MARKETING SCIENCE
CELEBRATING 30 YEARS
Vol. 30, No. 2, March–April 2011, pp. 249–273
ISSN 0732-2399 | EISSN 1526-548X | 11 | 3002 | 0249
DOI 10.1287/mksc.1100.0609

Hybrid Advertising Auctions

Yi Zhu
Marshall School of Business, University of Southern California, Los Angeles, California 90089, zhuy@usc.edu

Kenneth C. Wilbur
Fuqua School of Business, Duke University, Durham, North Carolina 27708, kenneth.wilbur@duke.edu

Facebook and Google offer hybrid advertising auctions that allow advertisers to bid on a per-impression or a per-click basis for the same advertising space. This paper studies the properties of equilibrium and considers how to increase efficiency in this new auction format. Rational expectations require the publisher to consider past bid types to prevent revenue losses to strategic advertiser behavior. The equilibrium results contradict publisher statements and suggest that, conditional on setting rational expectations, publishers should consider offering multiple bid types to advertisers. For a special case of the model, we provide a payment scheme that achieves the socially optimal allocation of advertisers to slots and maximizes publisher revenues within the class of socially optimal payment schemes. When this special case does not hold, no payment scheme will always achieve the social optimum.

Key words: advertising; auctions; Internet marketing; search advertising
History: Received: October 23, 2009; accepted: September 4, 2010; Eric Bradlow served as the editor-in-chief and Miklos Sarvary served as associate editor for this article. Published online in *Articles in Advance* January 21, 2011.

1. Introduction

Auctions are the dominant sales mechanism to allocate online advertising space. The ascendance of Internet auctions has been matched by a growing importance in the literature (Chen and He 2006; Goldfarb and Tucker 2011b; Katona and Sarvary 2008; Yao and Mela 2011). One type of auction is the cost-per-thousand impressions (CPM) auction in which advertisers bid for impressions and pay each time their ad is displayed on a Web page. CPM ad pricing is dominant in the market for Internet display advertising. Another type of auction is the cost-per-click (CPC) auction in which advertisers bid for clicks and pay only when their ad is clicked. CPC ad pricing is dominant in the market for Internet search advertising.

This paper analyzes the hybrid advertising auction, in which each advertiser must choose whether to use CPC bidding or CPM bidding. Bids of both types compete for the same advertising space.

Two major websites, Facebook and Google, currently use hybrid auctions. In August 2010, Facebook was the most visited site on the Internet: 500 million people used it for 46 minutes per day, on average, with half of the users logging in every day (Facebook.com 2010). It was believed that Facebook earned about $700 million in advertising revenues in 2009 (Eldon 2010). Figure 1 shows that Facebook advertisers are required to choose a CPC or CPM option to bid for ad space on the site. Google uses the hybrid auction to allocate ad space on its content network, which generated $7.2 billion in 2008, 30.5% of the company's advertising revenues (Google 2010). However, Google does not offer CPM bidding for ads displayed next to its search results. In principle, *any* seller of online advertising could use a hybrid auction to sell advertising space.

This paper has two goals. The first is to understand the properties of equilibrium in this new auction format. The second is to consider how advertising sellers ("publishers") might offer advertisers more efficient mechanisms within the class of hybrid auctions.

The model features brand advertisers and direct response advertisers competing for a set of ad slots sold by a publisher. In each period, each advertiser chooses its bid type, bid level, and whether to maximize click-through rates. The publisher chooses its click-through rate expectations without knowledge of advertisers' types. The results show how the publisher can use its click-through expectation to deter a type of moral hazard on the part of brand advertisers. Without this deterrence, brand advertisers could calibrate high click-through expectations by using CPM bids in conjunction with high click-through rates and then profit by switching to CPC bids with low click-through effort to lower advertising costs. The analysis here is somewhat general compared to much of the literature: it allows for a potentially large number of slots, many bidders, private information, and repeated interactions.

Figure 1 Step 3 of 3 in Facebook's "Create an Ad" Process

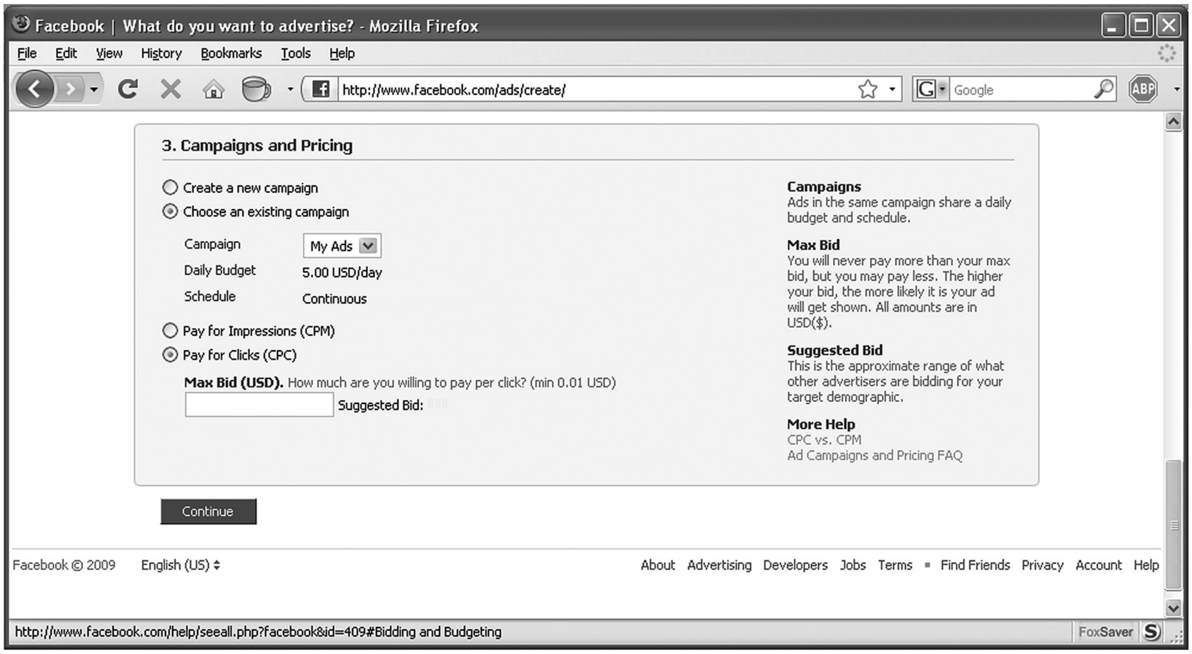

This paper derives several results of managerial interest, but the implications might matter most to those publishers—such as AOL, Microsoft, MySpace, and Yahoo!—that do not currently offer hybrid advertising auctions. These firms do not have the benefit of experience to inform a major change in their business model. Auction revenues depend critically on competition within the auction, so if offering multiple bid types can increase the number of bidders a platform attracts, its effect on publisher revenues may be substantial.

The next section discusses the academic literature to which this paper contributes. Section 3 presents the model's major assumptions. Section 4 outlines the game. Section 5 analyzes its equilibrium. Section 6 considers socially optimal payment schemes within the class of hybrid auction mechanisms. Section 7 concludes with managerial implications and directions for future research.

2. Academic Literature and Contributions

This paper adds to a quickly growing literature on online advertising. The pioneering treatments on equilibria in search advertising auctions are Edelman et al. (2007) and Varian (2007), which independently studied aspects of the auction mechanisms used by Google and Yahoo! known as the *generalized second-price* (GSP) auction. The GSP auction does not have a strictly dominant bidding strategy, but under intuitive refinements, advertisers with higher expected valuations per click occupy higher ad positions in equilibrium. Athey and Ellison (2008), Chen and He (2006), and Xu et al. (2008) studied how advertisers' bids are affected by interadvertiser competition. Recent analytical work has examined such topics as how to incorporate searcher and keyword characteristics into the advertising auction (Even-Dal et al. 2007), how CPC advertising auctions affect advertising's quality-signaling function (Feng and Xie 2007), the interplay between organic and sponsored search links and the publisher's optimal choice of paid links (Katona and Sarvary 2010), how to modify the position auction to account for externalities between advertisers at different positions (Kempe and Mahdian 2007), how to distribute available advertising space among bidding advertisers (Chen et al. 2009), and the effects of "click fraud" on search engine revenues (Wilbur and Zhu 2009).

There is also a rapidly expanding collection of empirical studies of search advertising markets. Ghose and Yang (2009) find that click-through and conversion rates decrease with ad position and that search engines account for both current bid price and prior click-through rates when allocating advertisements to ad slots. Goldfarb and Tucker (2010) find that pricing search advertisements separately across different keywords allows search engines to price discriminate among advertisers. Rutz and Bucklin (2007) show how to borrow information across a large number of keywords and regions to solve the optimal keyword selection and bidding problem. Rutz and Bucklin (2010) show that although generic keywords (e.g., "hotel Los Angeles") are often very expensive, they have spillover effects as consumers tend to begin shopping with a generic search and later use a branded search to purchase. Yao and Mela (2011) use a structural dynamic Bayesian model to analyze data from a product search engine. Among many findings, their study reports that frequent clickers place a greater emphasis on the position of the sponsored advertising link. They also find that a switch from a first-price to a second-price auction yields advertiser bids that are in line with willingness to pay, but this switch has a small impact on search engine revenue. Goldfarb and Tucker (2011a) provide evidence that both advertising targeting and obtrusiveness increase consumers' purchase intentions when used alone but interact negatively when used together.

Several recent papers examine questions similar to those posed here. Goel and Munagala (2009) propose a mechanism in which the publisher requires each advertiser to enter both a CPM *and* CPC bid. The extra information transferred by the advertisers to the publisher raises the publisher's revenue by allowing it to construct more efficient rankings. The fundamental difference in the Goel and Munagala paper is that they propose a new hybrid advertising auction mechanism that requires advertisers to enter two bids, whereas this paper analyzes single-bid auction mechanisms already in use.

Whereas the focus here is on hybrid advertising auctions with CPC and CPM bidding, a third bid type is the cost-per-action (CPA) model in which advertisers pay per purchase or lead. CPA bidding is used less frequently than CPC or CPM bidding (Nazerzadeh et al. 2008). The analysis below can be reinterpreted as a CPM/CPA hybrid advertising auction if one assumes that advertisers are choosing their conversion rate rather than their click-through rate. Edelman and Lee (2008) independently analyze a CPC/CPA hybrid advertising auction. They focus on characterizing equilibrium bids under intuitive refinements of the model and show that the publisher is weakly better off when it offers multiple types of bids to advertisers. Hu et al. (2010) show that CPC and CPA pricing models may conflict because of unobservable, noncontractible effort and adverse selection between brand and direct response advertisers. Hu (2004) uses a contract theory approach to show that performance-based pricing can align publisher and advertiser incentives when complete contracts are infeasible. Agarwal et al. (2009) describe a number of counterintuitive features of CPA auctions in contrast with CPC auctions. Jerath et al. (2010) study how a platform's choice of either a pure CPC or a pure CPM auction influences auction competition and resulting clicks. This work is also related to the literature on pay-per-lead and pay-per-conversion pricing in affiliate marketing; see, e.g., Libai et al. (2003).

The analysis in this paper is also related to signaling models. These were introduced as a possible resolution to inefficiency induced by asymmetric information. For example, in the classic labor market example of Spence (1973), a firm is willing to pay a higher wage to a good employee than a bad employee. However, the firm cannot identify employee type prior to hiring in the absence of a signal. The good employee has a lower cost of completing education, so she engages in additional years of schooling to signal her type to the employer. This general modeling framework has been used to study such marketing questions as the role of demand signaling in distribution channels (Chu 1992, Desai 2000), how price advertising impacts consumers' store price expectations (Simester 1995, Anderson and Simester 1998), and how uninformative advertising and money-back guarantees influence product quality perceptions (Mayzlin and Shin 2009, Moorthy and Srinivasan 1995).

The asymmetric information in our setting lies in the purpose of a new advertising campaign. The advertiser knows whether it is intended for branding or direct response purposes, but the publisher does not. Despite the presence of asymmetric information, there are two main differences between the model analyzed here and a typical signaling model. First, there is no assumption of a costly signal sent from the advertiser to the publisher. The publisher is able to mitigate all harmful effects of asymmetric information without requiring the advertiser to invest resources in a signal. Second, the problem of asymmetric information considered here may persist in many periods, whereas most signaling models assume static games. For example, the inefficiency in the example above may be mostly resolved if the firm hires the employee for a probationary period, which is long enough to determine her type. These two points highlight the differences between the current analysis and a standard signaling model.

This paper's primary contributions are to make a first statement about equilibrium strategies in hybrid advertising auctions and to develop an understanding of how to reach socially efficient outcomes in these auctions. The analysis differs from most of the literature in several key assumptions. It considers advertiser competition in the type of bid as well as bid level, and it allows for advertiser heterogeneity in payoff function as well as reservation price. It does so under the realistic assumptions of repeated interactions and private information about advertisers' types, profits, and click-through rates, whereas most of the literature considers static models of perfect information.

3. Key Assumptions and Empirical Support

The model makes several key assumptions that are motivated by real-world behavior. This section describes those assumptions, their support, and their impact on the results.

3.1. Advertiser Types and Revenues

There are two types of advertising campaigns: direct response and brand-focused. Direct response advertising seeks to stimulate immediate action, such as online purchases. Examples include large online merchants, such as Amazon.com and eBay.com, purchasing targeted ads related to the brands and products they sell. Many off-line businesses, such as mortgage brokers or law firms, engage in direct response advertising to generate registrations or sales leads. The ascendancy of the Google AdWords platform is generally attributed to its usefulness to direct response advertisers.

Brand advertising seeks to influence consumers' product perceptions by increasing awareness or influencing consumer attitudes. It is often employed to alter consumer decisions made in off-line environments, such as retail stores. Examples of brand advertising may include consumer package goods brands such as Scott paper towels or a luxury brand such as Gucci, seeking to influence consumers' brand associations.

Publishers typically do not know the purpose of individual campaigns. Online advertising is typically allocated to large numbers of advertisers via self-service automated processes. It would likely be expensive and perhaps impossible to verify manually whether each advertiser is a brand or direct response advertiser. Even if the publisher attempts manual verification of advertiser types (e.g., having a human employee visit each advertiser's website), it may not have enough information to discern those types perfectly, because a single advertiser's goals may vary from campaign to campaign. For example, Coca-Cola may seek to shape brand associations in one campaign; in another, it may seek to increase online enrollments in its "Coke Rewards" points loyalty program. Publisher uncertainty about campaign purpose is the source of asymmetric information in the model.

To formalize these assumptions, we model a set of $i = 1, \dots, N$ risk-neutral advertisers bidding for $k = 1, \dots, K \leq N$ ads offered by a publisher in each of $t = 0, \dots, T$ time periods.[1] There exist a set of brand advertisers B and a set of direct response advertisers D. Type B's payoff depends on exposures, whereas type D's depends on clicks. Brand advertiser i's profit per exposure is $r_{Bi} \sim F_B$, where F_B is a cumulative distribution function defined on the interval $(0, \infty)$. Direct response advertiser i's profit per click is $r_{Di} \sim F_D$, where F_D is a cumulative distribution function defined on the interval $(0, \infty)$.[2]

3.2. Advertiser Costless Effort and Click-Through Rates

Advertisers may influence click-through rates. As a very simple example, they can choose to encourage consumers to "click here." Another possible strategy is to include a "hard sell" in the ad, which might be effective in shaping off-line behavior but might discourage the consumer from clicking the ad. A third option is to alter the frequency with which new ads are introduced, which, in turn, may influence the likelihood of consumer clicks.

Many empirical studies support this assumption. Krishnamurthy (2000) suggests that the primary factors determining consumer response to banner ads are color, interactivity, and animation. Lohtia et al. (2003) confirm that these factors, along with emotion, influenced consumers' ad response in a field study of 8,725 banner advertisements in both business-to-business and business-to-customer settings. Robinson et al. (2007; see also many references therein) find that increasing the number of words in a banner ad from fewer than 6 to more than 15, holding other factors constant, can increase the click-through rate by more than 100%. Chandon et al. (2003) find that advertisement size, animation, and phraseology (e.g., "click here" or "online only") significantly influenced click-through rates; see also Baltas (2003). Yaveroglu and Donthu (2008) find that ad repetition has a significant effect on consumers' intent to click. Ghose and Yang (2009) show that click-through rates depend on whether an ad contains retailer or brand information. A full review of this literature is beyond the scope of this paper, but there is broad agreement that advertisement content influences consumer response to online advertising.

In addition to academic work, many sellers of online advertising offer tips on how to design ads to maximize click-through rates. They offer tools that facilitate experimentation to see which ads generate the highest click-through rates. An advertiser willing to expend effort to maximize its click-through rates could choose to employ strategies from a wide

[1] Our assumption that $K \leq N$ is in keeping with the literature and rules out the case of multiple-slot purchases by a single advertiser, greatly simplifying the analysis. In practice, the publisher controls K and therefore can set it equal to N when the number of bidders is less than the number of available ad slots.

[2] We have also solved models where each advertiser is characterized by a value for an ad exposure and a value for an ad click. The results in §5 go through mostly unchanged, but the case of discrete advertiser types is simpler.

range of studies, seller-generated tools, and its own experience.[3]

Could an online advertisement that generates a low click-through rate be profitable to a brand advertiser? Several studies support the idea that online advertisements with low click-through rates can be effective in building brands. Drèze and Husherr (2003) show that despite high rates of "ad blindness" (consumers' tendency to avoid focusing on the parts of Web pages where ads appear), consumers exposed to banner ads exhibit higher rates of aided and unaided brand recall regardless of whether they clicked on the ads. Danaher and Mullarkey (2003) demonstrate that time spent viewing a Web page increased the likelihood that a consumer would recall a brand whose banner ad appeared on that Web page.

Advertiser i's click-through rate is modeled as

$$\gamma_{it} = \underline{\gamma}_i + \eta_i x_{it}, \tag{1}$$

where $\underline{\gamma}_i \in (0, 1)$ is advertiser i's baseline click-through rate, which is determined by such exogenous factors as brand recognition and the match between the advertiser's Web page and the search term. Dummy variable x_{it} indicates whether advertiser i exerts costless effort to maximize its click-through rate in period t, and $\eta_i \in (0, 1 - \underline{\gamma}_i)$ is the advertiser-specific productivity of exerting costless effort.

