

# Data-Driven Reserve Prices for Social Advertising Auctions at LinkedIn

Tingting Cui  
LinkedIn Corporation  
tcui@linkedin.com

Lijun Peng  
LinkedIn Corporation  
lpeng@linkedin.com

David Pardoe  
LinkedIn Corporation  
dpardoe@linkedin.com

Kun Liu  
LinkedIn Corporation  
kliu@linkedin.com

Deepak Agarwal  
LinkedIn Corporation  
dagarwal@linkedin.com

Deepak Kumar  
LinkedIn Corporation  
dekumar@linkedin.com

## ABSTRACT

Online advertising auctions constitute an important source of revenue for search engines such as Google and Bing, as well as social networks such as Facebook, LinkedIn and Twitter. We study the problem of setting the optimal reserve price in a Generalized Second Price auction, guided by auction theory with suitable adaptations to social advertising at LinkedIn. Two types of reserve prices are deployed: one at the user level, which is kept private by the publisher, and the other at the audience segment level, which is made public to advertisers. We demonstrate through field experiments the effectiveness of this reserve price mechanism to promote demand growth, increase ads revenue, and improve advertiser experience.

## CCS CONCEPTS

• **Theory of computation** → **Computational pricing and auctions**; *Social networks*;

## KEYWORDS

Auctions, game theory, reserve price, online advertising

## 1 INTRODUCTION

Online advertising constitutes the economic basis for web-based business. Search engines such as Google serve millions of ads daily via keyword auctions [8]. Typically, advertising on search engines starts with advertisers specifying keywords of interest, along with bids indicating the maximum amount they are willing to pay. The cost to the advertisers is generally expressed in terms of cost per click (CPC), and is determined by the bid and the click-through rate (CTR). In this paper, we focus on an emerging form of online advertising, social advertising (specifically at LinkedIn), which allows advertisers to target a specific audience in a social network.

**Social Advertising.** Social advertising has evolved rapidly with the increased popularity in Online Social Networks (OSN). A distinct feature of social advertising is that OSNs usually have rich data on their users' identities. When joining an OSN, users provide information about themselves such as demographics, interests, education, profession, and social connections. Furthermore, when engaging with an OSN, users are usually required to log in. As a consequence, social advertising enables advertisers to target users directly through profile attributes, instead of through keywords. Hence, social advertising offers more effective targeting than traditional advertising. This unique advantage helps advertisers to improve brand awareness and drive quality leads.

**Auction Mechanism.** In online advertising, the number of advertisements which can be shown in a user's update stream is limited. In addition, different positions have different desirability: sponsored ads shown in top positions are more likely to be viewed or clicked by users. Therefore, OSNs need systems to allocate slots to advertisers, and auctions are a natural choice. Typically, multiple ad slots become available from a single user visit, and all ad slots are sold in a single auction. The prevailing auction format is the Generalized Second Price auction (GSP), where ads are ranked by their quality-weighted bids. The price of an ad in position  $i$  is determined by the next highest bid, *i.e.*, the bid of the ad in position  $i + 1$ . In practice, ad networks also assign *quality scores* to ads, and ads with higher quality scores are likely to be placed in a higher position, and pay lower CPC costs.

**Reserve Price.** A common approach for optimizing auction revenue is to apply reserve prices, which specify the minimum bid accepted by the auctioneer. Reserve prices have a great impact on revenues. Advertisers are discouraged from participating in auctions if reserve prices are too high, resulting in low sell-through rate and revenue. On the other hand, low reserve prices may offer poor support of OSNs' inventory when there is lack of competition.

At LinkedIn, reserve price optimization plays a particularly important role in its social advertising auctions. When LinkedIn launched its social advertising product Sponsored Content (SC) in 2013, segment-specific reserve prices were introduced in the auction. These reserve prices were modified from a rate card which was used to price LinkedIn's display ads inventory for guaranteed delivery. This rate card defined reserve prices for different geographic regions with additional markups for other profile attributes such as seniority, title, and industry. The SC market dynamics then evolved significantly, and the rate-card-based reserve prices became obsolete. Also, the scale of LinkedIn's social advertising imposes significant challenges in designing an effective system to compute and serve reserve prices. This motivated our development of a scalable data-driven approach to set reserve prices.

