

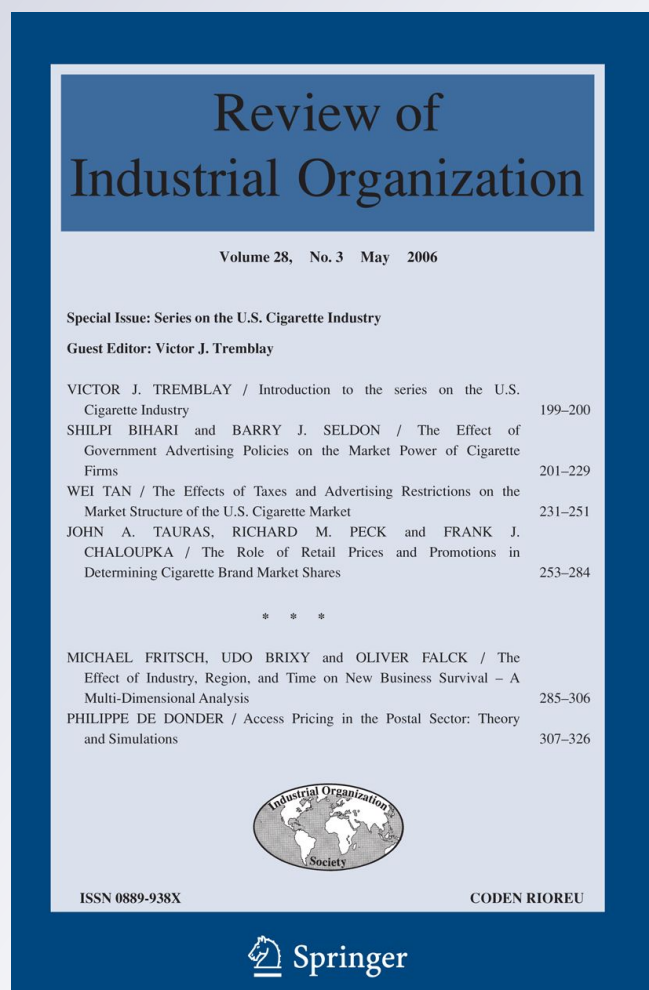
The Design of Advertising Exchanges

R. Preston McAfee

Review of Industrial Organization
An International Journal Published for
the Industrial Organization Society

ISSN 0889-938X
Volume 39
Number 3

Rev Ind Organ (2011) 39:169-185
DOI 10.1007/s11151-011-9300-1



Your article is protected by copyright and all rights are held exclusively by Springer Science+Business Media, LLC.. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your work, please use the accepted author's version for posting to your own website or your institution's repository. You may further deposit the accepted author's version on a funder's repository at a funder's request, provided it is not made publicly available until 12 months after publication.

The Design of Advertising Exchanges

R. Preston McAfee

Published online: 9 June 2011
© Springer Science+Business Media, LLC. 2011

Abstract Internet advertising exchanges possess three characteristics—fast delivery, low values, and automated systems—that influence market design. Automated learning systems induce the winner’s curse when several pricing types compete. Bidders frequently compete with different data, which induces randomization in equilibrium. Machine learning causes the value of information to leak across participants. Discrimination may be used to induce efficient exploration, although publishers (websites) may balk at participating. The creation of “learning accounts,” which divorce payments from receipts, may be used to internalize learning externalities. Under some learning mechanisms the learning account eventually shows a surplus. The solution is illustrated computationally.

Keywords Auctions · Winner’s curse · Machine learning · Display advertising · Internet advertising

The tools of market design—economics practiced as an engineering discipline—have seen widespread deployment in the past two decades. A diverse set of applications include spectrum auctions, physicians’ residencies, search keyword auctions, electricity auctions, secondary school placement, kidney exchange, and the sale of natural resources like mineral rights.¹ The total value of resources that have been directed using principles of market design has probably topped US\$200B at this point in time.

Display advertising—graphical advertisements on web pages—is primarily sold through contracts negotiated by humans. Increasingly, however, advertisements are

¹ See [Abdulkadiroglu et al. \(2005\)](#), [Edelman et al. \(2007\)](#), [McAfee et al. \(2010a\)](#), [McMillan \(1994\)](#), [Milgrom \(2000\)](#), [Roth et al. \(2005\)](#), [Roth \(2010\)](#), [Tietenberg \(2010\)](#), [Varian \(2007\)](#), and [Wilson \(2002\)](#).

run through exchanges, including Yahoo!'s Right Media Exchange (RMX), Google's Double-Click Exchange, and Microsoft's AD-ECN exchange. The design of these exchanges presents a remarkable set of challenges and interesting solutions. Some of the lessons that have been learned, especially regarding externalities in learning, are likely to be relevant in other settings, especially as the technological sophistication of exchanges grows. Many of the problems that are described herein also apply to the problem of sponsored search: advertising on search pages. A major distinction between the exchange environment and the search environment is that the search environment has a single publisher (seller) of advertising that represents most of the supply.

Every application of market design involves a setting with unique, salient features. Failures in balancing supply and demand in electricity markets either create transformer explosions or damage appliances; such extreme constraints limit the scope of market mechanisms. Lack of social acceptability prevents the pricing of human kidneys. In advertising auctions, three major features are: (i) the speed at which the auctions must be accomplished; (ii) the miniscule value and high volume of the items that are being traded; and (iii) the need to use automated systems for bidding, evaluation, and execution of the trades.

The speed of display advertising auctions is breath-taking. After a user clicks on a link and a new page starts to load, the new page itself calls for an advertisement, known as an impression. That call for an advertisement spawns a call to an exchange to supply the advertisement. The exchange then holds an auction for the right to show that particular user an ad. The auction is run, the advertisement selected, pulled from a database, and then sent to the page, all in a fraction of a second. Speed is of the essence, because slowly loading pages create a bad user experience. Moreover, many pages won't load the content until the ad loads, so that the user is left hanging until the ad is delivered.

