

Pay-for-performance (P4P) pricing schemes such as pay per click and pay per action have increased in popularity in Internet advertising. Meanwhile, pay-per-impression (PPI) schemes persist, and several publishers have begun to offer a hybrid mix of PPI and P4P schemes. Given the proliferation of pricing schemes, this study examines the optimal choices for publishers. The authors highlight two-sided information asymmetries in online advertising markets and the consequent trade-offs faced by a high-quality publisher using P4P schemes. Pay-for-performance schemes enable a high-quality publisher to reveal its superior quality; however, such schemes may incur allocative inefficiencies stemming from inaccurate estimates of advertiser qualities. The authors identify conditions under which a publisher may opt for a PPI, P4P, or hybrid scheme and, in doing so, provide theoretical explanations for the observed variations in the pricing schemes and the increasing popularity of hybrid schemes. Using a new “uncompromised” equilibrium refinement, the authors find that the hybrid scheme can emerge as an equilibrium choice in a variety of conditions. In addition, they provide prescriptive guidelines for firms choosing between different pricing schemes.

Keywords: online advertising, pay for performance, information asymmetry, hybrid pricing

Online Supplement: <http://dx.doi.org/10.1509/jmr.13.0074>

Information Asymmetry and Hybrid Advertising

Internet advertising has been increasing at a rapid pace. In the United States, revenues from Internet advertising reached \$36.5 billion in 2012 (Internet Advertising Bureau 2013) and surpassed print advertising revenues in 2010. In keeping with the rising popularity of Internet advertising,

there is increased interest among academics and practitioners in understanding the Internet advertising landscape. A notable aspect of Internet advertising is the proliferation of different pricing schemes. In addition to the traditional pay-per-impression (PPI; also known as cost per mille) pricing, many pay-for-performance (P4P) schemes such as pay per click (PPC; also known as cost per click), pay per sale (PPS), pay per action (PPA), and pay per lead are now prevalent. Under the PPI scheme, advertisers pay each time their advertisements are shown to an Internet user; under the P4P schemes, advertisers pay only for measurable outcomes, such as clicks, sales, and customer leads. In addition, many Internet advertising providers (which we refer to as “publishers”) offer both PPI and P4P schemes for advertisers to choose; examples include websites and networks such as Google Display Ads, Facebook, and AOL Advertising. The diversity of pricing schemes raises compelling questions regarding publishers’ motivations to use different

*De Liu is Associate Professor, Information and Decision Sciences Department, Carlson School of Management, University of Minnesota (e-mail: deliu@umn.edu). Siva Viswanathan is Associate Professor, Department of Decision, Operations & Information Technologies, Robert H. Smith School of Business, University of Maryland (e-mail: sviswana@rhsmith.umd.edu). The authors thank seminar participants at University of Minnesota; University of Texas at Austin; University of International Business and Economics; Tsinghua University; Southwestern University of Finance and Economics; Theory in Economics of Information Systems; Algorithmic Game Theory Workshop at University of California, Los Angeles; Workshop on Data Mining and Audience Intelligence for Online Advertising; and China Summer Workshop on Information Management for their valuable feedback. This research received support from the National Science Foundation of China under Grant No.71171052. Yuxin Chen served as associate editor for this article.

pricing schemes and which pricing schemes publishers should offer.

The choice of pricing schemes seems most relevant when publishers and advertisers are asymmetrically informed about each other. On the one hand, when advertisers use a PPI scheme, they must determine their willingness to pay per impression from a certain publisher, which inevitably involves estimation of the publisher's quality in terms of driving performance outcomes (e.g., clicks, sales). When these estimates deviate from a publisher's true quality, the publisher could be under- or overcompensated. A P4P scheme, on the other hand, can avoid this issue because it compensates publishers on the basis of measured rather than estimated performance outcomes. To this end, a high-quality publisher may find it advantageous to use P4P schemes to avoid being undercompensated by advertisers. However, publishers that offer P4P schemes must assume the risk of uncertain advertiser quality in terms of generating outcomes; after all, publishers are paid only when a measured outcome (e.g., a click) occurs. Inaccurate estimation of advertiser quality not only exposes publishers to payment fluctuations but also, more importantly, may cause inefficient allocation of advertising slots, which can result in significant revenue loss and is a key concern of the Internet advertising auction design (Lahaie 2006; Zhu and Wilbur 2010). In economics terminology, both of the aforementioned considerations—publishers' uncertainty about advertiser quality and advertisers' uncertainty about publisher quality—are cases of information asymmetry. The objective of this article is to examine the optimal choice of pricing schemes when there are information asymmetries between publishers and advertisers. In doing so, our study provides a theoretical explanation for the observed diversity in pricing schemes.

We analyze a publisher's pricing scheme choice using a game theoretic model in which the publisher announces a pricing scheme and advertisers compete for advertising slots in an auction using the announced pricing scheme. To capture information asymmetry, we assume that the two parties do not know each other's true qualities but observe noisy signals of them. By varying the informativeness of these signals, we accommodate a wide range of market conditions and examine the corresponding equilibrium pricing schemes. We use this model to answer the following specific research questions: Which pricing schemes should publishers offer in a market equilibrium? What are the theoretical explanations for the coexistence of multiple pricing schemes and the increasing popularity of hybrid schemes?

These research questions are partly motivated by varied practices and the lack of a systematic understanding of pricing

schemes in the online advertising market, particularly display advertising. Anecdotal evidence indicates varied practices among publishers. For example, whereas Facebook uses a hybrid scheme for its display advertisements, Amazon.com uses a PPI scheme for its counterpart. Some publishers (e.g., Tremor, Google Display Ads) have switched from one pricing scheme to another (see Table 1). Practitioners and industry experts often give different and sometimes contradictory opinions on which pricing schemes publishers should offer.¹ An article by *The Economist* considers PPS the "holy grail of advertising" (*The Economist* 2005), but PPS has not yet become the dominant pricing scheme. In addition to providing a theoretical basis for the variation in pricing schemes publishers adopt, our study also offers valuable prescriptive guidelines for advertisers and publishers choosing between the different pricing schemes.

We note that the nature and degree of information asymmetry may vary from market to market. Yet there are reasons to believe that information symmetry may exist and persist in many Internet advertising markets, especially with the new hypertargeted advertising. In traditional "share-of-voice" advertising, a publisher sells each advertiser a bulk share of ad requests. With hypertargeting, a publisher maximizes its revenue by selecting from a large number of advertisers to custom fit each ad request. The proliferation of hypertargeting increases the chances of a publisher dealing with unfamiliar advertisers while simultaneously facilitating many "long-tail" publishers to enter the advertising ecosystem. Indeed, currently advertisers often spread their ads across many Internet sites to reach the most valuable customers, understanding that their ads may appear on any number of these sites at any given moment. Responding to such a need, publishers often reach a broad set of advertisers by joining large ad networks. In such cases, the interaction between any advertiser–publisher pair is likely to be transient. In addition, with the dramatic increase in the number of advertisers, advertisements, publishers, and ad slots, matching advertisers (advertisements) with publishers (ad slots) can be a complex problem. An advertiser that is a good fit (and consequently, of high quality) for a particular publisher might be a bad fit (and thus, of lower quality) for a different publisher. These trends have increased the degree

¹For example, Debojyotipal.com (a blog catering to website and blog owners) recommends that PPC is for newcomers with low and medium traffic and should not be used on websites with large amounts of traffic. In contrast, Vaughn's Summaries (<http://www.vaughns-1-pagers.com/internet/internet-ad-networks.htm>), a popular source on the subject, states that PPC is the best choice in almost all situations.

Table 1
PRICING SCHEMES USED BY SOME PUBLISHERS

Search Advertising		Display Advertising
PPI		Major news and media sites (e.g., <i>The New York Times</i> , CBSNews.com, TV.com), ValueClick, Tribal Fusion, Amazon Banner ads, Twitter Promoted Tweets (at launch only), Google Display Ads (until 2005), Tremor (until 2011), AOL advertising (until about 2004)
P4P (PPC, PPA, etc.)	Google, Bing, Yahoo!	Amazon Product Ads, Shopping.com Merchant Program, Bidvertiser, Twitter Promoted Tweets, Tremor (since 2011)
Hybrid (PPI + P4P)		Google Display Ads (since 2005), Microsoft Display Ads, Facebook, LinkedIn, Clicksor, AOL Advertising (since approximately 2004), Kontera, Infolinks, Epom Market, Adcash, CPX Interactive, Chitika

of information asymmetry between the two parties. Such information asymmetry is likely to persist because of the rapid growth of the Internet and the capricious nature of Internet advertising: audiences for specific websites change continually, and ad copy is frequently terminated or replaced.

Our analyses yield the following key insights. First, we find that the choice of a P4P pricing scheme (e.g., PPC) can serve as a differentiation strategy for high-quality publishers. In general, a PPI scheme requires advertisers to estimate performance outcomes, which is a function of publisher quality. Because advertisers might not always distinguish low- and high-quality publishers, a low-quality publisher prefers a PPI scheme because it can masquerade as a high-quality publisher (i.e., it can “pool” with higher-quality publishers). A PPC scheme, in contrast, compensates the publisher on the basis of measured rather than advertiser-estimated outcomes, thus preventing a low-quality publisher from pooling with a high-quality one. Thus, a high-quality publisher may offer PPC as a way of distinguishing itself from a low-quality publisher. However, the PPC scheme is not without a cost: because publishers must use estimated advertiser qualities for allocating PPC ads, a PPC scheme typically results in a reduction in allocative efficiency, in the sense that advertising slots might not be allocated to those that value them the most. In some cases, the revenue loss that results from this allocative inefficiency outweighs the benefits from differentiation, and a high-quality publisher is better off choosing PPI. This insight offers several comparative static predictions: a high-quality publisher prefers a PPC scheme when there is a high degree of information asymmetry about publisher quality and when other high-quality publishers are rare—in which case, pooling with low-quality publishers would be very costly for a high-quality publisher. As a high-quality publisher improves at estimating advertiser quality, PPC becomes a more attractive strategy. This may explain why Twitter used PPI for its promoted tweets initially and then switched to P4P schemes after the initial launch. The trade-off between the need to reveal superior quality and the need to minimize allocative inefficiency drives high-quality publishers’ choice of pricing schemes.

Our second insight pertains to the advantage of the hybrid scheme. When a hybrid pricing scheme (e.g., a PPI and PPC hybrid) is offered, high-quality publishers will prefer the hybrid scheme to the pure PPI scheme because the former can lead to similar allocative efficiencies as a pure PPI scheme due to advertiser self-selection (into different pricing schemes). However, the pooling costs of a hybrid scheme are lower for high-quality publishers because only advertisers that choose the PPI portion of the hybrid scheme are required to estimate a publisher’s quality and can cause publisher pooling. When a high-quality publisher chooses a hybrid pricing scheme, a low-quality publisher must also follow suit to avoid being identified as low quality, causing the hybrid scheme to supplant the PPI scheme in equilibrium. Such a hybrid scheme equilibrium benefits high-quality publishers. Moreover, because the hybrid scheme has an advantage over the pure PPI scheme, the increased use of hybrid schemes reduces high-quality publishers’ reliance on costly pure P4P schemes and promotes overall allocative efficiency.

This article makes several contributions. First, to the best of our knowledge, we are the first to offer theoretical expla-

nations for the existence of different pricing schemes and the increasing popularity of the hybrid scheme. Second, we are the first to endogenize the choice of pricing schemes within an auction framework. Our analyses on pricing schemes add to the research on advertising auctions. Moreover, because we model the publisher as a mechanism designer with private information, we also contribute to the literature on mechanism design by informed principals (Julien and Mariotti 2006; Maskin and Tirole 1992). Third, we demonstrate a unique trade-off between allocative efficiency and quality separation as a driving mechanism for pricing scheme choices.

We organize the rest of the article as follows: In the next section, we review related research. Then, the following two sections outline the model and derive the main results. The final section concludes with a discussion of the implications of our findings.