The advertiser chooses whether to exert costless effort ($x_{it} = 1$) or not ($x_{it} = 0$). It is possible to understand this distinction by looking at an advertisement and trying to discern its purpose. If the ad copy takes the central route to persuasion (Tellis 2004), it is more likely that the advertiser is trying to generate clicks. Clicking is a conscious behavior, so rational arguments (like "free") will likely be more effective in generating this conscious behavior. If, on the other hand, the ad copy takes the peripheral route to persuasion, one might expect it to be a low-effort advertisement. An ad that does not try to engage the rational mind is less likely to generate a rational response like a click. This ad still may be of value to a brand advertiser by influencing consumers' latent attitudes and associations. That many online ads appear next to articles, blog posts, or social network content, often in places to which consumers do not pay conscious attention, suggests that quite a bit of display advertising may work by avoiding the central route to persuasion.

One could instead make x_{it} continuous or nonlinear. This would not change the primary results because costless effort will always lead advertisers to maximize or minimize effort within the range of feasible values. If effort is binary but costly, the results, again, will be mostly unchanged. Because advertising effort is a fixed cost, changing this cost only affects investments by firms whose advertising profits are relatively small, similar to changing the fixed cost of production in a standard model of oligopolistic price or quantity competition.

One could also consider an exposure rate variable to allow advertisers to invest effort to increase the probability a consumer is exposed to their ad. The main reason to leave this out is that although clicks can be easily tracked, currently consumer exposures can only be measured with eye-tracking technology. Such technology is not widely deployed, so advertisers' investments in exposure probabilities have limited effect on advertising costs or publisher revenues.

Note that all click-through-related variables are advertiser-specific. It is natural to think that direct response advertisers would likely have higher baseline click-through rates and higher returns to effort. No such distinction among advertiser types is made here, but the model is fully general and can easily be made more specific to accommodate such an assumption.

3.3. Publisher Information

We assume that the publisher knows $\underline{\gamma}_i$ and η_i. In general, the publisher likely has good information about click-through because it has data on many advertisers and campaigns. Each advertiser, on the other hand, typically only has information from its own past campaigns. The results presented in this paper do not require any advertiser knowledge of $\underline{\gamma}_i$ or η_i.

Advertisers, however, determine whether they exert costless effort to maximize click-through rates. Because this may vary from ad to ad, the publisher must anticipate effort levels. γ_{it}^E denotes the publisher's expectation of advertiser i's click-through rate in period t.

A publisher may not observe an advertiser's effort before the ad runs, but it may observe clicks after the ad starts running. It is likely that consumers are heterogeneous and that consumer response to ads may be stochastic. Therefore it would take the publisher some period of time both to deduce the signal from noise and to learn each advertiser's effort level with some predetermined degree of precision. The minimum amount of time needed to measure all advertisers' costless effort levels at some confidence level is the duration of one "period" in the repeated game. If click-through rates are sufficiently low, or if they are

[3] A recent study indicates that the percentage of consumers who click on at least one ad in a month fell from 32% in July 2007 to 16% in March 2009, and that 67% of all ad clicks come from just 4% of consumers (Loechner 2009). If this "clicking segment" is sufficiently homogeneous, it may become progressively easier for advertisers to choose ad copy to influence clicking probability.

sufficiently noisy, then this duration may be of some considerable length.

There are three ways that publishers may anticipate effort levels. The first is through past experience. A publisher is able to observe when an advertiser alters an ad. Consider the following scenario: advertisers bid and enter a set of ads W in period t. By the end of period t, the publisher has observed the effort level associated with each ad in W. If the *same* set of ads is entered in period $t+1$, the issues described in this paper will not apply in period $t+1$. Therefore the scope of the analysis is any period in which the set of ad creatives has changed from the previous period. Yet this scope may still be considerable in magnitude, because the number of keywords sold is very large, the duration of a single period may be considerable, and the number of ads entered by a single advertiser may be large. If the publisher has imperfect information about an advertiser's ad quality, and if that state of imperfect information is beneficial to the advertiser, the advertiser is likely to alter its ad creatives quite frequently to increase the frequency with which the publisher has imperfect information.[4] This is made even more likely by advertisers' ability to use software to generate ad creatives.

The second way in which a publisher can anticipate advertiser costless effort is to analyze the text or graphic content of the ad algorithmically and predict consumer response based on these "creative" elements. There are several reasons to think this is technologically infeasible. First, some publishers say they do not do this. For example, Google states in Figure 2 that its "AdWords system treats an edited ad like it's brand new and has no performance history." Second, if it were feasible, the publisher could suggest the click-maximizing ad text to any individual advertiser. No publisher currently offers this, though doing so would likely increase publisher revenues. If a publisher were able to produce this technology, the asymmetric information that motivates this paper would be fully resolved. However, as shown below, it is possible to fully mitigate any negative effects of strategic behavior of this asymmetric information even in the absence of this technology.

The third way to anticipate advertisers' choices of costless effort levels is to use economic reasoning, that is, to consider advertisers' equilibrium strategies. This is basically costless to the publisher. Section 5.1 proposes a way to do this.

[4] A counterargument would hold that the publisher could mostly deter this behavior by limiting the number of new ads an advertiser may use. The advertiser could circumvent this policy by opening multiple accounts. As long as advertiser identity is imperfectly observable, there is a threat of frequent ad introductions.

Figure 2 Google Help Page on Quality Scores for New Ads

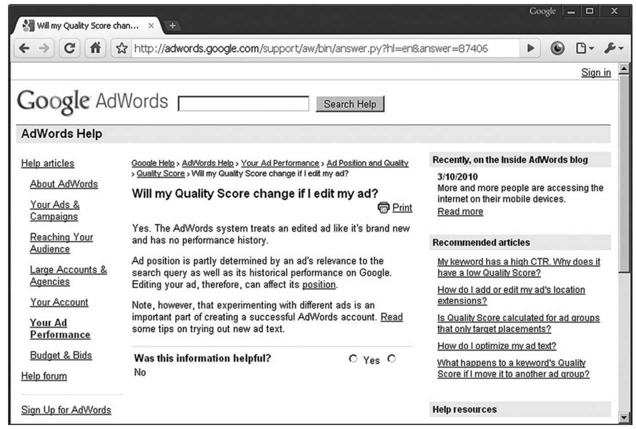

3.4. Position Effects

Following Katona and Sarvary (2010), an ad appearing in slot k has a position-dependent click-through multiplier X_k with $1 \geq X_1 > X_2 > \cdots > X_K > 0$. Position-dependent exposure multipliers $1 \geq Y_1 > Y_2 > \cdots > Y_K > 0$ allow an ad's position on the page to determine the likelihood that it is seen.

It is necessary to define page views, exposures, and impressions. A *page view* occurs when a consumer loads a Web page containing a set of ads. An *exposure* or *impression* occurs if the consumer processes an advertisement. In other words, a page view is a *potential* exposure or a *potential* impression. As discussed above, individual exposures are unobserved without eye-tracking technology. The measure of available consumer page views is normalized to 1 without loss of generality.

3.5. Hybrid Advertising Auction Rules

Assumptions about how the hybrid advertising auction works are based on (1) logic, (2) previous academic literature that has examined related auctions, and (3) public statements by companies that currently offer hybrid advertising auctions.

Assumption 1. *The publisher allows each advertiser i to enter either a CPC bid b_{it}^c or a CPM bid b_{it}^m in each period t and assigns advertisers to slots in order of total expected advertiser willingness to pay.*

If advertisers were charged their bids, an advertiser i with a CPM bid b_{it}^m would pay $Y_k b_{it}^m$ for slot k. If i instead entered a CPC bid b_{it}^c, its total expected payment for slot k would be $X_k \gamma_{it}^E b_{it}^c$. Assumption 1 implies that if advertiser i enters a CPM bid and advertiser j enters a CPC bid, i will be allocated to slot k and j to a less desirable slot if $Y_k b_{it}^m > X_k \gamma_{jt}^E b_{jt}^c$.

Assumption 1 is consistent both with Google's statement in Figure 3 that "neither type of ad [CPC or CPM] has a special advantage over the other," and

Figure 3 **Google Help Page on CPC and CPM Bid Competition**

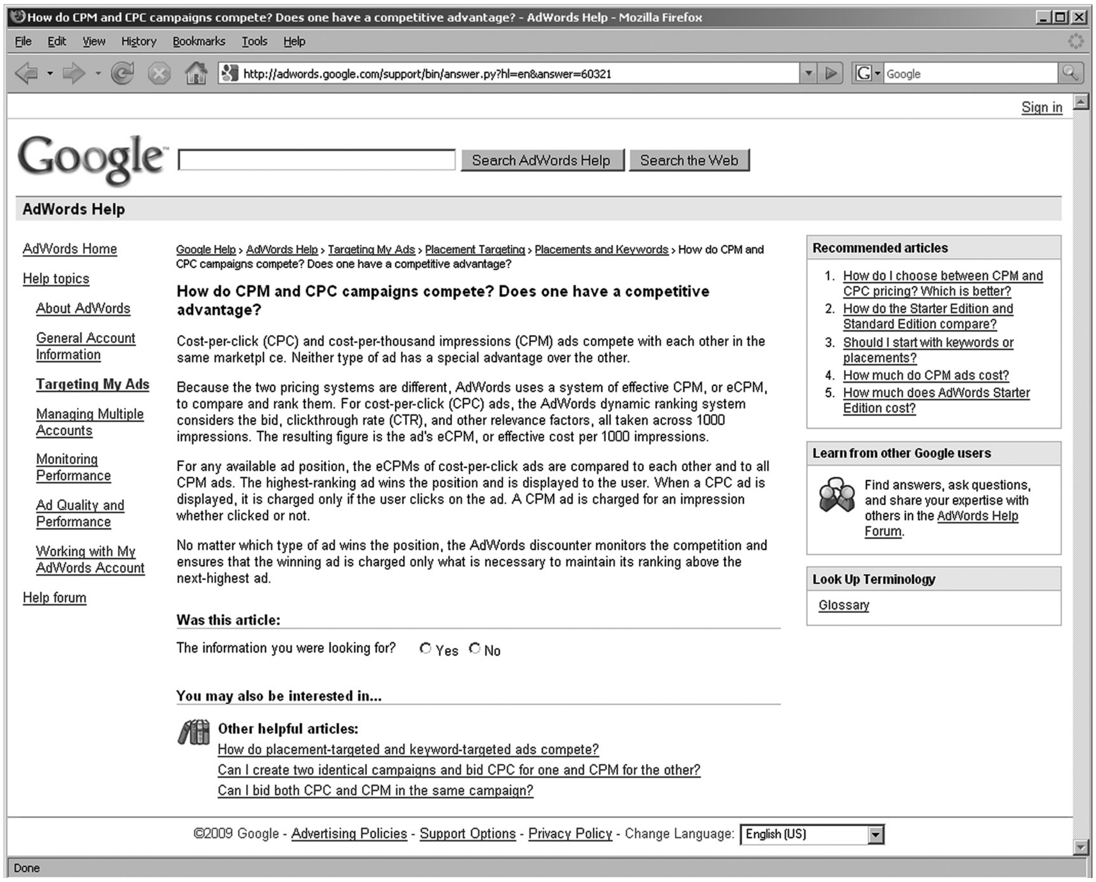
How do CPM and CPC campaigns compete? Does one have a competitive advantage? - AdWords Help - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://adwords.google.com/support/bin/answer.py?hl=en&answer=60321

Sign in

Google Search AdWords Help Search the Web

AdWords Help

AdWords Home
Help topics
About AdWords
General Account Information
Targeting My Ads
Managing Multiple Accounts
Monitoring Performance
Ad Quality and Performance
Working with My AdWords Account
Help forum

Google Help › AdWords Help › Targeting My Ads › Placement Targeting › Placements and Keywords › How do CPM and CPC campaigns compete? Does one have a competitive advantage?

How do CPM and CPC campaigns compete? Does one have a competitive advantage?

Cost-per-click (CPC) and cost-per-thousand impressions (CPM) ads compete with each other in the same marketpl ce. Neither type of ad has a special advantage over the other.

Because the two pricing systems are different, AdWords uses a system of effective CPM, or eCPM, to compare and rank them. For cost-per-click (CPC) ads, the AdWords dynamic ranking system considers the bid, clickthrough rate (CTR), and other relevance factors, all taken across 1000 impressions. The resulting figure is the ad's eCPM, or effective cost per 1000 impressions.

For any available ad position, the eCPMs of cost-per-click ads are compared to each other and to all CPM ads. The highest-ranking ad wins the position and is displayed to the user. When a CPC ad is displayed, it is charged only if the user clicks on the ad. A CPM ad is charged for an impression whether clicked or not.

No matter which type of ad wins the position, the AdWords discounter monitors the competition and ensures that the winning ad is charged only what is necessary to maintain its ranking above the next-highest ad.

Was this article:

The information you were looking for? Yes No

You may also be interested in...

Other helpful articles:
How do placement-targeted and keyword-targeted ads compete?
Can I create two identical campaigns and bid CPC for one and CPM for the other?
Can I bid both CPC and CPM in the same campaign?

Recommended articles
1. How do I choose between CPM and CPC pricing? Which is better?
2. How do the Starter Edition and Standard Edition compare?
3. Should I start with keywords or placements?
4. How much do CPM ads cost?
5. How much does AdWords Starter Edition cost?

Learn from other Google users
Find answers, ask questions, and share your expertise with others in the AdWords Help Forum.

Look Up Terminology
Glossary

©2009 Google - Advertising Policies - Support Options - Privacy Policy - Change Language: English (US)

Done

with Google's ad rank disclosures shown at the bottom of Figure 4. It is also consistent with Facebook's statement in Figure 5 that "for any available ad inventory, Facebook selects the best ad to run based on the cost per click or impression and the ad performance." If the publisher used any other ranking method, it would systematically bias advertisers toward using one type of bid over the other.

Assumption 2. *Each advertiser is charged the minimum amount necessary to keep its place in the ranking.*

Assumption 2 is in line both with the prior literature on CPC auctions (e.g., Edelman et al. 2007; Varian 2007, 2009) and with the common understanding of Google's pure CPC keyword auction. As Google states in Figure 2, "No matter which type of ad [CPC or CPM] wins the position, the AdWords discounter monitors the competition and ensures that the winning ad is charged only what is necessary to maintain its ranking above the next-highest ad."

Assume that advertiser i holds position k and advertiser j holds position $k+1$. Assumptions 1 and 2 imply that if both advertisers entered CPM bids, then i pays $Y_k b^m_{jt}$ total. If both advertisers entered CPC bids, then i pays $\gamma^E_{jt} b^c_{jt}/\gamma^E_{it}$ per click. If advertiser i entered a CPC bid and j entered a CPM bid, then i pays $Y_k b^m_{jt}/X_k \gamma^E_{it}$ per click. If i entered a CPM bid and j entered a CPC bid, then i pays $X_k \gamma^E_{jt} b^c_{jt}$.

To simplify advertiser profit functions in the next section, we write the CPM payment or the expected CPC payment of the advertiser in slot k in period t as

$$C_{kt} = \begin{cases} X_k \gamma^E_{i't} b^c_{i't} & \text{if advertiser } i' \text{ in slot } k+1 \\ & \text{submits a CPC bid} \\ Y_k b^m_{i't} & \text{if advertiser } i' \text{ in slot } k+1 \\ & \text{submits a CPM bid.} \end{cases}$$

We further assume that an advertiser only changes its bid type or effort level if it benefits from doing so. This assumption is consistent with any arbitrarily small nuisance cost of taking an action. It is a tie-breaking rule and can be thought of as "advertiser inertia." Pauwels (2004) finds that firms exhibit inertia in tactical decisions such as pricing and promotions even when inertia reduces profits. The advertiser inertia assumption is weaker than Pauwels' result, because it only presumes inaction when action cannot increase profits. The role of this assumption rules out degenerate expectation functions in the proof of Proposition 1.

Figure 4 **Google Help Page on Quality Scores and Ad Rank Formulas**

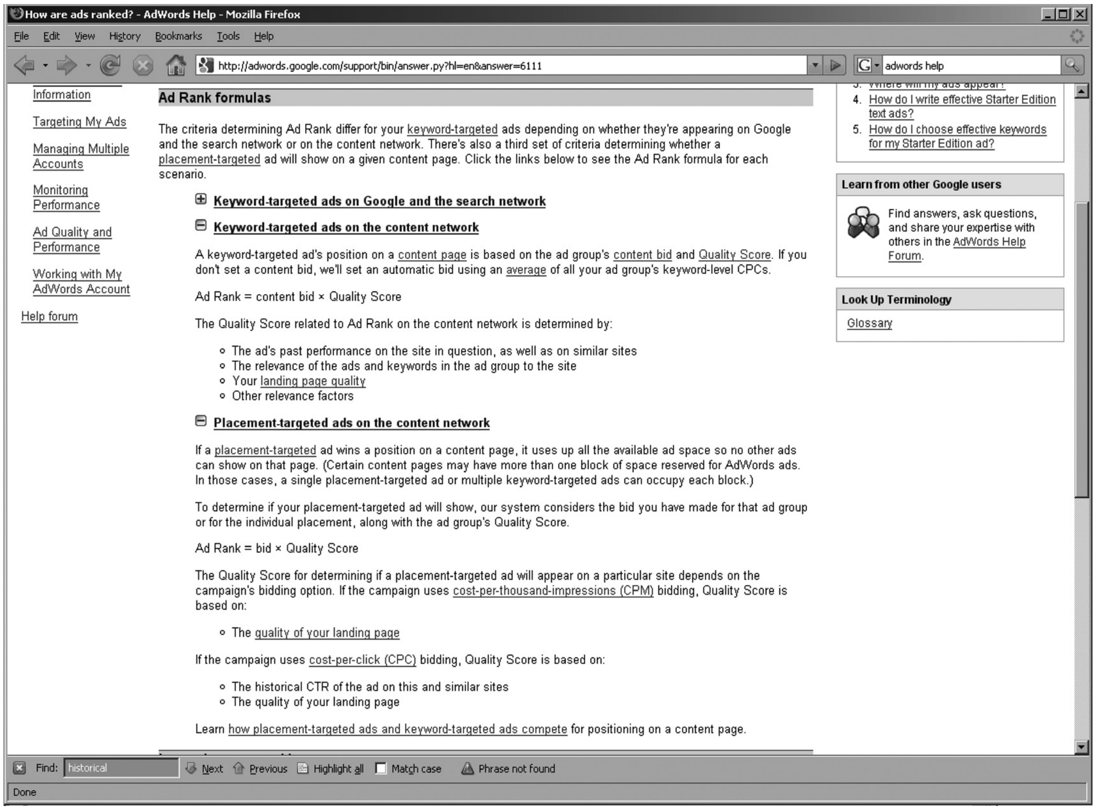
How are ads ranked? - AdWords Help - Mozilla Firefox

http://adwords.google.com/support/bin/answer.py?hl=en&answer=6111

Information
Targeting My Ads
Managing Multiple Accounts
Monitoring Performance
Ad Quality and Performance
Working with My AdWords Account
Help forum

Ad Rank formulas

The criteria determining Ad Rank differ for your keyword-targeted ads depending on whether they're appearing on Google and the search network or on the content network. There's also a third set of criteria determining whether a placement-targeted ad will show on a given content page. Click the links below to see the Ad Rank formula for each scenario.

