Title
Contextual Advertising

Journal
Marketing Science, 31(6)

ISSN
0732-2399

Authors
Zhang, Kaifu
Katona, Zsolt

Publication Date
2012-11-01

DOI
10.1287/mksc.1120.0740

Peer reviewed

Contextual Advertising

Kaifu Zhang and Zsolt Katona[1]

July 14, 2012

[1]Kaifu Zhang is Assistant Professor in Marketing at Cheung Kong Graduate School of Business, Beijing, China. Zsolt Katona is Assistant Professor of Marketing at the Haas School of Business, UC Berkeley, Berkeley, CA 94720. Email: kfzhang@ckgsb.edu.cn, zskatona@haas.berkeley.edu.

Contextual Advertising

Abstract

Contextual advertising entails the display of relevant ads based on the content that consumers view, exploiting the potential that consumers' content preferences are indicative of their product preferences. This paper studies the strategic aspects of such advertising, considering an intermediary who has access to a content base, sells advertising space to advertisers who compete in the product market, and provides the targeting technology. The results show that contextual targeting impacts advertiser profit in two ways: first, advertising through relevant content topics helps advertisers reach consumers with a strong preference for their product. Second, heterogeneity in consumers' *content* preferences can be leveraged to reduce *product* market competition, especially when competition is intense. The intermediary has incentives to strategically design its targeting technology, sometimes at the cost of the advertisers. When product market competition is moderate, the intermediary offers accurate targeting such that the consumers see the most relevant ads. When competition is high, the intermediary lowers the targeting accuracy such that the consumers see less relevant ads. Doing so intensifies competition and encourages advertisers to bid for multiple content topics in order to prevent their competitors from reaching consumers. In some cases, this may lead to an asymmetric equilibrium where one advertiser bids high even for the content topic that is more relevant to its competitor.

1 Introduction

> “ Google's toughest search is for a business model... In other words, can Google create a business model even remotely as good as its technology?” - New York Times, April 2002

Eight years after the above observation was made, Google's annual revenue had surpassed \$20 billion thanks to its immensely successful AdWord search advertising and AdSense affiliate advertising programmes. Years after the established success of Google, critics are casting doubts in a similar fashion on the emerging social media sites such as Facebook, YouTube and Twitter[1], questioning their ability to monetize their user base. Once again, advertising seems to be the answer. YouTube, Facebook and Twitter had all implemented their advertising programmes, and they have become some of the fastest growing advertising outlets on the Internet.

Why do successful search engines, video sharing websites, micro-blogging sites and social networks alike embrace advertising as their preferred business model? Other than their broad reach to Internet users, all the above-mentioned sites have offered *contextual targeting* as a major value proposition. Contextual advertising refers to the targeted delivery of advertisements according to the content each consumer views. Consider an example from YouTube. A video named ‘park ride’ features a stunt-performing cyclist. Viewers of the this video in early 2010 saw an overlay flash ad from the bike maker Lynskey, which makes customized performance bikes. Such accurate targeting is made possible by the wealth of user-generated videos on YouTube that cover a wide range of topics. Similarly, Google's AdSense[2] network attracts an enormous number of Internet content publishers who wish to monetize their websites. In early 2010, an ad for Dahon foldable bike was displayed on foldforum.com, which is a general interest discussion forum for foldable bike lovers. An ad from Organic Bikes,

[1] For example, see the article “YouTube is Doomed.” http://www.businessinsider.com/is-youtube-doomed-2009-4

[2] Google AdSense displays ads on third party publishers's sites. This program is not the same as Google AdWords which allows advertisers to bid to appear on Google's search page. However, the bidding process is similar as advertisers bid for keywords that are matched to the content provided by the third party publishers.

a company that sells bicycles made of bamboo, was displayed on the biking advocacy section of bikeforums.com, a popular gathering spot for environmentally minded city commuters. It is the heterogeneity in member publishers' content that allows the AdSense network to choose the most relevant content and deliver such highly targeted ads.

The idea of targeted advertising has a long history in the advertising industry. Advertising agencies sometimes offer media planning services to their clients, choosing the advertising medium (e.g., newspapers or TV channels) according to the type of product being advertised. However, the most exciting developments in contextual advertising have taken place in the online environment. There are two reasons for this. First, online intermediaries such as Google AdSense typically boast massively heterogeneous content bases, which can be leveraged to deliver finely targeted ads to a large Internet population. In addition, the development of sophisticated content analysis algorithms has automated the ad-content matching process and made it possible on a large scale. The goal of this paper is to study the phenomenon of contextual advertising, focusing on two unique aspects of the ecosystem:

- First, we explicitly model the idea that a consumer's content preference is indicative of her product preference. We assume that the consumers browse content and purchase products. Each consumer's product preference can be imperfectly inferred from what content topics she browses. Said differently, consumers' content preference and product preference are partially aligned. The advertisers can exploit this alignment - which we call *content relevance* - to target the consumers. We examine the link between consumers' preference structure and the effectiveness of contextual targeting.

- Second, we consider an independent, profit maximizing intermediary. The intermediary sells advertising spaces via second price auctions and implements the targeting technology. For example, Google Adsense relies on content analysis algorithms to determine the exact content of each webpage in order to deliver the relevant ad. We investigate the intermediary's incentives to strategically manipulate targeting accuracy.

We set up an analytical model with horizontally differentiated product firms. In order for the products to enter consumers' consideration sets, firms have to communicate their product information to potential consumers through advertising. Consumers encounter advertising when they browse content and their preferences for different content topics are (to some extent) correlated with their product preferences. A contextual advertising intermediary delivers an ad on each content page using its targeting technology. Firms bid for the advertising slots associated with different content topics and engage in price competition in the product market.

In the baseline analysis, we discuss two distinct roles of contextual advertising: first, by exploiting the alignment between the consumers' product and content preferences, advertisers can efficiently reach the consumers who have strong preference for their product. Second, by exploiting the *misalignment* between the consumers' product and content preferences, advertisers can create a type of 'informational differentiation', wherein some consumers only see one ad for one product although they like both products equally. Consequently, advertiser profits can either increase or decrease in content relevance, depending on the competitiveness in the product market.

These results resonate with some of the basic findings from the targeting and individual marketing literature (Chen, Narasimhan, and Zhang 2001, Iyer, Soberman, and Villas-Boas 2005). Our main contribution concerns the role and the strategic motives of the contextual advertising *intermediary*. In our model advertisers not only compete in the product market for consumers, but also compete for advertising space provided by the intermediary. We study the auctions conducted by the intermediary and the resulting equilibrium allocation of advertising slots when advertisers bid for the content topics in separate second price auctions.

In a duopoly setup, we find that each firm tends to advertise through the more relevant content topic as long as product market competition is moderate. When product market competition is sufficiently high, each firm has a strong incentive to bid for the content topic that is more relevant to her competitor. In this way, a firm is able to preempt the competitor from advertising to the consumers who like both products. The incentive of competitive reduction can be strong enough to lead to a 'topic shelving' equilibrium wherein one firm advertises through both the relevant content topic and less relevant content topic. This analysis suggests two distinct drivers of the advertisers willingness to pay for content topics: (1) to generate additional traffic; and (2) to reduce competition. When reducing competition is the main objective, an advertiser may have a high willingness to pay for a content topic that generates little additional traffic.

One of our most important results discusses the intermediary's optimal choice of targeting accuracy. We find that when product market competition is low, the intermediary should offer accurate targeting such that the firms can reach distinct, minimally overlapping audiences by advertising through different content topics. As long as content relevance is not perfect, the audience of each content topic includes both firms' strong preference consumers. Consequently, each advertiser is interested in acquiring as many content topics as possible since each ad slot brings in additional traffic. This drives up the equilibrium payments in the second price auction. When product market competition is high, however, the intermediary should intentionally offer inaccurate targeting technology which decreases the advertisers' ability to create informational differentiation. This further intensifies price competition and drives up the firms' incentives to bid for their competitors' content topics in order to reduce competition. In doing so, the intermediary may endogenously induce the topic shelving equilibrium which would not have taken place if targeting is accurate.

Extending our basic model, we consider an extension with multiple intermediaries to analyze how advertisers' incentives change if there are two alternatives for buying advertising space in the same topic. We find that in this model with two identical intermediaries, the results are largely similar to those of our basic model as topic shelving and the non-shelving outcomes occur in the same parameter regions. However, we find an interesting additional type of duopoly equilibrium wherein advertisers differentiate in which intermediary they advertise at. Only in this latter case, are the intermediary profits lower than half of what a single intermediary would make.

The rest of the paper is organized as follows. We summarize the related literature and our relative contribution in Section 2. In Section 3 we present the model and conduct the basic analyses in Section 4. We study the intermediary's strategic choice of targeting accuracy in Section 5, and consider the presence of multiple intermediaries in Section 6. Finally, we discuss the practical implications of our results and conclude in Section 7.