RELATED RESEARCH

This article is situated in several streams of research. First, it is related to a few existing studies of pricing schemes for Internet advertising. Libai, Biyalogorsky, and Gerstner (2003) examine the coexistence of pay-per-lead and PPS schemes in affiliate marketing. Hu (2004) and Hu, Shin, and Tang (2010) study pricing schemes as an optimal contract problem between one publisher and one advertiser and find that Internet advertising pricing should include appropriate performance-based elements to provide incentives for two parties to improve advertisement performance. Sundararajan (2003) examines pricing schemes as an optimal two-part tariff problem. He shows that because advertisers are risk averse, P4P pricing is always optimal, even when publishers are constrained to offer PPI in parallel. Our research is markedly different from previous studies in that we examine pricing schemes in advertising auctions, which are increasingly common. We bring a new and important dimension (i.e., allocative efficiency) to the issue of pricing schemes. Most recently, Dellarocas (2012) has offered a compelling argument for the suboptimality of P4P schemes. His model links product pricing and advertising pricing schemes and shows that performance-based pricing creates distortion in product prices and reduces publisher revenue and social surplus. We focus on a distinct and important reason for the suboptimality of P4P schemes: the resulting allocative inefficiencies. We also analyze hybrid pricing schemes that are not discussed in the aforementioned articles.

Our research is also related to an increasing literature stream on Internet advertising auctions. Edelman, Ostrovsky, and Schwarz (2007) and Varian (2007) analyze the equilibria in search advertising auctions as “generalized second-price” auctions. Several articles have examined advertising auction design (Chen, Liu, and Whinston 2009; Feng 2008; Katona and Zhu 2012; Liu, Chen, and Whinston 2010) and its interaction with consumer search (Athey and Ellison 2011; Xu, Chen, and Whinston 2010, 2011). In addition, empirical research on advertising auctions has increased steadily (Agarwal, Hosanagar, and Smith 2011; Edelman and Ostrovsky 2007; Zhang and Feng 2011). Several authors (Animesh, Ramachandran, and Viswanathan 2009; Rutz and Bucklin 2007; Yao and Mela 2011) have examined the relationship between an advertiser’s bid and outcomes such as rank, click-through rates (CTRs), and conversions,

whereas others (Animesh, Ramachandran, and Viswanathan 2009; Ghose and Yang 2009; Jeziorski and Segal 2009) have focused on consumer behavior. Extant research on advertising auctions has largely taken the pricing scheme for granted, whereas we study pricing schemes as an endogenous choice.

A few recent articles have analyzed Internet advertising auctions with the hybrid pricing scheme (or “hybrid auctions”). Zhu and Wilbur (2010) study the equilibrium bidding in hybrid advertising auctions under the assumption that advertisers can choose their qualities after slots are assigned. They show that hybrid auctions can achieve the same outcomes as PPC auctions provided that publishers form rational expectations about advertisers’ quality conditional on their choice of pricing schemes. Edelman and Lee (2008) compare PPC and PPS auctions with a PPC/PPS hybrid auction. They show that hybrid auctions may produce as much revenue as unweighted PPC and PPS auctions, but it is unclear what their conclusions would be if the PPC and PPS auctions were weighted as they are in practice. Goel and Munagala (2009) analyze a different hybrid auction in which advertisers submit a PPI bid and a PPC bid for the publisher to choose. The existing results assume the hybrid scheme to be an exogenous choice and do not explain why advertisers choose hybrid auctions in the first place.

Because we model the publisher as a mechanism designer with private information, our model falls into the literature of mechanism design by informed principals. Several researchers (Jullien and Mariotti 2006; Maskin and Tirole 1992) have laid the ground work in this domain and have investigated more general settings than ours. We contribute to this literature stream by analyzing pricing schemes as a novel design dimension and creating a distinct model with a multiplicative payoff structure.

The issue of pricing schemes for Internet advertising is broadly related to research on the optimal pricing for information goods (Choudhary 2010; Jain and Kannan 2002; Sundararajan 2004). Although parallels exist between usage-based pricing of information goods and P4P in Internet advertising, there is a fundamental distinction: consumption of information goods is nonexclusive, whereas Internet advertising slots are exclusive resources. Thus, the impact of pricing schemes on allocative efficiency must be taken into account.

THE MODEL

We consider a one-shot game with a single publisher and n advertisers, all risk neutral. The publisher is endowed with a single impression, which it allocates using an auction. In practice, publishers may offer multiple slots at once. We discuss a multislot extension in the Web Appendix, in which we show that the main intuition of this article can be extended to a multislot case. We use the term “publisher” to refer to either the owner of the impression (e.g., *The New York Times*) or an intermediary that sells impressions on behalf of its owners. Similarly, an advertiser is interpreted as a merchant or its agent.

We use clicks as an illustrative performance metric. We assume each advertiser has a private valuation per click (valuation for short), denoted as $v \in [0, 1]$. Advertisers’ valuations are independently and identically distributed

according to a distribution $F(v)$, which has a strictly positive and differentiable density function $f(v)$. The distribution function satisfies an increasing hazard rate condition (i.e., $f(v)/[1 - F(v)]$) is monotonically increasing in v . This condition is satisfied by common distributions, such as uniform, normal, and logistic distributions, and is in line with prior research (Jullien and Mariotti 2006; Liu, Chen, and Whinston 2010).

An advertisement’s click performance is measured by CTR, which is the probability that a user will click on the advertisement. The CTR of an advertisement is subject to uncertainty but is also a function of the advertiser and the publisher’s characteristics. We capture an advertiser’s contribution to CTR as advertiser’s quality, a , and a publisher’s contribution to CTR as the publisher’s quality, b , and assume that²

$$(1) \quad \text{CTR} = \text{advertiser's quality } a \times \text{publisher's quality } b.$$

An advertiser’s quality captures the advertiser’s contribution to the click performance. It may reflect the appeal of the advertiser’s product or service. For example, a digital SLR camera ad by Nikon may attract more clicks than an ad by a lesser-known brand. An advertiser’s quality may also reflect the advertiser’s expertise in crafting highly effective ads.

Publisher’s quality b captures a publisher’s contribution to the ad performance. In traditional “share-of-voice” advertising, in which an ad is shown to a share of all visitors to a website, publisher’s quality depends mainly on the quality of the raw web traffic. Thus, a highly reputable photography website is more effective at promoting digital cameras than are personal blogs because, on average, the former attracts more legitimate shoppers. However, in the new targeted advertising, publishers treat each visitor as a unique marketing opportunity and selectively target him or her with the best-fitting ads. In these cases, average quality of raw traffic still matters, but so does targeting effectiveness. For example, a high-quality publisher may predict with reasonable accuracy whether a visitor to a wedding photography blog is interested in digital cameras (as opposed to, say, wedding dresses) and then choose appropriate ads to display. By targeting Internet users with the most relevant ads, the publisher effectively offers selected impressions of high quality to advertisers, even if those impressions originate from sources that are low in average quality (e.g., a personal blog). We note that the concept of “high quality” is relative to the specific market. For example, a high-quality display advertising publisher is defined relative to other display advertising publishers rather than to, say, search advertising publishers.

For simplicity, we assume that each advertiser’s quality is a random draw from two levels, a_l and a_h ($a_l < a_h$). We use $x \in \{l, h\}$ to denote an advertiser’s quality type. Similarly, the publisher’s quality is a random draw from two levels, b_l and b_h ($b_l < b_h$), and $y \in \{l, h\}$ denotes the publisher’s quality type. The probabilities of drawing a_h and b_h are α and β , respectively.

²In reality, the measured CTR is also subject to measurement errors and random noises. We assume that such errors and noises cancel out and thus omit them from further consideration in a risk-neutral model framework.

We assume that the publisher does not know an advertiser's true quality but observes a signal $\tilde{x} \in \{\tilde{l}, \tilde{h}\}$ about the advertiser's quality. The probability of observing signal $\tilde{l}$ ($\tilde{h}$) while the true quality type is h (l) is

$$(2) \quad P(\tilde{x} = \tilde{l} | x = h) = P(\tilde{x} = \tilde{h} | x = l) = \gamma_A.$$

We interpret the parameter γ_A as the probability of misclassifying an advertiser's quality. As γ_A increases, the quality signal $\tilde{x}$ becomes less informative. We assume that $\gamma_A \leq .5$ because if $\gamma_A > .5$, we can simply swap $\tilde{h}$ and $\tilde{l}$. When $\gamma_A = 0$, there is no information asymmetry about advertiser quality; when $\gamma_A = .5$, the quality signals are pure noise.

We similarly assume that advertisers do not know the publisher's true quality but observe a signal $\tilde{y} \in \{\tilde{l}, \tilde{h}\}$; we let γ_P , $0 \leq \gamma_P \leq .5$ denote the probability of misclassifying a publisher:

$$(3) \quad P(\tilde{y} = \tilde{l} | y = h) = P(\tilde{y} = \tilde{h} | y = l) = \gamma_P.$$

We use parameters γ_A and γ_P to capture various degrees of information asymmetry for different market conditions. For example, a large γ_A and a small γ_P capture the case in which the publisher knows very little about advertiser qualities but advertisers know a great deal about the publisher's quality. Traditional display advertising at premium websites often falls into this category. A small γ_A and a large γ_P may capture a market of targeted advertising in which the publisher has good information on advertiser quality, but the publisher's quality is difficult to determine.

Information asymmetry may persist in a wide range of market conditions for several reasons. As we mentioned previously, the large number of advertiser–publisher pairings and the transient nature of website audience and ad copies make it difficult for advertisers and publishers to learn about each other's qualities. In practice, predicting ad performance is rather challenging because performance events (e.g., clicks, sales, calls) are typically sparse (Rutz and Bucklin 2007). For example, CTRs typically range from 1% to 3% for sponsored search and .02% to .05% for display advertising (Raehsler 2014). Many factors—such as search engine optimization, click fraud, idiosyncratic marketing promotions, and social media trends—can influence ad performance and hinder effective discovery of underlying qualities.

Expected Advertiser and Publisher Quality

The publisher can form expectations about advertiser qualities. We denote $\hat{a}_l$ ($\hat{a}_h$) as the expected quality of an advertiser conditional on the quality signal $\tilde{l}$ ($\tilde{h}$). The expected advertiser quality can be calculated using a Bayesian rule (see the Web Appendix). Similarly, advertisers can form expectations about the publisher's quality conditional on signal $\tilde{y}$ and the publisher's pricing scheme choice m (which may convey additional information). We denote $\hat{b}_{\tilde{y}}^m$ as the expected quality of a publisher that chooses pricing scheme m and has a quality signal $\tilde{y}$ (see the Web Appendix).

We denote $\mu(y|m)$, advertisers' belief about the probability of the publisher being y -type, conditional on the pricing scheme m (but before learning the quality signal $\tilde{y}$). The expected publisher quality $\hat{b}_{\tilde{y}}^m$ my is a function of advertisers' beliefs about the publisher's true type.

Auction Rules and Pricing Schemes

More than 70% of publishers use auctions to allocate Internet advertisements (Kantrowitz 2013). In keeping with practice, we assume that the publisher uses an auction to allocate the impression among advertisers. The publisher can choose from three stylized pricing schemes: PPI (I), PPC (C), and a PPI/PPC hybrid (H). We can easily reinterpret our results as a choice between PPI, PPS, and a PPI/PPS hybrid or between PPC, PPS, and a PPC/PPS hybrid. When a PPI or PPC scheme is used, advertisers place PPI or PPC bids respectively. When a hybrid scheme is used, advertisers can choose between a PPI bid and a PPC bid. Regardless of the pricing scheme, the auction follows two general rules. First, advertisers are ranked, and the highest-ranked advertiser gets the impression (the allocation rule). Second, the winner of the auction pays the least amount to keep its winning position (the payment rule). The specific allocation and payment rules are as follows:

- PPI: Advertisers place and are ranked by PPI bids. The winner pays the second-highest PPI bid.
- PPC: Advertisers place PPC bids and are ranked by PPC bid $\times$ expected advertiser quality (i.e., weighted PPC bids). The winner pays the lowest PPC price to keep its winning position.
- Hybrid: Advertisers place either PPI or PPC bids. Those who place PPI bids are ranked by PPI bids, and those who place PPC bids are ranked by weighted PPC bids (i.e., PPC bid $\times$ expected advertiser quality $\times$ expected publisher quality). The winner pays the lowest PPI or PPC price (depending on the bid format) to keep its winning position.

