

On the Optimality of Clickthrough Fees in Online Markets

Michael R. Baye*

Cathy Gao

Indiana University

University of Miami-Ohio

John Morgan

University of California, Berkeley

May 2011

Abstract

We study optimal fee setting decisions by a monopoly online platform connecting advertisers with potential buyers in two environments: (1) a simple model that captures stylized features of advertising on search engines, social networks, and ad-supported email, and (2) a richer model that is more relevant for “directed” search at price comparison sites. While the platform can choose to charge both for impressions as well as clicks, we show that the platform maximizes profits by using clickthrough fees exclusively. Our model offers a rationale for the evolving practice of relying purely on clickthrough fees for revenues in many online advertising markets.

Keywords: Two-sided markets, clickthrough fees (CPC), impression fees (CPM), online advertising

*Correspondent: Michael R. Baye, Department of Business Economics and Public Policy, Kelley School of Business, Indiana University, Bloomington IN 47405. Email: mbye@indiana.edu. We are grateful to Graeme Hunter, David Myatt, Daniela Simova, and Dana Sisak for comments on earlier drafts.

1 Introduction

The pricing of online advertising has undergone a sea change since its early days. The dominant old media model of paying per impression, the so-called CPM (cost per impression) model, has given way to a pricing model where most payments are made contingent on the viewer taking some action, typically clicking on the ad.¹ This is the so-called CPC (cost per click) model of pricing. In this paper, we investigate why platforms such as price comparison sites have shifted from CPM to CPC.

This change is usually seen as a response to pressure from advertisers worried about the performance of the new online media. Under the CPC model, an advertiser only pays when an ad is effective—based on the user action of clicking. Thus, the risk of performance is shifted from the advertiser to the site displaying the ads. While there is little doubt that advertisers were (and still are) skeptical about the returns to online advertising, this paper points out an additional motivation for the shift: The platform hosting advertising content—what we have previously dubbed (Baye and Morgan, 2001) the “information gatekeeper”—also stands to benefit from this pricing arrangement. Indeed, our main result is that, when CPC and CPM are both available and can be used jointly as pricing instruments, the platform maximizes its profits by using CPC exclusively.

One might speculate that the CPC and CPM instruments are analogous to two-part tariffs for a standard monopoly: Price at marginal cost and use a fixed fee to extract surplus from advertisers. Under this logic, one might erroneously conclude that the gatekeeper should optimally set the clickthrough fee at zero (its marginal cost) and use an impression fee to extract surplus from potential advertisers.² Where this logic goes wrong is that it ignores competition among advertisers. In the standard setting, one buyer’s decision to accept the monopolist’s offer does not impact the valuations of other buyers. In the present setting, one firm’s decision to have its ad displayed does impact the amount other firms are willing to pay, since it reduces the likelihood that any given ad will be viewed as “best” and hence clicked by consumers. An optimizing platform takes this “competition for clicks”

¹One may wonder why “cost per impression” is abbreviated as CPM rather than CPI. The “M” in CPM is a legacy of the old media world where ads were priced per thousand impressions, with M being the Roman numeral for 1,000.

²If the total number of impressions is known (or reasonably estimated) in advance, CPM is equivalent to the usual fixed fee.

into account in its pricing, and this makes the analysis of the optimal fee structure more complex than in the standard textbook case.

The tradeoffs between CPC and CPM can be clearly seen in the simple model of optimal platform pricing presented in Section 2, where ads exogenously differ in their attractiveness and relevance to consumers. Since consumers click on the most attractive and relevant ad, a less attractive advertiser might still hope to obtain clicks via relevance. Advertisers enter until the least attractive (marginal) advertiser is just indifferent between utilizing this channel or earning zero from its outside option.

Suppose that impression and clickthrough fees are adjusted to maintain the same number of advertisers. Then a \$1 increase in clickthrough fees requires a reduction in impressions fees in proportion to the chance that no other firm’s ad is displayed, i.e. an order statistic. The platform benefits from such an adjustment in proportion to the chance that at least one firm’s ad is displayed (i.e., deemed relevant by its algorithm) since otherwise the platform gets no clickthrough fees. It is harmed by the reduction in impression fees to *all* firms. The heart of the argument is establishing that this decrease is proportional to the chance that *at least* one firm’s ad is displayed. Since the former probability always exceeds the latter, a greater reliance on clickthrough fees raises platform profits and hence the exclusive use of clickthrough fees is optimal (Proposition 1). Put differently, while clickthrough fees operate on an order statistic—only the most attractive relevant firm pays—impression fees affect all firms and hence are a blunter instrument for extracting surplus.

While the simple model is a clean illustration of this intuition, one may worry that the result may be altered or reversed when the attractiveness of ads or the value of the outside option is endogenized. To investigate this possibility, Section 3 considers a richer model in the spirit of Baye and Morgan (2001), the first paper to examine the fee-setting behavior of a “gatekeeper” serving consumers and firms in a two-sided online market. This model closely corresponds to a setting where the platform is a price comparison site.³ The price charged by a firm determines the attractiveness of its ad, while the firm’s outside option (i.e. the payoff when not using the site) depends on the intensity of advertising at the site. Thus, both the attractiveness of ads and the value of eschewing the platform’s channel are

³Our model of price comparison sites derives from a much older literature of seller competition across channels (see, e.g. Varian (1980), Rosenthal (1980), Narasimhan (1988)). With the exception of Baye and Morgan (2001), this literature treats the cost of using various sales channels, including the platform, as exogenous.

endogenous.

Despite these modeling differences, CPC remains the superior instrument for capturing surplus in the market (Proposition 3). Once again, the key driver is the fact that shoppers click on the *best* of the listed advertisements. Consequently, a firm realizes that it is obliged to pay a clickthrough fee only if, *ex post*, its advertisement best matches the preferences of shoppers visiting the gatekeeper’s site. In contrast, an advertiser pays the impression fee regardless of whether it generates any clicks, and *before* any information is revealed about whether its advertisement is “best.” Finally, the gatekeeper’s expected profits from the impression fee depend on the average number of ads induced by the fee, while its profits from clickthrough fees depend only on the probability that at least one firm advertises (an order statistic). The contingent nature of clickthrough fees, coupled with this order statistic effect, makes them a superior instrument for extracting surplus from potential advertisers. These results obtain in a model where firms are *ex ante* symmetric and there is no moral hazard. Obviously, if firms were vulnerable to moral hazard on the part of the gatekeeper, the superiority of clickthrough fees over impression fees would be even greater.

The institutional structure of the richer model also allows us to study other features of online markets. Conversion rates, the chance that clicks are converted into sales, play a key role. We show that, while the exclusive use of clickthrough fees is optimal for any positive conversion rate, higher conversion rates lead to higher clickthrough fees and higher profits for the gatekeeper. Firms also benefit from higher conversion rates. An important insight to emerge from this analysis is that the platform only partly captures the gains from its investment in improving conversions. Thus, platforms will tend to underinvest in improvements. This perhaps helps to explain why conversions remain stubbornly low (about 5% at most) on these sites.

While both models share the same implication about the optimality of CPC pricing, Section 4 shows there are important differences. Both models assume that there are (potentially small) transactions costs to firms wishing to advertise on the platform. In the simple model, platform profits increase continuously as transactions costs fall to zero. In the richer model, the outside option is endogenous and this leads to a discontinuity in payoffs—platform profits exhibit an upward jump when transactions costs are eliminated. In a similar vein, exclusive contracts (i.e. contracts that offer a single advertiser the exclusive right to advertise on the platform) are of no benefit to the platform in the simple

model but are helpful in the richer model.

Our paper is related to the literature on optimal fee structures in two-sided markets when there is a monopoly platform.⁴ An important difference is that payoffs in these models are increasing in the number of agents on the opposite side of the platform and independent of the number of agents on the same side of the platform. In other words, competition among users of the platform is effectively absent. In the settings we study, it is natural to include competition and indeed it plays a key role in optimal platform pricing. Our paper is somewhat related to the literature on position auctions (see, e.g. Edelman, Ostrowsky, and Schwarz (2007), Varian (2007), Athey and Ellison (forthcoming), as well as Edelman and Schwarz (2006)). While these models also study clickthrough fees, their levels are determined by the auction form selected by the platform. Seller decisions in these models consist purely of how much to bid in the position auction.

Like many of these models, our model assumes that there is a single dominant platform in the market. While this potentially limits the applicability of these models, casual empiricism suggests that such market structures are, in fact, common. For instance, eBay is the dominant platform for online auctions in the US; Google enjoys overwhelming market share in search; Facebook is the dominant social networking platform. A burgeoning empirical literature (see, e.g. Liebowitz and Margolis (1994), Tellis, et al. (2009), and Brown and Morgan (2009)) documents that single (dominant) platforms are the rule in many markets. Recent experimental evidence (see Hossain and Morgan (2009) and Hossain, Minor, and Morgan (forthcoming)) paints a similar picture: Unless platforms are strongly horizontally differentiated, the most likely market structure to emerge is a monopoly platform.

