Information Asymmetry and Hybrid Advertising*

De Liu[†], Siva Viswanathan[‡]

May 15, 2012

Abstract

Pay-for-performance advertising schemes such as pay-per-click (PPC) and pay-per-sale (PPS) have grown in popularity with the recent advances in digital technologies for targeting advertising and measuring outcomes. Meanwhile, the traditional pay-per-impression (PPI) scheme persists, and several advertising providers have started to offer a hybrid mix of PPI and PPC schemes. Given the multiple pricing schemes - PPI, PPC, PPS, and hybrid - our study examines the optimal choices for providers. We highlight the role of pricing schemes as a means of leveraging private information available to providers and advertisers. In particular, our study demonstrates a trade-off between using pay-for-performance schemes to reveal superior quality and using the PPI scheme to minimize allocative inefficiencies. Our study identifies conditions under which providers find it optimal to offer PPI, pay-for-performance, or hybrid schemes. Our results provide insights into a number of observed provider strategies, including the growing popularity of hybrid pricing schemes. We discuss the implications for advertisers, advertising providers, and technology providers.

Keywords: Online Advertising, Pay-for-Performance, Information Asymmetry, Hybrid Pricing

1 Introduction and motivation

As the Internet permeates almost all walks of life, interest in advertising to Internet users has grown in recent years. According to the Interactive Advertising Bureau, Internet advertising revenues for

*The authors wish to thank seminar participants at University of Minnesota, University of Texas at Austin, Algorithmic Game Theory Workshop at UCLA, 2011 Workshop on Data Mining and Audience Intelligence for Online Advertising, 2010 China Summer Workshop on Information Management, University of International Business and Economics, Tsinghua University, and Southwestern University of Finance and Economics for their valuable feedback.

[†]University of Kentucky, e-mail: de.liu@uky.edu

[‡]University of Maryland, College Park, e-mail: sviswana@rhsmith.umd.edu

the first quarter in 2011 reached $7.3 billion, a 23% increase over the same period in 2010. Meanwhile, Internet advertising revenues surpassed print advertising revenues in 2010, according to a recent Pew Research Center study. In keeping with the growing popularity of Internet advertising, there is growing interest among academicians as well as practitioners in understanding the Internet advertising landscape. An important and interesting feature of Internet advertising is the greater diversity in pricing schemes compared to traditional advertising (TV, radio, and print). Traditional advertising is dominated by the pay-per-impression (PPI) pricing scheme, in which advertisers are charged according to the number of impressions allocated to them. The Internet, on the other hand, provides a low-cost way of tracking the performance of impressions in terms of clicks, purchases, registrations, among others, and in the process facilitating a variety of "pay-for-performance" schemes, such as pay-per-click (PPC) and pay-per-sale (PPS).[1] In the case of PPC the advertiser pays only when a user clicks on the advertisement. PPS requires that the advertisers pay only when the advertisement yields a sale. As we move from PPI to PPC to PPS, advertising expenses are more closely tied to advertising performance (i.e., the advertiser's sales and revenue). It is not surprising, then, that PPS is often considered the "holy grail of advertising." (The Economist, 2005) However, predictions that these pay-for-performance schemes (PPS, in particular) would supplant the traditional PPI pricing have not materialized. On the contrary, as shown in Table 1, these different pricing schemes coexist, with increasing number of Internet advertising providers, including search engines, advertising networks, and large websites who run their own advertising operations, choosing to offer a mix of PPC and PPI schemes (hereafter denoted as the *hybrid scheme*). Facebook, for example, allows advertisers to pay either by clicks or by impressions.[2] Table 1 shows that advertising providers differ in their choice of pricing schemes. For example, while Facebook uses the hybrid scheme for its display advertisements, Amazon uses the traditional PPI scheme for its counterpart. Some advertising providers, such as eBay Partner network, adBrite, and Google Display Ads, have experimented with multiple pricing schemes over time. Meanwhile practitioners and bloggers often give different and sometimes contradictory opinions on the subject.[3] Given the

[1]PPS is one kind of pay-per-action model, where the "action" can be anything from a purchase, a registration, a download, or an e-mail.

[2]See its FAQ at http://www.facebook.com/help/?faq=220734457954046

[3]A simple search of "CPC versus CPM" will yield numerous entries that cover a range of different opinions. For example, about 2007, Webmaster World debated whether PPS will completely replace PPC in Google display advertisements. Debojoytipal.com, a blog catering to website and blog owners, believes that PPC is for newcomers with low and medium traffic and should not be used on websites with huge traffic. The blog also believes that PPS is most profitable for websites with high quality traffic. On the other hand, Vaughn's Summaries (vaughns-1-pagers.com/internet/internet-ad-networks.htm), a popular source on the subject, believes that PPC is the best choice in almost all situations, and PPI is only for large websites who cannot run PPC.

	Search Advertising	**Display Advertising**
PPI		Major newspapers (e.g. NY Times), ValueClick, Tribal Fusion, Amazon Links and Banner ads, adBrite (till 2008)
PPC	Google, Bing, Yahoo!	Amazon Product Ads, Shopping.com Merchant Program, Google AdSense (till 2005), Bidvertiser, eBay Partner Network (since 2009)
PPS	Google (2007-2008)	Amazon Associates, Google Affiliate Advertising, Commission Junction, Linkshare, eBay Partner Network (till 2009).
Hybrid (PPI/PPC)		Google AdSense (since 2005), MSN, Facebook, Twitter Promoted Tweets, Linkedin, Clicksor, adBrite (since 2008)

Search and display advertising are two leading forms of Internet advertising. Search advertisements are text-based advertisements in search engine result pages while display advertisements are text, graphic, video, or interactive advertisements that are embedded in non-search pages.

Table 1: Pricing Schemes Used by Search and Display Advertising Providers

significance of choosing pricing schemes in Internet advertising and lack of understanding of the issue, the goal of this paper is to answer two related questions: First, can we explain the coexistence of multiple pricing schemes, especially the presence of hybrid pricing schemes in today's Internet advertising market? Second, given the choice of multiple pricing schemes, how should advertising providers (a.k.a publishers) choose pricing schemes?

This research is based on two fundamental observations about Internet advertising. First, significant information asymmetry exists regarding advertisement performances. The performance of an advertisement, in terms of number of clicks or sales generated, depends not only on the quality of the advertisement but also on how effectively an advertising provider can target the right Internet users at the right time. Advertisers are not perfectly informed about advertising providers' targeting quality, and advertising providers may not be informed about the quality of the advertisements provided by the advertiser. Pricing schemes matter in the context of information asymmetry because advertising providers may then use pricing schemes to differentiate themselves and to reveal their targeting quality. Second, because advertisers differ in their valuation of an impression, and the supply of impressions is limited, allocative efficiency matters in Internet advertising.[4] Different pricing schemes focus on different elements of an advertiser's valuation, and may therefore, result in different allocations of impressions. The impact of pricing scheme on allocative efficiency is most

[4] Allocative efficiency is defined as the ability of a market to allocate goods to those agents who value them the most.

evident when advertising impressions are allocated using auctions, which are increasingly adopted by advertising providers including Google, Facebook, and Linkedin (see Edelman et al. 2007 for an account of the auction format used by Google). Different pricing schemes may result in different rankings of advertisers in these auctions, and thus, different allocative efficiencies. In light of these observations, we study the optimality of the pricing schemes in an auction framework by modeling their impact on allocative efficiency, as well as on the ability of advertising providers to differentiate themselves.

Our analysis offers two important insights. First, the choice of a pure pay-for-performance pricing scheme (say PPC or PPS) can serve as a differentiation strategy for high-quality providers. In general, high-quality providers are more likely to offer pay-for-performance schemes as they reveal the provider's targeting quality, and low-quality providers prefer a PPI scheme. Naturally, a pay-for-performance scheme is more likely in the presence of many low-quality providers as pooling would be more costly for high-quality advertising providers. However, given the uncertainty about advertiser quality, pay-for-performance schemes typically result in allocative inefficiencies, and in some cases the cost of this inefficiency outweighs any benefits from better differentiation. Consequently, high-quality providers may find it optimal to pool with low-quality ones by choosing a PPI scheme. As providers get better at estimating advertiser quality, pay-for-performance schemes become a more attractive strategy for high-quality providers. The trade-off between revealing superior targeting quality and minimizing allocative inefficiency drives a high-quality provider's choice of pricing scheme.

Second, when advertising providers have a choice of offering a hybrid pricing scheme (e.g., offering both PPI and PPC for advertisers to choose), high-quality providers prefer the hybrid scheme to the pure PPI scheme and choose pure pay-for-performance schemes less often. A hybrid scheme can lead to allocative efficiencies similar to a pure PPI scheme because high-quality advertisers self-select into the PPI portion of the hybrid scheme. Yet, the costs to a high-quality provider from pooling with low-quality ones are reduced compared with a pure PPI scheme as only a portion of advertisers who choose the PPI portion of the hybrid scheme may misjudge a provider's targeting quality. For this reason, a hybrid scheme dominates the PPI scheme for high-quality providers. By adopting a hybrid pricing scheme, high-quality providers can force low-quality providers to follow suit, causing the hybrid scheme to supplant the PPI scheme in equilibrium. Moreover, in some cases high-quality providers may find it beneficial to substitute the hybrid scheme for the more costly pure pay-for-performance schemes. The spread of the hybrid scheme benefits high-quality providers and the Internet advertising industry as a whole because it achieves higher allocative efficiency than pure pay-for-performance schemes.