These allocation rules, which can be regarded as ranking by expected revenue, are in line with practice and prior literature on PPC and hybrid auctions (Lahaie and Pennock 2007; Liu, Chen, and Whinston 2010; Zhu and Wilbur 2010). Intuitively, publishers make the best use of their advertising resources by awarding a slot to the advertiser that generates the highest expected revenue. Google, for example, ranks different bids on the basis of expected revenue per impression (Google 2014). The payment rules for the three auction formats can be regarded as generalized second price (Edelman, Ostrovsky, and Schwarz 2007; Li, Liu, and Liu 2013) and are in line with practice. Other than the differences in allocation and payment rules, we maintain the same assumptions across three auction formats, including our informational assumption that the publisher and the advertisers only observe a quality signal about the other party. We illustrate the auction formats with the following example.

Example 1. Consider three advertisers with per-click valuations $v_1 = 5$, $v_2 = 3$, and $v_3 = 2$ and qualities $a_1 = a_3 = .02$, and $a_2 = .04$, respectively. If the publisher runs a PPI auction and the PPI bids are .05, .06, and .02, respectively, Advertiser 2 wins and pays .05. If the publisher runs a PPC auction, the PPC bids are 5, 3, and 2, and the expected advertiser qualities are .022, .032, and .022, respectively, Advertiser 1 wins with a weighted PPC of .11, paying .096/.022 = 4.36 per click. If the publisher runs a hybrid auction, Advertiser 2 places a PPI bid of .06 and Advertisers 1 and 3 place PPC bids of 5 and 2, respectively. Thus, the expected publisher quality is .5, and the weighting factor for Adver-

tisers 1 and 3 is .02. Advertiser 2 wins, paying a PPI price of .05 ($= .02 \times .5 \times 5$).

The game timeline. The game proceeds as follows (see Figure 1). The quality types of the advertisers and the publisher are drawn randomly. Next, the publisher chooses and announces a pricing scheme. The publisher then receives a signal about each advertiser's quality, and advertisers receive the same signal about the publisher. Next, the publisher calculates weighting factors for PPC bids. Advertisers learn the pricing scheme and their weighting factors before simultaneously placing their bids. A winner is then selected to fill the slot and pays according to the auction rules.

This game has two stages: in the second stage, a bidding game is played among advertisers. In the first stage, a game is played between the publisher and the advertisers in which the publisher strategically chooses a pricing scheme to influence the second-stage outcomes. We analyze the game using backward induction: we first solve the bidding equilibrium for each auction format and then solve the equilibrium pricing scheme.

We assume that the publisher maximizes its expected revenue by choosing a pricing scheme $m \in \{I, C, H\}$. Let $\sigma(m|y)$ denote a y -type publisher's probability of choosing pricing scheme m , where $\sigma(m|y) \in [0, 1]$ and $\sum_m \sigma(m|y) = 1$. We refer to $\sigma(m|y)$ as a publisher's strategy profile. When both publisher types play a pure strategy (i.e., choosing a payment scheme with probability 1), we simply denote the strategy profile as (m_l, m_h) , where m_l is the l -type's pricing scheme and m_h is the h -type's. For example, (I, C) represents that an l -type publisher chooses PPI and an h -type publisher chooses PPC.

Because the signal about the publisher could be either $\tilde{h}$ or $\tilde{l}$, it is useful to define the average expected quality (across two types of quality signals) for a y -type publisher as follows:

$$(4) \quad \hat{B}_y^m \triangleq \sum_{\tilde{y} \in \{\tilde{l}, \tilde{h}\}} P(\tilde{y}|y) \hat{b}_{\tilde{y}}^m.$$

When two publisher types choose a pricing scheme m with equal probability (say both choose PPI with probability 1), the pricing scheme provides no information about the publisher's true quality. We denote the average expected quality in this special case as $\bar{B}_y$.

We summarize our notations in Table 2. For notational simplicity, we use $a, \hat{a}, b, \hat{b}$, and $\bar{B}$ as shorthand symbols for $a_x, \hat{a}_x, b_y, \hat{b}_y^m$, and $\bar{B}_y$.

Table 2
NOTATIONS

Notation	Interpretation
v	Advertiser's valuation per click
a, b	Advertiser/publisher quality, $a \in \{a_l, a_h\}$, $b \in \{b_l, b_h\}$
x, y	Advertiser/publisher quality type, $x, y \in \{l, h\}$
α, β	Probability of an advertiser/publisher being a high-quality type
γ_A, γ_P	Rates of misclassifying an advertiser/publisher
$\tilde{x}, \tilde{y}$	Advertiser/publisher quality signal, $\tilde{x}, \tilde{y} \in \{\tilde{l}, \tilde{h}\}$
m	Publisher's pricing scheme choice, $m \in \{I, C, H\}$
$\hat{a}_x, \hat{b}_y^m$	Expected quality for an $\tilde{x}$ -signal advertiser/ $\tilde{y}$ -signal publisher
$\bar{B}_y^m$	Average expected quality for a y -type publisher
$\bar{B}_y$	Average expected publisher quality when two publisher types choose a pricing scheme with equal probability (say, when both choose PPI with probability 1)
$\mu(y m)$	Advertisers' belief about a publisher being y -type given its pricing choice m .
$\sigma(m y)$	Publisher's strategy profile—the probability of a y -type publisher choosing m
φ^m	Advertiser's equilibrium probability of winning under pricing scheme m
π^m	Expected publisher revenue under pricing scheme m

AUCTION OUTCOMES UNDER DIFFERENT PRICING SCHEMES

In this section, we discuss how pricing scheme choices can affect equilibrium bidding and auction revenues. We begin by analyzing the expected quality.

Preliminaries

Because of information asymmetry, the following relationship holds between expected and true advertiser qualities (see the Web Appendix):

$$(5) \quad a_l \leq \hat{a}_l \leq \hat{a}_h \leq a_h.$$


Thus, when expected qualities are used as weighting factors, a low-quality advertiser (which may receive a weighting factor of $\hat{a}_l$ or $\hat{a}_h$) benefits from information asymmetry, and the opposite is true for a high-quality advertiser.

Example 2: Expected advertiser quality. Continue with Example 1. Suppose that $\alpha = .7$ and $\gamma_A = .2$. The expected qualities of $\tilde{h}$ -signal and $\tilde{l}$ -signal advertisers are $\hat{a}_h = .32$ and $\hat{a}_l = .22$, respectively (recall that $a_h = .4$ and $a_l = .2$). For the publisher quality, we similarly have (see the Web Appendix)

$$(6) \quad b_l \leq \hat{B}_l^m \leq \hat{B}_h^m \leq b_h.$$

Thus, with information asymmetry, the average expected quality of a high-quality publisher is less than its true quality, and the opposite is true for a low-quality publisher. We use the example described in the next subsection to illustrate this relationship.

Figure 1
GAME TIMELINE



Example 3: Average expected publisher quality. Suppose $b_h = .8$, $b_l = .4$, $\beta = .6$, and $\gamma_p = .2$. Suppose also the advertisers' belief is $\mu(h|m) = .3$. The expected qualities of $\bar{l}$ - and $\bar{h}$ -signal publishers are $\hat{b}_l^m = .44$ and $\hat{b}_h^m = .65$, respectively. The average expected quality of $\bar{l}$ - and $\bar{h}$ -type publishers is $\hat{B}_l^m = .48$ and $\hat{B}_h^m = .61$, respectively. If, instead, the advertisers' belief is $\mu(h|m) = .6$, we have $\hat{b}_l^m = .51$, $\hat{b}_h^m = .74$, $\hat{B}_l^m = .56$, and $\hat{B}_h^m = .70$. Intuitively, Example 3 shows that the higher the belief that the publisher is an $\bar{h}$ -type, the higher the expected publisher qualities.

In the next three subsections, we derive the equilibrium auction revenues for the PPI, PPC, and hybrid schemes. We focus on how advertisers are ranked under each pricing scheme because, as per the well-known "revenue equivalence theorem" (Myerson 1981), the expected revenue of an auction mechanism is determined only by the mechanism's allocation rule. In other words, no matter how different two auction mechanisms may seem, if they allocate the objects in the same way, they will generate the same amount of expected revenue. Indeed, the revenue equivalence theorem enables us to directly derive the expected auction revenue from bidders' equilibrium winning probabilities (which reflect the allocation rule) without explicitly evaluating the equilibrium bidding functions (Myerson 1981). Thus, for each auction format, we first analyze the ranking rule and equilibrium winning probability and then present the expected auction revenue. All technical proofs appear in Appendices A and B.

The PPI Auction

An advertiser's true valuation per impression is va_b , but because the true publisher quality b is unknown, its valuation is $va\hat{b}$, where $\hat{b}$ is the expected publisher quality. As long as equilibrium bids increase with valuations,³ advertisers are ranked by their expected valuation-per-impression $va\hat{b}$ or, equivalently, by va (recall that $\hat{b}$ is the same for all advertisers).

It is important to note that because advertisers are ranked by va , the impression will be allocated to the advertiser that has the highest valuation. Thus, the PPI auction is efficient.

An advertiser wins if it has the highest va among all advertisers. Thus, its equilibrium winning probability, denoted as $\phi^I(v, a)$, can be calculated as follows:

$$(7) \quad \phi^I(v, a) = [\sum_{x \in \{l, h\}} P(x) F(va/a_x)]^{n-1},$$

where the term in brackets represents the probability that the advertiser (v, a) beats any other advertiser. Intuitively, if the opponent is an $\bar{l}$ -type, which occurs with probability $P(l) = 1 - \alpha$, it must have a valuation less than va/a_l . Similarly, if the opponent is an $\bar{h}$ -type, which occurs with probability $P(h) = \alpha$, it must have a valuation less than va/a_h .

With the equilibrium winning probability, we can directly derive the expected revenue:

Proposition 1: Given a publisher's strategy profile σ , a y -type publisher's expected revenue under the PPI auction is

$$(8) \quad \pi_y^I = B_y^I \pi_{base}^I,$$

³It is simple to show that this is the case. Note that a PPI auction, as modeled, is essentially a standard second-price auction, which has a truth-telling equilibrium.

where π_{base}^I , termed as the base revenue of the PPI auction, is given by

$$\pi_{base}^I \triangleq n \sum_{x \in \{l, h\}} P(x) \left[a_x \int_0^1 \phi^I(v, a_x) J(v) f(v) dv \right]$$

and $J(v) \triangleq v - 1 - F(v)f(v)$.

Proposition 1 suggests that the PPI auction revenue has a base revenue component π_{base}^I and a publisher quality component $\hat{B}_y^I$. The publisher quality component $\hat{B}_y^I$ is average expected publisher quality because advertisers use expected publisher quality in their PPI bids. The base revenue component is the sum of revenues from both high- and low-quality advertisers, which have to do with their winning probabilities.

Because advertisers are ranked by their true qualities, the base revenue component is not affected by information asymmetry. However, the publisher quality component is based on expected values; thus, when both high- and low-quality publishers choose PPI, the two types pool in the sense that advertisers cannot perfectly distinguish them. As a result, a high-quality publisher obtains a revenue of $\hat{B}_h^I \pi_{base}^I$ and a low-quality publisher obtains a revenue of $\hat{B}_l^I \pi_{base}^I$. Compared with complete-information revenues $b_h \pi_{base}^I$ and $b_l \pi_{base}^I$, we note that a high-quality publisher suffers from pooling with a low-quality publisher (because $\hat{B}_h^I \leq b_h$), whereas a low-quality publisher benefits from pooling with a high-quality one (because $b_l \leq \hat{B}_l^I$). Intuitively, because of the information asymmetry regarding the publisher's quality, advertisers, on average, undervalue a high-quality publisher (through their PPI bids) and overvalue a low-quality publisher.