In the sale of spectrum licenses, individual licenses often sell for hundreds of millions of dollars. Consequently, an enormous amount of thought goes into the behavior of participants: Humans do the bidding; sophisticated software is built to aid the bidders; the auction can do complex things; and the auction may drag on for weeks. The process involves a high degree of deliberation.

Display ad auctions represent the opposite end of the deliberation spectrum. Not only is there no time to do complex things, the items are individually worth very little—usually less than a penny. Prices are sufficiently low that they are quoted in price per thousand, and suppliers of all but the most valuable audiences (e.g., medical and finance) would view \$5 as a good price for a thousand impressions. In contrast to the old joke, however, a seller of advertising impressions *can* make it up in volume: The major exchanges trade billions of impressions per day.

The third characteristic of display advertising is the use of automated systems. Automated systems are needed not just because of the speed of the auctions—a human wouldn't be able to bid in a fraction of a second—but also because of the complex nature of the item that is being sold. The complexity arises because of the varied needs of advertisers. Some advertisers target demographics: age, gender, ethnicity, income, family status. Almost all have relevant geographic markets, ranging from a few nine-digit zip codes to a continent. Many advertisers target consumer interests such as sports cars, skiing, or Swiss cooking; and there are at least 5,000 distinct targeting

variables. Further complicating matters are dozens of standard ad sizes. There may be restrictions put on the ads, such as no moving images, no skin except hands and face, or color limitations.

For major advertisers, such restrictions create trillions of ad types, which could sell for different prices: e.g., a 200×300 pixel flash ad, shown to a 40–45 year old woman residing in Cambridge, MA, married, young children, family income \$50–75K, interested in fashion, family cars, cross-country skiing, and books, currently visiting a news page about politics, appearing on Valentine's day. Changing any one of those descriptors in principle might change the value to some of the advertisers, and hence the market price of the advertisement. Moreover, characteristics of the page and user are not the only relevant considerations: Advertisers have actually sought to advertise only in cities where the sun was currently shining, or only on days where the stock market was up over the previous close.

Given the complexity of the goods being sold, auctions are a natural way to transact. Advertisers can place bids for the types of target opportunities that they seek, and whoever values the opportunity the most will win. Certainly real-time auctions maximize the potential advertiser value and will tend to be more efficient than other transaction mechanisms. Auctions in this environment, with their fire hose speeds and volumes, present a variety of challenges, and the solutions to these challenges will likely prove useful as machine learning permeates exchanges.

Section 1 describes a problem related to the winner's curse, which arises when several pricing types compete in an auction. Section 2 explores the effects of bidders with different data about a common value and shows that randomization is part of the equilibrium. Section 3 provides an overview of machine learning: in particular, how learning about some items naturally spills over to other, related, items. Section 4 explores the effects of the externalities that are created by machine learning: when there is only one publisher or website owner, discrimination in the auction can be used to induce efficient learning. Section 5 considers the case where there are many publishers and advertisers, and proposes a general solution for internalizing the externalities. Section 6 illustrates the solution computationally. Section 7 concludes.

1 The Machine Learner's Curse

Advertisers may prefer to pay by the impression, or only when a click on the ad occurs, or when some action, like a sale, occurs. These are referred to as CPM, CPC, and CPA, for cost per impression, cost per click, and cost per action pricing, respectively. In some exchanges, and in Right Media in particular, all of these pricing tactics may co-exist simultaneously, with some advertisers bidding CPM and others CPC. (The analysis of CPC here applies to CPA as well.) When both CPM and CPC pricing tactics arise, the natural strategy is to estimate the probability of a click, and then compare the cost per impression with the expected cost per impression for the CPC campaign: specifically, the bid price per click times the probability of a click. This product is known as *eCPM*, for expected CPM. The strategy of choosing the highest value of *eCPM* is flawed, however, through our old friend: the winner's curse (Wilson 1969).

In a standard auction context, the winner's curse states that the bidder who overestimates the value of an item is more likely to win the bidding, and thus that the winner will typically be a bidder who over-estimated the value of the item, even if every bidder estimates in an unbiased fashion. The winner's curse arises because the auction selects in a biased manner, favoring high estimates. Savvy bidders adjust their bids in response.

In the advertising setting, however, a second form of winner's curse may arise even when advertisers' bids reflect value. Standard auctions will favor bidders whose click probability is over-estimated, even if the click probability was estimated in an unbiased fashion. Consider having estimates of the click probabilities for two CPC ads with similar true (but unknown) click probabilities. Typically one of the ads will be overestimated, by an amount proportional to its standard deviation of the estimate. Since the true click probabilities were similar, the overestimate will be selected as the best CPC ad, and its true click probability will be less than the estimated click probability. Generally, higher variance ads will have a larger overestimate, and a CPM ad—which doesn't require an estimate of the click probability—is the equivalent of a zero variance CPC ad. Thus, an unbiased prediction method will systematically favor the higher variance estimates, and the realized revenue from the campaigns will be less than the expected revenue.² This form of winner's curse—on the auction selection mechanism—is independent of the standard winner's curse that operates through the selection of bids.

As with the winner's curse, there is a simple fix for this problem: The exchange must adjust the estimated click probability to account for the expected bias. This adjustment requires little beyond the standard deviation of the estimate. In a binomial world, such an estimate is straightforward. However, the simple binomial isn't appropriate in a machine learning environment, for reasons that we explore in Sect. 3. Moreover, the option value of learning is a separate consideration from the winner's curse, and is one that suggests favoring risky payoffs.

2 Cherry-Picking, Data, and Randomization

The plethora of targeting criteria that are available in advertising exchanges presents a challenge for bidders. A bidder who wants to reach interested car buyers might advertise on auto pages. Some visitors to auto pages are clearly better than others. For example, visitors under 17 years of age rarely buy automobiles. Different advertisers have distinct databases; indeed there is a thriving business in the sale of data about customers to advertisers, with hundreds of suppliers. As a consequence, advertisers typically compete in auctions with other advertisers who have different data about the customers.

