International Journal of Industrial Organization 30 (2012) 80–88

Contents lists available at ScienceDirect

International Journal of Industrial Organization

journal homepage: www.elsevier.com/locate/ijio

Comparing predicted prices in auctions for online advertising

Eric Bax [a,*], Anand Kuratti [b], Preston Mcafee [c], Julian Romero [d]

[a] Yahoo! Marketplace Architecture, 3333 Empire Blvd., Burbank, CA 91504, USA
[b] Yahoo! Software Development India Pvt. Ltd., Torrey Pines, Embassy Golflink Business Parks, Off Koramangala-Indiranagar, Intermediate Ring Road, 560 071, India
[c] Yahoo! Research, 3333 Empire Blvd., Burbank, CA 91504, USA
[d] Purdue University 100 S. Grant St. West Lafayette, IN 47907, USA

ARTICLE INFO

Article history:
Received 21 December 2009
Received in revised form 9 June 2011
Accepted 9 June 2011
Available online 14 July 2011

JEL classification:
C13
C44
C45
D81
D84

Keywords:
Reversion
Validation
Bias
Auction
Prediction

ABSTRACT

Online publishers sell opportunities to show ads. Some advertisers pay only if their ad elicits a user response. Publishers estimate response rates for ads in order to estimate expected revenues from showing the ads. Then publishers select ads that maximize estimated expected revenue.

By taking a maximum among estimates, publishers inadvertently select ads based on a combination of actual expected revenue and inaccurate estimation of expected revenue. Publishers can increase actual expected revenue by selecting ads to maximize a combination of estimated expected revenue and estimation accuracy.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

Online publishers use auctions to sell opportunities to advertise, called *ad calls*, to online advertisers. There are two broad categories of online advertising auctions: search and display. In search advertising auctions the advertiser pays only if their ad elicits a click. In display advertising auctions, advertisers may select a basis for payment. Some advertisers pay when the ad is shown, others pay only when showing the ad elicits a user response such as a click or a purchase. (For details on auctions for online advertising, refer to Varian (2006, 2009), Edelman et al. (2007), and Lahaie and Pennock (2007).)

When advertisers pay per click or other user response, the revenue received by the publisher for showing an ad is random. Since user response rates are not known exactly but must be estimated, there is uncertainty in addition to randomness. The estimation accuracy of response rates varies. One reason is that the amount of historical data varies. Another reason is that the response rates themselves vary, and more data is required to estimate smaller rates with the same relative accuracy.

With randomness, a risk-neutral seller seeks to maximize expected revenue. Facing uncertainty, the seller may select an offer having maximum estimated expected revenue. However, this is not necessarily the best policy for maximizing actual expected revenue.

The reason is that selecting a maximum estimate selects for a combination of having an over-estimate and having a large actual expected revenue. Some classes of ads are more likely to have inaccurate estimates, such as ads with lower response rates and ads for which there is less historical data. Even if the individual response rate estimates are unbiased, these classes are more likely to have the largest response rate over-estimates. So selecting a maximum estimate can favor these classes even if they offer less expected revenue than other classes.

Having more buyers in the auction exacerbates the problem, because more estimates means more and more extreme over-estimates. However, having many buyers is not sufficient for selecting a maximum estimate to be a sub-optimal policy for maximizing expected revenue. Varying levels of uncertainty about revenue distributions is also required.

This paper is organized as follows. Section 2 describes related work. Section 3 presents some theory on selection bias for estimated offer values. Section 4 explores correcting selection bias for online display advertising auctions. Section 5 focuses on corrections for search advertising auctions. Section 6 discusses opportunities for future work.

* Corresponding author. Tel.: +1 626 296 1946.
E-mail addresses: ebax@yahoo-inc.com (E. Bax), kuratti@yahoo-inc.com (A. Kuratti), mcafee@yahoo-inc.com (P. Mcafee), jnromero@purdue.edu (J. Romero).

0167-7187/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.ijindorg.2011.06.001

2. Related work

There are a few areas of work related to this paper. One is work by Athey and Levin (2001) on U.S. Forest Service timber auctions. In that work, as in this paper, the seller selects an offer (ex ante) based on estimated values but is paid (ex post) based on actual values. The forest service work focuses on how buyers can use private information to exploit the seller's estimation and selection process.

Another area of related work, by Wilson (1969) and Thaler (1988), concerns the winner's curse. The winner's curse occurs when multiple bidders estimate the value of an item and submit bids based on those estimates. The auctioneer, by selecting the highest bid, tends to select a bid based on an overestimate of value. As a result, the winner tends to realize less value than their bid. Both the winner's curse and the revenue loss studied in this paper are the result of the difference between actual values and first order statistics of estimates of values. (For more on order statistics, refer to David and Nagaraja (2003).) The revenue loss studied in this paper is borne by the seller or market-maker, because the seller or market-maker must estimate the values of bids and bears the exposure from misestimation. When the bidders, rather than the market-maker, bear the risk, Wilson (1969) gives a method to correct for bias.

Another area of related work is machine learning, where uniform error bounds are used extensively to predict whether a model selected on the basis of limited training data is likely to fit as-yet-unseen data drawn from the same distribution. As limited training data is used to estimate the test performance of more models, it becomes less likely that a model that maximizes estimated expected performance will perform nearly as well as its estimate on test data. (For background on machine learning, see Duda et al. (1973), Valiant (1984), and Devroye et al. (1996). Work on uniform error bounds includes Vapnik and Chervonenkis (1971), Audibert et al. (2007), Langford (2005), and Bax and Callejas (2008).) This effect is similar to the gap between a maximum estimated expected revenue ad and the actual expected revenue from that ad. Both are manifestations of regression to the mean, studied by Galton (1886) and Samuels (1991).