**Main contributions.** Given the significance of reserve prices in auctions, we aim at developing data-driven reserve price models and infrastructure to optimize the joint benefit of LinkedIn's revenue and advertisers' satisfaction. In this paper, we develop a systematic, scalable and adaptive methodology to derive reserve prices at both user level and segment level, and demonstrate its effectiveness through field experiments. Specifically,

- We build a scalable regression model which predicts the distribution of bidders' valuations for each individual user

(i.e., the value derived by the advertiser from showing an ad to a targeted user, or from the user clicking on an ad). This allows us to derive revenue maximizing reserve prices at the user level.

- We design a novel mechanism that combines the user level reserve prices to produce the segment level reserve prices in different geographic markets. A built-in control knob enables us to balance the trade-off between our revenue and advertisers' satisfaction.
- We report on field experiments showing substantial increase in the sell-through rate from the emerging markets (e.g., Latin America), significantly higher bids and revenue per click (+36% lift for bids at floor and +2% lift for bids above floor) from the developed markets (e.g., United States), and great reduction in advertiser churn rates. The data-driven reserve prices prove to be beneficial for both short-term auction revenue and long-term auction health.

The remainder of this paper is organized as follows. In Section 2, we review related literature. In Section 3, we provide background information on the reserve price optimization in the social advertising domain. In Section 4, we discuss engineering implementations of reserve prices. The design of and results from our field experiments are presented in Section 5. Section 6 concludes.

## 2 LITERATURE REVIEW

There has been a good amount of work on the economic and algorithmic perspectives of internet ads auctions, with the majority developed in the domain of search advertising. Formal modeling and equilibrium analysis of GSP auctions are pioneered in [5], [18] and [19], where a subclass of Nash equilibria called envy-free equilibria is defined and shown to be outcome equivalent to VCG auctions for the full information game.

With regard to reserve prices in online advertising auctions, most of the existing literature considers GSP auctions for search ads. Edelman and Schwarz [6] model auctions for a keyword as a repeated game, where equilibrium selection is limited to the outcome of the stable limit of repeated rational play. In that specific setting, GSP with the Myerson reserve price [12] is shown to be revenue optimal. The equilibrium selection is relaxed in [11] to all stable outcomes of the auction, and GSP with an appropriate reserve price is proved to generate a constant fraction of the optimal revenue. In practice, Ostrovsky and Schwarz [13] is the only large scale field experiment we are aware of on reserve price optimization for online advertising auctions.

The theoretical and practical design of online social ads auctions, however, is considerably less developed. The pioneering work [10, 14, 16, 17] has shown effectiveness of online social ads and brought more visibility to social advertising. The distinct features of online social ads are that advertisers target a set of members through profile attributes, and that advertisers are usually budget-constrained. In the context of combinatorial auctions with budgets [4, 7], it is known that no truthful auction exists with respect to both valuation and budgets that produces a Pareto-efficient allocation. In Borgs et. al. [2], an asymptotically revenue-maximizing mechanism is presented, where the asymptotic parameter is a budget dominance ratio which measures the size of a single advertiser

relative to the maximum revenue. However, this mechanism does not work well when there is imbalanced revenue per advertiser, as is the case at LinkedIn. In addition, the static algorithm described in [2] is not practical as our supply of advertising impressions arrive in an online fashion.

