

The Impact of Privacy Policy on the Auction Market for Online Display Advertising

Garrett Johnson*

May 8, 2013

Abstract

The advent of online advertising has simultaneously created unprecedented opportunities for advertisers to target consumers and prompted privacy concerns among consumers and regulators. This paper estimates the financial impact of privacy policies on the online display ad industry by applying an empirical model to a proprietary auction dataset. Two challenges complicate the analysis. First, while the advertisers are assumed to publicly observe tracking profiles, the econometrician does not see this data. My model overcomes this challenge by disentangling the unobserved premium paid for certain users from the observed bids. In order to simulate a market in which advertisers can no longer track users, I set the unobserved bid premium's variance to zero. Second, the data provider uses a novel auction mechanism in which first-price bidders and second-price bidders operate concurrently. I develop new techniques to analyze these hybrid auctions. I consider three privacy policies that vary by the degree of user choice. My results suggest that online publisher revenues drop by 3.9% under an opt-out policy, 34.6% under an opt-in policy, and 38.5% under a tracking ban. Total advertiser surplus drops by 4.6%, 40.9%, and 45.5% respectively.

*I thank the chair of my dissertation committee Robert Porter for his assistance. I also thank the other committee members Eric Anderson, Benjamin Jones, and Elie Tamer. I thank Ivan Canay, Nicholas Della Penna, Jakub Kastl, and Christopher Ody for helpful discussions. Thanks as well to Mark Chicu, Arlene Chu, Simone Galperti, Megan Greenfield, Michael Powell, David Reiley, Katherine Rush, and Kim Singletary for additional assistance. All errors remain my own.

1 Introduction

Online advertisers use personal browsing data to improve ad effectiveness, but these methods have also spurred calls for industry-specific privacy regulation. Online display ads, or banner ads, are prominently displayed on webpages. Revenues from banner ads have increased from \$1.7 billion in 2002 to \$6.8 billion in 2011 (Interactive Advertising Bureau, 2012) in part due to the harnessing of tracking technology. Tracking technology allows advertisers to target specific consumers by mining data from their browsing histories. In response to privacy concerns, American regulators want to rein in tracking practices but are unsure how various policies would impact the online ad industry’s performance. This paper examines the effects of privacy regulation on the online ad industry by analyzing its auction market. Using a large proprietary dataset, I simulate a market governed by several different counterfactual policies, each allowing a different proportion of users to be tracked. My results suggest that all the policies would more than halve publisher and advertiser surplus among the segment of users who can not be tracked.

The online ad auctions operate as follows. The auctions sell individual ad *impressions*, which are defined as one view of the ad by one user on one webpage. A *user* is an individual who browses the Internet. In 2011, American users viewed an incredible 4.9 trillion online display ads (comScore, 2011). Websites are referred to as *publishers* and sell advertising space to *advertisers*. When a user creates an impression opportunity by visiting a webpage, the publisher chooses either to fulfill an advertiser’s standing bulk purchase (called a guaranteed contract) or sell the impression on the spot market. The spot market is an auction operated by an *ad exchange*. When queried, the ad exchange passes along the impression’s ad dimensions, publisher, and user information to potential advertisers. Once they determine how much the impression is worth to them, they calculate their bid. The exchange then determines who wins the right to display their ad. The entire process occurs in milliseconds. Note that advertisers must pay for each delivered impression regardless of whether the user clicks on, or even views, the online display ad.

The use of tracking differentiates online display ads from traditional advertising. Tracking improves the match between users and advertisers. It exploits technologies like Internet cookies, which are small text files stored on a user’s computer that allow publishers and advertisers to identify the user. Tracking enables an advertiser to remember users who visit its websites, and retarget them with advertising on external websites. Third-party companies called data collectors track users across a vast publisher network. They use data collected from browsing histories—combined with information from online and offline databases—to profile user interests, demographics,

and location. Data collectors then sell bundles of users to advertisers under headings such as: ‘Male ages 35–40’; ‘College educated’; ‘Earns \$60–80,000’; ‘Interested in gardening and Nascar’; ‘In market for a flight from Chicago to Tel Aviv’; ‘Married with children’; ‘Republican’; and ‘Likely to move in three months’. When advertisers use these profiles to target users, advertisers engage in what is called *behavioral targeting*. Auctions are useful in this industry because advertisers have asymmetric information regarding the impression’s worth, so auctions facilitate both allocative efficiency and price discovery.

Privacy advocates object to tracking for four main reasons. First, surveys find that at least two thirds of the American public oppose behaviorally targeted advertising (Turow et al., 2009; McDonald and Cranor, 2010; Morales, 2010; Pubmatic, 2011). Second, data collectors track potentially vulnerable groups like children, the sick, and the overweight (Center for Digital Democracy, 2012; Angwin, 2010). Third, tracking can be used to price discriminate or offer discriminate (‘web-lining’), which may make some consumers worse off (Mattioli, 2012; Angwin, 2010). Fourth, strong privacy laws regarding traditional media spotlight the dearth of privacy regulation in the digital media. For instance, the government needs a warrant to look at the books a person borrows from the library, but private companies can track everything a user does online (American Civil Liberties Union, 2011).

American regulators wish to respond to these privacy concerns without undermining the profitability of the ad-supported Internet sector. Six privacy bills are currently before Congress (Interactive Advertising Bureau, 2011), and the Obama administration is also pushing for legislation (White House, 2012). Industry groups like the Internet Advertising Bureau (2011) oppose privacy regulation and argue that it would unduly burden the industry. Meanwhile, privacy advocates like the American Civil Liberties Union (2011) claim the impact of regulation on the revenues of at least the publishers would be negligible. As Levin and Milgrom (2010) demonstrate, the narrow ad targeting in these markets reduces competition. Since a tracking ban could thicken ad markets, the net impact of privacy policy on publisher revenue is an open question that I address in this paper. The effect of regulation on advertiser surplus depends on whether the losses from reduced targeting exceed the benefits from reduced competition. Due to the lack of good estimates of the impact of privacy regulation on the online ad industry, regulators have been debating policy proposals in an empirical vacuum.

In this paper, I measure the welfare impact of privacy policy on advertisers and publishers in the auction market for online display ads. I consider three regulatory options: 1) an *opt-out* policy where users can be tracked by default; 2) an *opt-in*

policy where users cannot be tracked by default; and 3) a *tracking ban*. In Section 2, I predict that these policies will vary the portion of users to be considered Do Not Track (DNT) by about 10%, 90%, and 100% respectively. The ad auction market facilitates tracking by enabling advertisers to locate their target users through a single broad sweep across publishers. I measure the consequences of each of these policies for publisher revenues and advertiser surplus. Section 2 further discusses in more detail the market institutions, tracking methods, and their associated privacy issues, as well as the above policy alternatives.

In Section 3, I describe my proprietary ad exchange dataset and its two types of bidders. The data contain several hundred million online display ad auctions that include thousands of publishers and hundreds of advertisers. The data feature two types of bidders that participate in the auction differently and face separate auction rules. *Offline bidders* pre-specify a single bid amount that they wish to submit randomly for impressions that meet their criteria. In contrast, *real-time bidders* evaluate and bid on impressions as they arise in real time. These bidders employ computer algorithms that both evaluate an impression’s user and publisher characteristics, and anticipate competition. Both types of bidders can track users. In the hybrid auction, the offline bidders play by *second-price* rules and the real-time bidders play by *first-price* rules. The presence of two bidder types complicate the auction analysis.

In Section 4, I describe equilibrium bidding in a hybrid auction, which includes first- and second-price bidders. In a conditionally independent private value setting, a weakly dominant strategy for a second-price bidder is to bid his valuation. First-price bidders shade their bids below their valuations, as they trade off higher surplus—conditional on winning—with a lower chance of winning. An offline (second-price) bidder bids B with positive probability, so the real-time (first-price) bidder’s probability of winning discontinuously improves by bidding just above B . Consequently, first-price bidders should not bid in a dominated interval below B .

Section 5 models bidder behaviour and outlines the model’s identification and estimation approaches. In the model, bidders publicly observe a complete tracking profile of user characteristics. Conditional on these user characteristics, bidders have independent private values. However, I—as the econometrician—do not observe the user tracking characteristics. In auction terminology, the auctions are said to have unobserved heterogeneity at the user level. Bidders know more about the users than the econometrician, so their bids will show dependence within users. Offline bidders employ user tracking characteristics to select their target audience. I estimate this target audience’s share of the population using a mixture model. Without tracking, an offline bidder scales down his bid to reflect the average value of an untargeted

user. Real-time bidders collapse user tracking characteristics into a common, one-dimensional quality ranking that represents user-specific unobserved heterogeneity. A real-time bidder’s value for an ad impression is the product of this unobserved heterogeneity term and a private, idiosyncratic term. I disentangle the two distributions when bids below the seller’s reserve price are not observed. The between-user bid variation identifies the distribution of the user-specific unobserved heterogeneity. The within-user bid variation identifies the distribution of idiosyncratic tastes. I simulate the counterfactual policy market by turning off the variation in the user-specific component, while allowing the idiosyncratic component to vary.

In Section 6, I present my offline and real-time bidder estimates as well as my policy counterfactual estimates. I estimate the model on the subset of auctions that pertain to American users visiting the top three revenue-generating websites. Collectively, these represent over 100 million auctions and half the revenue from American users in the data. My results indicate that opt-out, opt-in, and tracking ban policies would respectively reduce publisher revenues by 3.9%, 34.6%, and 38.5%. The policies would lower advertiser surplus by 4.6%, 40.9%, and 45.5%.

Section 7 proposes several extensions of the model. The first extension proposes a way to account for self-selection by opt-out users. The other extensions discuss market adjustments that mitigate the impact of each of the privacy policies. Section 8 concludes the paper with a summary of findings.

1.1 Literature Review

This study is among the first to discuss tracking and user privacy in the online display ad industry. Goldfarb and Tucker (2011) use the introduction of Europe’s opt-in tracking policy as a natural experiment and find that online display ads are 65% less effective in terms of survey metrics as a result of the policy. Tucker (2012) discusses the economic perspective on privacy and online advertising. Tucker (2011) shows through an empirical example that users can be more receptive to behavioral targeting when they can modify a publisher’s privacy settings. Bailey and Farahat (2012)’s experimental study confirms that behavioral targeting improves ad effectiveness; they find a median improvement of 40% in brand searches and 65% in click through rates. This literature does not grapple with the welfare impact of privacy regulation on the Internet advertising industry.

Several papers in the economic literature discuss the online display ad industry but without addressing policy concerns. Evans (2009) provides an overview of the market. Levin and Milgrom (2010) discuss the market design trade-off between fine-grained ad targeting and thick markets. Abraham et al. (2011) examine the adverse

selection problem that arises when computer programs called bots trigger a worthless ad impression, and discuss the theoretical implications for a second-price auction. Mahdian et al. (2012) investigate a publisher’s incentives to share user cookies with advertisers and how these incentives change with the publisher market structure. Celis et al. (2012) propose a novel auction mechanism that increases revenues, and then simulate the mechanism’s performance using ad auction data. They assume a symmetric independent value setup, but do not allow for user-level correlation in bids because of data limitations.

I examine the online display advertising industry through the lens of empirical auction theory. Hendricks and Porter (2007), Athey and Haile (2007), and Paarsch and Hong (2006) provide a comprehensive introduction to the literature. In my analysis of real-time bidders, my identification strategy follows the support variation arguments outlined in D’Haultfoeuille and Février (2010) and my estimation approach resembles Brendstrup and Paarsch (2006)’s sieve maximum likelihood estimator. I model user tracking as a user-specific unobserved heterogeneity term in bidder’s valuations. The methods previously developed by Krasnokutskaya (2011) and Hu et al. (2009) to identify and estimate unobserved heterogeneity do not apply here because they preclude a binding reserve price. I extend these models to accommodate a binding reserve price when the unobserved heterogeneity is at the user level and we observe a long panel of users.

2 Background

2.1 Online Display Advertising Market Overview

Figure 1 illustrates the online display advertising market from the point of view of the seller, the publisher. The supply side is characterized by both high concentration and a ‘long tail’ containing millions of small publishers. The top 50 publishers account for 90% of online display ad revenues (IAB 2012). Facebook alone sold a quarter of all display ads in the U.S. in 2011, whereas Yahoo! sold 11%, and Microsoft sold 4.5% (comScore 2012). As Figure 1 demonstrates, the publisher must choose an advertiser to fill an ad impression within about 0.1 seconds of a user’s visit. The publisher must choose between satisfying its existing guaranteed contract advertisers or selling the impression on the exchange. A guaranteed contract with an advertiser is a bulk ad purchase that specifies the price and quantity, as well as the time frame and targeting criteria. Though publishers prefer the assured revenue of guaranteed contracts, publishers sell the remaining inventory—called remnant inventory—on a spot

market for two reasons. Either publishers cannot locate enough guaranteed contract buyers or publishers cushion their guaranteed contracts, which punish under-delivery, in the face of uncertain supply. Ghosh et al. (2009) discuss the publisher’s choice between selling to guaranteed contracts or on the exchange and the authors worry that advertisers on the exchange will ‘cherry-pick’ the valuable users. Nonetheless, ad exchanges counsel publishers to ignore this issue and let the guaranteed contracts determine their ad inventory’s reserve price. McAfee et al. (2010) describe how a publisher should choose among its guaranteed contract advertisers.

On the demand side, advertisers may enter on one or both sides of the market depending on their campaign goals. The demand side is more diffuse than the supply side. The top five advertisers in 2011 were AT&T, Experian (personal finance services), Verizon, Scottrade, and Google (comScore, 2012). Performance advertisers like direct marketers prefer the auction market since it allows them to balance the costs and benefits of their ads. Brand advertisers prefer guaranteed contracts to limit their brand’s exposure to risks like offensive content. Some examples of guaranteed contracts include: “200,000 impressions to US users on the New York Time’s finance related pages in July”, “all impressions on the Yahoo! homepage on Sept 21”, “300,000 impressions on AOL to a retailer’s existing customers in April.” Large publishers can accommodate guaranteed contracts with retargeting and behavioural targeting, however the contracting costs make this impractical for smaller publishers. The auction market enables large-scale retargeting and behavioural targeting across multiple publishers. Within the auction market, most advertisers currently use real time bidding, though some offline bidders persist due either to technological limitations or lower costs.

The market features about ten ad exchanges and the largest are Yahoo!’s RightMedia and Google’s DoubleClick. RightMedia hosts over 9 billion auctions daily (Ghosh et al., 2009). Most exchanges use a second-price auction rule¹, some use first-price auction rules, while others use a hybrid of the two rules. The exchanges typically charge the sellers a flat fee or a commission on the sale. Exchanges sold 13% of online display ads in 2011 and the figure is expected to grow (Angwin, 2012). Exchanges owe their growth in part to tracking since exchanges allow advertisers to locate narrowly targeted users across multiple publishers at larger scale. As such, tracking changed the exchange market from low-value remnant impressions to more

¹DoubleClick’s ‘optional’ second-price rule allows bidders representing multiple advertisers to submit their top two bids: bidders may be contractually obligated to do so by the advertisers they represent since bidders have no incentive to reveal their second highest internal bid truthfully (Mansour et al., 2012). The Microsoft Ad Exchange employs a standard second-price auction (Celis et al., 2012).

targeted and valuable impressions.

2.2 Tracking & User Privacy

Tracking is now widespread online. As early as 2000, advertisers used tracking to improve ad targeting (Federal Trade Commission, 2000). In 2011, tracking became so pervasive that the top 100 American publishers installed an average of 49 tracking cookies per visit (Hoofnagle et al., 2012). Specialized firms called data collectors² monitor users' visits on a vast network of publishers using a variety of tracking technologies. Tracking technologies extend beyond cookies to include some that explicitly counteract user efforts to avoid tracking (Hoofnagle et al., 2012). Data collectors typically profile users in terms of their age, gender, zip code, estimated income, marital status, home ownership, interests, past purchases, purchase intent, etc. Data collectors bundle users and sell them to advertisers, often for a few pennies per user (Angwin, 2010). Data collectors integrate with ad exchanges to facilitate these sales. Beales (2010)'s survey of advertising networks suggest that tracking commands a price premium: average CPMs were \$1.98 for untargeted "run of network" ads, \$4.12 for BT, and \$3.07 for retargeting.

Privacy advocates object to tracking for many reasons. As mentioned earlier, these reasons include: 1) two thirds of users say they object to tracking; 2) advertisers track children and behaviorally target users based on sensitive information like their weight and medical conditions; and 3) tracking information be used for price or offer discrimination. Tracking is unavoidable: existing privacy tools are incomplete, are easily circumvented, and companies are neither legally required to respect the user's preferences nor to offer an opt out (Privacy Rights Clearinghouse, 2010). The Federal Trade Commission only intervenes on behalf of users when a companies violate its own privacy statement (Federal Trade Commission, 2012b). Despite industry claims that tracking is anonymous, data collectors routinely use personally identifiable information to connect users to offline data like their terrestrial address, public records, and past sales (American Civil Liberties Union, 2011; Steel, 2010). Even when tracking is anonymous, browsing histories are often unique and user data can be deanonymized in cases (Olejnik et al., 2012; Ohm, 2010). Finally, several studies show that users struggle to understand current browser privacy controls and privacy policies (Leon et al., 2011; McDonald and Cranor, 2009; McDonald et al., 2009; Turow et al., 2007).

The internet advertising industry acknowledges privacy concerns but argues that

²Data collectors are also referred to as data providers, data brokers, or trackers.

it can regulate itself. Recently, the Digital Advertising Alliance (2012b) advanced a voluntary ‘Ad Choices’ program that includes a small icon within participating ‘enhanced notice’ ads. The ads link to a website that both explains ad tracking and allows users to opt out of behaviorally targeted ads. The Alliance delivers 900 billion enhanced noticed impressions monthly and 90% of the online behavioural advertising market has committed to its principles (Federal Trade Commission, 2012b). Unfortunately, this solution suffers from the technological conundrum of opt-outs: labeling users as ‘opt-out’ requires the same technologies used to track a user that the user may wish to disable. Internet browsers solve the technological problem of opt-outs by creating a Do Not Track (DNT) browsing option that communicates the user’s preferences to publishers. However, advertisers object to Internet Explorer’s plan to make DNT its default setting (Digital Advertising Alliance, 2012a) and advertisers circumvented Safari’s setting which restricts cookies by default (Federal Trade Commission, 2012a). A survey of data collectors revealed that 52% either did not offer an opt-out option or did not participate in the industry’s self-regulatory programs; moreover, only a single data collector offered to comply with browser-level DNT (TechJournal, 2012). Though the industry progressed on privacy issues, tracking is profitable and some firms may free-ride on these efforts.

2.3 Privacy Policy

America’s legislators, regulators, and executive are all concerned about online tracking and are moving to regulate the industry. Currently, six separate bills in the House of Representatives and Senate would restrict tracking (see Interactive Advertising Bureau 2011 for an overview). The Federal Trade Commission (2012b) and the Obama Administration (White House, 2012) provide a regulatory blueprint based on the “Fair Information Practice Principles” of Notice, Access, Security, and Choice. Within the online display ad industry, these principles mean that the industry must clearly communicate to the users how their data is used (Notice), allow users to view—and possibly correct—their tracking profiles (Access), safeguard user data (Security), and give users the choice not to be tracked (Choice). The first three principles could impose significant compliance costs on the industry (Hahn, 2001); I focus on the choice aspect of the policy as it would restrict the tracking that differentiates online display ads from traditional advertising.

American regulators wish to mitigate the adverse impact to the Internet industry, but they lack impact estimates. The FTC (2012b) and White House (2012) proposals contain no such estimates. One contribution of this paper is to measure the impact of privacy policy on publisher revenues and advertiser surplus.

