Optimal Reserve Prices in Upstream Auctions: Empirical Application on Online Video Advertising

Miguel Angel Alcobendas Lisbona
Yahoo! Inc.
lisbona@yahoo-inc.com

Sheide Chammas
Yahoo! Inc.
chammas@yahoo-inc.com

Kuang-chih Lee
Yahoo! Inc.
kclee@yahoo-inc.com

ABSTRACT

We consider optimal reserve prices in BrightRoll Video Exchange when the inventory opportunity comes from other exchanges (downstream marketplaces). We show that the existence of downstream auctions impacts BrightRoll floors decision. Moreover, it renders the classical derivation of the optimal floor inadequate and suboptimal. We derive the new downstream-corrected reserve price and compare its performance with respect to existing floors and the classical optimal price. The downstream-corrected optimal reserve price proves superior to both. The relevance of this study transcends its particular context and is applicable to a wide range of scenarios where sequential auctions exist and where marketplaces interact with each other.

1. INTRODUCTION

The importance of video advertising is growing. According to a PWC report 2015 [1], online video advertising revenue exhibits the fastest growth in internet advertising: rising from \$6.32bn in 2014 to a forecasted revenue of \$15.39bn in 2019.

Using data from BrightRoll Video Exchange, we estimate and test a structural model to compute optimal reserve prices. A reserve price corresponds to the lowest bid that the seller is willing to accept for his item. The novelty of this work is that we set the optimal reserve price in a marketplace (upstream) taking into account that, in order to deliver the video advertisement, the winner of the exchange competes against bidders in a second exchange, called downstream marketplace. The use of several exchanges may help publishers to increase demand and competition for the inventory opportunity, since different marketplaces may have different demand partners.

In this paper, we derive the reserve price, also called floor, that BrightRoll Video Exchange should charge in order to maximize its expected profits considering the existence of a downstream auction.

From an exchange perspective, video advertising is not very different from allocating static display advertising: the publisher has an inventory opportunity that wants to monetize using an auction conducted by a marketplace. The main difference between the classic banner and video ads is the format: we may have a video ad playing (pre-roll), during (mid-roll), or after (post-roll) the streaming content. Another example is having an ad running concurrently with the streaming content. As the reader can expect, the format affects the bidder's valuation of the opportunity.

In the presence of downstream auctions, when a publisher has an inventory opportunity and wants to allocate it through an auction, it sends an ad-request to a market-place (downstream). Then, the exchange will send a bid request to several real time bidders (RTB), that may be, for instance, demand aggregators (DSP), advertisers, or other marketplaces (upstream). The upstream marketplace will also auction the opportunity. The winner of the latter exchange will then compete against bids submitted downstream. Finally, the downstream winner will be the one delivering the advertisement.

We are not aware of any study quantifying the importance of downstream auctions. However, at BrigthRoll a significant amount of inventory opportunities comes from other exchanges. Given the relevant role of the company on video advertising, we believe that their implications deserve to be studied.

We show that for a given inventory opportunity, the optimal upstream exchange floor is greater than the classical derivation of the optimal reserve price for a monopolist (Riley and Samuelson [9]). The result is very intuitive: the existence of a downstream exchange makes harder to deliver the video ad, since the upstream winner has to compete against bidders downstream.

We use auction theory and statistical tools to overcome data limitations. In particular, we only have information about auctions with bids above the current reserve price, leading to a left-censored data problem. Moreover, conditional on auctions with bids above the floor, we only know the identity of the winning bidder, its bid, the second highest bid, and the current floor. We do not either have data about auctions conducted downstream, except a dummy variable indicating that the upstream winner also won the downstream auction and delivered the video ad.

Due to the large amount of data and expensive computations, we decided to implement the algorithm using Apache Spark, a general-purpose cluster computing system designed to deal with this type of tasks.

We measure the impact of the aforementioned classical

DOI: 10.1145/1235

optimal floor on inventory with and without a downstream exchange. Then, we use experiment data to derive the expected performance of optimal floors with the downstream correction.

Our test shows that implementing the classical optimal reserve price without correction increases the overall generated revenue on inventory with and without a downstream auction. In the particular case of inventory with BrightRoll Video Exchange as a unique marketplace, the expected revenue lift equals 39%, with a observed revenue increase in 77% of the different types of inventory, also called placements. If we look at the effects of the classical floor for a monopolist on inventory with downstream auctions, results are deceptive: while the overall expected revenue lift equals 25%, our recommendations only increased revenue in 67% of the analyzed placements. Results improve by correcting floors due to the existence of downstream auctions. In such a case, the expected revenue lift equals 29% and revenue increased in 77% of placements, outperforming the results from the classical monopolist optimal floor.

As expected, our recommendations work better when we set floors above the current ones. As previously mentioned, data is left-censored. As a result, we have to infer the distribution of bidders' valuation using data that is only observed when bids are above current floors.

The rest of the paper is organized as follows. Section 2 reviews previous work. Section 3 introduces general aspects of the type of problems we are trying to solve. Section 4 presents the model and optimality conditions for the reserve price. Section 5 outlines the estimation methodology. Section 6 describes the data and estimate results. Section 7 presents the experiment results. Finally, section 8 concludes.

2. RELATED WORK

Reserve prices have been largely studied in auction theory. However, empirical literature is scarce. One of the rare studies was conducted by Ostrovsky and Schwarz [3]. Their paper is one of the first large scale experiments measuring the impact of setting reserve prices for online advertisement. They combine previous work in auction theory by Myerson [7], Riley and Samuelson [9] and Varian [11] to derive optimal reserve prices for sponsored search. Their experiment shows that revenue substantially increases as a result of implementing optimal reserve prices. While they have full information about advertisers bids, we have to face the problem of limited data and sequential auctions.

Yuan et al. [12] empirically test several algorithms to set reserve prices. They propose a real-time control function approach to correct reserve prices by using the highest bid. They also test the classical optimal reserve price and two other algorithms based on the regression of the two highest bids. Based on results from their test experiment, the proposed control function approach works better than the optimal auction approach. Due to system requirements, we are not allow to change the reserve price at impression level, preventing us to test the Yuan et al. [12] approach. In any case, we are concerned about the longterm effects of their proposed algorithm. Having a reserve price that directly depends on the highest bid may give bidders incentives to shade their bids as in first price auctions (see V. Krishna [5]). Authors are aware of this possibility, but due to confounding factors (e.g. holidays at the beginning of the experiment, bidders being not aware of the experiment, ...) they could not show that bidding behavior were not affected by the algorithm. We could neither use the regression based models since we do not observe auction data when the highest bid is below the reserve price.