Keyword-targeted ads on Google and the search network

Keyword-targeted ads on the content network

A keyword-targeted ad's position on a content page is based on the ad group's content bid and Quality Score. If you don't set a content bid, we'll set an automatic bid using an average of all your ad group's keyword-level CPCs.

Ad Rank = content bid × Quality Score

The Quality Score related to Ad Rank on the content network is determined by:

- The ad's past performance on the site in question, as well as on similar sites
- The relevance of the ads and keywords in the ad group to the site
- Your landing page quality
- Other relevance factors

Placement-targeted ads on the content network

If a placement-targeted ad wins a position on a content page, it uses up all the available ad space so no other ads can show on that page. (Certain content pages may have more than one block of space reserved for AdWords ads. In those cases, a single placement-targeted ad or multiple keyword-targeted ads can occupy each block.)

To determine if your placement-targeted ad will show, our system considers the bid you have made for that ad group or for the individual placement, along with the ad group's Quality Score.

Ad Rank = bid × Quality Score

The Quality Score for determining if a placement-targeted ad will appear on a particular site depends on the campaign's bidding option. If the campaign uses cost-per-thousand-impressions (CPM) bidding, Quality Score is based on:

- The quality of your landing page

If the campaign uses cost-per-click (CPC) bidding, Quality Score is based on:

- The historical CTR of the ad on this and similar sites
- The quality of your landing page

Learn how placement-targeted ads and keyword-targeted ads compete for positioning on a content page.

4. How do I write effective Starter Edition text ads?
5. How do I choose effective keywords for my Starter Edition ad?

Learn from other Google users

Find answers, ask questions, and share your expertise with others in the AdWords Help Forum.

Look Up Terminology

Glossary

Figure 5 **Facebook CPC/CPM Help Page**

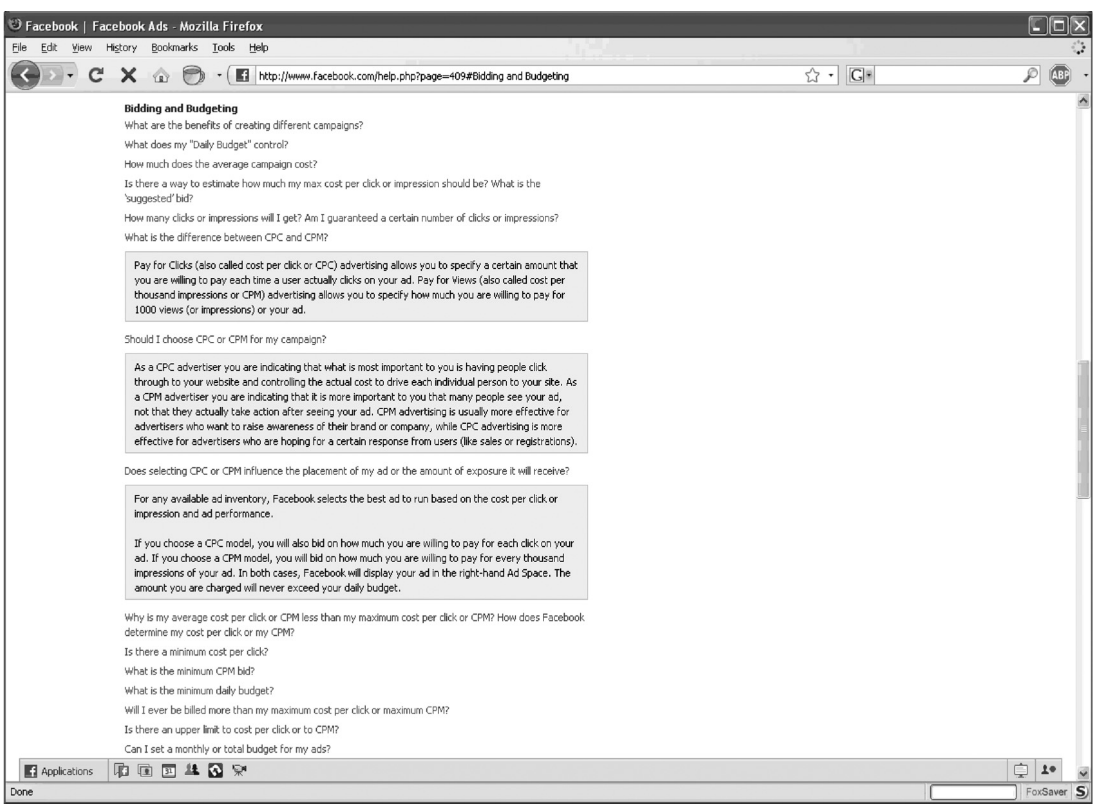
Facebook | Facebook Ads - Mozilla Firefox

http://www.facebook.com/help.php?page=409#Bidding and Budgeting

Bidding and Budgeting

What are the benefits of creating different campaigns?

What does my "Daily Budget" control?

How much does the average campaign cost?

Is there a way to estimate how much my max cost per click or impression should be? What is the 'suggested' bid?

How many clicks or impressions will I get? Am I guaranteed a certain number of clicks or impressions?

What is the difference between CPC and CPM?

Pay for Clicks (also called cost per click or CPC) advertising allows you to specify a certain amount that you are willing to pay each time a user actually clicks on your ad. Pay for Views (also called cost per thousand impressions or CPM) advertising allows you to specify how much you are willing to pay for 1000 views (or impressions) or your ad.

Should I choose CPC or CPM for my campaign?

As a CPC advertiser you are indicating that what is most important to you is having people click through to your website and controlling the actual cost to drive each individual person to your site. As a CPM advertiser you are indicating that it is more important to you that many people see your ad, not that they actually take action after seeing your ad. CPM advertising is usually more effective for advertisers who want to raise awareness of their brand or company, while CPC advertising is more effective for advertisers who are hoping for a certain response from users (like sales or registrations).

Does selecting CPC or CPM influence the placement of my ad or the amount of exposure it will receive?

For any available ad inventory, Facebook selects the best ad to run based on the cost per click or impression and ad performance.

If you choose a CPC model, you will also bid on how much you are willing to pay for each click on your ad. If you choose a CPM model, you will bid on how much you are willing to pay for every thousand impressions of your ad. In both cases, Facebook will display your ad in the right-hand Ad Space. The amount you are charged will never exceed your daily budget.

Why is my average cost per click or CPM less than my maximum cost per click or CPM? How does Facebook determine my cost per click or my CPM?

Is there a minimum cost per click?

What is the minimum CPM bid?

What is the minimum daily budget?

Will I ever be billed more than my maximum cost per click or maximum CPM?

Is there an upper limit to cost per click or to CPM?

Can I set a monthly or total budget for my ads?

This section closes with an observation and an associated assumption. Conditional on an equilibrium assignment, it is possible for the publisher to increase click-through expectations to increase higher advertisers' payments without changing the equilibrium assignment. Consider a simplified example to understand this action. Suppose a seller is offering one item to two bidders with privately held valuations of \$2 and \$1. The seller could use a mechanism such that (1) each bidder enters a bid and (2) after observing the bids, charge the higher bidder one penny less than her bid. This would result in revenues of \$1.99 rather than the second-price auction payment of \$1. In our setting, the publisher could achieve a very similar effect by manipulating click-through expectations to reduce the difference between advertisers' expected willingness to pay. We assume that the publisher does not do this for several reasons. First, it may be constrained by legal contracts or the threat of a class-action lawsuit. Second, it would likely reduce advertiser participation in the auction by reducing advertiser surplus. Third, it may be perceived as "unfair" and damage the publisher's reputation. Fourth, we do not know of any theoretical auction paper that considers strategies of this type.

4. The Dynamic Hybrid Advertising Auction

This section presents remaining model assumptions. Table 1 summarizes all notation.

Table 1 **Summary of All Notation**

Symbol	Definition
$\underline{\gamma}_i$	Minimal click-through rate for advertiser i
x_{it}	$=1$ if advertiser i exerts costless effort to maximize click-through rate in period t; $=0$ otherwise
η_i	Marginal number of clicks produced by advertiser i's costless effort
γ_{it}^E	Publisher's expectation of advertiser i's click-through rate in period t
g_{it}	$=c$ if advertiser i enters a CPC bid in period t
	$=m$ if advertiser i enters a CPM bid in period t
$b_{it}^{g_{it}}$	Advertiser i's bid of type g_{it} in period t
X_k	Slot-dependent click-through multiplier
Y_k	Slot-dependent exposure multiplier
r_{Bi}	Reservation value per click of brand advertiser i
r_{Di}	Reservation value per click of direct response advertiser i
F_B	CDF of brand advertisers' reservation values
F_D	CDF of direct response advertisers' reservation values
R_{ikt}	Total revenues of advertiser i in position k in period t
$\delta_{it}^{g_{it}} C_{kt}$	Total cost of position k to advertiser i in period t
$\pi_{ikt}^{g_{it}}$	One-period profits of advertiser i in position k in period t
β	Discount rate
Π_i^a	Discounted sum of advertiser i's expected profits in all periods
Π^p	Discounted sum of publisher's expected profits in all periods
σ	Possible publisher expectation function
H_{it}	Vector of advertiser i's past bid types up to period t and efforts up to $t-1$
T	Total number of periods

4.1. Profit Functions

For simplicity, advertiser reservation values (r_{Bi} and r_{Di}), baseline click-through rates ($\underline{\gamma}_i$), and returns to costless effort (η_i) do not change over time. Total revenues according to advertiser identity, type, and position are R_{ikt}. Let $g_{it} \in \{c, m\}$ indicate bid type (CPC or CPM), and

$$\delta^{g_{it}} = \begin{cases} 1 & \text{if } g_{it} = m \\ \dfrac{\gamma_{it}}{\gamma_{it}^E} & \text{if } g_{it} = c, \end{cases} \tag{2}$$

so that the one-period profits of advertiser i in slot k in period t can be written as

$$\pi_{ikt}^{g_{it}} = R_{ikt} - \delta_{it}^{g_{it}} C_{kt}. \tag{3}$$

Define N-vectors $g_t = (g_{it})_{i=1,\ldots,N}$, $b_t = (b_{it}^{g_{it}})_{i=1,\ldots,N}$, $\gamma_t^E = (\gamma_{it}^E)_{i=1,\ldots,N}$, $x_t = (x_{it})_{i=1,\ldots,N}$, and $\delta_t = (\delta_{it}^{g_{it}})_{i=1,\ldots,N}$. The discounted sum of expected advertiser profits is

$$\mathrm{E}\,\Pi_i^a = \max_{\{g_{it}\},\{x_{it}\},\{b_{it}^{g_{it}}\}} \sum_{t=0}^{T} \beta^t \mathrm{E}\pi_{ikt}\big(x_{it}, b_{it}^{g_{it}} \mid \gamma_{it}^E, b_{-it}^{g_{-it}}\big), \tag{4}$$

where β is the discount factor and $b_{-it}^{g_{-it}}$ is an $(N-1)$-vector of other advertisers' bids in period t. The publisher's one-period profits are the sum of payments that advertisers make for all occupied slots, so the discounted sum of expected publisher profits is

$$\mathrm{E}\,\Pi^p = \max_{\gamma_t^E} \sum_{t=0}^{T} \sum_{k=1}^{K} \beta^t \mathrm{E}\big[\delta_{it}^{g_{it}} C_{kt} \mid b_t, x_t\big]. \tag{5}$$

4.2. Timing of the Game

This paper analyzes a repeated game with private information because online advertising auctions are repeated with high frequency. Results that hold in a static model might not hold in a more realistic dynamic setting.

Within each period t, three sets of strategic actions are taken: advertisers choose bids (types and amounts), the publisher sets click-through expectations, and advertisers choose costless effort levels. After all actions are taken, the publisher's mechanism assigns advertisers to slots, consumers click ads, advertiser profits are realized, and transfers are made from advertisers to the publisher.

This paper's results require two assumptions regarding timing. The first assumption required for the results below is that the publisher does not observe whether the advertiser has exerted costless effort to maximize click-through rates prior to the assignment of advertisements to slots. This assumption preserves the moral hazard nature of the publisher/advertiser interaction. The full justification for this assumption is in §3.3.

Second, it must be the case that the publisher does not finalize its click-through expectations until after it observes advertisers' bids. This is necessary because the publisher must form expectations about every advertiser that joins the auction, and it only knows which advertisers are in the auction after the advertisers enter their bids. This is feasible for three reasons: (1) the publisher controls the computer system into which bids are entered, (2) the publisher controls the time lag between when bids are entered and when they affect the ad rankings, and (3) publishers' software can form click-through expectations reasonably quickly. If this assumption is violated, the publisher will not be able to use any information about an advertiser's bid to set click-through expectations.

Thus, bids must be entered before click-through expectations are formed, and click-through expectations must be formed prior to effort variables being revealed. Any sequence of actions meeting these requirements will produce the results below. It is perhaps simplest to think of this as a two-stage game within each period t. First, advertisers simultaneously enter bids and choose costless effort levels. The publisher then observes the bids, but not the effort levels, and sets its click-through expectations.

4.3. Equilibrium Concept

The equilibrium concept concludes the specification of the game.

Definition 1. A perfect Bayesian Nash equilibrium is defined by any set of publisher expectations σ function about advertisers' click-through rates $(\gamma_t^E)_{t=0,\ldots,T}$, any set of bids $(b_t)_{t=0,\ldots,T}$, and any set of costless effort levels $(x_t)_{t=0,\ldots,T}$ for which the following conditions hold:

(1) *Incentive compatibility*: the choice sequence of costless effort levels $(x_{i0},\ldots,x_{iT})$ and bids $(b_{i0}^{g_{i1}},\ldots,b_{iT}^{g_{iT}})$ maximize expected profits $\mathrm{E}\Pi_i^a$ for all advertisers $i=1,\ldots,N$.

(2) *Individual rationality*: for any advertiser i who wins slot k in period t, $\mathrm{E}\pi_{ikt}^{g_{it}} \geq 0$.

(3) *Publisher optimality*: the choice of expected click-through rates $(\gamma_{it}^E)_{i=1,\ldots,N,\,t=0,\ldots,T}$ maximize expected profits $\mathrm{E}\,\Pi^p$.

(4) *Consistency of publisher beliefs*: the publisher updates its belief using players' observed actions. For any publisher expected click-through rate function σ,

$$\Pr(\gamma_{it}^E=\gamma_i+\eta_i \mid \sigma, H_{it}) = \frac{\Pr(\gamma_{it}=\gamma_i+\eta_i, H_{it}\mid\sigma)}{\Pr(\gamma_{it}=\gamma_i+\eta_i, H_{it}\mid\sigma)+\Pr(\gamma_{it}=\gamma_i, H_{it}\mid\sigma)} \quad \text{for all } i=1,\ldots,n,$$

and $t=1,\ldots,T$, where H_{it} contains advertiser i's past bids up to and including period t and past click-through rates up to period $t-1$.

The first condition ensures that each advertiser chooses its bid type, bid level, and costless effort level in each period to maximize long-run profits. The second condition is a standard individual rationality constraint ensuring nonnegative profits for winning advertisers. The third condition ensures that the publisher sets rational expected click-through rates. The fourth condition means the publisher belief about advertisers' click-through rate is correct in equilibrium for every period after the start of the campaign; this condition is a standard in perfect Bayesian equilibrium analysis.

5. Equilibrium Analysis

This section characterizes equilibrium. It proves that equilibrium exists and that the publisher has a unique expectation function that prevents advertisers from reducing cost through strategic choice of bid types and effort, but this leads to the counterintuitive result that direct response advertisers always use CPM bids. Finally, it compares the hybrid auction with the GSP auction to show that any GSP equilibrium can be supported in the hybrid auction format.

Equilibrium is determined by the interplay of advertisers' strategies and publisher expectations. We first characterize publisher expectations and then derive equilibrium advertiser choices of bid types and effort under that expectation function. Then, given equilibrium advertiser strategies, Proposition 1 shows that equilibrium exists and the expectation function is unique.

5.1. Publisher Expectations

Publisher expectations are critical to produce the optimal assignment of advertisers to slots. They affect advertisers' costs, which, in turn, determine advertisers' optimal bid types.

The publisher must use its effort expectation function to prevent advertisers from "gaming the system." The threat the publisher faces comes primarily from brand advertisers, who, by definition, do not care about clicks. This risk arises because the publisher must assign advertisers to slots prior to observing their costless effort levels.

The risk is that a brand advertiser may use high effort levels in conjunction with CPM bidding to lead the publisher to expect a high effort level. The CPM bids would ensure that the advertiser does not pay any additional cost in connection with the high effort levels. Then, the brand advertiser could take advantage of the high click-through expectation to gain a cost advantage. It could do this by switching to a CPC bid with a low effort level to reduce its total payment for the same advertising space.

However, the publisher has a tool at hand. Advertisers' bids must precede the assignment of advertisers to slots. Therefore the publisher may base its expectations on the type of bid the advertiser enters. Because profitable effort reversals of the type described above require switching bid types from one period to the next, the publisher has an "early warning system." A publisher belief function that punishes an advertiser for switching bid types can prevent such opportunistic behavior. This gives rise to Bayesian Publisher Expectations, as formalized in Definition 2.