I consider three policy alternatives that enable user choice:

1. Opt-out: the industry can track users by default unless they signal that they do not want to be tracked;
2. Opt-in: the industry cannot track users by default unless it receives the users' explicit consent; and
3. Tracking Ban: the industry cannot track a user under any circumstance.

The Obama administration, the FTC, and most legislative proposals favor an opt-out policy. The European Union implemented an opt-in policy for online user data in 2002 Goldfarb and Tucker (2011), and explicitly clarified in 2010 that tracking requires opt-in permission (European Data Protection Authorities, 2010). An outright ban is more of a curiosity.

Although predictions vary, I predict the percentage of Do Not Track (DNT) users as 10% in the opt-out case and 90% in the opt-in case. These polar cases assume strong status quo bias and also reflect the estimated 39% of users who are unaware of tracking (Morales, 2010). Few users use existing privacy tools. Mozilla's Firefox reports that only 8.6% of its desktop users enable its DNT setting (Fowler, 2012); more broadly, one search engine reported that less than 3% of its searches arose from users with DNT enabled (Lindahl, 2012). The Digital Advertising Alliance serves more than 1 trillion enhanced notice ads monthly and more than one million users have opted out since the Ad Choice program began in 2010 (Digital Advertising Alliance, 2012a). Other predictors of DNT opt-out rates are higher. About 12% of users blocked third-party cookies in 2005 (Hoofnagle et al., 2012); an estimated 20% of users claim to value privacy as much as the Internet itself (Deighton and Quelch, 2009); and one online survey found that as high as 90% of users declared they intended to activate their browser's DNT setting (Bachman, 2012). I forecast that 90% of users would be DNT in an opt-out policy again because of status quo bias. This policy is difficult to predict because it would require the industry to adjust significantly. Publishers could mitigate this collapse if they insisted on tracking users as a pre-condition for granting them access. Consequently, an opt-in policy could have important competitive effects if smaller publishers are less able to install barriers to access.³ Furthermore, tracking is a tough sell to users if publishers cannot bury permission in a lengthy legal contract.

³Thus, an opt-in policy in particular would asymmetrically impact small publishers with little power to negotiate tracking for content and too few users to offer first party behavioural targeting. On the other hand, small publisher would be able to free ride if other publishers can convince users to opt-in to all sites rather than just their specific site. Along these lines, FTC Commissioner Rosch warns in his dissenting FTC statement that privacy policy could have unintended anti-competitive effects (Federal Trade Commission, 2012b).

	Offline Bidder	Real Time Bidder
Bidding Method	Pre-specified bids (offline)	Real time bids (online)
Observed Bids	Fixed bid, randomly submitted	Variable over time
Tracking	Crude or none	Detailed, or none
Prevalence	~90%	~10%
Rules	second-price	first-price

Table 1: A summary of the both bidder types

3 Data

I analyze a large proprietary dataset from an ad exchange that auctions online display advertising. To protect the identity of the data provider, I omit certain details like the exchange’s sales and customers. The dataset features several hundred million observations sampled from a single week in 2010. Each observation represents an auction arising from a single ad impression that a publisher sent to the ad exchange.

Each auction observation includes information about the ad, the seller, the bidders, and the user. The data contain the ad’s dimensions and its webpage URL. On the seller side, the data record the publisher’s identity, content category, and reserve price. On the bidder side, the data contain the top two bids above the reserve price and the corresponding bidder identities, but not any information about other bids. The data also include a user identifier and report the user’s browser and location (e.g. country and state). I do not observe any user tracking information.

Thousands of publishers participate on the exchange. The top ten publishers in impression terms represent 46% of the data. Hundreds of different unique ad creatives compete for ad inventory, and dozens of bidder intermediaries—each representing multiple advertisers—bid on their behalf. The two thirds of users are American. Millions of users who appear, with a mean of 7.5 impressions. Users who appear more than 1,000 times make up 4% of the data.

The publishers’ reserve prices are set high, and only a fifth of impressions are sold (i.e. receive one or more bids), and only a tenth receive two or more bids. The high reserve prices reflect the publishers’ outside option of delivering the impression to their guaranteed contract advertisers. On average, the sale price of the impressions appear to be low relative to the market.

The data feature two types of bidders that I analyze separately: offline bidders and real-time bidders. *Offline bidders* pre-specify rules that include their bid, target auction criteria, and a probability of submitting the bid.⁴ *Real-time bidders* evaluate and bid on auction impressions in real time. Table 1 summarizes the differences

⁴For this reason, offline bidders are also called rules-based bidders.

between the bidder types. An offline bidder’s bid criteria specifies the eligible publishers, content categories, and user characteristics that the bidder targets. For instance, General Motors may run a campaign that targets American users on automotive or sports-related websites for which it will bid \$1 with 40% probability (and \$0 with 60% probability). In practice, advertisers specify their bid, ad budget, and time frame; the ad system then calculates the bid probability that spreads the purchases over that time frame. Real-time bidders, on the other hand, employ computer algorithms that return bids within 0.1 seconds. These algorithms incorporate an ad’s past performance (e.g. clicks) on different publishers and user types, as well as the bidder’s competition. Both real-time and offline bidders may or may not employ user tracking information. Real-time bidders can use fine-grained tracking information since their bidding algorithms integrate that information; in contrast, offline bidders coarsely integrate tracking by manually refining their target audience in their bid criteria. Real-time bidders account for 10% of observed winning bids. Real-time bidding rates have grown since my sample period.⁵

Offline and real-time bidders’ bids exhibit distinct patterns. The scatterplots in Figures 2–4 illustrate the evolution of the different bidders’ observed high bid over the sample period. The real-time bidder depicted in Figure 2 demonstrates enormous variation in its bids at all points in time. In contrast, the offline bidder in Figure 3 submits a single bid throughout the entire week. In theory, the offline bidder can modify his bid and criteria at any point in time; however, such adjustments are costly. Only a quarter of offline bidders alter their bid in the sample, as is the case in Figure 4. Two thirds of all the observed highest bids are from offline bidders who left their bid unchanged.

The ad exchange employs a novel auction mechanism with hybrid rules: the offline bidders play by *second-price* rules and the real-time bidders play by *first-price* rules. The bidder with the highest bid wins the auction, regardless of her type, provided that her bid exceeds the reserve price. If an offline bidder wins, then he pays the second-highest bid, no matter of whether an offline or a real-time bidder submits the second highest bid. He pays the reserve price when he is the sole bidder. If a real-time bidder wins, then she simply pays her bid. If bidders are symmetric, the hybrid auction rules are suboptimal as they treat bidders asymmetrically. However, the bidders in the sample are heterogeneous. The rationale for the hybrid mechanism is unclear.

⁵ At DoubleClick, real time bidding rates grew from 8% to 64% within 2011 and reached 68% by May 2012 (Google 2011).

It may be designed to offset the information advantage of the real-time bidders. In another research project, I examine the revenue properties of the hybrid mechanism and in comparison it with a symmetric first-price and second-price mechanisms.

4 Theory

In this section, I develop a theory of equilibrium bidding in a hybrid auction in which some bidders play by first-price rules and others play by second-price rules. The goal is to map the observed bid data into the bidder valuations. This mapping is a first step towards disentangling the role of tracking in the next section and ultimately towards simulating a counterfactual market under privacy policy. I establish an equilibrium bidding bid function $\beta(v)$ that maps valuations v into bids b when a seller sets a reserve price r . The inverse bid function $\eta(b)$ then translates bids into the bidder valuations—the object of interest.

I begin by briefly summarizing the setting and results. I study hybrid auctions where offline bidders randomly submit their bid and play by second-price rules, and real-time bidders bid dynamically and play by first-price rules. I assume that bidders treat each auction separately and have independent private values. Under these assumptions, I show that a second-price bidder optimally bids his valuation and that a first-price bidder optimally bids less than her valuation; in other words, the logic of the pure second-price and pure first-price auctions respectively holds in this hybrid setting. Discontinuous first-price bid functions distinguish these auctions. When a first-price bidder enters the market she trades off her surplus conditional on winning the auction with her probability of winning given her bid. The probability of winning is given by the distribution of her competitors' highest bid. Figure 5 illustrates that the competitors' bid density is typically characterized by multiple mass points or positive probability bids repeatedly submitted by the offline bidders. The first-price bidders' best responses imply a stark prediction: the support of their bids should have gaps because the first-price bidder wants to exploit the jump in the probability of winning just above a mass point in the distribution of competitor bids.

Propositions 1 and 2 describe the general features of the bidder best response functions. Proposition 1 states that it is still optimal for the second-price bidder to bid his valuation in a hybrid auction.

Proposition 1. *In an independent private values auction with hybrid rules, a weakly dominant strategy for a second-price bidder is to bid his valuation, i.e. $\beta^{SP}(v) = v$.*

This result follows from the usual dominant strategy logic of a pure second-price auction (see e.g. Krishna 2009).

Proposition 2 describes the first-price bidder’s optimal bidding behavior.

Proposition 2. *In an independent private values auction with hybrid rules, the first-price bidder optimally*

1. *shades her bid below her valuation: $\beta^{FP}(v) \leq v$; and*
2. *avoids bidding in an interval $(b_L, B]$ for some b_L when she faces competition from a bidder who bids $B > r$ with probability $\Pr[B] > 0$. $\beta^{FP}(\cdot)$ is therefore discontinuous meaning that the support of the first-price bidder’s induced distribution of bids is not connected.*

I sketch the intuition behind Proposition 2 below and refer the reader to Appendix A. The first-price bidder i chooses her bid b , given her valuation v , to maximize

$$\max_b (v - b) \Pr[i \text{ wins} | b]$$

Thus, the first-price bidder trades off the higher surplus conditional on winning, and the diminished probability of winning as she raises her bid. As with pure first-price rules, Part 1 of Proposition 2 holds because bidding above (below) her valuation ensures a negative (positive) expected payoff. Part 2 states that a first-price bidder avoids bidding in an interval of *strictly dominated* bids below a competitor’s bid B when $\Pr[B] > 0$, a phenomenon I refer to as a *bid gap*. The intuition behind the bid gap is that the discontinuous improvement in the probability of winning at $B + \varepsilon$, for $\varepsilon > 0$ small, more than compensates for the diminished surplus conditional on winning. Consider for example a first-price bidder with valuation $v = \$1.25$ who faces a single rival who bids $B = \$1$ with probability $\frac{1}{2}$. The first-price bids of $b = \$0.99$ and $b = \$1.01$ yield expected utility $(\$1.25 - \$0.99) \frac{1}{2} = \$0.13$ and $(\$1.25 - \$1.01) 1 = \$0.24$ respectively; thus, outbidding B can increase the first-price bidder’s payoffs.

Before characterizing the equilibrium, I make two assumptions about the bidders.

Assumption 4.1. *The first-price bidder valuations v are independently distributed on an interval $[0, \bar{V}]$ according to the distribution $F(v)$ with continuously differentiable density $f(v)$, which is strictly positive on the interior of the support of $F(v)$.*

Assumption 4.2. *The second-price bidder has valuation v above the seller’s reserve price r , and bids B with probability $\Pr[B] = \alpha > 0$ and bids 0 with complimentary probability.*

Assumption 1 states that the first-price bidders have independent private values and makes regularity assumptions on the distribution of valuations $F(\cdot)$. Assumption

2 reflects a feature of the institutional setting: the second-price bidders are offline bidders who submit a fixed bid randomly on ad inventory that meets their criteria. I assume that the probability of submitting this bid is exogenous. The offline bidders bid randomly in order to spread their fixed advertising budget over time and across users, or ‘pulse’ their advertising in a simple and transparent manner. Strictly speaking, however, this strategy is suboptimal in this environment: the budget-constrained offline bidders should optimally bid on all auctions and shade his bid downwards (Gummadi et al., 2012).

Proposition 3 characterizes the equilibrium when first-price bidders compete with second-price bidders who submit bids with positive probability.

Proposition 3. *Suppose the auction has $n > 1$ symmetric first-price bidders who satisfy Assumption 4.1 and m asymmetric second-price bidders who satisfy Assumption 4.2. The seller sets a reserve price of $r \geq 0$. Then a pure, symmetric Bayesian Nash equilibrium exists in which*

1. *the first-price bidder’s equilibrium bidding function $\beta^{FP}(v)$ is strictly increasing with $\beta^{FP}(r) = r$ and has $K \leq m + 1$ segments of the form*

$$\beta_k^{FP}(x; c_k) = v - \frac{\int_{v'_k}^v F(u)^{n-1} du + c_k}{F(v)^{n-1}} \quad (4.1)$$

for $v \in (v'_k, v'_{k+1})$, constants c_k , and threshold values v'_k ; and

2. *second-price bidder j ’s equilibrium bidding function is*

$$\beta^{SP}(v_j) = \begin{cases} 0 & \text{with } \Pr = 1 - \alpha_j \\ v_j & \text{with } \Pr = \alpha_j \end{cases}$$

I sketch the intuition below and refer the reader to Appendix A for the full proof. The second-price bidders bid their valuations due to Proposition 1, but submit this randomly due to Assumption 2. The first-price bidders equilibrium behaviour is more complicated. Consider first the solution without the offline bidder. A unique symmetric equilibrium exists under these weak regularity conditions on $f(\cdot)$ (Athey and Haile, 2007). The equilibrium bidding function is $\beta^{FP}(v; 0)$ because the initial condition $\beta^{FP}(r) = r$ implies that $c = 0$. The first-price bidder’s first-order conditions define the optimal $\beta^{FP}(v)$ as the solution to a differential equation. When the competitors’ distribution has mass points, the solution shares the same first-order condition which defines the same differential equation as without the mass points. Due to Proposition

2, the equilibrium bidding function has multiple initial conditions. The first initial condition is still $\beta_1^{FP}(r) = r$ implying that $c_1 = 0$. Suppose the competition features a single positive bid B with $\Pr[B] > 0$. Proposition 2 states that the discontinuous improvement in the winning probability at $b = B + \varepsilon$ implies that this bid dominates bids in some interval $(b_L, B]$. The threshold b_L and its associated valuation x' such that $b_L = \beta(v')$ are then given by equating the expected payoffs at b_L and $B + \varepsilon$ as $\varepsilon > 0$ approaches 0. Beyond v' , $\beta^{FP}(\cdot)$ is given by $\beta_2^{FP}(v; c_2)$ where c_2 is pinned down by the initial condition $\lim_{\delta \rightarrow 0^+} \beta^{FP}(v' + \delta) = B$. Note that this phenomenon occurs because some bidders play by second-price rules. If all bidders play by first-price rules, then the bidder who submits B with positive probability could win as often by bidding $b_L + \varepsilon$ while reducing payment, so bidding B is not an equilibrium. However, this phenomenon is not limited to hybrid auctions: it also occurs in pure first-price auctions if the seller's reserve price is secret and not continuously distributed.

Figure 6 illustrates a simple example in which two first-price bidders with uniformly distributed values face an offline bidder who randomly bids $B = \frac{1}{4}$ with probability $\frac{1}{2}$ and bids 0 with remaining probability. The first-price bidder's optimal bidding function is given by

$$\beta^{FP}(v) = \begin{cases} \frac{v}{2} & \text{if } v \leq \frac{1}{3} \\ \frac{v}{2} + \frac{1}{36v} & \text{if } v > \frac{1}{3} \end{cases}$$

which implies the bid gap $b \notin (\frac{1}{6}, \frac{1}{4}]$. Figure 6 shows that $\beta^{FP}(v)$ coincides with the solution (dashed line) without the offline bidder ($\beta(v) = \frac{v}{2}$) until the indifference valuation $v' = \frac{1}{3}$ and its associated bid $b_L = \frac{1}{6}$. Figure 6's dotted indifference line equates the utilities at $B + \varepsilon$ and b_L . The indifference line intersects with the bidding function to pin down v' and b_L . Beyond this valuation, the bidding function is reinitialized above B at the initial condition $\beta^{FP}(\frac{1}{3}) = \frac{1}{4}$.

Though Proposition 3 establishes the existence of an equilibrium, the equilibrium may not be unique because, in cases with multiple offline bids, the gaps can overlap. In Appendix B.1, I establish that multiple equilibria can exist a simple example. In the example, two first-price bidders with uniformly distributed values face two offline bids (0.25, 0.3) with probabilities (0.1, 0.15). The first equilibrium has two bid gaps (0.227, 0.25] and (0.266, 0.3], and the second equilibrium has the single bid gap (0.243, 0.3]. Unfortunately, we typically observe bids in the theoretical gaps, so the data can not be used to guide which equilibrium is selected. My bidding function algorithm selects an equilibrium by moving through the vector of offline bids from the lowest to the highest component B_i and selects the highest component B_j such

that its gap includes B_i (i.e., $b_j^L \leq B_i$). In the example above, this would select the second equilibrium with the single bid gap. See Appendix B, for a more in depth discussion of multiple equilibria in the hybrid auction.

The theorized gaps in the bidding function is a stark prediction that could fail in the data for various reasons. The analysis assumes that all bidders have correct beliefs regarding their opponents' distribution of valuations. Bids in a predicted gap may reflect incorrect beliefs or experimentation and learning in an evolving marketplace. In reality, bidders must pay to learn about their opponents' distribution of bids by investing in better algorithms and incurring computational and data storage costs to update their model. So, these costs may exceed the loss in profits due to submitting a strictly dominated bid in a gap. Furthermore, a bidder must aggregate impression markets at some level to learn its competitors' bid distribution. Since the markets can be cut differently (by time, user and publisher characteristics), the bidder's beliefs would be incorrect if he aggregates the markets incorrectly. Nonetheless, the frequency and consistency of some offline bids are hard to ignore. These factors explain why we empirically observe real time bids in a theorized gap. As I elaborate in the model section, I abstract away from these factors and assume that bidders randomly bid in a gap with some positive probability but otherwise behave optimally.

Finally, I discuss some of the key assumptions that I use in this section, starting with the assumption that valuations are private. This assumption rules out common values meaning that, conditional on observables, the valuations of a bidder's opponents are irrelevant to the bidder's own valuation. In reality, the market has a common value flavor. Some impressions are generated by computer algorithms called bots that advertisers view as worthless 'lemons' in the language of Abraham et al. (2011). This is a common value setting if bidders receive a private signal that a user is a bot. However, identifying bots is expensive. My discussions with people in the industry suggest that only some of the largest ad buyers invest in bot detection. Advertisers instead reduce their bids to reflect the fact that a certain portion of their ad purchases are worthless. Overall, my discussions with the people in the industry suggest the common value problem is not top of mind among advertisers. It is an open issue whether a common value auction model is identified when only the top two bids are observed or observed with censoring due to the reserve price. Though I am unable to identify a common value auction, my full model in section 5 allows bidder's bids to be correlated within users.

I further assume independence for a given auction. Though I allow the distribution of private values to be conditional on characteristics observed by the bidder, the

distributions are otherwise independent across bidders. I also assume that the value distribution is independent across time. Independence is especially relevant for the real-time bidders, since it allows me to treat their decision problem as static. Real-time bidders strategies could be dynamic in the sense that they incorporate an ad’s past performance (e.g. clicks). However, the data do not contain any ad performance information. Celis et al. (2012) also assume independent private values in their empirical analysis of a second-price online display ad auction. I weaken their assumption of independence across bidders by allowing valuations to be correlated within users due to an unobserved heterogeneity term as I discuss in the model section.

5 Model

In this section, I model user tracking in the online display ad auction market as user-level unobserved heterogeneity. That is, auctions systematically differ by user characteristics that bidders observe due to tracking but that I—as the econometrician—do not observe. This section comprises two subsections that model: 1) the offline bidders; and 2) the real-time bidders.