Our research contributes to a growing literature on Internet advertising auctions. Edelman et al. (2007) and Varian (2007) analyze the equilibria in search advertising auctions as "generalized second-price (GSP)" auctions. Several papers examine advertising auction design (Chen et al., 2009; Feng, 2008; Liu et al., 2010; Weber and Zheng, 2007) and its interaction with consumer search (Athey and Ellison, 2011; Xu et al., 2010, 2011). In addition, empirical research on advertising auctions has grown steadily (Agarwal et al., 2011; Edelman and Ostrovsky, 2007; Zhang and Feng, 2011). Several authors (Rutz and Bucklin, 2007; Animesh et al., 2009; Yao and Mela, 2011; Abhishek et al., 2010) have examined the relationship between an advertisers' bid and outcomes such as rank, click-through rate, and conversions, while others (Animesh et al., 2011; Ghose and Yang, 2009; Jeziorski and Segal, 2009) focus on consumer behavior. Extant research on advertising auctions largely takes the pricing scheme for granted, and surprisingly the theoretical or empirical studies have failed to examine the choice of pricing schemes, or their impact on advertiser and provider strategies.

A few studies have examined pricing schemes used in early days of Internet advertising when auctions were not as common. Hu (2004) and Hu et al. (2010) study pricing schemes as an optimal contract design problem between one provider and one advertiser. They find that Internet advertising pricing should include appropriate performance-based elements to provide incentives for the provider and the advertiser to improve advertisement performance. Sundararajan (2003) examines pricing schemes as an optimal two-part tariff problem. He shows that, because advertisers are risk averse, pay-for-performance pricing is always optimal, even when providers are constrained to offer PPI in parallel. In contrast, our findings show that pay-for-performance may be less than optimal due to its allocative inefficiencies - a defining feature of the sponsored search auctions context. Furthermore, offering a hybrid scheme is an opportunity, rather than a liability, for high quality providers. A key distinction between this research and above papers is that we examine pricing schemes in an auction setting where allocative efficiency matters.

To the best of our knowledge, this study is the first to establish the rationale for hybrid pricing in Internet advertising auctions. A few recent papers have examined Internet advertising auctions with hybrid pricing. Zhu and Wilbur (2010) study the equilibrium bidding in hybrid advertising auctions but in a different setting from ours. In their model, advertisers can choose to induce more or fewer clicks *after* the slots are allocated. Their study takes the hybrid pricing for granted and does not explain why hybrid pricing is used in the first place. In their study, the number of clicks generated is affected by advertisers' ex-post actions. In contrast, we study an information asymmetry model in which the number of clicks generated is affected by ex-ante characteristics of advertisers and advertising providers. Edelman and Lee (2008) compare unweighted pure PPC and PPS auctions with a PPC/PPS hybrid auction. They show that advertisers will bid truthfully in the hybrid scheme and that hybrid auctions may produce no less revenue than unweighted pure PPC and PPS auctions since the latter two are special cases of PPC/PPS hybrid auctions. It is unclear though what their conclusion will be if PPC and PPS auctions are weighted by advertisers expected quality, which is the case studied in this paper. Goel and Munagala (2009) analyze a different hybrid auction where advertisers submit a PPI bid and a PPC bid for the provider to choose. They show that such a hybrid auction maintains at least $(1 - 1/e)$ of the weighted-PPC auction revenue. But their hybrid auction model differs from the real-world hybrid auctions in which advertiser submits a single bid. In sum, the existing research on hybrid auctions does not explain why hybrid auctions are used in the first place.

Our research is also related to the vast literature on optimal pricing of congestible services such as network infrastructure and computing services (Bashyam, 2000; Dewan and Mendelson, 1990; Gupta and Stahl, 1997; Gupta et al., 2011; Hosanagar et al., 2005; Masuda and Whang, 2006; Mendelson, 1985) and information goods (Choudhary, 2010; Jain and Kannan, 2002; Sundararajan, 2004). Both streams of research have compared fixed and usage-based pricing. Choudhary (2010), in particular, shows that two identical information goods sellers can differentiate on pricing schemes to avoid perfect competition, and sellers will offer both fixed and usage-based pricing only when their offerings are sufficiently differentiated. While parallels exist between usage-based pricing of information goods and pay-for-performance in Internet advertising, there are notable distinctions: information goods are non-rivaled (capacitated services are subject to limited congestion), whereas advertising resources are rivaled and subject to capacity constraints. More important, as a unique feature of our problem, advertising performance is jointly determined by both provider and advertiser, whereas the usage of informational goods and capacitated services is typically determined by user alone.

The rest of the paper is organized as follows. Section 2 outlines the modeling primitives and Section 3 details the equilibrium analyses. Section 4 discusses implications of our findings and concludes.

2 The Model

We consider two sets of risk-neutral players, advertising providers and advertisers. Each provider has a single impression to offer to one of the n advertisers.[5] A provider may be the owner of the impression or an intermediary who sells the impression on behalf of its owner. Similarly, an advertiser may be a marketer or its agent. The transaction between providers and advertisers proceeds in the following way. A single provider is randomly chosen to meet the n advertisers. The provider then uses an auction to allocate the impression among the advertisers.

We focus on clicks as a performance measure. Advertisers differ in valuation per click (*valuation* for short), denoted as $v \in [0, 1]$.[6] Valuations are private information, independently and identically distributed according to the distribution $F(v)$, which has a strictly positive and differentiable density function $f(v)$. We assume that the increasing hazard rate (IHR) condition holds, i.e., $f(v) / [1 - F(v)]$ is increasing. This condition is satisfied by common distributions such as uniform, normal, and logistic.

The click-through rate, defined as the probability of being clicked on, is a measure of an advertisement's *performance* and is jointly determined by the advertiser quality a and the provider quality b:

$$\text{Click through rate} = a \times b \tag{1}$$

An advertiser's quality a captures the attractiveness of the advertisement. An advertiser's quality is high if the advertiser has a competitive product or an appealing offer. All else equal, a high-quality advertiser has a higher click-through rate than a low-quality advertiser. A provider's quality b captures the effectiveness of the provider's targeting technology. A provider's quality is higher if the provider uses technologies to target the most relevant Internet users, so that they tend to click on advertisements from this provider more often. Such targeting technologies typically involve using information, such as time, location, demographics, contextual data (content being requested), and behaviorial data (click/shopping history), to guide advertisement placements. The provider's targeting quality should be distinguished from the quality of raw traffic. With the help of rigorous targeting, high-quality providers can find nuggets buried in low-quality traffic and present

[5]In practice, advertising providers tend to offer multiple slots simultaneously. Several authors (e.g., Edelman et al. 2007 and Liu et al. (2010)) offer detailed analysis of such multiple-slot auctions. Having multiple slots complicates the equilibrium bidding but does not alter the qualitative results for the purpose of this research.

[6]Our results will not change if the upper bound is an arbitrary positive number.

them to advertisers.

For simplicity, we assume two types of advertiser quality, a_l and a_h ($a_l < a_h$) and use $x \in \{l, h\}$ to index an advertiser's quality type. The probability of an advertiser being h-type (i.e. high-quality) is α. Similarly, we assume two types of provider quality, b_l and b_h ($b_l < b_h$), and use $y \in \{l, h\}$ to a provider's quality type. The probability of a provider being h-type (i.e., high-quality) is β.

Providers lack complete information about an advertiser's quality. We use a parameter γ_A to model the degree of information incompleteness. Specifically, providers are unaware of an advertiser's true type x but share a prior $\hat{x} \in \{l, h\}$ about the advertiser's type, which differs from the true type x with probability γ_A. That is,

$$P\left(\hat{x} = l | x = h\right) = P\left(\hat{x} = h | x = l\right) = \gamma_A,\ 0 \leq \gamma_A \leq 0.5.$$

We interpret γ_A as a provider's *probability of misclassifying an advertiser.* Similarly, advertisers share a prior $\hat{y} \in \{l, h\}$ about a provider's type (e.g., they observe the same evidence about the provider's quality) and γ_P is the advertisers' *probability of misclassifying a provider*, i.e.

$$P\left(\hat{y} = l | y = h\right) = P\left(\hat{y} = h | y = l\right) = \gamma_P,\ 0 \leq \gamma_P \leq 0.5.$$

By having parameters γ_A and γ_P, we allow various degrees of information asymmetry between advertisers and providers, ranging from complete information ($\gamma = 0$) to no information at all ($\gamma = 0.5$). In general, providers have a low misclassification rate γ_A when they accumulate much click data and uses such data to estimate advertiser quality. Advertisers have a low misclassification rate γ_P when they have extended experience with the providers. In addition, environmental uncertainty affects both rates (γ_A and γ_P): the more noise in the advertisement's performance, the more difficult it is for advertisers and providers to learn each other's quality.

Pricing schemes and auction formats. Advertising providers uses an auction to allocate their impressions. A provider can choose from three pricing schemes: PP<u>I</u>, PP<u>C</u>, and a PPI/PPC <u>h</u>ybrid.[7] Let $\mathcal{M} = \{I, C, H\}$ denote available pricing schemes.

The auction follows two general rules. First, all bidders are ranked by expected value *per impression*. This ranking rule approximates the actual ranking rules used by leading advertising providers[8] and is consistent with the ranking rules used in prior work (Wilbur and Zhu, 2009).

[7] We can easily re-interpret our results as a choice between PPI, PPS, and a PPI/PPS hybrid or between PPC, PPS, and a PPC/PPS hybrid.

[8] Google, for example, explains how different pricing schemes compete in the same auction on the basis of expected

Also, prior analyses have shown that such a ranking rule can be justified on the basis of maximizing allocative efficiency (Liu et al., 2010; Lahaie and Pennock, 2007). Second, the winner of the auction does not pay own bid but instead pays the minimum price to keep the winning position. This payment rule is again consistent with the generalized second price rule used in practice and with the prior literature (Zhu and Wilbur, 2010; Edelman et al., 2007). These two general rules, when combined with specific pricing schemes, result in the following specific auction formats:

- In the PPI auction, advertisers bid on PPI and the winner pays the second highest bid.
- In the PPC auction, advertisers bid on PPC, PPC bids are weighted by the expected quality of the advertisers, and the winner pays the minimum PPC price to keep the winning position.
- In the hybrid auction, advertisers bid on PPI or PPC. PPI bids are unweighted but PPC bids are weighted by the expected quality of the advertisers. The winner pays the minimum PPI or PPC price to keep the winning position.