The PPC Auction

Under the PPC auction, advertisers place PPC bids, which are weighted by their expected advertiser quality, $\hat{a}$. We establish that an advertiser $(v, \hat{a})$ wins a weighted PPC auction if its $v\hat{a}$ is the highest (see Appendix B). Thus, the advertiser's equilibrium winning probability is as follows:

$$(9) \quad \phi^C(v, \hat{a}) = [\sum_{\tilde{x} \in \{\bar{l}, \bar{h}\}} P(\tilde{x}) F(v\hat{a}/\hat{a}_{\tilde{x}})]^{n-1}.$$

Because advertisers are ranked by expected qualities, inefficient allocations may occur.

It is important to note that an advertiser's optimal PPC bid is not a function of expected publisher quality. This is because the publisher's quality (which affects click performance) is merely a scale factor in an advertiser's total payoff and therefore does not affect the advertiser's bid optimization problem.

Using the same technique, we derive the expected revenue for the PPC auction as follows:

Proposition 2: Given a strategy profile σ , a y -type publisher's expected revenue under the PPC auction is

$$(10) \quad \pi_y^C = b_y \pi_{base}^C$$

$$\text{where } \pi_{base}^C \triangleq n \sum_{\tilde{x} \in \{\bar{l}, \bar{h}\}} P(\tilde{x}) \left[\hat{a}_{\tilde{x}} \int_0^1 \phi^C(v, \hat{a}_{\tilde{x}}) J(v) f(v) dv \right].$$

The PPC revenue also has a base revenue component π_{base}^C and a publisher quality component b_y . The base revenue is affected by information asymmetry because the weighting scheme is based on expected advertiser qualities. Information asymmetry does not affect the publisher quality component b_y , because the publisher's quality does not affect the optimal PPC bids and because the publisher is compensated by measured quality. In other words, a PPC publisher is not affected by information asymmetry regarding its quality, but it incurs a cost of inefficient allocation when estimated advertiser qualities, as weighting factors, deviate from true advertiser qualities.

The Hybrid Auction

As we described in the previous section, the weighting factor for PPC bidders consists of expected advertiser quality and expected publisher quality. An efficient ranking demands that the publisher treat as equals a PPI bidder and a PPC bidder that have the same valuation and quality. Because the PPI bidder uses the expected publisher quality in its bid, the publisher should also use the expected (rather than true) publisher quality as a weighting factor for the PPC bid to maintain the efficient ranking. It is important to note that such a hybrid weighting scheme reveals no additional information about the true publisher quality.⁴ If all advertisers choose PPI (or PPC), the hybrid auction degenerates. We are interested in nondegenerate hybrid auctions in which there are mixes of PPI and PPC bids. When advertisers self-select into different pricing schemes, an advertiser's choice of pricing scheme may reveal additional information about its true quality. Thus, the expected advertiser quality is also a function of an advertiser's pricing scheme choice, which suggests that we must jointly determine weighting factors and advertisers' pricing scheme choices. Our next result describes a unique weighting scheme and advertiser price scheme choices for the nondegenerate hybrid auction.

Lemma 1: If the weighting factor for all PPC bidders is $a_1\hat{b}$, all high-quality advertisers prefer PPI and all low-quality ones weakly prefer PPC. This is the only nondegenerate weighting scheme that meets the hybrid auction specification in the previous section.

Intuitively, an advertiser with true quality a and a PPC weighting factor w would prefer a PPI bid when $a\hat{b} > w$ and a PPC bid when $a\hat{b} < w$, and it would be indifferent when $a\hat{b} = w$ (Appendix B). Thus, if $a_1\hat{b}$ is the PPC weighting factor, all high-quality advertisers will choose PPI and low-quality advertisers will choose PPC, resulting in a nondegenerate hybrid auction. Given the bidding behavior, the expected quality of PPC bidders is indeed a_1 , and thus the weighting scheme meets the description in the previous section.

Lemma 1 suggests that a hybrid publisher would infer that PPC bidders consist of low-quality advertisers only and would thus assign a low weighting factor. Zhu and Wilbur (2010) obtain a similar insight for a setting in which adver-

tisers can choose their quality after the allocation. The weighting scheme in Lemma 1 uniquely separates high- and low-quality advertisers while meeting the "ranking-by-expected-revenue" criterion. Hereinafter, we use the hybrid weighting scheme described in Lemma 1.

With the hybrid weighting scheme in Lemma 1, all advertisers are indeed ranked by $va\hat{b}$. Thus, the hybrid auction is efficient and the associated equilibrium winning probability, $\phi^H(v, a)$, is the same as the PPI auction $\phi^I(v, a)$. The publisher's expected revenue, however, is different from the PPI revenue, as the next result illustrates.

Proposition 3: Given a strategy profile σ , a y -type publisher's expected revenue under a hybrid auction is as follows:

$$(11) \quad \pi_y^H = \hat{B}_y \pi_{h, \text{base}}^H + b_y \pi_{l, \text{base}}^H,$$

where $\pi_{h, \text{base}}^H \triangleq n\alpha a_h \int_0^1 \phi^H(v, a_h) J(v) f(v) dv$ and $\pi_{l, \text{base}}^H \triangleq n(1 - \alpha) a_l \int_0^1 \phi^H(v, a_l) J(v) f(v) dv$ are the base revenues from high-quality (PPI) and low-quality (PPC) advertisers, respectively. Moreover, ceteris paribus, a high-quality publisher would prefer the hybrid scheme to PPI, and the opposite is true for a low-quality publisher.

Comparing the base revenues for PPI and hybrid schemes, we find that

$$(12) \quad \pi_{\text{base}}^I = \pi_{h, \text{base}}^H + \pi_{l, \text{base}}^H.$$

Again, this is because PPI and hybrid auctions rank advertisers the same way. The two auctions differ only in the publisher quality component: under the hybrid scheme, the publisher quality components for the PPC and PPI revenues are b_y and $\hat{B}_y$, respectively. Thus, if both publisher types choose the hybrid scheme, they will partially pool in the sense that information asymmetry regarding publisher quality only affects PPI bids (and not PPC bids). In contrast, the two types fully pool if they both choose the PPI scheme. Because a high-quality publisher suffers from pooling with a low-quality publisher, the high-quality publisher would prefer, ceteris paribus, to pool with a low-quality publisher under a hybrid scheme rather than a PPI scheme.

EQUILIBRIUM PRICING SCHEME

We now turn to equilibrium pricing schemes. When the publisher holds private information relevant to advertisers' payoffs, its pricing scheme choice can signal its private information. Thus, we analyze the pricing scheme choice as a game of signaling between the publisher and the advertisers. However, unlike the standard game of signaling, publisher pooling is not possible under the PPC scheme because PPC bidding is not sensitive to the publisher's quality. In other words, unlike the PPI scheme, even if both publisher types choose PPC, they are still considered separated in the sense that their payoffs are based on their respective true quality.

The equilibrium of the game is trivial when the PPC base revenue exceeds the PPI base revenue, as the following lemma shows:

Lemma 2: If $\pi_{\text{base}}^I < \pi_{\text{base}}^C$, (C, C) is the only pure-strategy equilibrium.

Intuitively, when the PPC auction also has higher base revenue, a high-quality publisher clearly prefers PPC. A

⁴Two main reasons are behind this assumption. First, a high-quality publisher may not be able to communicate its true type credibly using the weighting factors, because a low-quality publisher can mimic the high-quality publisher using the same weighting factors. Second, in practice, publishers do not report precise PPC weighting factors, so these weighting factors are unlikely to be used as signals for publisher quality.

low-quality publisher that is unable to masquerade as a high-quality publisher under PPC would still choose PPC because of its higher base revenue.

Lemma 2 depicts a case in which both publisher types strongly prefer PPC, which seems inconsistent with practice. The opposite of the condition for Lemma 2 is more likely. Recall that PPI is more efficient than PPC, meaning that a PPI auction generates more total surplus (i.e., realized valuation) than a PPC auction. Under modest conditions, a PPI auction also generates higher base revenue than a PPC auction.⁵ For these reasons, our subsequent analysis focuses on the more interesting case in which

$$(13) \quad \pi_{\text{base}}^I \geq \pi_{\text{base}}^C.$$

Following Athey and Ellison (2011) and Wilbur and Zhu (2009), we adopt the perfect Bayesian equilibrium (PBE) concept, which requires that advertisers' belief be determined by the Bayesian rule whenever possible. We focus on pure-strategy equilibrium.

No Information Asymmetry or One-Sided Information Asymmetry

We first examine benchmark cases in which $\gamma_A = 0$, $\gamma_P = 0$, or both. If $\gamma_P = 0$ but $\gamma_A > 0$, there is no information asymmetry about the publisher quality: the publisher quality components are identical (i.e., $b_l = \hat{b}_l^m = \hat{B}_l^m$ and $b_h = \hat{b}_h^m = \hat{B}_h^m$). However, because of asymmetric information about advertiser qualities, PPC suffers from revenue losses due to allocative inefficiency. Thus, both publisher types prefer PPI to PPC and hybrid schemes, and PPI prevails in this setting. This prediction is consistent with the observation that PPI is dominant in traditional share-of-voice display advertising at premium websites, where there is little uncertainty about the publisher's quality.

If $\gamma_A = 0$ and $\gamma_P > 0$, there is no information asymmetry about advertiser qualities (thus, $\pi_{\text{base}}^I = \pi_{\text{base}}^C = \pi_{\text{base}}^H$), but there is uncertainty about publisher quality. To avoid costly pooling, a high-quality publisher strictly prefers PPC to a PPI or hybrid scheme. This may explain the dominance of PPC in sponsored search, whereby advertiser qualities can be estimated with reasonable accuracy, but qualities of two sponsored search publishers are difficult to compare because of the highly complex nature of their targeting algorithms and the secrecy and continuous evolution of such algorithms.

When $\gamma_A = 0$ and $\gamma_P = 0$, both publisher and advertiser qualities are public knowledge. There is no distortion in either the publisher quality component or the base revenue component. As a result, the three auction formats are equivalent and the publisher is indifferent. This scenario is not very realistic, but it provides a benchmark for other scenarios.

⁵The revenue ranking could reverse only when a weighted PPC scheme extracts enough additional revenue from incentivizing high-quality advertisers (by distorting the weighting scheme in favor of low-quality ones) to offset the revenue loss that results from inefficient allocations (Liu and Chen 2006). Because the incentive effect comes at the cost of reducing total surplus, it may only cause revenue reversal when there is a large enough (but not too large) distortion in the weighting scheme. Moreover, the likelihood of the revenue reversal diminishes quickly as the number of advertisers increases because the presence of more advertisers intensifies competition, and the additional value of distorting the weighting scheme is low (Liu and Chen 2006).

We summarize the benchmark cases in the following proposition:

Proposition 4: If $\gamma_P = 0$ and $\gamma_A > 0$ and the condition in Equation 13 holds, (I, I) is the only equilibrium. If $\gamma_A = 0$ and $\gamma_P > 0$, (I, C), (H, C), and (C, C) are equilibrium strategy profiles. If $\gamma_A = 0$ and $\gamma_P = 0$, any strategic profile is an equilibrium.

Two-Sided Information Asymmetry Without the Hybrid Scheme

A more realistic case is two-sided information asymmetry (i.e., $\gamma_A = 0$ and $\gamma_P = 0$). We examine two subcases in which the publisher can choose from (1) only PPI and PPC and (2) PPI, PPC, and hybrid pricing. As with other signaling games, there are multiple equilibria in this game that prevent us from prescribing exactly what will happen. Therefore, we employ a few equilibrium refinement strategies to eliminate equilibria that are less plausible:

Assumption 1: A weakly dominated strategy is not played.

Assumption 2: A Pareto-dominated equilibrium is not played.

The rationale for Assumption 2 is that weakly dominated strategies are imprudent (because the player may be worse off and can never be better off playing such strategies) and players should therefore avoid them. An equilibrium Pareto dominates another if all players are not worse off and at least one player is strictly better off in the former equilibrium. A Pareto-dominated equilibrium may be avoided through, for example, better pregame coordination among players (in our case, among publishers). The Pareto-dominance refinement helps simplify our presentation but is nonessential to our main findings.⁶

Proposition 5: Under Assumptions 1 and 2 and the condition in Equation 13, if

$$(14) \quad b_h \pi_{\text{base}}^C \leq \bar{B}_h \pi_{\text{base}}^I,$$

the equilibrium strategy profile is (I, I). Otherwise, the equilibrium strategy profile is (I, C).