2 Related Research

This paper is broadly related to three literature streams. First, our paper is closely related to the earlier work on the targeting of advertising. Iyer et al. (2005) argue that targeted advertising can help advertisers mitigate product market competition and increase advertiser profits in a competitive industry. Similarly, Gal-Or and Gal-Or (2005) consider a setup where customized advertising is intermediated by a common media distributor. The authors found that the media distributor can implement monopoly pricing utilizing customized advertising. More recently, Bergemann and Bonatti (2011) study the differences between offline and online advertising markets, focusing on the differences in targeting ability. Similar to these articles, we focus on the *competitive* implications of contextual advertising, e.g., how contextual targeting help competing advertisers reduce their product market competition[3]. In addition, we model three unique institutional details of the contextual advertising market not considered in this literature. First, contextual advertising relies crucially on the existence of heterogeneous media content which consumers browse. We explicitly model consumers' correlated preference structures for the content and the products, and study content relevance as a key parameter. Second, we consider the ad space auction process that is typical of many online contextual advertising platforms. Finally, we consider the strategic choices of the contextual advertising intermediary that are unique to the contextual advertising industry, such as the provision of targeting technology. Conceptually, our notion of imperfect contextual targeting is related to the idea of individual targetability proposed by Chen et al. (2001). They show that advertiser profits may increase or decrease as targetability increases. We obtain a similar result

[3]Empirically, Dong, Manchanda, and Chintagunta (2009) show that accounting for firms' strategic behavior is important in quantifying the benefit of segment-level targeting, thus supporting the focus on competitive implication in the advertising targeting literature.

concerning the relevance of content in contextual advertising. Using this result as a building block, our contribution lies in modeling the entire contextual advertising ecosystem including an intermediary who has access to the content base and provides the targeting technology. In addition to determining the allocation of advertising slots for different content topics, we also uncover the strategic incentives of the intermediary to manipulate targeting accuracy.

Second, our paper is related to the growing literature on keyword auctions (Chen and He 2011, Edelman, Ostrovsky, and Schwarz 2007, Katona and Sarvary 2010, Varian 2007) in search advertising. Most papers in this stream focus on understanding the properties of the widely adopted keyword auction mechanism, such as the Generalized Second Price auction. The auction design used in contextual advertising is essentially identical. Therefore, our model of ad space auction is built upon the auction mechanisms and solution concept proposed by Edelman et al. (2007) and Varian (2007). Search advertising can be considered as an extreme example of contextual advertising, wherein each 'keyword' can be considered as a 'content topic' that is maximally relevant to the product. This paper offers two distinct additions to the literature. On one hand, we study an environment that is more general and less sterile than search advertising in that content topic preference do not express interest in specific products as clearly as search for specific keywords. On the other hand, we address the natural problem that most advertisers involved in contextual or search advertising face: how much to bid for different topics when engaging in product market competition with other advertisers? Aside from a few papers (Sayedi, Jerath, and Srinivasan 2011, Shin 2009) studying the specific question of bidding for a competitor's branded keyword in search advertising this problem has not been explored.

Third, by explicitly considering the contextual advertising intermediary as an independent market player, our model is generally related to literature on distributional channels (see for example Coughlan (1985), Coughlan and Lal (1992)) and, specifically, earlier studies on commercial media (Dukes 2004, Gal-Or and Dukes 2003, Gal-Or and Gal-Or 2005, Gal-Or, Geylani, and Yildrim 2012). With this literature, we share the general idea that product firms can reduce price competition via the differentiation created by other channel players (for example, competing manufacturers can create differentiation by selling through geographically distant retailers.). Different from these works, we study content relevance, ad space auction and intermediary strategic decisions that are unique to the contextual advertising industry.

On a technical level, our model of price competition follows the widely adopted formulation by Narasimhan (1988), Varian (1980). Our conceptualization of advertising follows the informational advertising paradigm, wherein advertising helps a product enter a consumer's consideration set (Grossman and Shapiro 1984, Iyer et al. 2005, Nelson 1974).

3 The Model

We consider a market with two horizontally differentiated firms[4] offering their products to a unit mass of consumers. We assume that consumers learn about the existence of the products through informative advertising such that they do not buy a product unless they are aware of it. We assume that the potential consumers are primarily Internet users and firms can reach them while they browse content on the Web. Firms have the option of advertising through a contextual advertising intermediary. The intermediary delivers targeted ads to consumers browsing certain content topics on behalf of the advertisers. Consumers are heterogeneous with respect to the content topics they browse and these preferences may be correlated with their product preferences.

3.1 Consumers

Product Preferences: We adopt a standard discrete horizontal differentiation model with three segments (Iyer et al. 2005, Narasimhan 1988, Varian 1980). Each firm has a segment of consumers who have strong preferences for its product; a third segment prefer the two products equally and are therefore 'price shoppers'. Consumers who have strong preferences for firm i's product value their preferred product at 1 and the other product at 0. The price shoppers value both products at 1. We assume a symmetric setup in which the fraction of

[4]We use the terms 'firm' and 'advertiser' interchangeably.

price shoppers is α_p, yielding a strong-preference segment of size $\frac{1-\alpha_p}{2}$ for each firm. That is $\frac{1-\alpha_p}{2}$ captures the heterogeneity in consumer product preferences whereas α_p essentially measures the competitiveness of the product market.[5]

Content Preferences: We assume that there are two different content topics[6] on the Internet that have some relation to the products sold.[7] Similarly to the product market, we assume that consumers have heterogeneous preferences for content: an α_c fraction of them browse both content topics, whereas $\frac{1-\alpha_c}{2}$ consumers browse exclusively topic 1 and another $\frac{1-\alpha_c}{2}$ browse only topic 2. Hence $\frac{1-\alpha_c}{2}$ captures the heterogeneity in consumers' content preferences.

Preference Alignment: We assume that a consumer's interest in a particular content topic may be indicative of her preference for a particular product. A consumer who only browses content topic 1 is more likely to have a strong preference for product 1. Formally, we use s_{ij} to measure the number of customers who have a preference for product i and content topic j. The indices i and j can take the values of 1, 2 and b, where b indicates that a consumer has positive valuations for both products or that she browses both content topics. For example s_{b1} is the number of consumers who have positive valuation for both products, but only browse content topic 1, whereas s_{2b} is the size of the segment that have strong preferences for product 2, but browse both content topics. Collectively, the nine s_{ij} values capture the relationship between consumer preference distribution in the product and content markets. For example, when consumer preferences for content topics are independent of product preferences, we have

$$s_{11}^I = s_{21}^I = s_{12}^I = s_{22}^I = \frac{1-\alpha_p}{2} \cdot \frac{1-\alpha_c}{2}, \quad s_{b1}^I = s_{b2}^I = \alpha_p \frac{1-\alpha_c}{2},$$

$$s_{1b}^I = s_{2b}^I = \alpha_c \frac{1-\alpha_p}{2}, \quad \text{and} \quad s_{bb}^I = \alpha_p \alpha_c.$$

In other words, a consumer's content preference is totally uninformative of her brand pref-

[5] We do not assume that consumers who have a positive valuation for a product are necessarily aware of it. Our interpretation of this preference structure simply intends to capture heterogeneity in consumers' willingness to pay for different product types (e.g., specifications or features), rather than brand loyalty as some other papers using the same model do. Our approach is identical to that of Iyer et al. (2005).

[6] In Section 3 of the Online Appendix, we consider a setting with three different keywords.

[7] 'Content topics' could refer to user-generated videos in different categories on Youtube or the different types of affiliated websites in the AdSense network. These videos or websites have heterogeneous themes and attract different viewers.

erence. On the other extreme, when $\alpha_p = \alpha_c = \alpha$ and the preference for content topics is perfectly aligned with the preference for products, we have

$$s_{11}^A = s_{22}^A = \frac{1-\alpha}{2}, s_{bb}^A = \alpha, \quad \text{and} \quad s_{12}^A = s_{21}^A = s_{b1}^A = s_{b2}^A = s_{1b}^A = s_{2b}^A = 0.$$

Put differently, a consumer is interested in product i if and only if she is interested in content topic i. Thus, a consumer's content preference is maximally informative about her brand preference. When $\alpha_p \neq \alpha_c$, product and content preferences cannot be perfectly aligned, since different percentages of consumers are interested in both products than in both content topics. We define maximally aligned product and content preferences, using

$$s_{11}^A = s_{22}^A = \min\left(\frac{1-\alpha_c}{2}, \frac{1-\alpha_p}{2}\right), \quad s_{bb}^A = \min(\alpha_p, \alpha_c), \quad s_{b1}^A = s_{b2}^A = \frac{\max(\alpha_p - \alpha_c, 0)}{2},$$

$$s_{1b}^A = s_{2b}^A = \frac{\max(\alpha_c - \alpha_p, 0)}{2}, \quad \text{and} \quad s_{12}^A = s_{21}^A = 0.$$

To simplify notation and measure alignment between product and content preferences using a single parameter, we introduce ρ and place constraints on the possible values of s_{ij} as follows:

$$s_{ij} = (1-\rho)s_{ij}^I + \rho s_{ij}^A. \tag{1}$$

When $\rho = 0$, content and product preferences are independent, whereas when $\rho = 1$, they are maximally aligned.

Browsing and Clicking Behavior: In order for a product to enter a consumer's consideration set, the consumer needs to receive an ad for this product. As in Iyer et al. (2005), the consumers' product preferences can be considered as their endowed preferences over product attributes; but the consumers do not know the available options in the market. Advertising informs the consumers about the existence of the products and their attributes.

