

Pricing of Advertisements on the Internet

Subramanian Balachander

Purdue University

Karthik Kannan

Purdue University

2007-001

Pricing of Advertisements on the Internet

Subramanian Balachander*

Karthik Kannan†

June 2006

Krannert School of Management

Purdue University, West Lafayette, IN 47907

*sbalacha@mgmt.purdue.edu

†kkarthik@mgmt.purdue.edu

Pricing of Advertisements on the Internet

Abstract

Online advertisements are increasingly becoming an attractive channel for advertising. Pricing for online ads has evolved over time, and different auction-based policies are employed by website publishers to sell online ads. In this paper, we analyze, under different scenarios, the following three auction-based policies commonly used by the publishers: (i) invite bids for an impression of an ad, display the highest bidder's ad, and charge the winning advertiser the submitted bid (ii) invite bids for a click of an ad, display the ad of the bidder submitting the highest bid for a click, and charge the winner the bid submitted, and (iii) invite bids for a click of an ad, use the bid and the click-through-rate of the ad to compute the expected publisher profit from the ad, display the ad generating the highest profit, and charge the winner the submitted bid. A distinguishing feature of our analysis is that it considers the effect of product-market positioning of sellers on their bids for online ads and the consequent implications for pricing policies. Our analysis provides a number of insights useful to the website publisher and the advertisers. First, interestingly, we find that the bid per impression policy and the bid per click policy that uses the expected profit to determine the winner generate identical outcomes. Second, the advertiser profits are found to be higher when the publisher is capable of targeting the consumers. From the first two results, it appears that advertisers' angst at pricing based on impressions may be misplaced, and that a more critical determinant of the advertiser's profitability is the publisher's ability to target. Third, we find the surprising result that among the two bid per click policies, the one that uses the expected publisher profit can generate lower publisher profits. Fourth, we show that, similar to that in the literature on maximizing channel profits in the traditional markets, non-linear pricing of ads leads the publisher to gain the maximum profits even in this vastly different advertising distribution channel. Fifth, we demonstrate that offline ads and online ads are synergistic.

Keywords: Advertising and Media, Pricing, Internet Marketing, Game Theory, E-commerce

1 Introduction

The Internet is increasingly becoming an attractive channel for advertising. Spending on online ads is expected to grow 19% in 2006 while only a 2% growth is expected for offline advertisements (Weisman, 2006). Further, the market for online ads in the US, pegged at \$12 billion in 2004, is expected to grow to \$26 billion by 2010 (Delaney, 2005). The beneficiaries of this growth in online advertising revenue are web publishers/portals such as Yahoo, NYTimes.com on the one hand, and search engines such as Google and Yahoo's Overture on the other. While internet users visit web publishers such as NYTimes.com for their information content, they visit websites of search engines to search for the internet addresses of publishers such as NYTimes.com or for the addresses of online sellers of specific products. Traditionally, publishers have sold display advertisements (hereafter, simply ads) such as banner ads to advertisers interested in reaching visitors to their websites. In this case, advertisers primarily relied on the number of impressions or "eyeballs" delivered by the publishers and the asking price for the ad space, in choosing where to advertise.

More recently, search engines such as Google have given a new impetus to online advertising by selling search-related ads that are placed alongside search results delivered to users. Advertisers in this case have much more than an average demographic characteristic of the search engine user to go by in gauging the usefulness of advertising. The search keyword such as "wall clock" or "Howard Miller wall clock" that is entered by the search engine user reveals a specific interest in the product or even the brand, and this makes the user a valuable prospect to sellers of the product. This advertising innovation has made search-related advertising, which amounted to 40% of all online advertising in 2004, the fastest growing component of online advertising (Baker, 2004a, Delaney, 2005). The success of search-related advertising has enabled Google and Yahoo (Overture) to extend this idea to placing context-related ads at the website of publishers such as NYTimes.com (Delaney and Guth, 2005). For example, an internet user reading an article on skiing at a newspaper website

may be shown an ad from a vendor of ski equipment by Google or Yahoo (Overture). In this way, these search engine firms have also become intermediaries between advertisers and web publishers.

Just as online advertising has evolved from banner ads to search- or context-related ads, pricing of online ads to advertisers has evolved over time. In the early days of the Worldwide Web, the pricing model for online ads imitated that used for traditional off-line media such as print and TV, with the publisher charging a flat fee for displaying a fixed number of impressions of an ad. Since this flat fee was usually computed based on thousand impressions, this pricing model came to be known as CPM (cost-per-thousand). Subsequently, CDNow introduced the idea of paying web publishers based on click-throughs generated from ads or links placed on the publishers' website (Hoffman and Novak, 2000). Since the advertiser is charged only when a consumer clicks on the displayed ad, this pricing model came to be known as Cost-per-click (CPC). In basing payment on clicks, the CPC pricing model capitalized on the ease of measuring consumer response in the internet medium.¹ Many advertisers preferred CPC over CPM because of the former's closer link to an ad's performance (Hu, 2004, Weisman, 2006). In 2003, another innovation was introduced by GoTo.com (which eventually became Overture, and now part of Yahoo) when it auctioned the ad space to the highest bidding advertiser. This auction model is now commonplace in the search-related advertising category, where specific search keywords such as "wall clock" are auctioned, and advertisers use automated bidding software to submit bids per click for millions of search keywords. Google recently made an interesting announcement that it would also accept bids per impressions for its context-related ad services. In sum, pricing models for online advertisements differ on whether payment is based on impressions or clicks, and on whether the price per click or impression is posted or decided through an auction.

In this paper, we focus our analysis on the auction model for pricing online ads. In this context,

¹A pricing model that is even more closely tied to performance is one where the publisher is paid only when the click results in a purchase. This model is also used in internet advertising.

we analyze the implications of alternate pricing policies for the seller of online ad spaces (also referred to as the publisher), and the advertiser. A distinguishing feature of our analysis is that it considers the effect of product-market positioning of online advertisers on their bids for online ads and the consequent implications for pricing policies. Our analysis considers the following research questions: (i) What are the implications of pricing based on impressions versus clicks? (ii) How is pricing for search- or context-related advertising different from other less precisely targeted advertising methods? (iii) What are the implications for volume discounts given for higher click-through rates (CTR) of an ad? (iv) What are the implications of alternate policies for choosing the winning bidder for advertisements? With respect to the last question, two policies for choosing the winning bidder are observed in the online advertising industry. In one policy, used by Yahoo's Overture for search-related advertising, the advertiser who bids the highest CPC is considered the winning bidder. We refer to this policy as HCPC in the paper. On the other hand, Google's policy for search-related advertising derives a rank score by multiplying the maximum CPC bid by an advertiser for a search keyword with a quality score that increases with the advertiser's CTR, and the winning advertiser is the one with the highest rank score. Yahoo Overture's policy of choosing the winning bidder solely on the basis of CPC may be puzzling in the context of payment by clicks, since the winning ad may produce lower revenues due to a lower CTR.² We investigate this issue in this paper. In our analysis, we use an abstraction of Google's policy in which the winning bidder is chosen based on the product of the CPC bid and the expected CTR, and we refer to this policy as HCXR. Finally, we consider the impact of offline ads by the advertiser on the profitability of online advertising to the advertiser and publisher. This is an important issue facing marketing managers as they decide how to coordinate offline and online advertising in their marketing mix. This issue is

²For example, consider two advertisers with one paying \$4 as the CPC and having a click-through-rate (CTR) of 0.5% and the second paying \$3 as the CPC for ads generating a CTR of 0.75%. It is clear that displaying the first ad, although it has the highest CPC, will not generate the highest expected profits. This illustrates the importance of considering click-through-rates (CTRs), which Google takes into account.

also of interest to managers of online and offline media concerned about how competition between the two media will affect advertisers' spending patterns.

The model that we consider has a publisher or a search engine auctioning space for an online ad. We assume that there are two advertisers who are bidding to advertise in this space. The competing products that they sell are related to the search keyword or the context where the ad will be placed. Given our interest in analyzing the effect of the product-market positioning of the bidders, we assume that the products of these advertisers are positioned to appeal more to one of two different consumer segments. Our main conclusions are as follows. Bidding based on impressions or clicks yield surprisingly equivalent results in terms of advertiser and publisher profits. On the other hand, advertising placed with knowledge of a search keyword or context can be more profitable for the advertiser. Thus, advertisers should place more emphasis on the technologies employed by the web publisher to learn more about the online user than on whether the payment is based on clicks or impressions. We also find that volume discounts based on number of clicks or impressions can increase the publisher profit while also maximizing the profit of the channel consisting of the publisher and advertiser. It is interesting that volume discounts, which have been found useful in other contexts (Jeuland and Shugan, 1983), have a useful role in the case of online advertising as well. Furthermore, we find that for search- or context-related ads, the HCPC policy can be better than the HCXR policy in terms of the profit to the publisher. This is surprising since the HCXR policy takes into consideration the expected profits of ads before choosing the winner, while HCPC does not. Lastly, offline ads can increase advertisers' profit from online advertising suggesting that a synergy exists between these forms of advertising.

The rest of the paper is organized as follows. Section 2 reviews the nascent literature on the pricing of online ads and positions our contribution. Section 3 presents the basic model and is followed by Section 4 which analyzes the basic model. In Section 5, we relax some of the assumptions

in the basic model. Section 6 extends the basic model to analyze the impact of offline advertising on online advertising, and section 7 concludes the paper with a discussion of the results.