Proposition 5 shows that with two-sided information asymmetry, a low-quality publisher chooses PPI, whereas a high-quality publisher may choose either PPC or PPI, depending on the condition in Equation 14. Borrowing the terminology of signaling games, we refer to (I, C) as a separating equilibrium and (I, I) as a pooling equilibrium. Intuitively, a high-quality publisher trades off between allocative inefficiency under PPC and pooling with a low-quality publisher under PPI. This trade-off depends on several underlying parameters, as summarized by the following corollary.

Corollary 1: A high-quality publisher's preference for (I, C) increases when (a) β (the probability of being a high-quality publisher) decreases, (b) b_h/b_l increases, (c) γ_P (the rate of misclassifying a publisher) increases, and (d) $\pi_{\text{base}}^I/\pi_{\text{base}}^C$ decreases.

⁶More specifically, as we show in the proof of Proposition 5, without the Pareto-dominant refinement, we would have a middle parameter range in which both (I, I) and (I, C) are equilibria. With the Pareto-dominant refinement, only (I, I) remains because it Pareto-dominates (I, C) in this parameter range. The fundamental trade-off does not change because of the Pareto-dominance refinement.

By Corollary 1a, in a market in which most publishers are low quality, a high-quality publisher has a strong incentive to separate itself by using PPC. The incentive to separate is strong when a high-quality publisher has superior quality to a low-quality one (Corollary 1b; e.g., when the publisher is a disruptive innovator), when advertisers are not well informed about a publisher's quality (Corollary 1c; e.g., in hypertargeted advertising and other new markets), and when the allocative inefficiency loss under PPC is small (Corollary 1d; e.g., when the publisher has good estimates of advertiser qualities).

Corollary 1's predictions are consistent with our anecdotal observations. For example, well-known traditional publishers such as CBSNews.com and TV.com choose PPI because there is little uncertainty about their quality but they face significant hurdles in estimating advertiser qualities. Yet new and innovative publishers with superior quality (e.g., Twitter) often use P4P schemes to differentiate themselves. In sponsored search advertising, for which there are much better estimates of advertisers' qualities but it is difficult to compare qualities of sponsored search publishers, publishers more often use PPC. Display advertising publishers, in contrast, tend to use PPI instead of PPC. One explanation for this is that clicks are fewer and noisier in display advertising, making it difficult to estimate advertiser qualities. Implementing PPC in such a case may result in revenue loss due to inefficient ranking of advertisers.

Figure 2 illustrates the publisher's equilibrium strategy profile as a function of misclassification probabilities γ_P . When γ_P is above a threshold, a high-quality publisher adopts PPC. A high-quality publisher adopts PPI when the probability of misclassifying advertisers γ_A is high (but not too high). The boundary between the pooling equilibrium (I, I) and the separating equilibrium (I, C) is not monotonic because at a high γ_A , the weighted PPC auction can extract more revenue from high-quality advertisers, and such additional revenue extraction may offset some of the revenue loss from allocative inefficiency.

Corollary 1 also suggests that the trade-off between PPC and PPI depends on the relative contribution of publisher

and advertiser to advertising performance. When there is a small difference in publisher qualities (i.e., b_h/b_l is small), it is more important to ensure efficient allocation of advertisers and, thus, PPI is more desirable. When there is a small difference in advertiser qualities (i.e., a_h/a_l is small), it is more important to avoid pooling with low-quality publishers and, thus, PPC is more desirable. The former situation could occur when all publishers in the domain adopt similar technologies or have similar Internet audiences. The latter situation could occur when the publisher carefully selects advertisers so that the variation of advertiser qualities is small.

Two-Sided Information Asymmetry with the Hybrid Scheme

When a publisher can choose from PPI, PPC, or a hybrid scheme, the number of pure-strategy profiles increases from four to nine. Again, when $\pi_{base}^I < \pi_{base}^C$, (C, C) is the only pure-strategy equilibrium. Therefore, we continue to focus on the case $\pi_{base}^I \geq \pi_{base}^C$.

In Appendix B, we show that (I, I), (I, C), (H, C), and (H, H) are all PBEs. However, a closer investigation suggests that not all four equilibria are plausible. Specifically, the equilibrium (I, I) requires advertisers to hold the belief that a publisher that deviates to an off-equilibrium scheme H is more likely to be low quality. However, a low-quality publisher is less likely to initiate such a deviation because it is worse off under (H, H) than under (I, I). To rule out such an implausible off-equilibrium belief, we introduce the *uncompromised equilibrium*:

Definition 1: An equilibrium is uncompromised if one assigns a probability of no less than p to player type t for an off-equilibrium action a , if t plays a with probability p in an alternative equilibrium in which t is better off.

In a signaling game, the uncompromised equilibrium requires signal recipients (advertisers) to interpret off-equilibrium actions as attempts to overturn the current equilibrium by sender (publisher) types that seek a better alternative for themselves. Given this interpretation, it is reasonable to assume that sender types that are better off under the alternative equilibrium are more likely to deviate than those that are worse off. Therefore, signal recipients should assign higher probabilities to sender types that are better off under the alternative equilibrium. That is, advertisers should believe that a deviating publisher is more likely to be better off by deviating to an off-equilibrium pricing scheme.


In a special case in which all sender types that play an off-equilibrium action a are better off in the alternative equilibrium, the uncompromised equilibrium requires us to assign probabilities to these sender types exactly according to the alternative equilibrium. This special case is the undefeated equilibrium defined by Mailath, Okuno-Fujiwara, and Postlewaite (1993). Therefore, the uncompromised equilibrium is a useful generalization of the undefeated equilibrium. The difference between the two concepts is that the uncompromised equilibrium also deals with cases (such as ours) in which an off-equilibrium action is played by both better-off and worse-off sender types in an alternative equilibrium.

With the uncompromised equilibrium concept, we propose the following:

Proposition 6: Under Assumptions 1 and 2 and the condition in Equation 13, if

Figure 2

PUBLISHER'S EQUILIBRIUM STRATEGY PROFILE WITHOUT THE HYBRID SCHEME



Notes: $n = 6$, $F \sim \text{uniform}[1, 2]$, $\alpha = .7$, $a_l = .1$, $a_h = .2$, $\beta = .7$, $b_l = .1$, $b_h = .12$.

$$(15) \quad b_h \pi_{base}^C \leq \bar{B}_h \pi_{h, base}^H + b_h \pi_{l, base}^H,$$

there is one uncompromised PBE: (H, H), with a belief that a PPI publisher is high quality with probability no more than μ_2 (defined in Appendix B). Otherwise, there are two uncompromised PBEs: (I, C), with a belief that a hybrid publisher is low quality, and (H, C), with a belief that a PPI publisher is low quality.

Proposition 6 suggests that the equilibrium can be either separating, in the form of (I, C) or (H, C), or pooling (H, H), depending on the condition in Equation 15. The separating equilibria (I, C) and (H, C) are equivalent for all parties because advertisers perfectly infer the publisher's quality type in each case and the hybrid scheme simply reduces to the PPI scheme (see also Proposition 3).

Remarkably, Proposition 6 suggests that (H, H) is the only pooling equilibrium. The intuition for Proposition 6 is as follows. We know from Proposition 3 that a high-quality publisher prefers (H, H), whereas a low-quality publisher prefers (I, I). Although both (H, H) and (I, I) are perfect Bayesian equilibria, (H, H) can survive the uncompromised equilibrium refinement but (I, I) cannot. To illustrate this, we note that a high-quality publisher stands to gain from replacing (I, I) with (H, H), whereas a low-quality one stands to lose. The uncompromised equilibrium requires advertisers to hold a belief that a high-quality publisher more likely deviates to the hybrid scheme and a low-quality one to PPI. Such an off-equilibrium belief reinforces the (H, H) equilibrium but compromises the (I, I) equilibrium.

The finding on the sustainability of the hybrid equilibrium seems consistent with the observation that many publishers adopt hybrid schemes, including both new (e.g., Facebook, LinkedIn) and traditional (e.g., AOL Advertising, Clicksor) publishers. To our knowledge, this is the first research to provide a rationale for the hybrid scheme as an equilibrium choice.

Figure 3 shows the equilibrium strategy profile as a function of misclassification rates γ_A and γ_P . In this figure, the solid line represents the boundary between the separating equilibrium (Region I) and the pooling equilibrium (Regions II and III), whereas the dashed line represents such a boundary without the hybrid scheme (i.e., the boundary in Figure 2). The shape of the new boundary is similar to that in Fig-

ure 2: a high-quality publisher will adopt PPC when the probability of misclassifying a publisher γ_P is high, and it will adopt a hybrid scheme when γ_P is low and γ_A is high (but not too high). The boundary shifts upward, suggesting that after the hybrid scheme becomes an option, two publisher types pool more often under the hybrid scheme. Intuitively, when there is a better pooling option (H, H), a high-quality publisher relies less on the costly PPC auction to differentiate itself. We summarize these findings in Corollary 2 and Table 3.

By adding the hybrid scheme to a publisher's choice set, we find that

Corollary 2a: If the opposite of the condition in Equation 15 holds (Figure 3, Region I), the equilibrium changes from (I, C) to (H, C) or (I, C), which leaves the publisher, advertisers, and overall allocative efficiency unaffected.

Corollary 2b: If the condition in Equation 15 and the opposite of the condition in Equation 14 hold (Figure 3, Region II), the equilibrium changes from (I, C) to (H, H), which benefits the publisher and improves overall allocative efficiency.

Corollary 2c: If the condition in Equation 14 holds (Figure 3, Region III), the equilibrium changes from (I, I) to (H, H), which benefits a high-quality publisher, hurts a low-quality publisher, and does not affect advertiser expected payoffs or overall allocative efficiency.

DISCUSSION AND CONCLUSIONS

Our study is the first to develop an analytical model of endogenous choice of pricing schemes in Internet advertising auctions that provides a theoretical explanation for the existence of multiple pricing schemes as well as the increasing popularity of hybrid schemes. In our model, pricing schemes not only affect the allocation of advertisers but also convey information about the publisher's quality. Because advertisers cannot perfectly distinguish low- and high-quality publishers, a low-quality publisher can benefit from pooling with a high-quality one under the PPI scheme; however, a P4P scheme prevents publisher pooling so that a high-quality publisher has an incentive to separate itself using a P4P scheme. A P4P publisher incurs a cost, however, because a P4P scheme requires the publisher to rank advertisers on the basis of inaccurate estimates of advertiser qualities, which can result in inefficient allocations and revenue losses. This trade-off between the costs and benefits of P4P pricing schemes, together with a novel analysis of hybrid schemes, lies at the heart of our findings.

Figure 3

PUBLISHER'S EQUILIBRIUM STRATEGY PROFILE WITH THE HYBRID SCHEME




Table 3
IMPACT OF THE HYBRID SCHEME

	Region I	Region II	Region III
Equilibrium without the hybrid scheme	(I, C)	(I, C)	(I, I)
Equilibrium with the hybrid scheme	(I, C) or (H, C)	(H, H)	(H, H)
Publisher's expected revenue (l-type, h-type)	#, #	↑, ↑	↓, ↑
Advertiser's expected payoff (l-type, h-type)	#, #	?, ?	#, #
Overall allocative efficiency	#	↑	#

Notes: ↑ = increase; ↓ = decrease; # = same; ? = increase or decrease.

Our analyses show that high-quality publishers find it optimal to choose a P4P scheme (e.g., PPC, PPA) when the benefits of differentiating themselves from low-quality publishers outweigh the loss from allocative inefficiency. A pure P4P scheme is most beneficial for a novel publisher that is of superior quality and can obtain accurate estimates of advertiser qualities. This offers an explanation for the adoption of PPC or PPA schemes at high-quality publishers such as Google Search Ads, Amazon Product Ads, Shopping.com Merchant Program, and Twitter Promoted Tweets. A publisher that holds ad inventory from well-known sources (in which case, there is little uncertainty about the publisher's quality) and lacks expertise and infrastructures for estimating advertiser qualities would be better off with the traditional PPI scheme. This explains why premium publishers such as CBSNews.com and TV.com often use PPI.