We assume that each consumer browses N webpages on the content topic(s) she is interested in. Each time a consumer browses a page, she sees an ad and clicks on it if she has positive valuation for the product featured in that ad.[8] Thus, a consumer's content preference

[8]We do not formally model the utility a consumer can expect from clicking on an ad. The assumption that a consumer always clicks when her valuation is positive seems rather strong, but an alternative formulation in which consumers click with a certain (fixed) probability yields similar results.

determines which pages she will browse. A consumer's product preferences determine which ads she will click on conditional on seeing them. For example, a consumer who is interested in both products will click on both ads if she sees them. However, she may not have the chance to click on a product's ad if it is never displayed on any of the N pages she browses. Which ad will be displayed on a content page depends on the outcome of the ad auction process and the intermediary's targeting technology, which we explain in the next subsection.

We assume that each firm can also reach the consumers through an 'outside' medium, such as TV, newspaper, or online channels other than what the contextual ad intermediary offers. For simplicity, we assume that a γ fraction of consumers follow this outside medium and that both firms advertise through it. This guarantees that each firm will always stay in the market even without advertising through the contextual advertising intermediary.

3.2 Advertisers and the Intermediary

Advertisers bid for the ad slots associated with different content topics in separate second price auctions.[9] If a firm is the highest bidder for a content topic, the intermediary promises to deliver that firm's ad on each content page belonging to that topic. However, the intermediary's targeting technology can be imperfect.

Ad Delivery and Targeting Technology: Suppose that firm 1 wins the ad slot for content topic 1 and firm 2 wins the ad slot for content topic 2. The intermediary is responsible for displaying firm 1's ad on every content page that belongs to topic 1. However, the intermediary can only achieve this objective imperfectly. Each time a consumer browses a webpage in content topic 1, there is a ϑ probability that she will (correctly) see the ad for product 1; with probability $1-\vartheta$ she will (incorrectly) see the ad for product 2. More generally, $1-\vartheta$ denotes the probability that the intermediary makes a mistake, in which case it displays an ad intended for content topic i on a page that is actually in content topic j. Clearly, $\vartheta = 1$ represents the case of a perfect targeting technology, while $\vartheta = \frac{1}{2}$ represents a case of totally

[9] In the basic model, we assume there is one advertising slot on each content page. We present analysis with multiple ad-slots in Section 2 of the Online Appendix.

ineffective targeting technology. We assume that the probability of mismatching is identical and independent across the webpages. Clearly, when a consumer browses N pages in her preferred content area, she has a $1 - \vartheta^N$ chance of seeing a 'incorrect' ad intended to be displayed for the other content topic.

In the case of AdSense, Google involves a large number of web pages in its content network. AdSense relies on sophisticated content analysis algorithms to infer the content topic of each page. ϑ essentially denotes the probability that the algorithm will correctly infer each page's content. If the algorithm makes an incorrect inference, an irrelevant ad may be displayed. This does not happen if $\vartheta = 1$, whereas $\vartheta = \frac{1}{2}$ would mean that Google randomly assigns ads to pages.

Content Topic (i.e., Ad Slot) Auction: We consider a second price Pay-per-click (PPC) auction with click-through rate (CTR) correction, which is the most standard format used in the industry. The intermediary runs an auction where each advertiser submits a bid for each content topic separately, and bids are corrected for expected CTRs when determining the winner.[10]

Equilibrium Concept: To solve the advertising space auction, we extend the so-called envy-free equilibrium concept (Edelman et al. 2007, Varian 2007) to multiple items. This type of equilibrium is a widely used concept for sponsored link auctions. The basic idea is that when a bidder considers deviating from her equilibrium strategies, and possibly acquiring a different slot, she evaluates the deviation by using the price that is currently paid for that slot. This is a stronger condition for profitable deviations than in a simple Nash equilibrium as the currently paid price may be higher than what the deviating bidder will eventually pay due to the possible change in order. Thus, the set of envy-free (or symmetric) equilibria is a subset of the Nash-equilibria. We generalize this concept to a setting with simultaneous auctions for

[10]Specifically, the intermediary determines the winner by ranking $PPC_i * CTR_i$. Suppose $PPC_i * CTR_i > PPC_{-i} * CTR_{-i}$. Then the winning advertiser i will pay the adjusted second price $\frac{PPC_{-i} * CTR_{-i}}{CTR_i}$ for each click she receives. In total, she will pay $PPC_{-i} * CTR_{-i}$. This is the auction procedure used by many online contextual advertising intermediaries, such as Google AdSense. When click-through rates can be perfectly estimated, bidding by PPC is equivalent to bidding by impression.

multiple items (in this case advertising slots). Specifically, in any envy-free equilibrium, each bidder considers deviations that entail acquiring or giving up a *set* of advertising slots at the sum of their current prices[11]. As is common in the literature, we evaluate auctioneer profit based on the minimal-revenue envy-free equilibrium.

Pricing: Firms inform their consumers through both contextual advertising and the 'outside' advertising medium. Once consumers learn about the products, firms set prices and consumers purchase. We normalize the marginal cost of production to zero.

3.3 Timing

To summarize, we present the key model parameters in Table 1 and describe the timing below:

1. *Intermediary Strategy:* the intermediary makes strategic decisions such as targeting precision (ϑ).
2. *Advertising Slot Allocation:* The contextual advertising intermediary organizes second-price auctions to allocate the advertising slots associated with each content topic to the advertisers.
3. *Browsing and Advertising:* The ads are delivered based on the auction results. Consumers browse the content pages they are interested in and see/click on the ads.
4. *Pricing:* Advertisers set prices for their products.
5. *Purchasing:* Consumers make purchase decisions and profits are realized.

4 Analysis

To use backward induction, we start by determining the equilibrium advertiser revenues in the subgame where each firm advertises through the more relevant content topic (i.e., firm i

[11] Although the auction involves multiple items, we do not consider the possibility of combinatorial auctions.

Parameter	Interpretation
Consumer Preference	
α_p	Product market competitiveness
α_c	Content preference homogeneity
ρ	Content relevance
Intermediary choice	
ϑ	Targeting accuracy
Browsing & Outside media	
N	Number of pages browsed
γ	Outside media coverage

Table 1: Summary of Model Parameters

advertises to and only to the consumers who browse content topic i). In order to determine equilibrium prices, we need to count the consumers who only consider one product and those who compare prices. The solution of such a pricing game is well known in the literature (Narasimhan 1988, Varian 1980) with firms employing symmetric mixed strategies. Each firm's profit equals to the number of consumers who only consider a single product. In our model, a consumer will consider only a single product either because she has a strong preference for that product or because she is only aware of that product. By evaluating the sizes of the different consumer segments, we obtain the following expressions for advertiser revenues:[12]

$$\Pi_d = (1-\gamma)S_c + \gamma\frac{1-\alpha_p}{2} \tag{2}$$

where S_c is defined as follows:

$$S_c = (1-(1-\vartheta)^N)s_{11}+(1-\vartheta^N(1-\vartheta)^N)s_{1b}+(1-\vartheta^N)s_{12}+(1-\vartheta)^N s_{b2}+\vartheta^N s_{b1}+\vartheta^N(1-\vartheta)^N s_{bb} \tag{3}$$

The first term in Equation 2 represents the profit from the consumers who do not browse the 'outside' media. The firms can reach these consumers only through contextual advertising. A fraction S_c of these consumers will only consider buying product 1, where:

- s_{11}, s_{1b} and s_{12} are those consumers who have strong preferences for product 1. These consumers will purchase product 1 as long as they learn about it. Due to imperfect

[12]The details are provided in the proof of Lemma 1.

targeting technology, these consumers will learn about product 1 with probability $1 - (1 - \vartheta)^N$ if they only browse content topic 1. They will learn about product 1 with probability $1 - \vartheta^N$ if they only browse content topic 2. When they browse both content topics, they learn about product 1 with probability $1 - \vartheta^N(1 - \vartheta)^N$.

- s_{b2}, s_{b1} and s_{bb} represent those consumers who prefer both products equally. Some of them will only see the ad for product 1 and therefore will only consider product 1. For example, s_{b1} consumers will only browse content topic 1 and a ϑ^N fraction of them will only see product 1's ad. s_{bb} consumers will browse both content topics, but some of them will only see one ad due to imperfect targeting technology.

Equations 2 and 3 reveal two distinct roles of contextual advertising. First, by advertising through relevant content, an advertiser can efficiently reach the consumers who have strong preferences for its product (s_{11}, s_{1b} and s_{12}). Second, contextual advertising creates 'informational differentiation' in addition to any pre-existing product market differentiation. Although some consumers may equally prefer the advertisers' products, they have heterogeneous preferences for content topics. These consumers will consider only a single product if they have only seen ads for that product. The s_{b2}, s_{b1} and s_{bb} terms in Equation 3 correspond to these consumers. In other words, the advertisers can leverage the differentiation in the content market to reduce price competition in the product market.

An interesting tension exists between the two roles of contextual advertising. Consider the case where each firm advertises through the more relevant content topic. For the purpose of *reaching strong preference consumers*, greater alignment between the consumers' product and content preferences is needed. When content is more relevant (ρ is higher), it is more likely that the consumers who have strong preferences for a firm's product will also browse the corresponding content topic. Interestingly, however, *informational differentiation* is created precisely when product preferences are not totally aligned with content preferences. In the case where a consumer is interested in brand i if and only if she is interested in content topic i, all the consumers who are interested in both products will also browse both content topics, and are likely to receive both ads.