2 A Simple Model

The intuition behind our central results can be illustrated in a simple model that captures stylized features of advertising on platforms such as search engines, social networks, or advertising-supported e-mail. Suppose there is a unit mass of consumers that exclusively use a given platform to make search queries, to connect with friends, or for some other activity not directly related to the viewing of ads. This activity provides the platform with some information about consumer preferences. Given what it learns from consumer

⁴See, e.g. Weyl (2009) as well as sections of Rochet and Tirole (2003), Caillaud and Jullien (2003) and Armstrong (2006).

activity on its site, the platform attempts to match consumers with “relevant” ads from a pool of $n \in \{1, 2, \dots, N\}$ firms that have opted to advertise through its channel. A given firm’s ad is identified as relevant (and hence “pushed” to consumers) with probability $\alpha \in (0, 1)$. This allows for the possibility that the platform may not be able to perfectly match ads to consumers (some relevant ads may not be displayed) owing to imperfections in its proprietary algorithm.

Figure 1 provides a screenshot that is representative of the types of ads considered in this section. Owing to heterogeneities in the creative abilities and long-run business decisions of firms, the content and appeal of even *relevant* ads differs across firms; some firms might advertise “free overnight shipping” or “20 percent discounts” while others emphasize product breadth or inventory depth. Let a summarize the overall quality of a given firm’s (multi-dimensional) ad, and note that the identity of the firm displaying a given ad (e.g., Amazon) may also impact the overall quality. Order the firms such that their ad qualities are $a_1 > a_2 > \dots > a_N$, which are common knowledge. A consumer who uses the platform is shown relevant ads, observes the quality of each, and clicks on the most attractive one (i.e., the ad with the highest value of a). The incremental value of a click, v , is assumed to be identical for each firm.⁵ Let $\varepsilon < \alpha v$ be a firm’s cost of delivering its ad copy to the platform, i.e., ε represents a transactions cost. This guarantees that the platform can profitably induce at least one firm to utilize its channel.

The platform announces an impression fee (T) and a clickthrough fee (c), and then firms make a decision to enter this advertising channel. Notice that, conditional on an ad being relevant (and hence being displayed), consumers click on the ad with the greatest overall quality (the highest value of a). Thus, firms with higher quality ads are more appealing to consumers and are expected to generate more clicks. Consequently, a firm’s incremental profit from entering the channel is increasing in a ; entry continues until the firm with the marginal a , denoted a_n , earns zero incremental profits by advertising in the channel:⁶

$$\alpha(v - c)(1 - \alpha)^{n-1} - \varepsilon - \alpha T = 0 \tag{1}$$

⁵This does not rule out the possibility that firms with larger a_i ’s have made more substantial long-term investments and thus may have higher overall costs. These sunk investments do not affect the incremental value of a click on this channel. The main conclusions in this section extend to the case of heterogeneous v ’s, provided that firms with higher-quality ads have higher v_i ’s.

⁶As will become clear in the sequel, the platform has an incentive to set fees such that the marginal firm earns exactly zero profits; the integer constraint does not bind.

Intuitively, the marginal firm only earns $v - c$ when its ad is relevant (a probability α event) *and* no other firms' ads are relevant (which happens with probability $(1 - \alpha)^{n-1}$). The firm pays ε to create its ad copy, and pays T when its ad is displayed (which occurs with probability α).

While the platform incurs costs in setting up and supporting its website as well as serving images to consumers, these costs typically do not vary with the number of advertisers on the site nor the fee structure to advertisers. As a consequence, we abstract away from this aspect of the platform's profit function. Thus, when n firms choose to use the channel, the platform's profits are

$$U = \alpha n T + (1 - (1 - \alpha)^n) c \quad (2)$$

Here, the first term reflects the platform's expected impression (or CPM) revenues while the second term is clickthrough (or CPC) revenues. Notice that revenues from impressions depend on the expected number of relevant ads ($n\alpha$) while clickthrough revenues depend on the probability that at least one ad is relevant, $(1 - (1 - \alpha)^n)$.

2.1 The Optimality of Clickthrough Fees

Our strategy is to first identify the optimal fee structure for a given number of advertisers using the platform, and then to determine the platform's optimal n . Proposition 1 identifies the optimal c and T for any given n , and shows that a platform maximizes profits by using clickthrough fees exclusively. Proposition 2 then identifies the number of advertisers the platform induces (by altering its fees) to maximize its profits.

The following lemma is central to the results in this and subsequent sections of our paper

Lemma 1 *Suppose $n > 1$. Then the function $\psi(x; n) \equiv 1 - x^n - nx^{n-1}(1 - x)$ is strictly positive and strictly decreasing in x for all $x \in [0, 1)$.*

Proof. First, notice that $\psi(0; n) = 1$ and $\psi(1; n) = 0$. Next, notice that

$$\frac{\partial \psi}{\partial x} = -n(n-1)(1-x)x^{n-2} < 0$$

for all $x \in (0, 1)$. ■

This lemma is intuitive if one views $1 - x$ as the probability that a firm's ad is relevant. In this case, ψ is simply the difference between the probability that at least one firm's ad is relevant and the probability that *exactly* one firm's ad is relevant. Obviously, the former is larger than the latter.

Proposition 1 *Suppose n firms advertise through the platform. Then:*

(a) *If $n = 1$, any combination of clickthrough and impression fees satisfying equation (1) maximizes the platform's profits. Hence, the exclusive use of clickthrough fees is weakly optimal.*

(b) *If $n \geq 2$, the platform maximizes profits by exclusively using clickthrough fees.*

(c) *The profit maximizing clickthrough fee is*

$$c(n) = \begin{cases} v - \frac{\varepsilon}{\alpha} & \text{if } n = 1 \\ v \left(1 - \frac{\varepsilon}{\alpha v} \frac{1}{(1-\alpha)^{n-1}}\right) & \text{if } n \geq 2 \end{cases} \quad (3)$$

Proof. First, notice that, when n firms enter, equation (1) must hold with equality. If the n th firm earned strictly positive profits (owing to the integer constraint), the platform could increase either c or T without affecting n and increase its profits.

When $n = 1$, the platform's profits are $U = \alpha T + \alpha c$ and, substituting for T using equation (1), we have

$$U = \alpha v - \varepsilon$$

which is independent of c . In this case the platform is indifferent between using CPC or CPM and there is a continuum of (c, T) 's that satisfy equation (1). Thus, the platform can do no better than to set $T = 0$ and charge a clickthrough fee of $c = v - \varepsilon/\alpha$.

When $n \geq 2$, the platform chooses c and T to maximize equation (2) subject to equation (1). Solving for T using equation (1) and substituting into equation (2) yields

$$U = \alpha n \left((v - c)(1 - \alpha)^{n-1} - \frac{\varepsilon}{\alpha} \right) + (1 - (1 - \alpha)^n) c$$

Hence

$$\begin{aligned} \frac{\partial U}{\partial c} &= -\alpha n (1 - \alpha)^{n-1} + (1 - (1 - \alpha)^n) \\ &= \psi(1 - \alpha; n) > 0 \end{aligned}$$

where the inequality follows from Lemma 1. Thus, when $n \geq 2$, the platform earns strictly more by exclusively using clickthrough fees. Solving equation (1) in terms of c yields the expression for $c(n)$ in the proposition. ■

Proposition 1 indicates that clickthrough fees are weakly optimal, and in fact are strictly optimal when two or more firms utilize the platform’s channel. The proof reveals that the platform’s marginal profit from impression fees is linear in the probability that a firm’s ad is deemed to be relevant. The marginal profit from clickthrough fees depends on an order statistic—the chance that at least one firm’s ad is relevant. When $n = 1$, these two probabilities coincide, and hence, we obtain a neutrality result reminiscent of Proposition 3 in Armstrong (2006). In contrast, when $n \geq 2$, consumers click the most attractive of the competing firms’ ads, which leads to a difference between the two probabilities. Expressed differently, given the linearity of the platform’s profit function in c and T , the marginal profit from T is αn while the marginal profit from c is $(1 - (1 - \alpha)^n)$. Comparing the T -versus- c slopes in equations (2) and (1), Lemma 1 implies that the exclusive use of clickthrough fees is optimal when $n \geq 2$. Notice that this order-statistic effect stemming from competition among advertisers is absent in Armstrong (2006) and other papers on two-sided markets. These papers assume that payoffs only depend on the number of users on the *other* side of the market—independent of the degree of competition on the *own* side of the market.