## 3 RESERVE PRICE MODELING

We start with a brief introduction to Sponsored Content (SC) at LinkedIn, which serves ads in the update streams of users. The update stream of a user contains content generated by the user's network connections, the groups the user belongs to, the companies the user follows, etc. SC is content that is sponsored by companies to be shown in users' update streams. The advertising process starts with advertisers creating campaigns that target a group of users, specified by targeting attributes from users' profiles such as occupation and skills. Each campaign also specifies a bid (maximum willingness to pay for a user action), along with a budget on the total amount to spend. LinkedIn offers two payment options: CPC (cost per click) and CPM (cost per thousand impressions). In this paper, we focus on CPC for reserve price modeling, as it accounts for the majority of SC bids. The reserve prices for CPM bids can be derived with a slightly modified model.

When a user visits LinkedIn, multiple advertising slots become available in the update stream. A GSP auction determines the placement and pricing of all competing ads. Among all ads that target the user (and have not exhausted their budget), the ad with the highest quality-weighted CPC bid is placed in the highest position, the one with the second highest quality-weighted bid is allocated in the second highest position, and so on. If the user clicks on an ad in position  $i$ , the advertiser is charged the amount equal to the next highest bid, i.e., the bid of the ad in position  $(i + 1)$ . In practice, we use quality-weighted bids to rank the ads, where ads considered high quality are likely to be placed in a higher position and are subject to a lower CPC cost. For simplicity, most of the results in this paper are derived assuming all ads are of the same quality, but they can be extended to the case where ads' quality scores differ [15].

### 3.1 Reserve Price Optimization

**Optimal Reserve Price in Social Advertising.** In Ostrovsky and Schwarz [13], Myerson's reserve price is derived and proved to be optimal in the sponsored search auction setting. However, as discussed in Section 1, some of the assumptions made by Ostrovsky and Schwarz [13] are too strong to hold in the social advertising domain. In general (without making strong assumptions of *envy-free* equilibrium), the optimal reserve price for GSP revenue depends on both the number of bidders and the number of available ad positions [15]. Nonetheless, we choose to stick with Myerson's reserve price and rationalize it from a different perspective as follows.

Unlike in search advertising where multiple ad positions are allocated side by side in the search result page, LinkedIn only allows one ad per page in the update stream. As a consequence, the click probability declines more dramatically by position (e.g., there is a significant drop in CTR between the first and second position). In this case, advertisers have an incentive to bid their true valuation in a GSP auction, and GSP is asymptotically equivalent to VCG [20] because the positional decline factor of click probability  $p_{\text{pos}}$  goes

to zero. To see this, let  $\mathbf{b} = (b_1, b_2, \dots, b_N)$ , ( $b_1 > b_2 > \dots > b_N$ ) be the bids submitted by the  $N$  advertisers, and  $1 > \alpha_1 > \alpha_2 > \dots > \alpha_M > 0$  be the click probabilities from the  $M$  available ad positions. Also, define  $K = \min\{N, M\}$ ,  $\alpha_{K+1} = 0$ , and  $b_{K+1} = 0$  if  $N \leq M$ . Assuming that all ads have the same quality, the VCG payment for an ad in position  $1 \leq k \leq K$  is

$$T_k^{VCG}(\mathbf{b}) = \sum_{\ell=k}^K \frac{\alpha_\ell - \alpha_{\ell+1}}{\alpha_k} b_{\ell+1}. \quad (1)$$

Now it becomes straightforward to see that

$$\lim_{\alpha_{k+1}/\alpha_k \rightarrow 0} T_k^{VCG}(\mathbf{b}) \rightarrow b_{k+1}. \quad (2)$$

To put simply, the VCG payment becomes the GSP payment as the positional decline factor of click probability goes to zero.

### 3.2 Bidder Valuation

A very important step in reserve price optimization is to estimate the valuation distribution, because the optimal reserve price is solely a function of it (as illustrated in Section 3.1). We make the following assumptions to estimate the bidder valuation.

- (1) The distribution of bidder valuations is known to the auctioneer as well as the bidders (necessary for Myerson's reserve price to hold).
- (2) We follow Ostrovsky and Schwarz's assumption that the valuation distribution is log-normal [13]. We observe for most campaigns, log-normal distributions are a reasonable fit for the bid distribution.
- (3) Advertisers bid their true valuation. This assumption is reasoned in Section 3.1. Thus, the observed bid distribution for a user is a good representation of advertiser valuation, which we use directly to derive the optimal reserve price.