2 Literature Review

The literature on the pricing of online ads is limited in view of its fairly recent origins. A related paper that looks at the auctioning of online ad space is Edelman, Ostrovsky, and Schwarz (2005). Their focus is on understanding the generic properties of the auction models employed by Google and Yahoo Overture in which advertisers submit a single bid and are assigned, based on their ranks, to advertising slots offering differing potential to generate click-throughs. Since the winning bidder of any advertising slot is charged the next highest bid, they term this auction a generalized second-price auction (GSP). They show that the GSP does not lead to a truth-telling equilibrium unlike a standard second-price auction (Vickrey, 1961). They then characterize the bidders' payments in equilibrium. Unlike Edelman et al. (2005), we consider the implication of product-market positioning of advertisers (bidders) on their profit and that of the publisher (auctioneer). This added structure placed on the bidding by the product-market positioning of bidders allows us to compare the efficacy of HCPC and HCXR as alternate policies to choose the winning bidder, and to compare bids based on clicks with those based on impressions. In contrast, these issues are not studied in Edelman et al. (2005) since they assume that the CTRs for the advertisers are the same. However, in order to focus on issues of interest to us, we abstract away from GSP-related issues by assuming advertisers compete for a single advertising slot.

Another paper that considers advertising auction models used by Google and Yahoo Overture is Feng, Bhargava, and Pennock (2005). They use simulations to compare search engine revenues from alternate algorithms, including HCPC and HCXR, for choosing the winning bidders for advertising slots. They conclude that HCXR weakly dominates HCPC in terms of revenue to the search engine.

In contrast, we find that search engine revenues under HCPC can equal or exceed those under HCXR for search-related advertising. Other work in this literature considers use of information about an advertiser's product performance in ranking bidders (Weber and Zheng, 2005).

In another work on the pricing of online ads, Hu (2004) argues that payment based on clicks rather than impressions is optimal with risk-averse publishers and advertisers since such a pricing policy mitigates moral hazard concerns.³ These concerns arise because the publisher exerts unobserved effort to improve targeting of the ad to consumers. In Hu (2004), an advertiser makes a price offer to the publisher who accepts or rejects it while we consider an auction model and assume risk-neutral players.

While we consider pricing of online ads, Dukes and Gal-Or (2003), consider the effect of price competition in the product market on negotiated rates for offline advertising. Apart from the difference in pricing mechanism, online ads differ from offline ads in that the former offers the consumer an opportunity to make an immediate purchase. However, as we note in our discussion of future research, there is scope for a single modeling framework that can capture the idiosyncratic aspects of online and offline advertising.

There is also an extensive literature on what can be considered as online price ads in the websites of intermediaries such as price comparison sites or shopbots (Chen, Iyer, and Padmanabhan, 2002, Iyer and Pazgal, 2003, Baye and Morgan, 2005). Unlike these papers, our analysis considers ads that inform or remind about existence of a seller than about price. Apart from this key difference, there are other significant institutional differences between these two types of advertisements that warrant a separate analysis for each.

³Hu (2004) also shows that it is optimal to make payment a function of observed purchase rate at the advertiser's website.

3 Model

Let P be a publisher or a search engine on whose website two advertisers, 1 and 2, selling in a product market, are interested in advertising. Let there be N consumers, where N is large, who visit the publisher's website and who are also in this product market. These consumers are of two types, 1 and 2, with an equal number of each. Advertiser i , $i \in \{1, 2\}$, has a value of v_h for a consumer of type i and a value of v_l for a consumer of type j , $j \in \{1, 2 : j \neq i\}$. The value of a consumer to an advertiser may depend on the expected profit from purchase by the consumer. Without the loss of generality, let $v_l < v_h$, which implies that an advertiser attaches a higher value to a consumer of his own type. The assumption implies that the advertiser's products appeal more to consumers of his own type. The difference between the valuations of the two types of consumers to the advertiser is a measure of the differentiation of the advertisers in the product market. A smaller difference indicates lesser differentiation.

The two types of consumers may also be expected to differ in their click-through rates (CTR) for ads from the advertisers. In order to understand how CTRs may differ across consumers and ads, it is convenient to think of two polar cases, one in which the consumer has no prior knowledge of either advertiser's type, and a second case where the consumer has complete prior knowledge of both advertisers' types.

First, consider the case where the consumer has no prior knowledge of the advertiser's type. In this case, the consumer's CTR should depend on the ad from the advertiser. An advertiser may design a tailored ad that appeals to a particular type (say j) by claiming that the advertiser is of that type (j). The consumer type to which the advertiser tailors his ads could even be his competitor's type. Thus, we do not preclude the advertisers from making false appeals to a consumer type for which his competitor is better positioned. If an ad is tailored to a consumer type (say j), we assume that the ad obtains a CTR of c_h from that type (type j) and a CTR of c_l from

the other consumer type (type i). We assume $c_h > c_l$, since consumers may reasonably be expected to click ads tailored to another type of consumer at a lower rate. Note that our assumptions imply that if advertiser i tailors ads to consumers of type j , he obtains a CTR of c_h . However, the value from attracting each of those consumers remains v_l .

Instead of tailored ads, an advertiser may choose to appeal to both consumer types with a blanket ad. As an example of a blanket ad, an advertiser whose selection consists primarily of Seth Thomas wall clocks and a few Howard Miller may claim to have “a good selection of wall clocks, including Seth Thomas and Howard Miller brands.” In this way, the advertiser attempts to appeal to the type of consumer interested in a Seth Thomas brand as well as one interested in a Howard Miller brand. In the case of a blanket ad, we assume that the consumer of both types have a CTR of c_h since the advertiser claims to cater to both types of consumers (alternate assumptions on this point are discussed later in footnote 7).

For the initial analysis, we assume that consumers have no prior knowledge of either advertiser’s products. Subsequently, we discuss the impact of relaxing this assumption by considering the second polar case where the consumer has complete prior knowledge of both the advertisers’ products.⁴ Throughout the paper, we assume the valuations and the CTRs to be common knowledge among the advertisers and the publisher.

Let the website publisher, P , auction an ad slot to the advertisers.⁵ The ad in the slot can be dynamically changed for each consumer visiting the website. Note that with two potential advertisers, it is sufficient to consider one slot. Based on the result shown in Edelman et al. (2005),

⁴Even with a prior knowledge about an online seller, 41% of web users used a search engine in 2003 to locate the seller on the Internet instead of typing the seller’s address into their browser (Oser, 2006).

⁵With common knowledge assumption, one can argue that the publisher can post prices rather than auction the ad slot. Although the valuations and the CTRs may be learned over time as advertisers repeatedly compete (thereby justifying the common knowledge assumption), the auction model cannot be abandoned after learning the valuations since valuations are subject to change. Thus, the use of an auction model in conjunction with common knowledge assumption may be justifiable and is also consistent with other papers in the advertisement-auction literature as well as the general auction literature (e.g., Edelman et al., 2005, Benoit and Krishna, 2001).

with two slots and two advertisers, the second slot generates zero profits and also imposes a negative externality on the publisher's profit from the first slot. Thus, it is not optimal for the publisher to offer two or more slots in this case.

The sequence of moves in our model is as follows. First, P announces the policy for selling the ad slot. It involves specifying how the advertisers will be charged for ads displayed (e.g., bid per impression or bid per click) and the manner in which the winner is chosen (e.g., bidder with the highest bid per impression, HCPC or HCXR). If the consumer type is known, the publisher accepts bids for each type. The advertisers then simultaneously submit their bids to the publisher. A bid submitted by an advertiser involves a bid price and a specific ad that the advertiser will be using. As mentioned earlier, the advertiser has two options in designing the ad: design a tailored ad to a consumer type or design a blanket ad to attract all consumer types. We assume that advertisers participate in the bidding only if the expected profit is non-negative. On receiving the bids, the publisher picks the winner in a manner consistent with the mechanism it specified in the first step and airs the ad accordingly to each of the N consumers. If there is a tie, we assume that the publisher randomly assigns $N/2$ consumers to each advertiser.

We assume that when accepting bids per impression, the advertiser with the highest bid is picked as the winner and pays the submitted bid. We refer to this policy as HCPM (highest cost per impression). On the other hand, while accepting bids per click, the publisher has the option of choosing one of two policies. Under the first policy, which we call HCPC (highest CPC), firm P airs the ad with the highest bid per click or CPC, and charges the advertiser the bid submitted. The second bid per click policy, which we call as HCXR, involves airing of the ad with the highest expected profit for the publisher and the winner is charged the bid submitted. The expected profit is computed as the product of the submitted bid and the CTR of the bidder's ad. Note that HCPC is similar to the pricing policies used by Yahoo while Google's policies are similar to HCXR. Note

also that each of three policies is similar to a first-price-sealed-bid auction mechanism (Wolfstetter, 1996). A second-price-sealed-bid equivalent of each of the policies, where the winner’s ad is displayed but the winner is only charged the second best bidder’s price, is feasible. With one ad slot available, both the first-price and the second-price mechanisms can be shown to be revenue-equivalent (the equivalence result will be similar to that demonstrated in the auction literature. See, Wolfstetter, 1996). For the ease of exposition, we restrict our analysis to the first price equivalent.

In our analysis, we first examine the policy where firm P invites bids on a per impression basis, resulting in a price that is linear in the number of impressions. Following that, we consider the policies where firm P seeks bids per click. Under each policy, we first consider the case where the publisher cannot identify the consumer’s type. In a search engine context, this may be because the consumer uses a general keyword rather than a more specific one. For example, a consumer most interested in a Howard Miller brand may search using the less specific term of “wall clock.” Alternately, the firm may not have the technology to know the consumer’s type. Subsequently, we consider the case where a consumer’s type is known with certainty, perhaps due to a consumer perfectly revealing his type by use of a specific keyword, or due to consumer tracking technology used by firm P . Toward the end of the analysis section, we also consider prices that are non-linear in the number of impressions or clicks.