To summarize, highly relevant content enhances the advertisers' ability to reach their strong-preference consumers, but reduces their ability to create informational differentiation. Depending on which effect is dominant, the advertiser revenues may either increase or decrease in content relevance ρ. Formally:

Lemma 1 *Suppose each advertiser acquires the content topic that is more relevant to its product. Then advertiser revenues are increasing in ρ iff* $\min(\alpha_p, \alpha_c) < \frac{\vartheta^N - (1-\vartheta)^N}{3\vartheta^N + 3(1-\vartheta)^N - 6\vartheta^N(1-\vartheta)^N}$. *Otherwise, revenues are decreasing in ρ.*

When consumers have heterogeneous product preferences as well as heterogeneous content preferences (α_p and α_c are small), advertiser profits increase in preference alignment. However, when there is a high enough proportion of consumers who are interested in both products or a high proportion interested in both content topics, higher alignment between consumers' product and content preferences actually reduces advertiser revenues.

In the former case, the majority of consumers have strong preferences for one of the products. Since competition is low, creating additional differentiation is less of a concern. Advertiser revenues increase as they become more effective in reaching their strong preference segments. In the latter case, consumer preferences are homogeneous and the need to create informational differentiation dominates. Consequently, advertiser revenues decrease in ρ. The above results are illustrated in Figure 1.

We shall point out that the above results resonate with the findings in Chen et al. (2001), where targeted pricing can either increase or decrease profits due to the trade-off between better price discrimination and intensified competition. Although our results concern targeted informational advertising, the two models share a number of mathematical similarities. We refer interested readers to the appendices of both papers.

The above analysis sheds light on how contextual advertising impacts advertisers' profits. It focuses on the subgame where each firm advertises through the more relevant content topics. We next examine the advertisers' bidding strategies and the equilibrium allocation of

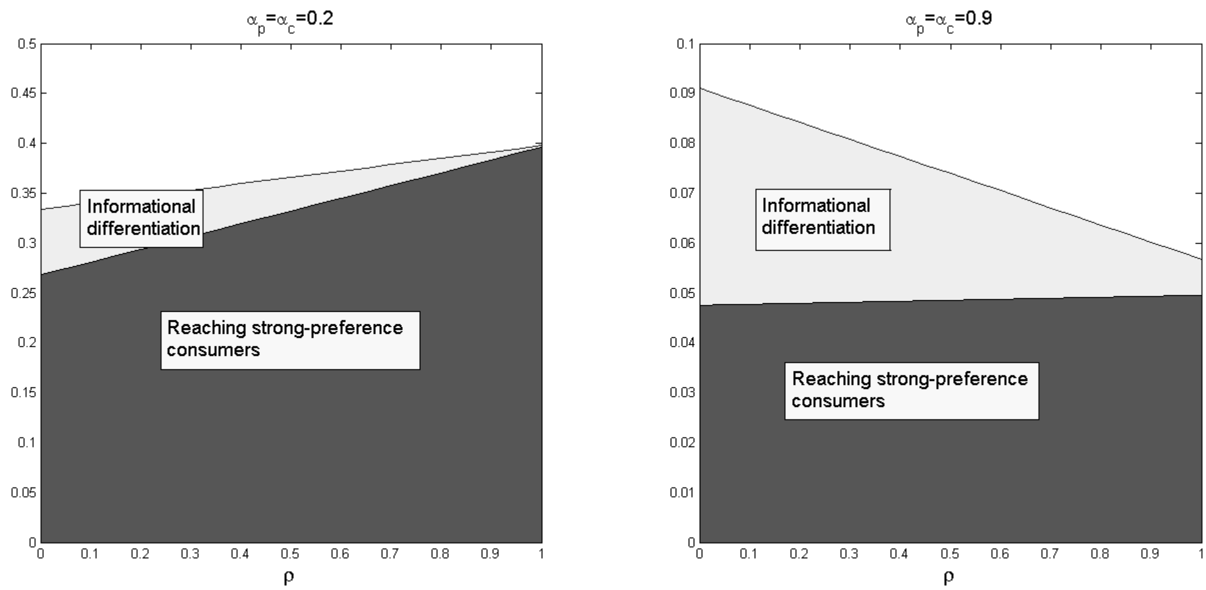

Figure 1: Decomposing Advertiser Revenue

content topics.

Proposition 1 (Content Topic Allocation and Intermediary Revenue) *The equilibrium content topic allocation is described by the following conditions:*

- *Each advertiser gets the more relevant content topic when* $\alpha_p \leq \overline{\alpha_p}(\alpha_c, \rho, \vartheta, N, \gamma)$. *The intermediary's revenue is*

  $$R^d = 2(1-\gamma)(\alpha_p+(1-\vartheta)^N s_{11}+\vartheta^N s_{12}+\vartheta^N(1-\vartheta)^N s_{1b}-(\vartheta^N+(1-\vartheta)^N)s_{b1}-\vartheta^N(1-\vartheta)^N s_{bb})$$

- *When* $\alpha_p > \overline{\alpha_p}(\alpha_c, \rho, \vartheta, N, \gamma)$, *one advertiser gets both content topics in the contextual advertising medium; the other advertiser only advertises through the outside media. The intermediary's revenue is*

  $$R^m = \frac{(1-\gamma)(1+\alpha_p-2\alpha_p\gamma)}{2}.$$

The threshold $\overline{\alpha_p}(\alpha_c, \rho, \vartheta, N, \gamma)$ *is uniquely defined as the solution of*

$$\frac{1+(1+2\gamma)\overline{\alpha_p}+4(1-\vartheta)^N(\overline{s_{11}}-\overline{s_{b1}})+4\vartheta^N(\overline{s_{12}}-\overline{s_{b2}})+4\vartheta^N(1-\vartheta)^N(\overline{s_{1b}}-\overline{s_{bb}})}{2(1-\overline{\alpha_p})} = 1$$

where $\overline{s_{ij}} = s_{ij}(\overline{\alpha_p}, \alpha_c, \rho)$ *are the corresponding segment sizes.*

Proposition 1 establishes two qualitatively different market outcomes. When product market competition is low (i.e., α_p is small), each advertiser wins the more relevant content topic. When product market competition is high (i.e., α_p is large such that s_{b*} is large), however, one firm will get both content topics in equilibrium. We refer to this latter outcome as the 'topic shelving' equilibrium. Since advertising through both content topics prevents the competitor from delivering its product information to some of the consumers, 'topic shelving' can be considered as a mean to reduce product market competition[13]. When competition is sufficiently intense, firm i may have a higher willingness to pay than firm $-i$ for content topic $-i$. 'Topic shelving' takes place as a consequence.

Proposition 1 also reveals two distinct determinants of intermediary revenue. Since the advertising slots are sold via second price auctions, the equilibrium payment to each content topic equals the second highest bid. When each firm advertises through the more relevant content topic, the equilibrium payment equals firm $-i$'s willingness to pay for content topic i. Firm i has two incentives to acquire content topic $-i$. First, firm i can generate additional demand by informing the consumers who browse pages in topic $-i$ but not in topic i. Second, firm i can reduce the pressure for price competition by ensuring that the consumers do not see firm $-i$'s ad. The equilibrium second price bid can thus be decomposed into two parts: the willingness to pay for *additional traffic* and the willingness to pay for *competition reduction*. To see this decomposition formally, consider the extreme case of $\vartheta = 1$. Firm 1's willingness to bid for content topic 2 corresponds to $(1-\gamma)\frac{1+\alpha_p}{2} - (1-\gamma)S_c = (1-\gamma)(s_{12} + s_{b2} + s_{bb})$. The component $s_{12} + s_{b2}$ represents the consumers who are interested in firm 1's product, but do not browse any pages in content topic 1. Therefore, firm 1 can inform these consumers (i.e., generate additional traffic) by advertising through content topic 2. The component s_{bb} corresponds to the consumers who learn about both products when both firms advertise. These consumers have already learned about firm 1's product since they browse pages in content topic 1. However, since they also learn about firm 2's product, the firms have to

[13] Due to the existence of the outside media, a firm can not completely block its competitor out of the market. However, each firm is able to marginally reduce product market competition by preventing its competitor from advertising through the contextual ad intermediary.

compete for these consumers in price. When firm 1 advertises through both content topics, these consumers will only learn about product 1. Thus, $(s_{12} + s_{b2})$ corresponds to firm 1's willingness to pay for additional traffic and s_{bb} corresponds to firm 1's willingness to pay for competition reduction.

The above decomposition has some interesting implications. For example, a firm may be willing to pay *more* for a content topic as it gets *fewer* unique click-throughs from that content topic. When α_p and ϑ are both high, firm i has a higher willingness to pay for content topic $-i$ as α_c gets higher. Higher α_c implies greater overlap between the content topics, and firm i gets less additional traffic from advertising through topic $-i$. However, bidding for topic $-i$ does help prevent the competitor from advertising. In fact, the need for competition reduction is stronger for higher α_c, since overlap in the content topics' audiences imply a larger s_{bb} segment and tougher product market competition.

The intermediary wants to design its targeting service so as to maximize the second price bids. Depending on the consumers' product and content preferences, the intermediary may pursue distinct strategies that focus on either maximizing the advertisers' willingness to pay for additional traffic or their willingness to pay for competition reduction. We elaborate on these results in Section 5.