### 3.3 Regression Model

In this section we present a regression model to predict the distribution of bidder valuations for a user being targeted by ad campaigns. We make an assumption that the bidder valuations for a user  $i$  are drawn from an i.i.d. log-normal distribution; *i.e.*, the log of bidder valuations are normally distributed with mean  $\mu_i$  and standard deviation  $\sigma_i$ . To make inference on  $\mu_i$  and  $\sigma_i$ , we make an additional assumption that  $\sigma_i = \sigma$  for all users, and utilize a linear regression model that fits the log of valuations (denoted by  $V$ ) against users' profile attributes (denoted by  $X$ ):

$$\log V = X\beta + \epsilon, \quad \epsilon \sim N(0, \sigma^2 I). \quad (3)$$

The design matrix  $X$  is a user-by-attribute binary matrix. Each row of  $X$  represents an indicator vector corresponding to the absence or presence of profile attributes for a user. Following the argument in Section 3.2 that bids are asymptotically equal to valuations, we use the observed bids (denoted by  $B$ ) in place of valuations (denoted by  $V$ ) in the regression. Let  $Y = \log B$  be the log bids, the least square error estimates of the parameters are given by

$$\hat{\beta} = (X^T X + \lambda I)^{-1} X^T Y, \quad (4)$$

where  $\lambda$  is the shrinkage parameter of the Ridge regression.

Given the size of LinkedIn's user base, and the abundance of user profile attributes, it is nontrivial to solve the linear regression at

scale. We build a custom regression solver running on our Hadoop clusters, which computes the matrix products  $\mathcal{A} = X^T X + \lambda I$  and  $\mathcal{B} = X^T Y$  in parallel, and solves a linear system  $\mathcal{A}\beta = \mathcal{B}$  on a single node. The computation of matrix products is greatly simplified under our binary design matrix, where  $\mathcal{A}_{ij}$  indicates the number of advertisement requests with feature  $i$  and feature  $j$  with regularization,  $\mathcal{B}_i$  represents the sum of log bids involving feature  $i$ . Each run of the regression uses the last 90 days of bid data, and is repeated regularly to reflect changes in the market dynamics.

To further improve model fitting, we build separate regression models for different geographic markets. This allows us to fine-tune modeling parameters (such as the shrinkage parameters and the sampling rates) to reflect different market dynamics. The fitting of these geographic models is reasonable, with R-square values ranging between 0.7 and 0.85.

### 3.4 User-Level and Campaign-Level Reserve Prices

**User-Level Reserve Prices.** With the estimation of bidder valuations, we can numerically solve an equation to derive the revenue maximizing reserve price

$$r - \frac{1 - F(r)}{f(r)} = 0, \quad (5)$$

where  $F$  and  $f$  are the CDF and PDF of the valuation distribution [12]. This gives us the optimal reserve price for an individual user (recall Section 3.1).

However, a social advertising campaign usually targets a segment of users, specified by their profile attributes. Typically a target segment consists of thousands or millions of users, making it difficult to publish the reserve prices for all targeted users to the advertiser. When SC was initially launched, reserve prices were set at the campaign level, calculated from a manually defined rate card specifying region-based prices and segment mark-ups. For example, the reserve price to target US users is  $\$a$  and that to target US users aged above 20 is  $\$(a + b)$ . This price enforces the minimum bid price a campaign has to set in order to participate in the auctions. As we move towards a data-driven approach, we decide to keep the public reserve price at the campaign level, to minimize the impact on advertiser experience.