From a theoretical perspective, our paper is related to sequential auctions and auctions with intermediaries. The existing theoretical work assumes that the distribution of the bidders' valuation and the number of competitors is known at each round of auctions. This is not our case since we do not have much information about what is happening downstream. For instance, McAfee and Vincent [6] derive optimal floors when the auctioneer is able to resell the object in a new auction round if bids are lower than the set floors. Other studies like Feldman, et al. [4], and Stavrogiannis et al. [10] characterize the equilibrium behavior of the players in an auction with intermediaries. In their models, there exists an exchange (downstream) that sends the bid request to other upstream exchanges (intermediaries). Given a predefined auction mechanism, Feldman et al. [4] focus on the derivation of the profit maximizing reserve price set by the downstream and upstream exchanges, and the optimal fee charged by the intermediary. On the other hand, Stavrogiannis, et al. [10] studies the optimal reserve price set by the downstream exchange and compares the intermediary profits depending on the auction design.

3. BACKGROUND

Figure 1 displays the analyzed business model. When a publisher has an inventory opportunity, it sends an advertisement request to a market-place denoted as $Downstream\ Marketplace$, which conducts an auction to allocate the opportunity. Then, the downstream marketplace will send a bid request to several real time bidders (RTB), that may be demand aggregators (DSP), advertisers, or other marketplaces. In Figure 1, we have two real time bidders: DSP1 and an exchange denoted $Upstream\ Marketplace$. Similar to the downstream exchange, the upstream marketplace will also conduct an auction to determine the winner and the bid passed to the downstream exchange. In our application the upstream exchange corresponds to the BrightRoll Video Exchange.

The upstream marketplace is assumed to conduct a second price auction and the downstream exchange uses a first price auction mechanism to allocate impressions. Both exchanges are able to set reserve prices in order to maximize their expected revenue. In this paper, we focus on the optimal reserve price set by the upstream exchange.

Only if the advertisement is delivered, the winner of the downstream marketplace pays the transaction price, and the exchange and publisher split the generated revenue. As we will see later, the optimality condition to set the upstream reserve price does not depend on fees charged by exchanges. For that reason, we will assume that the downstream marketplace does not charge any fee for conducting the auction. This is the case where publishers manage the allocation of their inventory using their own marketplace (e.g. Yahoo!, Facebook,...). On the other hand, the upstream marketplace will keep a percentage of the transaction price of the winner in its auction. As we will explain in detail later on, to ensure that the upstream exchange will receive the payment of the fee, it will keep the corresponding fee before passing the bid to the downstream auction.

As a illustration, imagine the scenario depicted in Fig-

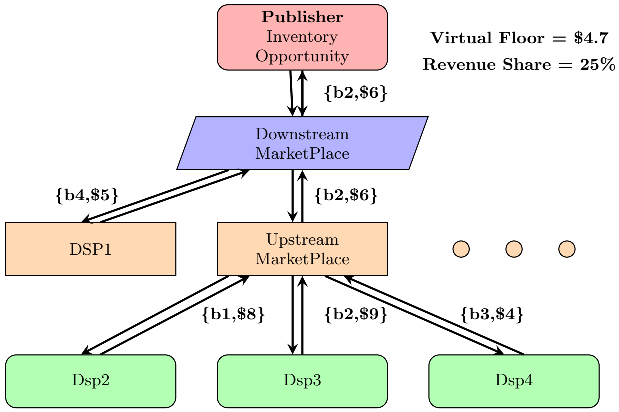

Figure 1: Marketplace Design

ure 1. The publisher has an inventory opportunity and sends the advertisement request to the downstream marketplace. The exchange will ask for bid requests to the upstream marketplace and DSP1. Then, the upstream marketplace will ask DSPs 2, 3 and 4 to submit a bid for the inventory opportunity. Assume that DSP2 submits a bid that equals 8 dollars eCPM, DSP3 bids 9, and DSP4 bids 4. Only bids above the floor are considered. In this example, the upstream floor is set to \$4.7. As a result, the bid submitted by DSP4 is discarded. Among the remaining bids, the winner of the upstream marketplace second price auction is b2, but we use the b1 bid, the set floor and the revenue kept by the marketplace to compute the bid passed to the downstream marketplace. The upstream transition price equals \$8 and corresponds to the maximum of b1 and the floor set at \$4.7. Assume that the upstream exchange keeps 25% of the transaction price in case that bidder 2 is able to deliver the advertisement. Given the fee, the bid submitted to the downstream marketplace equals \$6 and the potential revenue for the upstream marketplace is \$2 eCPM. The bid passed to the downstream auction (the tuple $\{b2, \$6\}$) competes with DSP1. Assuming that the downstream marketplace set a floor equal to zero, since the exchange runs a first price auction and b2's bid is higher than the one submitted by DSP1, b2 wins the auction and pays \$8, and the upstream marketplace makes \$2 profit from the whole transaction.

4. MODEL

We present a structural model to estimate the optimal reserve price that the upstream should set in order to maximize its expected revenues. In the upstream marketplace, we assume that there exist $\mathcal{N}$ potential bidders for a particular inventory opportunity. For each risk neutral bidder i, the valuation of the inventory opportunity (denoted as v_i) is assumed to be an independent and identically distributed realization of the random variable V over the interval $[0, \bar{v}]$, which has a cumulative distribution function $F_V(v)$ and probability density function $f_V(v)$.

In second price auctions with reserve prices, the dominant bidding strategy is to bid bidder's own valuation. The existence of a first-price auction downstream does not change the bidding strategy of bidders participating in the upstream exchange. As a result, truth-telling continues being a dominant strategy. That is,

$$b_i = v_i \text{ for } r \leq v_i$$

where b_i corresponds to the bid made by i and v_i corresponds to the aforementioned valuation of the inventory opportunity.