Definition 2. Bayesian Publisher Expectations (BPEs) imply the publisher's click-through expectations are based on past bid types:

$$\gamma_{it}^E = \begin{cases} \underline{\gamma}_i & \text{if } t = 0 \text{ or if } g_{it-s} = c \\ & \text{for any } s \in [0, t] \\ \gamma_{it-1} & \text{otherwise.} \end{cases} \tag{6}$$

BPE says that the publisher will expect all advertisers to exert low effort at the beginning of the game. It will only come to expect high effort after an advertiser has calibrated that expectation by exerting high effort in the previous period. Lemmas 1 and 2 show that direct response advertisers will always exert high effort under BPE, whereas brand advertisers will be indifferent between high and low effort. Proposition 1 establishes that Bayesian Publisher Expectations uniquely maximize publisher profits in equilibrium.

The advertiser strategy described above is a multistage strategy in which the advertiser invests with a CPM bid to calibrate high click-through expectations and a CPC bid to profit from those heightened expectations. Therefore BPE requires the publisher to expect low effort from an advertiser if it ever uses a bid type that indicates a payoff behavior. This is similar to a "grim trigger strategy" studied in the context of a repeated prisoner's dilemma game (e.g., Osborne 2004). If a strategic player can successfully commit to punish another player's action indefinitely, it can deter the other player from taking such an action.

Another way to understand BPE is to relate it to "Gresham's law," or "bad money drives out the good." Because low effort may yield a cost advantage under CPC bidding, the publisher has no choice but to assume that every advertiser using a CPC bid will enter a low click-through rate. This leads all direct response advertisers to optimally use CPM bids, as shown below.

5.2. Advertisers' Effort Levels

Publisher expectations in Definition 2 allow for some clear statements about advertiser behavior. We proceed by analyzing advertiser effort levels conditional on choice of bid type in Lemmas 1 and 2. Lemma 3 then characterizes optimal bid type choices. After solving for optimal advertiser actions conditional on publisher beliefs, Proposition 1 shows that BPE uniquely maximizes publisher revenues.

Lemma 1. *The dominant strategy for any direct response advertiser under BPE is to choose a high effort level.*

Proof. See Appendix A.

Lemma 1 ensures that direct response advertisers will maximize their click-through rates in all periods because clicks directly increase their revenues. Any rational bid that allocates a direct response advertiser to an ad slot is one that gives the advertiser a positive profit per click; therefore the advertiser will always seek to get as many clicks as it can. This ensures that the proposed definition of Bayesian Publisher Expectations does not harm this core constituency of advertisers.

Next, we consider brand advertisers' effort levels.

Lemma 2. *Any brand advertiser entering a CPM bid under BPE is indifferent between high and low effort. A brand advertiser entering a CPC bid will exert low effort.*

Proof. See Appendix A.

Brand advertisers, by definition, do not profit from clicks. When they use CPM bidding, they do not pay by the number of clicks received. Because the click level does not affect their one-period revenues or one-period costs, brand advertisers are indifferent between high and low effort. Brand advertisers using CPC bids will always minimize effort in a one-period horizon to minimize costs.

Whereas the one-period incentives are rather straightforward, the dynamic incentives could easily be different. This is where Bayesian Publisher Expectations play a key role. They ensure that the advertiser does not have the ability to profit in future periods by setting a particular costless effort level in period t. Section 7.3 describes, and Appendix C formally shows, that any publisher expectation function that does not exploit past bid-type information will be wrong in equilibrium.

It is important to note that Bayesian Publisher Expectations do not eliminate brand advertisers' motivation to avoid clicks. However, they do prevent the advertisers from strategically using this motivation to lower advertising costs.

Finally, consider advertisers' choices of bid types.

Lemma 3. *Under BPE, direct response advertisers always enter CPM bids. Brand advertisers are always indifferent between the best-possible CPM and best-possible CPC bids.*

Proof. See Appendix A.

Lemma 3 indicates that BPEs generate a strict preference among direct response advertisers for CPM bids. This comes from the BPE requirement that advertisers "earn" a high click-through expectation. A direct response advertiser is better off if it does not try to earn this high expectation in the first period, and given this decision, it is better off under CPM bidding in all subsequent periods. Brand advertisers are indifferent between optimal CPM and optimal CPC bids.

It should be noted that Lemmas 1–3 do not depend on any stability in other advertisers' bids. Even if other advertisers are following nonstationary bidding strategies, one can confidently predict the relationship between advertisers' optimal bid-type choices and effort levels.

Bayesian Publisher Expectations endogenously limit strategic behavior on the part of brand advertisers by punishing them forever if they ever use CPC bidding. This punishment is effective in that it removes any incentive for brand advertisers to choose CPM bidding over CPC bidding, or vice versa. An unintended consequence of this is its effect on direct response advertisers, who respond by using CPM bidding. However, this unintended consequence does no harm to the affected direct response advertisers.

5.3. Equilibrium Existence and Publisher Belief Uniqueness

Lemmas 1–3 characterized optimal advertiser behavior under BPE; Proposition 1 now shows that equilibrium exists and that BPE uniquely maximizes publisher revenues.

Proposition 1. *At least one equilibrium exists. In any equilibrium, the publisher uses BPE, and advertisers behave in accordance with Lemmas* 1–3.

Proof. See Appendix A.

The advertiser behavior under BPE described in Lemmas 1–3 satisfies the incentive compatibility and advertiser rationality constraints of Definition 1. Two conditions remain to be proven. The first is that BPE maximizes publisher revenues. The proof of Proposition 1 shows that any alternate expectation function either reduces publisher revenues earned from brand advertisers or is inconsistent with the advertisers' equilibrium click-through rates, and no alternate expectation function can increase expenditures paid by direct response advertisers. Second, it is shown that BPE is the only expectation function consistent with advertisers' equilibrium strategies.

It is important to emphasize that BPE uniqueness does not imply equilibrium uniqueness. It is common, in both multislot auctions and dynamic games, to have multiplicity of equilibria. Multiplicity arises here because of nonuniqueness of advertisers' optimal bid levels and brand advertisers' bid types. The ability to characterize a unique publisher belief that must hold in any equilibrium of the game indicates the robustness of the paper's main result.

5.4. Comparing the Hybrid Advertising Auction to the Generalized Second Price Auction

Many advertising publishers do not currently use hybrid advertising auctions. A natural question is whether auction outcomes in nonhybrid advertising auction mechanisms can be supported by a hybrid advertising auction setting.

The answer to this question depends naturally on what other auction mechanism is used. As a baseline, consider the GSP auction of Edelman et al. (2007). This auction is the one that Google and Yahoo! use to allocate advertisers to slots on their search results. The GSP auction is the same as the hybrid advertising auction mechanism described in §5.5 with the exception that it has no CPM bidding option.

Proposition 2. *Under BPE, any repeated GSP equilibrium assignment of advertisers to slots can also be supported in a repeated hybrid advertising auction.*

Proof. See Appendix A.

Proposition 2 shows that, under BPE, any GSP equilibrium can be supported by a hybrid advertising auction. When advertisers are offered a CPM bidding option, BPE compels direct response advertisers to take it, but doing so leaves publisher revenues unchanged. BPE make brand advertisers indifferent between CPC and CPM bidding, as discussed above.

Proposition 2 is somewhat reassuring in that a publisher currently using a GSP auction knows its previous auction outcomes could also be obtained in the hybrid advertising auction setting if it sets the proper expectations. If some advertisers have an unmodeled preference for a CPM bidding option, offering this extra feature could attract advertisers to the publisher's platform. This possibility is discussed further in §7.2.

6. Alternate Payment Schemes and Social Efficiency

This section considers how to increase efficiency within the class of hybrid advertising auction mechanisms. It is well known that the second-price auction (Vickrey 1961, Clarke 1971, Groves 1973; known as

VCG) maximizes social welfare. Three results emerge. First, under a particular restriction of the model, we present a hybrid payment scheme, which achieves the socially efficient VCG assignment of advertisers to slots. Second, this proposed payment scheme maximizes publisher revenues among the set of socially optimal payment schemes. Third, in §6.2, without this restriction of the model, it can be proven that no mechanism always achieves the VCG assignment.

6.1. Efficient Allocation of Advertisers to Slots

The basic idea of the VCG mechanism is to charge each advertiser for the externality its purchase imposes on the publisher. As a simple example, imagine a setting with two slots and three advertisers. Assume that advertiser 1 buys the first slot and advertiser 2 buys the second slot. Advertiser 1's purchase of the first slot reduces the publisher's revenue from both of the other advertisers relative to a scenario where advertiser 1 does not purchase the first slot. Advertiser 2 is pushed to the second slot, reducing its payment, and advertiser 3 is pushed out of the auction, eliminating its payment. A VCG mechanism corrects for this by charging advertiser 1 according to the publisher's revenue reduction from each of the other two advertisers. This charge then ensures the lowest-value advertisers do not purchase the highest-value slots.

It is well known that the benchmark auction used in the search advertising industry, the generalized-second-price auction of Edelman et al. (2007), is not guaranteed to reach the socially optimal assignment of advertisers to slots. Similarly, the hybrid advertising auction is not guaranteed to reach the VCG assignment. This follows directly from Proposition 2. Because any GSP equilibrium assignment can be supported in a hybrid advertising auction, and because some GSP equilibrium assignments do not achieve the social optimum, it must be the case that some hybrid advertising auction assignments do not achieve the social optimum. It is therefore interesting to consider how to modify the hybrid advertising auction mechanism to ensure that it achieves the socially efficient VCG assignment of advertisers to slots. Part of the analysis in this subsection extends the recent work by Aggarwal et al. (2006).

Under a restrictive assumption about the model primitives, one can construct a payment scheme that elicits truthful bids and produces the VCG assignment. The assumption needed is that the position-dependent exposure and click-through rate multipliers fall by the same amount for each slot k.

Uniform value depletion condition (*UVDC*): $Y_k - Y_{k+1} = X_k - X_{k+1}, \forall k = 1, \ldots, K$.

The uniform value depletion condition is a rather strong restriction on the model. The behavioral implications are as follows. Assume that the position-dependent exposure multipliers of the first three slots are 0.5, 0.45, and 0.4, respectively, and the position-dependent click-through multiplier of the first slot is 0.2. It must then be the case that the second slot has a click-through multiplier of 0.15 and that the third slot has a click-through multiplier of 0.1. Note, though, that actual click-through rates may still differ because they depend on the ads placed in these slots.

The uniform value depletion condition is useful for two reasons. First, for settings in which it holds, one can establish a socially optimal mechanism. This may be particularly likely when the total number of ad slots K is small. Second, it helps establish the intuition needed to get close to the VCG assignment when the UVDC does not hold.

Next, consider how to construct the socially optimal payment scheme under UVDC. There are three challenges in designing this payment scheme to reach the VCG outcome: (1) charging each advertiser for the externality it imposes on the publisher, (2) removing incentives for advertisers to misreport their valuations (truthful bidding), and (3) (conditional on achieving the VCG outcome) maximizing publisher revenues. Definition 2 gives the prices that accomplish this. For notational convenience, let $I(g_{kt} = c)$ be an indicator function that equals one when the advertiser in slot k in period t enters a CPC bid, and let $I(g_{kt} = m)$ be an indicator function that equals one when the advertiser in slot k in period t enters a CPM bid.

Definition 3. For $1 \le k \le K$, define $p_{ikt}^{g_{it}}$ as i's payment for slot k in period t given bid type $g_{it} \cdot p_{ikt}^c$ is a per-click payment charged when i enters a CPC bid, whereas p_{ikt}^m is per-exposure payment charged when i enters a CPM bid. For $k = 1, \ldots, K$, let

$$p_{ikt}^{g_{it}} = \begin{cases} \displaystyle\sum_{k'=k}^{K}\left(\frac{X_{k'} - X_{k'+1}}{X_k}\right)\frac{\gamma_{k'+1t}^E}{\gamma_{it}^E} b_{k'+1t}^c I(g_{k'+1t} = c) \\ \displaystyle\quad + \sum_{k'=k}^{K}\left(\frac{X_{k'} - X_{k'+1}}{X_k}\right)\frac{Y_k b_{k'+1t}^m}{\gamma_{it}^E X_k} I(g_{k'+1t} = m) \\ \hfill \text{if } g_{it} = c \\ \displaystyle\sum_{k'=k}^{K}(X_{k'} - X_{k'+1})\frac{X_k \gamma_{k'+1t}^E b_{k'+1t}^c}{Y_k} I(g_{k'+1t} = c) \\ \quad + \sum_{k'=k}^{K}(X_{k'} - X_{k'+1}) b_{k'+1t}^m I(g_{k'+1t} = m) \\ \hfill \text{if } g_{it} = m. \end{cases}$$

When an advertiser gets a slot in the hybrid auction, it pushes all lower bidders to lower slots. When these bidders are allocated to lower slots, they pay for fewer clicks and impressions, thereby reducing the amount of money the publisher receives from those advertisers. The payment scheme in Definition 3 essentially forces the advertiser in slot k to compensate the publisher for the revenue lost from all lower advertisers.

Proposition 3 establishes that the proposed payment scheme achieves the VCG assignment. Equilibrium truth-telling requires that each direct response advertiser i in each slot k will bid $b_{it}^c = r_{Di}$ or $b_{it}^m = \gamma_{it} r_{Di}$ and that each brand advertiser i in each slot k will bid $b_{it}^c = r_{Bi}/\gamma_{it}$ or $b_{it}^m = r_{Bi}$. Notice that, as in a second-price auction, the payment of the advertiser in slot k is not a function of that advertiser's own bid.

Proposition 3. *Under the uniform value depletion condition, the payment scheme in Definition 3 produces a unique equilibrium with truthful advertiser bids and the VCG allocation of advertisers to slots.*

Proof. See Appendix A.

Proposition 3 shows that truthful bidding is a strictly dominant strategy under the payment scheme proposed in Definition 3. This is proven by considering an advertiser with the kth highest expected payment. This advertiser will strictly prefer slot $k'+1$ to slot k' for any slot $k' < k$. The same advertiser will also strictly prefer $k'-1$ to slot k' for any slot $k' > k$. Applied recursively and to all slots and all advertisers, it can be seen that the proposed payment scheme will always allocate each advertiser to its socially efficient place in the ad ranking.

Proposition 4. *Under the uniform value depletion condition, no other truth-telling payment scheme produces higher revenues.*

Proof. See Appendix A.

Proposition 4 establishes that the proposed payment scheme maximizes publisher revenues among the class of payment schemes that achieve the VCG assignment of advertisers to slots. Interestingly, Propositions 3 and 4 do not rely on any assumptions about publisher expectations. Direct response advertisers always exert high effort, independent of publisher expectations. This is because Definition 3 ensures that they are charged less than their valuation, and therefore additional clicks always increase profits. Brand advertisers are indifferent between high and low effort. Proposition B1 in Appendix B formalizes these results.

A relatively minor operational concern in implementing this mechanism is that, although there have been assumed K slots throughout, Definition 3 requires values for X_{K+1} and Y_{K+1}. There are various ways to handle this. Perhaps the most natural would be to experiment with adding an additional advertising slot $K+1$, measuring X_{K+1} and Y_{K+1} for this slot, and then using these figures in the calculation of $p_{ikt}^{g_{it}}$. It also may be possible to endogenize K.

6.2. Nonexistence of Efficient Mechanisms Without the UVDC

When the uniform value depletion condition holds, one is able to construct a payment scheme that achieves the socially optimal assignment of advertisers to slots. But when this condition does not hold, there is no payment scheme that can reliably achieve this assignment in equilibrium.

Proposition 5. *If the uniform value depletion condition fails to hold for some slot k, no payment scheme will always achieve the VCG assignment of advertisers to slots.*

Proof. See Appendix A.

Proposition 4 is a nonexistence result. When UVDC does not hold, there is no payment scheme that always achieves the VCG assignment. Recall that the core principle of the VCG equilibrium is to charge each advertiser for the harm that its inclusion does to the publisher. When UVDC does not hold, the reduction in publisher revenues from moving advertiser i from slot k to slot $k+1$ depends on advertiser i's type because exposure rate multipliers and click-through rate multipliers decrease nonuniformly. Because advertiser types are private, the publisher cannot know its revenue loss from each advertiser and therefore cannot charge i for the externality generated by its inclusion.

Although there is no socially optimal payment scheme without UVDC, it may be possible for the publisher to design a payment scheme that gets closer to the VCG assignment. There are two key criteria to consider. First is the mix of advertiser types in the population of advertisers. For example, if 99% of revenues come from brand advertisers, then advertiser externalities are more likely to depend on the differences in exposure multipliers than the differences in click-through multipliers, and Definition 3 can be modified appropriately.

The second criterion is the rates at which exposure and click-through multipliers decrease across slots. For example, it could be argued that $Y_k > X_k$, $\forall\, k$ since an ad exposure is necessary but not sufficient for a click (in the absence of random clicking). If this is the case, and if UVDC does not hold, then it seems intuitive to expect that $Y_k - Y_{k+1} > X_k - X_{k+1}$. Therefore the harm done in moving a brand advertiser one slot lower in the ranking is likely to exceed the harm in moving a direct response advertiser one slot lower in the ranking. This difference in externality sizes across advertiser types could also be used to modify the payment scheme in Definition 3 in an attempt to get closer to the VCG assignment.

7. Discussion and Implications

This paper presents the first analysis, to our knowledge, of equilibrium in a hybrid advertising auction.

It has shown that the publisher's expectations play a key role in whether offering multiple bid types will reduce seller revenues. The analysis has produced several results that could influence publishers', advertisers', and policy makers' actions. These implications are especially relevant to publishers that do not currently offer hybrid auctions but may do so in the future.

7.1. Conventional Wisdom in Hybrid Advertising Auctions

CPM ad pricing has traditionally been associated with brand advertisers. CPM pricing has been the standard ad price metric in traditional advertising media (e.g., television, newspapers, billboards) for decades. These media were typically dominated by brand advertisers, leading to the association. CPM pricing continues to be standard in online display advertising (Evans 2008).

CPC pricing, by contrast, is relatively new. It was invented by GoTo.com in 1998 in a successful effort to lure advertisers from rival websites. Google used CPM pricing before adopting the quality-weighted CPC model in 2002 to prevent advertisers from purchasing prominent search ad positions with low-click ads (Battelle 2005).[5] Early search advertising was dominated by direct response advertisers. Because these advertisers can usually track consumer profitability at a fine level of granularity, CPC pricing allows them to compare their marginal profits and advertising cost at the level of the individual click. Direct response advertisers prefer CPC pricing and tend to be associated with its use.