The results on 'topic shelving' resonate with the empirical observations presented in Shin (2009), who observed that in some product categories advertisers will bid for their competitor's brand name, which is arguably the most relevant 'content topic' for that firm. Similarly, we observe a plethora of 'keyword spying' services emerging on the Internet during the recent years [14]. These companies provide their clients with technologies that analyze the ad space bidding behavior of their *product market* competitors. Some keyword spying companies even provide consulting services on what is the best bidding strategy to win these ad spaces from a competitor. There are multiple reasons why an advertiser might be interested in her competitor's content topic. Our analysis suggests that competitive reduction might be one explanation.

[14] See for example keycompete.com, spyfu.com, or keywordspy.com

5 Intermediary Strategy: Optimal Targeting Accuracy

In this section, we take the intermediary's perspective and consider the endogenous choice of ϑ. Recall from Section 3 that ϑ represents the probability that the intermediary delivers the right ad each time a consumer browses a content page. When $\vartheta = 1$, the intermediary correctly determines the content category of every single page and delivers the corresponding ad. When $\vartheta < 1$, the intermediary may incorrectly classify a page in content topic i as one in content topic $-i$. As a consequence, the intermediary may deliver an ad which is not intended for that page. To highlight the link between targeting accuracy and intermediary revenue, we focus on the simple case where the choice of ϑ is costless. Proposition 2 describes this result.

Proposition 2 (Optimal Targeting Technology) *When the choice of ϑ is costless, the intermediary's optimal*[15] *decision can be characterized as follows:*

- *When $\alpha_p \leq \widehat{\alpha_p}(\alpha_c, \rho, N, \gamma)$, the intermediary develops the most accurate targeting technology $\vartheta^* = 1$.*
- *When $\alpha_p > \widehat{\alpha_p}(\alpha_c, \rho, N, \gamma)$, the intermediary provides the least accurate targeting technology $\vartheta^* = \frac{1}{2}$.*

The threshold $\widehat{\alpha_p}(\alpha_c, \rho, N, \gamma)$ is implicitly defined as the solution of

$$\widehat{s_{12}} - \widehat{s_{b1}} - \frac{\widehat{s_{11}}}{2^N} - \frac{\widehat{s_{12}}}{2^N} - \frac{\widehat{s_{1b}}}{4^N} + \frac{2\widehat{s_{b1}}}{2^N} + \frac{\widehat{s_{bb}}}{4^N} = 0$$

where $\widehat{s_{ij}} = s_{ij}(\widehat{\alpha_p}, \alpha_c, \rho)$ are the corresponding segment sizes.

For sufficiently large N, the condition stated in the proposition converges to $s_{21} > s_{b1}$. This condition is more likely to be satisfied when market competition is low and content relevance is

[15] For sufficiently high α_p, topic shelving takes place and the intermediary is indifferent between a range of ϑ. Technically, Proposition 2 covers these cases as well since the intermediary weakly prefers $\vartheta^* = \frac{1}{2}$ or $\vartheta^* = 1$ when shelving takes place. We detail the relationship between intermediary choice and market outcome in Corollary 1.

low. In this case, the intermediary should offer an accurate targeting technology. When competition is high or content relevance is high, the intermediary has incentives to intentionally decrease targeting accuracy.

The above results can be best understood in light of the willingness-to-pay decomposition provided in Section 4. Consider the case where each firm advertises through the more relevant content topic. Recall that the equilibrium second price bids can be decomposed into a willingness to pay for *additional traffic* and a willingness to pay for *competition reduction*. When market competition is low, the intermediary should design its targeting technology to maximize advertisers' willingness to pay for additional traffic. This is achieved by offering accurate targeting, in which case those consumers who only browse content topic 2 are unlikely to receive firm 1's ad. As a consequence, firm 1 has to bid for content topic 2 in order to reach these consumers. This effect is particularly strong when consumers' content and product preferences are not highly aligned, such that many consumers who browse content topic 2 indeed have strong preferences for firm 1's product.

When product market competition is high, the intermediary should focus on increasing the advertisers' willingness to pay for competition reduction. To achieve this goal, the intermediary decreases the targeting accuracy in order to further intensify the competition between the advertisers. When targeting accuracy is low, consumers who only browse one content topic will likely see the 'wrong' ad and learn about both products. This diminishes informational differentiation and intensifies competition, which leads to higher willingness to pay for competition reduction. In fact, when α_p is sufficiently high, the intermediary may intentionally decrease ϑ such that topic shelving takes place in equilibrium. Corollary 1 summarizes the interaction between the intermediary's optimal choice and the equilibrium allocation of content topics.

Corollary 1 *As α_p increases from 0 to 1, we have the following scenarios:*

- *When α_p is small, the intermediary offers accurate targeting technology. Each advertiser advertises through the more relevant content topic.*

- *As α_p increases, the intermediary offers inaccurate targeting technology. Each advertiser advertises through the more relevant content topic.*

- *As α_p further increases, each advertiser will advertise through the more relevant content topic if targeting is accurate. However, the intermediary decreases ϑ such that topic shelving takes place in equilibrium.*

- *When α_p is sufficiently large, topic shelving takes place regardless of the intermediary's choice of ϑ.*

Figure 2 illustrates the results for two sets of parameter values. As expected, when product market competition is sufficiently high, topic shelving takes place regardless of the intermediary's choice. The intermediary is indifferent between different levels of ϑ. On the other extreme, when product market competition is low, each firm advertises through the more relevant topic regardless of the intermediary's choice. The intermediary may offer either accurate or inaccurate targeting depending on parameters, as explained in Proposition 2.

Interestingly, when α_p is an intermediate region, the market outcome depends on the intermediary's choice of ϑ. Each advertiser will get the more relevant content topic when targeting is accurate, but topic shelving will take place when it is inaccurate. We find that in this parameter region the intermediary should indeed decrease targeting accuracy and induce the topic shelving outcome. The intermediary makes higher profits by selling both content topics to the same advertiser.

Finally, we present a comparative statics result on how content relevance (ρ) impacts intermediary profit, given that the intermediary chooses the optimal targeting accuracy. Although ρ is not a decision variable for the intermediary, it is important to know how the equilibrium profits depend on different values of topic relevance. This might influence the intermediary's decision to introduce its services in certain markets depending on whether content and product preferences are aligned. Corollary 2 reveals that the parameter region in which intermediary profit is increasing in ρ coincide with the region in which the intermediary chooses accurate targeting technology $\vartheta^* = 1$.

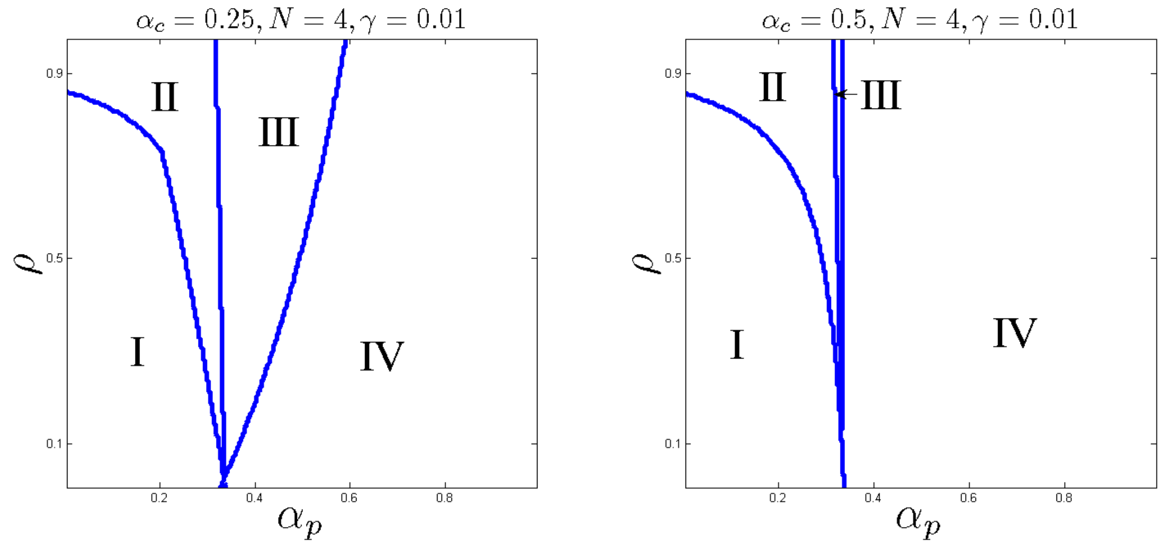

Figure 2: Intermediary choice and content topic allocation: (I) $\vartheta^* = 1$, each firm advertises through the more relevant topic. (II) $\vartheta^* = \frac{1}{2}$, each firm advertises through the more relevant topic. (III) $\vartheta^* = \frac{1}{2}$, topic shelving is induced. (IV) topic shelving takes place regardless of intermediary choice

Corollary 2 *Consider the parameter region in which topic shelving does not take place. When $\vartheta^* = 1$, intermediary profit is decreasing in content relevance. When $\vartheta^* = \frac{1}{2}$, intermediary profit is increasing in content relevance.*[16]

This result should also be understood in light of the willingness to pay decomposition in Section 4. The equilibrium $\vartheta^* = 1$ corresponds to the case where market competition is relatively low such that the intermediary focuses on maximizing the advertisers' willingness to pay for *additional traffic*. Firm 1 has a higher willingness to pay for content topic 2 when s_{12} is larger, such that advertising through topic 2 attracts firm 1's strong preference consumers who would otherwise not see firm 1's ad. This corresponds to a low ρ. When product market competition is high, the intermediary focuses on maximizing the advertisers willingness to pay for competition reduction. As the intermediary chooses $\vartheta^* = \frac{1}{2}$, it follows from Lemma 1 that advertiser competition will further intensify as ρ increases, as long as $\min(\alpha_p, \alpha_c) > 0$. Thus, their willingness to pay is also increasing in ρ.