4 Analysis

In this section, we assume that consumers have no prior knowledge about the advertisers. As mentioned in the previous section, we initially consider linear prices. All proofs are provided in a technical appendix.

4.1 Bid per impression, HCPM Policy

In this sub-section, we first consider the case where the consumer's type is not known. In this case, the result is given by the following proposition.

Proposition 1 *When the consumer's type is not known, and firm P invites bids per impression under the HCPM policy, there exists a unique equilibrium, which is symmetric. The optimal strategy for the advertisers is to use blanket ads and bid $\frac{v_h+v_l}{2}c_h$ per impression, firm P to randomize the airing of ads, making a profit of $N\frac{v_h+v_l}{2}c_h$ from consumers visiting the website. Each advertiser's profit is zero.*

Thus, when the consumer's type is not known, the advertisers can only bid up to the expected value per consumer. Since this expected value is *symmetric* for both advertisers, the publisher is able to extract all of the advertisers' profit. Also, advertisers rely on blanket ads since the consumer type is not known. Now consider the case where the consumer's type is known before airing of the ad. For example, the consumer's type may be revealed by the keyword submitted to a search engine by the consumer. In this case, the publisher can choose to accept bids based on the consumer's type. When such is the case, we have the following proposition.⁶

Proposition 2 *When firm P invites bids per impression based on the consumer's type, under the HCPM policy, there exists a unique equilibrium, which is symmetric. Each advertiser uses tailored ads and bids $v_l c_h + \varepsilon$ per impression for his consumers, and $v_l c_h$ per impression for his competitor's consumer, where ε is arbitrarily small. Firm P airs the ad of the highest bidder and makes a profit of $Nv_l c_h$ from consumers visiting the website. Each advertiser makes a profit of $\frac{N}{2}(v_h - v_l)c_h$.*

⁶In this case, there is another equilibrium which is trivial. The advertiser submits the same bid prices, but instead uses blanket ads. We ignore this trivial alternate equilibrium when we state that there is a unique equilibrium. Similar trivial equilibria exist for other Propositions and we ignore them as well.

Unlike in the previous proposition, when firm P invites bids based on the consumer's type, the advertisers' value for the consumer becomes asymmetric. This allows the advertisers to retain some of the surplus. The advertisers' higher profit in comparison to the case where the consumer type is not known is interesting and explains the general preference by advertisers for publishers such as Google who offer a more targeted advertising service. On the other hand, as the proposition shows, the publishers are worse off using this approach. However, providing the opportunity to target consumers through ads may have helped publishers such as Google attract advertisers away from traditional media.

A beneficial effect of knowledge of the consumer's type is that the combined profit of the advertising channel, i.e. the advertisers and the publisher, is maximized. Further, consumers receive truthful ads tailored to their type in equilibrium which leads to the best possible match between firm and consumer types from the channel profit point of view. Although we do not model consumer welfare, this outcome is likely to be favorable to consumer welfare as well due to the better match between consumer and firm.

The lower profits to firm P when inviting bids based on the consumer type raises incentive compatibility issues on the part of the publisher. In other words, firm P has an incentive not to invest in technology that will enable better targeting of consumers since better targeting reduces the publisher profits. Anecdotal evidence suggests that advertisers prefer paying for ads based on clicks rather than impressions or "eyeballs" since the former method is performance-based (See, for example, Hu, 2004, Weisman, 2006). It is possible that the incentive compatibility issue is one of the concerns underlying advertisers' preference for payment by clicks. It is therefore interesting to see if payment by clicks addresses this issue. We analyze such a pricing policy in the next section.

4.2 Bid per click

In this sub-section, we continue to assume that the publisher uses a linear pricing policy but that bids are invited per click rather than per impression. We first present the results when the consumer's type is not known.

Proposition 3 *When the consumer's type is not known and firm P invites bids per click,*

1. *under the HCPC policy, there exists a unique equilibrium, which is symmetric. The equilibrium is for each advertiser to use tailored ads appealing to his own consumer type and to bid $\frac{v_h c_h + v_l c_l}{c_h + c_l}$ per click. Firm P randomizes the airing of ads and makes a total profit of $N \frac{v_h c_h + v_l c_l}{2}$ from consumers visiting the website. Each advertiser's profit is zero.*
2. *under the HCXR policy, there exists a unique equilibrium, which is symmetric. The optimal strategy for the advertisers is to use blanket ads and bid $\frac{v_h + v_l}{2}$ per click, firm P to randomize the airing of ads, making a profit of $N \frac{v_h + v_l}{2} c_h$ from consumers visiting the website. Each advertiser's profit is zero.*

The proposition shows that the HCXR policy yields higher profit to firm P than the HCPC policy, which may be seen as a natural result of choosing the winning ad based on the higher expected profit per click. More interesting is the result that despite lacking information regarding consumer types, the equilibrium under HCPC always involves displaying to both consumer types ads that represent the advertiser's true type. On the other hand, the HCXR policy encourages blanket ads.⁷ Thus, the HCPC policy results in more truth in advertising. Consequently, the expected

⁷Our current analysis assumes no penalty in CTR for blanket ads by assuming that the CTR for a blanket ad is c_h . One can argue that consumers would be skeptical of blanket ads and will therefore click through them at a lower rate. Therefore, if we assume that the CTR from a blanket ad to be c_b such that $c_l \leq c_b < c_h$, only the equilibrium bids when the consumer types are not known are affected. More specifically, when c_b is sufficiently low (consumers become very skeptical of blanket ads), advertisers may use only tailored ads even under a HCXR policy. An alternate way of specifying c_b is through imposition of greater structure on consumers' prior beliefs about advertisers' types and to then derive the posterior beliefs. However, such a structural approach is likely to yield similar results while adding significantly to the complexity on the analysis.

profit (or bid) *per click* under HCPC is higher than it is under HCXR i.e., $\frac{v_h+v_l}{2} < \frac{v_h c_h+v_l c_l}{c_h+c_l}$.⁸

Comparison with the results in Proposition 1 suggests that when the consumer’s type is not known, the advertisers do not benefit from bidding based on clicks rather than impressions. The intuition is that the value of a click is just as symmetric across advertisers as it is for impressions when consumer types are not known. In such a case, advertiser profits are competed away.⁹ Finally, Proposition 1 establishes that the HCPM policy need not yield any less profit to firm P than bids per click as long as the publisher uses the HCXR policy with bids per click.

In the following proposition, we consider the situation where the consumer’s type is known.

Proposition 4 *When the consumer’s type is known and firm P invites bids per click, under either the HCPC or the HCXR policy, there exists an identical unique equilibrium, which is symmetric. In this equilibrium, each advertiser uses tailored ads appealing to his own consumer type, bids $v_l + \varepsilon$ per click for consumers of his type and bids v_l per click for consumers of his opponent’s type. Firm P chooses ads from advertiser i to serve consumers of type i which generates a profit of $Nv_l c_h$ from consumers visiting the website. Each advertiser’s profit is $\frac{Nc_h}{2}(v_h - v_l)$.*

Thus, when the consumer’s type is known, both the HCPC and HCXR policies surprisingly yield identical profits to the publisher. The rationale for this counterintuitive result is as follows. First, an advertiser showing ads tailored to his own type outbids any other advertiser strategy under either policy, since such tailored ads result in both the highest CTRs and the highest consumer value. Second, the value of the winning bid for a consumer type is determined by the value of

⁸Under the HCPC policy, it is possible for an advertiser to bid above v_h using a strategy of airing an irrelevant ad whose CTR is zero. We ignore this case because search engines have controls (both editorial and threshold CTRs) to prevent this strategy (Marckini, 2004).

⁹Note, however, that if advertisers were risk-averse, they would prefer to bid on clicks rather than impressions, if we assume uncertainty in the CTR. As seen in Proposition 1, advertisers bid based on the anticipated CTR as given by c_h . If c_h is uncertain, advertisers could make a loss if the actual value of c_h is lower than the value on which bids are based. On the other hand, if the bids are based on clicks, advertisers’ bids are independent of CTR under the HCXR policy. Thus, advertisers’ payoffs do not depend on the actual levels of CTR. The intuition is that bids are made conditional on a click in the case of bids-per-click which make the bids independent of the CTR. In the case of a HCPC policy with bids per click, the bids depend on the ratio c_h/c_l but not on the overall levels of c_h and c_l . Thus, advertisers’ profits are unaffected if CTR levels are lower than expected as long as c_h/c_l remains the same.

the second-best strategy, which is a tailored ad from an advertiser of opposite type in both cases. Thus, the winning ads and bids are identical under either policy resulting in identical profits to the publisher. This result may explain the use of both HCPC (Yahoo) and HCXR (Google) policies in the market for search-based ads.

Also, by comparing the publisher's profit under the HCXR policy in Proposition 4 with that in Proposition 3, we see that they are lower when the consumer's type is known. The advertisers' profits are correspondingly higher. This finding is similar to that obtained for the HCPM policy and the intuition once gain the asymmetric valuations of consumers to the advertiser, when the consumer type is known, allows the advertisers to retain more of the profit. Thus, the incentive compatibility issue with respect to firm P choosing better targeting technologies persists even when payment is based on clicks. However, the comparison of the expected profits show that under the HCPC policy, this issue may be less relevant particularly when $\frac{v_h}{v_l} < 2 - \frac{c_l}{c_h}$. This condition is satisfied when v_h (in relation to v_l) or c_l (in relation to c_h) is sufficiently small, which implies that asymmetry in the valuation of the consumer types by the advertisers is small. Under this condition, a firm P adopting the HCPC policy does better when the consumer's type is known. Thus, the HCPC policy provides the right incentives to the publisher under some conditions to learn the consumer's type and reveal it to the advertisers. As in the case of the HCPM policy, the combined profit of the advertising channel is maximized and truthful ads prevail in equilibrium when the consumer type is known. Thus, a policy that provides the right incentives to the publisher is desirable from the point of view of maximizing channel profit and getting the best possible match between consumers and advertisers.

