.397	0.048
# Impressions Supplied (in thousand)	$\mathbf{-1.84e-04}^{***}$	$\mathbf{5.65e-02}^{***}$	$\mathbf{9.44e-05}^{***}$
Treated	Y	Y	Y
Day	Y	Y	Y
Advertiser	Y	Y	Y
R-squared	0.212	0.197	0.188
Observations	3,635,899	5,382,829	5,382,829

Note: The dependent variables are all multiplied by a common, multiplicative constant for confidentiality. The unit of observation for the analysis is (advertiser-DSP-day-site-ad type). Bid CPM analysis uses 3,635,899 bid data observations from the opt-in advertisers. Ad impression and total payment analyses use 5,382,829 payment data observations from all advertisers. The P-value for testing the null hypothesis that the treatment has no effect is calculated using randomization inference (Athey and Imbens 2017). For randomization inference, we randomize treatment assignment at the experimental pair-level, the same way that the assignment was done in the experiment). ***denotes 1% and **5% significance.

4.3 Theoretical Rationale for Experimental Results

In addition to the bid CPM analysis reported in Table 5, the DiD analyses are also conducted on the number of impressions won and the total payment. The second and third columns in Table 5 respectively indicate that the change in impressions bought is not statistically significant, while the (scaled) total payments significantly increase by \$0.005, or about 29%. Combining these two results with the observation that bid CPM increases upon raising the reserve prices suggests that advertisers' bids are affected by practical constraints. Specifically, the results are most consistent with the theory predictions presented in the last row in Table 3 in which the minimum impression constraint binds for at least some advertisers facing a change from zero to positive reserve prices. Intuitively, when the reserve price increases, the probability of winning an impression decreases for an advertiser. The advertiser therefore increases bid CPMs in order to achieve at least the

[28] Recall that a bid CPM observation is a daily average (i.e., total bidding amount/#bids) for a given observational unit (advertiser-DSP-day-site-ad type). Thus the weighted DiD regression uses the number of bids for a given observational unit as the weight.

[29] It is also worth noting the significant and negative coefficient for the number of impressions supplied. With no binding constraint, advertisers would bid their true valuations regardless of the number of available impressions. The decrease in bid CPMs with respect to ad impressions is consistent with advertisers conditioning on multiple opportunities to bid under a binding constraint when forming their bidding strategies.

[30] Online Appendix B.2 and B.3 include robustness checks and associated discussion. Our findings are qualitatively robust to various model specifications and potential concerns, for example, arising from observing daily averages instead of the auction level outcomes.

minimum impression level set for an advertising campaign.[31] This bidding behavior increases the total advertiser payment and keeps the number of winning impressions the same. Were advertisers to consider both the minimum impression level (reach) and the budget in setting their ad campaign and bids, one or the other constraint would be more binding. Our data are consistent with the minimum impression constraint affecting advertisers' bids more tightly than the budget constraint.

Based on these findings from the experiments, the advertiser bidding model is estimated with the minimum impression constraint to allow advertisers to depart from the commonly adopted truth-telling strategies. In the next section, we outline a non-parametric approach for estimating advertiser valuations with the minimum impression constraint and discuss model identification.

5 Identification and Estimation

This sub-section discusses the identification and estimation strategies for the inference of advertiser valuations.

5.1 Identification

As the publisher currently imposes no reserve price, the identification strategy is described conditioned on no reserve price.[32] First we characterize the identification strategy when the impression constraint does not bind ($\mu^* = 0$), then we discuss the case of the binding constraint ($\mu^* > 0$).

5.1.1 Case1: $\mu^* = 0$

The equilibrium-bid function in Proposition 1 implies that advertisers bid truthfully when the minimum impression constraint does not bind. Thus, the probability density function of observed bids can directly be mapped to that of valuations. In other words, the distribution of bids identifies the distribution of valuations such that

$$f_V^0(v) = f_B^0(b)$$

$$F_V^0(v) = F_B^0(b)$$

[31] As bid CPMs incorporate advertisers' underlying valuations for the ad impressions, an alternative interpretation for our results could be that valuations change with reserves. This might occur were an increase in reserve prices to lead to fewer impressions being sold and therefore less ad clutter. Less clutter might increase users' attention or click-through-rate (CTR) for the displayed ads and therefore the advertiser valuations. We don't think this chain of reasoning is likely in our setting for two reasons. First, the publisher shows house ads (e.g., cross-promoting publisher's other brands/sites) in the unsold ad slots, thus keeping clutter relatively constant. Second, the DiD analysis of CTR shows no evidence of CTR change due to the increase in reserve prices. Less clutter would imply more clicks.

[32] When the reserve prices exit, the identification strategy will be similar, but the distribution of bids will represent a truncated distribution of valuations due to endogenous participation.

where f_V^0 and F_V^0 (f_B^0 and F_B^0) represent probability density function and cumulative density function of valuations (bids) respectively. The superscripts "0" denote the true population values.

Because only the payment data are used in estimation (see sub-section 2.2), the distribution of the observed payments needs to be linked to the distribution of valuations. Denoting D to be the winning payments, F_D^0 is defined as the distribution of the second highest bids. The distribution of order statistics implies that

$$\begin{aligned} F_D^0(v) &= n(n-1)\int_0^{F_V^0(v)} u^{n-2}(1-u)du \\ &\equiv \varphi\left(F_V^0(v)|n\right) \\ f_D^0(v) &= nf_v^0(v)(n-1)F_V^0(v)^{n-2}\left(1-F_V^0(v)\right) \end{aligned}$$

In the last line, $nf_v^0(v)$ indicates that one of the n advertisers (i.e., $\begin{pmatrix} n \\ 1 \end{pmatrix} = n$) draws v exactly, and $(n-1)F_V^0(v)^{n-2}\left(1-F_V^0(v)\right)$ indicates that $(n-2)$ out of the remaining $(n-1)$ advertisers (i.e., $\begin{pmatrix} n-1 \\ n-2 \end{pmatrix} = n-1$) draw valuations lower than v, $\left(F_V^0(v)^{n-2}\right)$, and 1 advertiser draws valuation higher than v, $\left(1-F_V^0(v)\right)$.

Since $\varphi(\cdot|n)$ is a strictly monotonic function given n, F_V^0 is identified from the distribution of observed winning payments F_D^0 when the number of potential bidders n is known (Paarsch, Hong, et al. 2006).

5.1.2 Case2: $\mu^* > 0$

When the minimum impression constraint binds, the distribution of bids identifies the distribution of valuations up to a location constant μ^* such that

$$\begin{aligned} f_W^0(w) &= f_B^0(b) \\ F_W^0(w) &= F_B^0(b) \end{aligned}$$

where $w = v + \mu^*$ and μ^* are the optimal Lagrangian multipliers. Thus $f_W(W) = f_W(v + \mu^*)$. Assuming we observe the advertiser type $\theta = (s, y)$, μ^* can be estimated from the conditions in Proposition 2, such that μ^* solves $y_\theta - \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{V + \mu_\theta^* \geq D\right\}\right] = 0$ if $y_\theta \geq \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{V \geq D\right\}\right]$.

Combining the argument in Case 1, F_V^0 is identified from the distribution of observed winning payments F_D^0, when the number of potential bidders n and the advertiser types (i.e., minimum impression level, campaign duration) are known.

5.2 Estimation

The estimation strategy is detailed in this subsection. As in the preceding section, we discuss the case when the constraint does not bind ($\mu^* = 0$) first, then incorporate the case of the binding constraint ($\mu^* > 0$).

5.2.1 Case1: $\mu^* = 0$

Under the truth-telling scenario, F_D^0 can be estimated by substituting the sample analogue for the population quantity as

$$\begin{aligned} \hat{F}_D(v) &= \frac{1}{T}\sum_{t=1}^{T} \mathbf{1}\left[d_t \le v\right] \\ &= n(n-1)\int_0^{\hat{F}_V(v)} u^{n-2}(1-u)du \end{aligned} \tag{3}$$

The asymptotic properties of the estimator $\hat{F}_D(v)$ (and the resulting $\hat{F}_V$) are discussed in Paarsch, Hong, et al. 2006.

Incorporating (discrete) covariates is possible as follows. For a given characteristic $z \in \boldsymbol{Z}$, the estimator of $F^0_{V|Z}(v|z)$ is specified as

$$\hat{F}_{D|Z}(v|z) = n(n-1)\int_0^{\hat{F}_{V|Z}(v|z)} u^{n-2}(1-u)du \tag{4}$$

Because the number of participants vary across auctions, n-specific non-parametric empirical cumulative distribution functions (ECDFs) are estimated first for a particular combination of the z (by n-specific, we mean an ECDF is inferred separately for each auctions with n bidders). Then $\hat{F}_{V|Z}(v|z)$ is obtained by kernel smoothing across the n-specific ECDFs to obtain a single ECDF. The exceptionally large number of payments observed given a particular combination of ad characteristics facilitates inference in this context of display ad markets.[33]

5.2.2 Case2: $\mu^* > 0$

The key to incorporating the minimum impression constraint is estimating $F_{V|Z}$ together with μ^* for a given z. To simplify notation, we drop the dependence on z. The estimation is done in two stages. In the first stage, we estimate location shifted distribution $F_V^0(w)$ where $w = v + \mu$. In the

[33] The observed characteristics affecting the valuation are discrete in our context, such as site, ad type (ad location, ad size, device), and time (month). Although the dimension of $\boldsymbol{Z}$ considered is large, we also have many observations per given particular combination of Z, enabling non-parametric estimation. If the covariates are continuous, semi-parametric approaches such as single-index models (e.g., the density-weighted derivative estimator in Powell et al. 1989, the maximum rank-correlation estimator in Han 1987) can instead be used to reduce the curse of dimensionality.

second stage, conditioned on the recovered distribution $F_V^0(w)$, μ is estimated using the condition in Proposition 2.

One challenge in estimating valuations when the impression constraint binds is that D, the distribution of maximum of competing bids, is a function of others' bids. As such, advertisers need to form beliefs about the bids of other advertisers. In estimation, this maximum is observed and equivalent to the rational expectation of the advertisers' regarding the maximum of the competitors' bids. The estimation procedure is outlined in the online Appendix C.1.

5.3 Institutional Details

This sub-section discusses four institutional aspects of the data that warrant additional attention: (i) metrics (e.g., payments, number of impressions won) are only available as daily averages instead of at the auction (impression) level, (ii) the number of potential bidders, n, are not directly observed, (iii) how the minimum impression level and campaign lengths are operationalized, and (iv) which data points are included in Z.