[16] Note that when ρ changes, the optimal intermediary choice may also change. We thus consider a fixed level of ρ and investigate how the intermediary profits changes locally.

6 Multiple Intermediaries

Throughout the paper we assumed that the intermediary is a monopolist. In this section we examine the basic forces that are at play when firms have the option of advertising through two very similar intermediaries. Since the model becomes substantially more complex by making this change, we make certain simplifying assumptions to obtain a basic understanding of advertiser incentives and equilibrium bids. We do not explore the otherwise interesting questions of how intermediaries compete for publishers and how they choose targeting accuracy under competition.

We assume that there are two identical intermediaries A and B, each of them offering ad space for two topics. They organize four auctions in total; for topics A1, A2, B1, B2, where topic A1 is browsed by the same consumers as B1 and A2 by the same consumers as B2. As before, ρ measures how much the pair of topics (A1,A2) is aligned with product preferences (which will be the same for B1,B2). We assume that there is no outside option to reach consumers ($\gamma = 0$) and targeting accuracy is perfect ($\vartheta = 1$) for both intermediaries. The four auctions are run simultaneously and we employ the same generalized envy-free equilibrium as in the basic model. As before, there are multiple equilibria, but the parameters uniquely determine the pattern of topic allocations between the two advertisers.

Proposition 3 *The equilibrium topic allocations are the following for* $0 < \rho < 1$*:*

1. *One advertiser gets all content topics iff* $1 \geq \alpha_p \geq \overline{\alpha}_p$ *where*

$$\overline{\alpha}_p = \overline{\alpha}_p(\alpha_c, \rho) = \max\left(\frac{1}{3}, 1 - \frac{2}{3} \cdot \frac{\alpha_c}{(\alpha_c + \rho(1/3 - \alpha_c))/\rho}\right)$$

is increasing in ρ *and decreasing in* α_c*.*

2. *Each advertiser gets the most relevant content topic at both intermediaries iff* $\overline{\alpha}_p \geq \alpha_p \geq \underline{\alpha}_p$ *where*

$$\underline{\alpha}_p = \underline{\alpha}_p(\alpha_c, \rho) = \min\left(\frac{1}{3}, 1 - \frac{2}{3} \cdot \frac{1 - \alpha_c}{1 - \alpha_c - \rho(1/3 - \alpha_c))/\rho}\right)$$

is decreasing in ρ *and increasing in* α_c*.*

3. *One advertiser gets both content topics at one intermediary, and the other gets both content topics at the other intermediary iff* $\underline{\alpha}_p \geq \alpha_p \geq 0$.

The proposition identifies three types of equilibria. The first two cases exactly match the two equilibrium types in Proposition 1. In the first case, which we call monopoly shelving, one firm acquires the right to advertise for all topics and becomes a monopolist in the product market. As before, we get this equilibrium, when the market is competitive enough, that is, when α_p and α_c are relatively high with weakly aligned product and content preferences. In the second type of equilibrium, matching our previous results, firms differentiate in which topic they advertise for and they each acquire the more relevant content topic at each intermediary. Interestingly, we find that a third, new type of, equilibrium emerges as the market competitiveness reaches its lowest levels. In this case, each firm gets both the more and less relevant topics to maximize their reach, but at different intermediaries. We call this outcome 'duopoly shelving'. The reason this happens is that firms can take advantage of the extra coverage that the less relevant topic provides as long as the product market is not very competitive (α_p is low). This is especially true when the topics are not much aligned as the less relevant topic might provide valuable consumers. In the extreme case, when topics are independent of product preferences firms are indifferent between this and the second type of outcome.

Comparing our results to the case when there is only one intermediary (Proposition 1), we can nicely match the types of equilibria. Topic shelving happens in the exact same parameter region, the difference is in the duopoly outcomes. In this way our original parameter region where the duopoly outcome was the equilibrium in Proposition 1 will be split between the two types of duopoly outcomes. An important question is how intermediary revenues change in these regions when we introduce a second intermediary. Since there are multiple equilibria and the actual payments and bids differ, we need to select a representative one. We already explained why we focus on the minimum revenue equilibrium that gives advertisers the highest surplus. To pin down the individual intermediary revenues, we assume that since the two intermediaries are symmetric (as are topic allocations with respect to the intermediaries) they make the same profits.

Corollary 3 *In the duopoly shelving case the total minimum revenue of the intermediaries is lower than if there was only one intermediary. Otherwise, the total revenues are the same.*

This result tells us that, with the exception of duopoly shelving, the intermediaries split the same revenue that a single intermediary covering the entire market would make. They do so because since both intermediaries use auctions to allocate their slots, traditional price competition is not present. The only exception is the case of duopoly shelving when advertisers exploit the possibility of each covering the entire market by using the two intermediaries. In this case, when the product market is not very competitive and topics are not much aligned with product preferences, intermediary revenues will be slightly lower.

Although we did not study the competitive choice of targeting accuracy, our results allow us to speculate about the nature of the competition. One could reach the conclusion that when one intermediary offers better targeting accuracy, advertisers would be more aggressive to bid for topics at that intermediary. This may be true to some extent, but our duopoly shelving equilibrium suggests that advertisers do not always have an incentive to fight for the best single topic. Rather, they can get both topics to cover the market, suggesting that the quality competition between intermediaries would be relatively soft, leading to not very high targeting accuracies.

7 Conclusion

We study contextual advertising, a business model that is being embraced by many successful user-generated content websites, social networking communities and affiliated advertising networks as the preferred means to monetize their vast content base. We focus on the idea that advertising targeting helps competing firms reduce their product market price competition, and discuss the implications for advertisers as well as contextual ad intermediaries. Our analysis reveals the importance of content relevance and targeting accuracy, and shows that targeting through more relevant content is not always beneficial for the advertisers. Competing advertisers have incentives to bid for less relevant content topics in order to preempt a competitor from reaching the consumers. The intermediary has an incentive to intentionally lower targeting accuracy at the expense of the advertisers.

Our results have important implications for both advertisers who wish to take advantage of contextual advertising and intermediaries who provide contextual ad services. It is crucial for advertisers to understand the competitive effects of advertising, especially when advertising through imperfectly relevant content. By uncovering the forces that govern the profitability of contextual advertising we offer guidelines to advertisers in how to bid for each slot and when it is better to settle for only the more relevant topic rather than to preemptively acquire advertising slots across all the different topics. Intermediaries also have an important role in this ecosystem. Clearly their choice of targeting technology is influenced by many factors from cost to technical feasibility and we do not model many of these. Nevertheless, intermediaries should understand the competitive implication of targeting accuracy, and choose to focus on either creating additional traffic or reducing informational differentiation, depending on the competitiveness of the product market.

The contextual advertising industry is a fast evolving sector with many exciting developments. Our stylized model takes a first step towards understanding this phenomenon, leaving many interesting questions to future research. First, many contextual advertising networks, including AdSense, rely on independent publishers to provide their content base. The intermediary operates a two-sided market where both advertisers and content publishers participate. Many Adsense publishers – particularly bloggers and small websites – are constantly tuning their content in order to improve their Adsense income. The potential conflict of interests between content publishers and intermediary is an interesting question for future research. In Section 1 of the Online Appendix, we provide some initial analysis about strategic publishers. We consider a case where the intermediary shares ad revenue with independent content publishers, and the publishers can optimize their content in order to maximize ad income. The results describe the optimal content structure the publishers would choose. We show that the intermediary is more likely to offer accurate targeting technology in the presence of strategic publishers. Second, despite the immense popularity of the AdSense-type content analysis algorithm, many alternative means are being developed to match ads with the relevant online content. For example, the ADSDAQ ad exchange sets up an open market where individual publishers can directly sell the ad space on their websites to advertisers. Such open markets exploit the private information each player holds, and may provide better matching than even the most sophisticated page analysis algorithm. Third, contextual targeting can be complemented by other forms of targeted ad delivery. While Google has become the dominant provider of traditional contextual advertising service, social media sites such as Facebook are offering innovative services that combine contextual targeting, demographic targeting and behavioral targeting. How different targeting methods interact poses an interesting yet complex question. We believe that these issues provide fruitful avenues for future research.

References

Bergemann, Dirk, A. Bonatti. 2011. Targeting in advertising markets: Implications for offline vs. online media. *Rand Journal of Economics* **42** 417–443.

Chen, Y., C. He. 2011. Paid placement: Advertising and search on the Internet. *The Economic Journal* forthcoming.

Chen, Y., C. Narasimhan, Z.J. Zhang. 2001. Individual marketing with imperfect targetability. *Marketing Science* **20**(1) 23–41.