Another relevant comparison is the one between bids per click and bids per impression. Interestingly, the outcomes are largely similar in terms of publisher and advertiser profits in the two cases. Thus, it appears that advertisers' angst at pricing based on "eyeballs" (impressions) may be

misplaced. A more critical determinant of the advertiser’s profitability is the publisher’s ability to identify the consumer type. This may explain why search-based ads are so popular among advertisers, since this technology helps narrow down the consumer’s specific need. The similar results with bidding based on impressions or clicks may explain why Google now offers an opportunity to bid based on impressions or clicks. In the first two rows of Table 1 on page 21, we have summarized the results so far (when consumers have no prior knowledge).

4.3 Bid per click and Volume subsidies

In this section we examine the effect of non-linearity in the cost of clicks or impressions to the advertiser caused by volume subsidies offered by firm P . We establish the result only for bid per click policies. The result for the HCPM policy is similar. We consider the case where the consumer’s type is known. As discussed later, volume subsidies are ineffective when the consumer’s type is not known.

Proposition 5 *When the consumer’s type is known and firm P invites bids per click with volume subsidies the following symmetric equilibrium exists under the HCPC or the HCXR policy (ε is arbitrarily small): (i) firm P offers a per click subsidy of $v_h - v_l - \varepsilon$, when the number of clicks exceed $\frac{N}{2}c_h$. (ii) each advertiser uses tailored ads and bids v_h for his consumer type and $v_h - \varepsilon$ per click for his opponent’s consumer type. Firm P airs ads of the highest bidder and makes a profit of Nv_hc_h . Each advertiser makes zero profit.*

Proposition 5 is interesting because volume subsidies enable firm P to increase the publisher profit to the maximum possible profit when the consumer’s type is known. Paradoxically, advertisers are hurt by volume subsidies as their profit decreases. The intuition is that the presence of volume subsidies encourages each advertiser to increase his bids for his competitor’s consumer type. Thus, the subsidies reduce the asymmetry in advertisers’ valuation of the consumers allowing the publisher

to appropriate the advertisers' profit. In equilibrium, advertisers end up winning bids for ads to be aired to their own consumer type thus matching advertisers and consumers of the same type. Therefore, neither advertiser actually benefits from the subsidy offered by the publisher. Thus, volume subsidies maximize the publisher's profit while also maximizing the channel profit. The utility of volume subsidies (quantity discounts) in maximizing channel profit (and manufacturer profit) in a product distribution channel is an important result in the marketing literature (Jeuland and Shugan, 1983). It is interesting that a parallel result prevails in a vastly different advertising distribution channel.

Note that the publisher needs to design the volume subsidy carefully so as to not give away the profit that advertisers realize from advertising to their own type. The publisher accomplishes this in Proposition 5 by offering a per-click subsidy only when an advertiser's volume of clicks exceeds $N_{c_h}/2$. Further, note that the publisher's profit when offering a volume subsidy with bid per click exceeds his profit when the consumer's type is not known (compare with Proposition 3). Thus, volume subsidies remove the incentive compatibility issues which may prevent a publisher from implementing technologies that enable learning and revealing of the consumer's type to advertisers.

Note also that results on volume subsidies also extend to the HCPM policy when consumer types are known. The subsidies are now based on the number of impressions shown instead of the number of clicks. Volume subsidies encourage each advertiser to increase his bid for impressions to consumers of his competitor's type so that the same results go through in this case as well. As a consequence, the HCPM policy can be shown to yield a similar payoff to the publisher as the HCPC or the HCXR policy.

Volume subsidies cannot, however, increase the publisher's profit when the consumer's type is not known. The intuition is readily seen for this assertion since all of the advertisers' profits are already being extracted by the publisher in this case (see Proposition 3).

5 Extensions

In this section, we focus on relaxing some of the assumptions made. First, we discuss relaxing the assumption of symmetric advertisers. Following that, we demonstrate how the assumption regarding consumers' prior knowledge of the advertisers' types affects the results.

5.1 Asymmetric Advertisers

In this sub-section, we discuss how asymmetry in the number of consumers of each type may affect the results in the previous section. In particular, let f_i represent the fraction of consumers of type i , where $f_i > \frac{1}{2}$. As before, we assume that advertisers i and j appeal more to consumers of type i and j respectively. Thus, advertiser i has a broader appeal than advertiser j , who is a niche player. It can be shown that the results are unchanged for bids per click or bids per impression when the consumer's type is known, and when there are no volume subsidies.¹⁰ With volume subsidies, the publisher improves his profit but is unable to extract all of the advertiser's profit. When the consumer's type is not known, the results are significantly different in the case of asymmetric advertisers. In this case, the advertiser with broader appeal outbids the niche advertiser, who is therefore unable to air his ad. (This may change if advertisers can change bids over time and if the majority consumer types exit the market after purchase.) Thus, search-based ads (consumer type known) enable smaller advertisers to participate in the online advertising market.

5.2 Consumer's Prior Knowledge about the Advertisers

We now consider how the results for the symmetric model with linear prices described in the previous section change if we relax the assumption that consumers have no prior knowledge about an advertiser's type. To do so, we consider the other polar case in which consumers have complete

¹⁰Results for the case of asymmetric advertisers are available from the authors upon request.

prior knowledge of an advertiser's type. Thus, an online advertisement serves mainly as a reminder or to provide a link to the advertiser's website in this case. In such a scenario, an advertiser receives a lower CTR (c_l) from a consumer of opposite type, resulting in a lower value of advertising to this consumer for this advertiser. This, in turn, lowers the advertiser's winning bid for online advertisements leading, in general, to a lower publisher profit. However, the most interesting results for this case pertain to the comparison of HCPC and HCXR policies when bids are based on clicks.¹¹ The following proposition describes these results.

Proposition 6 *Assume firm P invites bids per click and consumers have complete prior knowledge of each advertiser's type:*

1. *When the consumer's type is not known, the equilibrium results for the HCXR policy are identical to those for the HCPC policy. The equilibrium results for the latter policy are identical to those given in Proposition 3 (Part 1).*
2. *When the consumer's type is known, the equilibrium results for the HCPC policy are identical to those given in Proposition 4. However, under the HCXR policy, there exists a unique equilibrium, which is symmetric, in which each advertiser uses tailored ads appealing to his own consumer type, bids $\frac{v_l c_l}{c_h} + \varepsilon$ per click for consumers of his type and bids v_l for consumers of his competitor's type. Firm P chooses ads from advertiser i to serve consumers of type i which generates a profit of $N v_l c_l$ from consumers visiting the website. Each advertiser obtains a profit of $\frac{N}{2}(v_h c_h - v_l c_l)$.*

The above proposition highlights two main conclusions. First, the equilibrium results for the HCPC policy are unaffected by consumer's prior knowledge about each advertiser's type. This

¹¹Under the HCPM policy, the results are as follows. When the consumer type is not known, the bids submitted by the advertisers are lower ($\frac{v_h c_h + v_l c_l}{2}$), leading to a lower publisher profit. Advertiser profits continue to be zero while the publisher profit is $N \frac{v_h c_h + v_l c_l}{2}$. When the consumer type is known, the advertiser's winning bid is lower ($v_l c_l + \varepsilon$) than when consumers did not have prior knowledge, resulting in lower publisher profit ($N v_l c_l$) and greater advertiser profit ($N \frac{v_h c_h - v_l c_l}{2}$).

result is to be expected since the HCPC policy ignores differences in CTR, which is affected by consumer's prior knowledge. Moreover, the HCPC policy always encourages truth-telling on the part of the advertiser, and since this is equivalent to the consumer having complete prior knowledge of each advertiser's type, the prior knowledge of consumers does not affect the equilibrium results under HCPC.

The second conclusion from the above proposition is that HCPC becomes more attractive to the publisher in comparison to HCXR when consumers have complete prior knowledge. In particular, when the consumer type is not known, HCPC yields the same profit to the publisher as HCXR when consumers know each advertiser's type *a priori*. On the other hand, when the consumer's type is known, as is the case for search- or context-related ads, the publisher's profit with HCPC strictly exceeds that under the HCXR policy. This last result may appear to be more surprising than the equivalence between these two policies for the publisher that we obtained for the case where the consumer has no prior knowledge of the advertiser's type (cf. Proposition 4). The intuition is that when consumers have complete prior knowledge of the advertiser's type, an advertiser i cannot simply pretend to be of type j to consumers of type j . Thus, advertiser i can expect the lower CTR of c_l from consumers of type j even if he tailors his ad to this consumer type. Under the HCXR policy, the competing advertiser j can take advantage of this lower CTR of advertiser i from consumer j by winning advertisements for this consumer with a lower bid ($v_l c_l / c_h + \varepsilon$ instead of a bid of $v_l + \varepsilon$ under the HCPC policy or when consumers have no prior knowledge). Thus, advertisers make a higher profit at the expense of the publisher under the HCXR policy when consumers have complete prior knowledge of each advertiser's type.