We further show that high-quality publishers prefer a hybrid scheme to a pure PPI scheme. Using advertiser self-selection, a hybrid scheme can achieve similar allocative efficiency as a PPI scheme. Yet a hybrid scheme allows a high-quality publisher to partly avoid the cost of pooling because only the PPI portion of the hybrid revenue is subject to pooling. In contrast, a low-quality publisher prefers PPI to a hybrid scheme. Remarkably, we find that a hybrid pooling equilibrium (in which both publisher types adopt the hybrid scheme) can "compromise" a PPI pooling equilibrium but not vice versa. Our results predict that the hybrid scheme has the potential of supplanting PPI. Furthermore, the adoption of hybrid schemes is associated with reduced use of pure P4P schemes and improves allocative efficiency. Indeed, anecdotal evidence highlights the increasing popularity of hybrid schemes among leading publishers such as Google Display Ads, Facebook, LinkedIn, and AOL Advertising. Although the exact revenue numbers are not available, our research suggests that 14 of the top 20 display advertising publishers ranked by comScore (2013) use hybrid pricing schemes. Some of these publishers, such as Google Display Ads and AOL Advertising, have evolved from a pure PPI scheme to a hybrid scheme.

Our study highlights the role of information asymmetry in determining pricing schemes. In markets such as display advertising, in which advertiser performance data are sparse and noisy, pure P4P schemes can be overly costly because of asymmetric information about advertiser qualities. Therefore, PPI and hybrid schemes are mainstream in these markets. In contrast, in markets in which performance data on advertisers are abundant and less noisy (e.g., sponsored search, promoted tweets), PPC and other P4P schemes are popular. Information asymmetry is the greatest for new publishers that offer innovative and effective advertising models. We expect such publishers to be more likely to distinguish themselves with P4P schemes. Examples of such publishers include Advertising.com, which used PPC for its targeted contextual ads (before AOL purchased it), and Tweeter, which uses PPA pricing for its promoted tweets.

Our approach to pricing schemes is based on information asymmetry and its consequences. We note that there is also a popular view of pricing schemes based on risks. As per the risk-based view, PPI shifts the majority of the risk to advertisers, whereas P4P shifts it to publishers. Thus, a risk-averse advertiser would prefer P4P, whereas a risk-averse publisher would prefer PPI. However, absent information asymmetry, risk preferences alone do not seem to explain

why some advertisers choose PPI over PPC or why some publishers voluntarily offer P4P schemes. Our approach augments the risk-based view by showing that pricing schemes matter even in a risk-neutral environment (e.g., PPI creates a publisher pooling effect and P4P may lead to inefficient allocations and correlated revenue losses). Although including risk aversion may shift the equilibrium conditions, its impact is well understood, and it is unlikely to reverse the direction of our findings. Moreover, we highlight the impact of pricing schemes on allocative efficiency, which is a fundamental concern of Internet advertising that the risk-based view does not capture. Anecdotal evidence suggests that allocative efficiency is a key consideration for publishers, which are happy to run PPC or other P4P ads as long as advertisers pay more (Lee 2006). By capturing the key trade-offs arising from information asymmetry, our study complements the risk-based perspective of pricing schemes.

As a first step in understanding the optimal choice of pricing schemes in the evolving landscape of digital advertising, our study is not without limitations. We assume the publisher quality to be the same for advertisers, which may not be realistic for slots that attract vastly different advertisers. We limit ourselves to a single slot, but we show in the Web Appendix that extension to multiple slots does not alter the fundamental insights. In the Web Appendix, we briefly discuss what would happen if valuation per click is correlated with publisher quality, but more research is needed in this area. Our findings based on information asymmetry should be combined with other perspectives on pricing schemes: for example, the PPC scheme may suffer from additional inefficiencies due to advertiser obfuscation and click fraud (Athey and Ellison 2011; Wilbur and Zhu 2009), although research has shown that the effect of click fraud is ambiguous (Wilbur and Zhu 2009).

Our current findings pave the way for several extensions. A natural extension of our model is to examine side-by-side publisher competition. Another way to extend existing insights is to understand the dynamics of pricing scheme choices and to examine when a publisher should let ad networks manage its advertising. It would also be worthwhile to consider the impact of "hidden actions" by advertisers and publishers in addition to information asymmetry; Zhu and Wilbur (2010) and Chen, Liu, and Whinston (2009) provide some clues in this direction. Our model may be extended to study pricing schemes for affiliate marketing, another important form of Internet marketing. Affiliate marketing is also subject to two-sided information asymmetry because its performance depends on both the quality of the merchant and the affiliate's ability to drive high-quality traffic to the merchant's website. Finally, given the increased availability of data on Internet advertising, empirical tests of theoretical predictions would add to the field's understanding of this complex landscape.

APPENDIX A: EQUILIBRIUM AUCTION OUTCOMES

Proof of Proposition 1

A PPI auction is equivalent to a standard second-price auction in which advertisers have valuations v_{ab} . Because equilibrium bids are not a primary interest of this research, we follow Myerson's (1981) approach of deriving equilibrium revenue directly from equilibrium winning probabilities without explicitly evaluating equilibrium bidding func-

tions. Although this approach is overkill for the single-slot case, it is easier to extend to the multislot case.

We denote $U(v, v'|a, \hat{b})$ as the expected payoff of an advertiser (v, a) that places a PPI bid v' at $\hat{b}$:

$$U(v, v'|a, \hat{b}) = \phi^I(v', a)v\hat{b} - p(v'|a, \hat{b}),$$

where $\phi^I(v', a)$ and $p(v'|a, \hat{b})$ are, respectively, the winning probability and the expected payment of the advertiser. We denote $V(v|a, \hat{b}) = U(v, v'|a, \hat{b})$ as advertiser (v, a) 's equilibrium expected payoff. Note that

$$\begin{aligned} \frac{dV(v|a, \hat{b})}{dv} &= \frac{\partial}{\partial v} U(v, v'|a, \hat{b}) \Big|_{v'=v} + \frac{\partial}{\partial v'} U(v, v'|a, \hat{b}) \Big|_{v'=v} \frac{dv'}{dv} \\ &= \frac{\partial}{\partial v} U(v, v'|a, \hat{b}) \Big|_{v'=v} = \phi^I(v, a)\hat{b}, \end{aligned}$$

where the second step occurs because $\frac{\partial v'}{\partial v} \Big|_{v'=v} = 0$ in an incentive-compatible equilibrium. Under the assumption that the advertiser with $v = 0$ has zero expected payoff, we immediately have

$$(A1) \quad V(v|a, \hat{b}) = a\hat{b} \int_0^v \phi^I(u, a) du.$$

A publisher's expected revenue is the total expected payment from all advertisers. Noting that $p(v|a, \hat{b}) = \phi^I(v, a)v\hat{b} - V(v|a, \hat{b})$, we can express the expected revenue of a $\tilde{y}$ -signal publisher as

$$\begin{aligned} n \sum_x P(x) \int_0^1 p(v|a_x, \hat{b}) f(v) dv &= n \sum_x P(x) \int_0^1 [\phi^I(v, a_x)v\hat{b} - V(v|a_x, \hat{b})] f(v) dv \\ &= n\hat{b} \sum_x P(x) a_x \int_0^1 [\phi^I(v, a_x)f(v)v - \phi^I(v, a_x) \int_v^1 f(u) du] dv \\ &= n\hat{b} \sum_x P(x) a_x \int_0^1 \phi^I(v, a_x) \left[v - \frac{1-F(v)}{f(v)} \right] f(v) dv, \end{aligned}$$

where the second equality is due to integration by parts. Finally, for a y -type publisher, we take expectation of $\hat{b}$ with regard to $\tilde{y}$ and obtain Equation 7.

Proof of Proposition 2

We denote $t(v, \hat{a})$ as the equilibrium bidding function of an advertiser with valuation v and expected quality $\hat{a}$. We assume the bidding function is monotonic (proof available on request); that is,

$$(A2) \quad v' > v \Rightarrow t(v', \hat{a}) > t(v, \hat{a}).$$

We next show that advertisers $(v, \hat{a}_I)$ and $(v\hat{a}_I/\hat{a}_H, \hat{a}_H)$ tie in equilibrium; that is,

$$(A3) \quad t(v, \hat{a}_I)\hat{a}_I/\hat{a}_H = t(v\hat{a}_I/\hat{a}_H, \hat{a}_H).$$

With a slight abuse of notation, we denote $U(v, t\hat{a})$ as the per-click payoff of an advertiser $(v, \hat{a})$ that bids t per click.

The per-impression payoff is simply $\hat{a}U(v, t\hat{a})$, but because the scale factor $\hat{a}$ does not affect the advertiser's bid optimization problem, we focus on the per-click payoff. We denote $\phi^C(t, \hat{a})$ and $p(t, \hat{a})$ as the winning probability and expected payment per click of an advertiser with expected quality $\hat{a}$ and bid t . Next, let advertiser $(v, \hat{a}_I)$ bid t and advertiser $(v\hat{a}_I/\hat{a}_H, \hat{a}_H)$ bid $t\hat{a}_I/\hat{a}_H$.

$$(A4) \quad \begin{aligned} U(v\hat{a}_I/\hat{a}_H, t\hat{a}_I/\hat{a}_H|\hat{a}_H) &= \phi^C(t\hat{a}_I/\hat{a}_H, \hat{a}_H) v\hat{a}_I/\hat{a}_H - p(t\hat{a}_I/\hat{a}_H, \hat{a}_H) \\ &= \phi^C(t, \hat{a}_I) v\hat{a}_I/\hat{a}_H - p(t, \hat{a}_I) t\hat{a}_I/\hat{a}_H = U(v, t|\hat{a}_I), \end{aligned}$$

where the second equality exists because two advertisers tie and pay the same expected amount. Equation A4 implies that if t maximizes $U(v, t|\hat{a}_I)$, $t\hat{a}_I/\hat{a}_H$ must maximize $U(v\hat{a}_I/\hat{a}_H, t\hat{a}_I/\hat{a}_H|\hat{a}_H)$.

Equations A2 and A3 together imply that an advertiser $(v, \hat{a})$ wins only if its $v\hat{a}$ is the highest. By the revenue equivalence theorem (Myerson 1981), we can obtain the equilibrium payoff of advertiser $(v, \hat{a})$ as

$$(A5) \quad V(v, \hat{a}) = \int_0^v \phi^C(u, \hat{a}) du.$$

Because the expected payment per click from the advertiser is $p(v, \hat{a}) = \phi^C(v, \hat{a})v - V(v, \hat{a})$ and the expected payment per impression is simply $\hat{a}p(v, \hat{a})$, a y -type publisher's expected revenue from all advertisers is

$$\begin{aligned} \pi_y^C &= n \sum_x P(\tilde{x}) \hat{a}_{\tilde{x}} b_y \int_0^1 p(v, \hat{a}_{\tilde{x}}) f(v) dv \\ &= nb_y \sum_x P(\tilde{x}) \hat{a}_{\tilde{x}} \int_0^1 \left[\phi^C(v, \hat{a}_{\tilde{x}})v - \int_0^v \phi^C(u, \hat{a}_{\tilde{x}}) du \right] f(v) dv \\ &= nb_y \sum_x P(\tilde{x}) \hat{a}_{\tilde{x}} \int_0^1 \left[\phi^C(v, \hat{a}_{\tilde{x}})f(v)v - \phi^C(v, \hat{a}_{\tilde{x}}) \int_v^1 f(u) du \right] dv \\ &= nb_y \sum_x P(\tilde{x}) \hat{a}_{\tilde{x}} \int_0^1 \phi^C(v, \hat{a}_{\tilde{x}}) \left[v - \frac{1-F(v)}{f(v)} \right] f(v) dv, \end{aligned}$$

where the third equality is due to integration by parts.