The proof is very intuitive. Assume that in the upstream marketplace a potential bidder bids less than his valuation. Given such a choice, the potential bidder would risk losing the inventory opportunity to another who bids higher than him, but who has lower valuation. Note that the price paid by the winner is not determined by its own bid but the bid of the nearest opponent. So bidding below v_i is not optimal in the upstream auction. This result is reinforced by the existence of a downstream exchange, since bidding below bidder's valuation increases the probability that a downstream bidder bids slightly higher than the bidder but his valuation is still lower than v_i. On the other hand, bidding more than bidder's own valuation increases the chances of winning the upstream and downstream auctions. However, using the same argument as before, if the upstream bidder i wins, he may pay more than the inventory opportunity is worth to him. Since we have ruled out bids both above and below v_i, the only option is to bid his own valuation v_i. In the presence of a downstream auction, there is no win from deviating from truth-telling.

We extend the general formulation proposed by Riley and Samuelson [9] by capturing the effects of introducing a downstream auction. We will also follow their approach to derive the optimality condition for the reserve price.

Bidder 1's expected profit from participating in the upstream marketplace equals

$$\Pi(x, v_1) = v_1 F_V(x)^{\mathcal{N}-1} P_D(\bar{w}_T(r)) - S(x)$$

where v_1 is the value of the inventory opportunity for bidder 1, x corresponds to bidder 1's bid, $F_V(x)^{\mathcal{N}-1}$ is the probability that bidder 1 wins the upstream auction (i.e. everybody else bids less than bidder 1). $P_D(\bar{w}_T(r))$ is to the probability of winning the downstream auction. It depends on $\bar{w}_T(r)$ that equals the expected transaction price conditional on having at least one bidder above the reserve price $\bar{w}_T$. Finally, $S(x)$ is the expected payment given bid x.

Given bidder 1's profit function, we compute the expected revenue of the upstream marketplace and derive the optimal reserve price taking into account the existence of a downstream exchange.

The following first-order condition must be satisfied if the bidder wants to maximize his expected gain,

$$\frac{\partial \Pi(x, v_1)}{\partial x} = v_1 \frac{d}{dx}\left[F_V(x)^{\mathcal{N}-1} P_D(\bar{w}_T(r))\right] - S'(x)$$

Under truth-telling, the first-order condition for bidder 1 equals

$$\frac{\partial \Pi(v_1, v_1)}{\partial x} = v_1 \frac{d}{dx}\left[F_V(v_1)^{\mathcal{N}-1} P_D(\bar{w}_T(r))\right] - S'(v_1)$$

In the equilibrium, the derivative of the expected payment must satisfy

$$S'(v_1) = v_1 \frac{d}{dx}\left[F_V(v_1)^{\mathcal{N}-1} P_D(\bar{w}_T(r))\right]$$

If the bidder has a valuation equal to the reserve price r, he expects to pay

$$S(r) = rF_V(r)^{\mathcal{N}-1} P_D(\bar{w}_T(r))$$

In this case, the expected transaction bid is just the reserve price (i.e. $\bar{w}(r) = r$).

From bidder 1's perspective, the expected payment is then

$$\begin{aligned} S(v_1) &= rF_V(r)^{\mathcal{N}-1} P_D(r) + \int_r^{v_1} S'(u)du \\ &= rF_V(r)^{\mathcal{N}-1} P_D(r) + \int_r^{v_1} u \frac{d}{du}\left[F_V(u)^{\mathcal{N}-1} P_D(\bar{w}_T(r))\right] \end{aligned}$$

integrating by parts, the expected payment is

$$\begin{aligned} S(v_1) = &\; v_1 F_V(v_1)^{\mathcal{N}-1} P_D(\bar{w}(r)) \\ &- \int_r^{v_1} F_V(u)^{\mathcal{N}-1} P_D(\bar{w}_T(r))du \end{aligned}$$

The upstream marketplace expects to receive a percentage of what bidder 1 expects to pay. This revenue share is denoted as rev and treated as given. Since there are $\mathcal{N}$ symmetric bidders, the expected revenue of the upstream exchange (Π_{ups}) equals $\mathcal{N}$ times the revenue share times the expected payment,

$$\begin{aligned} \Pi_{ups} &= rev\, \mathcal{N} \int_r^{\bar{v}} P(u) f_V(u)du \\ &= rev\, \mathcal{N} \int_r^{\bar{v}} \Big[uF_V(u)^{\mathcal{N}-1} P_D(\bar{w}(r)) \\ &\quad - \int_r^{u} F_V(v)^{\mathcal{N}-1} P_D(\bar{w}_T(r))dv\Big] dF_V(u) \end{aligned}$$

Integrating by parts

$$\Pi_{ups} = rev\, \mathcal{N} P_D(\bar{w}_T(r)) \int_r^{\bar{v}} \left[uf(u) - 1 + F_V(u)\right] F_V(u)^{\mathcal{N}-1} du \quad (1)$$

Given the expected profits of the upstream marketplace, we compute the derivative of Π_{ups} with respect to r to find the optimality condition for r to be optimal.

$$\begin{aligned} &\left[rf_V(r) - 1 + F_V(r)\right] F_V(r)^{\mathcal{N}-1} P_D(r) = \\ &\frac{\partial P_D(\bar{w}_T(r))}{\partial r} \int_r^{\bar{v}} \left[uf(u) - 1 + F(u)\right] F_V(u)^{\mathcal{N}-1} du \end{aligned} \quad (2)$$

where the probability of winning in the downstream exchange is assumed to be increasing in r. Using the chain rule, the derivative of the probability of winning the downstream auction with respect to the reserve price equals

$$\frac{\partial P_D(\bar{w}_T(r))}{\partial r} = \frac{\partial P_D(\bar{w}_T)}{\partial \bar{w}_T} \frac{\partial \bar{w}_T(r)}{\partial r} \quad (3)$$

Looking at equation 3, assuming that $P_D(\bar{w}_T(r))$ is increasing in r is equivalent to assume that $P_D(\bar{w}_T(r))$ is increasing in $\bar{w}_T$, since by definition $\bar{w}_T$ increases with r.

The reserve price that follows from equation 2 leads to the optimal reserve price and it is denoted as ρ_c^*.

The following proposition summarizes previous results.

Proposition 1: Under the assumptions that bidders' valuation of the inventory opportunity is a i.i.d. realization of the random variable V, and bidders are risk neutral, the reserve price that maximizes upstream exchange expected revenue ρ_c^* satisfies

$$\begin{aligned} &\left[\rho_c^* f_V(\rho_c^*) - 1 + F_V(\rho_c^*)\right] F_V(\rho_c^*)^{\mathcal{N}-1} P_D(\rho_c^*) = \\ &\frac{\partial P_D(\bar{w}_T(\rho_c^*))}{\partial \rho_c^*} \int_{\rho_c^*}^{\bar{v}} \left[uf(u) - 1 + F(u)\right] F_V(u)^{\mathcal{N}-1} du \end{aligned} \quad (4)$$

Note that without downstream auctions, the optimality condition for the reserve price equals (see Riley and Samuelson [9])

$$\left[\rho_u^* f_V(\rho_u^*) - 1 + F_V(\rho_u^*)\right] = 0 \quad (5)$$

where ρ_u^* corresponds to the optimal reserve price without downstream correction. Contrary to expression 5, equation 4 depends on the number of competitors and probability of winning the downstream auction. Note however, that the revenue share kept by the upstream exchange (rev) does not have any impact on the optimal reserve price.

Proposition 2: Under the assumptions that bidders' valuation of the inventory opportunity is a i.i.d. realization of the random variable V, bidders are risk neutral, and the probability of winning the downstream auction is increasing in the reserve price, the reserve price that maximizes upstream exchange expected revenue ρ_c^* is greater than the uncorrected optimal floor ρ_u^*.

The proof is simple. Under the stablished assumptions, the payment of bidders upstream is increasing in r, and the optimal reserve price that maximizes the upstream expected revenues in the presence of a downstream exchange is ρ_c^*. As a result, using ρ_u^* in equation 2 is not optimal, lowering the expected revenue.

For a given inventory, the corrected optimal floor (ρ_c^*) is greater than the uncorrected one (ρ_u^*). This result is very intuitive. The existence of a downstream auction makes harder to show an impression since the winner of the upstream marketplace has to compete with bidders downstream.

As we will discuss in detail later on, we will analyze the impact of ρ_c^* and ρ_u^* using data from BrightRoll Video Exchange.

The last part of this section is devoted to further develop the expected transaction price conditional on having, at least, one bidder above the reserve price ($\bar{w}_T$).

By definition, $\bar{w}_T$ equals

$$\bar{w}_T(r) = rP_{W_T}(w = r) + \int_r^{\bar{v}} uf_{W_T}(u, r)du$$

where the first component on the right hand side is the expected payment when there is a single bidder above the reserve price, and the second term equals the expected payment when two or more bids are above the reserve price r. $P_{W_T}(w = r)$ corresponds to the probability that the transaction price is equal to the reserve price r, and $f_{W_T}(u, r)$ is the probability distribution function of the transaction price when there exist more than one bid above the reserve price. Both probabilities are conditional on having, at least, one bid above r. That is the conditional probability for $W = r$ corresponds to

$$P_{W_T}(w = r) = \left[\frac{\mathcal{N} F_V(r)^{\mathcal{N}-1}[1 - F_V(r)]}{1 - F_V(r)^{\mathcal{N}}} \right]$$

where the numerator indicates the probability that there is a single bid above the reserve price, and the denominator equals the probability that someone bids above the reserve price. Similarly, the density function of the transaction price conditional on having two or more bidders above the reserve price equals

$$f_{W_T}(u, r) = \left[\frac{\mathcal{N}(\mathcal{N}-1) F_V(u)^{\mathcal{N}-2}[1 - F_V(u)] f_V(u)}{1 - F_V(r)^{\mathcal{N}}} \right]$$

Remember that the upstream exchange runs a second price auction. Consequently, the numerator corresponds to the density of the second highest bid and follows from the standard formulation of the distribution of the second-highest-order statistic from an i.i.d. sample of size $\mathcal{N}$ from an arbitrary distribution f_V (see Arnold, Balakrishnan, and Nagaraja (1992)). Once again, the denominator equals the probability that there exists at least one bid above the reserve price.

After some algebra, it is easy to show that the corresponding derivative with respect the reserve price equals

$$\begin{aligned} \frac{\partial \bar{w}_T(r)}{\partial r} &= \frac{\partial}{\partial r} r P_{W_T}(w = r) + \frac{\partial}{\partial r} \int_r^{\bar{v}} u f_{W_T}(u, r) du \\ &= P_{W_T}(w = r) + r \frac{\partial P_{W_T}(w = r)}{\partial r} + \int_r^{\bar{v}} u \frac{\partial f_{W_T}(u, r)}{\partial r} du \end{aligned}$$

Derivative that will be used to compute the optimal reserve price (ρ_c^*) in equations 3 and 4.

In the next section we will discuss how to implement the previous results to get the optimal reserve price using data from BrightRoll Video Exchange.

5. ESTIMATION

In order to compute the optimal reserve price in equation 4 we need the distribution of bidders' valuation for the inventory opportunity (f_V, F_V), and the expected probability of winning the downstream auction (P_D). Given data constraints explained below, we estimate the aforementioned distributions using a parametric approach.

The BrightRoll Video Exchange dataset only provides information about the type of inventory (also called placement), the identity and the bid of the winner in the upstream auction, the transaction price (w), the current reserve price (r). On the other hand, we do not have information about upstream auctions where all bids are below the reserve price. Finally, we also have an indicator about which upstream auctions led to an impression in the downstream exchange, but little is known about downstream bidders' behavior.

Our structural model takes into account the available information in order to compute the optimal reserve price in the upstream auction. We will estimate the optimal reserve price in the upstream marketplace in two steps: first, we characterize the valuation distributions of the inventory (f_V and F_V), and estimate the parameters capturing the relationship between the probability of wining downstream and the expected transaction price upstream. In step two, we will use the estimates and expression 4 to compute the optimal reserve price.

In step 1, following Paarsch and Hong [8] we estimate the parameters of the model using the likelihood function of the transactions price. Given the ordered bidder's valuations in the upstream auction $v_1 > v_2 > \cdots > v_{\mathcal{N}}$ and knowing that truth-telling is an equilibrium, the bidder with highest valuation v_1 wins the upstream auction, and pays, conditional on winning the downstream auction, the maximum between what his nearest opponent is willing to pay v_2 and the reserve price r. That is bidder 1 pays w and equals

$$w = max\{v_2, r\}$$

Note that the actual payment is only made when the winner of the upstream auction delivers the advertisement.