Throughout this section, I maintain several assumptions. All advertisers publicly view user tracking characteristics. These characteristics can include user demographics, interests, and past purchases. Formally, each bidder views a vector of user u ’s characteristics c_u in some space \mathcal{C} . For instance, \mathcal{C} could be the product space of gender and income, $\mathcal{C} = \{M, F\} \times \mathbb{R}_+$. As in the previous section, bidders have independent private values conditional on observable auction characteristics and the unobservable (from our point of view) user tracking characteristics. This assumption allows bidder values to differ systematically between users. The assumption rules out common values which allow bidders to see private signals regarding user tracking characteristics that would be informative about their opponents’ valuations. It is not known whether a common value model is identified when only the two highest bids are observed with censoring. To focus on tracking, I condition on observable auction characteristics (e.g. website, user country and browser) so that auctions differ only in user tracking characteristics. Finally, bidders are exogenously assigned to be either offline or real-time bidders. I do allow the distribution of valuations to differ between the offline and real-time bidders. The differences between real-time and offline bidder behavior arise from their underlying preferences.

5.1 Offline Bidders

The offline bidder’s restricted strategy space includes a fixed bid, a set of targeting criteria, and a random bidding probability. Offline bidders are asymmetric and specify targeting criteria based on user tracking characteristics. Formally, offline bidder i chooses a subset of user characteristics $T_i \subset \mathcal{C}$ that defines its target audience. When user u satisfies bidder i ’s criteria, I write $u \in T_i$. In the model, we wish identify and estimate the proportion of users who satisfy i ’s criteria $\Pr[u \in T_i]$ and the probability that i bids $B_i > 0$ on his target audience, $\Pr[b_i = B_i | u \in T_i]$. I show the model is identified from only the winning bid when the offline bidder criteria overlap. Since users can satisfy several bidders’ criteria simultaneously, the offline bidder’s user targeting can be viewed as a mixture model. The mixture model is identified by repeated user observations and estimated using maximum likelihood. When bidders can not track users, their counterfactual bids reflect the average value of untargeted impressions.

Assumption 5.1.1 rationalizes the offline bidder’s fixed bid. Offline bidder i values user u ’s t th auction according to:

Assumption 5.1.1. $v_{iut}^{Off} = x_i \cdot y_{iu}$ where $y_{iu} = I[u \in T_i]$.

Assumption 5.1.1 states that offline bidder i ’s valuation is the product of the bidder-specific constant x_i and the user-bidder match term y_{iu} . y_{iu} is an indicator variable that equals one when user u meets i ’s criteria. x_i represents i ’s value for users in his target group. Recall that offline bidders play by second-price rules and bid their valuation. As such, offline bidder i bids $B_i = x_i$. Assumption 5.1.1 also asserts that users outside the target group are worthless. This assumption could be relaxed, since revealed preference implies that the value of untargeted users is less than the reserve price. y_{iu} is the unobserved heterogeneity term whose binary form allows for a richer model that matches bidders to their desired user tracking characteristics. In principle, offline bidders could enter multiple bids that specify different values x_i for different target groups. This is rare in the data, however, so I treat different bid amounts as different bidders for simplicity. I wish to relax this in the future.

Assumption 5.1.2. rationalizes the final component of the offline bidder’s strategy space: probabilistic bidding.

Assumption 5.1.2. a) *Offline bidder i ’s conditional bidding probability functions are*

$$\beta^{Off} \left(v_{iut}^{Off} | u \in T_i \right) = \begin{cases} 0 & \text{with } \Pr = 1 - \alpha_i \\ B_i & \text{with } \Pr = \alpha_i \end{cases}$$

$$\beta^{Off} \left(v_{iut}^{Off} | u \notin T_i \right) = 0 \text{ with } \Pr = 1$$

b) When i can no longer track users, its conditional bidding probability function are

$$\beta^{Off,CF} \left(v_{iut}^{Off,CF} \right) = \begin{cases} 0 & \text{with Pr } 1 - \alpha_i \\ B_i^{CF} & \text{with Pr } \alpha_i \end{cases}$$

Offline bidder i bids randomly on users that meet its criteria according to an exogenously given probability. This arises from an advertiser's fixed budget which it wishes to spread across time and users. I take the ad budget, and hence α_i as given. I maintain the same bidding probability α_i in the counterfactual. Given its rationale, we could alternatively modify α_i so that i expends his same budget in the counterfactual.

Proposition 4 asserts that, when the offline bidder can no longer track users, he will bid his expected valuation for untargeted users.

Proposition 4. *Under Assumptions 5.1.1 and 5.1.2b, when the offline bidder i cannot track users, i counterfactual (denoted by CF) is given by*

$$B_i^{CF} = E \left[v_{iut}^{Off} \right] = B_i \tau_i$$

where $\tau_i = \Pr [u \in T_i]$ and i bids randomly according to $\Pr [b_i^{CF} = B_i^{CF}] = \alpha_i$ and $\Pr [b_i^{CF} = 0] = 1 - \alpha_i$.

Proof. i 's counterfactual valuation can be decomposed using the Law of Total Probability

$$\begin{aligned} E [v_{iut}] &= E [v_{iut}|u \in T_i] \Pr [u \in T_i] + E [v_{iut}|u \notin T_i] \Pr [u \notin T_i] \\ &= x_i \Pr [u \in T_i] + 0 \cdot \Pr [u \notin T_i] \\ &= x_i \Pr [u \in T_i] \\ &= B_i \tau_i \end{aligned}$$

This follows from Assumption 5.1.1, which asserts that an offline bidder has 0 value for users outside his targeting criteria. Proposition 1 states that $B_i^{CF} = E [v_{iut}]$. i 's probabilistic bidding is assumed in 5.1.2b. \square

Proposition 4 highlights the role of the targeting probability τ_i and the conditional bidding probability α_i in describing the offline bidder's counterfactual behavior. For example, an offline bidder who bids \$1 on women with $\tau = \frac{1}{2}$ will bid \$0.50 when he no longer observes gender. Given that we assume $E [v_{iut}|u \notin T_i] = 0$ rather than $E [v_{iut}|u \notin T_i] \in [0, r]$, we can view our counterfactual estimates as a lower bound for the counterfactual revenues.

5.1.1 Identification

The offline bidder model is identified using repeated user auctions when viewed as a mixture model. We wish to identify the two model parameters: bidder i 's targeting probability $\tau_i = \Pr[u \in T_i]$ and conditional bidding probability $\alpha_i = \Pr[b_i = B_i | u \in T_i]$. I begin by explaining how these quantities are identified in the simpler case in which we observe all bids. I proceed to show that these quantities are identified in a setting with ordered bids and unconditionally dependent bids.

By first considering the case in which when all offline bids are observed, we see that the offline bidder model is a mixture model. Again, we wish to identify τ_i and α_i . We require that $B_i \geq r$ to identify $x_i = B_i$ as well as α_i and τ_i . We incompletely observe the set of users such that $u \in T_i$ since $b_{iut} = 0$ can occur when $u \notin T_i$ or when $u \in T_i$ with probability $1 - \alpha_i$. As such, $\Pr[b_{iut}]$ is given by a mixture model with the user two types $u \in T_i$ and $u \notin T_i$ and their associated conditional probability $\Pr[b_{iut} | u \in T_i]$ and $\Pr[b_{iut} | u \notin T_i]$.

If users appear only once in the data, we can still create informative bounds on the parameters α_i and τ_i . The set of users for which we observe $b_{iut} = B_i$ provide a lower bound for τ_i . Denote these users by $u \in T_i^{LB}$. As for the upper bound, all users could satisfy i 's criteria ($\tau_i^{UB} = 1$) because the i bids randomly on those users that meet its criteria. Similarly, α_i is bounded below by assuming all users satisfy i 's criteria ($\tau_i = 1$) and bounded above by assuming that τ_i equals its lower bound τ_i^{LB} . The bounds are thus

$$\begin{aligned} \tau_i &\in [\tau_i^{LB} = \Pr[u \in T_i^{LB}], 1] \\ \alpha_i &\in [\alpha_i^{LB} = \Pr[b_{iut} = B_i | \tau_i = 1], \alpha_i^{UB} = \Pr[b_{iut} = B_i | \tau_i = \tau_i^{LB}]] \end{aligned} \tag{5.1}$$

If we observe the same user across multiple auctions, the quantities α_i and τ_i are point identified. Intuitively, though we do not know if any user u with $b_{iut} = 0$ for all t satisfies $u \in T_i$, this becomes increasingly unlikely (and informs α_i) as we observe more auctions t' with $b_{iut'} = 0$. Suppose we observe the same user twice, then the probability of their observed bids is given by

$$\begin{aligned} \Pr[b_{iut}, b_{iut'}] &= \Pr[b_{iut}, b_{iut'} | u \in T_i] \Pr[u \in T_i] + \Pr[b_{iut}, b_{iut'} | u \notin T_i] \Pr[u \notin T_i] \\ &= \Pr[b_{iut} | u \in T_i] \Pr[b_{iut'} | u \in T_i] \tau_i + \Pr[b_{iut} | u \notin T_i] \Pr[b_{iut'} | u \notin T_i] (1 - \tau_i) \end{aligned}$$

using the Law of Total Probability. The expression simplifies whenever we observe an auction t' such that $b_{iut'} = B_i$ since $\Pr[b_{iut'} = B_i | u \notin T_i] = 0$. If we observe users

twice, the system of equations

$$\begin{aligned}\Pr [b_{iut} = 0, b_{iut'} = 0] &= (1 - \alpha_i)^2 \tau_i + 1 - \tau_i \\ \Pr [b_{iut} = 0, b_{iut'} = B_i] &= 2\alpha_i(1 - \alpha_i) \tau_i\end{aligned}\tag{5.2}$$

pins down τ_i and α_i .

Two factors complicate the identification of τ_i and α_i in our setting: 1) bidder targeting probabilities are dependent conditional on observables; and 2) we only observe the top two bids. Overlap in bidder criteria induces dependence in offline bidders' unobserved heterogeneity terms y_{ui}, y_{uj} . This is unavoidable since, for example, two advertisers could target the same gender (perfect positive correlation) or target opposite genders (perfect negative correlation). Since we observe only the top two bids and bids are not independent, lower bids are censored from above and are not censored at random. Proposition 5 overcomes these challenges and separately identifies each offline bidder's τ_i and α_i in the weaker setting where we only observe the winning bid.

Consider now the case where we only observe the winning bid W_{ut} as depicted in Figure 7. The winning bid simplifies to three possible cases: $W_{ut} < B_i$, $W_{ut} = B_i$ and $W_{ut} > B_i$ ⁶. b_{iut} is censored when $W_{ut} > B_i$, so we can simplify i 's competition to the effective competitors who outbid B_i and collectively denote them by C . Figure 7 illustrates a Venn diagram of bidder of i and C 's target subpopulations. i exclusively targets the subset of users satisfying the criteria $T_{i \setminus C}$. C exclusively targets the criteria $T_{C \setminus i}$, they both target $T_{i \cap C}$, and they both ignore T_\emptyset . Denoting $\tau_E = \Pr [u \in T_E]$, we seek the targeting probability $\tau_i = \tau_{i \setminus C} + \tau_{i \cap C}$. This targeting model is general in that it allows any dependence structure in the match indicators y_{iu} and y_{Cu} . Further, I allow i 's effective competition to vary its conditional bidding probability on the users it alone targets: $\Pr [W_{ut} > B_i | u \in T_{C \setminus i}] \equiv \alpha_{C \setminus i} \neq \alpha_{i \cap C} \equiv \Pr [W_{ut} > B_i | u \in T_{i \cap C}]$.

As before, some observed bids rule out possible user types. In particular, $W_{ut} = B_i$ implies that $u \notin T_{C \setminus i} \cup T_\emptyset$, and $W_{ut} > B_i$ implies that $u \notin T_{i \setminus C} \cup T_\emptyset$. As an example, if we see the same user twice with bids $W_{u1} < B_i$ and $W_{u2} > B_i$, the probability of the data is given by

$$\begin{aligned}& \Pr [W_{u1} < B_i, W_{u2} > B_i] \\ &= \Pr [W_{u1} < B_i | u \in T_{C \setminus i}; \alpha] \Pr [W_{u2} > B_i | u \in T_{C \setminus i}; \alpha] \tau_{C \setminus i} \\ & \quad + \Pr [W_{u1} < B_i | u \in T_{i \cap C}; \alpha] \Pr [W_{u2} > B_i | u \in T_{i \cap C}; \alpha] \tau_{i \cap C} \\ &= (1 - \alpha_{C \setminus i}) \alpha_{C \setminus i} \tau_{C \setminus i} + (1 - \alpha_{i \cap C}) (1 - \alpha_i) \alpha_{i \cap C} \tau_{i \cap C}\end{aligned}$$

⁶For ease of exposition, I assume away ties, so that $W_{ut} = B_i$ if and only if i wins the auction.

The winning bid data also represent the outcome of a mixture model where a user potentially belongs to at most four possible types.

Proposition 5 presents my identification result.

Proposition 5. *Under Assumptions 5.1.1 and 5.1.2, α_i and τ_i are identified from the winning bid when we observe some users who appear once and some who appear twice in auctions provided that the solution to the implied system of equations is unique.*

See Appendix A for the detailed proof. The proof includes a system of seven equations of the form in in equation (5.2). These equations pin down the six unknowns in the vectors $\alpha = [\alpha_i, \alpha_{i \cap C}, \alpha_{C \setminus i}]$ and $\tau = [\tau_{i \setminus C}, \tau_{i \cap C}, \tau_{C \setminus i}]$. These quantities are not identified, for instance, when $\alpha_{i \cap C} \equiv \Pr[W_{ut} > B_i | u \in T_{i \cap C}] = 1$. Specifically, if i 's competitors always bid on users that both i and C target, then i 's bids will always be censored. Then, i 's targeting probability and conditional bidding probability are not identified on this region of users. Since we actually observe the top two bids, we could handle at most one such bidder by dropping his bids.

This paper is among the first to use mixture models in an empirical auction setting. Lamy (2012) identifies an asymmetric first-price auction model when bidders are anonymous. Observed bid distributions are a mixture over the individual bid distributions. He explains that his model can also be applied to models with unobserved heterogeneity when its form is discrete. An et al. (2010) identify a first-price auction model where the number of bidders is possibly random and unknown to the econometrician but known to bidders. In their case, the mixture is over the number of bidders who generate the bid data.

5.1.2 Estimation

I estimate the offline bidder model using constrained maximum likelihood estimation. My estimation approach follows my identification argument. The parameters of interest are the targeting probability $\tau_i = \tau_{i \setminus C} + \tau_{i \cap C}$ and conditional bid probability α_i . The nuisance parameters are $\tau_{C \setminus i}$, $\alpha_{C \setminus i}$, and $\alpha_{i \cap C}$.

The log-likelihood function for the parameters $\theta \equiv \tau, \alpha$ has the form

$$\mathcal{L}(W; \theta) = \sum_U \log \left[\sum_E \tau_E \cdot \prod_{T_u} \Pr(W_{ut} | u \in T_E; \alpha) \right]$$

where U denotes the set of users, $E = \{i \setminus C, i \cap C, C \setminus i, \emptyset\}$ the possible user types, and T_u the total number of auctions for user u . The constrained maximum likelihood

estimation problem is written as follows

$$\begin{aligned} & \max_{\theta \equiv \tau, \alpha} \mathcal{L}(W; \theta) \\ s.t. & \sum_E \tau_E = 1 \\ & [\tau, \alpha] \equiv \theta \in [\underline{\theta}, \bar{\theta}] \end{aligned}$$

This includes the equality constraint that all users must satisfy a single set of targeting characteristics. The constraints also include rectangular bounds of the form (5.1). Though the likelihood function currently only makes use of the winning bid, I use the second-highest bid data to tighten these bounds on τ and α . The rectangular constraints also reduce computational time. The model simplifies the winning bid data into a few main cases that can be collapsed into a frequency weighted data set. This greatly reduces the computational burden of estimating the structural parameters in datasets with millions of auctions. In the future, I can see if an expectation-maximization algorithm approach improves performance.

In Monte Carlo simulations, the estimator generally performs well. When i wins infrequently, the estimator has trouble distinguishing between low τ_i and low α_i cases. This problem recedes as the number of users and the number of observation per user increase.

When the parameters are not on the boundary, I use the nonparametric bootstrap for inference. Inference is challenging when parameter estimates approach the boundary. The nonparametric bootstrap is known to perform poorly on the boundary (Horowitz 2001). In this mixture model, we can lose identification when parameters are on the boundary (e.g. when $\alpha_{i \cap C} = 1$, as discussed above). Inference in such cases remains an unsettled issue in the econometrics literature. The parametric bootstrap might be appropriate in this case.

5.2 Real-Time Bidders

In the model, the real-time bidders rank users in terms of their tracking characteristics. Real-time bidders publicly observe user tracking characteristics, which are unobserved by the econometrician. Tracking characteristics enter the real-time bidder's valuation as a common, user-level unobserved heterogeneity component. Their valuation also includes an idiosyncratic taste component. In the counterfactual, the unobserved user heterogeneity term collapses to its mean value across consumers while the idiosyncratic term is still allowed to vary. I separately identify the distribution of the idiosyncratic and unobserved user heterogeneity components using a combi-

nation of within-user and between-user variation in bids. To estimate the model, I isolate users who share a common realization of the unobserved user heterogeneity term. I then use maximum likelihood to estimate the idiosyncratic term’s distribution holding the unobserved heterogeneity realization fixed. Proposition 2 predicts that bidders should not bid in an interval below an offline bid, but observed bids do not respect this. To rationalize this behavior, I model real-time bids as arising from a probability mixture of optimal bids with probability $1 - \delta$ and sub-optimal bids in these dominated intervals with probability δ . Real-time bidders optimize their bids against their opponents’ entire bid distribution.

Assumption 5.2.1 describes real-time bidder i ’s value of user u ’s t th auction.

Assumption 5.2.1.

$$v_{iut}^{RT} = x_{it}y_u$$

where x_{it} and y_u satisfy the regularity conditions in Assumption 4.1 (though the support of y_u can be strictly positive). x_{it} and y_u are mutually independent as are x_{it} and the offline bidder’s y_{iu} .

Real-time bidders are symmetric. Their utility function is the product of a bidder- and auction-level idiosyncratic taste term x_{it} and a common, user-level unobserved heterogeneity term y_u . x_{it} captures the bidder’s private taste which can vary due to budget smoothing and ad performance. y_u captures the bidder’s shared quality ranking for users. Bidders share a function $f : \mathcal{C} \rightarrow \mathbb{R}_+$ that maps the multi-dimensional user tracking characteristics into a user quality scalar. We can think of this scalar as being increasing in user income and responsiveness to advertising. This utility function resembles that of Krasnokutskaya (2011) though with user-level rather than auction-level unobserved heterogeneity. The regularity assumptions on x_{it} and y_u ensure an equilibrium (see Proposition 3). They are also used to identify and estimate the model.

Assumption 5.2.2 lays out a real-time bidder’s counterfactual valuations (denoted by CF) when it cannot track users.

Assumption 5.2.2. $v_{iut}^{RT,CF} = x_{it}E[y_u]$

Without tracking, real-time bidders evaluate untargted users according to the expected value of their unobserved user heterogeneity. The counterfactual valuations still vary because x_{it} fluctuates as before. The identification argument pins down y_u and therefore $E[y_u]$.

Thus, when bidders cannot track, the new equilibrium is given by a corollary to Proposition 3.