One might wonder who wins and who loses from a shift from CPM to CPC (for a given $n \geq 2$). Consumers observe the same set of ads regardless of whether the platform uses clickthrough or impression fees, so their welfare is unaffected. Likewise, the marginal firm earns zero profits regardless, so it too is unaffected. However, inframarginal firms—those displaying more attractive ads—are strictly worse off. To see this, note that the profits of the i th most attractive firm ($i < n$) are

$$\pi_i = \alpha(v - c)(1 - \alpha)^{i-1} - \varepsilon - \alpha T$$

When the platform uses clickthrough fees exclusively, firm i earns

$$\begin{aligned} \pi_i(c) &= \alpha(v - c(n))(1 - \alpha)^{i-1} - \varepsilon \\ &= \varepsilon \left(\frac{(1 - \alpha)^{i-1} - (1 - \alpha)^{n-1}}{(1 - \alpha)^{n-1}} \right) \end{aligned}$$

In contrast, when the platform exclusively uses impression fees, equation (1) implies $T(n) = v(1 - \alpha)^{n-1} - \varepsilon/\alpha$ and firm i earns

$$\begin{aligned} \pi_i(T) &= \alpha v(1 - \alpha)^{i-1} - \varepsilon - \alpha T(n) \\ &= \alpha v \left((1 - \alpha)^{i-1} - (1 - \alpha)^{n-1} \right) \end{aligned}$$

Since $(T, c) \geq 0$, equation (1) implies $\alpha(1 - \alpha)^{n-1}v > \varepsilon$. Hence, for all $i < n$,

$$\begin{aligned} \pi_i(c) - \pi_i(T) &= \left((1 - \alpha)^{i-1} - (1 - \alpha)^{n-1} \right) \left(\frac{\varepsilon}{(1 - \alpha)^{n-1}} - \alpha v \right) \\ &< 0 \end{aligned}$$

Thus, the benefits to the platform of exclusively using clickthrough fees comes entirely out of the hides of inframarginal firms.⁷

Proposition 1 also reveals that, for a given number of advertisers, the optimal clickthrough fee is v less a discount that depends on two factors. The first factor, $\frac{\varepsilon}{\alpha v}$, accounts for the fact that the marginal firm bears transactions costs of creating ad copy to advertise on the platform. The platform thus finds it optimal to discount c by the fraction of total available surplus (αv) eaten up by these costs. The second factor, $1/(1 - \alpha)^{n-1}$, discounts for competition among advertisers, and represents a reduction in clickthrough fees to induce the marginal firm to use the platform. It is intuitive that the larger is n , the greater the discount in c required to induce a given level of entry into the channel, other things equal.

Notice that, regardless of how many advertisers the platform ultimately finds it optimal to attract, Proposition 1 indicates that it can optimally set $T = 0$ and increase c until the marginal firm is indifferent between entering or not. By further adjusting c , the platform can induce entry or exit to fully optimize its profits.

To determine the number of advertisers that maximizes the platform's profits, substitute the expression for $c(n)$ in equation (3) into the expression for U in equation (2) to obtain

$$U(n) = \begin{cases} \alpha v - \varepsilon & \text{for } n = 1 \\ (1 - (1 - \alpha)^n) \left(1 - \frac{\varepsilon}{\alpha v} \frac{1}{(1 - \alpha)^{n-1}} \right) v & \text{for } n \geq 2 \end{cases} \quad (4)$$

This expression gives the platform's maximal profits when n firms enter and the clickthrough fee is set optimally. In the Appendix we show that $\Delta(n + 1) \equiv U(n + 1) - U(n)$ is strictly decreasing, so that the incremental contribution to the platform of inducing another firm to enter declines as more firms use its channel. Since $\Delta(1) > 0$, the platform optimally induces at least one firm to enter. Since $\lim_{n \rightarrow \infty} \Delta(n + 1) < 0$, the optimal number of firms (from the platform's perspective) is finite. Finally, since there are only N potential entrants, the optimal number of advertisers cannot exceed N . Thus, the optimal number of advertisers, n^* , satisfies $1 \leq n^* \leq N$ and the optimal fees generating this level of entry are

⁷Since any (T, c) satisfying equation (1) is a convex combination of these two extreme contracts, the result holds locally as well—any shift from CPM to CPC pricing for a given n harms inframarginal firms.

as in Proposition 1 with $n = n^*$. The following proposition provides a closed form solution for n^* and summarizes the results of this section.

Proposition 2 *A profit maximizing platform induces n^* firms to advertise through its channel, where*

$$n^* = \begin{cases} 1 & \text{if } \frac{\varepsilon}{\alpha v} > 1 - \alpha \\ \min \left\{ \left\lceil \frac{1}{2} \frac{\log \left[\frac{\varepsilon}{\alpha v} \right]}{\log [1 - \alpha]} \right\rceil, N \right\} & \text{if } \frac{\varepsilon}{\alpha v} \leq 1 - \alpha \end{cases}$$

and $\lceil \cdot \rceil$ denotes the integer ceiling function. The platform's maximal profits, $U^* = U(n^*)$, are attained by setting $T = 0$ and charging a clickthrough fee of $c(n^*) > 0$.

Three aspects of this proposition are noteworthy. First, the optimal number of advertisers depends on $\frac{\varepsilon}{\alpha v}$, which represents the fraction of available surplus consumed by transactions costs. Proposition 2 implies that when this fraction is large, the platform finds it optimal to limit the number of competing advertisers by charging a relatively high clickthrough fee. Indeed, when this fraction exceeds the probability that a firm's ad is deemed irrelevant by the search engine ($1 - \alpha$), the platform finds it optimal to limit ads to a single firm. As this fraction declines, the profit maximizing clickthrough fee decreases and the number of firms choosing to advertise through the platform rises.

Second, as $\varepsilon \rightarrow 0$, the profit-maximizing number of advertisers tends to N , the universe of potential advertisers. This is intuitive: As transactions costs decline, more and more firms find it profitable to use the channel. Additionally, a reduction in transaction costs permits the platform to charge a higher clickthrough fee. Indeed, when N is unbounded, the optimal clickthrough fee converges to v as ε tends to zero (and hence, n^* tends to infinity). Intuitively, as the number of firms increases, the platform becomes increasingly certain that *at least one* firm's ad will be relevant. Since consumers only click on the best of the relevant ads, and a firm only pays the platform if its ad is clicked, the platform can extract the full value, v , from the firm that is clicked.

Finally, it is useful to discuss how changes in α —which in the present model measures the platform's "efficiency" in matching relevant ads to consumers—impact the platform and participants. Notice in Proposition 2 that, generically, $\log \left[\frac{\varepsilon}{\alpha v} \right] / \log [1 - \alpha]$ is not integer-valued and $\frac{\varepsilon}{\alpha v} \neq 1 - \alpha$. Hence, small changes in α generically have no effect on n^* , the equilibrium number of advertisers. Equation (2) implies that, for a given c , an increase in α increases platform profits. Since the platform is always free to leave c unchanged in

response to an increase in α , it follows that $\partial U^*/\partial\alpha > 0$. Consequently, small improvements in match efficiency generically benefit the platform.

3 A Richer Model

While the simple model demonstrates the optimality of using clickthrough fees exclusively, it suffers from several limitations. First, the outside option for firms is assumed to be fixed and independent of the success of the platform in attracting advertisers, shoppers, and clicks. While this may make sense in settings where consumers are *not* engaged in directed search, i.e. where the display of an ad is incidental to the intentions of consumers visiting the platform, in many settings consumer search might be directed toward a purchase decision. For instance, the search query “quality hotels in Bloomington” is presumably made with the intent of booking a hotel in that charming midwestern city and so advertisements about the hospitality possibilities are welcomed by consumers. The absence of such ads may well lead the consumer to utilize other channels, such as visiting the websites of hotel chains directly, to secure accommodation.

Second, the decision to display an ad is not under the control of firms—rather it is determined exogenously, presumably through the platform’s algorithms designed to match advertisers to (implicit) consumer needs. While this is reasonable when consumers are engaged in non-directed search, it seems problematic in circumstances where a consumer’s purchase intentions may be readily discerned. Third, both the attractiveness of advertisements and the value of a click are assumed to be outside of the control of firms. While this makes sense when these attributes are based on long-lived characteristics such as brand or reputation, it makes less sense when they depend on *flexible* characteristics of the ad, such as the price charged. In short, the previous analysis assumes that (1) the outside option for firms is exogenous; (2) a is exogenous; and (3) α and v are exogenous.

In this section, we offer a model that relaxes all of these assumptions and demonstrates that the exclusive use of clickthrough fees remains optimal in a directed search environment. Perhaps the purest form of directed search occurs at a price comparison site—a “virtual market” that consumers visit primarily to obtain price information. Such sites provide consumers a *complete* list of the prices charged by all firms advertising a given product on the platform. The platform may decide the order to present offers, but does not exclude any

advertised price. Regardless of the order, a single mouse click typically enables a shopper to sort prices from lowest to highest to easily identify the best price. Figure 2 offers a screenshot of the consumer experience for a typical price comparison site.

Over the past decade, empirical research has documented three main stylized facts regarding prices on such sites: (1) Seemingly identical advertised products exhibit considerable price dispersion; (2) firms' positions in the distribution of prices, as well as the identity of the firm offering the lowest price, change frequently; and (3) the number of firms advertising prices for a given product is highly variable, changing from day-to-day.⁸ Thus, the price dispersion in Figure 2 is the rule rather than the exception for products advertised at comparison sites. The model presented below comes from a class of models frequently used to rationalize these facts. Two key features of this class of models are: (1) advertised prices are the main driver of clicks—firms compete for price sensitive consumers on the site; and (2) the comparison site is not the only channel used to make purchases—consumers can (and do) visit firms' sites directly to make purchases.

The details of the richer model are as follows. There are $N > 1$ firms that sell similar products with a constant marginal cost of $m \geq 0$. Firms reach consumers through two sales channels: (1) Their individual websites, and (2) a price comparison site operated by a monopoly platform. To make clear the distinction between the simple model and the richer model presented in this section, we refer to the platform as the comparison site or gatekeeper.⁹ Consumers have unit demand and a maximal willingness to pay of $r > m$. There are two types of consumers: Loyals and Shoppers. Loyals view the products and/or service quality of the firms as highly differentiated and exclusively frequent their preferred firm's website. Shoppers, on the other hand, view firms as essentially identical; they visit the comparison site and click on the firm offering the lowest price. If an acceptable price is not found, shoppers visit the website of a randomly selected firm.¹⁰ Each firm attracts $L > 0$ loyals, and there are $S > 0$ shoppers.