Advertisers generally have access to three kinds of information about the opportunity to advertise. First, advertisers may have advertised to the user previously and written a "cookie" on the user's computer. A cookie is a small text file that, in principle, is accessible only by the party that created the file. The contents of this cookie can

² This discussion is based on [Bax et al. \(2011\)](#), which also provides a practical solution.

be accessed by the advertiser to form a bid. Second, advertisers may have access to common information, such as the website that the user is currently visiting and the user's IP address. IP addresses frequently provide the user's geographic location, often within a kilometer. Since IP addresses are relatively stable, the IP address may be used as an index, and information about the user is recorded by the advertiser. That is, the advertiser can record all the previous pages in which it encountered that IP address, from which it is often possible to infer characteristics about the user. The current website usually provides a referrer (the previous page visited). Third, access to the information in the first two types of information can be purchased from third parties.

Competing with bidders with superior information is fraught with challenges. One challenge on which I wish to focus is colloquially known as cherry-picking or cream-skimming, a type of adverse selection, which here entails a rival's bidding high on a high quality subset of another firm's target audience. Because there are trillions of categories of advertising opportunities, any bid will encompass hundreds of billions of categories, and it may be possible for rivals to identify subsets that are relatively high value and attempt to acquire these subsets via higher bids. For example, the credit-rating bureaus sell data on creditworthiness, and such data can be used to identify more likely prospects for major purchases like cars and vacations. Suppose that an advertiser bids \$5 per thousand to advertise on a travel site, earning a reasonable rate of return. A rival may start bidding \$6 for creditworthy individuals, thereby extracting most of the credit-worthy individuals and rendering the original \$5 bid unprofitable because, say, the remaining impressions were of lower value. If the bidders have similar values for the advertising opportunities, data purchases are socially wasteful.

Such cherry-picking drives bidders to increasingly refined bidding strategies, resulting ultimately in "real-time bidding," where potentially every opportunity gets a distinct bid that is computed on the fly. Cherry-picking is potentially destructive, however, because it forces bidders to follow strategies that are costly, both in data costs and in computational effort; this is a familiar concept from the economic literature on sorting costs: selling in imperfectly sorted packages (Kenney and Klein 1983) to reduce the costs of sorting.

There is an alternative strategy for the bidders that tends to limit the destructive force of cherry-picking: randomization. Rather than bid a fixed level, an advertiser might randomize the bid submitted. Randomization limits the return to the (better-informed, but not perfectly informed) cherry-picker, who only wins a portion of the inventory, and preserves some of the return of the randomizer. Randomization in advertising exchanges is analogous to the use of randomization in the stock market in the presence of insider trading (Manne 1966).

Let me illustrate the concept with a simple model. There are $n \geq 2$ potential bidders in a common-value, second-price auction. Bidders have data on an opportunity with probability α , $0 < \alpha < 1$. The data provide a signal about the value of the opportunity. For simplicity I assume that the actual value is revealed by the signal, and that it is a common value that is drawn from a cumulative distribution function G .³

³ Revealing the actual value appears inessential; the key element of the theory is that the uninformed bidders will typically make losses against informed bidders, and there is a positive probability that no bidder is informed.

As a consequence of second pricing, bidders who are informed bid the actual value. Bidders who are not informed will usually choose to bid something. To see why, suppose the uninformed bidders chose not to bid. A bid of a penny would win the item at a price of 0 when all the bidders were uninformed, and would otherwise almost always lose. Bidding a small amount produces approximately the mean value, with probability $(1 - \alpha)^{n-1}$. This logic shows that uninformed bidders must bid in equilibrium.⁴

Suppose the uninformed use the bidding distribution F , which may have point masses. Consider an uninformed bidder—say, bidder 1—who bids b . If at least one bidder is informed, bidder 1 wins only if the value is less than b . This means that if any bidder is informed, bidder 1 makes a negative expected profit. If k bidders are informed, the cumulative distribution of the maximum of bids by bidders 2, . . . , n is

$$\theta_k(x) = \begin{cases} F(x)^{n-1} & \text{if } k = 0 \\ G(x)F(x)^{n-1-k} & \text{if } k \geq 1 \end{cases}.$$

The function θ_k gives the distribution of the highest bid of rivals. From θ_k , we have the probability of winning, $\theta_k(b)$, as a function of the bid b , as well as the expected price conditional on winning, $1/\theta_k(b) \int_0^b x\theta'_k(x)dx = b - 1/\theta_k(b) \int_0^b \theta_k(x)dx$, for any given value of k .

Let $\mu = \int_0^\infty vg(v)dv$ be the average common-value. Bidder 1 earns

$$\begin{aligned} \pi(b) = & (1 - \alpha)^{n-1} \left(\mu F(b)^{n-1} - \int_0^b x dF(x)^{n-1} \right) + \sum_{k=1}^{n-1} \binom{n-1}{k} \alpha^k (1 - \alpha)^{n-1-k} \\ & \times \left(\int_0^b vg(v)dv F(b)^{n-1-k} - \int_0^b x dG(x)F(x)^{n-1-k} \right). \end{aligned}$$

With probability $(1 - \alpha)^{n-1}$, all the rivals are uninformed. Bidder 1 wins with probability $F(b)^{n-1}$, obtaining the average value and paying the price that is the highest bid from $n - 1$ independent draws from F . In contrast, there are $k \geq 1$ informed rivals with probability $\binom{n-1}{k} \alpha^k (1 - \alpha)^{n-1-k}$, and in this case all k rivals bid the true value. Thus 1 wins only when the true value is less than b , and the other $n - 1 - k$ uninformed bidders bid less than b as well. This gives a value that is the expected value, and a price that is a draw from θ_k , subject to its being less than b .