Corollary 6. *Suppose the auction has $n > 1$ symmetric real time (first-price) bidders who satisfy the regularity conditions in Assumption 5.2.1 and the preferences in Assumption 5.2.2. The auction also has m asymmetric offline (second-price) bidders who satisfy Assumption 5.1.1 and 5.1.2. The seller sets a reserve price of $r \geq 0$. Then a pure, symmetric Bayesian Nash equilibrium exists in which*

1. *the real-time bidder's counterfactual equilibrium bidding function $\beta_k^{RT,CF}(x; E[y_u])$ is strictly increasing with $\beta_k^{RT,CF}\left(\frac{r}{E[y_u]}; E[y_u]\right) = r$ and has $2 \leq K \leq m + 1$ segments of the form*

$$\beta_k^{RT,CF}(x; E[y_u]) = xE[y_u] - E[y_u] \frac{\int_{\frac{x'_k}{E[y_u]}}^x F(u)^{n-1} du + c_k(E[y_u])}{F(x)^{n-1}} \quad (5.3)$$

for $x \in \left(\frac{x'_k}{E[y_u]}, \frac{x'_{k+1}}{E[y_u]}\right)$, constants $c_k(E[y_u])$ that depend on Y_u and threshold values x'_k ; and

2. *the offline bidder j 's equilibrium counterfactual bidding function is*

$$\beta^{Off,CF}(v_{jt}^{Off,CF}) = \begin{cases} 0 & \text{with } \Pr = 1 - \alpha_j \\ v_{jt}^{Off,CF} & \text{with } \Pr = \alpha_j \end{cases}$$

where $v_{jt}^{Off,CF} = E[v_{jut}|u \in T_j] \Pr[u \in T_j] = x_j \tau_j$.

Corollary 6 assembles previous results and lays out the counterfactual market under privacy policy. In Section 6, I simulate the counterfactual market and compare it to the status quo in order to evaluate the policy. Corollary 6 also shows how the model's structural parameters enter the counterfactual. A real-time bidder's preferences depend on the distribution of her idiosyncratic taste term and the mean of her unobserved user heterogeneity term. Corollary 6 clarifies that the real-time bidder's counterfactual bids reflect not only her modified utility function due to Assumption 5.3.2, but also her response to the modified offline bids. An offline bidder's behavior is a function of his targeting probability (τ_i), his conditional bidding probability (α_j) and his value for targeted users, given by $x_i = B_i$. Below, I discuss how I identify and estimate the real-time model.

5.2.1 Identification

To construct the market counterfactual in Corollary 6, we first wish to identify the distribution of the idiosyncratic and unobserved user heterogeneity utility terms $F_x(\cdot)$ and $F_y(\cdot)$. Proposition 7 states that these distributions are identified.

Proposition 7. *Suppose the auction has $n > 1$ symmetric real-time bidders who satisfy Assumption 5.2.1, m asymmetric offline bidders who satisfy Assumption 5.1.1 and 5.1.2, and a seller who sets the reserve price r . $F_x(\cdot)$ and $F_y(\cdot)$ are identified from the support maximum or another quantile of user bids.*

Appendix A contains a detailed proof. I present an overview below.

The proof of Proposition 7 employs the maximum and occasionally another quantile of user bids to isolate the distributions of the idiosyncratic term x_{it} and unobserved user heterogeneity term y_u . Proposition 7's conditions generate the equilibrium in Proposition 3. This defines the real-time bid function $\beta^{RT}(x; Y_u)$, which depends on the realized Y_u . From Assumption 5.2.1, denote the bounded support of x_{it} and y_u by $[\underline{X}, \bar{X}]$ and $[\underline{Y}, \bar{Y}]$. By fixing a user u , we also fix u 's realized unobserved heterogeneity term Y_u . Let $b(u) = \beta(x; Y_u)$ denote a real-time bid for user u , where $b(u)$ is a random variable that varies due to x_{it} . Denote $b(u)$'s support by

$$[\underline{b}(u), \bar{b}(u)] = [\beta(\underline{X}; Y_u), \beta(\bar{X}; Y_u)]$$

Figure 8 conveys the intuition behind the proof. Figure 8 layers the supports of x_{it} and y_{it} and relates them to the support of bids. The within-user variation $b(u)$ in bids identifies $F_x(\cdot)$ by holding $y_u = Y_u$ constant. To connect $b(u)$ to x , I use the inverse bid function, that Proposition 3 delivers. The between-user variation in a user's maximum bid $\bar{b}(u) = \beta(\bar{X}; Y_u)$ primarily identifies $F_y(\cdot)$ by holding $x_{it} = \bar{X}$ constant. To connect $\bar{b}(u)$ to Y_u , I show that $\beta(\bar{X}; Y_u)$ defines a finite inverse correspondence that outputs Y_u . When this correspondence is not unique, I can pin down the correct Y_u using a bid quantile sufficiently close to the reserve price.

The support variation approach extends auction models with unobserved heterogeneity to settings where the reserve price censors the observed bid distribution. My identification argument resembles D'Haultfoeuille and Février (2010). Their result is more general in that it only requires three observed repeated measurements (bids) to identify the distributions of x and y . Similarly, Hu et al. (2009)'s allow $\beta(x; y)$ to be non-separable in y and also requires three bids. Krasnokutskaya (2011) needs only two bids when $\beta(x; y)$ is separable in y (e.g. $u_{it} = x_{it}y_{it}$ implies $\beta_i(u_{it}) = y_{it}\alpha_i(x_{it})$). Nevertheless, none of these approaches can accommodate censoring of the bid distribution, since the observed x and y combinations are no longer random. Moreover, bids are non-separable in y whenever the real-time bidders face a binding reserve price or an offline (positive probability) bid. D'Haultfoeuille and Février (2010), Krasnokutskaya (2011) and Hu et al. (2009) are also more general in the sense that they allow for auction-level than user-level unobserved heterogeneity. However, it is not known

if such models are identified when we observe ordered bids (e.g. the two highest bids) rather than random bids.

5.2.2 Estimation

We wish to disentangle and estimate the distributions of the idiosyncratic and un-observed user heterogeneity terms $F_x(\cdot)$ and $F_y(\cdot)$. Following the identification by support variation argument, I use frequent users' maximum bid to isolate y_u and estimate $F_y(\cdot)$. I use maximum likelihood to estimate $F_x(\cdot)$ holding y_u fixed. To rationalize all observed bids, I model bids as a probability mixture of optimal and sub-optimal bids in the dominated bid gap. Real-time bidders optimize against their competitors bids, which includes the sub-optimal bids. I discuss the estimation process below.

In order to locate the bid gaps, the estimator needs to first fix a distribution of valuations. Recall from Section 3 that real-time bidders play by first-price rules and face offline competitors who submit bids with positive probability. Proposition 2 predicts that real-time bidders should avoid bidding in an interval below an offline bid: the 'bid gap'. The rationale is that their probability of winning improves discontinuously by outbidding the offline bidder. We can locate the bid gaps by equating the real-time bidder's utility at the endpoints of a gap, provided we know the bidder's distribution of valuations. At each iteration, the estimator employs a given value distribution to derive the optimal bid function and its gaps. It then searches over a family of distributions to best fit the empirical bid distribution. Thus, the estimator reverses previous approaches that first estimate the distribution of bids, then back out the implied valuations, in settings with a continuous bid function (Guerre et al., 2000).

The data show little evidence that real-time bids respect the bid gaps predicted by the theory. Figure 2 depicts a typical real-time bidder's bids for one publisher, which shows no obvious gaps in its bid support. Figure 9 shows all bidder's bids on a single ad form (defined by the ad dimensions) and a single publisher over the week of data. As is typical, the real-time bids overlap with the big gaps, which are conservatively given by the gray bars.⁷ As it stands, my theory of bidding cannot rationalize a bid in the gap.

To account for bids in the gap, I assume the distribution of bids $G(b)$ represents a probability mixture of the optimal bids $G^*(b)$ and sub-optimal bids $G^{gap}(b)$ in the gap, given by

$$G(b; \delta) = (1 - \delta) G^*(b) + \delta G^{gap}(b)$$

⁷The bid gaps in Figure 9 assume a fixed 10% markdown ($b_{iut} = 0.9v_{iut}$).

I take the stand that bids in the gap represents sub-optimal behavior that is not rationalized by my model. This would occur for instance if the real-time bidder’s priors differ from the empirical distribution over my sample period. The dominant strategy logic of the bid gaps is hard to avoid, though it might be interesting to consider richer models of learning and dynamics that better fit the data. The δ approach has its own challenges: the optimal bid function endogenizes δ in that the real-time bidder optimizes against $G(b; \delta)$ rather than $G^*(b)$. Finally, I estimate δ as a structural parameter.

The real-time estimator includes three steps:

1. Isolate the unobserved heterogeneity term y_u
2. Estimate the idiosyncratic taste distribution $F_x(\cdot)$ and the bid gap probability δ
3. Estimate $E[y_u]$

In the first step, I select a group of users who share a common realized Y_u term among the set of users who appear frequently. In the second step, I estimate the within-user distribution of the idiosyncratic taste term x_{it} using maximum likelihood. In the final step, I estimate the distribution of the unobserved user heterogeneity term $F_y(\cdot)$ from the between-user variation. I use the estimated distribution to estimate $E[y_u]$, which governs counterfactual bidding. I discuss these steps in detail below.

Step 1: Isolate y_u

The first step holds y_u constant by selecting the appropriate the subset of data. This sample yields the within-user variation in x_{it} that feeds the estimator in Step 2. I choose the subset of frequent users who appear in more than T auctions, where T is sufficiently large that a user’s maximum bid approximates the upper limit of its support.⁸ I rank user u by his highest observed bid $b_{max}(u)$. I then isolate the unobserved heterogeneity component by choosing the set of users $u_{max} \in U_{max}$ such that $b_{max}(u_{max}) = \max_u b_{max}(u)$.⁹ I estimate $F_x(\cdot)$ on the U_{max} sample to hold $y_u = Y_{max}$ constant. This is the sample equivalent of fixing $y = \bar{Y}$ and examining the set of bids generated by $\beta(x; \bar{Y})$. In the future, I may estimate $F_x(\cdot)$ on the entire set of users who appear more than T times. However, this compounds the computational demands of Step 2, which must be computed for every realization of Y_u .

Step 2: Estimate $F_x(\cdot)$ and δ

⁸In application, T is chosen to be 400 or 500 auctions chosen at the publisher level.

⁹I can choose another set of users U' as long as their corresponding $Y_u \geq E[y_y]$. This ensures that the portion of the x distribution for which $x \geq \frac{r}{E[Y]}$ is identified, since this is used for the counterfactual.

I use maximum likelihood to estimate the idiosyncratic taste distribution given the maximum realized unobserved user heterogeneity term $F_x(\cdot; Y_{max})$ and the bid gap probability δ .

As discussed, I use the distribution of valuations $F_x(\cdot)$ as a starting point to estimate the distribution of bids. The distribution of X is not identified below $\frac{r}{Y_{max}}$, since bidders only bid when $x_{it}Y_{max} \geq r$. I separately estimate the censoring probability $F_x\left(\frac{r}{Y_{max}}\right)$ and the non-censored $F_x\left(\cdot|x \geq \frac{r}{Y_{max}}\right)$ distribution given by

$$F_x\left(x|x \geq \frac{r}{Y_{max}}\right) = \frac{F_x(x) - F_x\left(\frac{r}{Y_{max}}\right)}{1 - F_x\left(\frac{r}{Y_{max}}\right)}$$

I restrict $F_x\left(\cdot|x \geq \frac{r}{Y_{max}}\right)$ to a family of distributions parametrized by Θ . In principle, we can nonparametrically estimate $F_x\left(\cdot|x \geq \frac{r}{Y_{max}}\right)$ by choosing Θ to be a space of sieve functions (Brendstrup and Paarsch, 2006). In practice, Θ is the family of re-scaled Beta distributions. The Beta distributions use only two parameters, but flexibly describe many standard distribution shapes. A third, scaling parameter θ_3 ensures that non-censored valuations lie on the support $[r, r + \theta_3]$.

The likelihood function is constrained by the fact that the highest observed bid B_{max} must fall below the maximum bid $\beta(X_{max}; Y_{max}, \delta, \theta)$ as a function of θ and δ . As Athey and Haile (2002) point out, first-price bidders coordinate on the same maximum bid since a bidder is worse off whenever he outbids it. First-price auction estimates can therefore be sensitive to the constraint $B_{max} \leq \beta(X_{max}; Y_{max}, \delta, \theta)$. In the estimator, the scaling parameter θ_3 is an implicit parameter in the estimation that ensures that $B_{max} = \beta(X_{max}; Y_{max}, \delta, \theta)$.

The likelihood function combines n symmetric real-time bidders and m asymmetric offline bidders. The former have bid distribution $G(\cdot)$ and the latter bid $B = [B_1 \dots B_m]$ with probability $\alpha = [\alpha_1 \dots \alpha_m]$. I assume that the number of potential real-time bidders n is constant conditional on observed characteristics. I set n to be the total number of real-time bidders who ever submit an observed bid conditional on observed characteristics. The density of the top two bids $(b_k^{n-n:n}, b_j^{n:n})$ takes three main forms. If we observe no bids (\emptyset, \emptyset) , then the density is given by

$$f(\emptyset, \emptyset; r) = G(r)^n \prod_i^m (1 - \alpha_i)$$

If we observe only the highest bid y_j , then the density is given by

$$f(\emptyset, b_j^{n:n}; r) = \begin{cases} G(r)^n \alpha_j \prod_{i \neq j}^m (1 - \alpha_i) & j \in \text{Offline} \\ nG(r)^{n-1} g(b_j^{n:n}) \prod_i^m (1 - \alpha_i) & j \in RT \end{cases}$$

If we observe both the highest $b_j^{n:n}$ and second-highest bids x_k , then the density is given by

$$f(b_k^{n-n:n}, b_j^{n:n}; r) = \begin{cases} G(b_k^{n-n:n})^n \prod_{i \neq j, k}^m (1 - \alpha_i)^{I[B_i > b_k^{n-n:n}]} \cdot \alpha_j \alpha_k & k, j \in \text{Offline} \\ n(n-1) G(b_k^{n-n:n})^{n-2} \prod_i^m (1 - \alpha_i)^{I[B_i > b_k^{n-n:n}]} & k, j \in RT \\ g(b_k^{n-n:n}) g(b_j^{n:n}) & \\ nG(b_k^{n-n:n})^{n-1} \prod_{i \neq j}^m (1 - \alpha_i)^{I[B_i > b_k^{n-n:n}]} \cdot g(b_k^{n-n:n}) \alpha_j & k \in RT, j \in \text{Offline} \\ nG(b_k^{n-n:n})^{n-1} \prod_i^m (1 - \alpha_i)^{I[B_i > b_k^{n-n:n}]} \cdot \alpha_k g(b_j^{n:n}) & k \in \text{Offline}, j \in RT \end{cases}$$

Note that we treat the offline bids and the real-time bids as independent in the density function. This is valid because we have fixed $y = Y_{max}$ and we assume in Assumption 5.3.1 that x_{it} and the offline bidder's y_{iu} are independent. Evaluating the likelihood function is computationally intensive, so I use a grid search to find the $\hat{\theta}$ and $\hat{\delta}$ that maximize the likelihood.

The estimator in Step 2 has two components: a) estimate the offline bidding probability α and the real-time bidder's censoring probability $F_x\left(\frac{r}{Y_{max}}\right)$; and b) derive the bid distribution $G(b; Y_{max}, \theta, \delta)$ as a function of θ and δ .

Step 2A: Estimate α and $F_x\left(\frac{r}{Y_{max}}\right)$

Since Step 2B is computationally intensive, we wish to estimate as much outside that step as possible. These include the offline bid probabilities α and the real-time censoring probability $F_x\left(\frac{r}{Y_{max}}\right) = G(r; Y_{max})$. In principle, we can estimate the α_j components from using $\alpha_j = \frac{\lambda_j}{1 + \lambda_j}$ where $\lambda_j \equiv \frac{f(\emptyset, b_j^{n:n} = B_j; r)}{f(\emptyset, \emptyset; r)}$. In practice, I use the offline bidder estimates from section 5.2.2. Then, we can use the estimates to determine $F(r) = \left(\frac{f(\emptyset, \emptyset; r)}{\prod_j^m (1 - \alpha_j)}\right)^{\frac{1}{n}}$.

Step 2B: Derive $G(b; Y_{max}, \delta, \theta)$

To derive the distribution of $G(b; Y_{max}, \delta, \theta)$, the key is to determine the optimal bid function $\beta(x; Y_{max}, \delta, \theta)$. This bid function is monotonically increasing (Proposition 3) and defines the inverse $\eta(b; Y_{max}, \delta, \theta)$. $G(b)$ is then given by:

$$G(b) = F\left(\frac{r}{Y_{max}}\right) + \left(1 - F\left(\frac{r}{Y_{max}}\right)\right) \left((1 - \delta) G^* \left(b | b \geq \frac{r}{Y_{max}}; Y_{max}, \delta, \theta \right) + \delta G^{gap}(b) \right)$$

where

$$G^* \left(b \mid b \geq \frac{r}{Y_{max}}; Y_{max}, \delta, \theta \right) = F_x \left(\eta(b; Y_{max}, \delta, \theta) \mid x \geq \frac{r}{Y_{max}}; Y_{max}, \delta, \theta \right)$$

As shown in Proposition 3, we can calculate $\beta(x; Y_{max}, \delta, \theta)$ when $\delta = 0$ given the offline bids B . These gaps may overlap, so the number of gaps is weakly less than the number of offline bids. As in section 3, multiple equilibria can exist when the gaps overlap. My bid function algorithm (in the appendix) selects an equilibrium by moving up through the ordered B and choosing the gap with overlap from the highest possible component. When $\delta > 0$, a real-time bidder's chance of winning by outbidding B increases by the mass of real-time competitor bids in the gap. For simplicity, I assume that $G^{gap}(b)$ is uniform over the total support of the gaps. When $\delta > 0$, the size and arrangement of the gaps can change. I solve for $\beta(x; Y_{max}, \delta, \theta)$ using a fixed point algorithm. For more details on the bid function algorithm, see the appendix.

Step 3: Estimate $E[y_u]$

Given $\beta(X_{max}; Y_u, \hat{\delta}, \hat{\theta})$, I back out Y_u primarily from the maximum of the frequent user's bid support $b_{max}(u)$ (see the proof of Proposition 7). Towards, this we need to verify that $\beta(\bar{X}; Y_u)$ is strictly increasing in Y_u given the estimated parameters. Failing this, Proposition 7 explains that we will need to make use of another bid quantile to pin down Y_u in cases. The variation in Y_u between frequent users tells us the distribution of y_u . However, we are only interested in the mean of y_u for the counterfactual. If $F_y(\cdot)$ is censored by the reserve price, we could instead use the fact that $y \in [0, \frac{r}{X_{max}}]$ to bound $E[y_u]$. In application, I use the mean of the non-censored distribution since most of $F_y(\cdot)$ is observed.

In future work, I intend to compute confidence intervals for the real time estimates. I will need to show how the estimates depend on using Y_{max} to approximate \bar{Y} . The nonparametric bootstrap and subsampling are inappropriate for inference on an extremum estimator. We would need to restrict the distribution of y_u to be parametric in order to use the parametric bootstrap instead. Setting this aside, we can use the nonparametric bootstrap to compute confidence intervals for the estimates in Step 2. The estimates should be precise given that they typically use several thousand observations.

6 Results

I first present the structural auction estimates of the bidder value distributions. I use these estimates to simulate the counterfactual market under tracking restrictions to evaluate the effect of the privacy policy. Overall, my results suggest that publishers are unambiguously worse off under tracking restrictions. Though advertisers are worse off overall, some advertisers are better off because the resulting market is less competitive. My estimates suggest that a tracking ban would devastate the industry: publisher revenue falls by 38.5% and advertiser surplus falls by 45.5%.

I estimate the model for a sample of American users visiting the top three publishers in revenue terms on the exchange. These three publishers account for over 100 million auctions and half of all revenues arising from American users in the data. I intend to extend the analysis to additional publishers in the future.