⁸See, for instance, Baye, Morgan, and Scholten (2004a,b); Ellison and Ellison (2009); Brynjolfsson and Smith (2000); and Chevalier and Goolsbee (2003). Baye, Morgan, and Scholten (2006) provide an overview of these and many other papers that document these properties.

⁹The model presented here readily extends to environments with multiple comparison sites, provided that firms and shoppers single-home (for example, are loyal to a particular price comparison site).

¹⁰See Proposition 1 in Baye and Morgan for the conditions under which this is an optimal search strategy on the part of shoppers.

As a practical matter, clicks and visits to websites do not always translate into sales.¹¹ Consumer “shopping carts” are abandoned regularly. To account for this, we assume that visits and clicks translate into sales only some of the time. Specifically, a loyal customer purchases the product with probability $\lambda \in (0, 1]$, while a click from a shopper is converted into a sale with probability $\gamma \in (0, 1]$. In industry parlance, γ and λ are *conversion rates*.

As with the simple model, the gatekeeper operating the comparison site has two instruments to earn revenues: the impression fee ($T \geq 0$) and clickthrough fee ($c \geq 0$) charged to firms utilizing its channel.¹² In return for these fees, a firm’s price is displayed to shoppers visiting the gatekeeper’s site. The impression fee is paid regardless of whether a consumer clicks or ultimately purchases an item, and, since a firm’s ad is served to all shoppers, is equivalent to CPM pricing. The clickthrough fee is paid only if the consumer clicks and regardless of whether a purchase is made; it represents the CPC component of utilizing the channel.

Each firm independently determines its price and whether to advertise it at the comparison site.¹³ Let p_i denote firm i ’s price and $\alpha_i \in [0, 1]$ denote the (possibly degenerate) probability that firm i chooses to advertise on the comparison site. In addition to the fees paid to the gatekeeper, a firm that advertises on the comparison site also incurs a transactions cost of $\varepsilon > 0$. This represents the frictions associated with inputting the relevant price and product information into the gatekeeper’s system, verifying its accuracy, and so on. Obviously, if ε were too large, firms would never advertise on the comparison site even if doing so were free. We thus assume $\varepsilon < S((r - m)\gamma\frac{N-1}{N})$, which guarantees activity on the gatekeeper’s site. A firm that does not utilize the comparison site avoids all of these costs, but at the potential cost of failing to reach shoppers.

¹¹Conversion rates in online markets for most products average about 3%. See, for instance, “Comparison Search Engines Tested” at www.marketingexperiments.com.

¹²In principle, the gatekeeper might try to earn revenues from consumers by turning its site into a “paywall” and charging consumers for access. In practice, websites rarely use this instrument—presumably because it is in their interest to attract as many “eyeballs” as possible and imposing such fees would diminish the reach of their channel; see Proposition 5 in Baye-Morgan (2001). In any event, formally modeling this in the present setting would require that we account for content provided by the gatekeeper other than advertising, which is beyond the scope of the model. Hence, we assume this revenue stream is unavailable.

¹³While our formulation excludes price discrimination between loyal and shoppers, which is consistent with practice, our results would be qualitatively unchanged were one to admit price discrimination. See Baye and Morgan (2002) and Nahm (2003) for characterizations of firm behavior under price discrimination.

The timing of decisions is as follows. First, the gatekeeper announces fees (T, c) in an attempt to maximize its profits. Given the fee structure associated with this channel, firms simultaneously and independently choose (p_i, α_i) ; that is, they determine their prices and whether to advertise them at the comparison site. Finally, consumers browse the internet and make purchase decisions.

We solve the model by first characterizing symmetric Nash equilibrium firm behavior for a given $(T, c) \geq 0$, and then allow the comparison site to optimally determine the profit-maximizing level of its fees. Our focus on symmetric equilibria is motivated by several considerations. First, since the attractiveness of advertised prices is endogenous and firms are otherwise identical, it seems arbitrary to assume that firms somehow coordinate on entry based on their index number or some other criterion. Second, it makes the model consistent with the extant literature where symmetric mixed-strategy equilibria rationalize price dispersion and variation in the number of firms advertising on price comparison sites. Finally, the symmetry assumption permits us to study how advertising fees impact firms' demands for advertising in a parsimonious way and to examine the extent to which the optimality of clickthrough fees in the simple model stems from asymmetries.

3.1 Firm Behavior

Equilibrium firm behavior is similar to that in Baye and Morgan (2001) and has the following features (see the Appendix for details and a formal proof). First, consistent with Figure 2 and the empirical literature, advertised prices are dispersed according to an atomless distribution $F(p)$ on $[p_0, r]$, where $m < p_0 < r$. Since shoppers visiting the comparison site find lower advertised prices more attractive than higher ones, the quality of each firm's ad is endogenous. Intuitively, firms randomize their advertised prices to make them unpredictable, thereby preventing rival firms from profiting by systematically offering more attractive ads.

Second, each firm advertises a price on the comparison site with probability $\alpha \in (0, 1)$, so there is variation in the observed number of ads. A firm that does not advertise simply charges a price of r on its own website—the reservation price of its loyal customers.

Third, a firm's outside option depends on the advertising decisions of other firms. In particular, a firm that advertises a price of p on the comparison site earns expected profits of

$$\pi(p) = (p - m)\lambda L + ((p - m)\gamma - c)(1 - \alpha F(p))^{N-1} S - T - \varepsilon \quad (5)$$

while it earns

$$\pi_0 = (r - m)\lambda L + (r - m)\gamma(1 - \alpha)^{N-1} \frac{S}{N} \quad (6)$$

if it does not advertise. Since, in equilibrium, F is determined such that each firm is indifferent between advertising any price $p \in [p_0, r]$ on the comparison site and simply charging r on its own website, $\pi(p) = \pi_0$ for all $p \in [p_0, r]$. In particular, setting $\pi(r) = \pi_0$ determines the equilibrium advertising propensity, α . This is the analog to the entry condition (equation 1) of the simple model:

$$(r - m)\lambda L + ((r - m)\gamma - c)S(1 - \alpha)^{N-1} - T - \varepsilon = (r - m)\left(\lambda L + \gamma \frac{S}{N}(1 - \alpha)^{N-1}\right) \quad (7)$$

The left-hand-side is a firm's expected profits when it displays the worst possible ad on the comparison site (a price of r). The right-hand-side is its outside option—its expected profits when it exclusively sells through its own website. Notice that the value of the firm's outside option is decreasing in α , which is endogenous.

We may solve equation (7) for α to obtain the equilibrium advertising propensities of firms as a function of the comparison site's fee-setting decisions:

$$\alpha = 1 - \left(\frac{\varepsilon + T}{S((r - m)\gamma \frac{N-1}{N} - c)}\right)^{\frac{1}{N-1}} \in (0, 1) \quad (8)$$

for $T + \varepsilon < S((r - m)\gamma \frac{N-1}{N} - c)$. Notice that α is decreasing in T and c ; indeed, if the comparison site set fees so high that $T + \varepsilon \geq S((r - m)\gamma \frac{N-1}{N} - c)$, firms would completely shun the comparison site.

3.2 The Optimality of Clickthrough Fees

Given that each firm advertises its price on the comparison site with probability α , the price comparison site's profits are

$$U = \alpha NT + (1 - (1 - \alpha)^N)Sc \quad (9)$$

Similar to the analogous expression in equation (2), the first term reflects the comparison site's expected CPM revenues while the second term is clickthrough (or CPC) revenues.

Notice that impression revenues depend on the expected number of ads (αN) while click-through revenues depend on the probability that at least one firm advertises ($1 - (1 - \alpha)^N$) at the site. Despite the similarity in objective functions, the endogeneity of firms' outside options and α complicates the determination of the comparison site's optimal fee structure.

We are now in a position to show that the price comparison site maximizes profits by exclusively utilizing a clickthrough fee, and to identify the optimal level of that fee. Since the platform can induce firms to employ given (equilibrium) advertising propensities through differing combinations of impression and clickthrough fees, we use a two-step procedure similar to what we did in the simple model. In the first step, we show that for a given α , the platform maximizes profits by exclusively using clickthrough fees. In the second step, we pin down the advertising propensities that maximize the platform's profits. The optimal clickthrough fee is then the one that induces these advertising propensities.

It is convenient to rewrite equation (8) to obtain the locus of (T, c) choices that induce a given advertising propensity:

$$T = (1 - \alpha)^{N-1} S \left((r - m) \gamma \frac{N - 1}{N} - c \right) - \varepsilon \quad (10)$$

Substituting equation (10) into equation (9) yields (after simplification)

$$U = S c \psi(1 - \alpha; N) + \alpha (1 - \alpha)^{N-1} S (r - m) \gamma (N - 1) - \alpha N \varepsilon$$

Notice that, similar to the simple model, $\partial U / \partial c = S \psi(1 - \alpha; N)$. By Lemma 1, we know this is strictly positive when firms are active at the price comparison site (i.e., $\alpha > 0$). Hence, we have shown:

Proposition 3 *A profit-maximizing price comparison site uses clickthrough fees exclusively.*

Intuitively, the comparison site's expected revenues from impression fees depend on the average number of ads and the *ex ante* amount firms are willing to pay to advertise, while expected clickthrough revenues depend on the probability that at least one firm advertises and the amount a firm is willing to pay when it "wins" all the clicks. Thus, clickthrough fees are superior because of the same order statistic effect that occurred in the simple model.