When the upstream exchange conducts an auction we may have three possible scenarios: one where all bids are below the reserve price, the case where only one bidder is willing to pay more than the reserve price, and the outcome where two or more bidders submit a bid above the floor.

The probability that all bidders bid below the reserve price equals

$$Pr(w = 0) = F_V(r; \theta)^{\mathcal{N}}$$

where θ corresponds to the vector of parameters defining the shape and scale of the parametric distribution of the bidders' value of the inventory.

When only one bidder bids above the reserve price and the rest $\mathcal{N} - 1$ bidders have valuations below it, the probability is equal to

$$Pr(w = r) = \mathcal{N} F_V(r; \theta)^{\mathcal{N}-1}[1 - F_V(r; \theta)]$$

Finally, the corresponding probability density when two or more bidders value the inventory more than the reserve price is

$$\begin{aligned} f_W(w, n > 2; \theta) = \mathcal{N}(\mathcal{N}-1) F_V(w; \theta)^{\mathcal{N}-2} \\ \times [1 - F_V(w; \theta)] f_V(w; \theta) \end{aligned}$$

where n is the number of bids above the reserve price. This result corresponds to the distribution of the random variable V_2 that corresponds to the valuation of the second highest bid, and follows from the standard formulation of the distribution of the second-highest-order statistic from an i.i.d. sample of size $\mathcal{N}$ from an arbitrary distribution f_V (see Arnold, Balakrishnan, and Nagaraja [2]).

Figure 2 summarizes the structure of the density function of the transaction price in the upstream exchange. As previously mentioned, the distribution has three parts: w equals zero when the inventory is not allocated, w equals the reserve price when only one potential bidder bids above r, and w follows the distribution of the bidder with the second-highest valuation when more than one bidder values the opportunity above r.

As a result, the probability density function of w in the presence of a reserve price is:

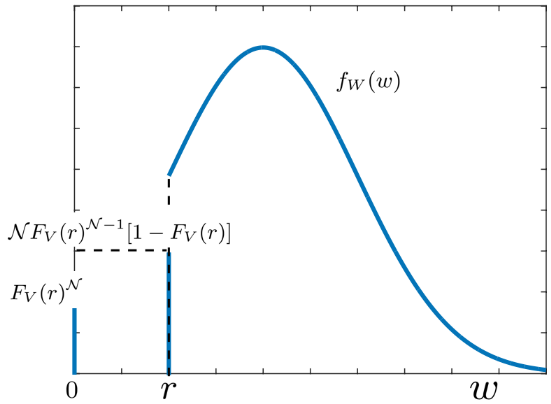

Figure 2: Truncated Probability Distribution of W

$$f^*_W(w;\theta) = \Big(F_V(r;\theta)^{\mathcal{N}}\Big)^{D_0} \Big(\mathcal{N}F_V(r;\theta)^{\mathcal{N}-1}[1-F_V(r;\theta)]\Big)^{D_1} \Big(\mathcal{N}(\mathcal{N}-1)F_V(w;\theta)^{\mathcal{N}-2}[1-F_V(w;\theta)]f_V(w;\theta)\Big)^{1-D_0-D_1} \tag{6}$$

where D_0 and D_1 are indicator variables. $D_0 = 1$ if all bids are below the reserve price and 0 otherwise. On the other hand, $D_1 = 1$ only if only one of the bids is above the reserve price.

Since the dataset does not contain information about auctions with the highest bid below r, we use the truncated version of equation 6 (see Amemiya 1985),

$$f_W(w;\theta) = \left(\frac{\mathcal{N}F_V(r;\theta)^{\mathcal{N}-1}[1-F_V(r;\theta)]}{1-F_V(r;\theta)^{\mathcal{N}}}\right)^{D_1} \left(\frac{\mathcal{N}(\mathcal{N}-1)F_V(w;\theta)^{\mathcal{N}-2}[1-F_V(w;\theta)]f_V(w;\theta)}{1-F_V(r;\theta)^{\mathcal{N}}}\right)^{1-D_1} \tag{7}$$

Given the probability density function in 7, we can construct the likelihood function as follows

$$L = \Pi_{q=1}^{Q} f_W(w_t;\theta|z_t)$$

where Q corresponds to the total number of auctions conducted for a particular type of inventory.

As previously noted, θ corresponds to the vector of parameters characterizing to the distribution functions F_V and f_V. Since we do not have information about auctions with the highest bid below the current floor, we need a parametric distribution to make out of sample predictions. For instance, it may be optimal for the upstream marketplace to decrease the current floor. This type of inference cannot be done using nonparametric approaches. However, choosing a parametric form has some risks linked to misspecification that can lead to poor results. We can assume different parametric forms as long as they have positive support. In our application we use the Weibull family, with density function defined as

$$f_V(v;\theta) = exp(\theta_1)\theta_2 v^{\theta_2-1} exp(-exp(\theta_1)v^{\theta_2})$$

Given the likelihood function in 8, we can estimate the vector of parameters θ by minimizing the negative log likelihood. That is

$$min_\theta - log(L) \tag{8}$$

In order to compute optimal reserve prices in equation 4, we also need to estimate the relationship between the probability of winning the downstream auction $P_D(\bar{w}_T)$, and the expected transaction price upstream $(\bar{w}_T)$. In this paper, we assume a simple linear relationship between both variables,

$$P_D(\bar{w}) = \gamma_0 + \gamma_1\bar{w}(r) + \epsilon \tag{9}$$

where γ_0 and γ_1 correspond to the intercept and slope respectively. ϵ captures possible measurement errors and unobserved factors for the upstream auction designer. Assuming that ϵ is i.i.d., and it is not correlated with $\bar{w}$ nor r, we estimate the parameters using ordinary least squares. If the assumptions hold, the resulting estimates of γ_0 and γ_1 are unbiased. However, we are aware that some correlation between $\bar{w}$ and the unobserved part may exists. For instance, the value of the inventory opportunity may depend on the hour of the day or type of user, affecting the probability of winning the downstream auction and the expected upstream transaction price. This omitted variable problem may be solved by using instrumental variables. Further research should be devoted to avoid this confounding problem.