6.1 Structural Model Estimate

6.1.1 Offline Bidders

I estimate the offline bidders individually. I split the data by publisher and ad form (ad dimensions). Given that the estimator is unreliable when the offline bidder wins infrequently, I only estimate the model on offline bidders who submit over 1,000 winning bids.¹⁰ To condition on observables, I define an auction to be eligible if the offline bidder ever bids on an auction with the same observable characteristics.¹¹

Table 2 summarizes the estimates of 184 offline bids on the top three publishers. The model parameters include the bidder targeting probability and conditional bidding probability. The targeting probability describes the proportion of users in the offline bidder’s audience who visit the publisher. The mean and median targeting probabilities are 53% and 63% respectively. The 25th quantile is 9%, so few bidders seem to finely target eligible users. Table 2 also shows the correlation between the targeting of the offline bidder and the bidders who outbid it. On average, the targeting correlation is positive though it is frequently negative as well. This shows the importance of allowing for arbitrary targeting correlation in the model. The condi-

¹⁰This cut-off could be modified. To be conservative in calculating publisher losses, I assume the offline bidders with less than 1,000 winning bids are not targeting ($\tau = 1$) and attribute their infrequent bids to a low bidding probability α . Since even 1,000 winning bids typically generate less than \$1 in revenues, this assumption has little bearing on overall revenues.

¹¹I presently define the set of observable characteristics to be the intersection of the user’s browser, state and Designated Marketing Area evaluated for every of the hour of the week. I am concerned however that the hour-level criterion is too narrowly defined. I intend to relax this in the future and to intersect all these variables over the entire data week. I would then exclude the hours during which the bidder does not bid, to allow for entrance and exit.

	Mean	Min.	Q=0.25	Median	Q=0.75	Max.
Targeting Probability	53.1%	0.3%	9.3%	62.5%	84.1%	100.0%
Targeting Correlation	11.6%	-36.2%	-6.3%	9.0%	24.1%	97.3%
Conditional Bidding Probability	35.9%	0.2%	18.2%	32.5%	49.6%	92.5%

Table 2: A summary of the offline bidder estimates for 184 frequent offline bids on the top three publishers. The targeting correlation describes the correlation between the offline bids target audience and that of the bidders who outbid him.

tional bidding probability describes how often the offline bidder bids on users who satisfy his targeting criteria. The purpose of the conditional bidding probability is to smooth the offline bidder’s ad budget. The conditional bidding probability estimates range from 0.2% to 93% with a mean of 36%. This shows that bidders selectively bid on observable users, though the mean of 36% is actually fairly high given the tremendous supply of impressions.

6.1.2 Real-Time Bidders

I estimate the distribution of valuations for the symmetric real-time bidders. This distribution is composed of the distributions of the idiosyncratic and user-level unobserved heterogeneity terms x_{it} and y_u . I estimate the distribution of the idiosyncratic utility term $F_x(\cdot)$ separately for each publisher’s combination of ad form and reserve price. I only estimate the model when these cuts of the data include over 750 real-time bids among the chosen set of frequent users with common Y_u realizations. For computational reasons, I currently estimate $F_x(\cdot)$ for a single Y_u realization though I intend to extend this estimation to all Y_u later.

Table 3 shows the estimates for the symmetric real-time bidder idiosyncratic taste distribution $F_x(\cdot)$ for various cuts of the data. I estimate the non-censored part of the distribution $F_x(\cdot | xY_u \geq r)$, by fitting it to a beta distribution parametrized by (θ_1, θ_2) . I scale this distribution to have support on $[r, \theta_3 + r]$ using the scaling parameter θ_3 . I also calculate the probability δ with which a bidder bids in the dominated ‘bid gap.’ The estimated $\hat{\delta}$ is 24% on average across the 18 density estimates. Figure 10 illustrates the $F_x(\cdot | xY_u \geq r)$ density plot for publisher A’s three different ad forms as indicated in Table 3. Figure 10 shows the two predominant shapes that all the fitted distributions take. When the fitted beta distribution has the form $\theta_1 > 1$ and $\theta_2 < 1$ (see Figure 10, Cut #1 and #3), the distribution has a reverse ‘L’ shape that puts weight on the right tail of the support. When both $\theta_1, \theta_2 < 1$ (see Figure 10, Cut #2), the fitted beta distribution has a ‘U’ shape.

The real-time bidder estimator has severable notable features. For example,

	θ_1	θ_2	θ_3	δ
Publisher A				
Ad 1 X r 1	8	0.01	6	0.1
Ad 2 X r 1	0.01	0.1	4	0.15
Ad 3 X r 1	6	0.1	5.9	0.3
Publisher B				
Ad 1 X r 1	2.5	0.1	7.6	0.4
Ad 1 X r 2	4	0.1	6.4	0.25
Ad 1 X r 3	0.1	0.01	5	0.1
Ad 1 X r 4	4	0.01	19.1	0
Ad 2 X r 1	1.8	0.3	8.9	0.4
Ad 2 X r 2	2.5	0.3	7.9	0.25
Ad 2 X r 3	1.8	0.3	11.2	0.05
Ad 2 X r 4	2.5	0.1	17.3	0
Publisher C				
Ad 1 X r 1	0.9	0.3	4.8	0.2
Ad 1 X r 2	8	0.01	6.5	0.4
Ad 2 X r 1	0.01	0.1	3	0.15
Ad 2 X r 2	0.7	0.01	2	0.5
Ad 2 X r 3	8	0.01	0.9	0.05
Ad 3 X r 1	2	0.01	8.6	0.5
Ad 3 X r 2	1.8	0.01	8.1	0.5

Table 3: The estimates of the real-time bidders' value distribution for several ad forms ('Ad') and reserve prices ('r') for the top three publishers. The beta distribution parameters are defined by (θ_1, θ_2) with scaling parameter θ_3 and the probability of bids in the gap is denoted δ .

the real-time bidder’s optimal bid function currently ignores infrequent offline bids ($\Pr[B_i] < 0.01$) since these should have little bearing on optimal bidding behavior. Since each set of beta distribution parameters implies a maximum possible bid, I rescale the bid support using parameter θ_3 to allow all beta parameters (θ_1, θ_2) to be considered as candidate estimates. θ_3 is set implicitly to ensure the maximum observed bid matches the theoretical maximum bid given the model parameters, $b_{max} = \beta(X_{max}; Y_{max}, \theta, \delta)$. Following Guerre et al. (2000), I rescale bids by $b^\dagger = \sqrt{b - r}$ to avoid the problem that the bid density will otherwise approach infinity for bids near the reserve price. For computational reasons, the estimator currently uses grid search. Despite expanding this grid, some estimated parameters remain on the boundary of the grid, which I will further expand if required after extending the estimator from a single Y_u to all realizations of y_u .

6.2 Privacy Counterfactual

I use the structural auction results to simulate counterfactual auctions in which bidders cannot track users. Corollary 6 outlines the counterfactual equilibrium in these simulations.

I compare the average advertiser and publisher outcomes in the status quo and policy simulations. I simulate 100,000 auction draws for each combination of ad form and reserve price in each of the top three publishers. Since bidders can no longer condition on user tracking information, the real-time and offline bids are now independent. We can then use the imputed counterfactual estimates without worrying about dependence based on unobservables. However, there could still be dependence based on observables. I currently do not allow this for all observable characteristics. In the future, I could draw users at random from the sample and assign bids based on their observable characteristics.

As welfare measures, I examine the advertiser’s bidder surplus and the publisher’s revenue. The offline bidder’s surplus is the difference between his bid (which equals his valuation) and the price paid. The real-time bidder’s surplus is her bid markdown $v - b$, and this calculation requires inverting her bid with the estimated inverse bid function $v = \eta(b; \hat{\theta})$. The publisher’s measure of producer surplus is more nuanced as it depends on the outside value of an impression on the unobserved guaranteed market. Moreover, the publisher’s outside value of a marginal impression sent to the exchange declines in impressions as demand in the guaranteed market falls and the publisher needs less flexibility to fulfill its guaranteed contracts. Currently, I use the publisher revenues as my measure of publisher surplus as an upper bound estimate. Since the reserve price provides an upper bound for the advertiser’s outside option,

	<i>Opt-Out</i>	<i>Opt-In</i>	<i>Ban</i>
Total Industry Surplus	-4.3%	-39.1%	-43.5%
Publisher Revenues	-3.8%	-34.6%	-38.5%
Total Advertiser Surplus	-4.5%	-40.9%	-45.5%
Offline Bidder	-4.7%	-42.6%	-47.3%
Real-Time Bidder	-4.5%	-40.7%	-45.3%

Table 4: The overall impact of privacy policy on the online display ad auction participants. The opt-in, opt-out, and tracking ban policies envision that the proportion of users who can not be tracked is 10%, 90%, and 100% respectively.

subtracting the reserve price from the sale price provides a lower bound for the level of producer surplus in the status quo. If we are willing to assume the guaranteed market weakly falls under the privacy policy, subtracting the status quo reserve price also provides a lower bound for the counterfactual publisher surplus. In the next revision, I will include this measure of publisher surplus. Since the decrease in revenues is the same in both cases, the lower bound estimates of surplus would have a larger percentage decline in welfare. I will also differentiate the proportion of revenue which accrue to the publisher and the commission collected by the exchange.

I consider the effect of opt-in, opt-out and tracking ban policies. As described in Section 2.3, I project that these policies correspond to Do Not Track (DNT) user percentages of 10%, 90%, and 100%. In the baseline model, I assume users opt out at random and that the buyers and sellers make no adjustments like changing reserve price. The policy’s impact is then proportional to the number of users who cannot be tracked. That is, for the status quo outcome w_{SQ} , counterfactual outcome w_{CF} and DNT user share γ_{DNT} , the impact is given by

$$\begin{aligned}
 \text{Impact}(\gamma_{DNT}) &= w_{SQ} - [\gamma_{DNT}w_{CF} + (1 - \gamma_{DNT})w_{SQ}] \\
 &= \gamma_{DNT}(w_{SQ} - w_{CF})
 \end{aligned}$$

Though the welfare impact scales linearly with the Do Not Track population percentage in the base line model, these estimates provide policymakers a yardstick to gage their policy aggressiveness.

The overall results suggest large losses for publishers and advertisers. Table 4 shows the welfare losses due to the opt-out, opt-in, and tracking ban policies for the different agents. These policies correspond to a drop in total surplus of 4.4%, 39.2%, and 43.5%. The burden is shared quite equally since publisher revenues falls 3.9%, 34.6%, and 38.5% advertiser surplus falls 4.6%, 40.9%, and 45.5%. Within the advertisers, offline and real-time bidder surplus falls about equally by -47.3%

	<i>Total</i>	<i>Publisher</i>	<i>Advertiser Surplus</i>		
	<i>Surplus</i>	<i>Revenues</i>	<i>Total</i>	<i>Offline</i>	<i>Real-Time</i>
Publisher A	-42.4%	-48.5%	-39.9%	-59.1%	-36.4%
Publisher B	-41.9%	-12.1%	-46.3%	-53.2%	-46.2%
Publisher C	-48.9%	-46.1%	-54.1%	-32.1%	-66.7%

Table 5: The impact of a tracking ban policy on publishers and advertisers for each of the top three publishers.

and -45.3% respectively under a tracking ban. The offline bidder surplus falls as some bidders exit the market and this offsets the benefit to the remaining bidders of reduced competition. While the total losses to publishers and advertisers seem large, they are on the same order of magnitude as Goldfarb and Tucker (2011)’s analysis of Europe’s opt-out policy that found a 65% decrease in ad effectiveness. To give a sense of the economic magnitudes, I perform a back-of-the-envelope calculation for the impact on publisher revenues. If auctions represent 20% of the industry’s \$6.8 billion in revenues, the three policies imply losses of \$52, \$471 and \$523 million respectively. This is a lower bound impact estimate since it does not include the impact on advertisers, the guaranteed contract side of the market or the compliance costs to the industry.

The results vary when we consider the three publishers separately. Under a tracking ban, Table 5 shows the revenues effect by publisher varies from -12.1% to -41.9%. The effect on advertiser surplus varies from -39.9% to -54.1%.

Note that the existing results will need to undergo some further refinements. The policy impact estimates are sensitive to the censoring probability $F(r)$ estimate. $F(r)$ is typically very close to 1. So, the real-time estimates and the simulated auction outcomes are sensitive to the $F(r)$ estimate. For the time being, I simulate the status quo industry welfare as a more reasonable baseline. At present, I compute the real-time surplus with respect to the subsample of users with a single unobserved heterogeneity realization $Y'_u > E[y_u]$. I can extend this to all frequent users in the future. In the meantime, we can think of the reduction real-time bidder’s surplus as an upper bound. For the time being, I also choose $E[y_u]$ to be the Y_u^* such that $\beta(X_{max}; Y_u^*) = E_y[\beta(X_{max}; y_u)]$. Later, I will update this to be the true mean estimate $E[y_u]$.

7 Extensions

In Section 7, I propose four extensions to my analysis. The first corrects for opt-out user self-selection by weighing the counterfactuals by current opt-out user browsing behaviour. The next three describe market adjustments: 1) the seller reduces his reserve price; 2) the portion of real-time bidders is increased to reflect current levels; and 3) the advertisers shift their ad budgets to websites where they can find their target audience. I describe the extensions in detail below.

7.1 Do Not Track user self-selection

The first extension corrects for the self-selection of Do Not Track (DNT) users in an opt-out policy. To date, I assume that users opt out at random, so that the subpopulation that opts out is identical to its complementary subpopulation. An opt-out policy would affect not only the proportion of DNT users but also the composition of users who self-select into DNT. An Internet Advertising Bureau study found that the segment of privacy concerned users were disproportionately lower income households and single (Deighton and Quelch, 2009).

To correct for self-selection, I intend to weight the policy’s impact to reflect the browsing patterns of existing DNT users. My data feature a tiny fraction of users (less than 0.1%) who opt out of tracking. The fraction is too small to reliably estimate the structural model. However, opt-out users systematically differ in the types of content they access. In this extension, I will account for self-selection by choosing users for the opt-out counterfactual whose browsing behavior resembles that of the current opt-out population. The extension indicates how publishers with different content types may be affected differently by an opt-out policy and how the aggregate market impact would differ. For instance, if opt-out users are more likely to visit low-value websites (e.g. file-sharing websites), the total impact of the opt-out policy would be blunted.

7.2 Seller adjustment

In this extension, I propose that the sellers reduce their reserve price to mitigate the policy’s effect on their revenues. To date, I assume that the sellers maintain the status quo reserve price in the privacy counterfactual despite the fact that revenues fall by 6% under a tracking ban. The assumption is realistic if the seller’s reserve price reflects the publisher’s outside value given by his guaranteed contracts and if

these contracts do not change.¹² However, I do not know how the guaranteed side of the display ad market will respond to privacy policy. Beyond the reserve price, I have no data on the other side of the market to inform a more general equilibrium view.

Instead, I propose that sellers reduce the reserve price to maintain the status quo ad exchange sell-through rate. That is, sellers wish to maintain the fraction of inventory that they sell on the exchange. This assumption delivers an estimate for the alternate counterfactual's impact. We can think of this exercise as indicating the order of magnitude of the effect if the supply side adjusts. If advertisers can not track users, the guaranteed delivery market's equilibrium price could lie somewhere between remaining constant and falling by as much as the auction market. Then the two estimates as bounds on the publisher's revenue loss. The constant sell-through rate counterfactual is also more applicable when publishers view the fraction of inventory they sell on the exchange to be remnant inventory. Publishers have remnant inventory if they struggle to find enough guaranteed buyers or if they were purposely leave a fraction of their inventory unsold to have some flexibility to meet their guaranteed contracts. One problem with reducing the reserve price is that the status quo bids below the reserve price are not identified, so I cannot predict bidder entry at a lower reserve price.

7.3 Increased real-time bidder prevalence

The third extension adjusts the proportion of real-time bids to reflect the higher level of real-time bidders in modern online display advertising auctions. Real-time bidders account for about 10% of sold impressions in the data. However, Google's DoubleClick ad exchange suggests that the proportion of real-time bidders grew rapidly since the time my data was sampled. By 2011, real-time bidders represented 68% of the auction's bids (Google, 2011). This extension would boost the share of real-time bidders to reflect their rise in prominence. In the process, this would also reduce the distortion created by offline bidders on the real-time bid function. Nonetheless, I can not know how the composition of the real-time entrants and the offline bidders who exit might differ.

¹²Alternately, we can view that publishers submit the optimal reserve price and back out the outside value that rationalizes the reserve price. We could then recompute the seller's optimal reserve price in the counterfactual market. However, the publisher market is so competitive that it is unclear if publishers have sufficient market power to price their inventory above their outside value.

7.4 Offline bidder adjustment

This extension allows the demand side to readjust by increasing the advertiser’s bidding probability on sites where it can find its target audience. For example, suppose an offline advertiser values its target audience at \$1 and bids on two websites where its target audience appears with probability 40% and 60%. Its counterfactual bids would therefore be \$.40 and \$.60. If all the publisher maintain a reserve price of \$.50, then the offline bidder would exit the first publisher’s auctions since only the second bid exceeds the reserve price. Instead, I propose that advertisers expend their ad budgets on sites where their target audience is better represented. This too would mitigate the collapse in the market and lead to interesting compositional effects. In particular, specific interest websites may do better than general interest publishers.

8 Conclusion

This paper estimates the previously unquantified impact of privacy policy on the auction market for online display advertising. I measure the responsiveness of publisher revenues and advertiser surplus in order to guide regulators in their policy selection. An opt-out policy allows privacy-conscious users to avoid tracking. If these users represent a minority of user (i.e., 10%), my estimates suggest that publisher revenues fall by 3.9% and advertiser surplus falls by 4.6%. The losses are substantial though perhaps manageable in an industry that has grown 20% annually over the past decade. Alternatives like an opt-in policy and an outright ban would nearly or entirely or eliminate the portion of users who can be tracked. In these cases, publisher revenues and advertiser surplus would fall by half or more, as many advertisers would exit the market. This paper solely examines the auction market for online advertising. Back-of-the-envelope calculations suggest that the publisher’s revenue losses alone in an opt-out, opt-in or tracking ban policies would be \$471, \$52, or \$523 million respectively. These costs must be weighed against the benefit to consumers of increased privacy protection net of the cost to consumers of worse ad targeting. This consumer welfare calculation is a challenging and interesting question for future research.

Beyond answering the policy question, this paper makes some additional methodological contributions. I demonstrate two novel approaches to auction models with unobserved heterogeneity where we observe a panel of users in auctions and the unobserved heterogeneity is at the user level. In the first model, the unobserved heterogeneity utility term takes a binary form but can accommodate bidder-to-user matching. This appears in a setting where asymmetric bidders bid by proxy on users using binary auction criteria. In the second model, the unobserved user heterogeneity

appears as a scalar term in the bidder’s utility. I extend unobserved heterogeneity models to auctions where the reserve price binds, but where we observe a long panel of users. I identify and estimate the model by exploiting both the between-user and within-user variation in bids. In addition, I describe the equilibrium of a novel auction where some bidders play by first-price rules and others play by second-price rules. This novel setting necessitates fresh estimation approaches for the structural auction model that can accommodate discontinuous bid functions. Finally, I analyze a very large online display ad auction dataset, an industry which represent an extensive application of auctions.