The distribution F represents an equilibrium bidding distribution if any bid getting positive weight by F maximizes π . The lowest bid in the support must give zero profits, because if it gave strictly positive profits, a bidder would be better off bidding slightly

⁴ This model is distinguished from Abraham et al. (2011), which precedes this work, primarily by my assumption that more than one bidder may be informed. This assumption allows independence, which simplifies the analysis. Several informed bidders is sensible in the advertising context, but not in Hendricks and Porter (1988) analysis of off-shore oil auctions, which analyzed first-price auctions with one informed bidder.

more to resolve ties in her favor. Consequently, in any equilibrium, bids in the support of F produce zero profits, and no bids produce positive profits for uninformed bidders.

If the support of the distribution F has a non-empty interior, $\pi'(b) = 0$ for all b in the interior of the support. In ‘‘Appendix’’, I show that $\pi'(b) = 0$ implies

$$\left(\frac{(1 - \alpha)F(b)}{\alpha + (1 - \alpha)F(b)} \right)^{n-2} = \frac{\int_0^b G(v)dv}{\int_b^\infty 1 - G(v)dv}. \tag{1}$$

For $n > 2$, the left-hand side is increasing in b and ranges from 0 to $(1 - \alpha)^{n-2}$ as b ranges from 0 to the top of the support. The right-hand side is increasing in b and ranges from 0 to 1 as b ranges from 0 to the mean μ of G .⁵ Consequently the uninformed bidders never bid more than the mean value, which is intuitive, since nothing suggests that the value exceeds the average. Provided that G is strictly increasing and $n > 2$, an equilibrium bidding function F is strictly increasing and unique. Moreover F increases, meaning bids fall, in both the probability of being informed, α , and the number of participants, n .

When $n = 2$, the competition to bidder 1 is either informed, in which case profits are exactly zero, or uninformed. Profits are zero when facing an informed rival, because that bidder bids the true value; if the uninformed bidder wins against an informed rival, the price is the true value. Therefore, since profits average zero, profits must also be zero when the uninformed bidder faces an uninformed bidder. This situation collapses to Bertrand competition, and the uninformed bidders bid the mean μ . For $n > 2$, however, there is a non-degenerate distribution of bids: the solution to (1).

If there are a large number of bidders, and we assume that the expected number of informed bidders $\alpha n \approx A$ is held constant, then

$$\frac{\int_0^b G(v)dv}{\int_b^\infty 1 - G(v)dv} = \left(1 - \frac{A/n}{A/n + (1 - A/n) F(b)} \right)^{n-2} \approx e^{-\frac{A}{F(b)}},$$

which gives a sharp, if somewhat unusual, closed form for the bidding distribution.

When G is $U[0,1]$, $\frac{\int_0^b G(v)dv}{\int_b^\infty 1 - G(v)dv} = \left(\frac{b}{1-b}\right)^2$, and thus, $F(b) = \frac{\alpha}{1-\alpha} \left(\frac{b^{2/n-2}}{(1-b)^{2/n-2} - b^{2/n-2}}\right)$.

Uninformed bidders will generally randomize their bids in response to the possibility of rivals with superior data. Randomization protects against cherry-picking by informed bidders, even perfectly informed bidders, by limiting the effects of the cherry-picking strategy. When a better-informed (but not perfectly informed) agent bids somewhat higher on a high-value subset of the inventory, it only wins a portion of that inventory rather than the entire subset. Yahoo! has implemented technology to share competing demands of high-value advertising campaigns using randomized bids along the lines described here, insuring that bidders who set broad targeting criteria receive a broad mixture of inventory types. In particular, randomized bids insure

⁵ The value at $b = \mu$ can be derived from the expression $\mu = \int_0^\infty 1 - G(x)dx$.

that campaigns win some inventory that would have commanded a high price in the exchange. The method is discussed in McAfee et al. (2010b).

I now turn to the incorporation of machine learning into the prediction of actions in exchanges.

3 A Sidebar on Machine Learning

Recent advances in machine learning are impressive indeed. We depend on them when we search the web to produce an appropriate list of relevant documents out of a population of hundreds of billions of documents. Just how impressive machine learning is can be seen from a market statistic concerning unique searches. When a user searches for an item, say “Windows XP error 1706,” all search engines standardize the query. Standardization involves removing some punctuation and capitalization, eliminating the plural on some terms, correcting spelling mistakes, and other adjustments to make it more likely for queries to be comparable. Even with standardization, the majority of queries, representing more than 10% of searches, searched in a month occur only once. Thus, there is very little direct data on user behavior concerning most of the queries!

In spite of the absence of data, search usually works pretty well. Even if a query like “Windows XP error 1706” had never occurred before, there have been lots of “Windows XP error” queries. Through various means (page rank, text matching, historical clicks), the domain microsoft.com came to rank highly for these nearby queries, and thus is favored over other sites that might have “error 1706” on them. Approximately speaking, a query that is “near” others will be matched to sites that work for the neighbors. Machine learning exploits similarity to perform well even when exactly matching data are lacking.

Machine learning works well overall, but text matching is not perfect. Slight differences in text may produce dramatic differences in meaning. For example, the auto-correct feature of iPhone text entry produces often hilarious changes.⁶

As a consequence of the structure of machine learning, it is often very difficult to estimate the actual variance with which a prediction is made. Thousands of behaviors across thousands of related terms, with an unknown matching strength or correlation, go into click probability estimation, and it is often not possible with current technology to estimate accurately the variance of the estimates. As a result, the commonly assumed environment in economics modeling of a known mean and variance is implausible in some machine learning settings. Moreover, it is even more difficult to assign “credit” to past learning for present accuracy. Thus, it appears impossible to quantify the value that is created by some event in making better future decisions.

4 Externalities from Machine Learning: The Value of Learning

As described above, information from one action spills over to others; for example, if we see a high click rate on Ford ads placed on *Car & Driver* magazine pages, the system will learn, thanks to common features, that Toyota ads are also likely to work

⁶ See <http://damnyouautocorrect.com/>.

well on *Road & Track* pages. This is a classic externality, and is most extreme with “new” ads, particularly ads with new, never-before-seen features. The early experience of these ads on a variety of pages offers future benefits for other exchange participants, primarily through learning how the new features interact with old features.