Given the parameter estimates, we use the optimality condition 4 to compute the upstream optimal reserve price when we account for the existence of a downstream exchange. Similarly, we will use condition 5 when there exist no downstream auctions.

6. DATA AND ESTIMATE RESULTS

We use data provided by BrightRoll Video Exchange. As we will explain in detail later on, we will estimate the optimal reserve price for inventory with and without a downstream marketplace. In the case where there is only one exchange managing the ad-request, BrightRoll marketplace is the only one deciding which bidder will deliver the advertisement. On the other hand, when there exists a downstream marketplace, we will estimate the reserve price taking into account that the winner of the BrightRoll exchange will compete against other bids in a downstream auction. In such a case, the BrightRoll Exchange will be considered the upstream marketplace.

We used 2 weeks data and experimented with inventory with and without downstream auctions. Inventory that captures a significant amount of revenue for BrightRoll. We studied 71 different placements for the case where we only have one exchange, and 30 placements with a downstream marketplace. Due to the large amount of data and expensive computations, we decided to implement the algorithm using Apache Spark, a general-purpose cluster computing system designed to deal with this type of tasks.

We estimated the optimal reserve prices for the first 3 days of the experiment and computed the average of the daily results weighted by the number of ad-requests. The resulting floors are used in the A|B test during the remaining days of the experiment.

In this paper, we assume that all downstream exchanges in our training set conduct a first price auction and upstream

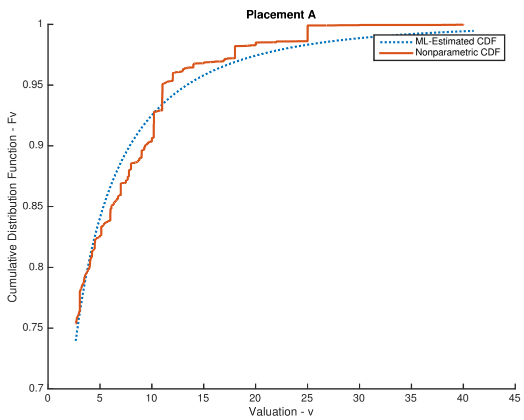

Figure 3: Inventory A: ML vs Nonparametric CDF

marketplaces use a second price auction. In reality, this is not the case for all analyzed placements. In some of them, exchanges work in the opposite way: the downstream uses a second price and the upstream marketplace a first price auction. In this scenario, truth-telling is still an equilibrium for the upstream bidders (the proof is similar than the case considered in the paper). For these placements, our results can be interpreted as a simulation exercise of the consequences of applying optimal floors if placements used the business model considered in the paper.[1]

For each of the placements, we estimate the distribution of bidders' valuation. As a illustration, Figure 3 compares the nonparametric cumulative distribution function for bidder's valuation of inventory A ($\hat{F}_V$) with the Weibull parametric approximation (F_V). Given the aforementioned limitations of the nonparametric approach, the comparison is only possible for the range of observed data. In this particular example, the non-parametric representation is limited to values above the current reserve price set at 2.67 dollars eCPM.

In order to compute the nonparametric distribution, we use the definition of cumulative distribution function of the second order statistic $\hat{F}_W$,

$$\hat{F}_W = \mathcal{N} * \hat{F}_v^{(\mathcal{N}-1)} - (\mathcal{N}-1) * \hat{F}_v^{\mathcal{N}} \qquad (10)$$

where $\mathcal{N}$ is known.[2] $\hat{F}_W$ can be constructed using observed data. Given this information we just need to find $\hat{F}_V$ such that the equality holds. In the particular case of inventory type A, the Weibull and the nonparametric distributions are pretty close. This is not always the case. We could test how close the distributions are and compare the behavior of different parametric families. However, in this paper we will assess the success of our model by comparing the expected lift of our recommended floors with the currently implemented ones using the A|B test.

Given estimates of the density and cumulative distribution functions, we use ordinary least squares to estimate the parameters that define the realtionship between P_D and $\bar{w}_T$ (equation 9). Further details about how we estimate the equation can be found in Appendix 1.

Once all the parameters of the model are estimated, computing reserve prices follows from equations 4 and 5.

7. A|B TEST RESULTS

We conducted an A|B test in order to assess the impact of recommended floors. For each type of inventory, also called placement, we randomly split the ad-request traffic in two: 5% of the traffic is randomly selected to be part of the group with the recommended floor (test group), and the rest of the traffic will be part of control group using the current floor.

In order for the reader to understand how we measure the effectiveness of floor recommendations and the effect of having a downstream marketplace, Table 1 displays two types of inventory denoted as placement B and placement C. While inventory B does not face a downstream auction, the winner of inventory C in the upstream auction has to compete against bids from other bidders in the downstream auction.

The *Floor* column in Table 1 shows the implemented reserve price for the test and control groups, and it is measured in dollars eCPM. The variable *Exp* indicates if the row corresponds to the test ($Exp = 1$) or control group ($Exp = 0$). *Nb Auctions* denotes the number of ad-requests for that particular type of inventory. *Nb Successful* indicates the number of ad-requests that ends up with a winner in the upstream auction. Column 6 corresponds to the ratio of successful auctions to the total number of ad-requests. As the name indicates, column *Nb Impressions* denotes the total number of shown ads. The difference between the number of impressions and successful auctions is displayed in column 8. Several reasons explain why the number of impressions is different from the number of successful auctions: first, having a winner in the upstream exchange does not guarantee that the bidder is going to win the auction conducted downstream. Second, video advertisement is characterized by having a high fallout rate. That happens when the winner of the ad-request is not able to deliver the advertisement. For instance, due to a creative errors or latency problems loading the video. For simplicity, we assume that the fallout rate is independent on the bid made by RTBs. Column *Revenue* corresponds to the revenue in dollars for each group. Note that the test and control groups are not directly comparable, since the test group only corresponds to 5% of all ad-requests. For that reason, we standardize the amount using the total number of ad-requests in the test and control groups. This amount corresponds to column *Standardized Revenue* ($). Finally, the *Revenue Lift* denotes the expected revenue lift as a result of implementing the recommended floor.