References

- Abraham, I., S. Athey, M. Babaioff, and M. Grubb (2011). Peaches, lemons, and cookies: Designing auction markets with dispersed information.
- American Civil Liberties Union (2011, March 16). The state of online consumer privacy. Statement Before the Senate Commerce, Science and Transportation Committee.
- An, Y., Y. Hu, and M. Shum (2010). Estimating first-price auctions with an unknown number of bidders: A misclassification approach. *Journal of Econometrics* 157(2), 328 – 341.
- Angwin, J. (2010, July 30). The web’s new gold mine: Your secrets. Wall Street Journal.
- Angwin, J. (2012, June 18). Online tracking ramps up. Wall Street Journal.
- Athey, S. and P. Haile (2002). Identification of standard auction models. *Econometrica* 70(6), 2107–2140.
- Athey, S. and P. A. Haile (2007). Nonparametric approaches to auctions. In J. Heckman and E. Leamer (Eds.), *Handbook of Econometrics*, Volume 6 of *Handbook of Econometrics*, Chapter 60. Elsevier.
- Bachman, K. (2012, March 20). Study: Internet user adoption of DNT hard to predict internet users less conflicted about ad-supported free content. <http://www.adweek.com/>.

- Bailey, M. and A. Farahat (2012). *Essays in applied economics: evidence from natural experiments*, PhD Dissertation How Effective is Targeted Advertising? Stanford University.
- Beales, H. (2010). The value of behavioral targeting. Network Advertising Initiative.
- Brendstrup, B. and H. J. Paarsch (2006). Identification and estimation in sequential, asymmetric, english auctions. *Journal of Econometrics* 134(1), 69 – 94.
- Celis, L., G. Lewis, M. Mobius, and H. Nazerzadeh (2012, June 17). Buy-it-now or take-a-chance: Price discrimination through randomized auctions.
- Center for Digital Democracy (2012, August 22). Consumer, privacy, and child health organizations charge six children’s websites with unfair and deceptive marketing. Press Release.
- comScore (2011, February). The 2010 U.S. digital year in review: A recap of the year in digital marketing.
- comScore (2012, February). U.S. digital future in focus 2012. <http://www.comscore.com/>.
- Deighton, J. and J. Quelch (2009). Economic value of the advertising-supported internet ecosystem. *IAB Report*.
- D’Haultfoeuille, X. and P. Février (2010). Identification of mixture models using support variations.
- Digital Advertising Alliance (2012a, October 9). DAA statement on DNT browser settings. <http://www.businesswire.com/>.
- Digital Advertising Alliance (2012b). <http://www.youradchoices.com/>.
- European Data Protection Authorities (2010, June 24). Opt-out is not sufficient: European Data Protection Authorities clarify EU rules on online behavioural advertising. Article 29 Data Protection Working Party Press Release.
- Evans, D. (2009). The online advertising industry: Economics, evolution, and privacy. *The Journal of Economic Perspectives* 23(3), 37–60.
- Federal Trade Commission (2000). Privacy online: Fair information practices in the electronic marketplace, a report to congress. <http://www.ftc.gov/>.

- Federal Trade Commission (2012a, August 9). Google will pay \$22.5 million to settle FTC charges it misrepresented privacy assurances to users of Apple’s Safari Internet browser. Press Release.
- Federal Trade Commission (2012b). Protecting consumer privacy: Making sure companies keep their privacy promises to consumers.
- Fowler, A. (2012, May 17). Do Not Track gains more support around the web.
- Ghosh, A., P. McAfee, K. Papineni, and S. Vassilvitskii (2009). Bidding for representative allocations for display advertising. *Internet and Network Economics*, 208–219.
- Goldfarb, A. and C. Tucker (2011). Privacy regulation and online advertising. *Management Science* 57(1), 57–71.
- Google (2011, July). The arrival of real-time bidding and what it means for media buyers. White Paper.
- Guerre, E., I. Perrigne, and Q. Vuong (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica* 68(3), 525–574.
- Gummadi, R., P. Key, and A. Proutiere (2012). Repeated auctions under budget constraints: Optimal bidding strategies and equilibria. Microsoft Research Working Paper.
- Hahn, R. (2001). An assessment of the costs of proposed online privacy legislation. *AEI-Brookings Joint Center for Regulatory Studies*.
- Hendricks, K. and R. H. Porter (2007). *An Empirical Perspective on Auctions*, Volume 3 of *Handbook of Industrial Organization*, Chapter 32, pp. 2073–2143. Elsevier.
- Hoofnagle, C. J., A. Soltani, N. Good, D. J. Wambach, and M. D. Ayenson (2012). Behavioral advertising: the offer you cannot refuse. *Harvard Law & Policy Review* 6.
- Hu, Y., D. McAdams, and M. Shum (2009). Nonparametric identification of auction models with non-separable unobserved heterogeneity. *Economics Working Paper Archive*.
- Interactive Advertising Bureau (2011, July). IAB Legislative and Regulatory Affairs Update.

- Interactive Advertising Bureau (2012, April). 2011 IAB Internet Advertising Revenue Report.
- Krasnokutskaya, E. (2011). Identification and estimation of auction models with unobserved heterogeneity. *The Review of Economic Studies* 78(1), 293–327.
- Krishna, V. (2009). *Auction theory*. Academic press.
- Lamy, L. (2012). The econometrics of auctions with asymmetric anonymous bidders. *Journal of Econometrics* 167(1), 113 – 132.
- Leon, P., B. Ur, R. Balebako, L. Cranor, R. Shay, and Y. Wang (2011). Why Johnny can't opt out: A usability evaluation of tools to limit online behavioral advertising. Carnegie Mellon University.
- Levin, J. and P. Milgrom (2010). Online advertising: Heterogeneity and conflation in market design. *American Economic Review* 100(2), 603–607.
- Lindahl, G. (2012, June 28). Do we have statistics on "Do Not Track" usage, except from Firefox? Quora.
- Mahdian, M., A. Ghosh, P. McAfee, and S. Vassilvitskii (2012). To match or not to match: Economics of cookie matching in online advertising. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pp. 741–753. ACM.
- Mansour, Y., S. Muthukrishnan, and N. Nisan (2012). Doubleclick ad exchange auction. Computing Research Repository.
- Mattioli, D. (2012, August 23). On Orbitz, Mac users steered to pricier hotels. Wall Street Journal.
- McAfee, P., K. Papineni, and S. Vassilvitskii (2010). Maximally Representative Allocations for Guaranteed Delivery Advertising Campaigns. Yahoo! Research manuscript.
- McDonald, A. and L. Cranor (2009). The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society* 4, 543–897.
- McDonald, A. and L. Cranor (2010). Beliefs and behaviors: Internet users' understanding of behavioral advertising.
- McDonald, A., R. Reeder, P. Kelley, and L. Cranor (2009). A comparative study of online privacy policies and formats. In *Privacy Enhancing Technologies*, pp. 37–55. Springer.

- Morales, L. (2010, December 21). U.S. internet users ready to limit online tracking for ads. <http://www.gallup.com/>.
- Ohm, P. (2010). Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA Law Review* 57(6), 1701.
- Olejnuk, L., C. Castelluccia, and A. Janc (2012). Why Johnny can't browse in peace: On the uniqueness of web browsing history patterns. Privacy Enhancing Technologies Symposium.
- Paarsch, H. and H. Hong (2006). An introduction to the structural econometrics of auction data. *MIT Press Books 1*.
- Privacy Rights Clearinghouse (2010, April 14). The proliferation of online information brokers and reports of abuses of consumer privacy: Submission to the federal trade commission. <https://www.privacyrights.org/>.
- Pubmatic (2011). Audience selling for publishers I. www.pubmatic.com/.
- Steel, E. (2010, October 24). A web pioneer profiles users by name. <http://online.wsj.com/>.
- TechJournal (2012, June 28). Online behavioral tracking pervasive, Google privacy practices often violated. <http://www.techjournal.org/>.
- Tucker, C. (2011). Social networks, personalized advertising, and privacy controls.
- Tucker, C. E. (2012). The economics of advertising and privacy. *International Journal of Industrial Organization* 30(3), 326 – 329.
- Turow, J., C. Hoofnagle, D. Mulligan, and N. Good (2007). The Federal Trade Commission and consumer privacy in the coming decade. *I/S: A Journal of Law and Policy for the Information Society* 3, 723.
- Turow, J., J. King, C. Hoofnagle, A. Bleakley, and M. Hennessy (2009). Americans reject tailored advertising and three activities that enable it. Departmental Papers (Annenberg School for Communication).
- White House (2012). Consumer data privacy in a networked world: A framework for protecting privacy and promoting innovation in the digital economy.

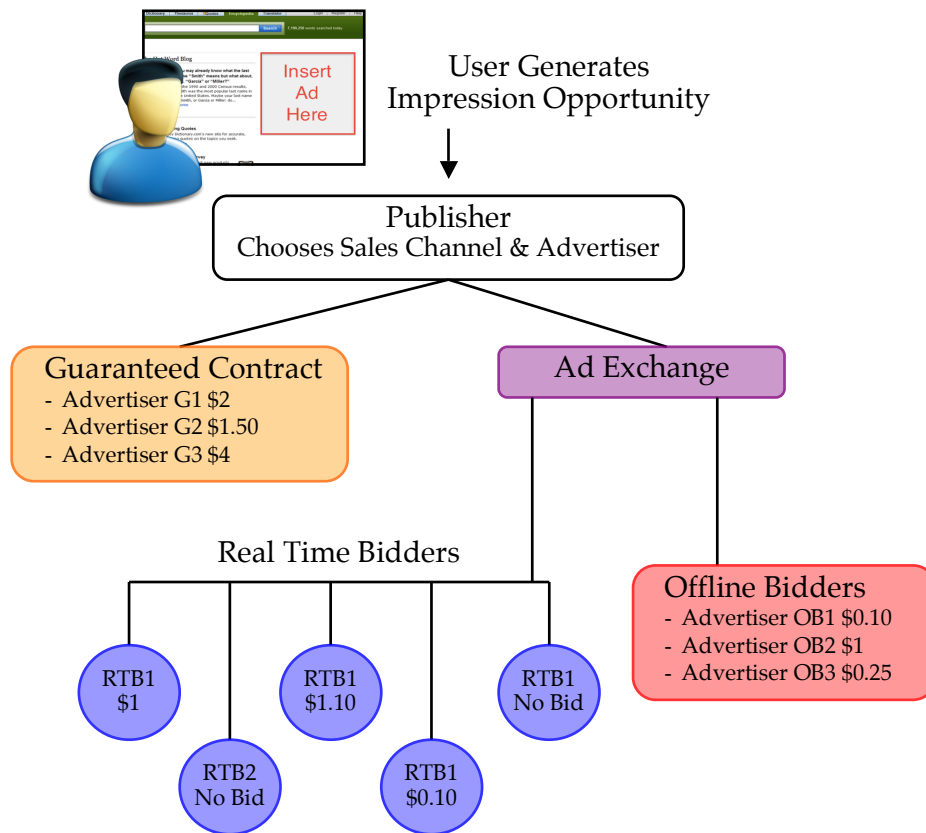


Figure 1: From the publisher’s point of view, a user generates an impression opportunity. The publisher chooses to sell to a standing customer by fulfilling a guaranteed contract or sell on an ad exchange. The ad exchange runs an auction with offline bidders that pre-specify their bids and real-time bidders that bid on each impression in real time.

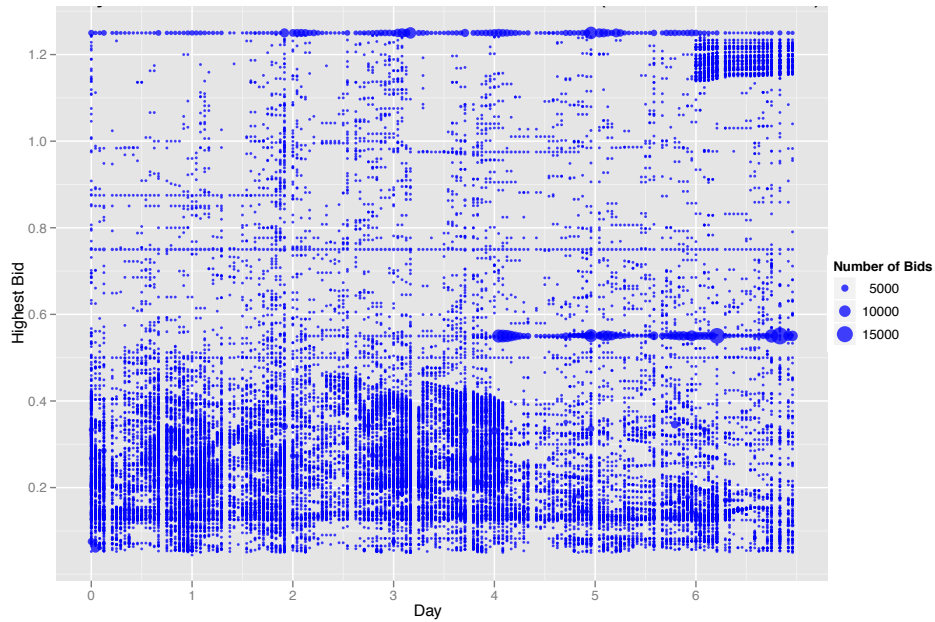


Figure 2: The hourly evolution of one real-time bidder's observed highest bid over a week demonstrates high variance. Bids are aggregated hourly and larger dots indicate more frequent bids.

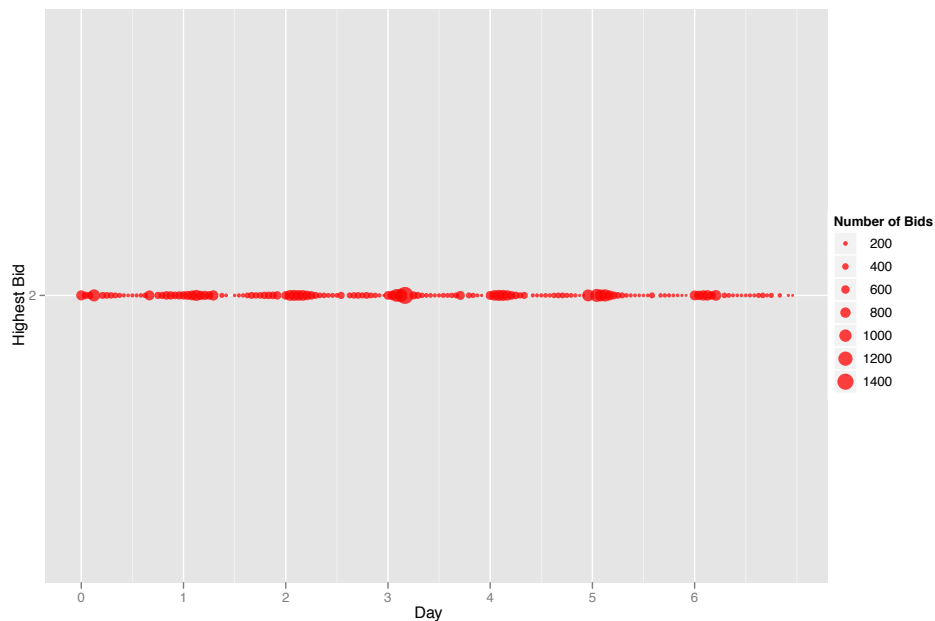


Figure 3: The hourly evolution of one offline bidder's observed highest bid over a week reveals that he left his fixed bid unchanged. Bids are aggregated hourly and larger dots indicate more frequent bids.

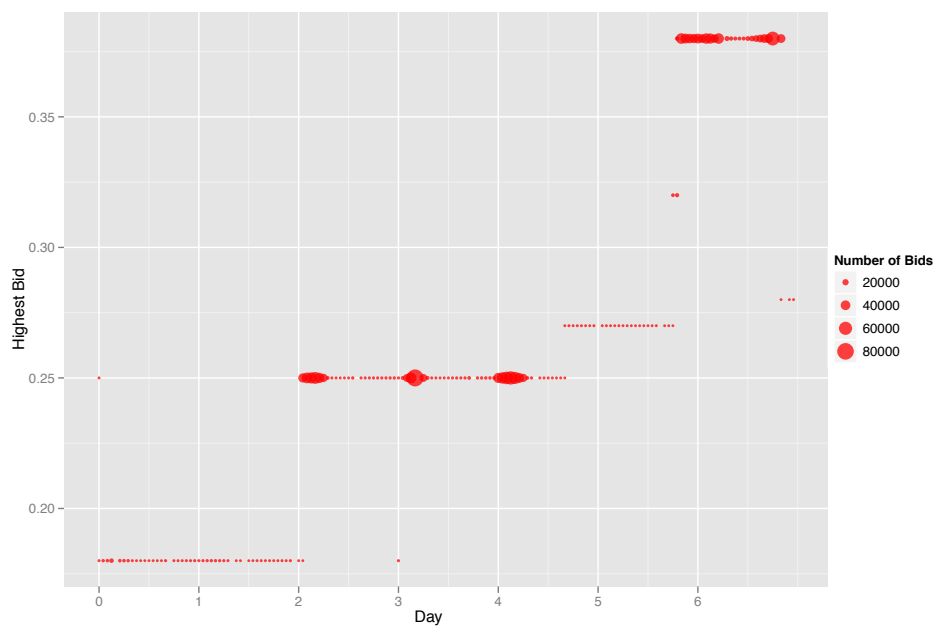


Figure 4: The hourly evolution of one offline bidder’s observed highest bid over a week reveals that she modified her fixed bid on a number of occasions. Bids are aggregated hourly and larger dots indicate more frequent bids.

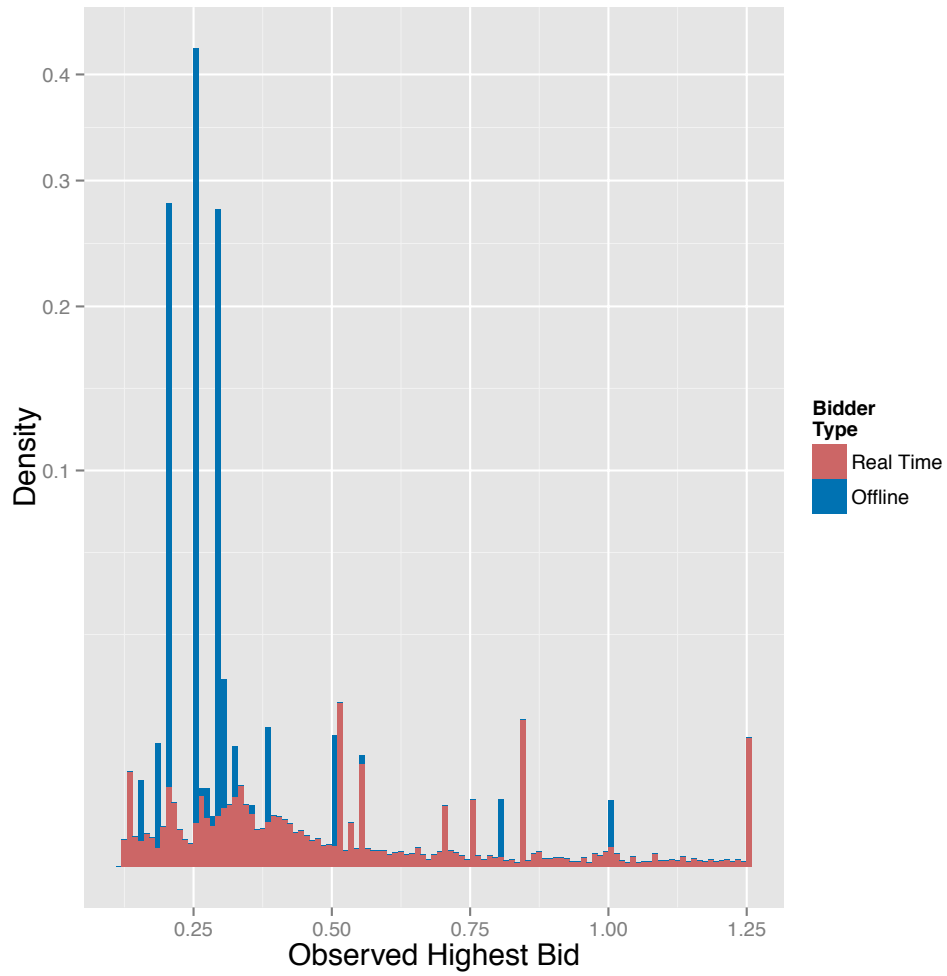


Figure 5: A histogram showing the density of observed highest bids for American users on a top publisher. The distinction between real time and offline bids shows they are characterized by discrete and continuous distribution respectively. Note that the y-axis uses a square root scale so that the real-time bids are visible given the very high probability on some offline bids.