**Campaign-Level Reserve Prices.** A natural choice of computing the campaign-level reserve price is to use a quantile function of the distribution on the user-level reserve prices. The reserve price for a campaign targeting a user segment  $S$  is defined as

$$\hat{r}_S = \sup \{r > 0 \mid \Pr\{R_S \leq r\} \leq p\}, \quad (6)$$

where  $R_S$  is a random variable that denotes the reserve price for a randomly selected user in segment  $S$ , and  $0 < p < 1$  is the quantile of choice. To illustrate the idea, suppose a campaign's targeting segment consists of a million users, and we set  $p = 0.5$ , then the campaign reserve price is equal to the median of the one million users' reserve prices in the target segment. In practice, we take the trade-off between revenue and efficiency into account and apply a discount factor depending on the sell-through rates of different regional markets. Because we start with existing reserve prices, we tend to overestimate the bidder valuations, as any valuation below the reserve price is not observed. The over-estimation is

likely to be more severe for emerging markets (e.g., Latin American countries), where market liquidity is low due to the relatively high existing reserve prices. This flexible discount scheme serves as a heuristic approach to adjust for possible over-pricing due to the over-estimation in valuations. An optimal but more complicated approach is proposed as future work in Section 6.

Although we don't disclose the user-level reserve prices to the advertisers, they remain effective in SC auctions. For example, let a campaign  $j$  targeting user segment  $S$  and bidding  $b_j$  compete for an impression from user  $i \in S$ , and let  $\hat{r}_i$  and  $\hat{r}_S$  be the user-level and campaign-level reserve prices respectively. The campaign will participate in the auction only if  $b_j \geq \hat{r}_i$  and  $b_j \geq \hat{r}_S$ . Its minimum payment is determined by  $\max\{\hat{r}_i, \hat{r}_S\}$ , the maximum of the user level and the campaign-level reserve prices.

As we discussed in Section 3.1, SC advertisers have little incentive to shade bids with respect to the ad positions, if each impression is auctioned off independently. In reality, a group of users with multiple targeted attributes may share the same advertiser budget, and the auction becomes combinatorial. In this setting, even the VCG auction loses its efficiency and truthfulness [4, 7]. If the budget of an advertiser is too small to pay for all user impressions in her target segment, she will have an incentive to bid lower than her true valuation in order to select a less competitive subgroup of users. Such strategic "bargain-hunting" behavior is constantly being observed in SC auctions. In the short term, a campaign-level reserve price helps to increase auction revenue by putting a lower bound on the advertiser bids. To optimize long term revenue, a formal model needs to be developed to study the equilibrium configuration of this combinatorial auction with public and private reserve prices. This remains an open subject for future research.

## 4 ENGINEERING DESIGN

The scale of LinkedIn's user base and ads business imposes significant challenges to the engineering design and implementation of our reserve price system. LinkedIn has more than 500 million registered users, each having up to several million profile features (i.e., each vector of  $X$  has several million entries). In this section, we describe the system architecture and major considerations of the engineering implementation.

The reserve price system consists of two major components: 1) an offline Hadoop pipeline and, 2) an online web service component. Figure 1 illustrates the system architecture of the data-driven reserve price system and its interaction with other components.

The offline data pipeline runs periodically, reading the latest member profile and ad auction logs, fitting the bidder valuation distributions, and computing user level reserve prices. The online service component is part of the campaign manager web service. It serves campaign creation and update requests from advertisers and returns the campaign-level reserve prices.

The offline pipelines can be further divided into two steps: predicting the bidder valuation distribution, and deriving the user level reserve prices from this distribution. After retrieving the latest member profile snapshot, the second step extracts the profile attributes, computes the reserve price for each user based on the predicted distribution, and stores the optimal reserve price for each user in an efficient datastore called Pinot.

Pinot is a real-time distributed online analytical processing (OLAP) datastore delivering scalable real time analytics with low latency, and is widely used in LinkedIn across different product teams [9]. The dimension columns of our Pinot table correspond to the list of targetable profile attributes (e.g., geographical locations, skills, industries, and companies). Since the user level reserve price only depends on the profile attributes, and two users with the same profile attributes have the same reserve price, we aggregate the user level reserve price data by profile attributes before loading it into Pinot. It is worth mentioning that although there can be trillions of combinations of the profile targeting, the number of data entries in the Pinot table is bounded by the number of unique user profiles.