The nested classes of classifiers used in support vector machines and other kernel classifiers are similar to classes of ads with different estimation accuracies in this paper. Kernel methods favor classifiers from classes with more certain bounds on test data performance, even if their estimated expected performance is slightly inferior to classifiers from classes with more uncertainty. For more on support vector machines, refer to Vapnik (1998). For other kernel methods, refer to Shawe-Taylor and Cristianini (2004).

In statistics, Hsu and Chen (1996), Wilcox (1984), and Bechhofer and Turnbull (1978) study procedures to select populations with maximum means among sets of populations. In this paper, offers play the role of populations and awarding an ad call to an offer plays the role of a sample. Their work focuses on determining the number of samples needed to confidently select a population with maximum mean, while this paper focuses on selecting an offer before any further sampling.

3. Theory of selection bias for estimated offers

This section shows that favoring offers that have more accurately estimated offer values improves revenue, under the following model. Actual offer values $\mu_1,\ldots,\mu_n$ are drawn i.i.d. from some distribution. The auctioneer does not know these actual values. Instead, the auctioneer receives unbiased estimates $X_1,\ldots,X_n$ of the offer values. The estimation errors are normal, and the auctioneer knows their standard deviations $\sigma_1,\ldots,\sigma_n$.

For simplicity, we assume a first-price auction in this section. Expected revenue is the actual offer value μ_i of the winning offer. In subsequent sections we focus on second-price auctions.

Suppose the auctioneer selects a parameter value c and selects an offer that maximizes $X_i - c\sigma_i$ as the winning offer. Let $r(c)$ be the expected revenue from the winning offer. Then

$$r(c) = E\left[\mu_{\arg\max_i (X_i - c\sigma_i)}\right],$$

where the expectation is over the distribution of $(\mu_1,\ldots,\mu_n,X_1,\ldots,X_n)$.

The following theorem shows that selecting an offer based on the combination of estimated value and accuracy of estimation $X_i - c\sigma_i$ increases expected revenue over simply selecting an offer with maximum estimated value X_i. Specifically, when there are near ties for $max(X_1,\ldots,X_n)$, expected first-price revenue increases if we break ties in favor of offers with lower σ_i.

Theorem 3.1. *Let unknown actual offer values $\mu_1,\ldots,\mu_n$ be i.i.d. random variables. Let estimated offer values $X_1 \sim \mathcal{N}(\mu_1,\sigma_1),\ldots,X_n \sim \mathcal{N}(\mu_n,\sigma_n)$ be normal random variables with actual offer values $\mu_1,\ldots,\mu_n$ as means and known standard deviations $\sigma_1,\ldots,\sigma_n$. Assume $n \geq 3$ and $\sigma_1,\ldots,\sigma_n$ are not all equal. Then*

$$\left.\frac{\partial r(c)}{\partial c}\right|_{c=0} > 0.$$

The proof is in Appendix A. Here is a sketch of the proof based on a small example. Let $n=3$, with $\sigma_1=0$, $\sigma_2=1$, and $\sigma_3=2$. Let μ_1, μ_2, and μ_3 be drawn independently and uniformly at random from $\{7,10\}$. Then $X_1 \sim \mathcal{N}(\mu_1,0)$, $X_2 \sim \mathcal{N}(\mu_2,1)$, and $X_3 \sim \mathcal{N}(\mu_3,2)$. Define $X^* = \max(X_1,X_2,X_3)$.

Informally, the theorem says that when X_i and X_j are nearly tied for X^*, if we break the near tie in favor of the value with lower σ, then we increase the expectation of the selected μ. We can ignore cases where $\mu_i = \mu_j$, because selecting either value produces the same winner's μ. (We also ignore three-way ties, because they have probability $O(c^2)$ as $c \to 0$.)

Without loss of generality, assume $\sigma_j < \sigma_i$. Let T be the condition that X_i and X_j are nearly tied for X^*. Let W be the condition that breaking the near tie in favor of lower σ selects the greater μ: $W = \{\mu_i = 7 \wedge \mu_j = 10\}$. Let $\overline{W}$ be the condition that we select the lower-μ value as winner: $\overline{W} = \{\mu_i = 10 \wedge \mu_j = 7\}$. We want to show

$$Pr\{W|T\} > Pr\{\overline{W}|T\}.$$

Using Bayes' Theorem:

$$Pr\{T|W\} = \frac{Pr\{W|T\}Pr\{T\}}{Pr\{W\}},$$

and

$$Pr\{T|\overline{W}\} = \frac{Pr\{\overline{W}|T\}Pr\{T\}}{Pr\{\overline{W}\}}.$$

Since μ_1, μ_2, and μ_3 are i.i.d., $Pr\{W\} = Pr\{\overline{W}\}$. So we only need to show

$$Pr\{T|W\} > Pr\{T|\overline{W}\}.$$

Consider near ties between X_1 and X_2. The correction breaks near ties in favor of X_1 since $\sigma_1 < \sigma_2$. Because $\sigma_1 = 0$, $X_1 = \mu_1$, which is 7 or 10, and any near tie is near X_1. A near tie with $\mu_1 = 10$ and $\mu_2 = 7$ is just as likely as a near tie with $\mu_1 = 7$ and $\mu_2 = 10$; in either case, X_2 is $3\sigma_2$ from μ_2. However, for X_1 and X_2 to be a near tie for X^* requires X_3 less than X_1 and X_2. This is more likely when X_1 and X_2 are near 10 than near 7. As a result, breaking ties in favor of X_1 over X_2 is more likely to occur when $X_1 = \mu_1 = 10$ than when $X_1 = \mu_1 = 7$.

دریافت فوری ← متن کامل مقاله

ISIArticles
مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات