



Competing Ad Auctions

Itai Ashlagi
Benjamin Edelman
Hoan Soo Lee

Working Paper

10-055

Copyright © 2010, 2011 by Itai Ashlagi, Benjamin Edelman, and Hoan Soo Lee

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Competing Ad Auctions

Itai Ashlagi^a, Benjamin Edelman^b, Hoan Soo Lee^b

^aSloan School of Management, MIT

^bHarvard Business School

Abstract

We present a two-stage model of competing ad auctions. Search engines attract users via Cournot-style competition. Meanwhile, each advertiser must pay a participation cost to use each ad platform, and advertiser entry strategies are derived using symmetric Bayes-Nash equilibrium that lead to the VCG outcome of the ad auctions. Consistent with our model of participation costs, we find empirical evidence that multi-homing advertisers are larger than single-homing advertisers. We then link our model to search engine market conditions: We derive comparative statics on consumer choice parameters, presenting relationships between market share, quality, and user welfare. We also analyze the prospect of joining auctions to mitigate participation costs, and we characterize when such joins do and do not increase welfare.

1 Introduction

Online ad auctions sell advertising placements on search engines and elsewhere—providing key funding for a variety of online resources. Advertisers sign up with one or more ad platforms, specify their advertising preferences (including conditions in which they want their ads to be shown, and how much they are willing to pay), and receive clicks from interested users.

In this paper, we explore competition among ad platforms that offer search engine advertising services. Our motivations are several. For one, we want to understand why some advertisers choose to use only certain ad platforms but not others. After all, if an ad network cannot attract a broad selection of advertisers, it will be unable to present ads related to users' requests, and it will also garner substantially lower revenue. Second, we want to explore competition across auction platforms. Auction competition is a relatively unexplored topic in the literature. The primary difficulty in considering auction competition is that classic competition models (e.g. Bertrand and Cournot)

Benjamin Edelman serves and Hoan Soo Lee previously served as consultant and research intern to Microsoft and Microsoft Research, respectively. But they write on their own behalf, not at Microsoft's request or for Microsoft's benefit.

view competition through the setting of market price, either directly or via production quantities. But ad auctions set prices through importantly different mechanisms: strategic interaction among advertisers yields price discovery, while an ad platform can influence outcomes through, e.g., setting of reserve prices.

Competition in online ad auctions also differs from models of auction competition examined in the literature. Pai (2009) explores optimal mechanisms for a seller employing a second-price auction with a reserve price in a general context. In contrast, we limit our auction competition to the search engine industry. This restriction puts useful structure on competition: in order to attract users, an ad platform seeks to reduce cost of search to users. But reducing the cost of search decreases the number of clicks users perform, thereby reducing payments from advertisers. Separately, Ellison et al. (2004) consider competition among auctions occurring at the level of strategic entry by auctioneers in two segmented markets with differing buyer-to-seller ratios. There, an auctioneer chooses between entering in a market saturated with buyers by expecting to compete with many sellers (competition effect) and selling to great magnitudes (scale effect). Applying this framework to online ad auctions implies that search engines compete by specializing on exclusive sets of keywords with correspondingly segmented sets of buyers. However, the leading ad platforms do not strategically limit keywords or even focus on distinctive sets of keywords. Rather, leading ad platforms sell clicks from web searchers searching for all manner of subjects, and ad platforms sell to a common pool of advertisers and would-be advertisers.

We model competing ad auctions in two stages. In the first stage, advertisers choose to enter one or more ad auctions after considering each search engine’s user base (capacity) and click-through rates (technology) in light of an exogenous cost of joining each ad auction. Using the symmetric Bayes-Nash equilibrium, we derive advertisers’ entry strategies that lead to the VCG outcome of the ad auctions. In the second stage, participating advertisers submit bids in the ad auction(s) they chose to enter, and ad platforms strategically compete on increasing capacity while preserving an optimal level of technology to maximize revenue. This follows a model of Cournot-style competition among ad platforms.

Closest to our paper is Liu and Chiu (2010) which studies competition across ad auctions with auction capacities endogenized by consumer choice. However, they abstract away from the mechanisms that allocate and price advertising positions. In contrast, our approach is grounded in the unusual mechanisms used in selling online advertising.

We proceed in five parts. In Section 2, we develop model fundamentals and notation. In Section 3, we model ad auctions with participation costs—offering an initial explanation of why not all advertisers use all ad platforms, and testing that explanation with available data. In Section 4, we analyze competing ad auctions in light of user preferences over search engines. In Section 5

we develop a concept of “joining” two ad auctions, and we identify conditions where joins increase or reduce advertiser welfare. In Section 6, we conclude.

2 Ad auctions - GSP and VCG

In an *ad auction*, there are advertisers $\mathcal{N} = \{1, \dots, N\}$ and K slots for sale. Each advertiser can receive at most one slot. The positions are sold for a single period of time. Each slot k has an expected *click-through rate* $\alpha_k > 0$. The auction has a known capacity $C > 0$. Thus, if an advertiser wins slot k , the advertiser will receive $C\alpha_k$ clicks in expectation. We assume that $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_K$. Define $\alpha_k = 0$ for every $k > K$.

The *value* per click for advertiser j is $v_j \in [0, 1]$. Advertisers are risk neutral, and the payoff to advertiser j for winning slot k is $C\alpha_k v_j$ minus its payments to the ad platform. We also assume that for each advertiser j , v_j is drawn from a commonly known distribution F , and the value of each advertiser is private information.

Each advertiser j is required to submit a bid b_j . We denote by $b_{(j)}$ the j^{th} highest bid. Similarly we denote by $g(j)$ the identity of the j^{th} highest advertiser. In case of ties, the order among those advertisers is determined randomly.