Daily Aggregate Data The metrics (e.g., payments, # impressions won) provided by the ad exchange are aggregated to day level, and are not available to this or other publishers at more granular levels. More specifically, an observational unit in the exchange data represents total payments, number of impressions won, and number of clicks attained for each given advertiser-DSP-day-site-ad type. For each observational unit, the average (daily) CPM paid is calculated as (total payments / number of impressions won). This average (daily) CPM paid is used as d_t in Equation (3) in forming the estimator for $\hat{F}_D(v)$.

As each observation represents a different number of impressions won, we weigh the average CPM paid by the number of impressions won when estimating the distribution of valuations. That is, a data point (y average CPM paid, x number of impressions won) is treated as if there are x number of observations with y CPM payment.[34]

Number of Potential Bidders The identification strategy discussed above requires that the number of potential bidders n is known. The number of potential bidders for a given observational

[34] The magnitude of the potential bias in using the aggregate data (as opposed to the impression level data) is to be examined. Our conjecture is that the variance of the valuation distribution will be biased downward when using the aggregate data, but its impact on the optimal reserve price level itself will be small (or minimal).

unit (advertiser-DSP-day-site-ad type) will be

$$n = n_1(\text{\# advertisers with CPM payment, i.e., positive \# impressions won})$$
$$+ n_2(\text{\# advertisers with bids submitted, but zero impressions won})$$

n_1 is observed in the data, while n_2 is observed only for the opt-in advertisers. Thus, to operationalize n, we use the the following proxy:[35]

$$\hat{n} = n_1(\text{\# advertisers with CPM payment, i.e., positive \# impressions won})$$
$$+ n_{2,opt-in}(\text{\# opt-in advertisers with bids submitted, but zero impressions won})$$

Minimum Impression Level and Campaign Length When advertisers do not face binding minimum impression constraints, advertisers bid truthfully and the distribution of valuations can be recovered without the knowledge of the minimum impression level or the campaign length. In the case the minimum impression constraints bind for some advertisers, the identification strategy discussed in sub-section 5.1.2 requires more information from the researcher. More specifically, in forming the estimator in Equation (10) in the online Appendix, $\tau_\theta = \left(\frac{y_\theta}{\eta s_\theta}\right)$ term is required as an input. This term reflects the minimum winning rate the advertiser aims to attain for a given impression for a campaign. Thus, τ_θ is advertiser-campaign specific. In the estimation (and the counterfactual), we allow τ to vary across advertisers, but hold constant the minimum winning rate within an advertiser across campaigns. Denoting k to be the advertiser, τ_k is proxied by

$$\hat{\tau}_k = min\left(I_{k1}, ...I_{km}, ...I_{kM}\right),\ I_{km} > 0\ \forall m$$
$$I_{km} = \frac{1}{\sum_{day\in m} \mathbf{1}\left[i_{k,day} > 0\right] u_{day}} \sum_{day\in m} i_{k,day}$$

where $i_{k,day}$ represents the impressions won by advertiser k and u_{day} represents the impressions available for sale on a given day. I_{km} in the second line constructs the winning rate for a given month m, conditional on participating in the exchange. In other words, $\hat{\tau}_k$ is computed as the minimum of the monthly winning rates observed in the data, conditioned on the advertiser participating in the exchange. Accordingly, the optimal bidding strategy profile $\boldsymbol{\mu}|\boldsymbol{Z}$, is recovered at the advertiser level, instead of advertiser-campaign level.[36]

[35] $n_{2,opt-out}$ will be small for two reasons. First, 82% of advertisers opted-in to share their bidding information (default setting). Second, the opt-out advertisers (18%) constitute about 84% of the total revenues, meaning opt-out advertisers are highly likely to be included in n_1, which is observed.

[36] The identification strategy discussed in subsection 5.1 relies solely upon the observational data. Alternatively, one could leverage the experimental data (i.e., the causal change in bid CPMs) to impute the minimum impression level. Specifically, the optimality condition in Proposition 2 implies a one-to-one mapping between i) the minimum impression level and ii) the causal change in bid CPM in the experimental data. However, our identification strategy

The solution to the publisher's optimal reserve price in Equation (2) also involves the term $\delta_\theta = p_\theta s_\theta$ which represents the share of each advertiser's type present in the auction (where type θ is uniquely defined by a campaign length and a minimum impression level). This share of advertiser types is determined by the product of the advertiser type's arrival probability, p_θ, and the campaign duration, s_θ; this product represents the fraction of advertisers in an auction that play bidding strategy profile $\boldsymbol{\mu}_\theta$. We do not observe p_θ and s_θ, but can nonetheless compute δ_θ at the advertiser-type level (that is, setting $\theta = k$). We do this by approximating the weight for each advertiser by its observed share of participation. Thus, we use below proxy for $\delta_{\theta=k}$

$$\hat{\delta}_k|Z = \frac{\sum_{day} \mathbf{1}\left[i_{k,day} > 0\right]}{\sum_k \sum_{day} \mathbf{1}\left[i_{k,day} > 0\right]}$$

where $i_{k,day}$ represents the impressions won by advertiser k on a given day. $\mathbf{1}\left[i_{k,day} > 0\right]$ takes value 1 if the advertiser k wins a positive amount of impressions on a given day. We include 'month' as the observable characteristics in Z, so the summation is done over days within a month. This $\hat{\delta}_k|Z$ is the weight advertiser k (with the bidding strategy profile μ_k) plays toward the platform's revenue given Z.

Observed Covariates The distribution of valuations is estimated separately for each combination of Zs to control for heterogeneity. The observed covariates considered in Z are

- Site: 20 U.S. based sites are considered.
- Ad location: these include above-the-fold (ATF), MID, below-the-fold (BTF), and no information available.
- Ad size: are represented by 300x250, (728x90, 970x66), 320x50, and 300x600. These are the sizes conforming to Interactive Advertising Bureau (IAB) standard guideline and are most commonly used by the advertisers.[37]
- Device: the various devices include desktop, mobile, and tablet.
- Month: controls are included for seasonality.

In sum, we estimate the distribution of valuations for each $11,520$ (20x4x4x3x12) combination of

based on the observational data enables us to (i) recommend optimal pricing schemes for those sites upon which we have not conducted experiments and to (ii) reserve experimental data for providing a validation for our model.

[37] https://www.iab.com/guidelines/

Zs.[38]

6 Results

This section outlines the advertiser bidding model results used to infer the advertiser valuation distribution. Recall, estimation recovers i) the distribution of the underlying advertiser valuations F_V, and ii) the vector of optimal Lagrangian multipliers $\boldsymbol{\mu}^* = (\mu_1^*, .., \mu_\theta^*, ..\mu_\Theta^*)$ which reflects the tightness of the advertisers' minimum impression constraints. The pair $(F_V, \boldsymbol{\mu}^*)$ is obtained for each combination of observables $\boldsymbol{Z}$ (discrete space of site-ad location-ad size-device-month). Below we report the estimates for the treatment groups in the experimental data in the month prior to the experiment (09/2017).

6.1 The Valuation Distribution, F_V

Figure 7: Advertiser Valuations Distribution: Experiment 1

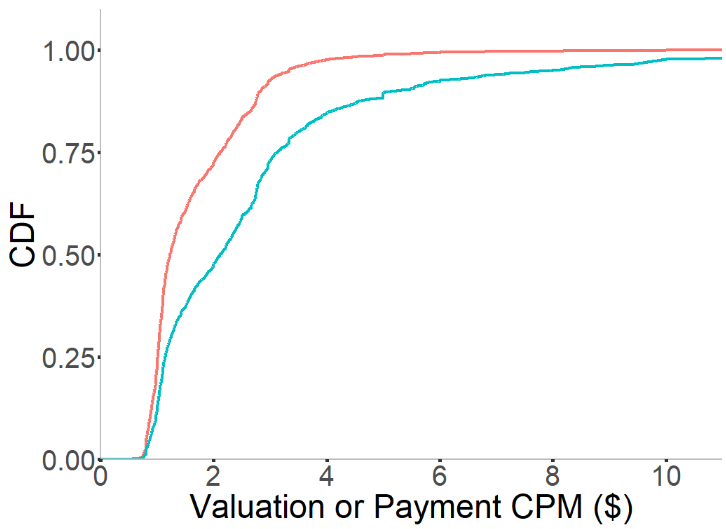

Note: The blue line is the recovered advertiser valuations distribution (F_V) and the red line is the observed payments (F_D).

The cumulative density function of the advertiser valuations, F_v, is plotted in Figure 7 for Experiment 1. The cdf of the advertiser valuations (F_V, blue line) is recovered from the observed payments (F_D, red line).[39] This figure shows how valuations exceed payments. For example, about 40% of advertisers have a valuation of less than \$2.00, and 60% of advertisers pay less than \$2.00.

[38]Advertiser (or DSP)-specific distribution of valuations will be explored in the future to control for the heterogeneity across advertisers (or DSPs).

[39]The recovered cdf of the advertiser valuations for all experiments are included in Figure 11 in online Appendix.

Table 6: Advertiser Valuations F_V: Mean and Standard Deviation

Experiment	Mean	SD
1	3.08	5.72
2	2.58	3.84
3	2.96	3.69
4	4.07	4.84
5	2.71	3.27
6	2.74	3.79
7	3.19	3.48
8	15.0	6.76
9	4.77	5.26
10	8.11	4.25
11	2.48	2.98
12	3.04	4.28
(Weighted) Mean	3.94	4.17
(Weighted) SD	2.84	0.98

Note: The weight used is the number of impressions won for each experiment.

In Table 6, the first and the second columns respectively represent the mean and the standard deviation of the advertiser valuations for each treatment group in the experiment. The standard deviation of the means for the distributions is large (2.84). In other words, the valuation distributions appear to vary by observables across cells such as (site-ad location-ad size-device-month), implying that different reserve prices should be set for different auctions.

6.2 The Impression Constraint, $\boldsymbol{\mu}^*$

Recall, μ^* is the premium advertisers pay to meet the minimum impression constraint (over the solution where there exists no such constraint). The first column in Table 7 reports the percentage of advertisers bound by the minimum impression constraint (i.e., those with positive Lagrangian multipliers $\boldsymbol{\mu_k}^* > 0$) at $r = 0$. On average, about 19% of the advertisers face binding minimum impression constraints during the prior experimental period when $r = 0$. Moreover, as the average advertiser valuation for an impression is about \$3.94, a "back of the envelope" calculation suggests that advertisers bid around 20% (0.78/3.94) higher than their true valuations because of the binding minimum impression constraint when there is no reserve, $r = 0$. Also of note, the standard deviations of $\boldsymbol{\mu}$ are high (see fifth column in Table 7), meaning the cost of the constraint varies significantly across advertisers.