If we assume that consumers' knowledge in reality may lie between the two extremes of no knowledge and complete knowledge, we may expect a consumer j to click an ad from advertiser i at a CTR that lies between c_h and c_l , say c . In this case, it can be easily seen that in comparison

Consumers have Prior Knowledge	Consumer Type	Comparison of Publisher Profits	Nature of Ads	Advertiser Profits	
				HCPM or HCXR	HCPC
No	Unknown	HCPM=HCXR > HCPC	Tailored only under HCPC.	0	0
	Known	HCPM=HCXR = HCPC	Always tailored.	$\frac{N(v_h - v_l)c_h}{2}$	$\frac{N(v_h - v_l)c_h}{2}$
Yes	Unknown	HCPM=HCXR=HCPC	Always tailored	0	0
	Known	HCPM=HCXR < HCPC	Always tailored.	$\frac{N(v_h c_h - v_l c_l)}{2}$	$\frac{N(v_h - v_l)c_h}{2}$

Table 1: Summary of results with linear prices.

to the HCPC policy, publishers are still worse off and advertisers are still better off with the HCXR policy as the winning bid for ads continue to be lower at $v_l c / c_h + \varepsilon$. Just as the advertiser gains from consumer knowledge under the HCXR policy, an advertiser can also gain from offline advertisement that has a similar effect, as will be seen later in section 7. It is to be noted that our result that HCPC dominates HCXR is a consequence of the positive correlation between the valuations and the CTR, i.e., a higher valuation has a higher CTR. Although a positive correlation between the valuations and the CTR is typically encountered, one can construct scenarios where the valuations and the CTRs may be negatively correlated. In those cases, HCXR dominates HCPC in terms of publisher profits. Table 1 presents the results so far excluding those relating to volume subsidies.

6 Managing Online and Offline Advertising

The advent of online advertising has raised important issues for marketing managers concerning the relative role of online advertising and traditional offline advertising in the marketing mix. In this section, we study how offline advertising may affect the bidding process for online advertising. For this purpose, we extend the model in Section 3 as follows. Assume that advertisers can reach consumers in an off-line media such as television, newspaper or magazines at the cost of α per consumer. We further assume that a key characteristic that distinguishes online from offline advertising is that the latter is less effective in inducing purchase. This is because online ads offer

the consumer the opportunity to make an immediate purchase compared to offline ads where the consumer needs to exert significant additional effort such as going online, making a call, or visiting a store to make a purchase. In this context, we make the extreme assumption that an offline ad cannot induce any purchase on its own. This assumption simplifies the analysis and enables us to focus on how offline ads affect online ads. Thus, our analysis cannot address the issue of whether and when an advertiser can use only offline ads. Lastly, we make the following additional stylized assumptions about how consumers respond to receiving both online and offline ads.

1. All offline ads by one or both advertisers are aired before the consumer sees an online ad at the publisher's website. The advertisers have the option of targeting their offline ads. We assume for simplicity that, when offline ads are targeted, consumers of a given type can be perfectly targeted. For example, an advertiser may choose specific television programs to reach the targeted consumer type.
2. We further assume that all consumers can be reached through offline ads. A consumer may see an offline ad from none, one or both advertisers depending on advertisers' offline targeting strategies, but will see the online ad of only the winning advertiser at the publisher's website.
3. An advertiser who succeeds in reaching a consumer through both online and offline ads experiences an increase in the CTR of his online ad. The assumption that advertising affects only the CTR but not the consumer value (v_h or v_l) is consistent with the hierarchy-of-effects model in marketing. Under this model, advertising is considered to be more effective in creating interest in a product, i.e., CTR, while its influence on the consumer's ultimate purchase decision (v_h or v_l) may be less significant. Specifically, we assume that if a consumer receives two identical messages, one offline and another online, that the advertiser is of a particular type, then his CTR for that advertiser increases by δ over his base CTR level for

that type of advertiser.¹² On the other hand, if a consumer receives two contrary messages, one indicating that the advertiser is of type i , and the other indicating that he is of type j , the consumer is assumed to react as if he received a blanket ad online (with a CTR of c_h).

We consider the case where the consumer's type is known for online ads, consumers have no prior knowledge about the advertisers, and bids are invited on a per-click basis. We initially do not allow for volume subsidies. The following proposition describes the result on offline advertising.

Proposition 7 *When the consumer's type is known for online ads and firm P invites bids per click without volume subsidies,*

1. *under the HCPC policy, both firms target their own types with offline advertising if $(v_h - v_l)\delta \geq \alpha$, and do not use offline advertising otherwise. When offline advertising is used, then, in the online publisher's website, each advertiser uses tailored ads and bids $v_l + \varepsilon$ for his consumer type and v_l per click for his opponent's consumer type, where ε is arbitrarily small. Firm P displays the ad of advertiser i to the consumer of type i and makes a profit of $Nv_l c_h$. Each advertiser makes a profit of $N[(v_h - v_l)(c_h + \delta) - \alpha]/2$.*
2. *under the HCXR policy, both firms target their own types with offline advertising if $v_h \delta \geq \alpha$, and do not use offline advertising otherwise. When offline advertising is used, then, in the online publisher's website, each advertiser uses tailored ads and bids $v_l c_h / (c_h + \delta) + \varepsilon$ for his consumer type and $v_l c_h / (c_h + \delta)$ per click for his opponent's consumer type, where ε is arbitrarily small. Firm P displays the ad of advertiser i to the consumer of type i and makes a profit of $Nv_l c_h / (c_h + \delta)$. Each advertiser makes a profit of $N[v_h(c_h + \delta) - v_l c_h - \alpha]/2$.*

¹²This implies that if consumer i receives two messages that an advertiser is of type i (j), his CTR increases to $c_h + \delta$ ($c_l + \delta$). Although it might seem more appropriate to assume that there would not be any increase in the consumer i 's CTR for an advertiser who sends two messages that he is of the opposite type j , the results are not sensitive to this assumption. We therefore make the simpler assumption for ease of exposition.

The above proposition results in three main insights. First, an advertiser chooses to target any offline advertising to his own type under either the HCPC or HCXR policy. Although an advertiser may increase the CTR of consumers of the competitor's type by advertising to them off-line, he is unable to appropriate the benefits because the rewards to the competitor from offline advertising to this consumer is even greater. Thus, the advertiser does not target his competitor's type with offline advertising in equilibrium. Second, the HCXR policy is more encouraging of offline advertising by the advertisers than the HCPC policy. In particular, the HCXR policy favors the use of offline advertising at higher costs (α), and offers higher profits to the advertisers from the use of offline advertising. These favorable results may be seen as an outcome of the HCXR policy's attention to CTR, which is the variable affected by offline advertising. The intuition is that the higher CTR resulting from offline advertising translates to a lower bid ($v_l c_h / (c_h + \delta) + \varepsilon$) being needed from an advertiser to advertise online to his own consumer type. This lower bid, coupled with the higher CTR due to the offline advertising, enables an advertiser to realize greater profits from offline advertising under HCXR. A third insight is that the publisher's profit are surprisingly reduced with offline advertising under the HCXR policy, even though this policy rewards ads with higher CTR. The intuition is that, in equilibrium, an advertiser enjoys a higher CTR only from his own consumer type as a result of the targeted offline advertising. Thus advertisers' valuations of consumers become even more asymmetric with offline advertising under the HCXR policy, thereby increasing the advertisers' profit at the expense of the publisher's profit. In contrast, the publisher's profit under the HCPC policy are unaffected by offline advertising.

We now discuss how the results of the proposition may change with alternate assumptions. If volume subsidies were allowed, a publisher can appropriate the advertiser's additional profit due to offline advertising (cf. Proposition 5). Thus, offline advertising is used in equilibrium even with volume subsidies but the benefits net of advertising costs go to the publisher. If targeting of offline

ads were to be impossible or imperfect, the results will be substantially similar to a case where the effective advertising cost of reaching a consumer is higher. The rationale is that ads that are misdirected to consumers of opposite type do not change the identity of the winning bidder for online ad, but merely increases the amount of wasted ads. However, due to an advertiser's ads sometimes reaching a competitor's consumer type, advertisers' evaluations of consumers become more symmetric. Thus, an inability to target offline advertising can increase the publisher's profit at the expense of the advertiser. Finally, if the consumer type is not known for online ads, publishers would benefit at the expense of the advertisers as seen earlier.

7 Discussion and Conclusion

The growth of online advertising, and in particular, search-related advertising has made it imperative for marketing and advertising managers to understand how to manage and integrate online advertising decisions into their marketing communications decisions. New specialized advertising agencies such as aQuantive have sprung up to advise and help firms in the management of their online advertising decisions (Baker, 2004b). On the other side, web publishers who sell advertising spaces on the Internet find a need to understand the implications of different pricing models for online advertisements. In this paper, we consider the implications for the publisher and for advertisers, of alternate pricing policies in the auctioning of online ads. Further, we examine the difference in outcomes for search or context-related advertising (which we model as the consumer type being known) versus other less targeted online advertising. We also analyze the implications of an advertiser's offline advertising for his profit from online advertising. We discuss below our main conclusions and the implications for managers with respect to the research questions considered by our analysis.