Unlike the simple model where the platform's gain from switching from CPM to CPC (for a given n) occurs solely at the expense of inframarginal firms, here the gains (for a given

α) arise purely at the expense of consumers, who end up paying higher prices as firms adjust their pricing behavior in response to the changing fee structure. To see this, recall that all firms are marginal in a symmetric equilibrium, so they are neither helped nor harmed from the shift to CPC for a given α . Thus, improved profits at the price comparison site come purely at the expense of consumers.

It remains to determine the advertising propensity induced by the comparison site. Proposition 3 implies that, for a given value of α , the optimal fee structure has $T = 0$. Thus, the clickthrough fee that induces a given α , denoted $c(\alpha)$, is the solution to equation (8) when $T = 0$:

$$c(\alpha) = (r - m) \gamma \frac{N - 1}{N} - \frac{\varepsilon}{(1 - \alpha)^{N-1} S} \quad (11)$$

Substituting this expression into the comparison site's payoff function when $T = 0$ yields the comparison site's profits solely as a function of the advertising propensity. It is helpful (and equivalent), to write this expression in terms of the complement, the chance a firm does not advertise, $z \equiv 1 - \alpha$:

$$U(z) = (1 - z^N) S (r - m) \gamma \frac{N - 1}{N} - \frac{1 - z^N}{z^{N-1}} \varepsilon \quad (12)$$

The fundamental tradeoff expressed in this equation is intuitive—increasing firms' advertising propensities (reducing z) benefits the price comparison site by reducing the chance that no firms list, but this comes at the cost of having to reduce the clickthrough fee to induce the higher advertising propensities.

The first-order condition for maximizing equation (12) is

$$\frac{\partial U}{\partial z} = -z^{N-1} S (r - m) \gamma (N - 1) + (N - 1) z^{-N} \varepsilon + \varepsilon = 0 \quad (13)$$

Notice that $\partial U / \partial z$ is positive for z arbitrarily close to zero, and negative when $z = 1$. Moreover, the second-order condition is satisfied, since

$$\frac{\partial^2 U}{\partial z^2} = -(N - 1)^2 z^{N-2} S (r - m) \gamma - N (N - 1) z^{-N-1} \varepsilon < 0 \quad (14)$$

Hence, there is a unique value of z that solves the first order condition. Using equation (11), we obtain:

Proposition 4 *A profit maximizing comparison site induces an advertising propensity of $1 - \zeta$, where $\zeta \in (0, 1)$ is the unique value of z solving equation (13). Its maximal profits,*

$U^* = S(1 - \zeta^N) c^*$, are attained by setting $T = 0$ and charging a clickthrough fee of

$$c^* = (r - m)\gamma \frac{N - 1}{N} - \frac{\varepsilon}{\zeta^{N-1}S}$$

3.3 Comparative Analysis

One of the advantages of tailoring the richer model to fit a particular online institution is that it permits us to examine how the platform and market participants change their behavior in response to changes in the parameters of the model. In this section, we show how changes in conversion rates and transactions costs impact the comparison site and firms.

Conversion Rates

Conversion rates—the chance that a click turns into a sale—are critically important to online advertisers as well as platforms. From the point of view of firms, it is hardly worthwhile to pay for clicks that don't lead to sales. Likewise, a platform that delivers clicks but no sales is unlikely to be able to sustain advertising revenues. Conversion rates depend on the types of consumers attracted to a site as well as the ease of use of the site. In the model, there are two conversion rates: those of shoppers (γ) and loyalists (λ).

We first examine how changes in the conversion rate of shoppers using the comparison site (γ) impact advertisers and the price comparison site. While γ is exogenous, comparative analysis of changes in conversion rates on the comparison site sheds light on its incentives to make long-term investments to improve conversion rates.

Proposition 5 *In equilibrium, an increase in the conversion rate of shoppers:*

- (a) *Raises the optimal clickthrough fee charged by the price comparison site and increases its profits; and*
- (b) *Leads firms to advertise more frequently and increases their profits.*

The intuition is as follows (a detailed proof is contained in the appendix). As the conversion rate of shoppers increases, the total surplus generated through the comparison site's channel increases. Other things equal, this induces firms to advertise more frequently. Thus, were the comparison site to respond by leaving the clickthrough fee unchanged, its profits would increase purely through more intensive advertising. Now, since the optimal clickthrough fee balances the marginal benefit of extracting surplus from advertisers with

the marginal cost of curtailing advertising propensities—and an increase in conversion rates increases these benefits and reduces these costs (through their effect on α)—it is optimal for the price comparison site to raise its clickthrough fee. This accounts for part (a) of the proposition. Part (b) of the proposition stems from the fact that the strategic reaction of the comparison site to raise the clickthrough fee in response to an increase in γ does not overwhelm the direct effect of raising the value of the channel through higher conversion rates.

Together, these results suggest that in an even richer model where the price comparison site could invest to improve the conversion rates of shoppers at its site, it would ignore this externality on firms, and therefore, have a tendency to underinvest on this aspect of its site. This perhaps explains why, despite years of lamenting the poor conversion rates in online markets, they remain stubbornly low.

How do higher conversion rates “off” the comparison site, those of loyals (λ), impact the platform and firms? The price comparison site’s optimal clickthrough fee and induced level of advertising, summarized in Proposition 1, are independent of λ , so ζ remains unchanged. Thus, the comparison site’s profits are unchanged. In contrast, the equilibrium profits of firms,

$$E\pi = (r - m) \left(\lambda L + \zeta^{N-1} \gamma \frac{S}{N} \right)$$

are increasing in the conversion rate of loyals, so firms benefit.

Proposition 6 *In equilibrium, an increase in the conversion rate of loyals:*

(a) *Has no effect on the comparison site’s optimal clickthrough fee or its equilibrium profits.*

(b) *Has no effect on firm advertising propensities but does raise their profits.*

This proposition suggests that firms also have incentives to make investments to improve clickthrough rates at their individual websites—that is, to make it easier for their “loyal” consumers to transform clicks into sales. Practical examples include loyalty programs, one click purchasing and the like. One might think that, since the relative value of the price comparison site channel has fallen, it would be induced to “compete” by lowering fees. The flaw in this reasoning is that the incremental value of the comparison site remains unchanged for a firm charging the highest price—and therefore equilibrium advertising propensities (and comparison site profits) remain unchanged. In contrast, a firm offering

the lowest price (p_0) now finds that price unattractive, and the equilibrium distribution of advertised prices increases as a result of the increase in λ . In short, increases in λ induce firms to advertise less attractive prices, such that the expected profits from advertising on the comparison site increase to that of their outside option.

Transactions Costs

Next, we study how transactions costs (ε) affect participants at the price comparison site. While the model treats these costs as exogenous, in the long-run the price comparison site could presumably invest to reduce these costs. Such investments might be in terms of ease of use of its site, better interoperability with clients, automating advertising requests, less need for prepaid accounts, and so on.

Proposition 7 *In equilibrium, an increase in transactions costs:*

- (a) *Reduces both the optimal clickthrough fee and profits of the price comparison site.*
- (b) *Leads firms to advertise less frequently and increases firm profits.*

While a formal proof is provided in the Appendix, the intuition is clear. Higher transactions costs reduce the value of the comparison site relative to the outside option. Were clickthrough fees left unchanged, firms would advertise less intensely and the comparison site's profits would decline. The comparison site optimally responds by decreasing its clickthrough fee in an attempt to mitigate the decline in advertising, but does not find it profitable to lower the fee to a level that induces the previous level of advertising. Since this fee reduction only partially compensates for the increase in transactions costs, firms advertise less frequently, which softens price competition. Both of these effects increase firm profits.

Proposition 7 reveals that firms and the price comparison site have divergent views about the benefits of reducing transactions costs. While the price comparison site benefits from a reduction, firms are harmed. This may partially explain the relative lack of cooperation between firms and price comparison sites in sharing information and making systems interoperable.

4 Differences Between the Two Models

While the models in Sections 2 and 3 are similar in that they both rationalize the switch from CPM to CPC that has occurred over the past decade, there are some potentially important differences. One key difference that we highlighted earlier concerned who won and who lost from this shift. In the simple model, increased platform profits came at the expense of inframarginal firms whereas the richer model implied that consumers ultimately bore the burden of this shift. This section highlights two additional differences: (1) In the simple model, the platform does not benefit from the ability to offer an exclusive advertising contract, while in the richer model it does (at least in the short-run); and (2) in the richer model, platform profits fall discontinuously when transactions costs increase from zero, while they decline continuously in the simple model. The endogenous outside option in the richer model drives this discontinuity. This section examines these differences, beginning with differences related to transactions costs.