The percent of advertisers bound by the constraint increases from 19 to 32 as the publisher increases the reserve prices. As a result, the cost of the constraint increases from \$0.78 to \$0.99 and the overall cost of the constraint across advertisers grows substantially, from 20% to 25%.

Table 7: The Impression Constraint, $\boldsymbol{\mu}^*$

Experiment	% Advertisers Constrained		$\boldsymbol{\mu}$: Mean		$\boldsymbol{\mu}$: SD	
	$r=0$	r^*	$r=0$	r^*	$r=0$	r^*
1	19.0	33.5	0.66	0.77	4.68	3.33
2	15.3	36.1	0.54	0.96	2.54	3.51
3	17.1	35.8	0.62	1.08	2.27	3.88
4	17.4	22.0	0.93	0.97	3.45	3.42
5	27.8	38.4	0.72	0.78	15.3	15.3
6	26.1	56.4	0.43	0.52	5.53	5.39
7	19.4	26.5	0.88	0.90	5.51	3.85
8	9.19	9.2	1.26	1.35	5.87	5.80
9	21.4	45.9	1.34	1.85	5.63	5.59
10	18.2	17.7	1.27	1.43	3.06	3.92
11	16.7	37.5	0.62	1.08	2.70	3.55
12	26.1	55.1	0.79	0.91	11.1	10.9
(Weighted) Mean	18.7	31.7	0.78	0.99	4.94	5.07
(Weighted) SD	4.35	10.7	0.24	0.22	3.91	3.66

7 Setting Reserve Prices

In this section, we compute the optimal reserve price to maximize the publisher's revenues in the ad exchange auctions. Based on the recovered advertiser distribution F_V, the optimal reserve can be obtained by solving the publisher's optimization problem prescribed in Equation (2).

7.1 Solving for the Optimal Reserve Price

This sub-section discusses the numerical approach used to solve for the optimal reserve price. For each experimental group, we compute two reserve prices. The first reserve price, r^*_{nc}, is calculated under the assumption that advertisers bid truthfully, that is, they face no impression constraint ($\boldsymbol{\mu}^* = 0$),. This can be done by solving the implicit function in Equation (11) in online Appendix.

The second reserve price, r^*, is calculated with the minimum impression constraint in place by solving the publisher's optimization problem in Equation (2). The second reserve price r^* takes into account advertisers' best responses with the minimum impression constraint in FMFE.

The computation of the second reserve price requires updating advertisers' beliefs on D to ensure that advertisers' beliefs are consistent with the bidding profile $\boldsymbol{\mu}$ at the new reserve price r considered. Thus solving the optimal reserve price r^* with the minimum impression constraint involves embedding the iterative best-response algorithm to find FMFE under the new reserve price. The detailed procedure is included in online Appendix C.2.2.

7.2 Optimal Reserves

Table 8 details the predicted reserve price bias and the profit loss arising from ignoring the minimum impression constraint across all the experimental cells. The reserve price bias is calculated as the

Table 8: Policy Simulation Results

Experiment	Reserve Price Bias $\left(\frac{r^*_{nc}-r^*}{r^*}\right)$	Profit Loss
1	-18.2	4.55
2	-59.3	29.3
3	-48.1	12.7
4	-28.6	5.90
5	-33.2	14.3
6	-5.71	5.50
7	-7.41	1.04
8	0.0	0.0
9	-40.0	17.7
10	-5.56	0.09
11	0.0	0.0
12	-4.98	4.98
(Weighted) Average	-27.2	9.09

(optimal reserve without the constraint - optimal reserve with the constraint)÷(optimal reserve with the constraint). The profit loss is calculated as the difference in profits when setting the optimal reserve with the constraint and without. On average, the optimal reserve price calculated with no constraint is about 27% lower than the optimal reserve calculated with the minimum impression constraint, but can be as high as 59%. Further, the profit loss in ignoring this constraint is calculated to be around 9% across experiments, but can be as high as 29%. From this, we conclude that the minimum impression constraint can lead to even more substantial gains in publisher revenues than a naive approach that assumes no constraint.

8 Model Validation

The experiment in Section 4 showed that setting reserves under the naive assumption of a one-shot auction with no constraints led to a 29% increase in auction revenues relative to the context of no constraint. While substantial, this lift raises the question of how much higher profit lift could have been if reserves were instead set using the minimum impression constraint model.

Based on the advertiser distribution F_V recovered in the previous section, it is possible to compute what the optimal reserve should have been by solving the publisher's optimization problem prescribed in Equation (2). Online Appendix D.2 details this calculation, and suggests that the use of the minimum impression constraint model would have led to an additional 9 percentage points of profit over the naive model (i.e., the lift would have increased from 29% to 38%).

While this forecast of what might have occurred had we used our optimal reserve pricing model is informative it remains a forecast. To obtain concrete empirical evidence that the minimum impression approach can increase revenues over a naive model with no constraint or a budget constraint model, this section reports the result of a validation exercise where we randomized reserve prices over a much larger range of experimental cells (than just 12).

8.1 Validation Experiment Design

8.1.1 Paired Randomized Experiments

The goal of the validation experiment was to gather new data and see how well our model predicts and if our reserve prices with minimum impression constraint is indeed optimal. A corollary benefit of the validation exercise is to assess the robustness of our model with respect to its various assumption. In general, to the degree the model assumptions differ from the true data generating mechanism, our predictions of the reserve prices where revenues are expected to be maximized will not be correct; that is, the observed optimal reserve will differ from the predicted optimal reserve.

The validation experimental data in Wave 3 (Figure 1) were collected for the period 04/01/2019 - 07/31/2019, where 04/01/2019 - 06/20/2019 constitute the 'Pre' period and 06/21/2019 - 07/31/2019 constitute the 'Post' period where the changes in reserve prices took place for the treatment group. Thirty one paired randomized experiments were conducted across all publishers' websites and the pairs were chosen to be closest in terms of contents, user demographics, revenues, and number of impressions (user visits).[40]

8.1.2 Setting Reserve Prices

We set a broad range of reserve prices for the treatment group that span both the naive optimal reserve prices, r^*_{nc} and the optimal reserve prices implied by our minimum impression constraint model, r^*. This broad span of reserve prices enables us to ascertain whether or not the reserve prices implied by the minimum impression constraint model generate higher revenues than any other reserve price levels. For the control group, the reserve prices were kept at the non-zero, ad-hoc reserve prices that were set by the publisher prior to the validation experiments.[41]

[40] The online Appendix B includes the details of the validation experimental pairs and the assignment of treatment/control groups. Each experimental pair is referred to by an Experiment Pair ID (i.e., Experiment 1 to 12). The unit in a pair (i.e., a treatment or a control) is henceforth referred to as an experimental cell. Also, the table of balance shows that none of the observables are statistically different between the treatment and control groups in the pre-treatment period for the validation experiment. The pre-trend assumption required for our identification is also discussed in this online Appendix section.

[41] Our results from the first set of experiments motivated the publisher to raise their the reserve prices to non-zero levels.

8.2 Validation Experiment Results

Figure 8 presents the results from the validation experiments. For each of the experimental cell, we first estimate the optimal reserve prices implied by our minimum impression constraint model and compare that with the reserve price level exogenously set during the experiment for the treatment group. The x-axis plots this deviation from the optimal constrained reserve price. On the y-axis, we plot the percentage DiD change in eCPM.

Figure 8: Validation Experiment Result

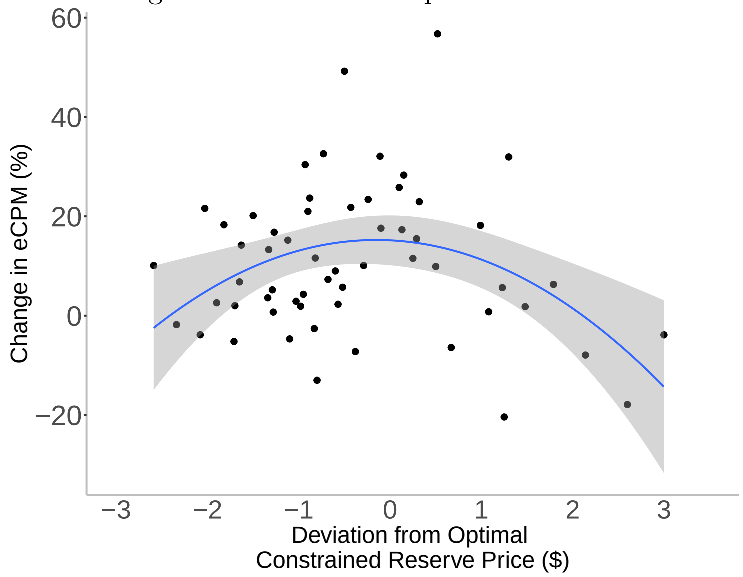

The data suggest that the minimum impression constraint solution is better because the peak of the curve is close to zero deviation. Conditioned on the relatively small sample available, this pattern (peak at zero deviation) is what we would hope to observe if the minimum impression constraint approach increases revenues over a naive model with no constraint.[42]

9 Conclusion

With the continued rapid growth in display advertising markets, there is an increasing value in characterizing the advertisers' valuations for ad impressions, and how advertiser strategies are affected by practical constraints (such as reach or budget) or by the reserve price in advertising

[42]Ideally, one would want to test if the change in eCPM with the minimum impression constraint approach is statistically different from the naive model. A formal test will involve calculating the loss in prediction error for each of these cells and conducting a statistical test such as KS test. This approach is not only computationally intensive, but also not practical in our context. Given we have 31 data points (31 paired randomized experiments) and given these experiments are highly costly for the publisher to run (e.g., the revenue risk they face), we have very low power and the difference will be statistically insignificant. Nevertheless, the peak of the curve is where we predict it would be, subject to all these caveats discussed.

exchange markets. Taking the perspective of the publisher, we consider how reserve prices should be set for auctions when selling display advertising impressions through these ad exchanges to advertisers play a repeated ad buying game and face practical constraints such as budget or impression counts..