Should pricing (bids) be based on clicks or on impressions? We find that pricing based on clicks

or impressions lead to substantially similar profit outcomes for the publisher and advertisers. In particular, while the results are identical for both pricing bases when HCXR is used for payment by clicks, there are minor differences with the HCPC policy.¹³ The intuition is that if the payment basis is impressions rather than clicks, advertisers simply incorporate the expected CTR in addition to the value of the clicks into their bids. In contrast, anecdotal evidence suggests that many advertisers dislike paying for impressions (Hoffman and Novak, 2000). There are several possible explanations for our contradictory result. First, initial pricing of online ads used posted prices rather than an auction model. It is therefore possible that publishers did not have enough information regarding the value of the ads to potential advertisers. Indeed, prices in terms of CPM dropped sharply from initial levels to be more commensurate with value to advertisers (Hoffman and Novak, 2000). Second, it should be noted that we assume that the advertisers know the CTRs of ads. Thus, our results apply to a post-learning phase when advertisers have learned the CTRs that can be expected. It can be argued that advertisers may prefer paying for clicks rather than impressions since the bids in the former case do not have to incorporate the CTR, which may be uncertain during the learning phase (compare Proposition 4 with Proposition 2, for instance). Third, risk-averse advertisers may prefer payment based on clicks rather than impressions (see also footnote 9). Risk aversion combined with moral hazard on the part of the publisher with respect to targeting can also lead to a preference for payment by clicks (Hu, 2004). However, publisher’s moral hazard may be less of an issue for search-related advertising where the targeting is explicitly known to be based on a search keyword or a context.

How does search-related advertising compare with other less targeted advertising? Our results suggest that search-related advertising is more profitable to advertisers than other less targeted

¹³Specifically, for less targeted online advertising and with the HCPC policy, the publisher’s profit is lower when payment is for clicks rather than impressions. The implication for managers is that if you are committed to the HCPC policy with bids per click, and the ads are less targeted (not search-related), then payment for impressions is a better option.

forms of online advertising (e.g., traditional banner ads). Conversely, search-related advertising is less profitable to the publisher. The intuition is that search-related advertising provides better knowledge of the specific need of the website user. This allows the advertiser who has strategically positioned to satisfy this specific need to outbid his competitor and still realize a significant profit. In our model, an advertiser can achieve a better positioning towards a particular need (such as a Howard Miller wall clock) by having a product line and website that is geared towards satisfying this need. In contrast, if the need of the website user is less precisely known, it puts competitors with different positioning strategies on a more level playing ground. This increases the intensity of the bidding competition reducing advertiser's profit while benefiting the publisher. Taken together with our results on impressions versus clicks, these results suggest that in making their advertising decisions, advertising managers should place more emphasis on the targeting abilities of the publisher rather than on whether payment is based on clicks or impressions.

What are the implications of HCPC and HCXR policies adopted by the publisher or search engine? Prima facie, it may appear that the HCPC policy is likely to be inferior to the HCXR policy for a publisher since the former ignores the CTRs of ads (see footnote 2 on this point). However, surprisingly, we find that for search- or context-related ads, the HCPC policy can generate the same profit as HCXR, and can even do better under some circumstances. Specifically, if consumers have no prior knowledge of the advertisers, HCPC and HCXR yield identical profits to the publisher, with an identical winning bid in either case. The intuition is that with no prior knowledge of the advertisers, the CTRs that will be commanded by either advertiser becomes independent of their positioning with regards to the specific need suggested by the search keyword. In other words, a consumer cannot identify based on the ad if the advertiser is the one who can best satisfy the need. (However, advertisers realize different payoffs once the consumer is at their website and understands the advertiser's positioning.) Thus, differences in CTR are non-existent in this case resulting in

the equivalence of HCPC and HCXR. However, if consumers have some prior knowledge, HCPC, interestingly generates higher profits for the publisher (and lower profits to the advertiser). The intuition is that a competing advertiser not ideally positioned to meet the need suggested by a search keyword will enjoy a lower CTR due to consumer's prior knowledge of the advertiser's positioning. This allows the advertiser who is more closely positioned towards the need to win with a lower bid under the HCXR policy (in comparison to the HCPC policy which ignores CTR), resulting in a higher profit for the advertiser and lower profit for the publisher. Given the different implications based on consumers' prior knowledge of advertisers, the question arises as to whether consumers are likely to have little or more prior knowledge of advertisers in real-world situations. Unless the advertisers are very small, it is possible that consumers have some prior knowledge of advertisers in real-world situations. Thus, it would appear that HCPC is a better policy from the publisher's point of view.

For less targeted advertising (and if search keywords sometimes lead to the wrong consumer), we find that HCXR leads to a higher publisher profit than HCPC. In this case, the need to outbid the competitor under HCPC leads each advertiser to tailor his ad narrowly to his target consumer since these consumers can support the highest possible CPC bid. However, this results in a lower profit from a lower overall CTR, as consumers who find the narrowly tailored ads irrelevant to them do not click on the advertisement. A benefit of the HCPC policy, however, is that consumers and advertisers are matched better under HCPC.

In sum, our analysis suggests that for search-related advertising, a publisher can do at least as well, if not better, with the HCPC policy in comparison to a HCXR policy. Conversely, an advertiser may realize lower profits when dealing with a publisher using HCPC. However, if the ad is targeted less precisely using demographic characteristic rather than a search keyword, or if the error in search-related consumer leads is sufficiently high, the publisher is likely to be better off

with the HCXR policy. Our analysis further suggests that for less targeted advertising, advertisers should tailor the ad narrowly to their target consumer when the publisher uses the HCPC policy, and use a blanket ad that appeal to all consumers when the publisher uses the HCXR policy.

Can volume discounts be a useful device for publishers? We find that volume discounts offered to advertisers based on the number of clicks or impressions can increase the publishers' profit while also achieving the maximum channel profit. However, advertisers' profit declines as the volume discounts induce them to compete more vigorously for advertising to consumers who might best be served by competitors. It is interesting to observe that volume discounts which have useful properties in the marketing of products (cf. Jeuland and Shugan, 1983) can be a beneficial tool to publishers in an online advertising context.

What are the implications of offline advertising by advertisers for their profitability from online advertising? If an advertiser's offline ads stimulates a higher CTR for his online ads, we find that the advertiser could benefit most by narrowly targeting his offline advertising to his best consumers. Such narrow targeting benefits all advertisers by increasing their profit from online advertising. An advertiser does not gain by unilaterally deviating to a strategy of also targeting his competitor's best consumers, since he cannot ultimately outbid the competitor to acquire these consumers in the online world. We also find that offline advertising can be more useful to the advertiser under the HCXR policy than under the HCPC policy. Conversely, the publisher's profits are hurt by offline advertising in the case of the HCXR policy. Thus, offline advertising has a synergistic benefit in the online advertising arena for advertisers.

What are the implications of asymmetry among competing advertisers? When the product market consists of niche sellers (advertisers) who appeal to a small segment of consumers, we find that such sellers are shut out of the online advertising market if the ads do not allow for precise targeting. However, with the more precise targeting afforded by search-related advertising, niche

sellers can more profitably compete for online advertising space.

We conclude with a discussion of the limitations of our study and with some suggestions for further research. A limitation is that we use a stylized model with two differentiated advertisers in our analysis. While others (e.g., Dukes and Gal-Or, 2003) also consider only two advertisers, it may be helpful to analyze a model with many more potential advertisers. It may also be useful to consider the impact of price competition in the product market as is the case in Dukes and Gal-Or (2003). In our analysis of offline advertising, we focused on the implications of an increased CTR resulting from offline advertising. A more general analysis can also consider the potential for offline advertising to generate direct visits to an advertiser's website leading to a single modeling framework for studying both online and offline advertising. Finally, we consider a search engine or publisher who auctions a single slot to potential advertisers. Extending our analysis to a situation with multiple ad slots and multiple advertisers may provide additional insights.

References

- Baker, S. (2004a, November). The Online Ad Surge. *Business Week*.
- Baker, S. (2004b, December 27). Where The Real Internet Money Is Made. *Business Week*.
- Baye, M. R. and J. Morgan (2005). Brand and Price Advertising in Online Markets. Working Paper, Indiana University.
- Benoit, J. and V. Krishna (2001, January). Multiple-Object Auctions with Budget Constrained Bidders. *Review of Economic Studies* 68(1).
- Chen, Y., G. Iyer, and V. Padmanabhan (2002). Referral Infomediaries. *Marketing Science* 21(4), 23–45.

- Delaney, K. J. (2005, May 3). Internet Ads Click With Firms; Some Shift Budgets. *Wall Street Journal*.
- Delaney, K. J. and R. A. Guth (2005, August 3). Yahoo, Microsoft Enhance Ad Focus. *Wall Street Journal*.
- Dukes, A. and E. Gal-Or (2003). Negotiations and Exclusivity Contracts for Advertising. *Marketing Science* 22(2), 222–245.
- Edelman, R., M. Ostrovsky, and M. Schwarz (2005). Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords. Working Paper, Harvard University.
- Feng, J., H. K. Bhargava, and D. M. Pennock (2005). Implementing Paid Placement in Web Search Engines: Computational Evaluation of Alternative Mechanisms. *INFORMS Journal of Computing*. Forthcoming.
- Hoffman, D. L. and T. P. Novak (2000, May-June). How to Acquire Customers on the Web. *Harvard Business Review* 78(3), 179–185.
- Hu, Y. (2004). Performance-based Pricing Models in Online Advertising. Working Paper, MIT.
- Iyer, G. and A. Pazgal (2003). Internet Shopping Agents: Virtual Colocation and Competition. *Marketing Science* 22(1), 85–106.
- Jeuland, A. P. and S. M. Shugan (1983). Managing Channel Profits. *Marketing Science* 2(3), 239–272.
- Marckini, F. (2004). Subtle PPC Differences Can Boost ROI. <http://www.clickz.com/experts/search/results/article.php/3335131>.

Oser, K. (2006). Why Marketers Need to Take Search Seriously. http://adage.com/searchmarketing/article?article_id=108777.

Vickrey, W. (1961). Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance* 16, 8–27.