If one side of the market is represented by a single player, as in keyword (paid search) auctions, a simple solution is to price the externality.⁷ Write the value from running advertisement i in the form

$$B_i = eCPM_i + VOL_i, \quad (2)$$

which represents the immediate expected value, $eCPM$, plus a mnemonically named value of learning term, which is the increased system or social future value from the better knowledge that is gained by running advertisement i now. Equation (2) represents a standard Bellman equation, which is familiar from dynamic programming (Luenberger 1979), which breaks the optimization problem into a current value and a future value. Unlike the Bellman equation, however, the VOL term does not represent the value of the future but the change in the value of the future that is associated with running the ad now, against, say, no ad. The magnitude of the future value depends on uncertainty about the actual performance of advertisement i . If the value of running ad i were known with certainty, there would be no future value of running ad i now.

With a single party on one side of the exchange, a natural solution to the problem of externalities is to run a second-price auction using values transformed by (2). This is done by ranking the ads from highest to lowest according to the value of B_i , then charging the advertiser with the highest B , say 1, an $eCPM$ price that just equates the highest B to the second highest B ; specifically the price is given by

$$p = eCPM_2 + VOL_2 - VOL_1. \quad (3)$$

Equation (3) gives the lowest value that 1 can bid, in $eCPM$ terms, and still win over 2.⁸

Provided that the bidders cannot influence the VOL terms, the discrimination that is implicit in pricing via $eCPM_2 + VOL_2 - VOL_1$ is incentive compatible; bidders will bid their value, because the price paid is independent of their bid, and the system-efficient advertisement is selected. Using the value of learning as part of the optimization has the direct effect of making the system more efficient, thus choosing advertisements more effectively. However, it is theoretically possible that revenue may fall, because the highest immediate value isn't selected.

There are two distinct reasons why revenue should rise: First, ads are selected more efficiently, so the overall value of running ads rises. This is not sufficient, however, to conclude that the value that is extracted by the publisher rises. Over a reasonably long time, the value that is extracted by the publisher must rise, because of a second reason: If both $eCPM$ and VOL are high for a given ad, that ad will run, which decreases VOL as

⁷ This approach is taken from Li et al. (2010).

⁸ For this analysis, it does not matter if payment is made per click or per impression. If payment is made per click, the per click price is just the $eCPM$ calculated in (3) divided by the click probability. Of course, the earlier discussion of the winner's curse still applies.

learning occurs. Consequently, it will eventually be the case that if $eCPM_1 > eCPM_2$, then $VOL_1 \leq VOL_2$. Thus, the net adjustment for the winning bidder, $VOL_2 - VOL_1$, will tend to be a surcharge to bidder 1 for not running ad number 2, about which less is known.⁹

A single party matters because that party can expect to capture most of the benefits of learning in the future. The search environment approximates the single party case; a variety of advertisers deal with a single publisher (the search engine), and the publisher can reasonably anticipate that following a sensible learning algorithm will more than pay in the future for any current foregone revenue. In contrast, with many publishers, the benefits of learning often will not accrue to the party that pays the costs.

5 Externalities from Machine Learning: The Learning Account

A single party on one side of the market matters, because that party can be expected to capture much of the gains from using a forward-looking approach to selecting advertisements. In contrast, in the advertising exchange setting, there are many publishers and advertisers. In many instances it may be optimal from the system or social perspective to select a low $eCPM$ advertisement because the value of learning is high, but in this case revenue is low, so the publisher will not benefit and will be disinclined to accept such advertisements. Learning involves positive externalities, and publishers may be loath to subsidize other parties by accepting low prices for future learning.

Moreover, any reasonable exchange must give the publisher controls that would allow it to avoid new advertisements. Publishers bear a cost from offensive ads; advertisements for home automation (“popunders”), weight loss, dental, and dizzying refinancing lead the list. But offensiveness is context-dependent. Semi-pornographic ads are acceptable only in some circumstances; oil company ads may be offensive on environmental pages. Publishers have a valid reason to want to control the types of advertisements that run on their pages. Such control, however, may be used to prevent low $eCPM$, high VOL ads from running.

Furthermore, it is not at all clear why a publisher should be asked to accept low $eCPM$ advertisements, even if this were feasible. The point of running such advertisements is to create a benefit for the system as a whole; imposing the social cost on a single publisher who happens to be available violates the principle that the parties receiving the benefits should pay the costs. This problem of accepting low $eCPM$ ads is potentially extreme, as the price given by (3) can easily be negative. While negative prices can readily be prevented with reserve prices, the possibility of negative prices illustrates the unreasonable nature of imposing the entire impact of the value of learning on the present publisher.

A possible solution to the problem of externalities induced by machine learning is to separate the payments made by advertisers from the payments made to publishers. Specifically, we charge advertisers based on (3): the second-price for the social benefit

⁹ The description here suppresses the most important contribution of Li et al. (2010), which is to produce a practical mechanism by approximating the VOL terms, and then testing using an experiment with live Yahoo! search traffic.

of running the advertisement. However, the payment made to publishers is the second highest value of e CPM. This payment system requires a learning account, operated by the mechanism, so that payments made and revenue collected can be different. The main question is whether the learning account runs a deficit or not; if the learning account runs a surplus, the solution to the problem of externalities is feasible without external subsidies.

The second highest value of the e CPM will not always come from the same advertiser as the second highest value of the social benefit B . For example, suppose advertiser 1 has an e CPM of 1 and VOL of 10, advertiser 2 has an e CPM of 1 and VOL of 5, and advertisers 3 and 4 both have e CPM equal to 2 and VOL equaling 0. In this case, the second highest e CPM is 2, associated with either advertiser 3 or 4, while the second highest advertiser payment arises with advertiser 2, and produces a second-price of $1 + 5 - 10 = -4$, as given by (3).