Proof of Lemma 1

We denote $\phi^H(z)$ as the equilibrium winning probability of an advertiser that bids a score of z (the score can be calculated from a PPI or PPC bid). We denote $s(z)$ as the expected second-highest score conditional on the highest score being z . Next, we consider an advertiser (v, a) with a PPC weighting factor w . If the advertiser submits a PPI bid $v'\hat{a}$, the expected payoff per impression is

$$(A6) \quad U^I(v, v') = \phi^H(v'\hat{a})[v\hat{a} - s(v'\hat{a})].$$

If the advertiser submits a PPC bid v' , the expected payoff (per impression) is

$$(A7) \quad U^C(v, v') = \hat{a}\phi^H(v'w)[v - s(v'w)/w].$$

When $\hat{a} = w$, $U^I(v, v') = U^C(v, v')$, which means that the advertiser is indifferent between PPI and PPC. Next, suppose $\hat{a} < w$. Let $v^*\hat{a}$ be the advertiser's optimal PPI bid. By

submitting a PPC bid of v^*ab/w , the advertiser gets an expected utility $U^C(v, v^*ab/w) = \varphi^H(v^*ab)[vab - s(v^*ab)ab/w] > U^I(v, v^*ab)$. Thus, the advertiser prefers PPC if $ab < w$. Similarly, if $ab > w$, the advertiser prefers PPI.

To show that the weighting scheme is unique, suppose that the weighting factor for $\tilde{h}$ -signal PPC bidders is $w_{\tilde{h}}$. If PPC bidders consist of both high- and low-quality advertisers, their expected quality (thus, the weight factor $w_{\tilde{h}}$) must be between $a_l\tilde{b}$ and $a_h\tilde{b}$. This implies that all high-quality $\tilde{h}$ -signal advertisers strictly prefer PPI—a contradiction. The only plausible case is that $\tilde{h}$ -signal advertisers consist of only low-quality advertisers, implying that $w_{\tilde{h}} = a_l\tilde{b}$. By the same logic, the weighting factor for $\tilde{l}$ -signal advertisers is $w_{\tilde{l}} = a_l\tilde{b}$.

Proof of Proposition 3

We denote $V_l^H(v)$ as the advertiser's equilibrium payoff per click from low-quality advertisers under the hybrid scheme. By the revenue equivalence theorem,

$$V_l^H(v) = \int_0^v \varphi^H(u, a_l) du,$$

and the expected payment per click from a low-quality advertiser is

$$\begin{aligned} & \int_0^1 [\varphi^H(v, a_l)v - V_l^H(v)] f(v) dv \\ &= \int_0^1 \left[\varphi^H(v, a_l)v - \int_0^v \varphi^H(u, a_l) du \right] f(v) dv \\ &= \int_0^1 \left[\varphi^H(v, a_l)vf(v) - \int_0^1 \varphi^H(u, a_l) \int_v^1 \varphi^H(f(v)) dv du \right] \\ &= \int_0^1 \varphi^H(v, a_l)J(v)f(v) dv, \end{aligned}$$

where the second equality is due to a change of the integration order. The publisher's expected revenue from low-quality advertisers is the sum of expected payment per impression of all low-quality advertisers. That is,

$$(A8) \quad n(1-\alpha)a_lb_y \int_0^1 \varphi^H(v, a_l)J(v)f(v) dv.$$

By a similar process, we can obtain the expected payment per impression from a high-quality advertiser as $a_h\hat{b}_y \int_0^1 \varphi^H(v, a_h)J(v)f(v) dv$.

A y-type publisher's expected revenue from all high-quality advertisers, averaged across all possible signals, is given by

$$(A9) \quad n\alpha a_h \hat{B}_y \int_0^1 \varphi^H(v, a_h)J(v)f(v) dv.$$

By adding Equations A8 and A9 together, we can get the y-type publisher's total expected revenue from all advertisers (Equation 11).

Next, we compare PPI and hybrid revenues. When advertisers hold the same belief about a PPI publisher and a hybrid publisher, we have $\hat{B}_y^I = \hat{B}_y^H = \bar{B}_y$. Thus, $\pi_y^H = \bar{B}_y\pi_{h,base}^H + b_y\pi_{l,base}^H$ and $\pi_y^I = \bar{B}_y\pi_{l,base}^I = \bar{B}_y\pi_{h,base}^I + \bar{B}_y\pi_{l,base}^I$. Because $b_l \leq \bar{B}_l \leq \bar{B}_h \leq b_h$, we have $\pi_l^I \leq \pi_l^H$ and $\pi_h^H \geq \pi_h^I$.

APPENDIX B: CHARACTERIZE THE PUBLISHER'S EQUILIBRIUM STRATEGY PROFILE

Proof of Lemma 2

The proof of Lemma 2 is straightforward, and thus, we have omitted it.

Proof of Proposition 5

Under the condition in Equation 13, a low-quality publisher gets an expected revenue of $\hat{B}_l\pi_{base}^I$ under PPI and $b_l\pi_{base}^C$ under PPC. Because $\hat{B}_l \geq b_l$, PPC is weakly dominated and, thus, not played by low-quality publishers. For this reason, we only consider the remaining two strategy profiles (I, C) and (I, I).

For (I, C) to be a PBE, we only need to show that a high-quality publisher optimally chooses PPC. A high-quality publisher's expected revenue under PPC is $b_h\pi_{base}^C$. By deviating to PPI, the publisher gets $b_l\pi_{base}^I$. A high-quality publisher will not deviate if

$$(B1) \quad b_h\pi_{base}^C > b_l\pi_{base}^I.$$

For (I, I) to be PBE, we only need to show that a high-quality publisher optimally chooses PPI. A high-quality publisher's expected revenue under PPI is $\bar{B}_h\pi_{base}^I$. By deviating to PPC, the publisher's expected revenue is $b_h\pi_{base}^C$. The high-quality publisher would not deviate if

$$(B2) \quad b_h\pi_{base}^C \leq \bar{B}_h\pi_{base}^I.$$

Because $b_l < \bar{B}_h$, the conditions in Equations B1 and B2 may hold at the same time. We next show that (I, C) is Pareto-dominated when both conditions hold. Under (I, I), the expected revenues for the high- and low-quality publishers are $\bar{B}_h\pi_{base}^I$ and $\bar{B}_l\pi_{base}^I$, respectively. Under (I, C), their expected revenues are $b_h\pi_{base}^C$ and $b_l\pi_{base}^C$, respectively. In accordance with the condition in Equation B2 and $\bar{B}_l > b_l$, (I, C) is Pareto-dominated by (I, I) and, thus, not played.

Proof of Corollary 1

The proof is straightforward from the condition in Equation 14; thus, we have omitted it.

Lemma 3

Under Assumptions 1 and 2 and the condition in Equation 13,

(a) (I, I) with $\mu(h|H) \leq \mu_1$ is a PBE when

$$(B3) \quad b_h\pi_{base}^C \leq \bar{B}_h\pi_{base}^I,$$

where $\mu(h|H) \leq \mu_1$ summarizes the following two conditions:

$$(B4) \quad \bar{B}_h\pi_{base}^I \geq \hat{B}_h\pi_{h,base}^H + b_h\pi_{l,base}^H, \text{ and}$$

$$(B5) \quad \bar{B}_l\pi_{base}^I \geq \hat{B}_l\pi_{h,base}^H + b_l\pi_{l,base}^H.$$

(b) (H, H) with a belief $\mu(h|I) \leq \mu_2$ is a PBE when

$$(B6) \quad b_h \pi_{base}^C \leq \bar{B}_h \pi_{h, base}^H + b_h \pi_{l, base}^H,$$

where $\mu(h|I) \leq \mu_2$ summarizes the following two conditions,

$$(B7) \quad \hat{B}_h \pi_{base}^I \leq \bar{B}_h \pi_{h, base}^H + b_h \pi_{l, base}^H, \text{ and}$$

$$(B8) \quad \hat{B}_l \pi_{base}^I \leq \bar{B}_l \pi_{h, base}^H + b_l \pi_{l, base}^H.$$

(c) (I, C) with a belief $\mu(h|H) = 0$ and (H, C) with a belief $\mu(h|I) = 0$ are PBEs when

$$(B9) \quad b_h \pi_{base}^C > \bar{B}_h \pi_{h, base}^H + b_h \pi_{l, base}^H.$$

By Lemma 3, (I, I) is an equilibrium when the condition in Equation A3 holds and advertisers believe that a hybrid publisher is most likely low quality (i.e., $\mu(h|H) < \mu_1$). Similarly, (H, H) is an equilibrium when the condition in Equation B6 holds and advertisers believe that a PPI publisher is most likely low quality. The equilibrium (I, C) requires that the condition in Equation B9 holds and advertisers believe that a hybrid publisher is low quality.

Proof of Lemma 3

Because PPC is weakly dominated by PPI and the hybrid scheme for a low-quality publisher, we can rule out (C, I), (C, H), and (C, C). (I, H) and (H, I) are not PBE either because the low-quality publisher is better off by mimicking the high-quality publisher in either case. We show that the remaining four strategy profiles are PBE under appropriate conditions.

(a) A low-quality publisher gets $\bar{B}_l \pi_{base}^I$ under (I, I) and $\hat{B}_l \pi_{h, base}^H + b_l \pi_{l, base}^H$ by deviating to the hybrid scheme. A high-quality publisher gets $\bar{B}_h \pi_{base}^H$ under (I, I), $\hat{B}_h \pi_{h, base}^H + b_h \pi_{l, base}^H$ by deviating to the hybrid scheme and $b_h \pi_{base}^C$ by deviating to PPC. the condition in Equations B3, B4, and B5 ensure that all deviations are unprofitable. Note that $\hat{B}_h$ and $\hat{B}_l$ increase with $\mu(h|H)$; therefore, we can combine Equations B4 and B5 into a requirement on belief $\mu(h|H)$.

(b) A low-quality publisher gets $\bar{B}_l \pi_{h, base}^H + b_l \pi_{l, base}^H$ under (H, H) and $\hat{B}_l \pi_{base}^I$ by deviating to PPI. A high-quality publisher gets $\bar{B}_h \pi_{h, base}^H + b_h \pi_{l, base}^H$ under (H, H), $\hat{B}_h \pi_{base}^I$ by deviating to PPI, and $b_h \pi_{base}^C$ by deviating to PPC. the conditions in Equations B6, B7, and B8 ensure that all deviations are unprofitable. Similar to (a), the conditions in Equations B4 and B5 can be represented as a requirement on belief $\mu(h|I)$.

(c) A low-quality publisher gets $b_l \pi_{base}^I$ under (I, C) and $\hat{B}_l \pi_{h, base}^H + b_l \pi_{l, base}^H$ by deviating to the hybrid scheme. Recall that $\pi_{base}^I = \pi_{h, base}^H + \pi_{l, base}^H$; the deviation is unprofitable only when $\hat{B}_l = b_l$, which implies that $\mu(h|H) = 0$. A high-quality publisher gets $b_h \pi_{base}^C$ under (I, C), $b_l \pi_{base}^I$ by deviating to PPI, and $b_h \pi_{h, base}^H + b_h \pi_{l, base}^H$ (note that $\mu(h|H) = 0$) by deviating to the hybrid scheme. For the high-quality publisher, deviating to PPI is clearly dominated by deviating to the hybrid scheme. The belief $\mu(h|H) = 0$ and the following condition ensure that all deviations are unprofitable. Thus,

$$(B10) \quad b_h \pi_{base}^C \geq b_l \pi_{h, base}^H + b_h \pi_{l, base}^H.$$

(d) A low-quality publisher gets $b_l \pi_{base}^I$ under (H, C) and $\hat{B}_l \pi_{base}^I$ by deviating to PPI. The deviation is unprofitable only when $\hat{B}_l = b_l$, which implies $\mu(h|I) = 0$. A high-quality publisher gets $b_h \pi_{base}^C$ under (H, C), $b_l \pi_{base}^I$ (with $\mu(h|I) = 0$) by deviating to PPI, and $b_h \pi_{h, base}^H + b_h \pi_{l, base}^H$ by deviating to the hybrid scheme. Because $\pi_{base}^I = \pi_{h, base}^H + \pi_{l, base}^H$, deviating to PPI is clearly dominated by deviating to the hybrid scheme. The belief $\mu(h|I) = 0$ and the condition in Equation B10 ensure that all deviations are unprofitable.