Weber, T. and E. Zheng (2005). A Model of Search Intermediaries and Paid Referrals. Working Paper, Stanford University.

Weisman, R. (2006). Virtual ads pose real threat to traditional media. *Boston Globe*, February 12, 2006.

Wolfstetter (1996). Auctions: An Introduction. *Journal of Economic Surveys* 10, 367–420.

A Technical Appendix

A.1 Proof of Proposition 1

Proof. We first consider symmetric equilibria. Consider first a symmetric equilibrium in which both advertisers employ blanket ads and bid b . Then, both advertisers get $\frac{N}{2}$ impressions each. Suppose $b < \frac{v_h+v_l}{2}c_h$, then an advertiser can improve his profit by bidding $b + \varepsilon$ and winning all N impressions, where ε is arbitrarily small. A $b > \frac{v_h+v_l}{2}c_h$ cannot be an equilibrium since it generates negative profits. With $b = \frac{v_h+v_l}{2}c_h$, neither advertiser gains by deviating on his bid since a higher bid generates profit and a lower bid does not improve profit. Further, an advertiser does not gain by deviating to a tailored ad since the maximum he would bid per impression with such an ad is $\frac{v_h c_h + v_l c_l}{2} < \frac{v_h+v_l}{2}c_h$. Thus, $b = \frac{v_h+v_l}{2}c_h$ with blanket ads is a symmetric equilibrium and each advertiser makes zero profit in this equilibrium. We assert that there can be no symmetric equilibrium in which each advertiser uses ads tailored to his consumer type and bids some b . Using arguments similar to the above, we can eliminate all values of b except for $b = \frac{v_h c_h + v_l c_l}{2}$ as a candidate for equilibrium. However, if $b = \frac{v_h c_h + v_l c_l}{2}$, an advertiser can gain by deviating to a blanket ad and bidding $b + \varepsilon$. Along similar lines, a symmetric equilibrium tailored to the competitor's type can also be shown to be unsustainable. Therefore, $b = \frac{v_h+v_l}{2}c_h$ with blanket ads is the only symmetric equilibrium. It can be readily shown that there can be no asymmetric equilibrium. Thus, the symmetric equilibrium identified above is the unique equilibrium. ■

A.2 Proof of Proposition 2

Proof. First, consider symmetric equilibria. Let the equilibrium bids for a consumer type (say i) be b_{self} from the advertiser who is of the same type (advertiser i) as the consumer and b_{opp} from the other type (type j). Advertiser j 's maximum profitable bid for an impression to consumer i

is $v_l c_h$ using an ad tailored to consumer type i (or if it is a blanket ad). Advertiser i values each impression shown to such a consumer at $v_h c_h$ if the ad is tailored to the i type (or if it is a blanket). Any $b_{\text{self}}, b_{\text{opp}}$ such that $b_{\text{self}} \neq b_{\text{opp}} \leq v_l c_h$ cannot be an equilibrium because the lower bidder can make positive profit by equaling the other's bid. Any $b_{\text{self}} = b_{\text{opp}} \leq v_l c_h$ is not sustainable as an equilibrium either because advertiser i , who now gets only $\frac{N}{4}$ consumers (of type i), can gain by increasing his bid by ε to get all $\frac{N}{2}$ consumers of type i . However, neither advertiser has an incentive to deviate when $b_{\text{self}} = v_l c_h + \varepsilon$ is bid with a tailored ad, and $b_{\text{opp}} = v_l c_h$. Since no other equilibrium exists for this consumer type, it also implies that the strategies constitute a unique symmetric equilibrium.¹⁴ It can be readily shown that there can be no asymmetric equilibrium.¹⁵ Thus, the symmetric equilibrium identified above is unique. ■

A.3 Proof of Proposition 3

A.3.1 Proof for Part 1

Proof. Consider the symmetric equilibria first. Let b be a symmetric equilibrium bid and let the advertisers tailor the ad to their respective consumers. In such a case, each of their ads is displayed to $\frac{N}{2}$ consumers randomly. $b \not\leq \frac{v_h c_h + v_l c_l}{c_h + c_l}$ since at least one advertiser has an incentive to deviate with a bid $b + \varepsilon$, where ε is arbitrarily small, so as to display its ad to all N consumers.

A bid $b > \frac{v_h c_h + v_l c_l}{c_h + c_l}$ only generates negative profits for each advertiser and is not viable. When $b = \frac{v_h c_h + v_l c_l}{c_h + c_l}$, neither advertiser gains by deviating with any other bid. Moreover, an advertiser

does not gain by deviating to blanket ads since the maximum feasible bid with a blanket ad is only

$\frac{v_h + v_l}{2} < \frac{v_h c_h + v_l c_l}{c_h + c_l}$. Thus, $b = \frac{v_h c_h + v_l c_l}{c_h + c_l}$ with tailored ads is a symmetric equilibrium.

¹⁴ In this case, there is another symmetric equilibrium which is trivial. The advertiser submits the same bid prices, but instead uses a blanket ad to tailor its own consumer type. We ignore this trivial alternate equilibrium when we state that there is a unique symmetric equilibrium. Similar trivial equilibria exist for other Propositions and we ignore them as well.

¹⁵ Similar to Footnote 14, we ignore the trivial asymmetric equilibrium that can be derived from the symmetric equilibrium identified above by assigning a blanket ad to one advertiser instead of a tailored ad.

Now consider a symmetric equilibrium where both advertisers bid b and use a blanket ad. A $b < \frac{v_h+v_l}{2}$ is not viable since incentive exists for at least one advertiser to deviate and outbid the rival by ε . A $b > \frac{v_h+v_l}{2}$ yields only negative profit to each advertiser. When $b = \frac{v_h+v_l}{2}$, neither advertiser has an incentive to deviate on its bid. However, an advertiser can gain by switching to tailored ads and bidding $\frac{v_h+v_l}{2} + \varepsilon$. Thus, a symmetric equilibrium with blanket ads is not viable. Along similar lines, a symmetric equilibrium tailored to the competitor's type can also be shown to be unsustainable. Finally, it can also be readily shown that no asymmetric equilibrium is feasible. Thus, the symmetric equilibrium identified above is unique. ■

A.3.2 Proof for Part 2

Proof. The proof is identical to that of Proposition 1 in Section A.1 except that the metric used by the bidders to compete is bids per click. ■

A.4 Proof of Proposition 4

A.4.1 Proof for HCPC

Proof. We begin by first focusing on symmetric equilibrium strategies. Let the equilibrium bids for a consumer type (say i) be b_{self} from the advertiser who is of the same type (advertiser i) as the consumer and b_{opp} from the other type (type j). Independent of the nature of the ads shown by the advertisers, advertiser j values a click from a consumer of type i at v_l and advertiser i values it at v_h . Any $b_{\text{self}}, b_{\text{opp}}$ such that $b_{\text{self}} \neq b_{\text{opp}} \leq v_l$ cannot be an equilibrium because the lower bidder can make positive profit by equaling the other's bid. Any $b_{\text{self}} = b_{\text{opp}} \leq v_l$ is not sustainable as an equilibrium either because advertiser i , who now gets only $\frac{N}{4}$ consumers (of type i), can gain by increasing his bid by ε to get all $\frac{N}{2}$ consumers of type i . However, neither advertiser has an incentive to deviate when $b_{\text{self}} = v_l + \varepsilon$ is bid with a tailored ad, and $b_{\text{opp}} = v_l$. Since no other

equilibrium exists for a specific consumer type, it also implies that the strategies above constitute a unique symmetric equilibrium. We can readily show that there can be no asymmetric equilibrium either. Thus, the symmetric equilibrium identified above is unique. ■

A.4.2 Proof for HCXR

Proof. This proof is identical to that in Section A.4.1 except that the publisher chooses the bidder with the highest expected profit. ■

A.5 Proof of Proposition 5

Proof. The logic for the proof is similar across both HCPC and HCXR except for minor differences, which we highlight. Let Δ be the discount offered by the publisher. In a symmetric equilibrium, there are three possible scenarios: (1) Each advertiser serves his competitor's type only, (2) Each advertiser serves both the consumer types equally, and (3) Each advertiser serves his consumer type only.

Consider the strategy of one advertiser, say i , in scenario (1). Since advertiser i already serves consumer type j , he is willing to bid up to $v_h + \Delta$ per click to also serve consumer type i . However, since we assume that advertiser j serves consumer i in equilibrium, j should be bidding $v_h + \Delta + \varepsilon$ per click in order to win. But, j values those consumers at $v_l \leq v_h$ each. Clearly, this generates negative profits. So advertisers are better off not participating and this equilibrium cannot be sustained.

Next, we consider scenario (2). Since both advertisers serve both consumer types in equilibrium, their bids should be equal for a given consumer type and across consumer types in a symmetric equilibrium. Since advertisers will not participate unless they obtain non-negative profit, there are two possible cases in equilibrium: (a) both advertisers make positive profit, and (b) both advertisers

make zero profit. Consider case (a). Then an advertiser should make a positive profit from at least one consumer type. Let i be one such consumer type and let the equilibrium bid be b for this consumer type. Then by bidding $b + \varepsilon$ for consumer type i , this advertiser gains because he gets $\frac{N}{2}$ rather than $\frac{N}{4}$ consumers of type i . Thus, case (a) cannot prevail in equilibrium. Consider case (b). If an advertiser makes a positive profit on one consumer type (say i) and makes an exact amount of loss on the other type (say j), he can gain in the same fashion as argued above by increasing his bid for consumer type i by ε . On the other hand, if an advertiser makes zero profit from both consumer types, by increasing the bid on one consumer type (say i) by ε , he loses $\frac{N\varepsilon}{2}$ from consumer type i but gains $\frac{Nc_h\Delta}{4}$ through the discount offered by the publisher since he serves $\frac{3N}{4}$ consumers. By making ε arbitrarily small, the advertiser can increase his profit. Thus, case (b) cannot prevail in equilibrium. Therefore, scenario (2) cannot prevail in equilibrium.