It shall be noted that in the hybrid auction, the choice of PPC weighting factors affects advertisers' choice between PPC and PPI bids, which in turn affects PPC weighting factors because PPC weighting factors, calculated as expected quality of PPC bidders, are influenced by the composition of PPC bidders. If PPC weighting factors are too high or too low, the hybrid auction can degenerate to a pure PPC or pure PPI auction (Edelman and Lee, 2008). We later show that there exists a unique PPC weighting factor such that the hybrid auction does not degenerate and the weighting factor reflects PPC advertisers' expected quality. We shall use this weighting factor for hybrid auctions.

The allocative efficiency in our setting is defined as the winner's valuation for the impression, namely

$$v^{(1)}a^{(1)}b$$

where $v^{(1)}$and $a^{(1)}$ denote the winning advertiser's valuation-per-click and quality, and b denotes the chosen provider's quality. We say an auction is efficient if it always chooses the winner who values the impression the most.

The game timeline. The game proceeds as follows. At first, nature draws the quality types for all advertisers and providers, who learn their own types but not others. Next, each provider

value per impression at its FAQ page http://support.google.com/adwords/bin/answer.py?hl=en&answer=113233.

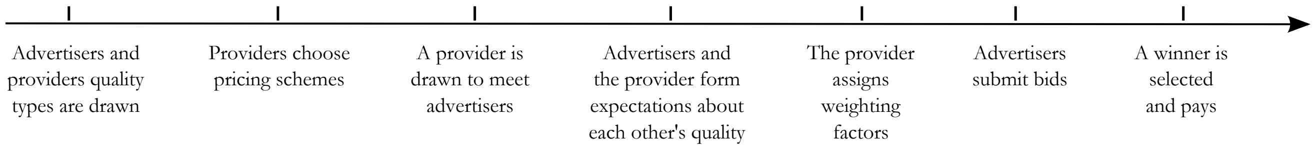

Figure 1: Game Timeline

chooses a pricing scheme. Then a single provider is chosen to auction its impression to advertisers and the provider annouces the pricing scheme. The provider and advertisers obtain priors about the other party and form expectations about the other party's quality. The provider sets the weighting factors (when applicable) based on expected quality of advertisers. Advertisers, after learning the provider's pricing scheme and their weighting factors, simultaneously submit their bids. The auction runs and a winner is selected and pays according to the payment rule. This game has two stages: at the second stage, after a provider is chosen and the pricing scheme is announced, a bidding game is played by the advertisers. At the first stage, a game is played by the providers who must choose a pricing scheme, anticipating the impact of their choice on the subsequent bidding by advertisers.

A provider chooses the pricing scheme $m \in \{I, C, H\}$ to maximize the expected revenue. Let $\sigma_y(m)$ denote the probability for a y-type provider to choose pricing scheme m. A strategy profile $\sigma = (\sigma_l, \sigma_h)$ describes the strategies of both l-type and h-type providers. When both types of providers play pure strategies (i.e., choose a pricing scheme with probability 1), we can simply write the strategy profile as $\sigma = (m_l, m_h)$. For example, (I, C) means that an l-type provider chooses PPI and an h-type provider chooses PPC.

3 Equilibrium Analysis

We use backward induction to analyze this game. We first characterize the advertisers' equilibrium bidding and derive a given provider's expected revenue under each pricing scheme. We then examine the game between providers and find their equilibrium strategy profiles.

3.1 Preliminaries

We begin by examining how a provider forms expectations about advertisers' quality. The provider does not know the advertisers' true types and thus can only form expectations based on priors. We denote $\hat{a}_{\hat{x}}$ as the expected quality of an advertiser with a quality prior $\hat{x}$. Naturally, because of misclassification, the expected quality of advertiser with an h-prior is less than a_h and the expected quality of an advertiser with an l-prior is more than a_l, i.e., (see Appendix A.A.1 for proofs of results in this section),

$$a_l \leq \hat{a}_l \leq \hat{a}_h \leq a_h \tag{2}$$

Advertisers form expectations of a provider's quality based not only on the prior $\hat{y}$ but also on the pricing scheme choice m, because different types of providers may inherently prefer different pricing schemes. We denote $\mu(\cdot|m)$ as advertisers' belief about a provider's quality type conditional on the announced pricing scheme m (but not on the prior). Given this belief and the prior $\hat{y}$, advertisers can calculate the *expected quality* of a provider $(\hat{y}, m)$, denoted as $\hat{b}_{\hat{y}}^{m}$.

Because advertisers may obtain a low or high prior about a provider, the provider also needs to know its *average expected quality*, denoted as $\hat{\hat{b}}_{y}^{m}$. When two types of providers choose a pricing scheme m with equal probability, advertisers' belief $\mu(\cdot|m)$ coincides with the natural probabilities of the two provider types. We use $\bar{b}_y$ to denote the average expected quality in this special case.

Intuitively, because advertisers may misclassify providers, the average expected quality of an h-type (l-type) provider is less (more) than the true quality, i.e.,

$$b_l \leq \hat{\hat{b}}_{l}^{m} \leq \hat{\hat{b}}_{h}^{m} \leq b_h, \forall m \in \{I, H\} \tag{3}$$

For notational convenience, we sometimes use a, $\hat{a}$, $\hat{b}$, and $\hat{\hat{b}}$ as shorthands for a_x, $\hat{a}_{\hat{x}}$, $\hat{b}_{\hat{y}}^{m}$, and $\hat{\hat{b}}_{y}^{m}$.

3.2 The Advertiser's Equilibrium Bidding

3.2.1 Equilibrium Bidding under the PPI Auction

The PPI auction can be viewed as a standard second-price auction in which a bidder's valuation for the impression is $va\hat{b}$. Because the expected quality of a provider $\hat{b}$ is the same across all advertisers, an advertiser wins the PPI auction when she has the highest valuation for the impression. Hence,

Lemma 1. *The PPI auction is efficient.*

In the PPI auction, we may characterize each advertiser using a pair (v, a). Because an advertiser wins when she has the highest va, an advertiser's *equilibrium winning probability*, denoted as

$\phi^I(v,a)$, can be calculated as,

$$\phi^I(v,a) = \left[\sum_x P(x) F\left(v\frac{a}{a_x}\right)\right]^{n-1} \tag{4}$$

where the expression $\sum_x P(x) F\left(v\frac{a}{a_x}\right)$ is the probability that advertiser (v,a) has a higher valuation for the impression than another advertiser who may be high- or low-quality.

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

$$\pi_y^I = \hat{\hat{b}}_y^I \pi_{base}^I \tag{5}$$

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

$$\pi_{base}^I \equiv n \sum_x P(x)\left[a_x \int_0^1 \phi^I(v,a_x) J(v) f(v)\, dv\right]$$

where

$$J(v) \equiv v - \frac{1-F(v)}{f(v)}$$

All proofs are deferred to the Appendix.

Recall that $\hat{\hat{b}}_y^I$ is a provider's average expected quality for choosing PPI. Proposition 1 suggests that the PPI auction revenue has a base revenue component π_{base}^I and a provider expected quality component $\hat{\hat{b}}_y^I$. Because advertisers' winning probabilities are based on the true quality, the provider's incomplete information does not affect the base revenue. On the other hand, the provider expected quality is clearly affected by advertisers' incomplete information. A high-quality provider suffers from advertiser's imperfect information and a low-quality provider benefits because of provider pooling (see (3)).

3.2.2 Equilibrium Bidding under the PPC Auction

Under the PPC auction, each bid is weighted by advertiser's expected quality $\hat{a}$. In the proof of Proposition 2, we establish that an advertiser $(v,\hat{a})$ wins the PPC auction if and only if the advertiser's weighted valuation (i.e., $v\hat{a}$) is the highest. So an advertiser $(v,\hat{a})$ has an equilibrium winning probability,

$$\phi^C(v,\hat{a}) = \left[\sum_{\hat{x}} P(\hat{x}) F\left(v\frac{\hat{a}}{\hat{a}_{\hat{x}}}\right)\right]^{n-1} \tag{6}$$

This winning probability is similar to its PPI counterpart (4) except a replacement of true quality a with the expected quality $\hat{a}$.

Because advertisers are ranked by $v\hat{a}$ instead of va, the PPC auction is generally inefficient.

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

$$\pi_y^C = b_y \pi_{base}^C \tag{7}$$

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

$$\pi_{base}^C \equiv n \sum_{\hat{x}} P(\hat{x}) \left[\hat{a}_{\hat{x}} \int_0^1 \phi^C(v, \hat{a}_{\hat{x}}) J(v) f(v) \, dv \right]$$

The PPC auction revenue similarly has a base revenue component π_{base}^C and a provider quality component b_y. In contrast with the PPI auction, the provider's incomplete information about advertisers affects the base revenue through the weighting factors; advertisers' incomplete information about the provider does not affect the provider quality component, as the provider's quality is transparent under the PPC auction.

3.2.3 Equilibrium Bidding under the Hybrid Auction

For the hybrid auction, we must also specify the weighting factors for PPC bids. We first examine how the weighting factor w for PPC bids affects an advertiser's choice.

Lemma 2. *Under a hybrid auction, an advertiser with quality a and weighting factor w will choose PPI when $a\hat{b} > w$, PPC when $a\hat{b} < w$, and be indifferent when $a\hat{b} = w$.*

When the weighting factor is high enough (i.e., $w > a_h \hat{b}$), an advertiser chooses PPC regardless of valuation or quality type. The hybrid auction degenerates into a pure PPC auction. When the weighting factor is low enough (i.e., $w < a_l \hat{b}$), an advertiser chooses PPI regardless of quality type or valuation and the hybrid auction degenerates into a pure PPI auction. Only when the weighting factor is at a moderate level, $w \in \left[a_l \hat{b}, a_h \hat{b} \right]$, two types of advertisers will separate: all high-quality advertisers choose PPI whereas all low-quality advertiser choose PPC.