Facebook's help file, shown in Figure 5, explains the industry's conventional wisdom:

> As a CPC advertiser you are indicating that what is most important to you is having people click through to your website and controlling the actual cost to drive each individual person to your site. As a CPM advertiser you are indicating that it is more important to you that many people see your ad, not that they actually take action after seeing your ad. CPM advertising is usually more effective for advertisers who want to raise awareness of their brand or company, while CPC advertising is more effective for advertisers who are hoping for a certain response from users (like sales or registrations).

It appears that many advertisers believe this conventional wisdom. For example, Newcomb (2005) quotes an advertising executive, saying, "If the search campaign is largely for branding purposes, we will migrate to the CPM pricing model and bid as high as we can afford. For direct response clients, we'll stick to CPC…." An extensive search produced no statements contradicting the conventional wisdom in 2008, when the first draft of this paper was written. However, since then, some practitioners seem to have understood the conventional wisdom was flawed; see, e.g., Keane (2010).

Lemmas 1–3 and Proposition 1 show that this conventional wisdom is inconsistent with equilibrium publisher beliefs. These results suggest that advertisers should think carefully about their optimal choice of bid types conditional on optimal publisher expectations. When advertisers face a rational publisher using Bayesian Publisher Expectations, direct response advertisers should always use CPM bidding. If advertisers purchase slots from a publisher that does not understand the analysis in this paper, it will likely be the case that direct response advertisers face lower costs under CPM bids with high costless effort, and brand advertisers will be more profitable using CPC bids with low costless effort. If an advertiser does not know the publisher's expectation function, it would be advisable to test an identical set of ads under two different campaigns, one using CPC bidding and the other using CPM bidding, to see which performs better.

7.2. Should Publishers Offer Hybrid Advertising Auctions?

Proposition 2 shows that any equilibrium outcome in a generalized second-price auction format can also be supported in a hybrid advertising auction. Equilibrium advertiser behavior guaranteed by Lemmas 1–3 and BPE suggest that offering multiple bid types might not reduce revenue relative to a pure CPC auction. However, this implication is subject to the caveats that the model requires perfectly rational advertisers and does not have a unique equilibrium. Although we do not see a reason that offering multiple bid types would lead advertisers to change their equilibrium strategies, it is impossible to rule the possibility out.

Publishers may have an additional, unmodeled incentive to offer multiple types of bid. As mentioned above, CPM ad pricing has been traditionally associated with brand advertisers, whereas CPC ad pricing was developed for direct response advertisers. It may be that these advertisers are less than fully rational, or that they have some cost of adopting a nonpreferred ad price metric; if so, then offering multiple bid types may attract some set of advertisers to a publisher's platform. For example, CPM bidding may facilitate price comparisons across advertising media. Offering multiple bid types could increase publisher revenues because auction prices depend critically on the degree of competition in the auction. This platform adoption argument suggests that publishers may find it worthwhile to experiment with offering multiple types of bids.

[5] We thank a reviewer for pointing this out.

At the time of this writing, the online advertising market still appeared to be converging to equilibrium. If advertisers become accustomed to having multiple bidding options, they may demand them from other websites. Because the hybrid auction has been adopted by the advertising sales leader (Google) and Internet traffic leader (Facebook), one might speculate that more companies will adopt it in the future. One barrier to adopting the hybrid auction may be the associated complexity in forming click-through expectations; a major purpose for the present analysis is to show that offering multiple bid types need not harm auction efficiency or publisher revenues. Because multiple bid options may offer advertisers some benefit of convenience, and have no obvious downside for publishers, it seems natural to speculate that more companies will adopt them in the future. We predict that adoption would be most likely by those publishers whose advertising inventory appeals to both brand advertisers and direct response advertisers.

7.3. Historical Publisher Expectations

This paper has heretofore considered a strategic publisher that sets click-through expectations to maximize its profits. This requires that the publisher *dis*believe the conventional wisdom described in §7.1. However, Facebook's statement in Figure 5 reinforces this conventional wisdom. In addition, Bayesian publisher expectations requires use of an advertiser's bid type in addition to its past click-through rates. Figure 4 makes it clear that Google's quality scores are *not* based on past bid types. This evidence suggests that publishers currently offering hybrid auctions may not set fully rational click-through expectations. In addition, publishers that might offer hybrid auctions in the future may not have a full appreciation for the intricacies of this auction format.

This section considers what may happen if the publisher is not fully rational. A natural way for a nonstrategic publisher to set click-through expectations would be to look at each advertiser's past click-through performance, as suggested by Google's statements in Figure 4. We call this Historical Publisher Expectations, or HPE. Appendix C proves a series of results based on a publisher expectation function that only considers past click-through rates. This section discusses the intuition for these results and what they imply for publishers.

First, we formalize the idea (mentioned above) of a *bid reversal* as a multistep process in which a brand advertiser first enters a series of CPM bids with high effort levels. Because the CPM bid option is used, the high effort does not increase the advertiser's cost, but it does increase the publisher's expectation of future click-through rates. In the final step, the advertiser can enter a CPC bid with a low effort level. Under HPE, this series of actions may lower its costs. A repeated bid reversal is called a *lattice strategy*.

Two main results emerge. First, when the publisher uses HPE, *every* brand advertiser will engage in at least one bid reversal during the course of the game. Second, when the publisher uses HPE, there are regions of parameter space in which brand advertisers will use bid reversals with very high frequencies. This second result is proven by employing the one-stage deviation principle of Blackwell (1965) to show that a two-period bid reversal strategy may be played repeatedly in equilibrium.

A clear implication of the model is its prediction that publishers should use past periods' bid types as an additional factor in setting click-through rate expectations. It is necessary to consider that if a bidder switches bid type several times, he may be following a lattice strategy, and the publisher should lower his click-through rate expectations appropriately. Understanding this insight is particularly important for sellers that do not currently offer hybrid advertising auctions to maximize hybrid auction revenues.

An alternate mechanism to prevent the use of the lattice strategy is to "reset" click-through expectations to zero for any new advertisement, even when it is a minor variation on an old ad, as Google does. (Facebook does not currently disclose its policy.) The downside to this strategy is that a low-click advertiser can repeatedly formulate new advertisements, and each new low-click ad will be given a "blank slate" and placed appropriately. Instead, it may be more profitable to borrow information across many similar ads to set click-through expectations and incorporate data on past bidding behavior, as required by Bayesian Publisher Expectations.

7.4. Public Policy Implications

It is conceivable that search engines could be subject to enhanced regulatory attention in the future. Google dominates the U.S. market with 66% of all clicks. Calculated using shares of total clicks, the industry has a Herfindahl Index of 0.47 (Munarriz 2010). This is well above the 0.18 threshold at which the U.S. Justice Department considers a market "concentrated" and subject to enhanced merger scrutiny. This dominance is even more pronounced in other countries. For example, Google's share of clicks exceeds 90% in Germany and France.

If policy makers regulate publishers' business models in the future, it may be relevant to ask what type of auction enhances social welfare. This paper has shown how to achieve the social optimum exactly when the uniform value depletion condition holds. It has also shown that without this condition, no payment scheme will always achieve the social optimum in a hybrid auction setting. But the discussion in §6.2 reveals how to get close to the socially optimal payment scheme, taking into account advertiser heterogeneity and the rates at which click-through and exposure multipliers decrease across ad slots.

These results may interest search engines even in the absence of regulatory attention. If advertisers' values per click and impression are correlated with viewers' utility of seeing advertisers' ads, implementing a payment scheme to get close to the social optimum would likely improve consumers' utility of the search platform. This may help a search engine attract eyeballs, which it can then monetize through additional ad sales.

7.5. Maximizing Publisher Auction Revenues

An important and interesting question not broached here is how to design a hybrid auction mechanism to maximize publisher revenues, without regard for the social optimum. This is difficult in the current setting because advertisers are heterogeneous in two dimensions: type and reservation value.

The starting point for this topic is the seminal work of Myerson (1981). For analytical tractability, Myerson (1981) assumed that bidders' values were drawn from a single distribution function and derived a symmetric equilibrium. Ülkü (2009) extended Myerson's framework to allow bidders to have one-dimensional private information and showed how to find an optimal auction mechanism with a "generalized" virtual value. Edelman and Schwarz (2010) applied this framework in the GSP under an assumption that advertisers only differ in one dimension: value per click. If the valuations of both brand advertisers and direct response advertisers can be collapsed into a single index of advertiser differentiation, their framework could be applied to hybrid auctions. To the best of our knowledge, optimal mechanism design when selling heterogeneous objects to bidders that differ in more than one dimension is an extremely difficult and unsolved problem; this, then, is a profitable direction for future research.

Acknowledgments

The authors thank Wilfred Amaldoss, Anthony Dukes, Ben Edelman, Chuan He, Matthew Selove, Woochoel Shin, and seminar participants at Facebook, Google, University of Southern California, the 2009 NET Institute Conference, and the 2009 INFORMS Marketing Science Conference for helpful comments and discussions. They gratefully acknowledge financial support from the NET Institute (http://www.netinst.org) and the Kauffman Foundation.

Appendix A. Proofs of All Lemmas and Propositions

Lemma 1. *The dominant strategy for any direct response advertiser under BPE is to choose a high effort level.*

Proof. When direct response advertiser i gets slot k in period t with a CPC bid, its profit is $\pi^c_{ikt} = X_k\gamma_{it}r_{Di} - \gamma_{it}/\gamma^E_{it}C_{kt}$. This is weakly increasing in x_{it} so long as $r_{Di} \geq C_{kt}/X_k\gamma^E_{it}$, which must be true if the advertiser bid rationally and won the auction. If the direct response advertiser enters a CPM bid, its profit is $\pi^m_{ikt} = X_k\gamma_{it}r_{Di} - C_{kt}$, which is strictly increasing in x_{it}.

Intertemporal profit maximization also favors high effort. If the direct response advertiser enters a CPM bid in period $t+1$, then BPE ensures that costless effort in period t does not affect profits in period $t+1$. If i enters a CPC bid in period $t+1$, its profits are strictly lower since $\gamma^E_{it+1} = \underline{\gamma}_i$. Thus low effort in period t would not increase its profit in period $t+1$ by the arguments in the previous paragraph. Q.E.D.

Lemma 2. *Any brand advertiser entering a CPM bid under BPE is indifferent between high and low effort. A brand advertiser entering a CPC bid will exert low effort.*

Proof. We begin with the first statement. The brand advertiser's current-period revenues R_{ikt} are not a function of its costless effort level x_{it}. When it enters a CPM bid, its current costs also are not a function of its effort level x_{it}. Because it has used a CPM bid, BPE ensures that the gatekeeper will always expect it to choose low effort; therefore, none of its future profits can be impacted by its current effort level. Because current and future profits are unaffected by effort, the advertiser is indifferent between high and low effort.

Now, we consider the second statement. When a brand advertiser submits a CPC bid in period t, its one-period incentive is to minimize its click-through rate to minimize its advertising cost. Moreover, the publisher's BPE expectation function ensures that there is no effect of x_{it} on advertiser payoffs in future periods. The advertiser's use of a CPC bid in period t leads the publisher to set low click expectations in all future periods regardless of the effort entered in period t. Q.E.D.

Lemma 3. *Under BPE, direct response advertisers always enter CPM bids. Brand advertisers are always indifferent between the best-possible CPM and best-possible CPC bids.*

Proof. Beginning with the first claim, we show that any direct response advertiser i has a strict preference for a CPM bid in period 0 and then show that it never switches bid types. From Lemma 1, we know that direct response advertisers will exert high effort regardless of bid type choice. Since $\gamma^E_{i0} = \underline{\gamma}_i$ under BPE, if direct response advertiser i gets slot k in period 0, its profit from a CPC bid is $\pi^c_{ik0} = X_k\gamma_{i0}(r_{Di} - C_{k0}/\gamma^E_{i0})$, strictly less than its CPM bid profit of $\pi^m_{ik0} = X_k\gamma_{i0}r_{Di} - C_{k0}$, leading to the preference for CPM bidding in period 0.

Given a CPM bid in period 0, direct response advertiser i faces $\gamma^E_{i1} = \underline{\gamma}_i + \eta_i$ in period 1. It gets $\pi^m_{ik1} = X_k\gamma_{i1}r_{Di} - C_{k1}$ with a CPM bid, or $\pi^c_{ik1} = X_k\gamma_{i1}(r_{Di} - C_{k1}/\gamma_i)$ with a CPC bid in period 1. Profits are strictly higher with a CPM bid. We can iterate forward to show that CPM bidding dominates CPC bidding in every period of the game.

Now, suppose advertiser i is a brand advertiser. Under BPE, its first-period profit is either $\pi^m_{ik0} = Y_k r_{Bi} - C_{k0}$ with a CPM bid or $\pi^c_{ik0} = Y_k r_{Bi} - \gamma_{i0}/\gamma^E_{i0}C_{k0}$ with a CPC bid. π^c_{ik0}

is decreasing with effort, and conditional on low effort, we have $\pi^m_{ik0} = \pi^c_{ik0}$. BPE ensures that the brand advertiser will continue to face $\gamma^E_{it} = \underline{\gamma}_i$ in subsequent periods and will therefore maintain its indifference between bid types. Q.E.D.

Proposition 1. *At least one equilibrium exists. In any equilibrium, the publisher uses BPE, and advertisers behave in accordance with Lemmas* 1–3.

Proof. We first prove equilibrium existence by addressing each point in Definition 1, and we then show that BPE is the unique equilibrium belief by contradiction.

(1) *Incentive compatibility*: Lemmas 1–3 describe the optimal strategy of brand and direct response strategy, which maximizes their profit under BPE and satisfies the incentive compatibility constraint.

(2) *Individual rationality*: Assumption 2 states that each advertiser is charged the minimum amount necessary to keep its place in the ranking, which guarantees that advertisers get nonnegative profits in equilibrium.

(3) *Publisher optimality*: We prove this by contradiction. If the publisher's expectation is $\gamma^E_{it}(\gamma_{it} \mid g_{it-s} = c) \in (\underline{\gamma}_i, \underline{\gamma}_i + \eta]$, straightforward extensions to Lemmas 1–3 will show that the dominant strategy for brand advertisers is to submit a CPC bid with low effort, and for direct response advertisers, it is to submit a CPM bid with high effort. Assume that, under BPE, brand advertiser i gets position k with a CPM bid $b^{m^*}_{it}$. Under our supposed expectation function, i's optimal bid is $b^{c^*}_{it} = Y_k b^{m^*}_{it} / (X_k \gamma^E_{it}(\gamma_{it} \mid g_{it-s} = c))$. Holding other advertisers' bids constant, i gets slot k and pays $\underline{\gamma}_i / \gamma^E_{it} C_k$. Thus, the publisher's revenue loss is $(1 - \underline{\gamma}_i / (\gamma^E_{it})) C_k > 0$. Every direct response advertiser gets an identical payoff if we change BPE, so the switch in publisher expectations does not alter direct response advertisers' equilibrium bids or publisher revenues from those bids. Thus, any deviation from BPE will reduce publisher revenues from brand advertisers, and no deviation can increase publisher revenues from direct response advertisers. This completes our proof that BPE maximizes publisher revenues.

(4) *Consistency of publisher beliefs*: Under BPE, the strategy profile for direct response advertisers is to submit CPM bids with high effort, and brand advertisers get the same payoff with CPC bids with low effort or CPM bids with any effort. The publisher's belief about advertiser i in time period t is given by Bayes' rule:

$$\begin{aligned}&\Pr(\gamma^E_{it} = \underline{\gamma}_i + \eta_i \mid \sigma, H_{it})\\&\quad= \frac{\Pr(\gamma_{it} = \underline{\gamma}_i + \eta_i, H_{it} \mid \sigma)}{\Pr(\gamma_{it} = \underline{\gamma}_i + \eta_i, H_{it} \mid \sigma) + \Pr(\gamma_{it} = \underline{\gamma}_i, H_{it} \mid \sigma)}.\end{aligned}$$

Advertisers can be divided into two groups: those that submit at least one CPC bid and those that do not. For this first group, BPE specifies $\gamma^E_{it} = \underline{\gamma}_i$ for the duration of the game. First, from the advertiser's strategy profile, we know advertiser i is a brand advertiser. Second, advertiser inertia removes the possibility that i randomly changes effort across period because BPE has removed the profitability of doing that. Therefore

$$\Pr(\gamma_{it} = \underline{\gamma}_i + \eta_i, H_{it} \mid \sigma) = 0, \quad \text{and}$$

$$\begin{aligned}&\Pr(\gamma^E_{it} = \underline{\gamma}_i + \eta_i \mid \sigma, H_{it})\\&\quad= \frac{\Pr(\gamma_{it} = \underline{\gamma}_i + \eta_i, H_{it} \mid \sigma)}{\Pr(\gamma_{it} = \underline{\gamma}_i + \eta_i, H_{it} \mid \sigma) + \Pr(\gamma_{it} = \underline{\gamma}_i, H_{it} \mid \sigma)} = 0,\end{aligned}$$

proving consistency of BPE publisher beliefs. For the other group of advertisers, there are two cases. In case (1), advertiser i is a direct response advertiser. In this case, i always submits a CPM bid with a high click-through rate in time period $t - 1 \gamma_{it-1} = \underline{\gamma}_i + \eta_i$, from Lemmas 1 and 3. Then $\Pr(\gamma_{it} = \underline{\gamma}_i + \eta_i, H_{it} \mid \sigma) = 1$ and $\Pr(\gamma^E_{it} = \underline{\gamma}_i + \eta_i \mid \sigma, H_{it}) = 1$, proving consistency. In case (2), i is a brand advertiser. Suppose advertiser i used high effort in period $t-1$ and $\gamma_{it-1} = \underline{\gamma}_i + \eta_i$. By advertiser inertia, it will put the same effort level in period t, so again we have $\Pr(\gamma^E_{it} = \underline{\gamma}_i + \eta_i \mid \sigma, H_{it}) = 1$. If i instead used low effort in period $t-1$, the same logic can be used to show consistency of publisher beliefs. This completes the proof that BPE satisfies condition 4.