If we look at inventory B in Table 1, we propose a 13.41 dollars eCPM reserve price instead of the current one set at $2. As previously noted, the number of auctions assigned to the test group is 5% of the total number of ad-requests. Among ad-requests, 87% of auctions in the test group are successful. The percentage is lower than the one of the control group (91%). This result is expected, since increasing the floor decreases the probability that an auction clears (i.e. auction with at least one bidder above the floor). If we look at the ratio of the number of impressions to the number of successful auctions, around 40% of successful auctions led to an impression in both groups. These impressions generated $411 for the test group and $7,819 for the control.

[1] We are aware that some biases may exist in the estimates linked to the downstream correction term (P_D).

[2] Since we know the identities of the upstream winners, we use the inverse of the Herfindahl index to compute the number of effective competitors $\mathcal{N}$.

Inventory	Floor ($)	Exp	Nb Auctions	Nb Successful	Nb Successful / Nb Auctions	Nb Impressions	Nb Impressions/ Nb Successful	Revenue ($)	Standardized Revenue ($)	Revenue Lift
B	13.41	1	68,872	60,151	87%	24,021	40%	$411	$7,819	8%
	2.00	0	1,308,445	1,194,116	91%	467,240	39%	$7,265		
C	7.69	1	1,770,254	281,262	16%	18,704	7%	$171	$3,247	101%
	13.33	0	33,613,967	844,257	2%	119,933	14%	$1,612		

Table 1: A|B Output

After standardizing the revenue, the expected revenue lift as a result of implementing the recommended floor equals 8%.

Similar analysis can be done for inventory type C. In this case, we recommend a floor below the current one ($7.69 instead of $13.33). Lowering the floor, the ratio successful vs total auctions in the upstream marketplace increases from 2% to 16%. Once again, this result is expected since lowering floors rises the probability that at least one bidder bids higher than the reserve price. In contrast to inventory B, in C the test group ratio of number of impressions to number of successful auctions is half of the control group. This is consistent with the existence of downstream auctions. As previously noted, the upstream winner of inventory C has to face the competition from other bids submitted in the downstream auction. Decreasing the reserve price, increases the number of auctions that clears upstream. However, it decreases the winner bid passed to the downstream auction, decreasing the probability of winning the downstream auction. This is not the case in auctions of type B, where the ratio of impressions to number of successful auctions is the same for the control and test groups. In placements without downstream auctions as in B, winners do not face competition from other marketplaces. If this is the case, the only reason that explains that the ratio is lower than one is fully attributed to fallout, that is assumed to be independent on the transaction price.

We computed and tested the impact of optimal floors on inventory with and without downstream auctions. Due to business requirements, in the A|B experiment we applied the floor resulting from expression 5 and denoted as ρ^*_u. We did not directly applied the correction for the existence of a downstream exchange as appearing in 4 and denoted by ρ^*_c. Instead, we use the data resulting from the experiment and the predicted downstream winning probability to evaluate the optimal floor with downstream correction. Further details about how we do compute the effects of the corrected optimal floors can be found in Appendix 2.

Table 2 shows the impact of implementing the recommended floors. Results are aggregated for confidentiality reasons. The first column describes the type of inventory: the first row corresponds to auctions without a downstream exchange. The second row describes the impact of implementing uncorrected optimal floors on inventory with downstream auctions (ρ^*_u). Finally, the third row corresponds to auctions with a downstream marketplace and corrected optimal reserve price (ρ^*_c). Column *Nb Placements* captures the number of placements in the experiment. We tested optimal floors in 71 different types of placements without downstream auctions. Similarly, we studied the impact of changing floors in 30 placements with downstream auctions. The fifth column describes the percentage of placements that lead to a lift in revenues. This is one of the main indicators to assess the validity of our predictions. For the case of placements without downstream auctions, 77% of our recommended floors (one for each placement) led to an increase in revenue. The performance decreases when applying the uncorrected optimal floor ρ^*_u to placements with downstream auctions. In this case, 67% of the recommendations led to an increase in revenues. The last row corresponds to cases where we use optimal floors with a correction for the existence of downstream auctions. In this case, 77% of the recommended floors led to an increase in revenue. The latter result indicates that the downstream correction outperforms the standard formulation of optimal floors without the correction. Finally, the last column in Table 2 shows the expected revenue lift as a result of the recommendations. This indicator is computed using the total revenue generated in the control group and the standardized revenue of the test group. Implementing optimal reserve prices increases revenue. In the case of placements without downstream marketplaces, the expected revenue lift equals 39%. Once again, the optimal reserve price with downstream correction outperforms the optimal floor without correction in terms of revenue lift: 29% lift using ρ^*_c vs 25% using the uncorrected optimal floor ρ^*_u.

Finally, Table 3 disaggregates previous results by type of recommendation: above and below the current reserve price. For inventory with no downstream auction, 24 out of 71 placements have a recommended floor above the current one, and 47 out of 71 have a lower recommendation. The performance of the recommended floors is linked to the available data. Remember that we only have data of auctions where the highest bid is above the reserve price. As a result, we have to infer the shape of the distribution of inventory valuation for values below the current floor, decreasing the performance of recommendations. As expected, the model behaves better when we recommend above: 88% of recommendations above the current floor led to positive revenue, and 77% in the case of recommendations below the current reserve price. When analyzing the impact of uncorrected optimal floors (ρ^*_u) on inventory with a downstream auction, all recommended floors above the current ones led to an increase in revenue. On the other hand, only 52% of recommendations below the current floor led to an increase in revenue. Finally, the last three rows in Table 3 displays the performance of recommended floors when we correct for the existence of a downstream auction (ρ^*_c). When we make recommendations above the current floor, we improved revenue in 100% of cases with respect to current floors. In the case of corrected floors that are below the current reserve price, we improved in 71% of cases, outperforming the percentage obtained using ρ^*_u.