We consider payments and outcomes under two distinct auction mechanisms. Modern ad platforms generally use the *Generalized Second Price* (GSP) structure, wherein an advertiser $g(j)$ receiving position k pays a total of $C\alpha_j b_{(j+1)}$. In the *Vickrey-Clarke-Groves* (VCG) ad auction, each advertiser pays its impact on all others’ social welfare, assuming bids equal values. Hence, under VCG, advertiser $g(j)$ receiving position k pays a total of $p^{VCG,k} = \alpha_k C \sum_{j=\min(K+1,n)}^{\min(k+1,n)} (\alpha_{j-1} - \alpha_j) b_{(j)}$. Both GSP and VCG ad auctions allocate the first position to the highest-bidding advertiser, the second position to the second-highest bidder, and so forth.

Bidding truthfully is a dominant strategy for every advertiser in the VCG ad auction. (See Vickrey (1961), Clarke (1971), Groves (1973).) Edelman et al. (2007) and Varian (2007) show that there exists an equilibrium in the GSP ad auction under complete information such that the GSP outcome coincides with the outcome of the VCG auction in which each advertiser bids truthfully. Furthermore, this equilibrium yields the lowest revenue for the seller, so it is in some sense best for the advertisers. Finally, Cary et al. (2007) shows that a reasonable myopic GSP bidding strategy converges to this equilibrium.

3 Ad Auctions with Participation Costs

Consider a GSP ad auction, and suppose there is a cost $Z > 0$ for each advertiser to participate in the auction; Z can be interpreted as an advertiser’s transaction cost in submitting its campaign into the ad platform. Elements of Z include creating an account, setting advertising parameters, monitoring effectiveness, adjusting bids, and paying bills. Z is not transferred to the auctioneer. Therefore, a high participation cost, relative to a platform’s capacity, will cause an advertiser to forego use of that platform, even if the advertiser would otherwise find it profitable to use that platform. For example, if the following inequality holds

$$\frac{Z}{C} > \alpha_1 v, \tag{1}$$

then an advertiser with value v will never enter the auction since the advertiser would realize negative utility even if he managed to receive the first slot with zero payment.

In this section we show that there exists a unique threshold function that determines whether an advertiser will participate in a given ad auction. In particular, consider the following two-stage game: In the first stage, each advertiser decides whether to participate in the auction, as a function of the advertiser’s value v . The value of each advertiser is private information (whether or not the advertiser enters the auction). In the second stage, all advertisers that decide to participate in the auction submit a bid. The advertisers who enter face the same conditions as in Edelman et al. (2007) and Varian (2007), and the outcome coincides with the VCG outcome of the ad auction with these advertisers.¹

A strategy for an advertiser j is a function of $s : [0, 1] \rightarrow \{0, 1\}$ where $s_j(v_j) = 1$ means that advertiser j enters the auction given that its value v_j , and $s_j(v_j) = 0$ means j does not enter. We employ symmetric Bayes-Nash equilibrium as our benchmark equilibrium concept when further analyzing advertisers’ strategies.

Denote by $U(v, n)$ the expected utility (before considering entry cost Z) of an advertiser with valuation v , who decides to participate in the auction given that exactly n other advertisers also participate. For simplicity, we assume that if an advertiser’s expected utility is 0 (including costs), the advertiser prefers to enter the auction. Thus if all advertisers use the strategy s^* , then an advertiser with valuation v will choose to enter the auction if and only if

$$E_n[U(v, n)|s^*] \geq Z. \tag{2}$$

After advertisers decide whether to enter, a set of advertisers B bid in the auction, and their equilibrium bids coincide with the GSP equilibrium in Edelman et al. (2007) and Varian (2007).

¹In Ashlagi et al. (2010) the authors do not model participating costs, but rather let advertisers choose a single ad auction to participate in.

Theorem 3.1. *Suppose $U(1, N) > Z$. Then there exists a unique strategy s^* which forms a symmetric Bayes-Nash equilibrium in the two-stage game. In particular there exists a threshold $v^* > 0$ such that for each $v \in [0, 1]$, advertiser j*

$$s_j^*(v) = 1 \quad (3)$$

if $v \geq v^*$ and

$$s_j^*(v) = 0 \quad (4)$$

otherwise.

To prove theorem 3.1, the following lemma will be useful. The lemma provides two monotonicity properties: the expected utility for an advertiser who enters the auction is non-increasing in the number of participants, and is increasing in the advertiser's value.

Lemma 3.2. *For any $n < N$, and any v :*

1. $U(v, n) > U(v, n + 1)$.
2. $\partial U(v, n) / \partial v > 0$.

Proof. Let $Q(v, n)$ and $P(v, n)$ denote the expected number of clicks and expected payment for an advertiser with value v that decides to enter auction l given that there are $n - 1$ other advertisers in the auction. Thus the expected utility of such an advertiser is

$$U(v, n) = vQ(v, n) - P(v, n). \quad (5)$$

Note that

$$Q(v, n) = C \sum_{1 \leq k \leq n} \binom{n-1}{k} \alpha_k \tilde{F}(v)^{n-k} [1 - \tilde{F}(v)]^{k-1}. \quad (6)$$

where $\tilde{F}(v)$ is the probability that a bidder that entered the auction has a valuation lower than v .²

Let X_n be a random variable distributed over $\{\alpha_1, \dots, \alpha_N\}$ such that $P(X_n = \alpha_k) = \binom{n-1}{k} p^{n-k} (1-p)^{k-1}$ where $p = \tilde{F}(v)$, and let $Y_n \sim \text{Bin}(n-1, p)$, i.e. $P(Y_n = k) = \binom{n-1}{k} p^{n-k} (1-p)^{k-1}$. Note that $P(X_n = \alpha_k) = P(Y_n = k)$. For every $m > 1$, define

$$\mu(m, n) = (\alpha_m - \alpha_{m-1}) P(Y_n \geq m - 1),$$

and let $\mu(1, n) = \alpha_1 P(Y_n \geq 0)$. Observe that $Q_l(v, n) = CE[X_n] = C \sum_{m \geq 1} \mu(m, n)$ and that $\mu(1, n) = \alpha_1$. By our assumption on click-through rates, $\alpha_{k+1} - \alpha_k \leq 0$. Moreover, $P(Y_n \geq k)$ is