We are interested in a hybrid scheme that does not degenerate. By Lemma 2, a weighting $w \in \left[a_l \hat{b}, a_h \hat{b} \right]$ will completely separate low and high-quality advertisers such that all PPC bidders are low-quality. By our auction ranking rule, the provider should assign a PPC weighting factor of

$w = a_l \hat{b}$.[9]

Assumption 1. *Under the hybrid auction, the weighting factor for PPC bids is $w = a_l \hat{b}$.*

Lemma 3. *Under the hybrid auction, if $w = a_l \hat{b}$, then (a) the equilibrium bid function strictly increases in v and (b) a high-quality advertiser (v_h, a_h) and a low-quality advertiser (v_l, a_l) tie in equilibrium if and only if*

$$v_h a_h = v_l a_l. \tag{8}$$

Lemma 3 suggests that a hybrid auction with PPC weighting factor $w = a_l \hat{b}$ is as efficient as the PPI auction because they rank the advertisers in the same way.

By Assumption 1 and Lemma 3, we can calculate the equilibrium winning probability of an advertiser (v, a) as,

$$\phi^H (v, a) = \left[\sum_x P(x) F\left(v \frac{a}{a_x} \right) \right]^{n-1} \tag{9}$$

Not surprisingly, the probability of winning for each advertiser under the hybrid scheme is exactly the same as that under the PPI auction, given that both auctions allocate the impression the same way. However, the provider's expected revenues are different under two auctions.

Proposition 3. *Given provider strategy profile σ, a y-type provider's expected revenue under a hybrid auction is:*

$$\pi_y^H = \hat{\hat{b}}_y^H \pi_h^H + b_y \pi_l^H \tag{10}$$

where

$$\begin{aligned} \pi_h^H &\equiv n\alpha a_h \int_0^1 \phi^H (v, a_h) J(v) f(v) dv \\ \pi_l^H &\equiv n(1-\alpha) a_l \int_0^1 \phi^H (v, a_l) J(v) f(v) dv \end{aligned}$$

are termed as the base revenues from high- and low-quality advertisers respectively.

Because the PPI auction and the hybrid auction allocate the impression the same way, the base revenues from two types of advertisers are in fact the same under two auctions. Using the notations

[9] Recall, as noted earlier, each advertiser's bid is weighted by her expected quality

for the hybrid auction, we may rewrite the PPI base revenue as:

$$\pi^I_{base} = \pi^H_h + \pi^H_l \tag{11}$$

A notable difference between the two auctions, though, is that under hybrid auction, the revenue from low-quality advertisers is not subject to provider pooling. Hence, we have the following result,

Corollary 1. Suppose both provider types must choose the same pricing scheme, a high-quality provider is better off under the hybrid scheme than under PPI, and a low-quality provider is worse off.

The intuition for Corollary 1 is as follows: when both provider types choose the hybrid scheme, they *partially pool* because provider types are transparent to low-quality advertisers (who choose PPC bids). In contrast, two provider types *fully pool* under the PPI scheme. High-quality providers, who suffer from pooling, prefer the partial pooling under the hybrid scheme to the full pooling under the PPI scheme. The converse is true for low-quality providers.

3.3 Provider's Equilibrium Pricing Strategy

From the previous analysis, high-quality providers have the incentive to differentiate from low-quality providers; whereas low-quality providers have the incentive to pool with high-quality ones. So high- and low-quality providers must choose their pricing schemes strategically and a game is played between providers of different quality types. This game resembles a signaling game except that in the case of PPC, a provider's type is "transparent" in the sense that the knowledge of it has no impact on advertisers' equilibrium bidding or auction revenue.

We focus on Perfect Bayesian Equilibrium (PBE) for the provider's game, which requires that the Bayesian rule determines advertisers' beliefs about a provider's quality, whenever possible. As is well known, signaling games often have multiple equilibria. To focus on a subset of equilibria that are most plausible, we make the following equilibrium refinement assumptions.

Assumption 2. A weakly-dominated strategy is not played.

Assumption 3. A Pareto-dominated equilibrium is not played.

The rationale for Assumption 2 is that weakly dominated strategies are imprudent and therefore players should avoid them. An equilibrium Pareto dominates another equilibrium if all players are not worse off and at least one is strictly better off under the former. One motivation for focusing on Pareto un-dominated equilibrium is that a Pareto-dominated equilibrium could be avoided by pre-game negotiations.

To illustrate the impact of the hybrid scheme, we examine and compare two cases: when the hybrid scheme is an option for providers and when it is not.

3.3.1 When the Hybrid Scheme is Not Allowed – the Benchmark Case

As a benchmark, we first examine the case in which providers can choose between PPI and PPC only. When $\pi^I_{base} < \pi^C_{base}$, both types of providers prefer PPC to PPI so the only equilibrium is (C, C) (see Appendix A.A.8). To focus on the more interesting case, we assume:[10]

Assumption 4. $\pi^I_{base} \geq \pi^C_{base}$

Proposition 4. Under assumptions 2, 3, and 4, (I, I) is a PBE when

$$b_h \pi^C_{base} \leq \bar{b}_h \pi^I_{base} \tag{12}$$

and(I, C) is a PBE when

$$b_h \pi^C_{base} > \bar{b}_h \pi^I_{base} \tag{13}$$

Borrowing the terminology of signaling games, we call (I, C) a *separating equilibrium* and (I, I) a *pooling equilibrium.*

In Proposition 4, a low-quality provider always chooses PPI, whereas a high-quality provider may choose either PPC or PPI depending on whether condition (13) is satisfied. The main trade-off for a high-quality provider is to pool with a low-quality provider or to separate by choosing a pay-for-performance scheme at the cost of inefficient allocations. To gain further insights on this tradeoff, we perform the following comparative static analysis:

Corollary 2. *The separating equilibrium (I, C) is more likely when*

(a) the probability of high-quality provider β is low,

(b) b_h/b_l is high,

(c) the probability of misclassifying a provider γ_P is low, and

[10] This case is more interesting for two reasons. First, because the PPI auction is more efficient than the PPC auction, the case $\pi^I_{base} \geq \pi^C_{base}$ is a first-order approximation. Although prior research (e.g., Liu and Chen 2006) argue that the opposite cases can also occur, they do so under a restrictive condition. Second, the equilibrium in which all provider types choose PPC seems inconsistent with the coexistence of multiple pricing schemes in practice.

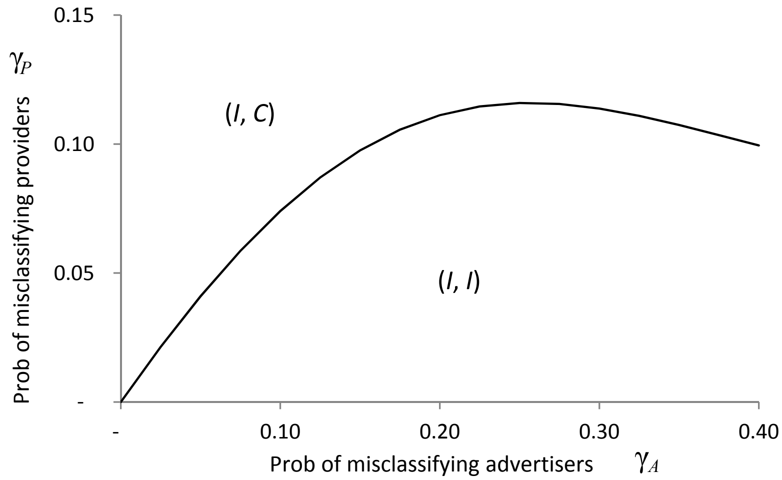

Figure 2: Provider's Equilibrium Strategy Profile without the Hybrid Scheme ($n = 6, \alpha = 0.7, a_l = 0.1, a_h = 0.2, \beta = 0.7, b_l = 0.1, b_h = 0.12$)

(d) $\pi^I_{base}/\pi^C_{base}$ *is low.*

By Corollary 2a, in a market where high-quality providers are rare, they are more likely to adopt PPC to avoid pooling with low-quality providers. Similarly, if the quality difference between two types of providers is large, a high-quality provider is more likely to adopt PPC (Corollary 2b). According to Corollary 2c, we are more likely to see PPC in a market where the providers are new and advertisers do not know their targeting qualities. Last, when the base revenue of the PPC auction is close to the PPI auction, PPC is more likely. This may happen, for example, when the weighted PPC auction effectively intensifies competition between high- and low-quality advertisers.[11]

Figure 2 illustrates the provider's equilibrium strategy profile as a function of misclassification probabilities. As the probability of misclassifying providers γ_P increases, a high-quality provider is more likely to adopt PPC. As the probability of misclassifying advertisers γ_A increases, a high-quality provider is generally more likely to adopt PPI. However, when γ_A is too high, a high-quality provider may once again prefer PPC because misclassifying advertisers creates a sizable incentive effect (see footnote 11), which offsets some of the revenue loss from inefficient allocation.

[11]In the PPC auction, because of misclassification, low-quality advertisers tend to receive a higher weighting factor and high-quality ones tend to receive a lower weighting factor. As a result, the competition between low- and high-quality advertisers intensifies, forcing high-quality advertisers to bid higher. Such an incentive effect is most prominent under low competition for high-quality advertisers.