Why are publishers and advertisers not treated symmetrically? The answer is that the appearance of a publisher is exogenous, while the appearance of an advertiser is not. That is, the publisher appears to the system, asking for an impression, and the system assigns the advertiser. Thus, efficient learning entails selecting the right advertisement for a given publisher, not in generating the right publisher/advertiser pairs, because the system cannot control which publishers become available at any given moment.

Divorcing the advertiser and publisher payments makes economic sense, because the value of learning involves a future value. Moreover, it appears to be impossible to trace the value of any specific bit of learning, because the use of knowledge is so diffuse, and depends so heavily on the hypothetical of what would have transpired had some different ad been run. In contrast, it is quite easy to establish what payment would have been made to the publisher, absent any VOL considerations, and just pay the publisher that amount. The incentive effects of such a system are good: Publishers get the immediate second-price; advertisers are selected to maximize social welfare; and the advertisers' bids reflect their actual value.

The only question, then, is whether the learning account requires a subsidy. If the system at least breaks even, the mechanism is sustainable. Unfortunately, the need for a subsidy appears dependent on the exact method of learning employed. There are reasons to be both optimistic and pessimistic about the eventual surplus in the learning account. On the pessimistic side, we have a system with private information on the part of advertisers; efficiency in the one-shot mechanism is possible but may have no slack: The mechanism assigns all of the gains from trade to the advertisers, and the publishers break even. This lack of slack may mean that efficiency requires getting the solution exactly right. On the optimism side, there are welfare gains from efficient assignment that should generate net revenues. Moreover, the single publisher analysis suggests that eventually the learning account produces positive net revenue profits for the system: Eventually $VOL_2 - VOL_1 > 0$.

For some machine learning algorithms, the mechanism produces positive profits. In particular, for *upper confidence bound* (UCB) learning with zero discounting, the mechanism makes a positive profit for sufficiently large durations and negligible discounting.

UCB learning entails fixing a constant k , estimating the standard deviation for each ad i (σ_i), and running the ad with the highest value of $eCPM_i + k\sigma_i$. One can think of UCB learning as approximating the value of learning by the parameter k times the standard deviation. In practice, UCB entails picking one ad, say 1, to run for a while, which yields learning about that ad, so that its standard deviation falls¹⁰ until such point that $eCPM_1 + k\sigma_1 = eCPM_2 + k\sigma_2$. At this point, UCB alternates between 1 and 2, keeping the value of $eCPM_i + k\sigma_i$ approximately equal. Note that $eCPM$ is continually updated during this learning. The value of $eCPM_i + k\sigma_i$ tends to fall because $eCPM_i$, while random, tends to the true value, and σ_i tends down to zero. At some point $eCPM_1 + k\sigma_1$ and $eCPM_2 + k\sigma_2$ may fall to the level of $eCPM_3 + k\sigma_3$, at which point ad 3 is brought into the rotation. Bad outcomes about one of the advertisements may cause it to fall out of the rotation.

This process continues until a winning ad is identified and its variance is approximately zero. At that point, the winning ad runs nearly 100% of the time, even though it remains tied with $eCPM_i + k\sigma_i$ for other ads in the rotation, whose frequency diminishes.

Under UCB learning, then, VOL terms are negatively correlated to $eCPM$ for all ads that are run; indeed, they approximately sum to a constant. Thus, the only systematic losses that are sustained by the system occur in the initial phase when there is a leader, which, if there are T periods, turns out to be less than a constant times \sqrt{T} ; this is a vanishingly small fraction of the total time.¹¹

Thus, at least for one popular learning technology—a mechanism that charges advertisers to induce efficient learning, but pays publishers based on the one-shot value of their impression—makes a profit. UCB has some attractive properties but is generally flawed in the advertising context because it doesn't handle features well, and feature-based learning mechanisms are critical to matching ads and opportunities.

6 Illustration of UCB

Consider four advertisers, with $eCPM$ values of 0.35, 0.37, 0.39, and 0.41. These $eCPM$ values are modeled as click-rates, with equal values conditional on a click, so that the $eCPM$ is the true probability of a Bernoulli random variable, the payoff from running the ad. We estimate the click-rates by running each ad twenty times, as an initialization. The estimate of the standard deviation is set at 1/2 divided by the square root of the number of times that the ad has run, which is appropriate for a binomial random variable. UCB suggests picking the ad with the highest estimated click-rate plus a constant times the standard deviation; the constant was set at 2.¹²

¹⁰ In the constant $eCPM$ case, the standard deviation is approximately a constant over the square root of the number of times that the ad is run. Alternately, the estimated standard deviation can be a weighted average of recent variation.

¹¹ John Langford developed this insight about UCB. The argument depends on there being a clear best ad; if there is an $eCPM$ tie, the argument for a net surplus becomes a zero net surplus, which means long-run losses are possible since the future revenues are expected to be zero, and therefore insufficient to cover early losses.

¹² This constant is sufficient to insure that every ad is trialed, since the upper bound on $eCPM$ is 1 and the initial value of $k\sigma$ is 1.




Fig. 1 UCB values

A sample run is presented in Fig. 1. The number of trials is on the horizontal axis, with the UCB values on the vertical axis. The best ad is graphed in a solid black line, while the others are graphed in gray. Note that the ads have similar values of UCB and thus all four ads are in the rotation, although a gray ad has quite a temporary advantage around the 1,000th period. Even after 10,000 periods, the system has not yet converged. Indeed, black, which is the highest payoff, is not run for a substantial interval around period 9,000, owing to a relatively high showing by another ad.

The state of learning—the estimate of $eCPM$ —is presented in Fig. 2. The average click rates become quite close to the true values. Flat spots in this graph indicate extended periods where the advertisement is not running because its UCB is not highest.