²We make no assumption on the entry strategy profile, except requiring that it induce a measurable distribution \tilde{F} .

decreasing in n because more trials with the same probability of success will lead to more successes. Therefore, for all $m > 1$, every $\mu(m, n)$ is non-increasing in n , implying that $Q(v, n) > Q(v, n + 1)$. Finally, by Myerson (1981),

$$U(v, n) = \int_0^v Q(x, n) dx, \quad (7)$$

implying that $U(v, n) > U(v, n + 1)$.

The second part follows because the derivative of $U(v, n)$ with respect to v is equal to $Q(v, n)$ by (7) and because $\partial Q(v, n)/\partial v > 0$. \square

Proof of Theorem 3.1:

Fix the entry strategies of all other advertisers except j . By Lemma 3.2, j is better off entering given some fixed value v^* , as long as $v_j \geq v^*$. Thus, every symmetric Bayes-Nash equilibrium is characterized by a threshold value: there exists a $v^* \in [0, 1]$ such that every advertiser will enter the auction if and only if its value is at least v^* .

Fix a symmetric strategy profile that is characterized by a threshold s^* . To show existence of the threshold value, define

$$G(v) = E_n[U(v, n)|s^*] - Z.$$

By Lemma 3.2, $G(v)$ is continuous and strictly increasing in v and by assumption. Furthermore, $G_l(0) = -E_n[P(0, n)|s^*] - Z < 0$, and by assumption $G_l(1) = E_n[U(1, n)|s^*] > U(1, N) > Z > 0$. Therefore there exist v^* such that $G(v^*) = Z$. Since $G_l(v)$ is strictly increasing in v , v^* is unique. \square

The preceding proof indicates that the cut-off value v^* is a function of the platform's primitives, $Z, C, \alpha_1, \alpha_2, \dots, \alpha_k$ which uniquely define the auction. Comparative statics follow directly:

$$\frac{\partial v^*}{\partial C} < 0 \quad (8)$$

$$\frac{\partial v^*}{\partial Z} > 0 \quad (9)$$

$$\frac{\partial v^*}{\partial \alpha_k} < 0 \quad \forall k. \quad (10)$$

3.1 Advertiser Size and Multi-homing

Consider now a set of $L > 0$ GSP auctions, indexed by $l = 1, \dots, L$. Auctions may differ in both click-through rates and capacity. For simplicity, let participation cost $Z > 0$ remain fixed, although our results can easily be extended to consider varying participation costs. We add a subscript l for relevant parameters of each auction l .

Observe that the symmetric equilibrium in each auction form a symmetric equilibrium in the *extended game* in which each advertiser chooses which auction(s) to enter. Furthermore, the argu-

ment in the previous section implies the uniqueness of this equilibrium. In particular, advertisers’ entry decisions are independent.

Corollary 3.3. *Let v_1^*, \dots, v_L^* be the cut-off values corresponding to the unique symmetric Bayes-Nash equilibria as in Theorem 3.1 and let B_l be advertisers entering auction l in this equilibrium. For each $v_l^*, v_{l'}^* \in V^*$ and $B_l, B_{l'} \in B$, if $v_l^* \geq v_{l'}^*$, then $B_l \subseteq B_{l'}$.*

Proof. Pick any v_l^* and $v_{l'}^*$ in V^* such that $v_l^* \geq v_{l'}^*$. Then any advertiser with value $v_j \geq v_l^*$ will enter auction l , hence $j \in B_l$. Moreover, by assumption, $v_j \geq v_{l'}^*$ hence $j \in B_{l'}$. \square

Corollary 3.3 offers a testable implication: Large advertisers multi-home because they can spread participation cost Z across a large volume of ad purchases, whereas small advertisers find the participation costs too large to justify signing up with smaller platforms.

We test these claims with data from multiple services that track and preserve advertising at multiple ad platforms. We report normalized advertiser sizes based on independent observations by two different ad monitoring services.³ Based on data from a first data collection service, Table 1 compares the size (impression count, normalized with maximum value set to 1,000) of advertisers that use one, two, or all three of the ad platforms we examine. Based on data from a second service, Table 2 measures the size of multi-homing and non-multihoming advertisers by other metrics: Proportion of sites achieving an Alexa ranking (data available only for approximately the 25 million most popular sites on the web), average rank (bottom-coding unranked sites at a rank of 40 million), as well as average “reach” (number of users visiting the site) and average page views.⁴ Figure 1 plots the distribution of reach by multi-homing and non-multihoming advertisers.

By each metric, these tables are consistent with Corollary 3.3 and our model of advertisers’ participation costs. In particular, the advertisers that purchase ads from all three platforms are strikingly larger than the advertisers that purchase ads only from one or two platforms: They buy more ad impressions (Table 1) and are more likely to be ranked by Alexa, achieve a lower average rank (i.e. greater traffic), larger reach, and more page-views (Table 2) Meanwhile, the advertisers who choose to use only Google are the smallest by far—further confirming that small advertisers tend not to multi-home. Finally, Figure 1 confirms that triply-multihoming advertisers have pointwise larger reach than double-multihoming advertisers which in turn have pointwise larger reach than single-homing advertisers. We also verify this relationship through a simple ordered probit regression. First, we rank combinations of search engines to form a presence rank—reflecting, for example, that Google is larger than Yahoo and Microsoft together. See Table 3. Then, we run an ordered probit of each advertiser’s presence rank on its web site rank, reach,

³By agreement with our data sources, we do not report their names.