3.3.2 When the Hybrid Scheme is Allowed

Lemma 4. Under assumptions 2, 3, 4, and 1,

(a)(I, I) with belief that a hybrid provider is high quality with probability no more than μ_1 (i.e., $\mu(h|H) \leq \mu_1$) is a PBE when

$$b_h \pi^C_{base} \quad \leq \quad \bar{b}_h \pi^I_{base} \tag{14}$$

where μ_1 is the maximum probability to satisfy the following two conditions:

$$\bar{b}_h \pi^I_{base} \quad \geq \quad \hat{\hat{b}}^H_h \pi^H_h + b_h \pi^H_l \tag{15}$$

$$\bar{b}_l \pi^I_{base} \quad \geq \quad \hat{\hat{b}}^H_l \pi^H_h + b_l \pi^H_l \tag{16}$$

(b)(H, H) with belief that a PPI provider is high quality with probability no more than μ_2 (i.e., $\mu(h|I) \leq \mu_2$) is a PBE when

$$b_h \pi^C_{base} \quad \leq \quad \bar{b}_h \pi^H_h + b_h \pi^H_l \tag{17}$$

and where μ_2 is the maximum probability to satisfy the following two conditions,

$$\hat{\hat{b}}^I_h \pi^I_{base} \quad \leq \quad \bar{b}_h \pi^H_h + b_h \pi^H_l \tag{18}$$

$$\hat{\hat{b}}^I_l \pi^I_{base} \quad \leq \quad \bar{b}_l \pi^H_h + b_l \pi^H_l \tag{19}$$

(c) (I, C) with belief that a hybrid provider is low quality (i.e.,$\mu(h|H) = 0$) and (H, C) with belief that a PPI provider is low quality (i.e, $\mu(h|I) = 0$) are PBE when

$$b_h \pi^C_{base} \quad > \quad \bar{b}_h \pi^H_h + b_h \pi^H_l \tag{20}$$

By Lemma 4, (I, I) is an equilibrium when condition (14) holds and advertisers believe a hybrid provider is most likely low-quality (i.e., $\mu(h|H) < \mu_1$). Similarly, (H, H) is an equilibrium when condition (17) holds and advertisers believe a PPI provider is most likely low-quality. The equilibrium (I, C) requires that condition (20) holds and advertisers believe a PPI provider is low quality.

Among the three types of equilibria, (I, I) is "unstable" in the following sense. We can infer from condition (15) that (I, I) requires advertisers to believe that a low-quality provider is more likely to deviate to the hybrid scheme than a high-quality one.[12] However, this is implausible because a low-quality provider prefers (I, I) and a high-quality provider prefers (H, H) (Corollary 1). We formalize this notion of stability in an "uncompromised equilibrium" concept.

Definition 1. An *uncompromised equilibrium* requires that if an off-equilibrium action a in the current equilibrium is played by a player type t with probability p in an alternative equilibrium and t is better off in that equilibrium, then other players should assign a probability of no less than p to player type t for playing a in the current equilibrium. If the current equilibrium fails to hold under such an off-equilibrium belief, we say the current equilibrium is *compromised* by the alternative equilibrium.

The uncompromised equilibrium is a generalization of Mailath et al.'s(1993) *undefeated equilibrium,* which says that if *all* the player types who play the off-equilibrium action in an alternative equilibrium are strictly better off in that equilibrium, the probability assigned to each player type must be *exactly* according to the alternative equilibrium. Both the undefeated equilibrium and the uncompromised equilibrium view off-equilibrium messages as attempts to overturn the current equilibrium by player types who seek a better alternative. The difference lies in that our definition extends to the case where an off-equilibrium action is played by both better-off and worse-off player types in an alternative equilibrium.

Proposition 5. Under assumptions 2, 3, 4, and 1, (H, H) with belief that a PPI provider is high quality with probability no more than μ_2(i.e., $\mu(h|I) \leq \mu_2$, where μ_2 is defined in (4)) is an uncompromised PBE when

$$b_h \pi^C_{base} \quad \leq \quad \bar{b}_h \pi^H_h + b_h \pi^H_l. \tag{21}$$

(I, C) with belief that a hybrid provider is low quality and (H, C) with belief that a PPI provider is low quality are uncompromised PBE when

$$b_h \pi^C_{base} \quad > \quad \bar{b}_h \pi^H_h + b_h \pi^H_l. \tag{22}$$

[12]To see this, we need to show $\mu(h|H) < \beta$. If $\mu(h|H) \geq \beta$, we have $\hat{\hat{b}}^H_h \geq \bar{b}_h$. So $\bar{b}_h \pi^I_{base} = \bar{b}_h \left(\pi^H_h + \pi^H_l\right) \leq \hat{\hat{b}}^H_h \left(\pi^H_h + \pi^H_l\right) \leq \hat{\hat{b}}^H_h \pi^H_h + b_h \pi^H_l$ (one of the inequality must be strict), which contradicts (15).

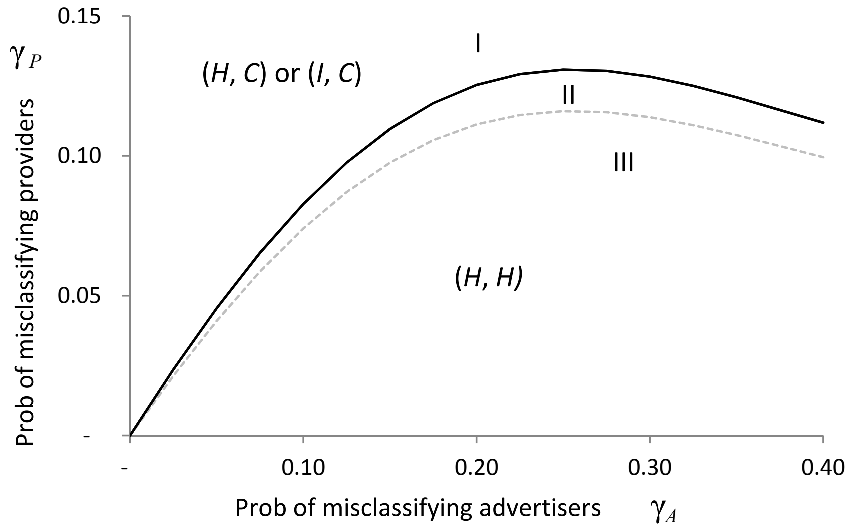

Figure 3: Provider's Equilibrium Strategy Profile with the Hybrid Scheme

As in the benchmark case (Proposition 5), the provider's game has a separating equilibrium, in the form of (I, C) or (H, C), and a pooling equilibrium, in the form of (H, H). The separating equilibria (I, C) and (H, C) are effectively the same for all parties because in either case, advertisers correctly infer a provider's quality from the pricing scheme, and the two auctions allocate the advertising impressions exactly the same way and generate the same revenues by Proposition 3.

Remarkably, Proposition 5 suggests that with the hybrid scheme, (H, H) replaces (I, I) as the new pooling equilibrium. This implies that the hybrid scheme can dominate the PPI scheme in equilibrium, even though a hybrid auction is no more efficient or profitable than a PPI auction when examined as a standalone auction. To our knowledge, this is the first paper to rationalize the use of the hybrid scheme as an equilibrium choice. The intuition for this result is as follows. The hybrid auction is as efficient as the PPI auction (Proposition 3), but the cost to pool for a high-quality provider is lower under the hybrid scheme because of partial pooling. So a high-quality provider prefers (H, H) whereas a low-quality provider prefers (I, I). (H, H) is selected in equilibrium because advertisers believes that a low-quality provider is more likely to deviate from (H, H) to (I, I) and such a belief deters deviation. In contrast, the (I, I) equilibrium is compromised by the belief that a deviator is more likely a high-quality provider.

Proposition 5 also suggests that the two provider types pool more often than without the hybrid scheme – note that condition (21) is more relaxed than condition (12). Intuitively, because high-quality providers find it less costly to pool under (H, H), they rely less on the costly PPC auction to differentiate.

	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)
l- & h-type provider's expected revenue	$=, =$	$+, +$	$-, +$
l- & h-type advertiser's expected payoff	$=, =$	+/–, +/–	$=, =$
overall efficiency	$=$	$+$	$=$

$+$: increase, $-$: decrease, +/–: increase or decrease,$=$: same

Table 2: Impact of the Hybrid Scheme

Figure 3 shows provider's equilibrium strategy profile as a function of misclassification rates. 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 the boundary in the benchmark case. The shape of the new boundary (solid line) is similar to that in the benchmark case: a high-quality provider is more likely to adopt PPC when the probability of misclassifying a provider γ_P increases and when the probability of misclassifying advertisers γ_A decreases (except when it is very high). The shift in the boundary before and after the introduction of the hybrid scheme suggests that two provider types pool more often after the hybrid scheme is introduced.

We summarize the impact of the hybrid scheme in the following Corollary and table.

Corollary 3. *By adding the hybrid scheme to provider's options, we find that*

(a) (I, C) is replaced by (H, C) or (I, C) in some cases (region I of Figure 3), which leaves providers, advertisers, and overall allocative efficiency unaffected.

(b) (I, C) is replaced by (H, H) in other cases (region II of Figure 3), which benefits all the providers and improves overall allocative efficiency.

(c) (I, I) is replaced by (H, H) (region III of Figure 3), which benefits a high-quality provider, hurts a low-quality provider, and does not affect advertisers' expected payoff or overall allocative efficiency.

4 Discussion and Conclusions

Our study is among the first to examine the optimality of different pricing schemes in the presence of information asymmetry between providers and advertisers. Our findings illustrate how providers can use different pricing schemes not only to rank advertisers differently but also to reveal their targeting quality. Because advertisers have incomplete information about provider quality, low-quality providers can benefit from pooling using PPI, whereas high-quality providers benefit from separating themselves using more transparent pay-for-performance schemes. On the other hand, given the incomplete information about advertiser quality, the use of pay-for-performance schemes creates allocative inefficiencies. This tradeoff between the cost and benefits of the different pricing schemes lies at the heart of our analyses.

Our analyses show that high-quality providers find it optimal to choose the pay-for-performance scheme (e.g., PPC and PPS) when the benefits of separating from low-quality providers outweigh the loss from allocative inefficiency. Thus, a nascent high-quality provider may benefit from the transparency afforded by pay-for-performance schemes, as its quality is largely unknown to advertisers. An established high-quality provider may also benefit from pay-for-performance schemes, but for a different reason: it may accumulate enough data to make good estimates about advertiser quality so that the cost of differentiation is low. Examples of the latter include Google Search Ads and Amazon Product Ads.