In a series of experiments, setting the reserve price under the assumption advertisers play a one-shot game without any constraints is shown to increase publisher's revenue substantially, by 35% (notably, at no additional cost to the publisher). By manipulating the reserve in this fashion, we test the extent to which reach or budget constraints (across multiple auctions) affect advertisers' bidding behaviors. Experimental findings indicate that increasing the reserve price increases advertisers' bid CPMs and total payments, while it does not have an impact on the total number of impressions won by the advertisers. We show these patterns are most consistent with advertisers playing repeated auctions with minimum impression constraint.

Subsequently, we construct an advertiser bidding model that incorporates the minimum impression constraint. The model builds on the notion of a fluid mean-field equilibrium developed in Balseiro et al. 2015, which well approximates the rational behavior of thousands of advertisers competing in repeated auctions with some constraints. We extend this theoretical framework to incorporate the minimum impression constraint, and suggest estimation and identification strategies in applying it to our empirical context. Our counterfactual results show that the reserve price solved without imposing the minimum impression constraint is 27% lower than the optimal reserve level with the minimum impression constraint. Because ignoring the minimum impression constraint biases the solution downward, the profit loss in ignoring this constraint is found to be about 9%. We then conduct an additional set of experiments to validate this pricing policy and find that the prices set most closely to the predicted optimal price yield the highest revenues on average.

While this paper addresses a question of growing economic importance with a novel dataset, a number of additional extensions are possible. In particular, publishers often sell advertising inventory via direct sales as well as RTB. Direct selling involves advance sale of a bundle of impressions directly to the advertiser at a fixed price. Extending our research to consider optimal joint pricing (e.g., fixed price in the direct, reserve price in the ad exchange), and whether inventory should first be made available to one channel or another, are interesting directions for future research.[43] A

[43] We consider reserve prices, taking direct sales decisions made as given (e.g., price, number of impressions bought and sold in direct). Taking the direct channel as given is an assumption that mirrors the structure of the market we consider, where the impressions are sold in the ad exchange after they do not sell in the direct sales channel. Thus, our

second question of interest is motivated by noting that advertiser valuations are incumbent upon the information available to the advertisers about the impression. This raises the question of whether and how much information a publisher should share with an exchange. Another interesting extension would be to extend our analysis to the context of a first price auction (FPA). To explore how the optimal reserve would change under an FPA, one could conduct a counterfactual analysis using a FPA.

To the best of our knowledge, this paper is among the first to empirically consider the issues of pricing in display advertising markets. Given our initial results and the growth in these markets, we hope this and future research will continue to yield economically meaningful implications for publishers in these rapidly growing markets and lead to more research in this area.

solution to the optimal reserve prices can be viewed as the sub-game perfect equilibrium solution to the second-stage, taking the first-stage decisions as given.

References

Angrist, J. D., G. W. Imbens, and D. B. Rubin. 1996. "Identification of causal effects using instrumental variables." *Journal of the American statistical Association* 91 (434): 444–455.

Athey, S., and G. W. Imbens. 2017. "The econometrics of randomized experiments." In *Handbook of economic field experiments,* 1:73–140. Elsevier.

Backus, M., and G. Lewis. 2016. *Dynamic demand estimation in auction markets.* Technical report. National Bureau of Economic Research.

Balseiro, S. R., O. Besbes, and G. Y. Weintraub. 2015. "Repeated auctions with budgets in ad exchanges: Approximations and design." *Management Science* 61 (4): 864–884.

Balseiro, S. R., and Y. Gur. 2019. "Learning in repeated auctions with budgets: Regret minimization and equilibrium." *Management Science.*

Barajas, J., R. Akella, M. Holtan, and A. Flores. 2016. "Experimental designs and estimation for online display advertising attribution in marketplaces." *Marketing Science* 35 (3): 465–483.

Cai, H., K. Ren, W. Zhang, K. Malialis, J. Wang, Y. Yu, and D. Guo. 2017. "Real-time bidding by reinforcement learning in display advertising." In *Proceedings of the 10th international conference on web search and data mining,* 661–670. ACM.

Celis, L. E., G. Lewis, M. Mobius, and H. Nazerzadeh. 2014. "Buy-it-now or take-a-chance: Price discrimination through randomized auctions." *Management Science* 60 (12): 2927–2948.

Choi, H., C. F. Mela, S. R. Balseiro, and A. Leary. 2020. "Online display advertising markets: A literature review and future directions." *Information Systems Research* 31 (2): 556–575.

Coey, D., B. J. Larsen, K. Sweeney, and C. Waisman. 2021. "Scalable optimal online auctions." *Marketing Science* 40 (4): 593–618.

Despotakis, S., R. Ravi, and A. Sayedi. 2021. "First-price auctions in online display advertising." *Journal of Marketing Research* 58 (5): 888–907.

Gallego, G., and G. Van Ryzin. 1994. "Optimal dynamic pricing of inventories with stochastic demand over finite horizons." *Management science* 40 (8): 999–1020.

Ghosh, A., B. Rubinstein, S. Vassilvitskii, and M. Zinkevich. 2009. "Adaptive bidding for display advertising." In *Proceedings of the 18th international conference on world wide web,* 251–260. ACM.

Goldfarb, A., and C. Tucker. 2011. "Privacy regulation and online advertising." *Management science* 57 (1): 57–71.

Han, A. K. 1987. "Non-parametric analysis of a generalized regression model: the maximum rank correlation estimator." *Journal of Econometrics* 35 (2-3): 303–316.

Hendricks, K., and A. Sorensen. 2015. *The role of intermediaries in dynamic auction markets.* Technical report. Working paper, University of Wisconsin.

Hoban, P. R., and R. E. Bucklin. 2015. "Effects of internet display advertising in the purchase funnel: Model-based insights from a randomized field experiment." *Journal of Marketing Research* 52 (3): 375–393.

Iyer, K., R. Johari, and M. Sundararajan. 2014. “Mean field equilibria of dynamic auctions with learning.” *Management Science* 60 (12): 2949–2970.

Johnson, G. A. 2013. “The impact of privacy policy on the auction market for online display advertising.” Working paper.

Johnson, G. A., R. A. Lewis, and E. I. Nubbemeyer. 2016. “The online display ad effectiveness funnel & carryover: A meta-study of predicted ghost ad experiments.” Working paper.

Johnson, G. A., R. A. Lewis, and D. Reiley. 2016. “Location, location, location: Repetition and proximity increase advertising effectiveness.” Working paper.

Krusell, P., and A. A. Smith Jr. 1998. “Income and wealth heterogeneity in the macroeconomy.” *Journal of political Economy* 106 (5): 867–896.

Lee, K.-C., A. Jalali, and A. Dasdan. 2013. “Real time bid optimization with smooth budget delivery in online advertising.” In *Proceedings of the 7th international workshop on data mining for online advertising,* 1. ACM.

Lewis, R. A., and J. M. Rao. 2015. “The unfavorable economics of measuring the returns to advertising.” *The Quarterly Journal of Economics* 130 (4): 1941–1973.

Ostrovsky, M., and M. Schwarz. 2011. “Reserve prices in internet advertising auctions: A field experiment.” In *Proceedings of the 12th conference on electronic commerce,* 59–60. ACM.

Paarsch, H. J., H. Hong, et al. 2006. “An introduction to the structural econometrics of auction data.” *MIT Press Books* 1.

Powell, J. L., J. H. Stock, and T. M. Stoker. 1989. “Semiparametric estimation of index coefficients.” *Econometrica: Journal of the Econometric Society:* 1403–1430.

Rafieian, O., and H. Yoganarasimhan. 2021. “Targeting and privacy in mobile advertising.” *Marketing Science* 40 (2): 193–218.

Riley, J. G., and W. F. Samuelson. 1981. “Optimal auctions.” *The American Economic Review* 71 (3): 381–392.

Sahni, N. S. 2015. “Effect of temporal spacing between advertising exposures: Evidence from online field experiments.” *Quantitative Marketing and Economics* 13 (3): 203–247.

Sayedi, A. 2018. “Real-time bidding in online display advertising.” *Marketing Science* 37 (4): 553–568.

Tucker, C. 2014. “Social networks, personalized advertising, and privacy controls.” *Journal of Marketing Research* 51 (5): 546–562.

Tunuguntla, S., and P. R. Hoban. 2021. “A near-optimal bidding strategy for real-time display advertising auctions.” *Journal of Marketing Research* 58 (1): 1–21.

Weintraub, G. Y., C. L. Benkard, and B. Van Roy. 2008. “Markov perfect industry dynamics with many firms.” *Econometrica* 76 (6): 1375–1411.

Xu, J., K.-c. Lee, W. Li, H. Qi, and Q. Lu. 2015. “Smart pacing for effective online ad campaign optimization.” In *Proceedings of the 21th international conference on knowledge discovery and data mining,* 2217–2226. ACM.

Yuan, S., J. Wang, and X. Zhao. 2013. "Real-time bidding for online advertising: Measurement and analysis." In *Proceedings of the 7th international workshop on data mining for online advertising,* 3. ACM.

Online Appendix

A Model

A.1 Proofs

A.1.1 Proposition 1

The proof for Proposition 1 follows the steps established in Balseiro et al. 2015 Section A.1. First, the dual of the primal problem in Equation (1) is introduced using a Lagrange multiplier for the minimum impression constraint. Second, the first-order conditions are derived to determine the solution for the dual problem.

Step 1: The Lagrangian for type θ is denoted as

$$\begin{aligned}\mathcal{L}_\theta(b,\mu) &= \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{b(V)\geq D\right\}(V-D)\right]+\mu\left[\eta s_\theta E_{V,D}\left[\mathbf{1}\left\{b(V)\geq D\right\}\right]-y_\theta\right] \\ &= \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{b(V)\geq D\right\}(V-D+\mu)\right]-\mu y_\theta\end{aligned} \tag{5}$$

where a Lagrange multiplier for the minimum impression constraint is $\mu \geq 0$. The dual problem (converting from maximizing the advertiser's objective function given its minimum impression constraint in Equation (1) to minimizing the Lagrangian multipliers while maximizing the objective function) is given by

$$\begin{aligned}\Psi_\theta(\mu) =& \inf_{\mu\geq 0}\sup_{b(\cdot)}\mathcal{L}_\theta(b,\mu) \\ =& \inf_{\mu\geq 0}\left\{\eta s_\theta \sup_{b(\cdot)}\left\{E_{V,D}\left[\mathbf{1}\left\{b(V)\geq D\right\}(V+\mu-D)\right]\right\}-\mu y_\theta\right\} \\ =& \inf_{\mu\geq 0}\left\{\eta s_\theta E_{V,D}\left[\mathbf{1}\left\{V\geq D-\mu\right\}(V-(D-\mu))\right]-\mu y_\theta\right\}\end{aligned} \tag{6}$$

the inf is the Lagrangian minimization step and the sup is the goal maximization step. The equality in the last line comes from the fact that the inner optimization problem is similar to an advertiser's problem who faces value $v+\mu$ and seeks to maximize its expected utility in the second-price auction so that bidding truthfully becomes optimal (consider Equation 1 without the constraints). That is for any given multiplier $\mu \geq 0$, the inner expectation term is maximized with the policy $b(V) = V + \mu$. Further, the term within the expectation in the last line is convex in μ, and the expectation preserves convexity, leading to a convex dual problem.