8. CONCLUSION

In this paper, we derive the reserve price that BrightRoll Video Exchange should charge in order to maximize its ex-

Type of Placement	Nb Placements	% with Positive Revenue Lift	Expected Revenue Lift (%)
No Downstream Auction (ρ_u^*)	71	77%	39%
Downstream Auction: No Correction (ρ_u^*)	30	67%	25%
Downstream Auction: Correction (ρ_c^*)	30	77%	29%

Table 2: Aggregated A|B Test Results

Type of Placement	Nb Placements	% with Positive Revenue Lift	Expected Revenue Lift (%)
No Downstream Auction (ρ_u^*)			
- Above Current Floor	24	88%	38%
- Below Current Floor	47	72%	40%
Downstream Auction: No Correction (ρ_u^*)			
- Above Current Floor	9	100%	92%
- Below Current Floor	21	52%	11%
Downstream Auction: Correction (ρ_c^*)			
- Above Current Floor	13	100%	88%
- Below Current Floor	17	71%	22%

Table 3: A|B Test Results by Recommendation Type

pected profits when the inventory opportunities come from a downstream auction. We prove that the classical approach to derive the reserve price is suboptimal and that the downstream-corrected reserve price increases the expected revenue of the marketplace with respect to the current floor and the classical optimal price.

The model can easily accommodate features. As a result, we can derive different floors depending on supply and demand characteristics (e.g. hour of the day, user characteristics, video format,...). Moreover, the relevance of this study transcends its particular context and is applicable to a wide range of scenarios where sequential auctions exist and where marketplaces interact with each other. Finally, further research can be devoted to analyze the endogeneity problem when deriving the relationship between the probability of winning downstream and the expected transaction price.

9. REFERENCES

[1] Global entertainment and media outlook 2015-2019. http://www.pwc.com/outlook.

[2] B. C. Arnold, N. Balakrishnan, and H. N. Nagaraja. A first course in order statistics. In *New York: John Wiley & Sons*, 1992.

[3] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *National Bureau of Economic Research*, 2005.

[4] J. Feldman, V. Mirrokni, S. Muthukrishnan, and M. M. Pai. Auctions with intermediaries: extended abstract. In *the ACM EC*, 2010.

[5] V. Krishna. Auction theory. In *Massachusetts: Elsevier Academic Press, 2nd edition*, 2010.

[6] R. P. McAfee and D. Vincent. Sequentially optimal auctions. *Games and Economic Behavior*, 18(2):246–276, 1997.

[7] R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[8] H. J. Paarsch and H. Hong. An introduction to the structural econometrics of auction data. In *Cambridge, Massachusetts: MIT Press, 1st edition*, 2006.

[9] J. G. Riley and W. F. Samuelson. Optimal auctions. *The American Economic Review*, pages 381–392, 1981.

[10] L. C. Stavrogiannis, E. H. Gerding, and M. Polukarov. Auction mechanisms for demand-side intermediaries in online advertising exchanges. In *the 13th international conference on autonomous agents and multiagent systems. AAMAS*, 2014.

[11] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.

[12] S. Yuan, J. Wang, B. Chen, P. Mason, and S. Seljan. An empirical study of reserve price optimisation in real-time bidding. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining.* ACM, 2014.

10. APPENDIX

10.1 Appendix 1: Estimation of the Winning Downstream Auction Probability

This section show further details about the estimation of $P_D(\bar{w}_T)$ in equation 9. In order to estimate δ_0 and δ_1 we first need to construct the dependent variable P_D and and covariate $\bar{w}_T(r)$. We trim the dataset discarding auctions within the 4th quantile ordered by the highest bid. As a result, the remaining dataset only contains auctions with a highest bid between the reserve price and the 3rd quantile. Then we select twenty equidistant cutoff points. For each cutoff, we compute the revenue and the ratio of number of impressions to the number of successful auctions resulting from auctions with a highest bid greater than the cutoff point. Each cutoff point is like imposing a reserve price, since auctions below each point does not clear. As a result, we will have twenty $\{P_D(\bar{w}_T)_l, \bar{w}_{Tl}\}$ for $l \in \{1, ..., 20\}$ pairs. Given the resulting pairs, we will use ordinary least squares to have an estimate of δ_0 and δ_1.

10.2 Appendix 2: Estimation of the Winning Downstream Auction Probability Evaluated at the Optimal Floor

For each tested inventory with downstream auctions, we use equations 4 and 9 to compute the optimal reserve price with downstream correction ρ_c^* and the predicted probability of showing the impression evaluated at the optimal floor $(P_D(\rho_c^*))$. For a given inventory, the corrected optimal floor ρ_c^* is greater than the uncorrected one ρ_u^*. This result allows us to use data from the test group to evaluate the performance of $P_D(\rho_c^*)$. As a result, we can use the data from the test group to approximate the impact of floors with downstream correction. We have information about all auctions that successfully cleared even the ones that did not lead to an impression. From the model, we are able to compute the floor with downstream correction ρ_c^* and estimate the probability of winning the downstream auction for that particular floor $P_D(\bar{w}_T(\rho_c^*))$.

In order to compute the expected revenue as a result of imposing the corrected optimal floor ρ_c^*, we use the expression 2, the parameter estimates resulting from minimizing equation 8, and estimates from equation 9.

We also need to compute $P_D(\bar{w}_T(\rho_c^*))$ to estimate the number of impressions resulting from increasing the probability of winning the downstream auction. Using the test group data, we remove successful auctions with the highest bid below ρ_c^* ($Successful_c$). We sum the total number of impressions from the trimmed dataset ($Impressions_c$). Using expression 6 we compute the expected transaction price given ρ_c^*. Given $\bar{w}_T(\rho_c^*)$ and using equation 9 and its corresponding estimates, we compute the expected probability of winning the downstream auction $\hat{P}_D(\bar{w}_T(\rho_c^*))$.

Given $\hat{P}_D(\bar{w}(\rho_c^*))$, $Impressions_c$, and $successful_c$, we need to find the extra number of impressions x, resulting from increasing the probability of winning downstream as follows,

$$\hat{P}_D(\bar{w}(\rho_c^*)) = \frac{Impressions_c + x}{Successful_c} \tag{11}$$

Given the total number of impressions, we can compute the corresponding revenue. While for $Impressions_c$ the computation is straightforward, for the particular added impressions (x) the transaction price will be equal to the floor with correction. The argument here is as follows: if x impression clears, it must be as a result of increasing the floor. So the clearing price will be the floor.

For instance, imagine that the estimated ratio of impressions to number of successful auctions for the corrected floor equals 0.14. We know that we have 1,000 successful auctions, and 100 impressions that qualify with the new floor. As a result, the number of new impressions (x) clearing at the new floor is 40, and the transaction price of each of this extra impressions equals the aforementioned corrected floor.