The state of the learning account, the cumulative value expressed as a proportion of the payments made by advertisers to date, is illustrated in Fig. 3. The learning account initially makes money. This is because the standard deviations have been set equal to start, so that whatever begins in the lead pays more than the advertiser is paid in round 2; that is, $VOL_1 < VOL_2$. However, losses mount, in this run reaching 5% of revenue. Eventually, however, the best ad is run most of the time and has the lowest VOL, owing to having run most of the time, so the learning account becomes positive around period 7,000. The learning account eventually goes to a zero fraction of the revenue, although this takes a very long time indeed.

There is an enormous variance in the proceeds of the learning account, based only on the randomness of the outcomes. Some runs never produced negative revenue, and a few had system earnings of 15% of advertiser payments for extended periods—as long as 100,000 periods. Other runs produced substantial, sustained losses for extended periods of time—over 20,000 periods. It may be that profits in the limit




Fig. 2 Average experienced click-rates

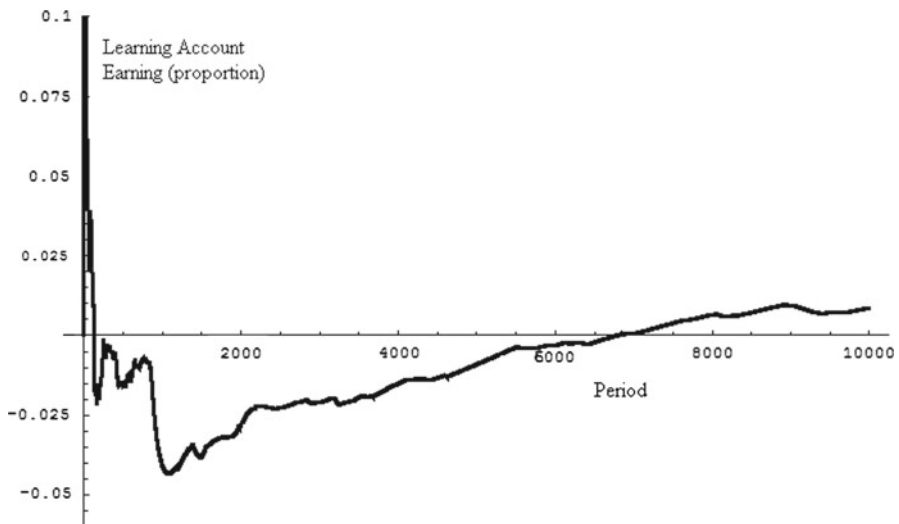


Fig. 3 Learning account (proportion of advertiser payments)

are cold comfort, especially in a world with discounting, given an apparent high variation in outcomes. However, an advantage of advertising as an application is that the flow of items is enormous—billions per day—so that limit results are perhaps more reasonable than in other settings. Discounting limits the value of learning, so that the losses should be attenuated by discounting.

7 Conclusion

The problem of matching advertisements and opportunities to advertise on web pages presents a remarkable opportunity to practice economics as an engineering discipline. The scale of the problem is unprecedented: arranging billions of transactions per day. The complexity of the problem is unprecedented: there are trillions of potentially-relevant product types. The speed of transactions is necessarily nearly instantaneous. The individual value of transactions is typically very small, requiring a “very low overhead” system.

This market design problem has already turned up several novel problems, which are problems that are likely to be important in other contexts. First, there is a winner’s curse type problem that arises naturally when products with differing pricing methods compete; some pricing methods naturally have higher variance than others, and these high variance estimates are more likely to be mis-estimated. Second, with dispersed data about common-values of items for sale, randomization becomes a natural part of behavior. It is worthwhile to build randomization into the system for the benefit of bidders with limited information. Third, the statistical uncertainty of machine learning creates externalities between transactions, and efficient markets require internalizing these externalities. One method of doing so is to separate payments from charges, which implies creating a learning account. Whether the learning account eventually reaches a surplus is as yet undetermined, although under one popular machine learning method the system does make a profit on average.

There are a variety of other quite important issues to confront in exchange design that is applied to advertising exchanges. For example, if the exchange is to make money, or even just recover costs, how should the exchange charge for its services? The major threat to an exchange is that other exchanges may attract the participants. Consequently, where possible it is useful for the exchange to charge where it adds value: charging participants only what they get by virtue of being in the exchange. In Yahoo!’s Right Media Exchange, many of the participants are themselves exchanges (or ad networks). The form that value-add pricing takes is to charge participants the surplus over what they would have gotten had their ad network not joined the exchange. Such a pricing policy creates a dominant strategy to join the exchange, although it may create other problems, such as incentives to form cartels.

Machine learning is opaque. A modern system may involve millions of variables, and there is no succinct answer to the question “why did my ad lose the bidding?” The lack of transparency of machine learning algorithms has recently spawned a vigorous discussion around search engines, which can be critical to online businesses, especially when adjustments to the algorithms dramatically change the fortunes of companies that appear in the results. In such circumstances, there is a tension between transparency and algorithmic accuracy, because complex algorithms will not appear transparent. In the search environment, the concerns typically revolve around fairness and bias.

In an exchange, in contrast, the tension may involve effective bidding versus effective algorithms. Typically, bids are optimized to the state of the system, which includes the algorithm as part of the environment. Increased complexity of the algorithm will degrade the effectiveness of the bidding, and specifically the speed at which the algorithm is revised will influence the appropriateness of the bids submitted. Algorithms

depend on bids that reflect value to induce efficiency. If improvements to the algorithm degrade the accuracy of the bids, the improvements to the algorithm may be lost.

It follows that improving bidding—making the bidders’ lives easier—is an important goal of market design. One simple way to improve the efficacy of bidding is to provide marketplace statistics like average prices over time. Such statistics provide a level of comfort to bidders, who don’t have to worry that their choices are extreme, and also guide optimization, by identifying relatively good deals. A second way to improve the efficacy of bidding is to provide counterfactuals, such as “had you bid b , this would have happened.” In providing both marketplace statistics and counterfactuals, it is important not to reveal inadvertently the behavior of any one participant, because that would adversely influence behavior.