4.1 “Small” vs. Zero Transactions Costs

Let $c^*(\varepsilon)$ and $U^*(\varepsilon)$ denote the optimal clickthrough fee and platform profits when transactions costs are $\varepsilon > 0$, and let $c^*(0)$ be the optimal clickthrough fee when transactions costs are zero. Our next proposition shows that, in the richer model, the case where $\varepsilon = 0$ is materially different than the ε -small case.

Proposition 8 *In the richer model, there is a discontinuity in both the optimal clickthrough fee and platform profits at $\varepsilon = 0$ (i.e., when there are no transactions costs). There are no such discontinuities in the simple model. Specifically,*

(a) *In the richer model,*

$$\begin{aligned} c^*(0) &= (r - m)\gamma \text{ but } \lim_{\varepsilon \rightarrow 0} c^*(\varepsilon) = (r - m)\gamma \frac{N - 1}{N} \\ U^*(0) &= S(r - m)\gamma \text{ but } \lim_{\varepsilon \rightarrow 0} U^*(\varepsilon) = S(r - m)\gamma \frac{N - 1}{N} \end{aligned}$$

(b) *In the simple model,*

$$\begin{aligned} c^*(0) &= \lim_{\varepsilon \rightarrow 0} c^*(\varepsilon) = v \\ U^*(0) &= \lim_{\varepsilon \rightarrow 0} U^*(\varepsilon) = \left(1 - (1 - \alpha)^N\right) v \end{aligned}$$

The reader might wonder why the richer model leads to a discontinuity at $\varepsilon = 0$ while the simple model displays no such discontinuity. The answer is that the simple model assumes the outside option of advertisers (the right-hand-side of equation 1) is independent of activity on the platform's site, while the richer model captures an environment where the outside option depends on activity at the comparison site (the right-hand-side of equation 7). Since the outside option in the richer model depends on the equilibrium advertising propensity, α , changes in transactions costs affect the value to firms both on *and off* the price comparison site, owing to cross-channel effects. When $N \rightarrow \infty$, the cross-channel effects disappear in the richer model, and the value of the outside option converges to $(r - m)\lambda L$, which is independent of α . Consequently, the discontinuity identified in Proposition 8 vanishes for this case. In short, endogenizing the outside option fundamentally changes the role of transactions costs in the model.

Proposition 8 reveals that small listing frictions may lead to a discontinuous decline in both the clickthrough fees and profits of the price comparison site. Thus, while transactions costs on the internet may be small (and most certainly are compared to offline markets), ignoring them completely may lead one to overstate the market power enjoyed by even a monopoly platform. Notice that the size of the discontinuity in the richer model at $\varepsilon = 0$ depends on N , and declines as N increases.

4.2 Exclusive Contracts

Up until now, we have assumed that the contracts offered by the platform were non-exclusive: All firms were offered the same contract, (T, c) , and able to advertise at the site under these terms. Would a platform benefit from the ability to offer a single firm the exclusive right to advertise on its platform?

In the simple model, the answer is no. In effect, that analysis permits the platform to offer an exclusive contract by regulating the number of entrants, n , through its choice of the clickthrough fee. By selecting a high enough clickthrough fee, the platform attracts only the firm with the highest quality ad, and the clickthrough fee is set in such a way that this single firm is indifferent between participating or not. Clearly, the platform can do no better under an exclusive arrangement. However, Proposition 2 showed that, generically, the platform does not wish to restrict entry to a single firm. Thus, the ability to offer an exclusive contract is not beneficial in the simple model.

The same is not true of the richer model. Suppose first that payments under an exclusive contract are contingent on the exclusive firm's decision to advertise its price at the price comparison site. In this case, if the price comparison site offered such a contract, the exclusive firm could sign it (thus foreclosing rivals from advertising on the site) and then choose not to advertise its price. This would guarantee the exclusive firm $1/N$ of the shoppers visiting its own website and eliminate any payments to the price comparison site. Thus, the price comparison site must moderate its fee under such an exclusive contract to make this option unattractive.

For this reason, when the price comparison site solely relies on clickthrough fees under an exclusive contract, its optimal fee is

$$c^E = (r - m)\gamma \frac{N - 1}{N} - \frac{\varepsilon}{S}$$

This leaves the exclusive firm indifferent between advertising a price of r at the comparison site and not advertising, and thus an optimizing exclusive firm can do no better than advertise a price of r with probability one. Similarly, if the price comparison site used only an impression fee in an exclusive arrangement, then the profit-maximizing fee is

$$T^E = S(r - m)\gamma \frac{N - 1}{N} - \varepsilon$$

Unlike optimal non-exclusive contracts, where the price comparison site has a strict incentive to solely use clickthrough fees, exclusive contingent contracts lead to a neutrality result: The price comparison site earns the same expected profits, $U^E = Sc^E$, regardless of whether it uses optimal clickthrough fees, impression fees, or a combination of the two.

To see that exclusivity benefits the price comparison site, recall that its profit under the optimal non-exclusive contract is $U^* = (1 - \zeta^N)Sc^*$, where c^* is

$$c^* = (r - m)\gamma \frac{N - 1}{N} - \frac{\varepsilon}{\zeta^{N-1}S}$$

Since $\zeta > 0$, it is immediate that $c^E > c^*$ and $U^* < U^E$, i.e., profits are higher under the exclusive arrangement. Intuitively, by committing to eliminate competition between the exclusive firm and its rivals, the comparison site is able to charge a higher clickthrough fee and earn higher profits.

To summarize:

Proposition 9 *In the simple model, the platform does not benefit from using an exclusive contract, whereas in the richer model it does.*

One might worry that this difference between the two models is driven by restrictions on the scope of exclusive contracts. In particular, the analysis above assumes that contracts consist only of *contingent* impression or clickthrough fees. While this is helpful in isolating the pure effect of offering an exclusive in our earlier setting, it is arguably unrealistic. Suppose that platforms were permitted to offer an exclusive in exchange for a non-contingent fee; that is, an advertiser would agree to pay an up front fixed amount A for the right to be the exclusive advertiser on the platform.

In the simple model, this amendment makes no difference—the platform still does not benefit from exclusives. To see this, recall that the platform sets the clickthrough fee such that it extracts the full incremental expected surplus from the marginal advertiser. Consequently, if the platform charged a non-contingent fee to an exclusive advertiser, it can do no better than to capture all of the expected surplus.

In the richer model, the platform may benefit from using non-contingent, *exclusive*, fees. To see this, suppose that the price comparison site offered firm 1 (say) the following exclusive contract: Pay a fixed amount $A = (r - m) \gamma S - \varepsilon$ for the exclusive right to list on the site. (Note that this fee is paid regardless of whether the firm advertises its price, and that neither clickthrough nor impression fees are needed under this contract.) Were firm 1 to accept this contract, it will (optimally) list a price of r and earn profits of $(r - m) \lambda L$. If the firm believes that some other firm will accept a similar contract if it declines, then firm 1 has a weak incentive to accept the exclusive contract.¹⁴ This incentive can be made strict by reducing A by an arbitrarily small amount. Under this contract, the price comparison site captures all of the surplus from shoppers. Notice that this is an improvement over an exclusive with contingent fees and, of course, over non-exclusive contracts as well.

If exclusives are superior (at least in the richer model), why aren't they more often observed in practice? First, in many settings the simple model is an apt description, and here exclusives confer no benefits. In settings where the richer model is appropriate, the benefits from exclusion derive from the “stickiness” of shoppers on the price comparison site. Specifically, shoppers are implicitly assumed to visit the price comparison site regardless of the value of the price information obtained there. This degree of loyalty to a price comparison site seems dubious. Presumably shoppers will soon realize that a comparison

¹⁴Such beliefs could be sustained were the gatekeeper able to make a (potentially infinite) sequence of exclusive offers to other firms—for instance, following with the same offer to firm 2, firm 3, ... firm N and then restarting with firm 1, *ad infinitum*.

site running an exclusive ad for a given product is a poor source of price information and bargains, and desert it in favor of another site. The combination of unhappy consumers and disenfranchised firms is a volatile mix for a comparison site seeking to maintain its competitive position.

5 Conclusion

It is by now a hackneyed phrase to refer to the internet as a dynamic environment. Yet, in terms of the pricing of advertising, it is hardly an exaggeration. The so-called CPM model of advertising—pricing ads based on the number of eyeballs or impressions—has been standard for over 100 years in old media models of advertising. This is mainly a reflection of the fact that making pricing contingent on the effectiveness of advertising was either impossible or prohibitively costly. As retail mogul John Wanamaker famously quipped, “Half the money I spend on advertising is wasted; the trouble is I don’t know which half.”

Given the prevalence and longevity of the CPM model, it is hardly surprising that it was embraced at the dawn of the internet. While the internet effected many dramatic changes in the business landscape—breaking down geographic barriers, remaking supply chains, democratizing the publication process—a seemingly minor change, the ability to see that a consumer clicked on an ad, was arguably the most revolutionary. This change allowed pricing contingent on advertising effectiveness, the CPC model of pricing. Here, firms were only required to pay if a consumer clicked, i.e. if an ad were effective.

The conventional view is that this change largely benefited advertisers. Under the CPC model, Wanamaker would no longer have to pay for the half of his advertising that was wasted, nor would he have to guess which half this was. It is obvious that the CPC model shifts the risk of ineffective advertising to the platform and away from advertisers. Moreover, the CPC model helps mitigate technological contrivances such as impressions consisting purely of armies of bots created by the platform. The puzzle then is why platforms were so eager to embrace this new model.