Coughlan, A.T. 1985. Competition and cooperation in marketing channel choice: Theory and application. *Marketing Science* 110–129.

Coughlan, A.T., R. Lal. 1992. Retail pricing: Does channel length matter? *Managerial and Decision Economics* **13**(3) 201–214.

Dong, X., P. Manchanda, P.K. Chintagunta. 2009. Quantifying the benefits of individual-level targeting in the presence of firm strategic behavior. *Journal of Marketing Research* **46**(2) 207–221.

Dukes, A. 2004. The Advertising Market in a Product Oligopoly. *Journal of Industrial Economics* **52**(3) 327–348.

Edelman, B., M. Ostrovsky, M. Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review* 242–259.

Gal-Or, E., A. Dukes. 2003. Minimum differentiation in commercial media markets. *Journal of Economics & Management Strategy* **12**(3) 291–325.

Gal-Or, Esther, Mordechai Gal-Or. 2005. Customized advertising via a common media distributor. *Marketing Science* **24**(2) 241–253.

Gal-Or, Esther, T. Geylani, T. P. Yildrim. 2012. The impact of advertising on media bias. *Journal of Marketing Research* **49**(1) 92–99.

Grossman, G.M., C. Shapiro. 1984. Informative advertising with differentiated products. *The Review of Economic Studies* **51**(1) 63–81.

Iyer, G., D. Soberman, M. Villas-Boas. 2005. The targeting of advertising. *Marketing Science* **24**(3) 461.

Katona, Zsolt, Miklos Sarvary. 2010. The race for sponsored links: Bidding patterns for search advertising. *Marketing Science* **29**(2) 199–215.

Narasimhan, C. 1988. Competitive promotional strategies. *Journal of Business* **61**(4) 427–449.

Nelson, P. 1974. Advertising as information. *Journal of Political economy* **82**(4).

Sayedi, Amin, Kinshuk Jerath, Kannan Srinivasan. 2011. Competitive Poaching in Sponsored Search Advertising and Strategic Impact on Traditional Advertising. *working paper* .

Shin, W. 2009. The Company that You Keep: When to Buy a Competitors Keyword. *working paper* .

Varian, H.R. 1980. A model of sales. *The American Economic Review* **70**(4) 651–659.

Varian, H.R. 2007. Position auctions. *International Journal of Industrial Organization* **25**(6) 1163–1178.

Appendix

Proof of Lemma 1: Consider the case where each firm advertises through one content topic. The consumers will either consider one of the products or consider both products depending on their product preferences and which ads they receive. Assume S consumers will only consider one product while R consumers will consider both firms' product. From Narasimhan (1988), the pricing equilibrium has both firms playing mixed strategies in the interval $\left[\frac{S}{R+S}, 1\right]$. The equilibrium profits are S for each firm. Let $S = S_c + S_o$ where S_c refers to the consumers who learn about the product through contextual advertising and S_o refers to the consumers who learn about the products through the outside media. We determine the size of S_c first by considering each of the s_{ij} segment in turn:

- s_{11}: these consumers will consider only product 1 as long as they see ad 1 at least once. Since each time a consumer browses a page in content topic 1 there is a $1-\vartheta$ probability that she will see the wrong ad (for product 2), the probability that a consumer from s_{11} will miss ad 1 is $(1-\vartheta)^N$. Thus $(1-(1-\vartheta)^N)s_{11}$ consumers will see ad 1 and only consider product 1.

- s_{1b}: these consumers will consider only product 1 as long as they see ad 1 at least once. Each consumer from this segment browses N pages in content topic 1 and N pages in content topic 2. The probability that she misses ad 1 for $2N$ times is $\vartheta^N(1-\vartheta)^N$. Thus $(1-\vartheta^N(1-\vartheta)^N)s_{1b}$ consumers will see ad 1 and only consider product 1.

- s_{12}: Similarly, $1-\vartheta^N$ of these consumers will see ad 1 at least once and will therefore consider only product 1.

- s_{b1}: These consumers will consider only product 1 only if they never see an ad for product 2. Each time the consumer browses a page in content topic 1 this would happen with probability ϑ. Thus a ϑ^N fraction of s_{b1} will only consider product 1.

- s_{b2}: Similarly, a $(1-\vartheta)^N$ fraction of these consumers will never see ad 2 and consider only product 1.

- s_{bb}: Similarly, a $(1-\vartheta)^N\vartheta^N$ fraction of these consumers will never see ad 2 and consider only product 1. These consumer only see ad 1 while browsing the N content pages in topic 1 and the N pages in topic 2.

- s_{21}, s_{22} and s_{2b} will never consider product 1.

Note that $s_{b1} = s_{b2}$. Summing the above terms, we obtain the stated expression for S_c. Among the consumers who consume the outside media, $\frac{1-\alpha_p}{2}$ of them will only consider product 1.

We obtain the expression for Π_d by plugging in s_{11}, s_{1b}, s_{12}, s_{b2} and s_{bb} as detailed in Section 3. Taking derivative with respect to ρ:

$$\begin{aligned}\frac{d\Pi_d}{\rho} =& \gamma(4\max\{\alpha_p,\alpha_c\} - 3\alpha_c - 3\alpha_p + 2\vartheta - 8\max\{\alpha_p,\alpha_c\}\vartheta + 3\alpha_c\alpha_p + 6\alpha_c\vartheta + 6\alpha_p\vartheta + \\ & 6\max\{\alpha_p,\alpha_c\}\vartheta^2 - 6\alpha_c\vartheta^2 - 6\alpha_p\vartheta^2 - 6\alpha_c\alpha_p\vartheta + 6\alpha_c\alpha_p\vartheta^2 - 1)/4\end{aligned}$$

The profit function is increasing in ρ when $\frac{d\Pi_d}{\rho} > 0$. Solving the inequality gives the condition specified in Lemma 1. $\square$

Proof of Proposition 1: The extended envy-free equilibrium defines a function $p() : \{\{1\},\{2\},\{1,2\}\} \to \mathbb{R}$ that assigns a price to any subset of content topics and a content topic allocation, such that no firm has any incentives to deviate given the current allocation and prices. We always restrict $p^*(\{1,2\}) = p^*(\{1\}) + p^*(\{2\})$. This captures the important institutional detail that most contextual advertising intermediaries do not use combinatoric auctions. Deviation is defined as obtaining a content topic that is not won by this firm in the current allocation, or giving up a content topic that is won by this firm in the current allocation.

First, observe that due to symmetry, there are three possible types of outcomes: (1) One advertiser wins both content topics; (2) each advertiser wins her more relevant content topic; (3) each advertiser wins her less relevant content topic. We rule out case (3) based on the fact that it is always Pareto dominated by case (2). Note that we do not consider the case where one content topic is not sold. Since in a two-bidder two-item auction with positive valuations, both content topics will always be sold in equilibrium. Put differently, when we evaluate the 'envy-free' conditions, we assume that if an advertiser gives up a content topic, that content topic will be automatically acquired by the other advertiser.

The firms' profits when each firm advertises through its more relevant content topic Π_d. When one firm advertises through both content topics, the profits are asymmetric. The firm which advertises through both content topics makes a profit of $\Pi_m = \gamma(s_{11} + s_{12} + s_{1b} +$

$s_{bb} + 2s_{b1}) + (1-\gamma)\frac{1-\alpha_p}{2}$. This can be seen from the fact that all the consumers who browse online content will receive product 1's ad. The firm which advertises only through the outside media makes a profit of $\Pi_0 = \gamma\Pi_m$ (Narasimhan (1988)).

We first provide envy-free conditions for case (1) to be an equilibrium. In equilibrium, each advertiser wins the more relevant content topic. The envy-free conditions entail:

$$\Pi_m - p^*(\{1,2\}) \le \Pi_d - p^*(\{i\}), \text{ and } \Pi_d - p^*(\{i\}) > \Pi_0, \text{ for } i = 1,2 \tag{4}$$

The first inequality in (4) states that each firm does not outbid its competitor for the less relevant content topic, in which case the firm will win both content topics and monopolize the market (making a profit Π_m). The second inequality in (4) states that each firm does not prefer to give up the more relevant content and stays out of the market. Thus, we obtain that $\Pi_m - \Pi_d \le p^*(\{i\}) \le \Pi_d - \Pi_0$, $i = 1,2$.