It remains to be shown that there is no other equilibrium belief. We do this by checking deviations from each part of BPE. We first check whether there exists any belief that is different from $\gamma^E_{it}(\gamma_{it} \mid g_{it-s} = c) = \underline{\gamma}_i$ for any $s \in [0, t]$. Suppose $\gamma^E_{it}(\gamma_{it} \mid g_{it-s} = c) \in (\underline{\gamma}_i, \underline{\gamma}_i + \underline{\eta}]$. Point 3 above shows this would violate publisher rationality, so any equilibrium belief must contain $\gamma^E_{it}(\gamma_{it} \mid g_{it-s} = c) = \underline{\gamma}_i$ for any $s \in [0, t]$. Second, we check whether there exists any belief that is different from γ_{it-1} if $g_{it-s} \neq c$ for every $s \in [0, t]$. If advertiser i is a direct response advertiser, then we know from Lemma 1 that $\gamma_{it} = \underline{\gamma}_i + \eta$, $\forall t$. Therefore, any deviation $\gamma^E_{it} \neq \gamma_{it-1}$ violates consistency of publisher beliefs. Having ruled out any alternate bid-type-contingent beliefs, we complete the consideration of any alternate belief function by looking at beliefs that are not contingent on bid types. The only possible equilibrium belief that is independent of bid type and maximizes publisher revenue is $\gamma^E_{it} = \underline{\gamma}_i$, $\forall t = [0, T]$. However, under this belief, the strategy profile for direct response advertisers is to choose to submit a CPM bid with high effort. This gives $\Pr(\gamma_{it} = \underline{\gamma}_i + \eta_i, H_{it} \mid \sigma) = 1$ and $\Pr(\gamma^E_{it} = \underline{\gamma}_i + \eta_i \mid \sigma, H_{it}) = 1$, so this belief violates Bayes' rule. Q.E.D.

Proposition 2. *Under BPE, any repeated GSP equilibrium assignment of advertisers to slots can also be supported in a repeated hybrid advertising auction game.*

Proof. First, we need to discuss in brief the repeated GSP auction equilibrium definition. The conditions in Definition 1 also define the GSP equilibrium, although advertisers have only one bid type (CPC).

Now, we prove the claim. Assume that there is a set of CPC bids $(b^c_t)_{t=1,\ldots,T}$ for every advertiser i and every time t that constructs equilibrium in a GSP auction. We need to prove that giving advertisers the option to use CPM bidding does not violate any of the conditions in Definition 1. By Lemma 3, brand advertisers are indifferent between the best-possible CPC and best-possible CPM bids. Therefore, brand advertisers would have no incentive to change their bids from CPC to CPM.

By Lemma 3, direct response advertisers strictly prefer CPM bidding with high effort in all periods in the hybrid advertising auction. If direct response advertiser i gets slot k in period t with a CPC bid b^c_{it} in the GSP auction, it can choose a CPM bid, which yields an equivalent total willingness to pay in the hybrid advertising auction: $b^m_{it} = X_k / Y_k (\underline{\gamma}_i + \eta_i) b^c_{it}$. If all direct response advertisers follow this strategy, then all advertisers will be allocated to the same slots in the hybrid advertising auction as they were in the repeated GSP auction. Q.E.D.

Proposition 3. *Under the uniform value depletion condition, the payment scheme in Definition* 3 *produces a unique equilibrium with truthful advertiser bids and the VCG allocation of advertisers to slots.*

Proof. An advertiser may have two reasons to bid untruthfully: to change its equilibrium slot assignment and/or to change its payment. Notice that the payment scheme in Definition 1 ensures that every advertiser's payment is unrelated to its bid, removing the motivation to change a payment. Furthermore, if a truthful bid would place the advertiser in its most preferred slot, then uncertainty about other advertisers' private valuations ensures that it will bid exactly its valuation, misreporting risks being allocated to a suboptimal slot without any cost reduction. Therefore, we only need to show that bidding truthfully places each advertiser in its most profitable advertising slot.

We start by supposing direct response advertiser i enters a truthful CPC bid $b^c_{it} = r_{Di}$, and we then show it has no incentive to deviate. Suppose i is assigned to slot k in equilibrium and considers raising its bid to get a lower slot $k' < k$. Its profit would be $\pi^c_{ik't} = X_{k'}\gamma_{it}(r_{Di} - p^c_{ik't})$. Suppose i compared this option to an adjacent slot $k'+1$, which yields $\pi^c_{ik'+1t} = X_{k'+1}\gamma_{it}(r_{Di} - p^c_{ik'+1t})$. Taking the difference between these profits $\pi^c_{ik'+1t} - \pi^c_{ik't}$ and substituting in for the per-click payments, we find

$$\pi^c_{ik'+1t} - \pi^c_{ik't} = \gamma_{it}\left[(X_{k'+1} - X_{k'})r_{Di} - \frac{1}{\gamma^E_{it}}(X_{k'+1}\gamma^E_{it}p^c_{ik'+1t} - X_{k'}\gamma^E_{it}p^c_{ik't})\right],$$

$$X_{k'+1}\gamma^E_{it}p^c_{ik'+1t} - X_{k'}\gamma^E_{it}p^c_{ik't} = \begin{cases} (X_{k'+1} - X_{k'})\gamma^E_{k'+1t}b^c_{k'+1t} \\ \quad \text{if advertiser in slot } k'+1 \text{ set } g_{k'+1t} = c \\ (X_{k'+1} - X_{k'})\dfrac{Y_k b^m_{k'+1}}{X_k} \\ \quad \text{if advertiser in slot } k'+1 \text{ set } g_{k'+1t} = m, \end{cases}$$

$$\pi^c_{ik'+1t} - \pi^c_{ik't} = \begin{cases} \gamma_{it}(X_{k'+1} - X_{k'})\left(r_{Di} - \dfrac{\gamma^E_{k'+1t}b^C_{k'+1t}}{\gamma^E_{it}}\right) \\ \quad \text{if advertiser in slot } k'+1 \text{ set } g_{k'+1t} = c \\ \gamma_{it}(X_{k'+1} - X_{k'})\left(r_{Di} - \dfrac{Y_k b^M_{k'+1t}}{X_k\gamma^E_{it}}\right) \\ \quad \text{if advertiser in slot } k'+1 \text{ set } g_{k'+1t} = m. \end{cases}$$

We know that $X_{k'+1} < X_{k'}$, so we will have $\pi^c_{ik'+1t} > \pi^c_{ik't}$ if and only if the second term in parentheses in the previous equation is negative. We also know that $k' < k$ and $k'+1 \le k$. Truthful bidding and Assumption 1 (assignment of advertisers to slots in order of expected total payment) therefore imply that

$$X_k\gamma^E_{it}r_{Di} = X_k\gamma^E_{it}b^c_{it} \le \begin{cases} X_k\gamma^E_{k'+1t}b^c_{k'+1t} & \text{if } g_{k'+1t} = c \\ Y_k b^m_{k'+1t} & \text{if } g_{k'+1t} = m. \end{cases} \tag{A1}$$

Rearranging terms in (A1) shows that it must be the case that $\pi^c_{ik'+1t} > \pi^c_{ik't}$. Since this holds for all $k' < k$, we can do this comparison recursively to show that the advertiser does not prefer any slot $k' < k$.

Now, suppose advertiser i considers lowering its bid to get a higher slot $k' > k$. We compare the profit when the advertiser gets slot $k'-1$ with slot k'. We can show that

$$\pi^c_{ik'-1t} - \pi^c_{ik't} = \begin{cases} \gamma_{it}(X_{k'-1} - X_{k'})\left(r_{Di} - \dfrac{\gamma^E_{k'-1t}b^C_{k'-1t}}{\gamma^E_{it}}\right) \\ \quad \text{if advertiser in slot } k'+1 \text{ entered } g_{k'+1t} = c \\ \gamma_{it}(X_{k'-1} - X_{k'})\left(r_{Di} - \dfrac{Y_k b^M_{k'-1t}}{X_k\gamma^E_{it}}\right) \\ \quad \text{if advertiser in slot } k'+1 \text{ entered } g_{k'+1t} = m. \end{cases}$$

Similar to the above, this implies $\pi^c_{ik'-1t} < \pi^c_{ik't}$, which, when applied recursively, indicates that advertiser i is strictly worse off in any slot $k' > k$.

If direct response advertiser i considers submitting a CPM bid to get slot $k' < k$,

$$\pi^m_{ik'+1t} - \pi^m_{ik't} = (X_{k'+1} - X_{k'})\gamma_{it}r_{Di} - (p^m_{ik'+1t} - p^m_{ik't}) = \begin{cases} (X_{k'+1} - X_{k'})\left(\gamma_{it}r_{Di} - \dfrac{X_k\gamma^E_{k'+1t}b^C_{k'+1t}}{Y_k}\right) \\ (X_{k'+1} - X_{k'})(\gamma_{it}r_{Di} - b^m_{k'+1t}). \end{cases}$$

We know that $\pi^m_{ik'+1t} - \pi^m_{ik't} \ge 0$ since

$$Y_k\gamma_{it}r_{Di} = Y_k b^m_{it} \le \begin{cases} X_k\gamma^E_{k'+1t}b^C_{k'+1t} \\ \quad \text{if advertiser in slot } k'+1 \text{ set } g_{k'+1t} = c \\ Y_k b^m_{k'+1t} \\ \quad \text{if advertiser in slot } k'+1 \text{ set } g_{k'+1t} = m. \end{cases}$$

Applied recursively, this indicates that the advertiser does no better in any slot $k' < k$. A similar analysis shows that a CPM bid for any lower slot $k' > k$ also decreases profits.

Now consider what happens when brand advertiser i enters a CPC bid to get slot $k' < k$.

$$\pi^c_{ik'+1t} - \pi^c_{ik't} = \begin{cases} (Y_{k'+1} - Y_{k'})r_{Bi} - (X_{k'+1} - X_{k'})\dfrac{\gamma_{it}}{\gamma^E_{it}}\gamma^E_{k'+1t}b^c_{k'+1t} \\ \quad \text{if advertiser in slot } k'+1 \text{ set } g_{k'+1t} = c \\ (Y_{k'+1} - Y_{k'})r_{Bi} - (X_{k'+1} - X_{k'})\dfrac{\gamma_{it}}{\gamma^E_{it}}\dfrac{Y_k b^m_{k'+1}}{X_k} \\ \quad \text{if advertiser in slot } k'+1 \text{ set } g_{k'+1t} = m, \end{cases}$$

and we know that

$$X_k\gamma^E_{it}\frac{r_{Bi}}{\gamma_{it}} \le \begin{cases} X_k\gamma^E_{k'+1t}b^c_{k'+1t} \\ \quad \text{if advertiser in slot } k'+1 \text{ set } g_{k'+1t} = c \\ Y_k b^m_{k'+1} \\ \quad \text{if the advertiser in slot } k'+1 \text{ set } g_{k'+1t} = m. \end{cases}$$

Since $k' < k$, we again can deduce that $\pi^c_{ik'+1t} \geq \pi^c_{ik't}$ for any $k' < k$. The structure of the proof when brand advertiser i submits a CPM bid to get a lower slot $k' > k$, submits a CPC bid to get a higher slot $k' < k$, or submits a CPC bid to get a lower slot $k' > k$ is similar and therefore omitted.

Next we show advertiser i in slot $k \in [1, K]$ will have nonnegative profits. Assume that advertiser i is a direct response advertiser with r_{Di} profit per click, truthful bid

$$b^{g_{it}}_{it} = \begin{cases} r_{Di} & \text{if } g_{it} = c \\ \gamma_{it} r_{Di} & \text{if } g_{it} = m, \end{cases}$$

and profit

$$\pi^{g_{it}}_{ikt} = \begin{cases} X_k \gamma_i r_{Di} - X_k \gamma_i p^c_{ikt} \\ X_k \gamma_i r_{Di} - p^m_{ikt}. \end{cases}$$

If i submits a CPC bid, then

$$\pi^c_{ikt} = X_k \gamma_i \Bigg[r_{Di} - \Bigg(\sum_{k'=k}^{K} \left(\frac{X_{k'} - X_{k'+1}}{X_k} \right) \frac{\gamma^E_{k'+1}}{\gamma^E_i} b^C_{k'+1} I(g_{k'+1} = c) + \sum_{k'=k}^{K} \left(\frac{X_{k'} - X_{k'+1}}{X_k} \right) \frac{Y_k b^M_{k'+1}}{\gamma^E_i X_k} I(g_{k'+1} = m) \Bigg) \Bigg].$$

Since $k' \in [k, K]$, we know that

$$X_k \gamma^E_i b^c_i \geq \begin{cases} X_k \gamma^E_{k'+1} b^c_{k'+1} & \text{if } g_{k'+1} = c \\ Y_k b^m_{k'+1} & \text{if } g_{k'+1} = m. \end{cases}$$

Thus

$$\begin{aligned} &\sum_{k'=k}^{K} \left(\frac{X_{k'} - X_{k'+1}}{X_k} \right) \frac{\gamma^E_{k'+1}}{\gamma^E_i} b^c_{k'+1} I(g_{k'+1} = c) \\ &\quad + \sum_{k'=k}^{K} \left(\frac{X_{k'} - X_{k'+1}}{X_k} \right) \frac{Y_k b^M_{k'+1}}{\gamma^E_i X_k} I(g_{k'+1} = m) \\ &\leq \sum_{k'=k}^{K} \left(\frac{X_{k'} - X_{k'+1}}{X_k} \right) \frac{\gamma^E_i}{\gamma^E_i} b^c_i I(g_{k'+1} = c) \\ &\quad + \sum_{k'=k}^{K} \left(\frac{X_{k'} - X_{k'+1}}{X_k} \right) \frac{X_k \gamma^E_i b^c_i}{\gamma^E_i X_k} I(g_{k'+1} = m) \\ &= \sum_{k'=k}^{K} \left(\frac{X_{k'} - X_{k'+1}}{X_k} \right) b^c_i I(g_{k'+1} = c) \\ &\quad + \sum_{k'=k}^{K} \left(\frac{X_{k'} - X_{k'+1}}{X_k} \right) b^c_i I(g_{k'+1} = m) \\ &= \sum_{k'=k}^{K} \left(\frac{X_{k'} - X_{k'+1}}{X_k} \right) b^c_i = b^c_i \left(1 - \frac{X_{K+1}}{X_k} \right). \end{aligned}$$

Therefore,

$$\begin{aligned} &r_{Di} - \Bigg(\sum_{k'=k}^{K} \left(\frac{X_{k'} - X_{k'+1}}{X_k} \right) \frac{\gamma^E_{k'+1}}{\gamma^E_i} b^c_{k'+1} I(g_{k'+1} = c) \\ &\quad + \sum_{k'=k}^{K} \left(\frac{X_{k'} - X_{k'+1}}{X_k} \right) \frac{Y_k b^M_{k'+1}}{\gamma^E_i X_k} I(g_{k'+1} = m) \Bigg) \\ &\geq r_{Di} - b^c_i \left(1 - \frac{X_{K+1}}{X_k} \right), \end{aligned}$$

implying that $\pi^c_{ikt} \geq 0$ since $X_{K+1} \geq 0$ and $r_{Di} = b^c_i$. If i submits a CPM bid,

$$\begin{aligned} \pi^m_{ikt} = X_k \gamma_i r_{Di} - \Bigg[&\sum_{k=i}^{K} (X_{k'} - X_{k'+1}) \frac{X_k \gamma^E_{k'+1} b^C_{k'+1}}{Y_k} I(g_{k'+1t} = c) \\ &+ \sum_{k=i}^{K} (X_{k'} - X_{k'+1}) b^M_{k'+1} I(g_{k'+1t} = m) \Bigg], \end{aligned}$$

which increases with γ_i, so the optimal click-through rate is $\gamma_i = \underline{\gamma}_i + \eta_i \geq \gamma^E_i$. Since $k' \in [k, K]$, we know that

$$Y_k b^m_i \geq \begin{cases} X_k \gamma^E_{k'+1} b^c_{k'+1} & \text{if } g_{k'+1} = c \\ Y_k b^m_{k'+1} & \text{if } g_{k'+1} = m, \end{cases}$$

and

$$\begin{aligned} &\sum_{k=i}^{K} (X_{k'} - X_{k'+1}) \frac{X_k \gamma^E_{k'+1} b^m_{k'+1}}{Y_k} I(g_{k'+1t} = c) \\ &\quad + \sum_{k=i}^{K} (X_{k'} - X_{k'+1}) b^m_{k'+1} I(g_{k'+1t} = m) \\ &\leq \sum_{k=i}^{K} (X_{k'} - X_{k'+1}) \frac{Y_k b^m_i}{Y_k} I(g_{k'+1t} = c) \\ &\quad + \sum_{k=i}^{K} (X_{k'} - X_{k'+1}) b^m_i I(g_{k'+1t} = m) \\ &= \sum_{k=i}^{K} (X_{k'} - X_{k'+1}) b^m_k = (X_k - X_{K+1}) b^m_k \leq X_k \gamma_i r_{Di}. \end{aligned}$$

Therefore, $\pi^m_{ikt} = X_k \gamma_i r_{Di} - p^m_{ikt} \geq X_k \gamma_i r_{Di} - X_k \gamma_i r_{Di} = 0$. The proof that a brand advertiser would get nonnegative profits is similar.

We have shown that truth-telling is strictly dominant conditional on advertisers being ranked in order of their total willingness to pay per slot and that advertisers would rationally participate in the auction. We now use a contradiction to prove that this is the only possible equilibrium. Suppose there is an equilibrium in which advertiser i in slot k had a higher willingness to pay than advertiser i' in slot $k' < k$. By the arguments presented above, this would imply $\pi^{g_{it}}_{ik't} > \pi^{g_{it}}_{ikt}$, so i would be strictly better off if it increases its bid until it wins slot k'. Therefore we have found our contradiction and this must not be an equilibrium. Q.E.D.