We have shown that platforms, in fact, benefit from the CPC model. The key intuition is that CPM rates impact platform revenues through the advertising propensity of all firms whereas CPC pricing affects revenues through an order statistic—the chance that at least one firm advertises. This latter effect leads to a more muted reduction in advertising

propensities from higher rates, thus making the *exclusive* use of clickthrough fees optimal. We first showed this result in a stark environment where the attractiveness of ads and the value of an advertiser's outside option (eschewing the advertising channel) are both exogenous. We also showed that the same intuition holds in a richer model where attractiveness is under the control of firms and where the value of the outside option is endogenous.

This richer model allowed us to study other aspects of online advertising, such as conversion rates (the chance that a click turns into a sale). The low level of conversions (only about 5% at most) is a perennial worry of online platforms and advertisers. We showed that both the platform and advertisers benefited from investments to improve conversions, but that neither party stands to recover the full value of such an investment. Thus, there is a familiar hold-up problem present. Amazon's momentous decision to vertically integrate, becoming both a platform and a seller, may then be seen as a strategic response to this dilemma.

While the main intuition for the optimality of CPC pricing prevails in both models, there are important differences. One essential difference relates to transactions costs of online advertising. While the internet is rightly viewed to have dramatically lowered frictions associated with advertising, they have not been eliminated entirely. In the simple model, the platform's profits increase continuously as transactions costs decrease. Endogenizing the value of the outside option, however, produces different effects. Here we showed that there is a discontinuous increase in the profits of the platform when transactions costs are eliminated entirely. Thus, while transactions costs online are certainly small, ignoring them entirely can yield erroneous conclusions about market power.

To summarize, we see our main contribution as shedding light on how a seemingly trivial technological change—the ability to (imperfectly) observe advertising effectiveness—fundamentally changed how advertising is priced. The value unlocked by this shift in pricing is enormous. For instance Google, whose revenues are almost entirely derived through CPC pricing, did not exist at the dawn of the internet but now enjoys a market capitalization of \$172 billion—a testament to the efficacy of CPC.

A Appendix

Proof of Proposition 2

Proposition 2 relies, in part, on the following two lemmas.

Lemma 2 $U(n+1) - U(n)$ is strictly decreasing in n for all $n \geq 1$.

Proof. To establish the claim, we show that the difference in differences is decreasing in n .

For all $n \geq 2$, we have

$$\begin{aligned}
 & U(n+1) - U(n) - (U(n+2) - U(n+1)) \\
 = & \alpha v \left((1-\alpha)^n - \frac{1}{(1-\alpha)^n} \frac{\varepsilon}{\alpha v} - \left((1-\alpha)^{n+1} - \frac{1}{(1-\alpha)^{n+1}} \frac{\varepsilon}{\alpha v} \right) \right) \\
 = & \alpha v \left((1-\alpha)^n \alpha + \frac{\alpha}{1-\alpha} \frac{1}{(1-\alpha)^n} \frac{\varepsilon}{\alpha v} \right) \\
 > & 0
 \end{aligned}$$

When $n = 1$, we have

$$\begin{aligned}
 & U(2) - U(1) - (U(3) - U(2)) \\
 = & \alpha v \left((1-\alpha) - \frac{\varepsilon}{\alpha v} - \left((1-\alpha)^2 - \frac{1}{(1-\alpha)^2} \frac{\varepsilon}{\alpha v} \right) \right) \\
 = & \alpha v \left((1-\alpha) \alpha + \frac{\varepsilon}{\alpha v} \left(\frac{1}{(1-\alpha)^2} - 1 \right) \right) \\
 > & 0
 \end{aligned}$$

Obviously, $U(0) = 0$; thus, it is immediate that $U(1) - U(0) - (U(2) - U(1)) > 0$ since $\alpha \in (0, 1)$. ■

Next, we will show that a finite value of n is always optimal.

Lemma 3 *There exists a unique integer $1 \leq n^* < \infty$ that maximizes platform profits. This is the largest integer such that*

$$U(n) - U(n-1) > 0 > U(n+1) - U(n)$$

Proof. First, notice that

$$\begin{aligned}
 & \lim_{n \rightarrow \infty} U(n+1) - U(n) \\
 = & \lim_{n \rightarrow \infty} \alpha v \left((1-\alpha)^n - \frac{1}{(1-\alpha)^n} \frac{\varepsilon}{\alpha v} \right) \\
 = & -\infty
 \end{aligned}$$

Thus, for sufficiently large values of n , adding an additional firm reduces the platform's profits. Next, notice that $U(1) - U(0) > 0$, since $\varepsilon < \alpha v$. Thus, $n^* \geq 1$. These observations together with Lemma 2 imply that there exists a unique finite value of n that maximizes platform profits. ■

To complete the proof of Proposition 2, we establish the closed form solution n^* . First, notice that $n^* = 1$ if and only if

$$1 - \frac{\varepsilon}{\alpha v} > 0 > (1 - \alpha) - \frac{\varepsilon}{\alpha v}, \text{ or}$$

$$1 > \frac{\varepsilon}{\alpha v} > 1 - \alpha$$

Next, notice that $n^* = 2$ is optimal iff

$$(1 - \alpha) - \frac{\varepsilon}{\alpha v} > 0 > (1 - \alpha)^2 - \frac{1}{(1 - \alpha)^2} \frac{\varepsilon}{\alpha v}, \text{ or}$$

$$1 - \alpha > \frac{\varepsilon}{\alpha v} > (1 - \alpha)^2 - \left(\frac{1}{(1 - \alpha)^2} - 1 \right) \frac{\varepsilon}{\alpha v}$$

It is easy to see that this interval is non-empty. Similarly, for all $n > 2$, we obtain

$$\alpha v \left((1 - \alpha)^{n-1} - \frac{1}{(1 - \alpha)^{n-1}} \frac{\varepsilon}{\alpha v} \right) > 0 > \alpha v \left((1 - \alpha)^n - \frac{1}{(1 - \alpha)^n} \frac{\varepsilon}{\alpha v} \right), \text{ or}$$

$$(1 - \alpha)^{n-1} - \left(\frac{1}{(1 - \alpha)^{n-1}} - 1 \right) \frac{\varepsilon}{\alpha v} > \frac{\varepsilon}{\alpha v} > (1 - \alpha)^n - \left(\frac{1}{(1 - \alpha)^n} - 1 \right) \frac{\varepsilon}{\alpha v}$$

where again, the interval is non-empty by inspection. Using the right-hand inequality and taking logs yields

$$2n \log [1 - \alpha] < \log \left[\frac{\varepsilon}{\alpha v} \right], \text{ or}$$

$$n < \frac{1}{2} \frac{\log \left[\frac{\varepsilon}{\alpha v} \right]}{\log [1 - \alpha]}$$

Now, since n^* is an integer, we obtain as its solution:

$$n^* = \left\lceil \frac{1}{2} \frac{\log \left[\frac{\varepsilon}{\alpha v} \right]}{\log [1 - \alpha]} \right\rceil$$

where $\lceil \cdot \rceil$ denotes the integer ceiling function.

Performing the same exercise for the left-hand inequality reveals

$$n^* = \left\lfloor \frac{1}{2} \frac{\log \left[\frac{\varepsilon}{\alpha v} \right]}{\log [1 - \alpha]} \right\rfloor + 1$$

$$= \left\lceil \frac{1}{2} \frac{\log \left[\frac{\varepsilon}{\alpha v} \right]}{\log [1 - \alpha]} \right\rceil$$

where $\lfloor \cdot \rfloor$ denotes the integer floor function. This completes the proof.

Characterization of Firm Behavior in the Richer Model

First, observe that a firm that does not advertise finds it optimal to price at r . To see this, notice that for any symmetric advertising propensity $\alpha \in [0, 1]$, a firm's profits from not advertising are:

$$(p - m) \left(\lambda L + (1 - \alpha)^{N-1} \gamma \frac{S}{N} \right)$$

which is maximized at $p = r$. This yields the value of the outside option given in equation (6). The following lemma establishes the other claims regarding firm behavior in the richer model.

Lemma 4 *Given (T, c) , symmetric equilibrium firm behavior is as follows:*

- I. *When $T \geq S \left((r - m) \gamma \frac{N-1}{N} - c \right) - \varepsilon$, no firms utilize the price comparison site.*
- II. *When $T < S \left((r - m) \gamma \frac{N-1}{N} - c \right) - \varepsilon$:*
 - (a) *Each firm advertises its price on the price comparison site with probability*

$$\alpha = 1 - \left(\frac{\varepsilon + T}{S \left((r - m) \gamma \frac{N-1}{N} - c \right)} \right)^{\frac{1}{N-1}} \in (0, 1)$$

- (b) *Conditional on advertising a price at the comparison site, a firm's advertised price may be viewed as a random draw from*

$$F(p) = \frac{1}{\alpha} \left(1 - \left(\frac{(r - p) \lambda L + \frac{N(T + \varepsilon)}{((r - m) \gamma (N - 1) - Nc)} \left((r - m) \gamma - c \right)}{S \left((p - m) \gamma - c \right)} \right)^{\frac{1}{N-1}} \right) \quad (15)$$

on $[p_0, r]$, where

$$p_0 = m + \frac{1}{(S\gamma + L\lambda)} \left(\lambda L (r - m) + \frac{N(T + \varepsilon)}{((r - m) \gamma (N - 1) - Nc)} \left((r - m) \gamma - c \right) + Sc \right) \in (m, r).$$

- (c) *Each firm earns expected profits of*

$$E\pi = (r - m) \lambda L + \frac{T + \varepsilon}{N \left(1 - \frac{c}{(r - m) \gamma} \right) - 1}$$

Proof. Part I.