Step 2: The first order condition of $\Psi_\theta(\mu)$ with respect to μ (that is, the FOC for the $\inf_{\mu\geq 0}\{\cdot\}$) is given by

$$(d/d\mu)\Psi_\theta(\mu) = \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{V\geq D-\mu\right\}\right]-y_\theta = 0 \tag{7}$$

To explain the solution to this FOC, we begin by noting that $\Psi_\theta(\mu)$ is convex in μ. If the constraint does not bind (i.e., $\eta s_\theta E_{V,D}\left[\mathbf{1}\left\{V\geq D\right\}\right]\geq y_\theta$), then $(d/d\mu)\Psi_\theta \geq 0$ at $\mu = 0$. This condition implies the function $\Psi_\theta(\mu)$ is increasing in μ for all $\mu \geq 0$, such that the function is minimized at $\mu^* = 0$ (has a corner solution). Intuitively, the Lagrangian multiplier can be interpreted as the cost of the impressions constraint; If the constraint does not bind, the constraint is costless.

On the other hand, when the constraint binds (i.e., the optimal unconstrained number of impressions is less than the minimum impression level, that is $\eta s_\theta E_{V,D}[\mathbf{1}\{V \geq D\}] < y_\theta$), then $(d/d\mu)\Psi_\theta(\mu)$ takes a negative value at $\mu = 0$. As $\mu \to \infty$, $(d/d\mu)\Psi_\theta(\mu)$ converges to a positive value $\eta s_\theta - y_\theta > 0$, because $lim_{\mu\to\infty} E_{V,D}[\mathbf{1}\{V \geq D - \mu\}] = 1$.[44] Thus, there exists a unique interior solution $\mu^* > 0$ for $(d/d\mu)\Psi_\theta(\mu^*) = 0$ as Ψ_θ is a convex function in μ.

Moreover, the complementary slackness conditions hold with the bidding function $b^*(V) = \beta_\theta^F = v + \mu^*$ and the optimal multiplier μ^* such that:

$$\mu^* \left[\eta s_\theta E_{V,D}\left[\mathbf{1}\left\{\beta_\theta^F(V) \geq D\right\}\right] - y_\theta\right] = 0,$$

That is, either i) the minimum impression constraint binds or ii) the Lagrangian multiplier μ^* is 0.

Lastly, there is no duality gap. That is, there is no difference between the primal (1) and dual (6) values. The bid function β_θ^F is primal feasible from the first-order conditions of the dual, and the primal objective value calculated at the proposed bid function ($b(V) = \beta_\theta^F$) is then

$$\begin{aligned}
&\eta s_\theta E_{V,D}\left[\mathbf{1}\left\{\beta_\theta^F \geq D\right\}(V - D)\right] \\
=&\left(\eta s_\theta E_{V,D}\left[\mathbf{1}\left\{\beta_\theta^F \geq D\right\}(V - D)\right] + \mu^*\left[\eta s_\theta E_{V,D}\left[\mathbf{1}\left\{\beta_\theta^F \geq D\right\}\right] - y_\theta\right]\right) - \mu^*\left[\eta s_\theta E_{V,D}\left[\mathbf{1}\left\{\beta_\theta^F \geq D\right\}\right] - y_\theta\right] \\
=&\mathcal{L}_\theta(\beta_\theta^F, \mu^*) - \mu^*\left[\eta s_\theta E_{V,D}\left[\mathbf{1}\left\{\beta_\theta^F(V) \geq D\right\}\right] - y_\theta\right] \\
=&\mathcal{L}_\theta(\beta_\theta^F, \mu^*) \\
=&\Psi_\theta(\mu^*)
\end{aligned}$$

where the third equality follows from the complementary slackness conditions, as $\mu = \mu^*$ is the value for which $\mu^*\left[\eta s_\theta E_{V,D}\left[\mathbf{1}\left\{\beta_\theta^F(V) \geq D\right\}\right] - y_\theta\right] = 0$. Finally, the last equality follows from the fact that $\Psi_\theta(\mu^*) = \sup_{b(\cdot)} \mathcal{L}_\theta(b, \mu^*)$ and the optimal bid function β_θ^F solves the latter problem.

A.1.2 Proposition 2

Once we establish the optimal bidding function and the optimal multiplier as above, Proposition 2 follows from Proposition 4.1 in Balseiro et al. 2015, in which they characterize the equilibrium. Since we analyze the minimum impression constraint without the participation constraint, the existence and the characterization of the equilibrium are valid only when the participation constraints do not bind in equilibrium. Readers are referred to their proof in the supplemental material (http://dx.doi.org/10.1287/mnsc.2014.2022).

A.2 Theoretical Predictions

In this appendix, we outline the intuition for the predictions in Table 3 in Section 3. These predictions indicate how advertiser behaviors change as the reserve prices increase from $r = 0$ (i.e., not reserve price) to a reserve price level $r^*_{nc} > 0$, that is the optimal reserve price under the assumption that advertisers do not face binding constraints.

In the subsequent analysis, we consider several cases: i) when the constraint binds neither at 0 nor at r^*_{nc} (not bind, not bind), ii) when the constraint does not bind at 0 but does bind at

[44] $(\eta s_\theta - y_\theta)$ is a finite positive number as we constrain $\eta s_\theta > y_\theta$ when defining the optimization problem in sub-section 3.3 (i.e., the minimum impression level is lower than the total available impressions).

r^*_{nc} (not bind, bind), and iii) when the constraint binds at both the 0 and r^*_{nc} reserve price levels (bind, bind). We consider two types of constraints: a maximum budget constraint and a minimum impression constraint. The minimum impression constraint and the maximum budget constraint are each considered in isolation. That is when we consider the budget constraint, we assume that the minimum impression constraint does not bind in both ($r = 0$) and ($r^*_{nc} > 0$).

A.2.1 No Binding Impression or Budget Constraint (Not Bind, Not Bind)

When the underlying state is (not bind, not bind), the advertiser bidding model collapses to the standard second-price auction without the constraint. In this case, advertisers bid their true valuations, and the distribution of bids will be invariant regardless of the reserve price level. The probability of winning an impression decreases at higher r^*_{nc}, but the total payment increases as r^*_{nc} maximizes publisher's revenues.

A.2.2 When Impression Constraints Bind ((Not Bind, Bind) or (Bind, Bind))

Impression Constraints: The Effect of the Reserve Price on the Optimal Bid The equilibrium Lagrangian multiplier μ^* increases monotonically with the increase in r (until the participation constraint binds). When the minimum impression constraint binds, the optimal $\mu^* > 0$ satisfies $y_\theta - \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{V + \mu^* \geq D\right\}\right] = 0$ (see Equation 7). An increase in the reserve price will increase D (the steady-state maximum of the competitors' bids), and to offset this effect, μ^* needs to increase as well to satisfy the equality constraint. As the optimal bidding strategy is derived as $\beta^F_\theta = v + \mu^*$, an increase in r increases μ^*, which in turn increases bids for advertisers participating both at $r = 0$ and r^*_{nc}.

Impression Constraints: The Effect of the Reserve Price on Number of Impressions Won and the Total Payment When at least some advertisers face the case of (not bind, bind), the number of impressions won by advertisers will decrease as the reserve prices increase. This is because some advertisers who were able to buy in excess of the minimum impression constraint at $r = 0$ cannot buy as many impressions as they used to at r^*_{nc}.

The direction of the total payment is ambiguous when reserve prices increase in the (not binding, binding) condition because there are opposing effects. Although the number of impressions won decreases, advertisers increase their bids with the increase in the reserve price leading to a higher payment CPM per impression sold.

Finally, when the advertiser faces the (bind, bind) condition, the number of impressions won do not change, as the advertiser buys the minimum number of impressions when the constraint binds.[45] Accordingly, the total payment will increase with the increase in advertisers' bids as the reserve price increases.

[45]The number of impressions can decrease if the participation constraint binds for some advertisers.

A.2.3 When Budget Constraints Bind ((Not Bind, Bind) or (Bind, Bind))

Budget Constraints: Optimal Bidding Strategy Balseiro et al. 2015 establish that the optimal bidding strategy when advertisers face the maximum budget constraint is

$$\beta_{\theta}^{F}\left(v|F_{D}\right)=\frac{v}{1+\mu^{*}} \tag{8}$$

where μ^* is

$$\begin{cases} \mu^{*}=0 & \text{if } b_{\theta}>\eta s_{\theta} E_{V,D}\left[\mathbf{1}\left\{V \geq D\right\} D\right] \\ \eta s_{\theta} E_{V,D}\left[\mathbf{1}\left\{V/(1+\mu^{*}) \geq D\right\} D\right]-b_{\theta}=0 & \text{if } b_{\theta} \leq \eta s_{\theta} E_{V,D}\left[\mathbf{1}\left\{V \geq D\right\} D\right] \end{cases}$$

and b_θ is the maximum (constrained) budget.

Budget Constraints: The Effect of Reserve Price on Optimal Bid When the advertiser faces the case of a (not bind, bind) budget constraint as the reserve price increases from 0 to r_{nc}^*, the bidding strategy will change from v to $\frac{v}{1+\mu^*}$ where $\mu^* > 0$. Thus, the bid will decrease in this case.

When the advertiser faces (bind, bind) constraint as the reserve price increases from 0 to r_{nc}^*, the effect on the bid is ambiguous. This is because $\frac{\partial \mu^*}{\partial r}$ is not monotonic. For example, Figure 9 shows the change in μ^* (y-axis) with respect to the change in r (x-axis), when $v \sim U[0,1]$ and $b = 0.1$, when there is one advertiser. The plot shows that μ^* first increases then decreases in the range $[0, r_{nc}^*] = [0, 0.5]$.