The online component resides in LinkedIn's campaign manager web service, determining the campaign-level reserve prices for each campaign creation or update request in real-time. For each request, the online component queries the Pinot table and the campaign level reserve price is derived based on a pre-configured quantile. Pinot's storage scalability and query optimization mechanism allows us to compute campaign level reserve prices very quickly (within 10 ms on average), ensuring a smooth advertiser experience.

## 5 EXPERIMENT

In this section, we discuss how we apply the methodology developed in Section 3 to set reserve prices for SC auctions in a real social advertising platform that handles over one hundred million ad requests daily, and we present results from our field experiments.

We adopt two different experiment designs to test the new data-driven reserve prices. The whole SC marketplace is divided into **emerging markets** where sell-through-rates are relatively low (e.g., Asia and Latin America), and **developed markets** where sell-through-rates are relatively high (e.g., US and Canada). For the emerging markets, the data-driven approach suggests lower reserve prices than the legacy rate-card based approach. As revenue risk is small from these markets, we choose to roll out the new reserve prices at once, which allows us to quickly observe the demand upside from the lower reserve prices. For the developed markets, which contribute a majority of SC revenue, we launch an A/B test that only reveals the data-driven reserve prices to a randomly selected group of advertisers.

### 5.1 Results from the Emerging Markets

For the emerging markets, we ran a 6-week long test in Q3 2015, during which the new data-driven reserve prices were rolled out to every advertiser. The average campaign reserve prices dropped by 20-60% depending on geographic market, and as a consequence, we observed significant increase in demand. The percent of auctions with at least one participant increased by 30-60% in these markets comparing to the 6-week average before the test. The impact of introducing the data-driven reserve prices has been significant as shown in Figure 2.

When we launched the reserve price test in the emerging markets, we expected short-term revenue loss from the lowered reserve prices. However, the increased demand quickly made up for the lower unit price, and we found revenue to be higher compared to the pre-test period in the second week of the test.




Figure 1: Architecture of Reserve Price System.




Figure 2: Percentage of auctions with at least one participant in emerging markets, normalized so the starting value is 1.0.

## 5.2 Results from the Developed Markets

For the developed markets, we randomly selected a group of advertisers to receive the data-driven reserve prices (treatment), while the rest receive the rate-card-based reserve prices (control). We report results from CPC campaigns targeting the US market only. (We observed similar results from CPM campaigns and campaigns targeting other developed markets.) In this analysis, we focus on campaigns newly created in the first 4 weeks, with over 40,000 and 10,000 campaigns in control and treatment, respectively. The average reserve prices for the treatment group increased by 26% comparing to the control group. We continued the experiment for three months and observed similar metric lifts.

We consider two sets of metrics: *revenue-related* and *advertiser-centric*. Revenue-related metrics measure the direct revenue impact from the reserve prices, and include campaign bid amount and revenue per click. Advertiser-centric metrics are designed to measure impact on advertiser experience, and include

- **abandonment rate** measuring how many campaign creation attempts are abandoned out of all attempts;
- **new campaigns created per advertiser**, the average number of campaigns created by an advertiser in a given period;
- **churn rate**, which measures how many campaigns become inactive out of all previously active campaigns.

**Revenue-Related Metrics.** To compare revenue-related metrics, we further divide the treatment/ control campaigns into two subgroups: a) those bidding exactly at the reserve prices; and b) those bidding above the reserve prices for a better understanding of how campaigns with different bidding patterns behave. After we perform stratified sampling to balance advertiser’s type and remove the outlier campaigns, we find bids and revenue per click from the treatment group to be significantly higher than in the control group (Figure 3). For campaigns bidding at reserve prices, both

Table 1: Changes in median bid and revenue per click, treatment v.s. control.

campaign group	increase in median bid	increase in median revenue per click
bid at reserve price	36.0%	36.0%
bid above reserve price	1.7%	2.2%

Table 2: Advertiser-Centric Metrics, normalized so that the control group always have values of 1.0.

campaign group	abandonment rate	churn rate	new campaigns per advertiser
treatment	1.03	0.83	1.07
control	1.00	1.00	1.00

median bid and median CPC increased by 36%; and for campaigns bidding above reserve prices, median bid increased by 1.7%, while median CPC increased by 2.2% (Table 1). We use median metrics to summarize the comparison since the campaign bid distributions are heavy-tailed.