When case (2) is an equilibrium such that one advertiser wins both content topics:

$$\Pi_m - p^*(\{1,2\}) \ge \Pi_0,\ \Pi_m - p^*(\{1,2\}) \le \Pi_0, \text{ and } p^*(\{i\}) \ge \Pi_d - \Pi_0, \text{ for } i = 1,2 \tag{5}$$

The first inequality in (5) states that firm 1 doesn't give up both content topics and exit the market. The second inequality in (5) states that firm 2 has no incentives to outbid firm 1 and win both content topics. These two conditions, collectively, imply $p^*(\{1,2\}) = \Pi_m - \Pi_0$. In particular, $p^*(\{1\}) = \Pi_m - \Pi_d$ and $p^*(\{2\}) = \Pi_d - \Pi_0$ is an equilibrium. Clearly, case (1) and (2) represent mutually exclusive conditions. Thus, each equilibrium is also unique when it exists. In each case, the intermediary profit equals to $p^*(\{1,2\})$, which corresponds to $2(\Pi_m - \Pi_d)$ if each firm wins one content topic and $\Pi_m - \Pi_0$ if one firm wins both content topics. □

Proof of Proposition 2 and Corollary 1: From Proposition 1, the intermediary profit is $R^d = 2(\Pi_m - \Pi_d)$ when $\Pi_m + \Pi_o < 2\Pi_d$ and $R^m = \Pi_m - \Pi_o$ when $\Pi_m + \Pi_o \ge 2\Pi_d$. Combining the two conditions, we obtain $R = \min(2(\Pi_m - \Pi_d), \Pi_m - \Pi_o)$. Since $\Pi_m - \Pi_o$ does not depend on the choice of ϑ, $\arg\max_\vartheta R = \arg\max_\vartheta(-\Pi_d)$. From Lemma 1, we have

$-\Pi_d|_{\vartheta=\frac{1}{2}} < -\Pi_d|_{\vartheta=1}$ iff $s_{12} - s_{b1} - \frac{s_{11}}{2^N} - \frac{s_{12}}{2^N} - \frac{s_{1b}}{4^N} + \frac{2s_{b1}}{2^N} + \frac{s_{bb}}{4^N} < 0$. Thus the threshold conditions provided in Proposition 2 are necessary for $\vartheta^* = 1$ or $\vartheta^* = \frac{1}{2}$.

To further establish that the threshold condition is sufficient, we need to rule out the possibility of any interior solution $\vartheta^* \in [\frac{1}{2}, 1]$. This can be established in several steps. First, observe that

$$\frac{\partial - \Pi_d}{\partial \vartheta} = N((s_{b1}-s_{11})(1-\vartheta)^{(}N-1)+\vartheta^{(}N-1)(s_{12}-s_{b1})+(2\vartheta-1)\vartheta^{(}N-1)(1-\vartheta)^{(}N-1)(s_{bb}-s_{s1b})).$$

The derivative is always smaller than zero for $\vartheta = \frac{1}{2}$. Three cases are possible:

- $s_{b1}-s_{11} < 0$ and $s_{12}-s_{b1} < 0$. In this case, $(s_{b1}-s_{11})(1-\vartheta)^{(}N-1)+\vartheta^{(}N-1)(s_{12}-s_{b1}) < 0$ for $\vartheta \in [\frac{1}{2}, 1]$.

- $s_{b1} - s_{11} < 0$, $s_{12} - s_{b1} > 0$ and $|s_{b1} - s_{11}| > |s_{12} - s_{b1}|$. In this case, $(s_{b1} - s_{11})(1 - \vartheta)^{(}N - 1) + \vartheta^{(}N - 1)(s_{12} - s_{b1})$ is strictly increasing.

- $s_{b1} - s_{11} > 0$, $s_{12} - s_{b1} < 0$ and $|s_{b1} - s_{11}| < |s_{12} - s_{b1}|$. In this case, $(s_{b1} - s_{11})(1 - \vartheta)^{(}N - 1) + \vartheta^{(}N - 1)(s_{12} - s_{b1})$ is strictly decreasing.

For an interior solution to exist, $\frac{\partial - \Pi_d}{\partial \vartheta}$ has to cross zero at least twice. Since it is a continuous function, the first extremum is a minimum and the second extremum is a maximum. Thus, $\frac{\partial - \Pi_d}{\partial \vartheta}$ has to be (1) first increasing then decreasing for $\vartheta > \frac{1}{2}$. (2) be greater than zero for at least some ϑ. Given the first and third case above, $(2\vartheta - 1)\vartheta^{(}N - 1)(1 - \vartheta)^{(}N - 1)(s_{bb} - s_{1b})$ has to be first increasing then decreasing in ϑ. This implies $s_{bb} > s_{s1b}$. Since the coefficient of $(2\vartheta - 1)\vartheta^{(}N - 1)(1 - \vartheta)^{(}N - 1)$ is positive, we can establish the following upper bound for $\frac{\partial - \Pi_d}{\partial \vartheta}$ in the first and third case:

$$\frac{\partial - \Pi_d}{\partial \vartheta} < \begin{cases} N\vartheta^2(s_{b1} - s_{11} + s_{12} - s_{b1} + s_{bb} - s_{s1b}) & s_{b1} - s_{11} > 0 \\ N(1 - \vartheta)^2(s_{b1} - s_{11} + s_{12} - s_{b1} + s_{bb} - s_{s1b}) & s_{12} - s_{b1} < 0 \end{cases}$$

Thus, a necessary condition for $\frac{\partial - \Pi_d}{\partial \vartheta} > 0$ for at least some ϑ is $s_{b1} - s_{11} + s_{12} - s_{b1} + s_{bb} - s_{1b} > 0$. However, this condition is sufficient for the topic shelving outcome to take place. Therefore, there is no interior solution ϑ^* that does not induce topic shelving. Similarly, in the second case, no shelving implies $s_{b1} - s_{11} + s_{12} - s_{b1} + s_{bb} - s_{1b} < 0$, which implies $(1-\vartheta)^N(s_{b1} - s_{11} + s_{bb} - s_{1b}) < 0$. This is a sufficient condition for the first derivative to be strictly increasing over $\vartheta \in [\frac{1}{2}, 1]$, hence, an interior extremum does not exist. To summarize, we conclude that when keyword shelving does not take place at ϑ^*, ϑ^* must not be an interior solution. □

Proof of Corrolary 2: Using the expressions provided in Proposition 1, it is clearly that intermediary revenue increasing in ρ in the first and decreasing in the second case. □

Proof of Proposition 3: We start by examining the possible outcomes with respect to the allocation of the four keywords. Let Π_m denote the monopolist's profit, that is, when one firm gets all the topics. Let Π_h and Π_l denote the product market profits of the two players respectively if the first player get three topics and the second player gets the remaining one topic. Let Π_d denote the duopoly profits when one player gets $A1, B1$, whereas its competitor gets $A2, B2$. Finally, let $\Pi_d{}'$ denote duopoly profits when one player gets $A1, A2$, its competitor $B1, B2$ or when one player gets $A1, B2$, and its competitor $A2, B1$. To determine the equilibrium topic allocations we need to determine the total profits of the two players. We show that the equilibrium topic allocation will be the one corresponding to the highest of $\Pi_m + 0, \Pi_h + \Pi_l, 2\Pi_d, 2\Pi_d{}'$. That is, a topic allocation is an equilibrium if it maximizes total product market profits. To see this, for example, in the case when $2\Pi_d$ is the highest, recall that an equilibrium is described by envy-free payment function $p^*()$ defined on the set of keywords. Consider a potential deviation from the duopoly outcome where player 1 gets keywords $A1, B1$. For a deviation that would add keyword $A2$ to be non-profitable we have to have $p^*(\{A2\}) \geq \Pi_h - \Pi_d$. For the same deviation not to be profitable from player 2's perspective we need $\Pi_d - \Pi_l \geq p^*(\{A2\})$, leading to $\Pi_d - \Pi_l \geq \Pi_h - \Pi_d$, further yielding $2\Pi_d \geq \Pi_h + \Pi_l$. Repeating the above for any possible deviation, we get that total profits have to be maximized in equilibrium. To determine the allocation that maximizes total profits, let us rewrite the profits, using the s_{ij} formulas. Recall that $\Pi_m = s_{11} + s_{1b} + s_{12} + 2s_{b1} + s_{bb}$ and $\Pi_d = s_{11} + s_{1b} + s_{b1}$. Furthermore, one can derive $\Pi_h = s_{11} + s_{1b} + s_{12} + s_{b1}$, $\Pi_l = s_{11} + s_{1b}$, and $\Pi_d{}' = s_{11} + s_{1b} + s_{12}$. Note that $\Pi_h + \Pi_l = \frac{2\Pi_d + 2\Pi_d{}'}{2}$, that is, the total profits of the three-one allocation is always between the two duopoly outcome, therefore it cannot be the maximum. To examine the remaining three possibilities - whether Π_m, $2\Pi_d$, or $2\Pi_d{}'$ is the maximal, we need to plug in the s_{ij} expressions. Straightforward calculations give the stated results. □

Proof of Corrolary 3: Similarly to the proof of Proposition 1, we can give a lower bound on the sum of prices that will be paid for each topic in equilibrium. As before in the monopoly case this will be Π_m, but here it is the total minimum revenue for all four (not only two) topics. In the basic duopoly case, the total payment for $A1$ plus $B1$ will be at least $\Pi_m - \Pi_d$, which is the same as for $A2$ plus $B2$. Following the exact same logic in the duopoly shelving case the total for $A1$ plus $A2$, as well as for $B1$ plus $B2$ will be $\Pi_m - \Pi_d{}'$. Although this way we can determine the minimum total revenue for a combination of keywords, there are multiple solution in terms of the distribution between the individual keywords. With the exception of the last, duopoly shelving case, we need to rely on the assumption that the two intermediaries have identical revenues due to their symmetric setup. Therefore, the minimum individual symmetric equilibrium revenues will be $R^m = \frac{\Pi_m}{2}$, $R^d = \Pi_m - \Pi_d$, and $R^{d'} = \Pi_m - \Pi_d{}'$. Plugging in the formulas for s_{ij}, it is clear that in the first two cases revenues are the same as in Proposition 1, whereas in the last case the revenues are weakly lower as this configuration is an equilibrium only if $\Pi_d{}' \geq \Pi_d$. □