Figure 9: Optimal Bidding Strategy with a Maximum Budget Constraint

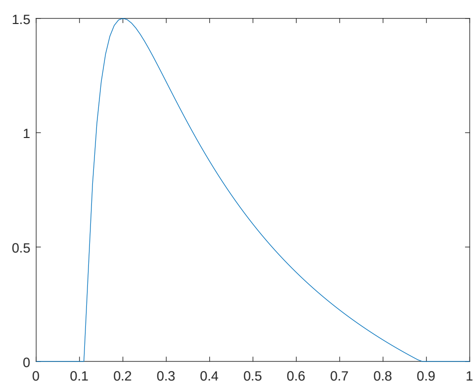

Note: The figure shows the change in μ^* (y-axis) with respect to the change in r (x-axis).

Budget Constraints: The Effect of Reserve Price on Number of Impressions Won and Total Payment When the advertiser faces the case of a (not binding, binding) budget constraint as the reserve price increases from 0 to r_{nc}^*, the number of impressions won will decrease as this advertiser now shades bids and the reserve price increases (see Equation 8). The total payment across advertisers increases, because advertisers spend more and the budget constraint becomes binding.

When the advertiser faces the case of a (binding, binding) budget constraint, the total payment will stay the same (at the binding budget level), but the effect on the number of impressions won is ambiguous because bid CPM may or may not increase with respect to the change in the reserve price.

B Experimental Evidence of a Constraint

This Appendix section contains two components. First, it details the design of the experiment used in Section 4 to explore the effect of reserve prices on advertiser bidding. Second, it summarizes a number of robustness checks pertaining to the analysis of the experimental outcomes.

B.1 Experiment Design

B.1.1 Treatment and Control Groups

The results from sub-section 2.3 show that advertisers' bid CPMs vary by site, device, ad location, and size. Therefore, the experiments were designed to run across different ad types including (i) site (site1 and site2), (i) device (desktop, mobile, tablet), (iii) ad location (above-the-fold/ATF, MID, below-the-fold/BTF, no info, front door), and (iv) size (300x250, 728x90, 970x66, 320x50, 300x600, 970x250). Twelve experiments were constructed; each constituted a pair. The two units in a pair were closest in ad characteristics and were randomized either into the treatment or the control group. For example, (site1, desktop, BTF, 300x250) and (site1, desktop, ATF, 300x250) were paired for the first experiment, and the randomly chosen (site1, desktop, BTF, 300x250) was assigned to the treatment group, whereas (site1, desktop, ATF, 300x250) was assigned to the control group. Table 9 shows the full list of the experiments and the corresponding treatment and control group characteristics.

Table 9: Treatment and Control Groups

Experiment Pair ID	Device	Site and Position		Inventory Size	
		Treatment	Control	Treatment	Control
1	Desktop	Site1 BTF	Site1 ATF	300x250	300x250
2	Desktop	Site2 MID	Site2 ATF, BTF	300x250	300x250
3	Mobile	Site1 No Info	Site2 MID	300x250	300x250
4	Mobile	Site2 ATF	Site2 BTF	300x250	300x250
5	Desktop	Site1 ATF	Site1 BTF	728x90, 970x66	728x90, 970x66
6	Tablet	Site1 ATF, BTF	Site1 No Info	728x90, 970x66	728x90, 970x66
7	Mobile	Site2 MID, BTF	Site1 No Info, Site2 ATF	320x50	320x50
8	Desktop	Site1 ATF	Site2 ATF	300x600	300x600
9	All	Site1 Front Door BTF	Site1 Front Door BTF	300x250	728x90
10	Desktop	Site1 ATF	Site1 BTF	970x250	970x250
11	Tablet	Site1 No Info	Site2 ATF, MID, BTF	300x250	300x250
12	Desktop	Site2 BTF	Site2 ATF	728x90, 970x66	728x90, 970x66

B.1.2 Balance of Observables

Table 10 reports the balance of observables between the treated and control groups in the pre-period. The observable metrics are calculated for each experimental cell (12 treatment cells and 12 control cells) and the p-values are obtained from two-sided paired t-tests of differences in treatment vs. control groups. The observables are not statistically significantly different between the treatment and control groups. There is little power with 12 observations, nonetheless all but one of the p-values exceed 50%.

Table 10: Table of Balance

Variable	P-value for Difference
Revenue	0.96
# Ad Requests	0.91
# Impressions Sold	0.68
# Advertisers	0.92
CPM paid	0.65
Click Through Rate	0.40
Bid CPM	0.56
eCPM	0.91
Sell Through Rate	0.91
#Bids / # Ad Requests	0.92

Note: Variables are calculated at the experimental cell level (12 treatment cells and 12 control cells). P-values are from two-sided paired t-tests of differences in treatment vs. control groups ($N = 12$).

B.1.3 Pre-Trend

The DiD analysis is predicated upon parallel pre-trends. In Figure 10, we plot pre-trends for the metrics used in the DiD analysis; eCPM, average bid CPM, number of ad impressions won, and total payment. Although control and treatment groups have different mean levels for some metrics, in general the trends appear to be parallel.

B.1.4 Calculating Reserve Prices

When inducing exogenous variation in the reserve pricing experiment, we selected price levels predicated on the assumption of no binding constraints in the hope of finding a level that would increase the publisher's revenues (had we set reserve prices far too high, revenues would have fallen compared to the historical zero reserve price levels). That is, the reserve prices set in the experiment were calculated presuming advertisers bid truthfully and that the observed bids reflect their underlying valuations.

Under this setting, the publisher can choose the reserve price r to maximize the publisher's revenue by solving

$$\max_r r\{1 - F_V(r)\} \tag{9}$$

where F_V is the cdf of advertiser valuation distribution (Riley and Samuelson 1981). $\{1 - F_V(r)\}$ is the probability that the ad impression will be sold at the reserve price r.[46]

Let b_i be the average bid CPM and L_i be the number of bids submitted for a given observational unit i (advertiser-DSP-day-site-ad type). We estimate $\{1 - F_V(r)\}$ in Equation (9) by the empirical sample analogue

$$\{1 - \hat{F}_V(r)\} = \frac{\sum_{i=1}^{n} \mathbf{I}(b_i > r)L_i}{\sum_{i=1}^{n} L_i}$$

[46] Under the symmetric independent private value paradigm, the optimal reserve price is independent of the number of bidders (Riley and Samuelson 1981). Considering the case where the publisher faces a single bidder; Equation (9) represents the expected revenue from selling the ad impression at the reserve price r.

Figure 10: Pre-Post Trends

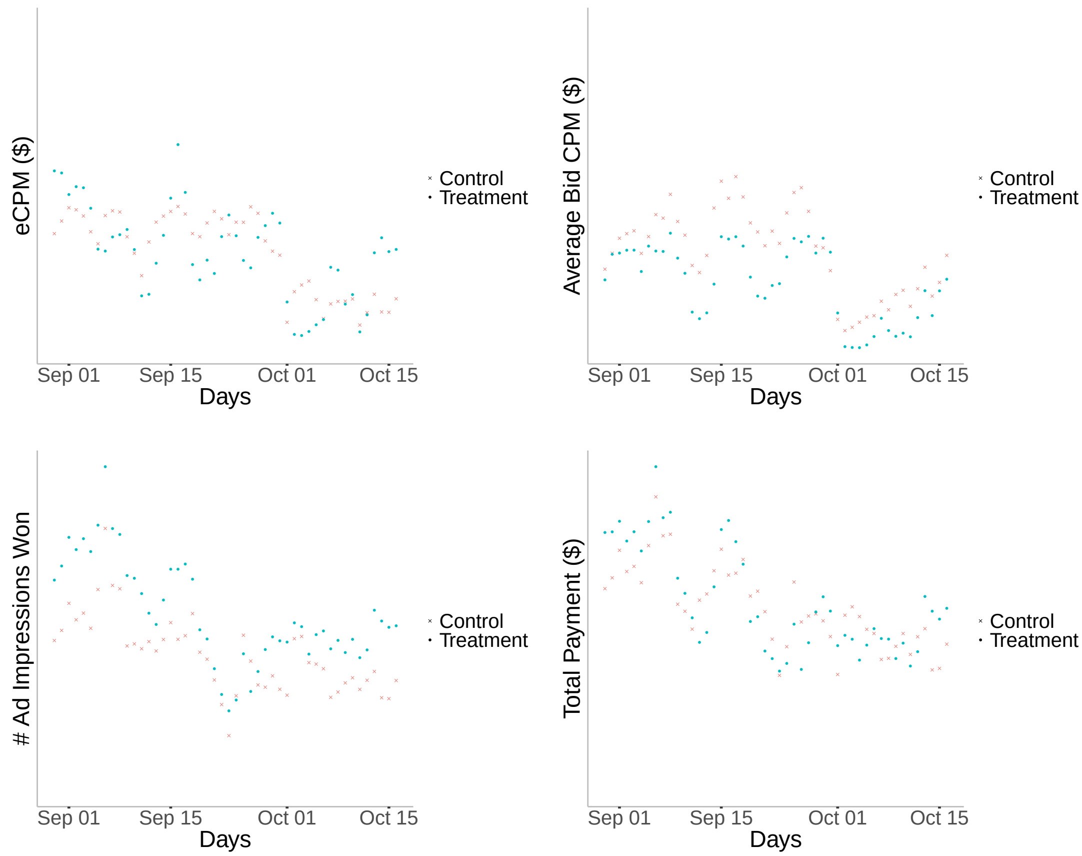

Note: Y-axis levels are not displayed for confidentiality.

using the pre-period data. Equation (9) is optimized with respect to r to find the optimal reserve prices for this unconstrained, naive model. We use these reserve price levels for the treatment group in our experiments (Section 4).

B.2 Experimental Results: Robustness Checks

B.2.1 The Effect of Reserve Prices on eCPM

Next, various DiD specification results are reported with various control variables. The outcome measure considered is *eCPM* (multiplied by a common, multiplicative constant for confidentiality). First, we present the analysis using the experimental cell as the unit of analysis, which has the virtue of transparency. In estimating the DiD regression, we weigh each observation by its ad impression requests (= number of ad impressions supplied to ad exchange) to account for the substantial amount of heterogeneity in experimental cell sizes. The result in Table 11 indicates that the causal increase in eCPM from setting reserve prices is about \$0.12 and not particularly sensitive to specification.