**Advertiser-Centric Metrics.** The comparison of advertiser-centric metrics between the treatment and control groups are summarized in Table 2. We notice that advertisers did NOT abandon the campaign creation process much more often (a 3% increase in the treatment group) despite the substantial increase in reserve prices. We also observe a 7% increase in campaigns created per advertiser, indicating advertisers are creating as many new campaigns as before despite the higher reserve prices. Finally, we see a 17% reduction in churn rate from the treatment group. This is mainly attributed to the group of campaigns bidding at the reserve price, as they now tend to submit higher (and more realistic) bids, and are thus more likely to win in the auctions and stay active. We consider the reduction in churn rate to be a great improvement in advertiser experience, especially for small and inexperienced advertisers.

It is worth mentioning that our results from A/B testing only capture the short term impact from the reserve prices, not the long term effects on advertisers’ bidding behavior. The increase in downside metrics, such as abandonment rates and churn rates, is likely over-estimated since advertisers’ short-term reactions to the higher reserve price is expected to reflect a worse impact than their long-term reaction. Our analysis indicates that advertiser experience is not adversely affected by the higher reserve price.

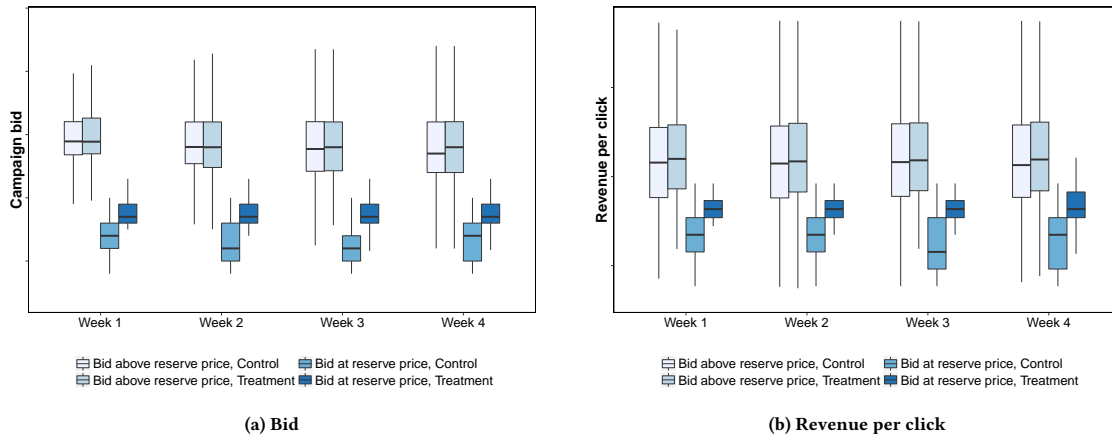


Figure 3: Box plots of bid and revenue per click among campaign groups in the first 4 weeks.

## 6 CONCLUSION AND FUTURE WORK

Reserve prices are an important tool for revenue maximization in auction mechanism design. Although the theoretical subject matter is well studied using stylized models, little practical work has been reported that is applicable in the setting of large-scale online social advertising. We propose a methodology that derives Myerson’s optimal reserve price at the user level, and aggregated reserve price at the campaign level. We deploy this pricing methodology in LinkedIn’s large-scale social advertising platform.