Our findings also suggest that high-quality providers may prefer a hybrid pricing scheme to a pure pricing scheme. While offering a hybrid scheme reduces a high-quality provider's ability to separate, it still benefits the provider. First, any loss from pooling with low-quality providers is reduced because only a few advertisers - those who submit PPI bids - need to estimate the quality of the provider. Secondly, the allocative efficiency of a hybrid scheme is similar to that of a pure PPI scheme because low- and high-quality advertisers self-select into different pricing schemes. A high-quality provider's adoption of hybrid schemes also forces low-quality providers to follow suit, leading to an equilibrium where the hybrid scheme prevails. That said, the hybrid scheme is more complex than pure pricing schemes and upfront costs, including setting up the automated auction platform and communicating to advertisers about the changes, can be a challenge for low-quality providers who have no experience with auction platforms or pay-for-performance pricing. This only makes hybrid schemes more attractive for high-quality providers as they can then use the hybrid scheme to differentiate themselves. The increased adoption of hybrid schemes in recent years (e.g., Google Display Ads, Facebook, Twitter, and adBrite) seems to confirm our prediction that the hybrid scheme may be a strong alternative to pure schemes among high-quality providers.

Our research highlights the role of information asymmetry in determining pricing schemes. In a world of complete information, where providers perfectly know advertiser's quality and vice versa, different pricing schemes would be equivalent.[13] With information asymmetry, however, pricing

[13]This can be easily seen from our existing results (Propositions 1, 2, and 3), noting that when estimated quality coincides with true quality, three auctions order advertisers in the same way and generate identical revenues for advertising providers.

schemes can be used as a means of leveraging private information: PPI leverages advertisers' private information; PPC leverages providers'; A hybrid scheme leverages both parties' to some degree. Thus pricing schemes differ in how the responsibility of monitoring (or estimating) is divided among providers and advertisers. This theoretical perspective provides a useful alternative to the conventional "risk-sharing" model of pricing schemes where the choice of pricing schemes is driven by the degree of risk aversion of players.

Our framework and findings highlight two important technologies in Internet advertising pricing: the technology for targeting Internet users with the most relevant advertisements and the technology for monitoring the quality of advertisers. The targeting technology is necessary for providers to become high quality. Thus, a large premium portal (e.g., a major newspaper site) without targeting is considered low quality, and prefers the conventional PPI scheme. In contrast, an advertising network of many small, "inferior" websites that uses superior targeting technologies (e.g., Google AdSense) is high quality; such networks frequently use pay-for-performance schemes. Meanwhile, the monitoring technology makes pay-for-performance schemes less costly by minimizing allocative inefficiency. Google can afford pure PPC for its search advertising partly because it has data and technology prowess to achieve reliable estimates of advertiser quality.

In contrast with provider technologies, technologies that help advertisers track provider quality have an opposite effect: as provider quality becomes more transparent, the need for high-quality providers to separate diminishes; all providers would opt for conventional PPI. However, the effect of provider technologies likely subdues this effect because in the Internet advertising industry, the concentration is much higher on the provider side, giving providers advantages in adopting and refining their technologies.

Whether a high-quality provider should adopt a pure pay-for-performance scheme would depend not only on the monitoring technology, but also on two parties' relative contribution to advertisement performance. If advertisement performance is influenced more by provider's quality, revealing one's high quality would be more important than efficiently sorting advertisers, thus pay-for-performance scheme is more likely to be used. Conversely, if advertiser quality has more impact on advertisement performance, PPI or a hybrid scheme is more likely used to leverage advertisers' private information and preserve allocative efficiency. Search engines and many broad-based advertising networks would not offer PPS because advertisers play a much more important role than advertising providers in converting clicks to sales.

Our study is a first step in understanding the role of information asymmetry in the evolving landscape of pricing schemes for digital advertising. Our findings pave the way for a number of possible extensions. Our model focuses on a representative provider drawn from a heterogeneous pool. A natural extension of our model is to examine the role of side-by-side provider competition. Another way to extend existing insights is to understand the dynamics of pricing scheme choices. Our study models issues of hidden information rather than hidden action. It would be interesting to consider the impacts of hidden action in addition to hidden information; Zhu and Wilbur (2010) and Chen et al. (2009) provide some clues in this direction. Additionally, other factors may also affect pricing scheme choices. For example, Athey and Ellison (2011) shows that PPC auctions suffer additional inefficiencies due to advertiser obfuscation; some argue that click fraud threatens the PPC pricing although research shows (Wilbur and Zhu, 2009) that the impact of click fraud on advertiser and search engine revenues may be ambiguous. In using our findings, it is also important to combine our insights with risk preference considerations: ceteris paribus, risk-averse advertisers would prefer PPC, and risk-averse providers, PPI. While including risk aversion is likely to shift the optimality conditions, it has a predictable impact and is unlikely to reverse the broad findings of our analyses. Finally, given the increased availability of data on advertising pricing schemes, empirical tests of these theoretical models would add to our understanding of this complex landscape.

References

Vibhanshu Abhishek, Kartik Hosanagar, and Peter S. Fader. Identifying the Aggregation Bias in Sponsored Search Data. *Working Paper*, 2010.

A. Agarwal, K. Hosanagar, and M.D. Smith. Location, location, location: An analysis of profitability of position in online advertising markets. *Journal of Marketing Research*, 2011.

A. Animesh, S. Viswanathan, and R. Agarwal. Competing "Creatively" in Sponsored Search Markets: The Effect of Rank, Differentiation Strategy, and Competition on Performance. *Information Systems Research*, 22(1):153–169, February 2011.

Animesh Animesh, Vandana Ramachandran, and Siva Viswanathan. Quality Uncertainty and the Performance of Online Sponsored Search Markets: An Empirical Investigation. *Information Systems Research*, 21(1):190–201, June 2009.

S. Athey and G. Ellison. Position auctions with consumer search. *Quarterly Journal of Economics*, 126(3):1213–1270, 2011.

TCA Bashyam. Service design and price competition in business information services. *Operations Research*, 48(3):362–375, 2000.

Jianqing Chen, De Liu, and Andrew B Whinston. Auctioning Keywords in Online Search. *Journal of Marketing*, 73(4):125–141, July 2009.

Vidyanand Choudhary. Use of pricing schemes for differentiating information goods. *Information Systems Research*, 21(1):78–92, March 2010.

Sanjeev Dewan and Haim Mendelson. User delay costs and internal pricing for a service facility. *Management Science*, 36(12):1502–1517, 1990.

B. Edelman and M. Ostrovsky. Strategic bidder behavior in sponsored search auctions. *Decision support systems*, 43(1):192–198, 2007.

Benjamin Edelman and HS Lee. CPC/CPA hybrid bidding in a second price auction. *Harvard Business School Working Paper*, 2008.

Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review*, 97(1):242–259, March 2007.

Juan Feng. Optimal mechanism for selling a set of commonly ranked objects. *Marketing Science*, 27(3):501–512, May 2008.

A. Ghose and S. Yang. An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets. *Management Science*, 55(10):1605–1622, July 2009.

Ashish Goel and Kamesh Munagala. Hybrid Keyword Search Auctions. *Proceedings of the 18th international conference on World Wide Web*, pages 1–18, July 2009.

A Gupta and DO Stahl. A stochastic equilibrium model of Internet pricing. *Journal of Economic Dynamics and Control*, 21:697–722, 1997.

A. Gupta, B. Jukic, D. O. Stahl, and a. B. Whinston. An Analysis of Incentives for Network Infrastructure Investment Under Different Pricing Strategies. *Information Systems Research*, 22 (1):215–232, January 2011.

Kartik Hosanagar, Ramayya Krishnan, John Chuang, and Vidyanand Choudhary. Pricing and Resource Allocation in Caching Services with Multiple Levels of Quality of Service. *Management Science*, 51(12):1844–1859, December 2005.

Y. Hu, J. Shin, and Z. Tang. Pricing of Online Advertising: Cost-Per-Click-Through Vs. Cost-Per-Action. *Proceedings of the 43rd Hawaii International Conference on System Sciences*, 2010.

Y.J. Hu. Performance-based pricing models in online advertising. *Working paper*, 2004.

S. Jain and PK Kannan. Pricing of information products on online servers: Issues, models, and analysis. *Management Science*, 48(9):1123–1142, 2002.

P. Jeziorski and I. Segal. What makes them click: Empirical analysis of consumer demand for search advertising. *Working Paper, Stanford University*, 2009.

Sébastien Lahaie and David M. Pennock. Revenue analysis of a family of ranking rules for keyword auctions. *Proceedings of the 8th ACM conference on Electronic commerce - EC '07*, page 50, 2007.

D Liu and J Chen. Designing online auctions with past performance information. *Decision Support Systems*, 42(3):1307–1320, December 2006.

D. Liu, J. Chen, and A. B. Whinston. Ex Ante Information and the Design of Keyword Auctions. *Information Systems Research*, 21(1):133–153, November 2010.

G.J. Mailath, M. Okuno-Fujiwara, and A. Postlewaite. On belief based refinements in signaling games. *Journal of Economic Theory*, 60:241–276, 1993.

Yasushi Masuda and Seungjin Whang. On the Optimality of Fixed-up-to Tariff for Telecommunications Service. *Information Systems Research*, 17(3):247–253, September 2006.

Haim Mendelson. Pricing Computer Services : Queueing Effects. *Communications of the ACM*, 28 (3):312–321, 1985.

R. B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, February 1981.

O. Rutz and R.E. Bucklin. A model of individual keyword performance in paid search advertising. *SSRN eLibrary*, 2007.

Arun Sundararajan. Pricing digital marketing: Information, risk sharing and performance. *Leonard N. Stern School of Business Working Paper*, 2003.

Arun Sundararajan. Nonlinear Pricing of Information Goods. *Management Science*, 50(12):1660–1673, December 2004.

The Economist. Pay per sale: The holy grail of advertising is within reach. *The Economist*, (950), 2005.

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

T. A. Weber and Z. Zheng. A Model of Search Intermediaries and Paid Referrals. *Information Systems Research*, 18(4):414–436, December 2007.