One can conduct similar DiD analyses using more granular-level data to further control for ad characteristics. Table 12 considers data whose unit of observation is (day-site-ad type). The

Table 11: Treatment Effect on eCPM ($): Cluster Level Analyses

	(1)		(2)		(3)	
DV= eCPM ($)	Estimate	SE	Estimate	SE	Estimate	SE
Treated × Post	0.13	(0.16)	0.11	(0.11)	**0.12***	(0.06)
Treated	−0.00	(0.12)	0.00	(0.15)	–	
Post	−0.01	(0.12)	0.00	(0.10)	−0.005	(0.04)
Experiment	–		y		y	
Treated × Experiment	–		–		y	
R-squared	0.03		0.39		0.94	
Observations	48		48		48	

Note: There are 12 treatment and 12 control cells, each with pre- and post- observations, giving us a total of 48 observations for this analysis. *Denotes 10 % significance.

effect of implementing reserve prices on eCPM is again estimated to be positive and significant, \$0.11 $\sim$ \$0.13, yielding 30% $\sim$ 35% increase in revenue across specifications from the baseline eCPM for the treatment group in the pre-period, \$0.37.

Table 12: Treatment Effect on eCPM ($): Unit Level Analyses

DV= eCPM ($)	(1)	(2)	(3)	(4)	(5)
Treated × Post	**0.13****	**0.13****	**0.11****	**0.12*****	**0.11*****
P-Value (Randomization Inference)	0.04	0.04	0.04	0.00	0.01
Treated	y	y	y	–	–
Post	y	–	–	–	–
Day	–	y	y	y	y
Experiment	–	–	y	y	y
Treated × Experiment	–	–	–	y	y
Site-Ad Type	–	–	–	–	y
R-squared	0.02	0.05	0.28	0.64	0.82
Observations	3,442	3,442	3,442	3,442	3,442

Note: There are 5,382,829 observations in the experiment at the (advertiser-DSP-day-site-ad type) level. Aggregation is necessary for the eCPM DiD analyses, because the denominator in the eCPM (i.e., impressions supplied to ad exchange) can only be defined at (day-site-ad type) level. Thus, the data are aggregated across advertisers and DSPs, leaving 3,442 observations. The p-value for testing the null hypothesis that the treatment has no effect is calculated using randomization inference, randomizing treatment assignment at the experimental pair-level.*** denotes 1% and ** 5 % significance.

B.2.2 The Effect of Reserve Prices on Bidding Behaviors

Table 13 reports the effect of reserve prices on advertiser bidding behaviors and its sensitivity to the inclusion of additional control variables. Column (1) is the specification reported in the main body of the paper (Table 5). Overall, the finding that the bid CPM and total payment increase while the number of impressions bought stays the same is qualitatively robust to various model specifications.

B.3 Additional Empirical Considerations

B.3.1 Bid Truncation

The publisher's ad exchange partner reports auction outcomes at the daily level, and available metrics are daily averages for the given observational units (advertiser-DSP-day-site-ad type). When the experimental reserve prices exceed an advertiser's valuation, the advertiser might not submit a bid in the treatment condition (perhaps if they perceive submitting bids to be costly, though in reality bids are automated so there is little reason not to bid). In this instance, all bids will be

Table 13: Treatment Effect on Bidding Behaviors: Robustness Checks

DV= Bid CPM ($)	(1)	(2)	(3)	(4)	(5)
Treated × Post	**0.039****	**0.050****	**0.049****	**0.049****	**0.049****
P-Value (Randomization Inference)	0.037	0.026	0.019	0.033	0.035
# Impressions Supplied (in thousand)	y	y	y	y	y
Treated	y	y	–	–	–
Day	y	y	y	y	y
Advertiser	y	y	y	y	y
Experiment	–	y	y	y	y
Treated × Experiment	–	–	y	y	y
Site-Ad Type	–	–	–	y	–
DSP	–	–	–	–	y
R-squared	0.212	0.223	0.226	0.229	0.230
Observations	3,635,899	3,635,899	3,635,899	3,635,899	3,635,899

DV= # Impressions Won	(1)	(2)	(3)	(4)	(5)
Treated × Post	0.351	0.220	0.321	0.314	0.282
P-Value (Randomization Inference)	0.397	0.437	0.422	0.447	0.436
# Impressions Supplied (in thousand)	y	y	y	y	y
Treated	y	y	–	–	–
Day	y	y	y	y	y
Advertiser	y	y	y	y	y
Experiment	–	y	y	y	y
Treated × Experiment	–	–	y	y	y
Site-Ad Type	–	–	–	y	–
DSP	–	–	–	–	y
R-squared	0.197	0.197	0.198	0.198	0.199
Observations	5,382,829	5,382,829	5,382,829	5,382,829	5,382,829

DV= Total Payment ($)	(1)	(2)	(3)	(4)	(5)
Treated × Post	**0.005****	**0.005****	**0.005****	**0.005****	**0.005****
P-Value (Randomization Inference)	0.048	0.045	0.022	0.020	0.020
# Impressions Supplied (in thousand)	y	y	y	y	y
Treated	y	y	–	–	–
Day	y	y	y	y	y
Advertiser	y	y	y	y	y
Experiment	–	y	y	y	y
Treated × Experiment	–	–	y	y	y
Site-Ad Type	–	–	–	y	–
DSP	–	–	–	–	y
R-squared	0.188	0.189	0.189	0.189	0.194
Observations	5,382,829	5,382,829	5,382,829	5,382,829	5,382,829

Note: The dependent variables are all multiplied by a common, multiplicative constant for confidentiality. The unit of observation for the analysis is (advertiser-DSP-day-site-ad type). Bid CPM analysis uses 3,635,899 bid data observations from the opt-in advertisers. Ad impression and total payment analyses use 5,382,829 payment data observations from all advertisers. The P-value for testing the null hypothesis that the treatment has no effect is calculated using randomization inference (Athey and Imbens 2017). For randomization inference, we randomize treatment assignment at the experimental pair-level, the same way that the assignment was done in the experiment). ***denotes 1% and **5% significance.

accounted toward daily averages for the control group with zero reserve prices, but only the bids exceeding the reserve prices will be observed and used for daily averages for the treatment group in the post-period. The 'average' bid CPMs will appear higher for the treatment group due to truncation, even were the untruncated distribution of bids remains invariant to the manipulation of reserve prices.

We explore this truncation issue using a DiD regression of the number of bids. Were advertisers to cease submitting bids below the reserve prices, one would expect to see a decrease in the number of bids in the post-period for the treatment group. The first two columns in Table 14 show the results from DiD regressions of the number of bids. The model (1) includes advertiser fixed effects to illustrate the within-advertiser change. The model (2), on the other hand, considers the extensive margin. In both cases, we find no significant decrease in the number of advertiser bids; in fact the coefficients are positive. Thus, there is no significant evidence of truncation in our context, presumably because the cost of submitting bids is so small as they are typically submitted by algorithms (and submitting bids below the reserve prices does not decrease advertisers' utility).

Table 14: Treatment Effect on Number of Bids and # Impressions Won

DV (Scaled)	# Bids		# Impressions Won	
	(1)	(2)	(1)	(2)
Treated × Post	2.86	0.74	0.351	0.033
P-value (Randomization Inference)	0.235	0.388	0.397	0.768
# Impressions Supplied (in thousand)	Y	Y	Y	Y
Treated	Y	Y	Y	Y
Day	Y	Y	Y	Y
Advertiser	Y	–	Y	–
R-squared	0.080	0.003	0.197	0.001
Observations	3,635,899	3,635,899	5,382,829	5,382,829

Note: The dependent variables are multiplied by a common, multiplicative constant for confidentiality. Bids analysis uses 3,635,899 bid data observations from the opt-in advertisers. Ad impression analyses use 5,382,829 payment data observations from all advertisers. The P-value for testing the null hypothesis that the treatment has no effect is calculated using randomization inference (Athey and Imbens 2017). For randomization inference, we randomize treatment assignment at the experimental pair-level, the same way that the assignment was done in the experiment). ***denotes 1% and **5% significance.

B.3.2 SUTVA Assumption

In this sub-section, we discuss the stable unit treatment value assumption (SUTVA). The assumption of no spillovers between different units implies that potential outcomes are invariant to random treatment assignment of others (Angrist et al. 1996).

Violation of this assumption may occur if, for example, advertisers strategically substitute the ad inventory (purchase more ad impressions) in the control group for those in the treatment group due to the increase in reserve prices in the treatment group ads. Revisiting Tables 14, we find that the number of submitted bids and the number of ad impressions won stay the same with respect to the increase in reserve prices. The first (last) two columns show the results from DiD regressions of the number of bids (the number of ad impressions won). The model (1) looks at the within-advertiser changes, whereas the model (2) excludes advertiser fixed effects to look at the extensive margin.

As their is no negative interaction between the treatment and post, strategic substitution does not seem to be a major factor in our setting.

B.3.3 Experiments 8 and 10

Recall, for most experiments, the reserve prices for the control group were zero (i.e., no reserve prices were set). However, for experiment pairs 8 and 10, the publisher set surprisingly high reserve prices in the pre-period that were the same for the treatment and control groups. In the post-period, the reserve prices were experimentally decreased to the unconstrained, naive optimal levels for the treatment group and decreased to zero for the control group.

For experiments 8 and 10, one would still expect increases in the eCPM. Moving from extremely high reserve prices to the unconstrained, naive solution levels (experimental treatment setting) should yield a bigger lift in *revenues* than moving to the zero reserve prices (experimental control setting). Considering the last row in Table 3 in which the minimum impression constraint binds, one would again expect increases in *bid CPM* and *total payments* but no change in the *number of impressions* bought for experiments 8 and 10. Moving from extremely high reserve prices to zero would relax the minimum impression constraint to a larger extent so advertisers would decrease their bid CPM to a larger extent, compared to moving to the naive solution level with smaller decrease in their bid CPM. Hence, the difference in difference in bid CPMs should remain positive for experiments 8 and 10. Similarly for the total payments, the number of impressions won would be similar in the (bind, bind) condition, but the treatment group would be bidding higher amounts under higher reserve prices, which raises the total payment (as it is the product of impressions won and the second highest bids).

C Estimation and Policy Simulation

This section overviews the computational algorithms used in estimation and policy simulations.

C.1 Computational Steps in Estimation

The estimation proceeds as follows:

Stage 1 Denoting $w = v + \mu$, $\hat{F}_W(w)$ and $\hat{F}_D(d)$ are non-parametrically estimated as described in Equations (3) and (4).

Stage 2 Using the optimality condition in Proposition 2, $\boldsymbol{\mu}^*$ is solved for the FMFE using the following algorithm.