⁴Reach and page-views are reported per thousand users.

and page-views. See Table 4. The coefficients on reach, rank, and page-views are all positive and significantly different from zero.

4 Competing Auctions

In this section, we explore the micro-foundation of ad auction competition. Our approach is grounded in a fundamental tradeoff facing each search engine: On one hand, a search engine must provide high-quality results to satisfy users’ requirements as quickly as possible. On the other hand, a search engine will reap greater revenues if it designs its listings to cause users to click more advertisements. Google early recognized this tradeoff: In their seminal 1998 paper, Google co-founders Brin and Page flagged an “inherent” incentive for a search engine to reduce the quality of its algorithmic results or otherwise encourage users to click on advertisements (Brin and Page (1998)). The tradeoff remains timely: For example, when launching its new “Instant” search service, Google touted a savings of 2 to 5 seconds per search. Such time-savings could benefit users, but some advertisers found that Instant Search reduced clicks on their advertisements (?). Meanwhile, (Edelman (2010)) points out that for some search terms, advertisements can be distracting or affirmatively harmful—providing a further tradeoff between satisfying users and increasing revenue.

We formalize search engines’ response to the quality/revenue tradeoff via a *search cost* parameter, s . A greater search cost implies more advertisement clicks and more revenue to the search engine. Conversely, a search engine with a lower search cost gives users the “right” links more quickly, yielding fewer advertisement clicks but increasing user satisfaction and attracting, all else equal, more users.

Following Hotelling’s model of consumers distributed on a circle, we build a model of consumer choice over search engines. Through this model, we endogenize a search engine’s user base (capacity), and we incorporate a search cost parameter to capture the tradeoff between increasing consumer welfare versus increasing search engine revenue.

4.1 Consumer Choice

Suppose consumers are uniformly distributed around a circle with unit circumference. Suppose two search engines, A and B , occupy diametrically opposite locations on the circle. (The result can be extended to three or more engines.) A consumer who accesses search engine $l \in \{A, B\}$ receives value δ_l , less the search engine’s cost of search, s_l . To access search engine l , the consumer incurs a cost td (where d reflects the consumer’s distance from l and t gives a coefficient on distance). Thus, a consumer chooses engine A if the consumer is located at distance d such that:

$$\delta_A - s_A - td \geq \delta_B - s_B - t(1/2 - d). \tag{11}$$

Given a uniform distribution of consumers, the fraction of consumers choosing search engine A is $\frac{1}{t}(\Delta\delta - \Delta s + \frac{1}{2}t)$ where $\Delta\delta = \delta_A - \delta_B$. Similarly, the fraction of consumers choosing search engine B is $\frac{1}{t}(\Delta s - \Delta\delta + \frac{1}{2}t)$. As a result, A 's capacity is $C_A = \frac{C}{t}(\Delta\delta - \Delta s + \frac{1}{2}t)$ where C is the total number of consumers in the market. Thus, search engine A 's user base decreases as search cost increases, but increases in the value derived from use.

4.2 Ad Platform Selection of Search Cost

At the start of the second stage, advertisers have committed to the bidding process, and the number of entrants, $n^* = |B_l|$, and the pool of advertisers B_l (with valuations $\{v_1, \dots, v_{n^*}\}$) are fixed. Then the revenue to search engine l is:

$$C_l(s_l) \sum_{j=1}^{n^*} \sum_{k=j}^{n^*} (\alpha_{k,l} - \alpha_{k+1,l}) v_{(k+1)}$$

which is equivalent to:

$$C_l(s_l) \sum_{k=1}^{n^*} (\alpha_{k,l} - \alpha_{k+1,l}) k v_{(k+1)}.$$

Since search cost is proportional to the expected number of clicks required to reach the desired information, any weighted aggregate click-through rate is strictly increasing in s_A : $\frac{\partial \sum_{k \in K} w_k \alpha_{kl}}{\partial s_A} > 0$ for any $\sum w_k = 1$. Moreover, assume that this function is linear, i.e. $\frac{\partial^2 \sum_{k \in K} w_k \alpha_{kl}}{\partial s_A^2} = 0$.

For tractability, suppose the search engine's technology is exponential (subsection 5.2). Then $\alpha_{k,l} = \beta_l^k$. Define $\bar{V} = \sum_{k=1}^{n^*} k v_{(k+1)}$ and $w_k = k v_{(k+1)} / \bar{V}$.

Then we can rewrite the revenue as:

$$C_l(s_l) (1 - \beta_l) \bar{V} \sum_{k=1}^{n^*} w_k \beta_l^k.$$

By the linearity assumption, we can parameterize the sum as $\sum_{k=1}^{n^*} w_k \beta_l^k = a_l s_l + b_l$, yielding the simplified expression:

$$C_l(s_l) (1 - \beta_l) \bar{V} (a_l s_l + b_l).$$

This revenue function has a unique internal maximum attained by the following first order conditions:

$$s_l = \frac{1}{1 + a_l} (\Delta\delta + \frac{1}{2}t - b_l + s_{l'}) \quad (12)$$

$$s_{l'} = \frac{1}{1 + a_{l'}} (-\Delta\delta + \frac{1}{2}t - b_{l'} + s_l). \quad (13)$$

4.3 Comparative Statics

We now consider two special cases: two search engines have pointwise equal technologies ($a_l = a_{l'}, b_l = b_{l'} = b$), and one search engine pointwise dominates the other ($a_l > a_{l'}, b_l > b_{l'}$).