Kenneth C. Wilbur and Yi Zhu. Click Fraud. *Marketing Science*, 28(2):293–308, March 2009.

Lizhen Xu, Jianqing Chen, and Andrew Whinston. Oligopolistic Pricing with Online Search. *Journal of Management Information Systems*, 27(3):111–142, January 2010.

Lizhen Xu, Jianqing Chen, and Andrew Whinston. Price Competition and endogenous Valuation in search advertising. *Journal of Marketing Research*, XLVIII(June):566–586, 2011.

Song Yao and Carl F. Mela. A dynamic model of sponsored search advertising. *Marketing Science*, 30(3), 2011.

X. Zhang and J. Feng. Cyclical Bid Adjustments in Search-Engine Advertising. *Management Science*, 57(9):1703–1719, August 2011.

Yi Zhu and K.C. Wilbur. Hybrid Advertising Auctions. *Marketing Science*, 30(2), 2010.

A Appendix

A.1 Formulas Used in Section 3.1

Let $P(\hat{x}|x)$ denote the probability that an x-type advertiser has a prior $\hat{x}$. The probability of an $\hat{x}$-prior advertiser is

$$P(\hat{x}) = \sum_x P(\hat{x}|x) P(x)$$

The expected quality of an $\hat{x}$-prior advertiser is calculated as

$$\hat{a}_{\hat{x}} = \frac{1}{P(\hat{x})} \sum_x P(\hat{x}|x) P(x) a_x \tag{23}$$

Because $\hat{a}_{\hat{x}}$ is a weighted average of a_l and a_h, we have $a_l \leq \hat{a}_{\hat{x}} \leq a_h$. We can also easily verify that $\hat{a}_l \leq \hat{a}_h$ as long as $\gamma_A \leq 0.5$.

Table 3 shows a numerical example of advertiser's expected quality.

x or $\hat{x}$	$P(x)$	a_x	$P(\hat{x})$	$\hat{a}_{\hat{x}}$
l	0.7	0.1	0.62	0.11
h	0.3	0.2	0.38	0.163

Table 3: Advertiser's Expected Quality ($\gamma_A = 0.2$)

We now turn to a provider's expected quality. When a pricing scheme m is played in the equilibrium, we can calculate the probability of a provider being a y-type, conditional on choosing m, using the Bayesian rule:

$$\mu(y|m) = \frac{P(y)\sigma_y(m)}{\sum_{y'} P(y')\sigma_{y'}(m)}.$$

Given an advertiser's belief $\mu(y|m)$, the probability of a provider being $\hat{y}$-prior, conditional on the pricing scheme m, is

$$P(\hat{y}|m) = \sum_y \mu(y|m) P(\hat{y}|y)$$

and the expected quality of a $\hat{y}$-prior provider who chooses pricing scheme m is

$$\hat{b}_{\hat{y}}^{m} = \frac{1}{P(\hat{y}|m)} \sum_{y} \mu(y|m) P(\hat{y}|y) b_y \tag{24}$$

A provider's average expected quality, given the true type y and the pricing scheme choice m, is calculated as

$$\hat{\hat{b}}_{y}^{m} = \sum_{\hat{y}} P(\hat{y}|y) \hat{b}_{\hat{y}}^{m} \tag{25}$$

and $\bar{b}_y$ is a special case of $\hat{\hat{b}}_{y}^{m}$ with $\mu(h|m) = \beta$.

The proof of $b_l \leq \hat{\hat{b}}_{l}^{m} \leq \hat{\hat{b}}_{h}^{m} \leq b_h$ is similar to that for advertiser quality and thus omitted.

Table 4 shows a numerical example of how advertiser's belief $\mu(h|m)$ affects provider's average expected quality.

			$\mu(h\|m) = 0.3$			$\mu(h\|m) = 0.7$		
y or $\hat{y}$	$P(y)$	b_y	$P(\hat{y})$	$\hat{b}_{\hat{y}}$	$\hat{\hat{b}}_y$	$P(\hat{y})$	$\hat{b}_{\hat{y}}$	$\hat{\hat{b}}_y$
l	0.7	0.1	0.62	0.11	0.12	0.38	0.137	0.148
h	0.3	0.2	0.38	0.163	0.152	0.62	0.19	0.180

Table 4: Provider's expected quality and average expected quality ($\gamma_A = 0.2$)

A.2 Proof of Lemma 1

From the preceding discussion, advertisers are ranked by va. Thus, the PPI auction always selects the advertiser who has the highest valuation for the impression. By definition, the PPI auction is efficient.

A.3 Proof of Proposition 1

Omitted. Available on request.

A.4 Proof of Proposition 2

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), i.e.,

$$v' > v \Rightarrow t(v', \hat{a}) > t(v, \hat{a}). \tag{26}$$

We next show that advertisers$(v, \hat{a}_l)$ and $\left(\frac{\hat{a}_l}{\hat{a}_h} v, \hat{a}_h\right)$ tie in equilibrium, i.e.,

$$\frac{\hat{a}_l}{\hat{a}_h} t(v, \hat{a}_l) = t\left(\frac{\hat{a}_l}{\hat{a}_h} v, \hat{a}_h\right) \tag{27}$$

With a slight abuse of notation, we denote $U(v,t|\hat{a})$ as the per-click payoff of an advertiser $(v,\hat{a})$ who bids t per click. We also denote $\phi^C(t,\hat{a})$ and $p(t,\hat{a})$ as the winning probability and expected payment per click of an advertiser with expect quality $\hat{a}$ and bid t. Now let advertiser$(v,\hat{a}_l)$ bid t and advertiser$\left(\frac{\hat{a}_l}{\hat{a}_h}v,\hat{a}_h\right)$ bid $\frac{\hat{a}_l}{\hat{a}_h}t$.

$$\begin{aligned}
U\left(\frac{\hat{a}_l}{\hat{a}_h}v,\frac{\hat{a}_l}{\hat{a}_h}t|\hat{a}_h\right) &= \phi^C\left(\frac{\hat{a}_l}{\hat{a}_h}t,\hat{a}_h\right)\frac{\hat{a}_l}{\hat{a}_h}v - p\left(\frac{\hat{a}_l}{\hat{a}_h}t,\hat{a}_h\right) \\
&= \phi^C(t,\hat{a}_l)\frac{\hat{a}_l}{\hat{a}_h}v - \frac{\hat{a}_l}{\hat{a}_h}p(t,\hat{a}_l) \\
&= \frac{\hat{a}_l}{\hat{a}_h}U(v,t|\hat{a}_l) \qquad (28)
\end{aligned}$$

where the second equality is because two advertisers tie and pay the same expected amount. (28) implies that if t maximizes $U(v,t|\hat{a}_l)$ then $\frac{\hat{a}_l}{\hat{a}_h}t$ must maximize $U\left(\frac{\hat{a}_l}{\hat{a}_h}v,\frac{\hat{a}_l}{\hat{a}_h}t|\hat{a}_h\right)$, thus proving (27).

(26) and (27) together imply that an advertiser $(v,\hat{a})$ wins if and only if the advertiser's $v\hat{a}$ is the highest.

By the revenue equivalence theorem (Myerson, 1981), we can obtain the equilibrium payoff (per click) of advertiser $(v,\hat{a})$ as,

$$V(v,\hat{a}) = \int_0^v \phi^C(u,\hat{a})\,du \qquad (29)$$

Because the expected payment per click from the advertiser is

$$p(v,\hat{a}) = \phi^C(v,\hat{a})\,v - V(v,\hat{a}),$$

a y-type provider's expected revenue from all advertisers is

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

where the third equality is due to integration by parts.

A.5 Proof of Lemma 2

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

$$U^I\left(v,v'\right) \;=\; \phi^H\left(v'a\hat{b}\right)\left[va\hat{b}-s\left(v'a\hat{b}\right)\right] \tag{30}$$

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

$$U^C\left(v,v'\right)=a\hat{b}\phi^H\left(v'w\right)\left[v-s\left(v'w\right)/w\right] \tag{31}$$

Clearly, when $a\hat{b}=w$, $U^I(v,v')=U^C(v,v')$, which means that the advertiser is indifferent between the two pricing schemes. Now suppose $a\hat{b}<w$. Let $v^*a\hat{b}$ be the advertiser's optimal PPI bid. By submitting a PPC bid of $v^*a\hat{b}/w$, the advertiser gets an expected utility:

$$U^C\left(v,v^*a\hat{b}/w\right)=\phi^H\left(v^*a\hat{b}\right)\left[va\hat{b}-s\left(v^*a\hat{b}\right)a\hat{b}/w\right]>U^I\left(v,v^*a\hat{b}\right)$$

So the advertiser prefers PPC when $a\hat{b}<w$. Similarly, we can show that when $a\hat{b}>w$, the advertiser prefers PPI.

A.6 Proof of Lemma 3

By Lemma **2**, when $w=a_l\hat{b}$, high-quality advertisers choose PPI while low-quality advertisers choose PPC. The proof for a strictly increasing equilibrium bidding function under the hybrid scheme is analogous to that under PPC and is available on request. We focus on (b). Consider a high-quality advertiser (v_h,a_h) who submits a PPI bid $ta_l\hat{b}$ and a low-quality advertiser (v_l,a_l) who submits a PPC bid t. By equations (30) and (31), the two advertisers' expected payoffs, denoted as U_h and U_l respectively, are,

$$\begin{aligned} U_h &= \phi^H\left(ta_l\hat{b}\right)\left[v_ha_h\hat{b}-s\left(ta_l\hat{b}\right)\right] && (32)\\ U_l &= a_l\hat{b}\phi^H\left(ta_l\hat{b}\right)\left[v_l-s\left(ta_l\hat{b}\right)/\left(a_l\hat{b}\right)\right] \\ &= \phi^H\left(ta_l\hat{b}\right)\left[v_la_l\hat{b}-s\left(ta_l\hat{b}\right)\right] && (33)\end{aligned}$$

When $v_la_l=v_ha_h$, we have $U_h=U_l$. So if t is the optimal bid for the low-quality advertiser, then $ta_l\hat{b}$ must be the optimal bid for the high-quality advertiser, suggesting that two advertisers will tie in equilibrium. Given that bidding functions strictly increases, there cannot be an advertiser with (v_l',a_l), $v_l'\neq v_l$ that ties with (v_h,a_h). So the condition (8) is also a necessary condition.