Compared to LinkedIn’s legacy static pricing, this data-driven approach suggests lower reserve prices for emerging markets where demand is low, and higher reserve prices for developed markets where demand is high. Our field experiment shows that lowering reserve prices in the emerging markets significantly increases demand, leading to higher revenue in the longer term. For the developed markets, our A/B testing results indicate that most advertisers previously bidding at the reserve prices are likely to increase bids to stay active in the auctions. For advertisers previously bidding above the reserve prices, their bidding behavior doesn’t seem to change in the short term. As a consequence, we observe increases in revenue per click, which is moderate from advertisers bidding above the reserve prices, and substantial from advertisers bidding at the reserve prices. In the long term, we expect even higher lift in revenue, as strategic advertisers increase their bids in response to the increased competition resulting from higher reserve prices.

In summary, we present a new methodology to set reserve prices in social advertising auctions, guided by auction theory and adapted to our specific domain. The methodology proves to work efficiently and effectively in LinkedIn’s large-scale social advertising platform.

In future work, we plan to improve the approach to estimate the distribution of bidders’ valuations. The existing bids used to estimate the new reserve prices are left-censored due to the fact that only bids above old reserve prices are observed. Given this, a truncated distribution might represent the bidders valuation more suitably. We would like to explore Cox model [3] and probit model [1] to better estimate bidders’ valuations.

## REFERENCES

- [1] T. Amemiya. The estimation of a simultaneous equation generalized probit model. *Econometrica*, 46(5):1193–1205, 1978.
- [2] C. Borgs, J. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. Multi-unit auctions with budget-constrained bidders. In *Proceedings of the 6th ACM conference on Electronic commerce*, pages 44–51. ACM, 2005.
- [3] D. R. Cox. Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological)*, 34(2):187–220, 1972.
- [4] S. Dobzinski, R. Lavi, and N. Nisan. Multi-unit auctions with budget limits. *Games and Economic Behavior*, 74(2):486–503, 2012.
- [5] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. Technical report, National Bureau of Economic Research, 2005.
- [6] B. Edelman and M. Schwarz. Optimal auction design and equilibrium selection in sponsored search auctions. *The American Economic Review*, pages 597–602, 2010.
- [7] A. Fiat, S. Leonardi, J. Saia, and P. Sankowski. Single valued combinatorial auctions with budgets. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 223–232. ACM, 2011.
- [8] Google. Google adwords, <https://www.google.com/adwords/>.
- [9] LinkedIn. Pinot, <https://github.com/linkedin/pinot/wiki>.
- [10] Y. Liu, C. Kliman-Silver, R. Bell, B. Krishnamurthy, and A. Mislove. Measurement and analysis of osn ad auctions. In *Proceedings of the Second ACM Conference on Online Social Networks*, COSN ’14, pages 139–150, New York, NY, USA, 2014. ACM.
- [11] B. Lucier, R. Paes Leme, and E. Tardos. On revenue in the generalized second price auction. In *Proceedings of the 21st international conference on World Wide Web*, pages 361–370. ACM, 2012.
- [12] R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- [13] M. Ostrovsky and M. Schwarz. Reserve prices in internet advertising auctions: A field experiment. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 59–60. ACM, June 2011.
- [14] D. Saez-Trumper, Y. Liu, R. Baeza-Yates, B. Krishnamurthy, and A. Mislove. Beyond cpm and cpc: Determining the value of users on osns. In *Proceedings of the Second ACM Conference on Online Social Networks*, COSN ’14, pages 161–168, New York, NY, USA, 2014. ACM.
- [15] D. R. Thompson and K. Leyton-Brown. Revenue optimization in the generalized second-price auction. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 837–852. ACM, 2013.
- [16] C. Tucker. Social advertising. Available at SSRN 1975897, 2012.
- [17] C. Tucker. Social networks, personalized advertising, and privacy controls. *Journal of Marketing Research*, 51(5):546–562, 2014.
- [18] H. R. Varian. Position auctions. *international Journal of industrial Organization*, 25(6):1163–1178, 2007.
- [19] H. R. Varian. Online ad auctions. *American Economic Review*, 99(2):430–34, 2009.
- [20] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.