Proposition 4. *Under the uniform value depletion condition, no other truth-telling payment scheme produces higher revenues.*

Proof. We prove this by contradiction. Below, we use the following property of our payment scheme implied by Definition 3:

$$X_{k+1} \gamma^E_{it} p^c_{ik+1t} - X_k \gamma^E_{it} p^c_{ikt} = \begin{cases} (X_{k+1} - X_k) \gamma^E_{k+1t} b^c_{k+1t} \\ \quad \text{if advertiser in slot } k+1 \text{ set } g_{k+1t} = c \\ (X_{k+1} - X_k) \dfrac{Y_k b^m_{k+1t}}{X_k} \\ \quad \text{if advertiser in slot } k+1 \text{ set } g_{k+1t} = m. \end{cases} \tag{A2}$$

Suppose there is a payment scheme in a truth-telling mechanism that defines $q_{ikt}^{g_{it}}$ as the price advertiser i pays when allocated to slot k in time t given bid type g_{it}. Assume that direct response advertiser i enters a CPC bid to get slot k in period t. In equilibrium, it must be the case that i is weakly more profitable in slot k than in slot $k+1$. This implies

$$X_k \gamma_{it}^E q_{ikt}^c - X_{k+1} \gamma_{it}^E q_{ik+1t}^c \le (X_k - X_{k+1}) \gamma_{it}^E r_{Di}. \tag{A3}$$

The right-hand side of (A3) is minimized if we take the smallest value of r_{Di} that preserves the same assignment of advertisers to slots.

$$r_{Di} = \begin{cases} \dfrac{\gamma_{k+1t}^E b_{k+1t}^c}{\gamma_{k+1t}^E} + \varepsilon \\ \quad \text{if advertiser in slot } k+1 \text{ set } g_{k+1t} = c \\ \dfrac{Y_k b_{k+1t}^m}{\gamma_{k+1t}^E X_k} + \varepsilon \\ \quad \text{if advertiser in slot } k+1 \text{ set } g_{k+1t} = m. \end{cases} \tag{A4}$$

This value of r_{Di} guarantees i appears in position k if $\varepsilon > 0$ and $b_{it}^c = r_{Di}$ (which is implied by the assumed mechanism producing truth-telling in equilibrium). Substituting (A4) into (A3) and taking the limit as ϵ approaches 0, we have

$$\begin{aligned} &X_k \gamma_{it}^E q_{ikt}^C - X_{k+1} \gamma_{it}^E q_{ik+1t}^c \\ &\quad \le \begin{cases} (X_k - X_{k+1})(\gamma_{k+1t}^E b_{k+1t}^c + \gamma_{it}^E \varepsilon) \\ \quad \text{if advertiser in slot } k+1 \text{ set } g_{k+1t} = c \\ (X_k - X_{k+1}) \left(\dfrac{Y_k b_{k+1t}^m}{X_k} + \gamma_{it}^E \varepsilon \right), \\ \quad \text{if advertiser in slot } k+1 \text{ set } g_{k+1t} = m, \end{cases} \\ &\quad = X_k \gamma_{it}^E p_{ikt}^c - X_{k+1} \gamma_{it}^E p_{ik+1t}^c, \end{aligned} \tag{A5}$$

where the last equality comes from applying (A2). Since it must be that $p_{K+1} = 0$ and $q_{K+1} = 0$, (A5) implies $q_{ikt}^c \le p_{ikt}^c, \forall k \in [1, K]$.

We have shown that no VCG payment scheme can produce higher payments at any slot than our proposed mechanism when a direct response advertiser uses a CPC bid. The structure of the proof when direct response advertiser i submits a CPM bid, when brand advertiser i submits a CPM bid, and when brand advertiser i submits a CPC bid is very similar and therefore omitted. Q.E.D.

Proposition 5. *If the uniform value depletion condition fails to hold for some slot k, no payment scheme will always achieve the VCG assignment of advertisers to slots.*

Proof. We prove this by contradiction. To simplify the proof, we assume that $X_k - X_{k+1} < Y_k - Y_{k+1}$ for some slot k, but the same arguments presented here will apply if this inequality is reversed.

Assume that there exists a payment scheme that achieves the VCG assignment. Assume that this payment scheme defines $q_{kt}^{g_{it}}$ as the per-exposure payment of any advertiser in slot k in period t. Suppose brand advertiser i with value-per-exposure r_{Bi} occupies slot k in equilibrium. Because the assumed mechanism yields the VCG assignment, it must have equilibrium truth-telling, so $b_{ikt}^m = r_{Bi}$. A condition of equilibrium is that i is assigned to its most preferred slot, so we have $(Y_k - Y_{k+1}) r_{Bi} \ge q_{kt}^m - q_{k+1t}^m$.

Now suppose we replace brand advertiser i with a direct response advertiser j with a value-per-click $r_{Dj} = r_{Bi}/\gamma_{jt}$ and expected click-through rate γ_{jt}^E. If j enters a truthful CPM bid, it must be the case that $b_{jt}^m = \gamma_{jt} r_{Dj} = r_{Bi} = b_{it}^m$, and advertiser j is assigned to slot k. However, for

$$r_{Dj} \in \left(\frac{(q_{kt}^m - q_{k+1t}^m)}{\gamma_{jt}(Y_k - Y_{k+1})}, \frac{(q_{kt}^m - q_{k+1t}^m)}{\gamma_{jt}(X_k - X_{k+1})} \right),$$

this assignment would be unprofitable since we have

$$(X_k - X_{k+1}) \gamma_{jt} r_{Dj} < q_{kt}^m - q_{k+1t}^m. \tag{A6}$$

Therefore the presumed mechanism will not always induce direct response advertisers to enter truthful CPM bids.

Next we consider what happens if j submits a truthful CPC bid $b_{jt}^c = r_{Dj}$. Suppose we have a brand advertiser i' in slot $k+1$ with value per impression

$$r_{Bi'} \in \left(\frac{(q_{kt}^m - q_{k+1t}^m)}{(Y_k - Y_{k+1})}, \min\left(\frac{X_k \gamma_{jt} r_{Dj}}{Y_k}, \gamma_{jt} r_{Dj} \right) \right), \qquad \gamma_{i't} = \gamma_{jt},$$

and $\gamma_{i't}^E = \gamma_{jt}^E$. Given truthful bidding and a VCG assignment of advertisers to slots, it must be that

$$\begin{aligned} X_k \gamma_{jt}^E r_{Dj} &= X_k \gamma_{jt}^E b_{jt}^c \\ &\ge \begin{cases} X_k \gamma_{k+1t}^E b_{k+1t}^c \\ \quad \text{if the advertiser in slot } k+1 \text{ set } g_{k+1t} = c \\ Y_k b_{k+1t}^m \\ \quad \text{if the advertiser in slot } k+1 \text{ set } g_{k+1t} = m, \end{cases} \end{aligned}$$

which guarantees advertiser j gets slot k with bid b_{jt}^c. This requirement is easily satisfied because we have not placed restrictions on X_k and $Y_k \cdot Y_k$. For example, it is strictly satisfied if $\gamma_{jt}^E = \gamma_{jt}$ and $X_k = Y_k \cdot j$'s revealed preference implies

$$(X_k - X_{k+1}) \gamma_{jt} r_{Dj} \ge \gamma_{jt}(X_k p_{kt}^c - X_{k+1} q_{k+1t}^c). \tag{A7}$$

Combining (A6) and (A7), we have the following inequality:

$$\gamma_{jt}(X_k q_{kt}^c - X_{k+1} q_{k+1t}^c) \le (X_k - X_{k+1}) \gamma_{jt} r_{Dj} < q_{kt}^m - q_{k+1t}^m,$$

which simplifies to

$$\gamma_{jt}(X_k q_k^c - X_{k+1} q_{k+1}^c) < q_{kt}^m - q_{k+1t}^m. \tag{A8}$$

Advertiser i' should be put lower than advertiser j if we have a truthful mechanism. However, advertiser i' cannot submit a CPM bid with a bid $b_{i't}^m = r_{Bi'}$ to get slot $k+1$ since $(Y_k - Y_{k+1}) r_{Bi'} > q_{kt}^m - q_{k+1t}^m$. Advertiser i' has incentive to increase its bid and get slot k. The only way to preserve the truth-telling property of our assumed mechanism is to assume that advertiser i' submits a truthful CPC bid $b_{i't}^c = r_{Bi'}/\gamma_{i't}$ and gets slot $k+1$. This implies that the following condition has to be satisfied:

$$(Y_k - Y_{k+1}) r_{Bi'} \le \gamma_{i'}(X_k q_{ikt}^c - X_{k+1} q_{k+1}^C) = \gamma_j (X_k q_{kt}^c - X_{k+1} q_{k+1t}^c).$$

Together, we have

$$q_{kt}^m - q_{k+1t}^m < (Y_k - Y_{k+1}) r_{Bi'} \le \gamma_j (X_k q_{kt}^c - X_{k+1} q_{k+1t}^c).$$

This gives us an inequality

$$q^m_{kt} - q^m_{k+1t} < \gamma_j (X_k q^c_{kt} - X_{k+1} q^c_{k+1t}),$$

which directly contradicts inequality (A8). Therefore we have proven that our assumed VCG mechanism cannot exist. Q.E.D.

Appendix B. Advertiser Behavior in §6.1

Here, we show how advertisers behave under the payment scheme described in Definition 3. For this appendix only, we assume that $Y_k > X_k, \forall k$ because, in the absence of random clicks, an ad exposure is necessary but not sufficient for an ad click.

Proposition B1. *Under BPE and the payment scheme in Definition 3, all advertisers submit CPM bids. All direct response advertisers exert high effort, whereas brand advertisers are indifferent between low and high effort.*

Proof. To show that CPM bidding dominates, we calculate the difference in total payments between the two bid types for advertiser i in slot k in period t:

$$\begin{aligned} X_k \gamma^E_{it} p^c_{ikt} - p^m_{ikt} &= \sum_{k'=i}^{K} (X_{k'} - X_{k'+1}) \left(1 - \frac{X_k}{Y_k}\right) \gamma^E_{k'+1t} b^c_{k'+1t} I(g_{k'+1t} = c) \\ &\quad + \sum_{k'=i}^{K} (X_{k'} - X_{k'+1}) \left(\frac{Y_k}{X_k} - 1\right) b^m_{k'+1t} I(g_{k'+1t} = m). \end{aligned}$$

Since $Y_k > X_k$ for all k, so $Y_k/X_k - 1 > 0$ and $1 - X_k/Y_k > 0$. Therefore $X_k \gamma^E_{it} p^c_{ikt} - p^m_{ikt} > 0$, so advertisers always prefer CPM bidding.

The results regarding effort follow directly from the proofs of Lemmas 1 and 2, which can be directly applied under the new payment scheme. Q.E.D.

Appendix C. Historical Publisher Expectations

In this section we consider a nonrational publisher whose click-through expectations are only set as a function of past click-through rates, rather than considering current or past bid types. Definition C1 describes our publisher's expectations.

Definition C1. HPEs are based on click-through rates in the previous τ periods: $\gamma^E_{it} = G_{it}(\gamma_{i0}, \ldots, \gamma_{it-1})$ with the properties (i) $EG_{it}(\gamma_{i0}, \ldots, \gamma_{it-1}) = a$ if and only if $\gamma_{it-s} = a$, $\forall\, 0 < s < t$; and (ii) $\partial G_t / \partial \gamma_{it-s} \geq 0$ for all $s \in (0, 1, \ldots, t)$.

HPE incorporate the intuitive properties that publisher expectations are degenerate if and only if an advertiser has never changed its costless effort level and expected click-through rates are nondecreasing in advertisers' past click-through rates. This allows for a wide range of expectation functions, including weighted averages and stochastic functions.

We characterize our notion of equilibrium in Definition C2. Note that, because we now consider a nonstrategic publisher, we require a different equilibrium concept.

Definition C2. A subgame-perfect Nash equilibrium (SPNE) is defined by any set of bids $(b_t)_{t=1,\ldots,T}$ and any set of costless effort levels $(x_t)_{t=1,\ldots,T}$ for which the following conditions hold:

(1) *Incentive compatibility*: the choice sequence of effort levels $(x_{i1}, \ldots, x_{iT})$ and bids $(b^{g_{i1}}_{i1}, \ldots, b^{g_{iT}}_{iT})$ maximize expected profits $E\Pi^a_i$ for all advertisers $i = 1, \ldots, N$.

(2) *Advertiser rationality*: for any advertiser i who wins slot k in period t, $E\pi^{g_{it}}_{ikt} \geq 0$.

The difference between Definition C2 and Definition 2 is that here the publisher is not required to set click-through expectations to maximize its long-run profits, essentially modeling it as a nonstrategic player.

We define a bid reversal as an advertiser using a CPM bid with high effort in period t, followed by a CPC bid with low effort in period $t+1$. We begin by showing that historical publisher expectations lead to strategic bid reversals in equilibrium.

Proposition C1. *Under HPE, every brand advertiser engages in at least one bid reversal.*

Proof. First, we show that T consecutive CPM bids would be suboptimal. If brand advertiser i uses pure CPM bidding, it can costlessly exert high effort in at least one period prior to t, so $E\gamma^E_{it} = EG_t(\gamma_{i0}, \ldots, \gamma_{it-1}) > \underline{\gamma}_i, \forall t$. Its expected payoff is

$$\begin{aligned} E\,\Pi_{ikt} &= E(\pi_{ikt} + \pi_{ikt+1} + \pi_{ikt+2} + \cdots) \\ &= (R_{ik} - C_k) + \beta(R_{ik} - C_k) + \beta^2(R_{ik} - C_k) + \cdots. \end{aligned}$$

However, if i exerts high effort in any period prior to period t and uses a CPC bid in period t with low effort, then its expected payoff is

$$E\,\Pi'_{ikt} = \left(R_{ikt} - \frac{\underline{\gamma}_i}{EG_{it}} C_{kt}\right) + \beta(R_{ikt} - C_{kt}) + \beta^2(R_{ikt} - C_{kt}) + \cdots,$$

which is strictly higher than $E\Pi_{ikt}$. Thus T consecutive CPM bids cannot be optimal.

Second, we show T consecutive CPC bids would also be suboptimal. Assume that brand advertiser i exerts low effort in every period, $\gamma^E_{it} = \underline{\gamma}_i, \forall t$. Its expected payoff is

$$\begin{aligned} E\Pi_{ikt} &= E(\pi_{ikt} + \pi_{ikt+1} + \pi_{ikt+2} + \cdots) \\ &= (R_{ikt} - C_{kt}) + \beta(R_{ikt} - C_{kt}) + \beta^2(R_{ikt} - C_{kt}) + \cdots. \end{aligned}$$

If the advertiser submits a CPM bid with a high effort at time t, then

$$\begin{aligned} E\Pi'_{ikt} &= E(\pi'_{ikt} + \pi'_{ikt+1} + \pi'_{ikt+2} + \cdots) \\ &= (R_{ikt} - C_{kt}) + \beta E\left(R_{ikt} - \frac{\underline{\gamma}_i}{\gamma^E_{it+1}} C_{kt}\right) \\ &\quad + \beta^2 E\left(R_{ikt} - \frac{\underline{\gamma}_i}{\gamma^E_{it+2}} C_{kt}\right) + \cdots. \end{aligned}$$

We know that

$$\begin{aligned} \gamma^E_{it+q} &= G_{it+q}(\gamma_{i0}, \ldots, \gamma_{it}, \gamma_{it+1}, \ldots, \gamma_{it+q-1}) \\ &= G_{it+q}(\underline{\gamma}_i, \ldots, \underline{\gamma}_i + \eta, \underline{\gamma}_i, \ldots, \underline{\gamma}_i) \\ &> G_{it+q}(\underline{\gamma}_i, \ldots, \underline{\gamma}_i, \underline{\gamma}_i, \ldots, \underline{\gamma}_i) \,\forall q, \end{aligned}$$

so $\underline{\gamma}_i/\gamma^E_{it+q} < 1\,\forall\, q$. $E\Pi'_{ikt}$ is strictly higher than $E\Pi'_{ikt}$, indicating that T consecutive CPC bids would be suboptimal for any brand advertiser. Thus any brand advertiser must use at least one bid reversal in equilibrium. Q.E.D.

Proposition C1 establishes that, under HPE, brand advertisers will *strictly* prefer a mix of CPC and CPM bidding over time. Could it be the case that every brand advertiser is incented to make just one bid reversal over the course of the game, yielding minimal harm to publisher revenues? We now make it clear that bid reversals can be employed with maximally high frequencies under some conditions. We first define the set of strategies the advertiser would follow and then show that members of this set may be used in equilibrium under HPE.

Definition C3. A lattice strategy is a repeated strategy profile in which an advertiser submits a CPM bid with a high effort level in some periods and enters an equivalent CPC bid with a low effort level in other periods.

A lattice strategy is a tactic in which a brand advertiser repeatedly manipulates the publisher's historical expectations. It uses CPM bidding with a high click-through rate to calibrate a higher expected click-through rate in future periods. It later profits from this heightened click-through expectation by switching to CPC bidding with low effort to reduce its advertising costs.

To show the lattice strategy may be played in equilibrium, we employ the one-stage deviation principle of Blackwell (1965), a method of testing whether a sequential strategy profile in a dynamic game is subgame perfect. This principle states that a multistage strategy is subgame perfect in a dynamic game if and only if no player has incentive to deviate from this strategy profile in exactly one stage. Fudenberg and Tirole (1991) provide a detailed discussion. This substantially simplifies the process of showing a strategy is subgame perfect because it is not necessary to rule out deviations in every single stage of the game.

An infinite number of lattice strategies exist. We consider here a two-stage lattice strategy in which a brand advertiser uses CPM and CPC bids in alternating periods. Besides being simple relative to some other lattice strategies, the two-stage lattice strategy would lead to the most frequent bid reversals. This shows that strategic bid reversals could occur with high frequency.

To simplify the proof of Proposition C2, we assume that a particular form of HPE in which the publisher sets click-through expectations stochastically according to each advertiser's observed click-through rates in the previous τ periods. That is, we assume that

$$\gamma^E_{it} = \begin{cases} \underline{\gamma}_i + \eta_i & \text{with probability } \phi_{it} \\ \underline{\gamma}_i & \text{with probability } 1-\phi_{it}, \end{cases}$$

where

$$\phi_{it} = \tau^{-1}\sum_{s=1}^{\tau} x_{it-s} \quad \text{for } \tau \geq t, \quad \text{or}$$

$$\phi_{it} = t^{-1}\sum_{s=1}^{t} x_{it-s} \quad \text{for } \tau < t.$$

We also assume, for simplicity, that the assignment of advertisers to slots remains constant over periods t.

Proposition C2. *When*

$$\frac{\beta^2}{\tau}\frac{1-\beta^{\tau-1}}{1-\beta^2} < \frac{1}{2} < \frac{\beta}{\tau}\frac{1-\beta^{\tau-1}}{1-\beta^2},$$

an SPNE may exist in which all brand advertisers employ the two-stage lattice strategy under HPE.

Proof. Because the lattice strategy is a two-stage strategy, by the one-stage deviation principle we can prove it is an SPNE strategy if we can show that i has no incentive to deviate from it in either stage given other advertisers' actions. We prove that, for advertiser i, the lattice strategy strictly dominates other options in the first period (stage $t-1$). Then, we show that it is also strictly dominant in the second period (t). We assume here that $t > \tau$ for expositional convenience; the arguments are easily applied to smaller t.