4.3.1 Pointwise Equality

Consider a search engine that is identical to its competitor ($a_l = a_{l'} = a, b_l = b_{l'} = b$) except that it enjoys a single advantage: a higher value delivered to consumers during search ($\delta_l > \delta_{l'}$). Such a search engine can leverage that strength by raising search cost while retaining higher capacity. The following theorem formalizes that advantage:

Theorem 4.1. *In the two-stage game with two auction platforms l and l' with pointwise identical baseline technologies, if $\delta_l \geq \delta_{l'}$, then in equilibrium: $s_l \geq s_{l'}, C_l \geq C_{l'}, \sum_{k \in K} w_k \alpha_{kl} \geq \sum_{k \in K} w_k \alpha_{kl'}, v_l^* \leq v_{l'}^*$ and $B_l \supseteq B_{l'}$.*

Proof. The first order conditions (12) and (13) require $s_l = \frac{\Delta\delta}{2+a} + \frac{\frac{1}{2}t-b}{a}$ and $s_{l'} = \frac{-\Delta\delta}{2+a} + \frac{\frac{1}{2}t-b}{a}$. The resulting capacities are:

$$C_l^* = C \left(\frac{1}{2} + \frac{a\Delta\delta}{t(2+a)} \right) \quad (14)$$

$$C_{l'}^* = C \left(\frac{1}{2} - \frac{a\Delta\delta}{t(2+a)} \right). \quad (15)$$

Thus, capacity, technology and search cost are all monotonic in δ ; if $\delta_j \geq \delta_{j'}$ then $C_l \geq C_{l'}$, $\sum_{k \in K} w_k \alpha_{kl} \geq \sum_{k \in K} w_k \alpha_{kl'}$ and $s_l \geq s_{l'}$.

The threshold values and advertiser set relation follow from Corollary 3.3. \square

4.3.2 Pointwise Dominance

Consider a search engine l that enjoys technology superior to its competitor l' : $a_l > a_{l'} \geq 0, b_l > b_{l'} \geq 0$. In that case, search engine l receives greater market share even though its value to consumers matches its competitor. In particular, l and l' pick search costs s such that

$$\Delta s = \frac{1}{a_{l'} + a_l + a_{l'} a_l} \left[(a_{l'} + a_l) \Delta\delta + (a_{l'} - a_l) \frac{t}{2} - a_{l'} b_l + a_l b_{l'} \right]. \quad (16)$$

Suppose $\Delta\delta = 0$. Then $\Delta s < 0$ but $\Delta C > 0$ since $a_l > a_{l'} > a_{l'} \left(\frac{t-2b_l}{t-2b_{l'}} \right)$. That is, l sets lower search cost and enjoys greater market share.

Alternatively, suppose a search engine l enjoys pointwise dominant technology (as in subsection 4.3.1), but gives consumers less value than its competitor l' . (That is, $\delta_l < \delta_{l'}$.) Then l will receive lower market share ($C_l < C_{l'}$) if its technology is not sufficiently greater than l' 's. Formally, $\Delta C < 0$ if and only if $2a_{l'} a_l \Delta\delta + a_l < a_{l'} \left(\frac{t-2b_l}{t-2b_{l'}} \right)$.

5 Joining Auctions

In this section, we establish a concept of “joining” auctions such that their available positions are pooled, and all advertisers participating in one auction automatically participate in the other. What happens to advertiser welfare and ad platform revenue if two ad auctions are joined? These questions take on special relevance in light of the 2009 partnership between Microsoft and Yahoo, as well as a 2008 proposed partnership between Google and Yahoo (ultimately aborted after antitrust regulators raised concerns).

We begin with several definitions.

Definition 5.1. *A set of click-through rates has the property of diminishing differences if $\alpha_k - \alpha_{k+1} \geq \alpha_{k+1} - \alpha_{k+2}$ for each $k \leq K$.*

Definition 5.2. *Auction l and auction l' join to form auction \tilde{l} if auction \tilde{l} has capacity $C_{\tilde{l}}$ such that*

$$\max\{C_l, C_{l'}\} \leq C_{\tilde{l}} \leq C_l + C_{l'} \quad (17)$$

and if auction \tilde{l} has click-through rate

$$\alpha_{k\tilde{l}} = \max\{\alpha_{kl}, \alpha_{kl'}\} \quad (18)$$

where α_{kl} ($\alpha_{kl'}$) is the click-through rate of slot k in auction l (l').

By taking the capacities of each auction as the size of the set of consumers choosing the engine, inequality (17) bounds $C_{\tilde{l}}$. Because some consumers use multiple search engines, we allow for overlap of capacity between two ad platforms—meaning a joined ad platform might have less capacity than the sum of capacities of its contributors. Furthermore, we assume that no consumer of either engine is lost upon join.

When auctions join, what click-through rates result? We envision ad platforms choosing the best components of each contributor, which implies click-through rates given by the stronger of the joining platforms. Hence the approach in (18).

5.1 Joining Auctions to Make All Advertisers Better Off

In the following theorem, we provide a condition in which a joined auction offers a sufficient improvement in capacity and technology to make every advertiser weakly better off ex post.

Theorem 5.3. *Suppose auctions l and l' join to form auction \tilde{l} . Let $C_l \geq C_{l'}$ and $\alpha_{kl'} \geq \alpha_{kl}$ for each $k \leq K$. Then every advertiser is weakly better off ex post if:*

$$C_{\tilde{l}} \geq \frac{C_{l'}\alpha_{kl'} + C_l\alpha_{kl}}{(N-k)\alpha_{k+1,l'} - (N-k-1)\alpha_{kl}} \quad \forall k \leq K. \quad (19)$$

Proof. Suppose auction l and l' join to form \tilde{l} and, without loss of generality, let $n_{\tilde{l}}^* \geq n_{l'}^*$.