Finally, consider scenario (3). Consider the strategy of one advertiser, say j . Since advertiser i serves consumer type i , he is willing to bid up to $v_l + \Delta$ per click to also serve consumer type j . However, since we assume that advertiser j serves consumer j in equilibrium, j should be bidding $v_l + \Delta + \varepsilon$ per click in order to win. Note that $v_l + \Delta + \varepsilon \leq v_h$; otherwise, it generates negative profits for advertiser j and violates the participation condition. With the above bids, the publisher generates a profit of $Nc_h(v_l + \Delta + \varepsilon)$, an expression increasing in Δ . The publisher's profits are maximized when $\Delta = v_h - v_l - \varepsilon$, for ε arbitrarily small. When $\Delta > v_h - v_l$, an equilibrium with nonnegative profits for advertisers does not exist. ■

A.6 Proof for Proposition 6

A.6.1 Proof for Part 1

Proof. The proof for the HCPC and the HCXR policies are identical to that in Proposition 3 in Section A.3.1. The key point to note is that, independent of the nature of the ads, the maximum

CTR of ads from the competitor's type is c_l instead of c_h . Given this, the part of the proof relating to the symmetric equilibrium involving blanket ads will not be required. ■

A.6.2 Proof for Part 2

Proof. The proof for the HCPC policy is identical to that for Proposition 4 in Section A.4. Even though the maximum CTR of ads from the competitor's type is c_l instead of c_h , none of the details of the proof change. In the following, we only prove the equilibrium for the HCXR policy. We begin by first focusing on symmetric equilibrium strategies. Let the equilibrium bids for a consumer type (say i) be b_{self} from the advertiser who is of the same type (advertiser i) as the consumer and b_{opp} from the other type (type j). Independent of the nature of the ads shown by the advertisers, advertiser j values a click from a consumer of type i at v_l , and the ad from j generates a profit of $b_{\text{opp}}c_l$ to the publisher. Advertiser i values a click from consumer type i at v_h and his ad generates an expected publisher profit of $b_{\text{self}}c_h$. Any $b_{\text{self}}, b_{\text{opp}}$ such that $b_{\text{self}} \neq b_{\text{opp}} \leq v_l \frac{c_l}{c_h}$ cannot be an equilibrium because the lower bidder can make positive profit by equaling the other's bid. Any $b_{\text{self}} = b_{\text{opp}} \leq v_l \frac{c_l}{c_h}$ is not sustainable as an equilibrium either because advertiser i , who now gets only $\frac{N}{4}$ consumers (of type i), can gain by increasing his bid by ε to get all $\frac{N}{2}$ consumers of type i . However, neither advertiser has an incentive to deviate when $b_{\text{self}} = v_l \frac{c_l}{c_h} + \varepsilon$ is bid with a tailored ad, and $b_{\text{opp}} = v_l \frac{c_l}{c_h}$. Since no other equilibrium exists for a specific consumer type, it also implies that the strategies above constitute a unique symmetric equilibrium. We can readily show that there can be no asymmetric equilibrium either. Thus, the symmetric equilibrium identified above is unique. ■

A.7 Proof for Proposition 7

Proof. In any equilibrium, there can be four possible kinds of consumers based on their exposure to

offline advertising.: (i) a consumer who is exposed to neither advertiser's offline ads; (ii) a consumer who is exposed to both i 's and j 's offline ads; (iii) a consumer who is exposed only to advertiser i 's offline ads but not j 's; and (iv) a consumer who is exposed only to advertiser j 's offline ads but not i 's. In each of the above cases, the consumer exposed to the ads may be of type i or type j .

Let us first consider the HCPC policy. ■

A.7.1 Proof for the HCPC policy

Proof. The equilibrium and the payoffs for case (i) are identical to those in Proposition 4. Each advertiser obtains a profit of $c_h(v_h - v_l)$ from its own consumer type and 0 from the competitor's type.

Next consider case (ii) and let the consumer be of type i . In equilibrium, advertiser i bids $v_l + \varepsilon$ and j bids v_l (the equilibrium calculations similar to that in Section A.4.1). The profits for advertiser i are $((v_h - v_l)(c_h + \delta) - \alpha)$ and for advertiser j are $(-\alpha)$. The equilibrium for the case when the consumers of type j are exposed to both offline ads is symmetric.

Next, consider case (iii) and consider first a consumer of type i . Along the lines of the proof in Section A.4.1, the equilibrium is for advertiser i to bid $v_l + \varepsilon$ and for advertiser j to submit v_l . Advertiser i generates a profit of $((v_h - v_l)(c_h + \delta) - \alpha)$ and advertiser j generates zero profit. Next consider the situation when the consumer exposed to in case (iii) is of type j . The equilibrium is for advertiser i to bid v_l and for advertiser j to bid $v_l + \varepsilon$. It is important to note that since advertisers i 's targeting of j -type consumer only changes the CTRs for advertiser i 's online ad but not the value of this consumer to advertiser i , the equilibrium continues to remain the same as it was when advertiser i did not target j -type consumers. Therefore, the profits for advertiser j is $(v_h - v_l)c_h$ and, for advertiser i , it is $(-\alpha)$. Due to the symmetry in the model, the analysis for case (iv) is identical to that for case (iii).

Player j 's action \rightarrow Player i 's action \downarrow	Target	Not to Target
Target	$((v_h - v_l)(c_h + \delta) - \alpha, -\alpha)$	$((v_h - v_l)(c_h + \delta) - \alpha, 0)$
Not to Target	$((v_h - v_l)c_h, -\alpha)$	$(c_h(v_h - v_l), 0)$

Table 2: Normal form representation of the advertiser's decisions with respect to consumer type i . In each cell, the profit of the advertiser i is first presented, and then j .

Table 2 represents in a normal form the advertiser profits under the HCPC policy from the offline-ad decisions. From this table, it is easy to see that when $(v_h - v_l)\delta \geq \alpha$, the unique equilibrium is for advertiser i to target consumer type i , and for advertiser j to not target this consumer type. However, when $(v_h - v_l)\delta < \alpha$, neither advertiser type targets consumer i with offline advertising in the unique equilibrium. A similar conclusion applies for targeting of consumer type j by the advertisers. ■

A.7.2 Proof for the HCXR policy

Proof. Similar to the proof for the HCPC policy, we consider each case separately. Note that the equilibrium and the payoffs for case (i) is identical to those in Proposition 4. Now consider case (ii). When consumers of type i are exposed to offline ads from both advertisers, advertiser i bids $v_l + \varepsilon$ and advertiser j bids v_l in equilibrium. The payoff for advertiser i is $((v_h - v_l)(c_h + \delta) - \alpha)$, and for advertiser j , it is $(-\alpha)$. The equilibrium for the case when the consumers of type j are exposed to both offline ads is symmetric.

Next consider case (iii) and consider a consumer of type i . In that case, the equilibrium bid prices are as follows: advertiser i bids $v_l \frac{c_h}{c_h + \delta} + \varepsilon$ and advertiser j bids v_l . The payoff for advertiser i is $((v_h - v_l)c_h + v_h\delta - \alpha)$ while advertiser j generates zero profits. Now consider a consumer of type j . Depending on whether $v_l(c_h + \delta) > v_h c_h$ or not the equilibrium is different. If $v_l(c_h + \delta) > v_h c_h$, advertiser i 's bid of $v_h \frac{c_h}{c_h + \delta} + \varepsilon$ while advertiser j 's bid is v_h . The profits for advertiser i and advertiser j are $(v_l(c_h + \delta) - v_h c_h - \alpha)$ and zero respectively. On the other hand, if $v_l(c_h + \delta) \leq v_h c_h$,

advertiser j bids $v_l \frac{c_h + \delta}{c_h} + \varepsilon$ and advertiser i bids v_l . The profits then are $((v_h - v_l)c_h - v_l\delta)$ for advertiser j and $(-\alpha)$ for advertiser i .

Due to the symmetry in the model, the analysis for case (iv) is identical to that for case (iii).

Player j 's action \rightarrow Player i 's action \downarrow	Target	Not to Target
Target	$((v_h - v_l)(c_h + \delta) - \alpha), -\alpha$	$((v_h - v_l)c_h + v_h\delta - \alpha), 0$
Not to Target	$(0, (v_l(c_h + \delta) - v_h c_h - \alpha))$ if $v_l(c_h + \delta) > v_h c_h$ $((v_h - v_l)c_h - v_l\delta), -\alpha$ if $v_l(c_h + \delta) < v_h c_h$	$(c_h(v_h - v_l), 0)$

Table 3: Normal form representation of the advertiser's decisions with respect to consumer type i . In each cell, the profit of the advertiser i is first presented, and then j . In the cell corresponding to the actions (Target, Not to Target), the profits differ depending on whether $v_l(c_h + \delta) > v_h c_h$ or not.

Table 3 provides in a normal form the advertiser profits from the offline-ad decisions under the HCXR policy. From this table, it is easy to see that when $v_h\delta \geq \alpha$, the unique equilibrium is for advertiser i to target consumer type i , and for advertiser j to not target this consumer type. However, when $v_h\delta < \alpha$, neither advertiser type targets consumer i with offline advertising in the unique equilibrium. A similar conclusion applies for targeting of consumer type j by the advertisers.

■