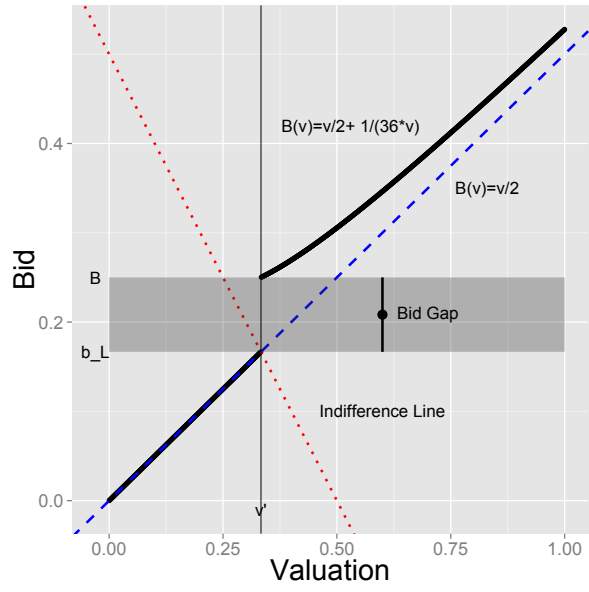


Figure 6: The optimal bid function when two uniform first-price bidders face an offline bidder with $\Pr[B] = \frac{1}{2} = \Pr[0]$. The dashed blue line represents the optimal bid function in the absence of the offline bidder. The dotted red line represents the indifference relation given by $u(x', b_L) = \lim_{\delta \rightarrow 0^+} u(x' + \delta, B)$. The interaction between the dotted and dashed lines determines the indifference valuation x' and $b_L = \beta(x')$. The gap in the bid support is highlighted in grey.

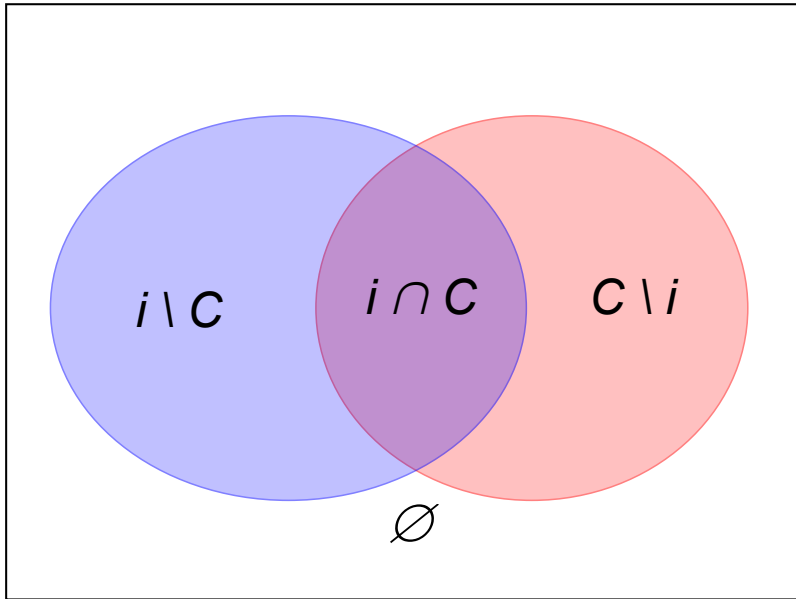


Figure 7: A Venn diagram illustrates the targeting of bidders i and its effective competitors C on the space of users.

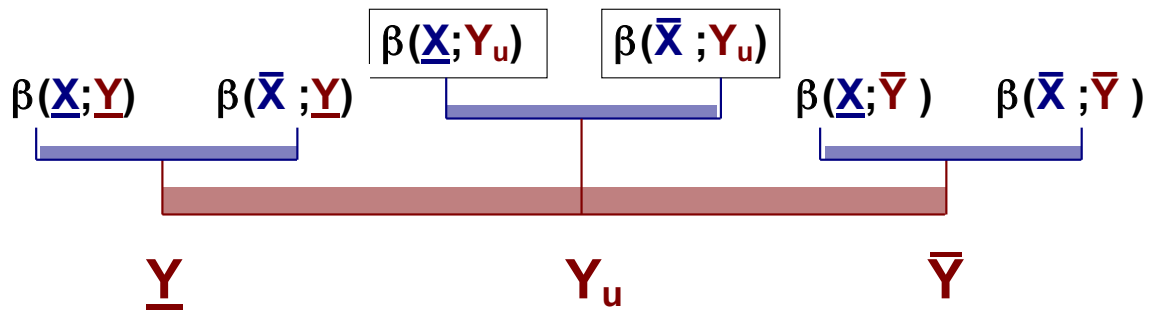


Figure 8: Demonstrates how the support of each user's bids corresponds to the underlying support of the idiosyncratic taste component x and the unobserved user heterogeneity term y in the real-time bidder's valuation.

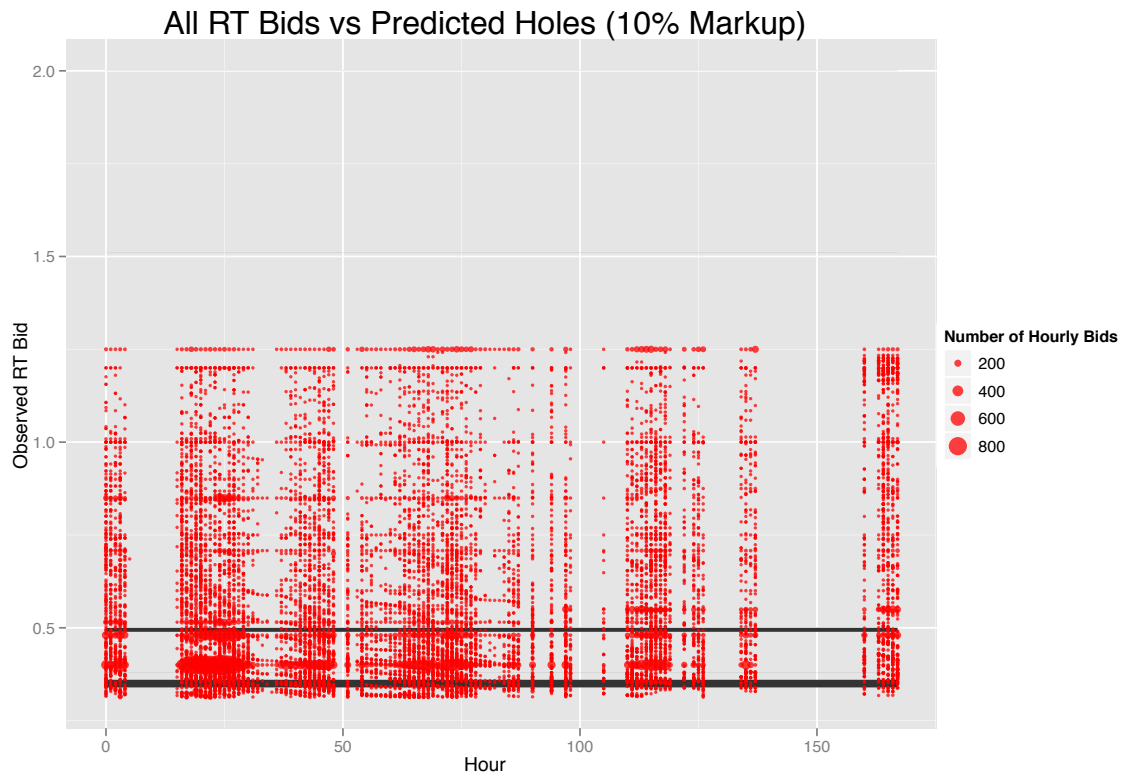


Figure 9: Illustrates the bids of several real-time bidders with respect to the predicted gaps in the bid function. The data is restricted to a single ad form (defined by its dimensions) for a single publisher and American users who appear more than 100 times. The bid gaps are represented by the dark gray lines and reflect a 10% markup in the bid function $b = 0.9v$. I will correct the time periods with missing data in a later revision.

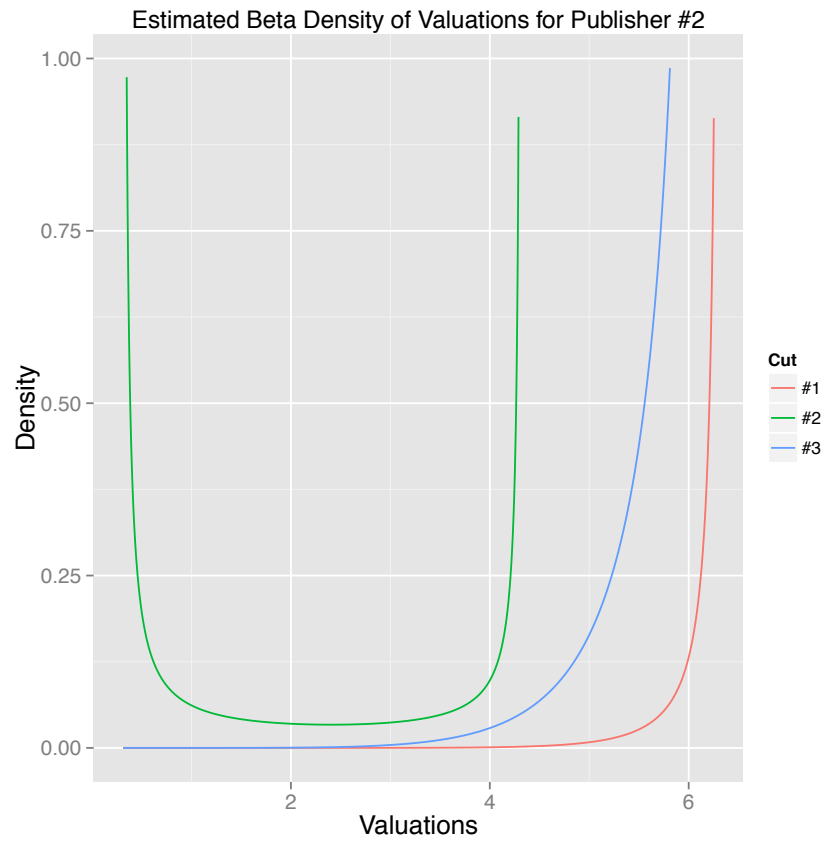


Figure 10: Real-time bidder value distribution estimates for Publisher A ad types 1-3 with beta distribution parameters (θ_1, θ_2) and rescaling parameter θ_3 given by #1 = (8, 0.01, 6), #2 = (0.01, 0.1, 4), and #3 = (6, 0.1, 5.9).

Appendix

A Proofs

Proposition 1

Proof. See Proposition 2.1 in Krishna (2009). □

Proposition 2

Proof. Part 1 holds because a first-price bidder i chooses the bid b , given his valuation x to solve

$$\max_{b \leq x} (x - b) \Pr [i \text{ wins} | b]$$

Thus, bidding above (below) her valuation ensures a negative (positive) expected payoff. For Part 2, we wish to show that i 's bids in an interval $(b_L, B]$ are strictly dominated for some b_L when it faces competition from a bidder who bids $B > r$ with positive probability ($\alpha = \Pr [B] > 0$). We consider the case where $\alpha < 1$, since that case is isomorphic to setting the reserve price at B . Write $\Pr [i \text{ wins} | b] = H(b) (1 - \alpha)^{I[B > b]} \frac{\alpha^{I[B = b]}}{2}$ where $H(b)$ denotes the distribution of competing first-price bidders (equals 1 if none exist).

Suppose that i 's valuation satisfies $x > B$. Then we choose $b_L = B - \varepsilon$ where ε must first satisfy $x - B > \varepsilon > 0$. Then, compare the payoffs for bid $B + \varepsilon$ and $b_L = B - \varepsilon$. b_L is dominated

$$\frac{u(b_L)}{u(B + \varepsilon)} = \frac{x - (B - \varepsilon)}{x - (B + \varepsilon)} \frac{H(B - \varepsilon) (1 - \alpha)}{H(B + \varepsilon)} < 1$$

for some ε since $\frac{H(B - \varepsilon)(1 - \alpha)}{H(B + \varepsilon)} < 1$ and $\lim_{\varepsilon \rightarrow 0} \frac{x - (B - \varepsilon)}{x - (B + \varepsilon)} = 1$. Moreover, this inequality holds for any $b \in [b_L, B]$ provided that ε satisfies

$$\frac{u(b_L)}{u(B + \varepsilon)} < \frac{x - (B - \varepsilon)}{x - (B + \varepsilon)} \frac{H(B) \frac{\alpha}{2}}{H(B + \varepsilon)} < 1$$

Now, suppose that i 's valuation satisfies $x \leq B$. If $x < B - \varepsilon = b_L$ then any bid $b \in [b_L, B]$ is obviously strictly dominated by some $b' < b_L$ provided that $\Pr [i \text{ wins} | b] > 0$. Now suppose that $B - \varepsilon \leq x \leq B$. We must also choose ε such that $b \in [B - \varepsilon, B]$ is strictly dominated for some $b' < b_L$. Such an ε must exist or else $b = \beta^*(x) = x$ for some x , which we know is not optimal so long as $\Pr [i \text{ wins} | b'] > 0$ for $b' < b_L$. □

Proposition 3

Proof. In an equilibrium, the second-price bidders bid their valuations due to Proposition 1, but submit their bids randomly due to Assumption 2.

Theorem 2.1 from Athey and Haile (2007) delivers a unique equilibrium in pure strategies with a strictly increasing differentiable bid function $\beta(x)$ in a independent private value auction with symmetric first-price bidders under the regularity conditions on the value distribution $F(\cdot)$. Below, we show that the same theorem yields an equilibrium here because the problem is unchanged but for the discontinuity described in Proposition 2.

The first-price bidder i chooses the bid b , given his valuation x , to solve

$$\max_{b \leq x} (x - b) \Pr [i \text{ wins} | b]$$

Suppose the equilibrium bid function $\beta(x)$ is strictly increasing, so that we can define the inverse bid function $\eta(b) = \beta^{-1}(b)$. Suppose that the first-price bidders face $m = 1$ second-price bidders with $B > r$ and $\Pr[B] = \alpha > 0$ ($\Pr[0] = 1 - \alpha$). Then we can rewrite the maximization problem as

$$\max_{b \leq x} (x - b) F(\eta(b))^{n-1} (1 - \alpha)^{I[b < B]} \frac{\alpha^{I[b=B]}}{2}$$

Note that the objective function for $b > B$ is the same as the problem with first-price bidders alone $\pi(b, x) = (x - b) F(\eta(b))^{n-1}$. For $b < B$, the optimization problem is the same because the objective function is scaled by the constant $(1 - \alpha)$.

By taking first order conditions, we can see that optimal bid function is the solution to the ordinary differential equation

$$\beta'(x) + (n - 1) \frac{f(x)}{F(x)} \beta(x) = (n - 1) \frac{f(x)}{F(x)} x$$

Solving this, we have the optimal bidding function

$$\beta(x; c) = x - \frac{\int^x F(u)^{n-1} du - c}{F(x)^{n-1}} \tag{A.1}$$

The initial condition satisfies $\beta(r; c) = r$, so $c = 0$. See Paarsch and Hong (2006) for a detailed proof.

From Proposition 2, we know the solution will satisfy $\beta^*(x) = \beta(x; c = 0)$ until some threshold x' such that $\beta^*(x) > B$ for $x > x'$.¹³ Proposition 2 tells us we also

¹³No such x' exists if B is sufficiently large so as to not affect the real-time bidders behaviour.

expect a gap in the bid support at $(b_L, B]$ for some b_L . This threshold is determined by an indifference condition that equates payoffs for bids above and below B . The indifference condition equates the payoff at x' for b_L with the limiting payoff for $b = B + \varepsilon$ as $\varepsilon \rightarrow 0^+$:

$$\begin{aligned} (x' - b_L)(1 - \alpha) F(x')^{n-1} &= \lim_{\varepsilon \rightarrow 0^+} (x' - B - \varepsilon) F(x' + \varepsilon)^{n-1} \\ (x' - b_L)(1 - \alpha) F(x')^{n-1} &= (x' - B) F(x')^{n-1} \end{aligned}$$

Solving, we have

$$b_L = \frac{B - \alpha x'}{1 - \alpha}$$

This provides a downward sloping relationship between b_L and x' . We intersect this with the increasing relationship between x and b given by $\beta^*(x) = \beta(x; c_1 = 0)$ to pin down b_L and x' . Above the threshold x' , the first order conditions tell us the solution still has the form $\beta^*(x) = \beta(x; c)$. To solve for c_2 , we use the ‘limiting’ initial condition $\lim_{\delta \rightarrow 0^+} \beta(x' + \delta; c_2) = B$. I use the word ‘limiting’ to remind us that the bidder does not actually optimally submit a bid at B , but just above it. From (A.1), this means that

$$c_2 = (B - x') F(x')^{n-1}$$

The equilibrium first-price bidding function therefore has the form

$$\beta(x) = \begin{cases} \beta(x; c_1 = 0) & x \leq x' \\ \beta(x; c_2) & x > x' \end{cases}$$

as desired.

Finally, consider the case with $m > 1$ offline bidders that submit an ordered vector of bids $B = [B_1, B_2, \dots, B_m]$ satisfying $r < B_1 < B_2 < \dots < B_m$ with the associated probability vector $\Pr[B] = \alpha = [\alpha_1, \alpha_2, \dots, \alpha_m] > 0$. Then first-price bidder’s problem is given by

$$\max_{b \leq x} (x - b) F(\eta(b))^{n-1} \prod_j^m (1 - \alpha_j)^{I[B_j > b]} \frac{\alpha_j^{I[B_j = b]}}{2}$$

In the simplest case, the logic of the single bid carries forward and the first-price equilibrium bid function has $m + 1$ components and a vector x' of m indifference

Let \bar{X} denote the x ’s support maximum. No such x' exists if $\beta(\bar{X}, 0) < \frac{B - \alpha \bar{X}}{1 - \alpha}$.

valuations that govern the thresholds where the bid function is discontinuous

$$\beta(x) = \begin{cases} \beta(x; c_1 = 0) & x \leq x'_1 \\ \beta(x; c_2) & x > x'_1 \\ \vdots & \vdots \\ \beta(x; c_{m+1}) & x > x'_m \end{cases}$$

However, the discontinuities in the bid function can sometimes overlap for B components and form a single discontinuity. Appendix B outlines the structure of the problem and how the possibility of overlapping bid gaps can create multiple equilibria. \square

Proposition 5

Proof. Fix an offline bidder i with effective competition C . Given that we only observe the winning bid, consider the following three possible events for the winning bid. Denote i 's positive bid by B_i , the winning bid by W_{ut} with the winning bidder's identity I_{ut} , user u and t denotes that user's t th auction:

1. $W_{ut} < B_i$ (includes $W_{ut} < r$): i 's bid is known with certainty to be 0 since it can only take 2 values.
2. $W_{ut} = B_i, I_{ut} = i$: B_i is known with certainty
3. $W_{ut} \geq B_i, I_{ut} \in C$: This is the case where B_i is censored.