Consider an advertiser who does not receive any position in either auction before the join. Such an advertiser is ranked in position $k > n_{\tilde{l}}^*$. Because that advertiser already achieves utility of 0, the joined auction \tilde{l} cannot make it worse off.

Consider an advertiser who receives a position in each of the auctions before the join. We introduce new notation to characterize ex post utility: denote by $u_l(v, k)$ the utility of a advertiser with value v who wins position k in auction l . Under VCG, we have:

$$u_l(v, k) = C\alpha_{kl}v - C \sum_{n^* > j \geq k} (\alpha_{jl} - \alpha_{j+1,l})v_{(j+1)} \quad (20)$$

where $\alpha_{kl} = 0$ for $k \geq K$.

Prior to the join, an advertiser with valuation v who receives positions in both auctions must receive position $k \leq n_{l'}^*$ in auction l' , yielding utility $u_l(v, k) + u_{l'}(v, k)$.

Using assumption (19) we get:

$$\begin{aligned} C_{\tilde{l}}\alpha_{k\tilde{l}}v - C_l\alpha_{kl}v - C_{l'}\alpha_{kl'}v &\geq C_{\tilde{l}}(N-k)(\alpha_{k\tilde{l}} - \alpha_{k+1,\tilde{l}})v \geq \\ &C_{\tilde{l}} \sum_{n_{\tilde{l}}^* > j \geq k} (\alpha_{j\tilde{l}} - \alpha_{j+1,\tilde{l}})v_{(j+1)} \geq C_{\tilde{l}} \sum_{n_{\tilde{l}}^* > j \geq k} (\alpha_{j\tilde{l}} - \alpha_{j+1,\tilde{l}})v_{(j+1)} - S, \end{aligned} \quad (21)$$

where $S = C_{l'} \sum (\alpha_{jl'} - \alpha_{j+1,l'})v_{(j+1)} + C_l \sum (\alpha_{jl} - \alpha_{j+1,l})v_{(j+1)}$.

The last inequality (21) implies:

$$\begin{aligned} C_{\tilde{l}}\alpha_{k\tilde{l}}v - C_{\tilde{l}} \sum (\alpha_{j\tilde{l}} - \alpha_{j+1,\tilde{l}})v_{(j+1)} &\geq \\ C_l\alpha_{kl}v - C_l \sum (\alpha_{jl} - \alpha_{j+1,l})v_{(j+1)} + C_{l'}\alpha_{kl'}v - C_{l'} \sum (\alpha_{jl'} - \alpha_{j+1,l'})v_{(j+1)}, \end{aligned}$$

which is equivalent to:

$$u_{\tilde{l}}(v, k) \geq u_l(v, k) + u_{l'}(v, k). \quad (22)$$

Therefore an advertiser receiving a position in both auctions is better off after the join.

Consider an advertiser with valuations v who receives position k in auction l with $n_{\tilde{l}}^* \geq k > n_{l'}^*$. Such an advertiser will have ex post utility of $u_l(v, k)$. Since we have already shown (22), such an advertiser, who wins a single position, is also better off after the join. \square

The conditions in Theorem 5.3 stipulate intuitive requirements for advertisers to gain from a joined ad auction: the resulting click-through rate and auction capacity must be sufficiently

improved relative to the offerings of the ad auctions when separate. First, the auction with fewer advertisers must add value to the join through a point-wise larger click-through rate. Second, there must be minimal overlap between the two auctions, so that $C_{\tilde{l}}$ is sufficiently larger than both C_l and $C_{l'}$. It is necessary for $C_{\tilde{l}}$ to be sufficiently large so that advertisers in auction l (who already face higher prices due to more advertisers in l) gain sufficiently from joining the two auctions.

5.2 Example with Exponential Click-through Rates

Suppose click-through rates $\alpha_{kl} = \beta_l^k$ for each $k \leq K$ where $\beta_l < 1$. Note that these exponential click-through rates obey both monotonic ordering of α_{kl} and the *diminishing differences* property. Following the framework in Theorem 5.3, we set $\beta_{\tilde{l}} = \beta_{l'} > \beta_l$. Condition (19) then becomes:

$$C_{\tilde{l}} \geq \frac{C_{l'}\beta_{l'}^k + C_l\beta_l^k}{(N-k)\beta_{l'}^{k+1} - (N-k-1)\beta_{l'}^k} \quad \forall k \leq K.$$

Simplifying and rearranging:

$$C_{\tilde{l}} - C_{l'} - C_l \frac{\beta_l^k}{\beta_{l'}^k} \geq C_{\tilde{l}}(1 - \beta_{l'})(N - k). \quad (23)$$

Because $\beta_l < \beta_{l'}$, we know:

$$C_{\tilde{l}} - C_{l'} - C_l \frac{\beta_l}{\beta_{l'}} \leq C_{\tilde{l}} - C_{l'} - C_l \frac{\beta_l^k}{\beta_{l'}^k}$$

and

$$C_{\tilde{l}}(1 - \beta_{l'})(N - 1) \geq C_{\tilde{l}}(1 - \beta_{l'})(N - k).$$

Thus the generalized restrictions in equation (23) become:

$$C_{\tilde{l}} - C_{l'} - C_l \frac{\beta_l}{\beta_{l'}} \geq C_{\tilde{l}}(1 - \beta_{l'})(N - 1) \quad (24)$$

because this condition implies the required restrictions for each $k \leq K$.