A.7 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, we get

$$V_l^H(v)=\int_0^v \phi^H(u,a_l)\,du \tag{34}$$

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

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

The provider's expected revenue from low-quality advertisers is the sum of expected payment per impression of all low-quality advertisers, i.e.,

$$n(1-\alpha)\,a_l b_y \int_0^1 \phi^H(v,a_l)\,J(v)\,f(v)\,dv \tag{35}$$

By a similar process, we can obtain the expected payment per impression from a high-quality advertiser as

$$a_h \hat{b}_{\hat{y}} \int_0^1 \phi^H(v,a_h)\,J(v)\,f(v)\,dv$$

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

$$n\alpha a_h \hat{\hat{b}}_y^H \int_0^1 \phi^H(v,a_h)\,J(v)\,f(v)\,dv \tag{36}$$

Adding (35) and (36) together, we can get the y-type provider's total expected revenue from all advertisers (10).

A.8 When $\pi^I_{base} < \pi^C_{base}$, (C,C) is the only PBE

A high-quality provider gets $b_h\pi^C_{base}$ under PPC and $\hat{\hat{b}}_h\pi^I_{base}$ under PPI. Because $\hat{\hat{b}}_h \le b_h$ and $\pi^I_{base} < \pi^C_{base}$, it is a dominant strategy for a high-quality provider to choose PPC. Given high-quality provider's strategy, a low-quality provider gets $b_l\pi^I_{base}$ under PPI and $b_l\pi^C_{base}$ under PPC. So a low-quality provider strictly prefers PPC. This suggests that (C,C) is the only PBE.

A.9 Proof of Proposition 4

Under Assumption 4, a low-quality provider gets an expected revenue of $\hat{\hat{b}}_l\pi^I_{base}$ under PPI and $b_l\pi^C_{base}$ under PPC. Because $\hat{\hat{b}}_l \ge b_l$, PPC is weakly dominated and thus not played by low-quality providers. So 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 provider optimally chooses PPC. A high-quality provider's expected revenue under PPC is $b_h\pi^C_{base}$. By deviating to PPI, the provider gets $b_l\pi^I_{base}$. A high-quality provider will not deviate if

$$b_h\pi^C_{base} > b_l\pi^I_{base} \tag{37}$$

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

$$b_h \pi^C_{base} \leq \bar{b}_h \pi^I_{base} \tag{38}$$

Because $b_l < \bar{b}_h$, conditions (37) and (38) may hold at the same time. We next show that (I, C) is Pareto dominated when both (37) and (38) hold. Under (I, I), the expected revenues for the high- and low-quality providers are $\bar{b}_h \pi^I_{base}$ and $\bar{b}_l \pi^I_{base}$, respectively. Under (I, C), their expected revenues are $b_h \pi^C_{base}$ and $b_l \pi^I_{base}$ respectively. Because $\bar{b}_h \pi^I_{base} \geq b_h \pi^C_{base}$ by condition (38) and $\bar{b}_l > b_l$ by (3), (I, C) is Pareto dominated by (I, I) and thus not played by Assumption 3.

A.10 Proof of Corollary 2

The proof of these results is straightforward from condition (13), thus omitted.

A.11 Proof of Lemma 4

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

(a) A low-quality provider gets $\bar{b}_l \pi^I_{base}$ under (I, I) and $\hat{\hat{b}}^H_l \pi^H_h + b_l \pi^H_l$ by deviating to the hybrid scheme. A high-quality provider gets $\bar{b}_h \pi^I_{base}$ under (I, I), $\hat{\hat{b}}^H_h \pi^H_h + b_h \pi^H_l$ by deviating to the hybrid scheme, and $b_h \pi^C_{base}$ by deviating to PPC. The conditions (14), (15), and (16) ensure all deviations are unprofitable. Notice that $\hat{\hat{b}}^H_h$ and $\hat{\hat{b}}^H_l$ increase with $\mu(h|H)$ so we can combine (15) and (16) into a requirement on belief $\mu(h|H)$.

(b) A low-quality provider gets $\bar{b}_l \pi^H_h + b_l \pi^H_l$ under (H, H) and $\hat{\hat{b}}^I_l \pi^I_{base}$ by deviating to PPI. A high-quality provider gets $\bar{b}_h \pi^H_h + b_h \pi^H_l$ under (H, H), $\hat{\hat{b}}^I_h \pi^I_{base}$ by deviating to PPI, and $b_h \pi^C_{base}$ by deviating to PPC. Conditions (17), (18), and (19) ensure that all deviations are unprofitable. Similar to (a), (15) and (16) can be represented as a requirement on belief $\mu(h|I)$.

(c) A low-quality provider gets $b_l \pi^I_{base}$ under (I, C) and $\hat{\hat{b}}^H_l \pi^H_h + b_l \pi^H_l$ by deviating to the hybrid scheme. Recall that $\pi^I_{base} = \pi^H_h + \pi^H_l$, the deviation is unprofitable only when $\hat{\hat{b}}^H_l = b_l$, which implies $\mu(h|H) = 0$. A high-quality provider gets $b_h \pi^C_{base}$ under (I, C), $b_l \pi^I_{base}$ by deviating to PPI, and $b_l \pi^H_h + b_h \pi^H_l$ (notice that $\mu(h|H) = 0$) by deviating to the hybrid scheme. For the high-quality provider, deviating to PPI is clearly dominated by deviating to the hybrid scheme. The belief

$\mu(h|H) = 0$ and the following condition ensure all deviations are unprofitable

$$b_h \pi_{base}^C \geq b_l \pi_h^H + b_h \pi_l^H. \tag{39}$$

(d) A low-quality provider gets $b_l \pi_{base}^I$ under (H, C) and $\hat{\hat{b}}_l^I \pi_{base}^I$ by deviating to PPI. The deviation is unprofitable only when $\hat{\hat{b}}_l^I = b_l$, which implies $\mu(h|I) = 0$. A high-quality provider gets $b_h \pi_{base}^C$ under (H, C), $b_l \pi_{base}^I$ (with $\mu(h|I) = 0$) by deviating to PPI, and $b_l \pi_h^H + b_h \pi_l^H$ by deviating to the hybrid scheme. Because $\pi_{base}^I = \pi_h^H + \pi_l^H$, deviating to PPI is clearly dominated by deviating to the hybrid scheme. The belief $\mu(h|I) = 0$ and condition (39) ensure all deviations are unprofitable.

When (17) and (39) hold simultaneously, (I, C), (H, C), and (H, H) are all PBEs. The equilibrium revenues under three PBEs are $\left(b_l \pi_{base}^I, b_h \pi_{base}^C\right)$, $\left(b_l \pi_{base}^I, b_h \pi_{base}^C\right)$, and $(\bar{b}_l \pi_h^H + b_l \pi_l^H, \bar{b}_h \pi_h^H + b_h \pi_l^H)$ respectively. By condition (17), (H, H) Pareto-dominates (I, C) and (H, C) thus we can revise the condition for (I, C) and (H, C) to (20).

A.12 Proof of Proposition 5

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

(a) (I, I). When (14) holds, both (I, I) and(H, H) exist. By Corollary 1, a high-quality provider is better off under (H, H) than under (I, I) whereas a low-quality provider is worse off. The uncompromised equilibrium requires $\mu(h|H) \geq \beta$, which contradict $\mu(h|H) \leq \mu_1$ because when $\mu(h|H) \geq \beta$, $\hat{\hat{b}}_h^H > \bar{b}_h$, and thus $\bar{b} \pi_{base}^I < \hat{\hat{b}}_h^H \pi_h^H + b_h \pi_l^H$, contradicting condition (15). So (I, I) is compromised by (H, H).

(b) (H, H). When (14) holds, both (I, I) and (H, H) exist. Similar to (a), the uncompromised equilibrium requires that $\mu(h|I) \leq \beta$, which clearly does not contract $\mu(h|I) \leq \mu_2$ So (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) only coexists with (H, C) under condition (20). But both provider 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.

A.13 Proof of Corollary 3

The impact on provider revenues. The result that the high-quality provider is better off under (H, H) than under (I, I) and the low-quality provider is worse off follows from Corollary 1. To see that both provider types are better off when (I, C) is replaced by (H, H), notice that this occurs when (by Propositions 4 and 5)

$$\bar{b}_h \pi_h^H + b_h \pi_l^H \geq b_h \pi_{base}^C \geq \bar{b}_h \pi_{base}^I \tag{40}$$

The equilibrium revenues for low- and high-quality providers are $\left(\bar{b}_l \pi_h^H + b_l \pi_l^H,\ \bar{b}_h \pi_h^H + b_h \pi_l^H\right)$ under (H, H) and $\left(b_l \pi_{base}^I, b_h \pi_{base}^C\right)$ under (I, C). Clearly the (H, H) revenues are higher than the (I, C) ones, noting that $b_l < \bar{b}_l$, (40), and (11).

The impact on advertiser payoffs. In the case where (I, C) is replaced by (H, C), all advertisers are indifferent because the PPI auction and the hybrid auction allocate the impression in the same way and in both cases, advertisers perfectly infer the provider's quality. In the case where (I, I) is replaced by (H, H) , advertisers are indifferent too. 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 the case where (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 provider is high quality (so that PPC would be used under and (I, C) equilibrium), a high-quality advertiser can be better off or worse off under PPC depending on whether the prior about the advertiser's quality is high or not.

Impact on allocative efficiency. We know from earlier proofs that PPI and hybrid auctions are more efficient than PPC auctions. The results on the overall efficiency follow immediately.