1. Start with an arbitrary vector of multipliers $\boldsymbol{\mu}$. That is $\mu_\theta^0 = \mu_\theta, \forall \theta \in \Theta$

2. Repeat

   (a) Using the estimates $\left(\hat{F}_D(d),\, \hat{F}_W(w)\right)$ obtained in Stage 1, N_{sim} simulated values are drawn from the estimated distributions $v \sim \hat{F}_V = \hat{F}_W(v+\mu_\theta^i)$ and $d \sim \hat{F}_D(d)$ to construct

$$\hat{h}(\mu_\theta; \boldsymbol{\mu}^i) = \frac{y_\theta}{\eta s_\theta} - \frac{1}{N_{sim}} \sum \left[\mathbf{1} \left\{ v_{sim} + \mu_\theta \geq d_{sim} \right\} \right] \tag{10}$$

This equation follows from the advertiser's optimal bidding strategy condition characterized in Proposition 2. Note that the first term is the advertiser's minimum auction winning rate for a given impression and a campaign as computed in Section 5.3, and the second term is the predicted winning rate conditional on μ_θ.

(b) μ_θ^{i+1} solves

$$\arg\min_{\mu_\theta \geq 0} \hat{h}(\mu_\theta; \boldsymbol{\mu}^i), \quad \forall \theta \in \Theta$$

where minimizing $\hat{h}(\mu_\theta; \boldsymbol{\mu}^i)$ over $\mu_\theta \geq 0$ ensures that the conditions in Proposition 2 are satisfied in equilibrium

(c) Compute the difference $\triangle = \left\| \boldsymbol{\mu}^{i+1} - \boldsymbol{\mu}^i \right\|$ and update $i = i + 1$

3. Until $\triangle < \epsilon$

C.2 Computational Approach to Computing Optimal Reserve Prices

This sub-section discusses the numerical approach used to solve for the optimal reserve price. We discuss the case when the constraint does not bind ($\mu^* = 0$) first, then incorporate the case of the binding constraint ($\mu^* > 0$).

C.2.1 Case1: $\mu^* = 0$

For purposes of the first set of experiments (Waive 2), advertiser valuations are estimated assuming advertisers use truth-telling strategies, and the optimal reserve prices are calculated conditioned on this assumption. Under the standard second-price, sealed-bid auction, where advertisers bid truthfully, their valuation distribution can be identified (sub-section 5.1.1) and estimated (sub-section 5.2.1) from the observed payment data for each given ad characteristics.

The publisher can maximize the revenue from the ad exchange by choosing the reserve price optimally. The optimal reserve price r^* can be expressed as

$$r^* = c + \frac{[1 - F_V(r^*)]}{f_V(r^*)} \tag{11}$$

where c is publisher's valuation (Riley and Samuelson 1981). F_V and f_V are cdf and pdf of advertiser valuation distribution. For each of the twelve experiments, we use the estimated distribution $\hat{F}_{V|Z}$ and $\hat{f}_{V|Z}$ to calculate $r^*_{nc}(Z)$.

C.2.2 Case2: $\mu^* > 0$

In the policy simulation, we need to find the new FMFE under the considered reserve price level. To do so, the advertisers' beliefs on D (which reflects the bids of competing advertisers) need to be updated to ensure that the advertisers' new beliefs are consistent with the bidding profile $\boldsymbol{\mu}(D, r)$ at the new reserve price. Thus solving the optimal reserve price with the minimum impression constraint involves embedding the iterative best-response algorithm. That is, given the recovered F_v,

1. Start with an arbitrary r^j for $j = 0$ (we start with $r^0 = r^*_{nc}$, the optimal reserve when advertisers bid truthfully)

2. Repeat: r-step

   (a) Start with an arbitrary vector of multipliers $\boldsymbol{\mu}$. That is $\mu^0_\theta = \mu_\theta, \forall\theta \in \Theta$

   (b) Repeat: $\boldsymbol{\mu}$-step

      i. Obtain $F_D\left(\cdot|\boldsymbol{\mu^i}\right)$ using Equation (3).

      ii. μ^{i+1}_θ solves

$$\arg\min_{\mu_\theta \geq 0} h(\mu_\theta; \boldsymbol{\mu}^i), \quad \forall\theta \in \Theta$$
$$h(\mu_\theta; \boldsymbol{\mu}^i) = \frac{y_\theta}{\eta s_\theta} - E_{V,D}\left[\mathbf{1}\left\{V + \mu_\theta \geq D\right\}\right]$$

      which finds the bidding strategy profile, $\boldsymbol{\mu}_\theta$, that minimizes the difference between the minimum winning rate the advertiser aims to attain for a given impression for a campaign (as computed in sub-section 5.3) and the predicted winning rate, $E_{V,D}\left[\mathbf{1}\left\{V + \mu_\theta \geq D\right\}\right]$, given $\boldsymbol{\mu}_\theta$.

      iii. Check the advertiser's participation constraint

      For θ with $0 > E_{V,D}\left[\mathbf{1}\left\{V + \mu_\theta \geq D\right\}(V - D)\right]$, update $\mu^{i+1}_\theta = \bar{\mu}_\theta$ where

$$\bar{\mu}_\theta = \text{argmax}_{\mu_\theta \geq 0}\left[0 \leq \eta s_\theta E_{V,D}\left[\mathbf{1}\left\{V + \mu_\theta \geq D\right\}(V - D)\right]\right]$$

      In other words, if the expected utility to the advertiser is less than 0 (the first equation), the bids are reduced to the point where the advertiser surplus becomes zero (the second line).

      iv. Compute the difference $\triangle = \left\|\boldsymbol{\mu}^{i+1} - \boldsymbol{\mu}^i\right\|$ and update $i = i + 1$

   (c) Until $\triangle < \epsilon$

   (d) Optimize Equation (2) to compute the new reserve price r^{j+1}

3. Until the global maximum is found

D Results

D.1 The Valuation Distribution, F_V

The cumulative density function of the advertiser valuations, F_v, is plotted in Figure 11. The cdf of advertiser valuations (F_V, blue line) is recovered from the observed payments (F_D, red line). The shape of the distributions vary substantially. In other words, the valuation distributions appear to vary by observables such as (site-ad location-ad size-device-month), implying that different reserve prices should be set for different auctions.

Figure 11: Advertiser Valuation Distributions

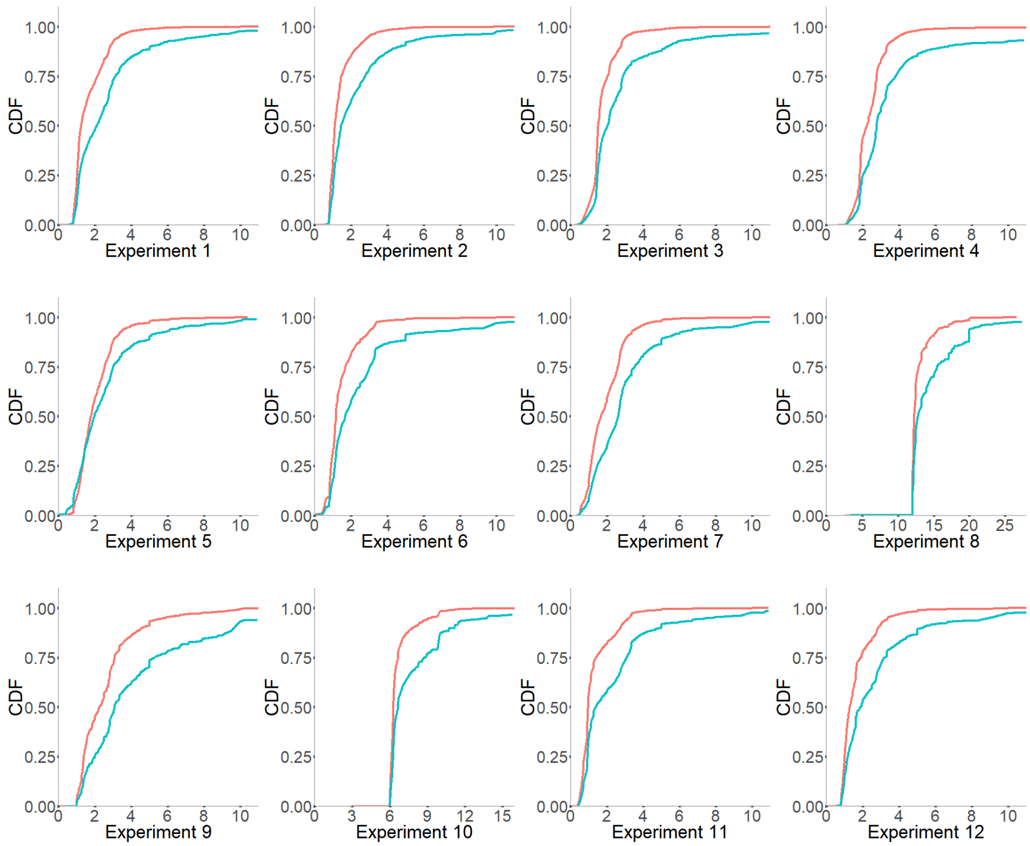

Note: The blue line is the recovered advertiser valuations distribution (F_V) and the red line is the observed payments (F_D).

D.2 Forecast Increase in Experimental Revenues Using Optimal Reserve

Table 15 details the predicted reserve price bias and the profit loss arising from ignoring the minimum impression constraint across all the experimental cells in Section 4. The reserve price bias is calculated as the (optimal reserve without the constraint - optimal reserve with the constraint)÷(optimal reserve with the constraint). The profit loss is calculated as the difference in profits when setting the optimal reserve with the constraint and without. On average, the optimal reserve price calculated with no constraint is about 27% lower than the optimal reserve calculated with the minimum impression constraint, but can be as high as 59%. Further, the profit loss in ignoring this constraint is calculated to be around 9% across experiments, but can be as high as 29%. From this, we conclude that the minimum impression constraint can lead to even more substantial gains in publisher revenues than a naive approach that assumes no constraint.

Table 15: Policy Simulation Results

Experiment	Reserve Price Bias $\left(\frac{r^*_{nc}-r^*}{r^*}\right)$	Profit Loss
1	-18.2	4.55
2	-59.3	29.3
3	-48.1	12.7
4	-28.6	5.90
5	-33.2	14.3
6	-5.71	5.50
7	-7.41	1.04
8	0.0	0.0
9	-40.0	17.7
10	-5.56	0.09
11	0.0	0.0
12	-4.98	4.98
(Weighted) Average	-27.2	9.09