By the subadditive property of joined auction \tilde{l} and equation (24):

$$C_{l'} + C_l \geq C_{\tilde{l}} \geq C_{\tilde{l}}(1 - \beta_{l'})(N - 1) + C_{l'} + C_l \frac{\beta_l}{\beta_{l'}}$$

which requires:

$$C_{l'} + C_l \geq C_{\tilde{l}}(1 - \beta_{l'})(N - 1) + C_{l'} + C_l \frac{\beta_l}{\beta_{l'}}.$$

Rearranging:

$$\frac{C_l}{C_{\tilde{l}}} \geq (N - 1) \frac{1 - \beta_{l'}}{1 - \frac{\beta_l}{\beta_{l'}}} \quad (25)$$

Since $C_l \leq C_{\bar{l}}$ and $N \geq 2$ (so that the RHS of (25) is not trivially 0), inequality (25) requires:

$$1 - \frac{\beta_l}{\beta_{l'}} > 1 - \beta_{l'}.$$

Rearranging yields the requirement:

$$\beta_{l'} > \sqrt{\beta_l}. \quad (26)$$

Condition (26) requires that the auction with superior technology have $\beta_{l'}$ sufficiently larger than β_l . Furthermore, if N is very large, then $\beta_{l'}$ must approach 1 in order to satisfy the requirement in (25). Note that these are necessary conditions, but not sufficient conditions.

5.3 Joining Auctions that Make Some Advertisers Worse Off

In this section we show that it a joined auction can negatively affect overall advertiser welfare. Consider the following definition.

Definition 5.4. *Auction l is uniformly stronger than auction l' if and only if $C_l \geq C_{l'}$, and $\alpha_{kl} \geq \alpha_{kl'}$ for each k . We denote this by $A_l \geq_{us} A_{l'}$.*

If one of the auctions is uniformly stronger than the other and if the resulting capacity remains equal to its original capacity (i.e. the uniformly weaker auction represents a subset of consumers of the stronger auction), then joining the auctions will make some advertisers weakly worse off. The following theorem identifies sufficient conditions that make advertisers worse off:

Theorem 5.5. *Suppose $A_l \geq_{us} A_{l'}$, and auctions l and l' join in the manner of Definition 5.2. If $C_{\bar{l}} = C_l$, then any advertiser that wins in both auctions is worse off.*

Proof. The assumption implies that $\alpha_{k\bar{l}} = \alpha_{kl} \geq \alpha_{kl'}$ for each $k \leq K$ and $C_{\bar{l}} = C_l$. Then the joined auction will be identical to auction l , and thus $u_{\bar{l}}(v, k) = u_l(v, k)$ for any v and k . Prior to the join, if an advertiser wins position $k \leq n_{l'}^*$ (and thus the advertiser receives a placement in both auctions), then its total pre-merger utility of $u_l(v, k) + u_{l'}(v, k)$ is greater than its post-merger utility $u_{\bar{l}}(v, k)$. \square

5.4 Joint Auctions with Endogenous Capacity and Technology

We now return to the foundations of consumer choice to more closely model the change in capacities when joining auctions. Suppose there are three engines $l \in \{A, B, C\}$ positioned on the unit circle with uniformly distributed consumers. Individually, each engine solves the following optimization:

$$\max_{s_l} C_l(\hat{s})(1 - \beta_l)\bar{V}(a_l s_l + b_l) \quad \forall l \quad (27)$$

where \hat{s} represents a vector of three search costs. Denote the solution triple as $(\tilde{s}_A, \tilde{s}_B, \tilde{s}_C)$.

Without loss of generality, let engines A and B join; the new platform's technology will be the pointwise maximum of A and B , and the new platform will retain the sum of the prior capacities. Moreover, consider a case in which one engine, A , has a technology that pointwise dominates that of engine B ($a_A > a_B$ and $b_A > b_B$) and $C_A < C_B$. The optimization program under a joined platform will be:

$$\max_{s_A, s_B} (C_A(\hat{s}) + C_B(\hat{s})) (1 - \beta_A) \bar{V}(a_A s_A + b_A). \quad (28)$$

When A and B join, the joined platform replaces the technology of B with A . Then the marginal benefit of increasing search cost on engine B is the increase in C_A (increase in technology no longer applies since technology B has been replaced) while the marginal cost is a reduction in C_B . By noting that $C - C_C(\hat{s}) = C_A(\hat{s}) + C_B(\hat{s})$ and $\frac{\partial C_l}{\partial s_{l'}} < 0 \quad \forall l \neq l'$, we can conclude that total capacity will only decrease in response to increasing s_B . Then the solution to the joint maximization becomes $s_B^* = 0 \leq \tilde{s}_B$.

On the other hand, the marginal cost of increasing s_A is exactly offset by the increase in C_B while the marginal benefit to technology remains. Then the joined platform will increase s_A such that $s_A^* \geq \tilde{s}_A$.

What will be the effect on the resulting individual and aggregate capacity and technology? In order to predict engine C 's best response to the join between A and B , we examine the first order condition of (27) for C :

$$C_C(s_c, s_{-c}) = -\frac{1}{a_C} \frac{\partial C_C}{\partial s_c} (a_C s_C + b_C). \quad (29)$$

Note that since consumers' utilities are linear in the cost of search, the partial derivative on the RHS is a constant. Point x_1 in Figure 2 illustrates the equilibrium market share of C prior to join. After the join, A and B adjust their search costs from \tilde{s}_{-c} to s_{-c}^* , effectively reducing C 's market share. Thus, C would find it undesirable to retain the former level of \tilde{s}_C would be too high (point x_2), and C best-responds by reducing its search cost to s_C^* . At location x_3 , C 's market share has decreased while the combined market share of A and B has increased. Therefore

$$C_A(s^*) + C_B(s^*) \geq C_A(\tilde{s}) + C_B(\tilde{s}) \geq \max\{C_A(\tilde{s}), C_B(\tilde{s})\}.$$

Thus, both A and B enjoy increased capacity. Moreover, A 's superior technology replaces B 's technology, so more advertisers will enter B . Similarly, by retaining the same level of technology while increasing capacity, A attracts more advertisers. Because A and B each attract more advertisers, the expected payment of advertisers on A and B will increase.