When B7 and B10 hold simultaneously, (I, C), (H, C), and (H, H) are all PBEs. The equilibrium revenues under three PBEs are $(b_l \pi_{base}^I, b_h \pi_{base}^C)$, $(b_l \pi_{base}^I, b_h \pi_{base}^C)$, and $(\bar{B}_h \pi_{h, base}^H + b_h \pi_{l, base}^H, \bar{B}_h \pi_{h, base}^H + b_h \pi_{l, base}^H)$, respectively. By the condition in Equation B6, (H, H) Pareto-dominates (I, C) and (H, C); thus, we can revise the condition for (I, C) and (H, C) to Equation B9.

Proof of Proposition 6

We show that (I, I) is compromised by (H, H) but (H, H), (I, C), (H, C) are uncompromised.

- (a) (I, I). When the condition in Equation B3 holds, both (I, I) and (H, H) are PBEs. According to Proposition 3, a high-quality publisher is better off under (H, H) than under (I, I) whereas a low-quality publisher is worse off. The uncompromised equilibrium requires $\mu(h|H) \geq \beta$. However, when $\mu(h|H) \geq \beta$, $\bar{B}_h \geq \bar{B}_l$, and note also $b_h > \bar{B}_h$, we have $\bar{B}_h \pi_{base}^I < \bar{B}_h \pi_{h, base}^H + b_h \pi_{l, base}^H$, contradicting the condition in Equation B4. Therefore, (I, I) is compromised by (H, H).
- (b) (H, H). When the condition in Equation A3 holds, both (I, I) and (H, H) are PBEs. Similar to (a), the uncompromised equilibrium requires that $\mu(h|I) \leq \beta$, which clearly does not contradict $\mu(h|I) \leq \mu_2$. Thus, (H, H) is not compromised by (I, I). (H, H) does not coexist with any other equilibrium, so (H, H) is uncompromised.
- (c) (I, C) and (H, C). (I, C) and (H, C) are both PBEs under the condition in Equation B9. But both publisher types are indifferent between (I, C) and (H, C). So the uncompromised equilibrium refinement has no force. (I, C) and (H, C) are thus uncompromised.

Proof of Corollary 2

The result that the high-quality publisher is better off under (H, H) than under (I, I) and the low-quality publisher is worse off follows from Proposition 3. To show that both publisher types are better off in case (b), note that this occurs when

$$(B11) \quad \bar{B}_h \pi_{h, base}^H + b_h \pi_{l, base}^H \geq b_h \pi_{base}^C \geq \bar{B}_h \pi_{base}^I.$$

The equilibrium revenues for low- and high-quality types are $(\bar{B}_l \pi_{h, base}^H + b_l \pi_{l, base}^H, \bar{B}_h \pi_{h, base}^H + b_h \pi_{l, base}^H)$ under (H, H) and $(b_l \pi_{base}^I, b_h \pi_{base}^C)$ under (I, C). The (H, H) revenues are higher than the (I, C) revenues, noting $b_l < \bar{B}_l$, the condition in Equation B11, and Equation 12.

We next turn to advertiser payoffs. In case (a), all advertisers are indifferent because the PPI and hybrid auctions allocate the impression in the same way and advertisers perfectly infer the publisher's quality in both equilibria. In case (c), in which the equilibrium changes from (I, I) to (H, H), advertisers are indifferent as well. Again, because PPI and hybrid auctions allocate the impression in the same way, their winning probabilities and expected payoffs are the same under either equilibrium. In case (b), in which (I, C) is replaced by (H, H), advertisers' winning probabilities are different, and we are unable to draw a definitive conclusion on advertiser payoffs because when the publisher is high quality (so that PPC would be used under the (I, C) equilibrium), a high-quality advertiser can be better off or worse off under PPC depending on whether the signal about the advertiser's quality is $\bar{h}$ or $\bar{l}$. Finally, we know from previous proofs that PPI and hybrid auctions are more efficient than PPC auctions.

REFERENCES

- Agarwal, Ashish, Kartik Hosanagar, and Michael D. Smith (2011), "Location, Location, Location: An Analysis of Profitability of Position in Online Advertising Markets," *Journal of Marketing Research*, 48 (December), 1057–73.
- Animesh, Animesh, Vandana Ramachandran, and Siva Viswanathan (2009), "Quality Uncertainty and the Performance of Online Sponsored Search Markets: An Empirical Investigation," *Information Systems Research*, 21 (1), 190–201.
- Athey, Susan and Glenn Ellison (2011), "Position Auctions with Consumer Search," *Quarterly Journal of Economics*, 126 (3), 1213–70.
- Chen, Jianqing, De Liu, and Andrew B. Whinston (2009), "Auctioning Keywords in Online Search," *Journal of Marketing*, 73 (July), 125–41.
- Choudhary, Vidyanand (2010), "Use of Pricing Schemes for Differentiating Information Goods," *Information Systems Research*, 21 (1), 78–92.
- comScore (2013), "comScore Media Matrix Ranks Top 50 U.S. Web Properties for May 2013," press release, (June 24), (accessed June 13, 2014), [available at http://www.comscore.com/Insights/Press_Releases/2013/6/comScore_Media_Matrix_Ranks_Top_50_US_Web_Properties_for_May_2013].
- Dellarocas, Chrysanthos (2012), "Double Marginalization in Performance-Based Advertising: Implications and Solutions," *Management Science*, 58 (6), 1178–95.
- The Economist (2005), "Pay per Sale: The Holy Grail of Advertising Is Within Reach," (September 29), (accessed June 12, 2014), [available at <http://www.economist.com/node/4462811>].
- Edelman, Benjamin and Hoan Soo Lee (2008), "CPC/CPA Hybrid Bidding in a Second Price Auction," working paper, Harvard Business School.
- and Michael Ostrovsky (2007), "Strategic Bidder Behavior in Sponsored Search Auctions," *Decision Support Systems*, 43, 192–98.
- , ———, and Michael Schwarz (2007), "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97 (1), 242–59.
- Feng, Juan (2008), "Optimal Mechanism for Selling a Set of Commonly Ranked Objects," *Marketing Science*, 27 (3), 501–512.
- Ghose, Anindya and Sha Yang (2009), "An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets," *Management Science*, 55 (10), 1605–1622.
- Goel, Ashish and Kamesh Munagala (2009), "Hybrid Keyword Search Auctions," in *Proceedings of the 18th International Conference on World Wide Web*. New York: Association for Computing Machinery, 221–30.
- Google (2014), "Understanding Bidding Basics," (accessed June 13, 2014), [available at <https://support.google.com/adwords/answer/2459326?hl=en>].
- Hu, Yu (2004), "Performance-Based Pricing Models in Online Advertising," unpublished manuscript.
- , Jiwoong Shin, and Zhulei Tang (2010), "Pricing of Online Advertising: Cost-Per-Click-Through vs. Cost-Per-Action," in *2010 43rd Hawaii International Conference on System Sciences*. Washington, DC: IEEE Computer Society, 1–9.
- Internet Advertising Bureau (2013), "IAB Internet Advertising Revenue Report: 2013 First Six Months' Results," research report, (October), (accessed June 12, 2014), [available at http://www.iab.net/media/file/IAB_Internet_Advertising_Revenue_Report_HY_2013.pdf].
- Jain, Sanjay and P. K. Kannan (2002), "Pricing of Information Products on Online Servers: Issues, Models, and Analysis," *Management Science*, 48 (9), 1123–42.
- Jeziorski, Przemyslaw and Ilya Segal (2009), "What Makes Them Click: Empirical Analysis of Consumer Demand for Search Advertising," working paper, Stanford University.
- Jullien, Bruno and Thomas Mariotti (2006), "Auction and the Informed Seller Problem," *Games and Economic Behavior*, 56 (2), 225–58.
- Kantrowitz, Alex (2013), "IAB: 72% of Publishers Are Now Using Programmatic Ad Technologies," *Advertising Age*, (November 4), (accessed June 12, 2014), [available at <http://adage.com/article/digital/72-publishers-programmatic-technology/245090/>].
- Katona, Zsolt and Yi Zhu (2012), "Quality Score That Makes You Invest," working paper.
- Lahaie, Sébastien (2006), "An Analysis of Alternative Slot Auction Designs for Sponsored Search," in *Proceedings of the 7th ACM Conference on Electronic Commerce*. New York: Association for Computing Machinery, 218–27.
- and David M. Pennock (2007), "Revenue Analysis of a Family of Ranking Rules for Keyword Auctions," in *Proceedings of the 8th ACM Conference on Electronic Commerce*. New York: Association for Computing Machinery, 50–56.
- Lee, Kevin (2006), "The CMO's Guide to CPA, CPC, and CPM," *ClickZ: Marketing News & Expert Advice*, (November 3), (accessed June 12, 2014), [available at <http://www.clickz.com/clickz/column/1696108/the-cmos-guide-cpa-cpc-cpm>].
- Li, Jun, De Liu, and Shulin Liu (2013), "Optimal Keyword Auctions for Optimal User Experiences," *Decision Support Systems*, 56, 450–61.
- Libai, Barak, Eyal Biyalogorsky, and Eitan Gerstner (2003), "Setting Referral Fees," *Journal of Service Research*, 5 (4), 303–315.
- Liu, De and Jianqing Chen (2006), "Designing Online Auctions with Past Performance Information," *Decision Support Systems*, 42, 1307–1320.
- , ———, and Andrew B. Whinston (2010), "Ex Ante Information and the Design of Keyword Auctions," *Information Systems Research*, 21 (1), 133–53.
- Mailath, George J., Masahiro Okuno-Fujiwara, and Andrew Postlewaite (1993), "On Belief Based Refinements in Signaling Games," *Journal of Economic Theory*, 60 (2), 241–76.
- Maskin, Eric and Jean Tirole (1992), "The Principal-Agent Relationship with an Informed Principal, II: Common Values," *Econometrica*, 60 (1), 1–42.
- Myerson, Roger B. (1981), "Optimal Auction Design," *Mathematics of Operations Research*, 6 (1), 58–73.
- Raehsler, Lisa (2014), "What Is a Good Click-Through Rate for PPC?" *ClickZ: Marketing News & Expert Advice*, (May 7), (accessed June 12, 2014), [available at <http://www.clickz.com/clickz/column/2186867/click-rate-ppc>].
- Rutz, Oliver J. and Randolph E. Bucklin (2007), "A Model of Individual Keyword Performance in Paid Search Advertising," working paper.
- Sundararajan, Arun (2003), "Pricing Digital Marketing: Information, Risk Sharing and Performance," working paper, Leonard N. Stern School of Business, New York University.
- (2004), "Nonlinear Pricing of Information Goods," *Management Science*, 50 (12), 1660–73.
- Varian, Hal R. (2007), "Position Auctions," *International Journal of Industrial Organization*, 25 (6), 1163–78.
- Wilbur, Kenneth C. and Yi Zhu (2009), "Click Fraud," *Marketing Science*, 28 (2), 293–308.
- Xu, Lizhen, Jianqing Chen, and Andrew Whinston (2010), "Oligopolistic Pricing with Online Search," *Journal of Management Information Systems*, 27 (3), 111–42.
- , ———, and ——— (2011), "Price Competition and Endogenous Valuation in Search Advertising," *Journal of Marketing Research*, 48 (June), 566–86.
- Yao, Song and Carl F. Mela (2011), "A Dynamic Model of Sponsored Search Advertising," *Marketing Science*, 30 (3), 447–68.
- Zhang, Xiaoquan and Juan Feng (2011), "Cyclical Bid Adjustments in Search-Engine Advertising," *Management Science*, 57 (9), 1703–1719.
- Zhu, Yi and Kenneth C. Wilbur (2010), "Hybrid Advertising Auctions," *Marketing Science*, 30 (2), 249–73.

Copyright of Journal of Marketing Research (JMR) is the property of American Marketing Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.