The probability of a realization of the winning bid can be written as the following sum:

$$\Pr[W_{ut}] = \sum_{T \in \{T_{i \setminus C}, T_{i \cap C}, T_{C \setminus i}, T_\emptyset\}} \Pr[W_{ut} | u \in T] \Pr[u \in T]$$

This expression uses the law of total probability $\Pr[W_{ut}]$ into its conditional probability given the targeting criteria that a user satisfies. The probability of the three above cases is then given by

$$\begin{aligned} \Pr[W_{ut} < B_i] &= (1 - \alpha_i) \tau_i + (1 - \alpha_i) (1 - \alpha_{i \cap C}) \tau_{i \cap C} + (1 - \alpha_{C \setminus i}) \tau_{C \setminus i} + \tau_\emptyset \\ \Pr[W_{ut} = B_i, I_{ut} = i] &= \alpha_i \tau_{i \setminus C} + \alpha_i (1 - \alpha_{i \cap C}) \tau_{i \cap C} \\ \Pr[W_{ut} \geq B_i, I_{ut} \in C] &= \alpha_{i \cap C} \tau_{i \cap C} + \alpha_{C \setminus i} \tau_{C \setminus i} \end{aligned}$$

where $\alpha_{C \setminus i}$ denotes the probability that C submits a bid on the overlapping subset of targeted users $\Pr[b_C \geq B_i | u \in T_{i \cap C}]$, and $\alpha_{i \cap C}$ denotes the equivalent for C 's

exclusively targeted users $\Pr [b_C \geq B_i | u \in T_{i \cap C}]$. Thus, the two need not be equal. However, by Assumption 5.1.2, i 's bidding probability does not vary by C 's targeting criteria:

$$\alpha_i = \Pr [b_i = B_i | u \in T_{i \cap C}] = \Pr [b_i = B_i | u \in T_{i \setminus C}]$$

The probabilities of the winning simplify because observed bids can rule out certain targeting criteria. For instance, the event $[W_{ut} = B_i, I_{ut} = i]$ implies $u \notin T_{C \setminus i}, T_\emptyset$ so that $\Pr [W_{ut} = B_i, I_{ut} = i | u \in T_{C \setminus i} \cup T_\emptyset] = 0$.

Now, suppose we see two auctions per user. Here as well, some combinations of observed winning bids rule out certain user types. For instance, the event $\{W_{u1} = B_i, I_{u1} = i, W_{u2} > B_i, I_{u2} > i\}$ implies $u \in \tau_{i \cap C}$. Consider the probabilities of the six possible events when we observe two auctions (t and t') per user. To economize on notation, I rule out the case where multiple bidders submit bids equal to b_i , so that the identity of W_{it} is self-evident:

$$\begin{aligned} \Pr [W_{ut} < b_i, W_{ut'} < b_i] &= (1 - \alpha_i)^2 \tau_{i \setminus C} + ((1 - \alpha_i)(1 - \alpha_{i \cap C}))^2 \tau_{i \cap C} \\ &\quad + (1 - \alpha_{C \setminus i})^2 \tau_{C \setminus i} + 1 - \tau \\ \Pr [W_{ut} = b_i, W_{ut'} = b_i] &= \alpha_i^2 \tau_{i \setminus C} + (\alpha_i(1 - \alpha_{i \cap C}))^2 \tau_{i \cap C} \\ \Pr [W_{ut} > b_i, W_{ut'} > b_i] &= \alpha_{i \cap C}^2 \tau_{i \cap C} + \alpha_{C \setminus i}^2 \tau_{C \setminus i} \\ \Pr [W_{ut} < b_i, W_{ut'} = b_i] &= 2 \cdot [(1 - \alpha_i) \alpha_i \tau_{i \setminus C} + (1 - \alpha_i)(1 - \alpha_{i \cap C})^2 \alpha_i \tau_{i \cap C}] \\ \Pr [W_{ut} < b_i, W_{ut'} > b_i] &= 2 \cdot [(1 - \alpha_i)(1 - \alpha_{i \cap C}) \alpha_{i \cap C} \tau_{i \cap C} + (1 - \alpha_{C \setminus i}) \alpha_{C \setminus i} \tau_{C \setminus i}] \\ \Pr [W_{ut} = b_i, W_{ut'} > b_i] &= 2 \alpha_i (1 - \alpha_{i \cap C}) \alpha_{i \cap C} \tau_{i \cap C} \end{aligned}$$

Since these six probabilities sum to 1, we have a system with 5 equations and 6 unknowns $\{\alpha_i, \alpha_{i \cap C}, \alpha_{C \setminus i}, \tau_{i \setminus C}, \tau_{i \cap C}, \tau_{C \setminus i}\}$. However, we have assumed that we observe a fraction of users who appear only once. Thus, we can add the 2 additional equations from above for a total of 7 equations. Thus, the parameters are over-identified provided that the system of equations admits a unique solution. \square

Proposition 7

Proof. The real-time bidder's valuation is given by

$$v_{iut} = x_{it} y_u$$

We wish to identify the distributions of x_{it} and y_u , $F_x(\cdot)$, $F_y(\cdot)$ from the distributions of observed bids $G(\cdot)$. Since x_{it} and y_u are continuously distributed with bounded support, we denote $\text{supp}(x) = [\underline{X}, \overline{X}]$ and $\text{supp}(y) = [\underline{Y}, \overline{Y}]$. Let $\beta(x; y)$ denote the optimal bidding function delivered by Proposition 3.

I identify $F_x(\cdot)$ using the within user variation in observed bids. Fixing user u fixes the corresponding y_u realization Y_u . User u 's bids are denoted by the random variable $b_u = \beta(x_{it}; Y_u)$, which is a function of the random variable x_{it} . Then, denote u 's bid support maximum by $\bar{b}_u = \beta(\bar{X}; Y_u)$. Finally, let \bar{U} denote the set of users with $Y_{\bar{u}} = \bar{Y}$. Users $\bar{u} \in \bar{U}$ can be identified from the bid data because they satisfy $\bar{b}_{\bar{u}} = \beta(\bar{X}; \bar{Y}) = \max_u \bar{b}(Y_u)$. Fix this set of users, we can identify the valuations x_{it} from the bids $b_{\bar{u}} = \beta(x_{it}; \bar{Y})$ using the relationship ,

$$x = \frac{\eta(b_{\bar{u}}; \bar{Y})}{\bar{Y}}$$

where $\eta(b; \bar{Y}) = \beta^{-1}(b; \bar{Y})$ is delivered by Proposition 3. $F_x(\cdot)$ is then identified by the distribution of the recovered x for $x \geq \frac{r}{\bar{Y}}$.

To recover $F_y(\cdot)$, I exploit the between user variation in observed bid quantiles. I begin by showing that $\beta(X; y)$ is strictly increasing in y when $\beta(X; y)$ is continuous at y . With this, I can define the inverse correspondence for a user's bid support maximum $y \in \gamma(\beta(\bar{X}; y))$. Then, I will show that we can use another quantile of $G(\cdot; y)$ to pin down y when this correspondence is not unique.

When we add the unobserved heterogeneity component, the symmetric first-price bidding function from Proposition 3 has K segments of the form

$$\beta_k(x; y) = xy - y \frac{\int_{x'_k(y)}^x F_x(u)^{n-1} du + c_k(y)}{F_x(x)^{n-1}}$$

where $x'_k(y)$ is the indifference point between segments k for $k > 1$ as a function of y . For $k = 1$, $x'_1(y) = r$ and $c_1(y) = 0$, because $\beta\left(\frac{r}{y}; y\right) = r$. For $k > 1$,

$$c_k(y) = \left(x'_k(y) - \frac{B_k}{y}\right) F(x'_k(y))^{n-1}$$

which satisfies the initial condition $\beta_k(x'_k(y); y) = B_k$. Thus, the bidding function can be written for $k > 1$ as

$$\beta_k(x; y) = xy - y \frac{\int_{x'_k(y)}^x F_x(u)^{n-1} du}{F_x(x)^{n-1}} - \frac{(yx'_k(y) - B_k) F_x(x'_k(y))^{n-1}}{F_x(x)^{n-1}}$$

For $X \in [r, x'_2(y))$, I show that $\frac{\partial \beta_1(x; y)}{\partial y} \Big|_{x=X}$ is positive:

$$\begin{aligned}
\frac{\partial \beta_1(x; y)}{\partial y} \Big|_{x=X} &= \frac{X - \frac{\int_{\frac{r}{y}}^X F_x(u)^{n-1} du + y \left(- \left(-\frac{r}{y^2} \right) F_x\left(\frac{r}{y}\right)^{n-1} \right)}{F_x(X)^{n-1}}}{F_x(X)^{n-1}} \\
&= \frac{XF_x(X)^{n-1} - \int_{\frac{r}{y}}^X F_x(u)^{n-1} du - \frac{r}{y} F_x\left(\frac{r}{y}\right)^{n-1}}{F_x(X)^{n-1}} \\
&= \frac{\int_{\frac{r}{y}}^X [F_x(X)^{n-1} - F_x(u)^{n-1}] du}{F_x(X)^{n-1}} \\
&= \frac{\frac{r}{y} \left(F_x(X)^{n-1} - F_x\left(\frac{r}{y}\right)^{n-1} \right)}{F_x(X)^{n-1}} \\
&\geq 0
\end{aligned}$$

and the inequality is strict for $X > \frac{r}{y}$.

For $k > 1$, I also show that we can show that $\frac{\partial \beta_1(x; y)}{\partial y} \Big|_{x=X} > 0$ for $X \in (x'_k(y), x'_{k+1}(y))$:

$$\begin{aligned}
\frac{\partial \beta_k(x; y)}{\partial y} \Big|_{x=X} &= \frac{XF_x(X)^{n-1} - \int_{x'_k(y)}^X F_x(u)^{n-1} du + y \frac{dx'_k(y)}{dy} F_x(x'_k(y))^{n-1}}{F_x(X)^{n-1}} \\
&\quad - \frac{\left(x'_k(y) + y \frac{dx'_k(y)}{dy} \right) F_x(x'_k(y))^{n-1}}{F_x(X)^{n-1}} \\
&\quad - \frac{(yx'_k(y) - B_k)(n-1) F_x(x'_k(y))^{n-2} f_x(x'_k(y)) \frac{dx'_k(y)}{dy}}{F_x(X)^{n-1}} \\
&= \frac{\int_{x'_k(y)}^X [F_x(X)^{n-1} - F_x(u)^{n-1}] du}{F_x(X)^{n-1}} \\
&\quad + \frac{x'_k(y) [F_x(X)^{n-1} - F_x(x'_k(y))^{n-1}]}{F_x(X)^{n-1}} \\
&\quad - \frac{(yx'_k(y) - B_k)(n-1) F_x(x'_k(y))^{n-2} f_x(x'_k(y)) \frac{dx'_k(y)}{dy}}{F_x(X)^{n-1}} \\
&\geq 0
\end{aligned}$$

which is strict for $X > x'_k(y)$ if $\frac{dx'_k(y)}{dy} < 0$.

I will show $\frac{dx'_k(y)}{dy} < 0$ by induction when $\beta(X; y)$ is continuous in y at Y . $\beta(X; y)$ is continuous at Y when the same components B^* of B below Y are *activated*, meaning that $B_j^* \in B^*$ defines a β_k segment that avoids the offline bid B_j^* . I show this for $x'_2(y)$ using the implicit function theorem. Towards this, I define the function $\zeta(x, y) =$

$\beta_1(x; y) - b_2^L(x, y) = 0$ where $b_2^L(x, y) = \frac{B_j^* - (1 - \prod_{i=1}^j (1 - \alpha_i))xy}{\prod_{i=1}^j (1 - \alpha_i)}$. Then, by application of the implicit function theorem, there is a $x'_2(y)$ such that $\zeta(x'_2(y), y) = 0$ and

$$\begin{aligned} \frac{dx'_2(y)}{dy} &= \frac{\zeta_x}{\zeta_y} \Big|_{(x,y)} \\ &= \frac{\frac{\partial \beta_1}{\partial x} - \frac{\partial b_2^L}{\partial x}}{\frac{\partial \beta_1}{\partial y} - \frac{\partial b_2^L}{\partial y}} \Big|_{(x,y)} \\ &> 0 \end{aligned}$$

since we have shown that $\frac{\partial \beta_1}{\partial x}, \frac{\partial \beta_1}{\partial y} > 0$ and $\frac{\partial b_2^L}{\partial x}, \frac{\partial b_2^L}{\partial y} < 0$. Consequently, $\frac{\partial \beta_2(x; y)}{\partial y} > 0$. An analogous argument shows that $\frac{\partial \beta_3(x; y)}{\partial y} > 0$ and the rest follows by induction.

Hence, I have shown that $\beta(X; y)$ is increasing in y whenever $\beta(X; y)$ is continuous in y for a given ‘regime’ of activated B^* . Changes in y can change the regime to say B^{**} but the increasing nature of $\beta(X; y)$ on continuous segments will still hold. We can show however that $\beta(X; y)$ can sometimes fall in y at a discontinuity brought on by a regime change. For instance, suppose that two first-price bidders with $F_x \sim \text{Beta}(2, 2)$ face two offline bids $B = (0.25, 0.3)$ with $\Pr[\beta] = \alpha = (0.1, 0.15)$ and reserve price $r = 0$. As we vary Y from 0 to 1, we see a discontinuous decrease in $\beta(\bar{X}; y)$ near $y = 0.511$ (using the equilibrium selection algorithm defined in Appendix B.2).

Thus, in some cases, the inverse correspondence for a user’s bid support maximum $y \in \gamma(\beta(\bar{X}; y))$ will not return a unique y . However, the correspondence will return finite such y because $\beta(\bar{X}; y)$ is increasing in y except for finite discontinuity points. We can pin down the correct y however using bid quantiles on the increasing segment $\beta_1(\cdot; \cdot)$. Suppose for example that $\{y_1, y_2\} = \gamma(b)$ (generalizing to finite such y is straightforward). We can choose a sufficiently small bid quantile Q' with associated $x_{Q'}$ such that $r < b_{Q'}^i \equiv \beta(x_{Q'}; y_i) = \beta_1(x_{Q'}^i; y_i)$ for $i = 1, 2$ meaning that $b_{Q'}^i$ lies on the first $\beta(\cdot; \cdot)$ segment. Then, the user’s empirical bid quantile $b_{Q'}$ identifies the correct y_i since $\beta(x_{Q'}; y_1) < \beta(x_{Q'}; y_2)$ if and only if $y_1 < y_2$.

Finally, we can identify $F_y(\cdot)$ from the distribution of the recovered y across users for $y \geq \frac{r}{X}$. □

B Multiple Equilibria

In this section of the appendix, I describe in greater detail how the equilibrium real-time bid function works with multiple offline bids. I also include real time bids in the gap with probability δ uniformly distributed on the total interval of the gap. First, I

show in a simple example that multiple equilibria can exist with multiple offline bids. Second, I define my selection algorithm that chooses a unique bid function.

We can also conceive of alternate equilibria of the form in Chapter 8 of Krishna (2009) in which the second-price bidders do not bid their valuation. However, in the data, we the observed first-price bids are typically distributed with positive probability on their entire support. This rules out equilibria where the second-price bidders do not bid their valuations—at least on the support of first-price bids.

B.1 Multiple Equilibria Example

Suppose that two uniform first-price bidders face competition of the form $B = (0.25, 0.3)$ and $\alpha = \Pr[B] = (0.1, 0.15)$. This has a symmetric equilibrium with two gaps given by the bid function

$$\beta^*(x) = \begin{cases} \frac{x}{2} & \text{if } x \leq \frac{5}{11} \\ \frac{x}{2} + \frac{5}{484x} & \text{if } \frac{1}{3} < x \leq 0.4906 \\ \frac{x}{2} + 0.0268x & \text{if } x > 0.4906 \end{cases}$$

with the gaps $(\frac{5}{22}, 0.25]$ and $(0.2664, 0.3]$. It also has a symmetric equilibrium with a single gap given by the bid function

$$\beta^*(x) = \begin{cases} \frac{x}{2} & \text{if } x \leq \frac{120}{247} \\ \frac{x}{2} + 0.0278x & \text{if } x > \frac{120}{247} \end{cases}$$

with the gap $(\frac{60}{247}, 0.3]$ where $\frac{60}{247} = 0.2430$. All decimals are approximations shown to four digits. Note that the algorithm described above selects the second of these two equilibria.

B.2 Equilibrium Selection Algorithm

I define my equilibrium bid function selection mechanism below. The conditions of Proposition 3 are assumed to hold including the presence of $n > 1$ symmetric real-time (first-price) bidders. The distribution of their bids is denoted by $G(\cdot)$. Denote the vector of positive offline bids by the ordered $B = [B_1 \dots B_m]$ with associated probability vector $\alpha = [\alpha_1 \dots \alpha_m]$. Denote the j th bid gap (the dominated bid intervals) by $(b_j^L, B_j]$ with corresponding indifference valuation x_j' . Multiple equilibria are possible because bid gaps for consecutive offline bids may overlap. My algorithm selects a single equilibrium by moving sequentially through B and choosing for each

component B_i the highest possible B_j such that the bid gaps overlap. That is, B_j dominated interval is given by $(b_j^L, B_j]$ and $B_i \in (b_j^L, B_j]$. The algorithm selects the location of the gaps, but the form of the equilibrium bid function is otherwise given by Proposition 3.

Specifically, for each ordered component B_i in B , I apply the following algorithm:

1. Define $j = i$.
2. While $j \leq m$

- (a) Suppose $\delta = 0$. If the indifference valuation x'_j implies that the $b_j^L = \beta(x'_j) \leq B_i$, then the indifference valuation is given by

$$(x'_j - b_j^L) \prod_{k=i}^j (1 - \alpha_k) F_x(x')^{n-1} = (x'_j - B_j) F_x(x')^{n-1}$$

which again implies a decreasing relationship between b_j^L and x'_j

$$b_j^L = \frac{B_j - \left(1 - \prod_{k=i}^j (1 - \alpha_k)\right) x'_j}{\prod_{k=i}^j (1 - \alpha_k)}$$

Now, if $\delta > 0$, the indifference equation is instead given by

$$(x'_j - b_j^L) \prod_{k=i}^j (1 - \alpha_k) G(b_j^L)^{n-1} = (x'_j - B_j) G(B_j)^{n-1}$$

which implies the bid gap

$$b_j^L = \frac{B_j + x'_j \left(\prod_{k=i}^j (1 - \alpha_k) \frac{G(b_j^L)^{n-1}}{G(B_j)^{n-1}} - 1 \right)}{\prod_{k=i}^j (1 - \alpha_k) \frac{G(b_j^L)^{n-1}}{G(B_j)^{n-1}}}$$

where $G(\cdot)$ depends on δ .¹⁴

¹⁴Specifically,

$$\frac{G(b_j^L)^{n-1}}{G(B_j)^{n-1}} = \frac{G(b_j^L)^{n-1}}{\left[G(b_j^L) + (1 - F(r)) \delta \frac{B_j - b_j^L}{\sum_k B_k - b_k^L} \right]^{n-1}}$$

since $G(b_j^L) = (1 - F(r))((1 - \delta)G^*(b) + \delta G^{gap}(b))$. Since $G^{gap}(b)$ is assumed to be uniform, the additional weight accrues to a $\frac{B_j - b_j^L}{\sum_k B_k - b_k^L}$ assuming the equilibrium bid gaps $(b_k^L, B_k]$ are known. Since this information is unknown, I solve for this using a fixed point loop initiated with $\delta = 0$.

(b) Redefine $j = j + 1$.

3. Collect all the b_j^L that satisfy $b_j^L < B_i$.

Note that if the set of b_j^L are satisfy $b_j^L < B_i$ are unique, then the equilibrium is unique. If not, we can potentially have multiple equilibria as in the example below.