6 Conclusion

Increasing search engine concentration makes it important to understand how advertisement auctions compete. If each buyer were restricted to entering exactly one auction, then auction platforms could compete on reserve price or on the cost of entry. But in fact advertisers can multi-home, albeit with additional costs for each ad platform they join. With multi-homing possible, an advertiser enters an ad auction if the expected benefit to entry outweighs the associated cost, independent of outcomes in other auctions. Hence traditional considerations of competition on reserve price or entry cost do not apply. We therefore turn to the supply side of ad auctions: search engines compete on capacity (each user picks only one search engine) and compete by setting search cost to balance advertisement clicks (auction real estate) against user utility.

Using the model's first stage equilibrium entry strategies, we find empirical evidence for the result that multi-homing advertisers are associated with greater traffic and larger reach. In the second stage, comparative statics around consumer choice parameters offer a stylized view of search engine market conditions: Google has greater market share and (some suggest) superior technology, while Yahoo retains higher market share than Microsoft despite (some say) inferior technology.

Our approach informs understanding of the advertiser welfare implications of joining auction platforms. Joining ad platforms can attract substantial regulatory attention: In November 2008, the Department of Justice planned to file antitrust charges to stop the proposed Google-Yahoo transaction. Then, in 2009, the Department of Justice approved the proposed partnership between Microsoft and Yahoo. At first glance it might seem paradoxical for Microsoft or the Department of Justice to claim that the Google-Yahoo transaction is undesirable, for advertisers and for the economy as a whole, while the Microsoft-Yahoo transaction offers net benefits. But our analysis suggests that that conclusion is entirely possible. In particular, by creating a joined ad platform of larger size than Microsoft or Yahoo alone, the transaction lets advertisers spread participation costs over a larger purchase—making it worth the while of small to midsize advertisers to sign up with the joined Microsoft-Yahoo platform.

We view our contribution as threefold. First, whereas standard models of online advertising take advertisers' participation as exogenous, we explicitly model an advertiser's decision to use or ignore a given ad platform, and we provide empirical support for our model. Second, we offer a model of search cost to demonstrate competition across platforms, and we provide empirical support for our model. Third, we analyze the prospect of joining auctions and characterize when such joins do and do not increase welfare.

References

- I. Ashlagi, D. Monderer, and M. Tennenholtz. Simultaneous ad auctions. *Mathematics of Operations Research*, 2010. Forthcoming.
- S. Brin and L. Page. The anatomy of a large-scale hypertextual web search engine. 1998.
- M. Cary, A. Das, B. Edelman, I. Giotis, K. Heimerl, A. Karlin, C. Mathieu, and M. Schwarz. Greedy bidding strategies for keyword auctions. In *Proc. 9th ACM Conference on Electronic Commerce (EC)*, 2007.
- E. Clarke. Multipart pricing of public goods. *Public Choice*, 18:19–33, 1971.
- B. Edelman. Adverse selection in online ‘trust’ certifications and search results. *Electronic Commerce Research and Applications*, 2010.
- B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.
- G. Ellison, D Fudenberg, and M. Mobious. Competing auctions. *Journal of the European Economic Association*, 2(1):30–66, 2004.
- T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- J. J. Liu and D. M. Chiu. Mathematical Modeling of Competition in Sponsored Search Market. Working paper, 2010.
- R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 21:58–73, 1981.
- M. Pai. Competing auctioneers. Working paper, 2009.
- H. Varian. Position auctions. *International Journal of Industrial Organization*, 25:1163–1178, 2007.
- W. Vickrey. Counterspeculations, auctions, and competitive sealed tenders. *Journal of Finance*, 16:15–27, 1961.

	Normalized Impression Count	Normalized Advertiser Count
Google only	17.34	1000.00
Yahoo only	26.01	338.65
Microsoft only	80.92	84.21
Google and Yahoo	663.58	209.26
Google and Microsoft	78.03	24.61
Yahoo and Microsoft	*	*
Google, Yahoo, and Microsoft	1000.00	65.76

Table 1: Advertiser Size and Multi-homing Status (Source 1)

* - Our data source did not identify any advertisers using Yahoo and Microsoft but not Google.

	Proportion Ranked	Average Rank	Average Reach	Average Page-Views
Google only	0.696	17,428,974	4.19	0.22
Yahoo only	0.735	14,784,742	5.20	0.40
Microsoft only	0.705	15,838,598	4.40	0.41
Goole and Yahoo	0.862	8,305,741	17.11	1.18
Google and Microsoft	0.888	7,234,154	4.70	0.26
Yahoo and Microsoft	0.871	8,325,335	3.07	0.15
Google, Yahoo, and Microsoft	0.940	3,803,684	62.94	5.64

Table 2: Advertiser Size and Multi-homing Status (Source 2)

	Presence Rank
Microsoft only	1
Yahoo only	2
Yahoo and Microsoft	3
Google Only	4
Google and Microsoft	5
Google and Yahoo	6
Google, Yahoo, and Microsoft	7

Table 3: Advertiser Presence Rank

	Reach	Rank	Page Views
Presence Rank	0.0832 *** (0.0208)	-1.01e-08 *** (4.29e-10)	0.8360 *** (0.2551)
Pseudo R^2	20.64	0.0079	0.0002
Model p -value	<0.001	<0.001	0.001
Observations	26286	26286	26286

Table 4: Ordered Probit Regressions of Advertiser Presence Rank on Reach, Rank and Page Views

* - $P < 0.10$

** - $P < 0.05$

*** - $P < 0.01$




Figure 1: Advertiser Reach and Multi-homing Status (Source 2)




Figure 2: Joint Auctions with Endogenous Capacity and Technology