Acknowledgments This paper is based on my keynote address to the International Industrial Organization Conference, April 9, 2011. I thank Kishore Papineni, David Reiley, and Serguei Vassilivskii for helpful comments.

Appendix: Derivation of (1)

$$\begin{aligned}
 0 &= \pi'(b) = (1-\alpha)^{n-1} (\mu-b) (n-1)F(b)^{n-2} f(b) + \sum_{k=1}^{n-1} \binom{n-1}{k} \alpha^k (1-\alpha)^{n-1-k} \\
 &\quad \times \left(bg(b)F(b)^{n-1-k} + \int_0^b vg(v)dv(n-1-k)F(b)^{n-2-k} f(b) - bd \left[G(b)F(b)^{n-1-k} \right] \right) \\
 &= (1-\alpha)^{n-1} (\mu-b) (n-1)F(b)^{n-2} f(b) \\
 &\quad + \sum_{k=1}^{n-1} \binom{n-1}{k} \alpha^k (1-\alpha)^{n-1-k} (n-1-k)F(b)^{n-2-k} f(b) \left(\int_0^b vg(v)dv - bG(b) \right) \\
 &= (1-\alpha)^{n-1} (\mu-b) (n-1)F(b)^{n-2} f(b) \\
 &\quad - (1-\alpha)(n-1)f(b) \left(\int_0^b G(v)dv \right) \sum_{k=1}^{n-2} \binom{n-2}{k} \alpha^k (1-\alpha)^{n-2-k} F(b)^{n-2-k} \\
 &= (1-\alpha)^{n-1} (\mu-b) (n-1)F(b)^{n-2} f(b) - (1-\alpha)(n-1)f(b) \\
 &\quad \times \left(\int_0^b G(v)dv \right) \left(\sum_{k=0}^{n-2} \binom{n-2}{k} \alpha^k (1-\alpha)^{n-2-k} F(b)^{n-2-k} - (1-\alpha)^{n-2} F(b)^{n-2} \right) \\
 &= (1-\alpha)^{n-1} \left(\mu-b + \int_0^b G(v)dv \right) (n-1)F(b)^{n-2} f(b) \\
 &\quad - (1-\alpha)(n-1)f(b) \left(\int_0^b G(v)dv \right) (\alpha + (1-\alpha)F(b))^{n-2}
 \end{aligned}$$

Thus,

$$(1-\alpha)^{n-2} \left(\mu - b + \int_0^b G(v)dv \right) F(b)^{n-2} = \left(\int_0^b G(v)dv \right) (\alpha + (1-\alpha)F(b))^{n-2}$$

or,

$$\begin{aligned} \left(\frac{(1-\alpha)F(b)}{\alpha + (1-\alpha)F(b)} \right)^{n-2} &= \frac{\int_0^b G(v)dv}{\mu - b + \int_0^b G(v)dv} = \frac{\int_0^b G(v)dv}{\int_0^\infty 1 - G(v)dv - \int_0^b 1 - G(v)dv} \\ &= \frac{\int_0^b G(v)dv}{\int_b^\infty 1 - G(v)dv} \end{aligned}$$

References

- Abdulkadiroglu, A., Parthak, P., Roth, A., & Sonmez, T. (2005). The Boston Public School match. In *AEA papers and proceedings* (pp. 368–371).
- Abraham, I., Athey, A., Babaioff, M., & Grubb, M. (2011). *Peaches, lemons, and cookies: Designing auction markets with dispersed information*. Harvard University (unpublished).
- Bax, E., Kuratti, A., McAfee, P., & Romero, J. (2011). *Comparing predicted prices in auctions for online advertising* (unpublished).
- Edelman, B., Ostrovsky, M., & Schwarz, M. (2007). Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97, 242–259.
- Hendricks, K., & Porter, R. (1988). An empirical study of an auction with asymmetric information. *American Economic Review*, 78, 865–883.
- Kenney, R., & Klein, B. (1983). The economics of block booking. *Journal of Law & Economics*, 26, 497–540.
- Li, S., Mahdian, M., & McAfee, P. (2010). Value of learning in sponsored search auctions. In *Proceedings of the 6th international workshop on internet and network economics* (pp. 294–305).
- Luenberger, D. (1979). *Introduction to dynamic systems: Theory, models and applications*. New York: Wiley.
- Manne, H. G. (1966). *Insider trading and the stock market*. New York: Free Press.
- McAfee, P., McMillan, J., & Wilkie, S. (2010a). The greatest auction in history. In J. J. Siegfried (Ed.), *Better Living Through Economics* (pp. 168–184). Cambridge, MA: Harvard University Press.
- McAfee, P., Papineni, K., & Vassilvitskii, S. (2010b). *Maximally representative allocations for guaranteed delivery advertising campaigns* (unpublished).
- McMillan, J. (1994). Selling spectrum rights. *Journal of Economic Perspectives*, 8, 145–162.
- Milgrom, P. (2000). Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108, 245–272.
- Roth, A. (2010). Deferred-acceptance algorithms: History, theory, practice. In J. J. Siegfried (Ed.), *Better Living Through Economics* (pp. 206–222). Cambridge, MA: Harvard University Press.
- Roth, A., Sonmez, T., & Unver, U. (2005). Kidney exchange. *AEA Papers and Proceedings*, 95, 376–380.
- Tietenberg, T. (2010). The evolution of emissions trading. In J. J. Siegfried (Ed.), *Better Living Through Economics* (pp. 42–58). Cambridge, MA: Harvard University Press.
- Varian, H. (2007). Position auctions. *International Journal of Industrial Organization*, 25, 1163–1178.
- Wilson, R. (1969). Competitive bidding with disparate options. *Management Science*, 15, 446–448.
- Wilson, R. (2002). Architecture of power markets. *Econometrica*, 70, 1299–1340.