We start by showing that the lattice strategy is strictly dominant for brand advertisers in period $t-1$. In period $t-1$, i has three other options: low effort/CPM bid, high effort/CPC bid, or low effort/CPC bid. Low effort with a CPM bid would leave i's payoff in period $t-1$ unchanged, but expected profits in period t are decreasing in γ_{it-1}. We can also quickly rule out a high effort/CPC bid as this payoff is strictly dominated by a high-effort/CPM bid. The final possible deviation in period $t-1$ is switching to a CPC bid with low effort, which would yield $\pi^c_{ikt-1} = R_{ik} - (\gamma_{it-1})/(\gamma^E_{it-1})C_{kt}$. Then i's total payoff from the switch would be

$$\begin{aligned} E\Pi'_{ikt-1} &= E(\pi'_{ikt-1}+\pi'_{ikt}+\pi'_{ikt+1}+\cdots) \\ &= \left[(R_{ik}-C_k)\frac{1}{2}+\left(R_{ik}-\frac{\gamma_i}{\underline{\gamma}_i+\eta_i}C_k\right)\frac{1}{2}\right] \\ &\quad +\beta\left[(R_{ik}-C_k)\frac{\tau/2+1}{\tau}+\left(R_{ik}-\frac{\gamma_i}{\underline{\gamma}_i+\eta_i}C_k\right)\frac{\tau/2-1}{\tau}\right] \\ &\quad +\beta^2(R_{ik}-C_k)+\beta^3\left[(R_{ik}-C_k)\frac{\tau/2+1}{\tau}\right. \\ &\quad \left. +\left(R_{ik}-\frac{\gamma_i}{\underline{\gamma}_i+\eta_i}C_k\right)\frac{\tau/2-1}{\tau}\right]+\cdots. \end{aligned}$$

Then i's total payoff from not switching at period $t-1$ is

$$\begin{aligned} E\Pi_{ikt-1} &= E(\pi_{ikt-1}+\pi_{ikt}+\pi_{ikt+1}+\cdots) \\ &= (R_{ik}-C_k)+\beta\left[(R_{ik}-C_k)\frac{1}{2}+\left(R_{ik}-\frac{\gamma_i}{\underline{\gamma}_i+\eta_i}C_k\right)\frac{1}{2}\right] \\ &\quad +\beta^2(R_{ik}-C_k)+\beta^3 \\ &\quad \cdot\left[(R_{ik}-C_k)\frac{1}{2}+\left(R_{ik}-\frac{\gamma_i}{\underline{\gamma}_i+\eta_i}C_k\right)\frac{1}{2}\right]+\cdots. \end{aligned}$$

Advertiser i has no incentive to deviate from the lattice strategy in period $t-1$ if and only if

$$E\Pi_{ikt-1}-E\Pi'_{ikt-1} \geq 0,$$

$$\begin{aligned} E\Pi_{ikt-1}-E\Pi'_{ikt-1} &= -\frac{1}{2}\frac{\eta_i}{\underline{\gamma}_i+\eta_i}C_k+\beta\frac{\eta_i}{\tau(\underline{\gamma}_i+\eta_i)}C_k \\ &\quad +\beta^3\frac{\eta_i}{\tau(\underline{\gamma}_i+\eta_i)}C_k+\cdots+\beta^{\tau-1}\frac{\eta_i}{\tau(\underline{\gamma}_i+\eta_i)}C_k \end{aligned}$$

$$= \frac{\eta_i}{\underline{\gamma}_i+\eta_i} C_k\left[-\frac{1}{2}+\frac{1}{\tau}(\beta+\beta^3+\cdots+\beta^{\tau-1})\right]$$
$$= \frac{\eta_i}{\underline{\gamma}_i+\eta_i} C_k\left(-\frac{1}{2}+\frac{\beta}{\tau}\frac{1-\beta^{\tau-1}}{1-\beta^2}\right).$$

Thus, if $\beta/\tau(1-\beta^{\tau-1})/(1-\beta^2) > 1/2$, the lattice strategy is strictly dominant in period $t-1$.

In period t, i can again deviate by choosing a high effort/CPC bid, low effort/CPM bid, or high effort/CPM bid. We can quickly rule out these first two options because the low effort/CPC bid dominates both the high effort/CPC and low effort/CPM bid options. The main question is whether the deviation to the high effort/CPM bid option is profitable. In this case, its total expected payoff in period t is

$$\begin{aligned} E\Pi'_{ikt} &= E(\pi'_{ikt}+\pi'_{ikt+1}+\pi'_{ikt+2}+\cdots) \\ &= (R_{ik}-C_k)+\beta(R_{ik}-C_k) \\ &\quad +\beta^2\bigg[(R_{ik}-C_k)\frac{\tau/2-1}{\tau} \\ &\qquad +\left(R_{ik}-\frac{\gamma_i}{\underline{\gamma}_i+\eta_i}C_k\right)\frac{\tau/2+1}{\tau}\bigg]+\cdots. \end{aligned}$$

The expected payoff if the advertiser does not deviate at period t is

$$\begin{aligned} E\Pi_{ikt} &= E(\pi_{ikt}+\pi_{ikt+1}+\pi_{ikt+2}+\cdots) \\ &= \left[(R_{ik}-C_k)\frac{1}{2}+\left(R_{ik}-\frac{\gamma_i}{\underline{\gamma}_i+\eta_i}C_k\right)\frac{1}{2}\right]+\beta(R_{ik}-C_k) \\ &\quad +\beta^2\bigg[(R_{ik}-C_k)\frac{\tau/2-1}{\tau} \\ &\qquad +\left(R_{ik}-\frac{\gamma_i}{\underline{\gamma}_i+\eta_i}C_k\right)\frac{\tau/2+1}{\tau}\bigg]+\cdots. \end{aligned}$$

The difference in payoff is

$$\begin{aligned} \Delta\Pi_{it} &= E\Pi_{ikt}-E\Pi'_{ikt} \\ &= \frac{1}{2}\frac{\eta_i}{\underline{\gamma}_i+\eta_i}C_k-\beta^2\frac{\eta_i}{\tau(\underline{\gamma}_i+\eta_i)}C_k \\ &\quad -\beta^4\frac{\eta_i}{\tau(\underline{\gamma}_i+\eta_i)}C_k-\cdots-\beta^\tau\frac{\eta_i}{\tau(\underline{\gamma}_i+\eta_i)}C_k \\ &= \frac{\eta_i}{\underline{\gamma}_i+\eta_i}C_k\left[\frac{1}{2}-\frac{1}{\tau}(\beta^2+\beta^4+\cdots+\beta^\tau)\right] \\ &= \frac{\eta_i}{\underline{\gamma}_i+\eta_i}C_k\left(\frac{1}{2}-\frac{\beta^2}{\tau}\frac{1-\beta^{\tau-1}}{1-\beta^2}\right). \end{aligned}$$

So if

$$\frac{\beta}{\tau}\frac{1-\beta^{\tau-1}}{1-\beta^2} > \frac{1}{2} \quad \text{and} \quad \frac{\beta^2}{\tau}\frac{1-\beta^{\tau-1}}{1-\beta^2} < \frac{1}{2},$$

the lattice strategy is strictly dominant in period t. Q.E.D.

Essentially, the lattice strategy involves an investment behavior and a payoff behavior. The investment takes place in the periods in which the brand advertiser enters a CPM bid with a high effort level. The payoff occurs in the periods when the advertiser uses a CPC bid with a low costless effort level. The conditions needed for SPNE existence require that (1) the net present value of future investment returns are high enough to prevent the advertiser from taking payoffs during investment periods and (2) that current-period payoff incentives are high enough that the advertiser does not optimally invest during payoff periods.

More sophisticated lattice strategies can also be analyzed, such as "bid CPM with high effort for w periods, followed by CPC with low effort for y periods," or "draw a binomial random variable w with probability p, and play CPM with high effort whenever $w=1$; play CPC with low effort when $w=0$." The parameter space in which a lattice strategy may be played in equilibrium increases in the length of the number of periods of CPM bidding included in the strategy. The profitability of all lattice strategies can be eliminated by BPE.

References

Agarwal, N., S. Athey, D. Yang. 2009. Skewed bidding in pay-per-action auctions for online advertising. *Amer. Econom. Rev.* **99**(2) 441–447.

Aggarwal, G., A. Goel, R. Motvani. 2006. Truthful auctions for pricing search keywords. *Proc. 7th ACM Conf. Electronic Commerce, Ann Arbor, MI*, ACM, New York, 1–7.

Anderson, E. T., D. I. Simester. 1998. The role of sale signs. *Marketing Sci.* **17**(2) 139–155.

Athey, S., G. Ellison. 2008. Position auctions with consumer search. Working paper, Massachusetts Institute of Technology, Boston.

Baltas, G. 2003. Determinants of Internet advertising effectiveness: An empirical study. *Internat. J. Market Res.* **45**(4) 505–513.

Battelle, J. 2005. *The Search: How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture*. Penguin Books, London.

Blackwell, D. 1965. Discounted dynamic programming. *Anal. Math. Statist.* **36**(1) 226–235.

Borgers, T., I. Cox, M. Pesendorfer, V. Petricek. 2008. Equilibrium bids in sponsored search auctions: Theory and evidence. Working paper, University of Michigan, Ann Arbor. http://www-personal.umich.edu/~tborgers/EquiBids.pdf.

Chandon, J. L., M. S. Choutou, D. R. Fortin. 2003. Effects of configuration and exposure levels on responses to Web advertisements. *J. Advertising Res.* **43**(2) 217–229.

Chen, Y., C. He. 2006. Paid placement: Advertising and search on the Internet. NET Institute Working Paper 06-02, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=936472.

Chen, J., D. Liu, A. B. Whinston. 2009. Auctioning keywords in online search. *J. Marketing* **73**(4) 125–141.

Chu, W. 1992. Demand signalling and screening in channels of distribution. *Marketing Sci.* **11**(4) 327–347.

Clarke, E. H. 1971. Multipart pricing of public goods. *Public Choice* **11**(1) 17–33.

Danaher, P. J., G. W. Mullarkey. 2003. Factors affecting online advertising recall: A study of students. *J. Advertising Res.* **43**(3) 252–267.

Desai, P. S. 2000. Multiple messages to retain retailers: Signaling new product demand. *Marketing Sci.* **19**(4) 381–389.

Drèze, X., F.-X. Husherr. 2003. Internet advertising: Is anybody watching? *J. Interactive Marketing* **17**(4) 8–23.

Edelman, B., H. S. Lee. 2008. CPC/CPA hybrid bidding in a second price auction. HBS Working Paper 09-074, Harvard Business School, Boston.

Edelman, B., M. Schwarz. 2010. Optimal auction design and equilibrium selection in sponsored search auction. *Amer. Econom. Rev.* **100**(2) 597–602.

Edelman, B., M. Ostrovsky, M. Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *Amer. Econom. Rev.* **97**(1) 242–259.

Eldon, E. 2010. Facebook revenues up to $700 million in 2009, on track towards $1.1 billion in 2010. *Inside Facebook* (March 2), http://www.insidefacebook.com/2010/03/02/facebook-made-up-to-700-million-in-2009-on-track-towards-1-1-billion-in-2010/.

Evans, D. S. 2008. The economics of the online advertising industry. *Rev. Network Econom.* **7**(3) 359–391.

Even-Dal, E., M. Kearns, J. Wortman. 2007. Sponsored search with contexts. X. Deng, F. C. Grahan, eds. *Internet and Network Economics*. (*WINE 2007*). *Lecture Notes in Computer Science*, Vol. 4858. Springer, Berlin, 312–317.

Facebook.com. 2010. Press room. Retrieved August 8, 2010, http://www.facebook.com/press/info.php?statistics.

Feng, J., J. Xie. 2007. Performance-based advertising: Price and advertising as signals of product quality. Working paper, University of Florida, Gainesville. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1027296.

Fudenberg, D., J. Tirole. 1991. *Game Theory*. MIT Press, Cambridge, MA.

Ghose, A., S. Yang. 2009. An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Sci.* **55**(10) 1605–1622.

Goel, A., K. Munagala. 2009. Hybrid keyword search auctions. *Proc*. 18*th World Wide Web Conf.* (*WWW'09*), ACM, New York, 221–230.

Goldfarb, A., C. Tucker. 2010. Search engine advertising: Channel substitution when pricing ads to context. Net Institute Working Paper 07-23. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1021451.

Goldfarb, A., C. Tucker. 2011a. Advertising bans and the substitutability of online and offline advertising. *J. Marketing Res*. Forthcoming.

Goldfarb, A., C. Tucker. 2011b. Online display advertising: Targeting and obtrusiveness. *Marketing Sci.*, ePub ahead of print February 9, http://mktsci.journal.informs.org/cgi/content/abstract/mksc.1100.0583v1.

Google Inc. 2010. 2009 Annual report. Accessed August 8, 2010, http://investor.google.com/pdf/2009_google_annual_report.pdf.

Groves, T. 1973. Incentives in teams. *Econometrica* **41**(4) 617–631.

He, C., Y. Chen. 2006. Managing e-marketplace: A strategic analysis of non-price advertising. *Marketing Sci.* **25**(2) 175–187.

Hu, Y. J. 2004. Performance-based pricing models in online advertising. Working paper, Purdue University, West Lafayette, IN. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=501082.

Hu, Y., J. Shin, Z. Tang. 2010. In search of the holy grail of online advertising: Cost per click versus cost per action. Working paper, Purdue University, West Lafayette, IN.

Jerath, K., L. Ma, Y.-H. Park, K. Srinivasan. 2010. A "position paradox" in sponsored search auctions. Johnson School Research Paper 36-09, Carnegie Mellon University, Pittsburgh. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1464545.

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

Keane, M. 2010. Facebook's ad changes could cost CPM advertisers. *Econsultancy* (May 3), http://econsultancy.com/blog/5849-facebook-s-ad-changes-could-cost-cpm-advertisers.

Kempe, D., M. Mahdian. 2007. A cascade model for externalities in sponsored search. C. Papadimitrou, S. Zhang, eds. *Internet and Network Economics*. (*WINE* 2008). *Lecture Notes in Computer Science*, Vol. 5385. Springer, Berlin, 585–596.

Krishnamurthy, S. 2000. Deciphering the Internet advertising puzzle. *Marketing Management* **9**(3) 35–39.

Libai, B., E. Biyalogorsky, E. Gerstner. 2003. Setting referral fees in affiliate marketing. *J. Service Res.* **5**(4) 303–315.

Loechner, J. 2009. 8% of Internet users account for 85% of all clicks. *MediaPost Blogs* (October 13), http://www.mediapost.com/publications/?fa=Articles.showArticle&art_aid=115210.

Lohtia, R., N. Donthu, E. K. Hershberger. 2003. The impact of content and design elements on banner advertising click-through rates. *J. Advertising Res.* **43**(4) 410–418.

Mayzlin, D., J. Shin. 2009. Uninformative advertising as an invitation to search. Working paper, Yale University, New Haven, CT.

Milgrom, P. 2003. Putting auction theory to work: The simultaneous ascending auction. *J. Political Econom.* **108**(2) 245–272.

Moorthy, S., K. Srinivasan. 1995. Signaling quality with a money-back guarantee: The role of transaction costs. *Marketing Sci.* **14**(4) 442–466.

Munarriz, R. A. 2010. Bing and Google gang up on Yahoo! *Motley Fool* (March 10), http://www.fool.com/investing/high-growth/2010/03/10/bing-and-google-gang-up-on-yahoo.aspx.

Myerson, R. 1981. Optimal auction design. *Math. Oper. Res*. **6**(1) 58–73.

Nazerzadeh, H., A. Saberi, R. Vohra. 2008. Dynamic cost-per-action mechanisms and applications to online marketing. *Proc. 17th Internat. Conf. World Wide Web (WWW'08)*, ACM, New York, 179–188.

Newcomb, K. 2005. Google's CPM ads meet lukewarm reception. *ClickZ* (April 27), http://www.clickz.com/3500856.

Nielsen. 2009. Global faces and networked places. Accessed August 8, 2010, http://blog.nielsen.com/nielsenwire/wp-content/uploads/2009/03/nielsen_globalfaces_mar09.pdf.

Osborne, M. J. 2004. *An Introduction to Game Theory*. Oxford University Press, New York.

Pauwels, K. 2004. How dynamic consumer response, competitor response, company support, and company inertia shape long-term marketing effectiveness. *Marketing Sci.* **23**(4) 596–610.

Robinson, H., A. Wysocka, C. Hand. 2007. Internet advertising effectiveness. *Internat. J. Advertising* **26**(4) 527–541.

Rutz, O. J., R. E. Bucklin. 2007. A model of individual keyword performance in paid search advertising. Working paper, Yale University, New Haven, CT.

Rutz, O. J., R. E. Bucklin. 2010. From generic to branded: A model of spillover dynamics in paid search advertising. *J. Marketing Res.* Forthcoming.

Simester, D. 1995. Signalling price image using advertised prices. *Marketing Sci.* **14**(2) 166–188.

Spence, M. 1973. Job market signaling. *Quart. J. Econom.* **87**(3) 355–374.

Tellis, G. J. 2004. *Effective advertising: Understanding When, How, and Why Advertising Works*. Sage, Thousand Oaks, CA.

Ülkü, L. 2009. Optimal conbinatorial mechanism design. Working paper, Instituto Technológico Autónomo de México, Mexico City. http://allman.rhon.itam.mx/∼levent.ulku/OCMDAugust2009.pdf.

Varian, H. 2007. Position auctions. *Internat. J. Indust. Organ.* **25**(6) 1163–1178.

Varian, H. 2009. Online ad auctions. *Amer. Econom. Rev. Papers Proc.* **99**(2) 430–434.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* **16**(1) 8–37.

Wilbur, K., Y. Zhu. 2009. Click fraud. *Marketing Sci.* **28**(2) 298–308.

Xu, L., J. Chen, A. B. Whinston. 2008. To place better or price cheaper? Bidding and pricing under keyword advertising. Presentation, Seventh Workshop on e-Business (WeB), December 13, Paris.

Yao, S., C. F. Mela. 2011. A dynamic model of sponsored search advertising. *Marketing Sci.* Forthcoming.

Yaveroglu, I., N. Donthu. 2008. Advertising repetition and placement issues in online environments. *J. Advertising* **37**(2) 31–43.