By way of contradiction, suppose there is a symmetric equilibrium in which firms are active under the stated conditions. By the usual reasoning (see Baye and Morgan, 2001), this equilibrium must be in atomless mixed-strategies with an upper support of r . Letting

F denote the distribution of prices induced in this putative symmetric equilibrium, equation (5) implies that a firm advertising a price of r earns $\pi(r)$. Using equation (6), a firm that eschews the comparison site earns profits of π_0 . Since $F(r) = 1$ and, by hypothesis, $\alpha > 0$, we thus have

$$\begin{aligned}\pi_0 - \pi(r) &= T + \varepsilon - (1 - \alpha)^{N-1} S \left((r - m) \gamma \frac{N-1}{N} - c \right) \\ &> T + \varepsilon - S \left((r - m) \gamma \frac{N-1}{N} - c \right) \\ &\geq 0\end{aligned}$$

since $T \geq S \left((r - m) \gamma \frac{N-1}{N} - c \right) - \varepsilon$. Hence, any firm is strictly better off eschewing the price comparison site, which contradicts the hypothesis that $\alpha > 0$.

Part II.

The derivation of part (a) was provided in the text. Notice that $\alpha \in (0, 1)$ is a well-defined probability under the conditions stated, since

$$0 < T + \varepsilon < S \left((r - m) \gamma \frac{N-1}{N} - c \right)$$

To prove part (b), notice that as in Baye and Morgan (2001), an equilibrium with active participation at the comparison site has the following two key properties: (1) A firm must be indifferent between advertising its price at the clearinghouse or not; and (2) a firm must earn the same expected payoff from advertising any price $p \in [p_0, r]$ at the comparison site. Equating $\pi(p)$ and π_0 in equations (5) and (6) of the text and solving for $(1 - \alpha F(p))^{N-1}$ yields

$$(1 - \alpha F(p))^{N-1} = \frac{(r - p) \lambda L + \frac{N(T+\varepsilon)}{((r-m)\gamma(N-1)-Nc)} ((r - m) \gamma - c)}{S((p - m) \gamma - c)} \quad (16)$$

Solving equation (16) for $F(p)$ gives the form in the lemma. It is clear that $F(r) = 1$; the expression for p_0 may be readily verified by setting $F(p_0) = 0$ and solving for p_0 . Finally, it is a routine matter to verify that F is strictly increasing on $[p_0, r]$ and that a player cannot improve his payoff by advertising a price $p \notin [p_0, r]$.

To prove part (d) substitute for $(1 - \alpha)^{N-1}$ using equation (16) into equation (6) to obtain the expression in the lemma. ■

Proof of Propositions 5, 6 and 7

We first apply the Implicit Function Theorem to equation (13) at $z = \zeta$, the complement of the optimal equilibrium advertising propensity, to establish the claims regarding the

impact on firms' advertising propensities.

$$\partial\zeta/\partial\lambda = -\frac{\frac{\partial^2 U}{\partial\lambda\partial z}}{\frac{\partial^2 U}{\partial z^2}} = 0 \quad (17)$$

$$\frac{\partial\zeta}{\partial\gamma} = -\frac{\frac{\partial^2 U}{\partial\gamma\partial z}}{\frac{\partial^2 U}{\partial z^2}} = \frac{-z^{N-1}S(r-m)(N-1)}{(N-1)^2 z^{N-2}S(r-m)\gamma + N(N-1)z^{-N-1}\varepsilon} < 0 \quad (18)$$

and

$$\frac{\partial\zeta}{\partial\varepsilon} = -\frac{\frac{\partial^2 U}{\partial\varepsilon\partial z}}{\frac{\partial^2 U}{\partial z^2}} = \frac{(N-1)z^{-N} + 1}{(N-1)^2 z^{N-2}S(r-m)\gamma + N(N-1)z^{-N-1}\varepsilon} > 0 \quad (19)$$

We next establish the effects on the optimal clickthrough fee. Since the optimal clickthrough fee is given by

$$c = (r-m)\gamma\frac{N-1}{N} - \zeta^{1-N}\frac{\varepsilon}{S}$$

we have

$$\begin{aligned} \frac{dc}{d\lambda} &= \frac{\partial c}{\partial\zeta}\frac{\partial\zeta}{\partial\lambda} + \frac{\partial c}{\partial\lambda} \\ &= (N-1)\frac{\varepsilon}{S}\zeta^{-N}\frac{\partial\zeta}{\partial\lambda} \\ &= 0 \end{aligned}$$

by equation (17). Next, notice that

$$\begin{aligned} \frac{dc}{d\gamma} &= \frac{\partial c}{\partial\zeta}\frac{\partial\zeta}{\partial\gamma} + \frac{\partial c}{\partial\gamma} \\ &= \left(\frac{\varepsilon}{S}(N-1)z^{-N}\right)\left(\frac{-z^{N-1}S(r-m)(N-1)}{(N-1)^2 z^{N-2}S(r-m)\gamma + N(N-1)z^{-N-1}\varepsilon}\right) + (r-m)\frac{N-1}{N} \\ &= (r-m)\frac{N-1}{N}\left(\frac{(N-1)z^{N-1}S(r-m)\gamma + Nz^{-N}\varepsilon - \varepsilon N}{(N-1)z^{N-1}S(r-m)\gamma + Nz^{-N}\varepsilon}\right) \end{aligned}$$

Using the fact that equation (13) implies $(N-1)z^{N-1}S(r-m)\gamma = (N-1)z^{-N}\varepsilon + \varepsilon$, we can write this as

$$\begin{aligned} \frac{dc}{d\gamma} &= (r-m)\frac{N-1}{N}\left(\frac{(N-1)z^{-N}\varepsilon + \varepsilon + Nz^{-N}\varepsilon - \varepsilon N}{(N-1)z^{N-1}S(r-m)\gamma + Nz^{-N}\varepsilon}\right) \\ &= (r-m)\frac{N-1}{N}\varepsilon\left(\frac{(N-1)(z^{-N}-1) + Nz^{-N}}{(N-1)z^{N-1}S(r-m)\gamma + Nz^{-N}\varepsilon}\right) > 0 \end{aligned}$$

since $z < 1$. Similarly,

$$\begin{aligned}\frac{dc}{d\varepsilon} &= \frac{\partial c}{\partial \zeta} \frac{\partial \zeta}{\partial \varepsilon} + \frac{\partial c}{\partial \varepsilon} \\ &= \left((N-1) z^{-N} \frac{\varepsilon}{S} \right) \left(\frac{(N-1) z^{-N} + 1}{(N-1)^2 z^{N-2} S (r-m) \gamma + N (N-1) z^{-N-1} \varepsilon} \right) - z^{1-N} \frac{1}{S} \\ &= z^{-N} \frac{1}{S} \left(\frac{(N-1) z^{-N} \varepsilon + \varepsilon}{z^{N-2} S (r-m) \gamma (N-1) + N z^{-N-1} \varepsilon} - z \right)\end{aligned}$$

Once again using equation (13), we have

$$\begin{aligned}\frac{dc}{d\varepsilon} &= z^{-N} \frac{1}{S} \left(\frac{(N-1) z^{N-1} S (r-m) \gamma}{N z^{N-2} S (r-m) \gamma \frac{N-1}{N} + N z^{-N-1} \varepsilon} - z \right) \\ &= z^{-(N-1)} \frac{1}{S} \left(\frac{1}{1 + z^{-2N+1} \frac{N\varepsilon}{S(r-m)\gamma(N-1)}} - 1 \right) < 0\end{aligned}$$

Turning to the claims regarding firm profits, notice that equilibrium firm profits are given by π_0 in equation (6) evaluated at $z = \zeta$:

$$E\pi = (r-m) \lambda L + (r-m) \gamma \frac{S}{N} \zeta^{N-1}$$

Hence

$$\begin{aligned}\frac{dE\pi}{d\lambda} &= \frac{\partial E\pi}{\partial \lambda} + \frac{\partial E\pi}{\partial \zeta} \frac{\partial \zeta}{\partial \lambda} \\ &= (r-m) L > 0\end{aligned}$$

and

$$\begin{aligned}\frac{dE\pi}{d\varepsilon} &= \frac{\partial E\pi}{\partial \varepsilon} + \frac{\partial E\pi}{\partial \zeta} \frac{\partial \zeta}{\partial \varepsilon} \\ &= (r-m) \gamma \frac{S}{N} (N-1) \frac{\partial \zeta}{\partial \varepsilon} > 0\end{aligned}$$

by equation (19). Finally,

$$\begin{aligned}\frac{dE\pi}{d\gamma} &= \frac{\partial E\pi}{\partial \gamma} + \frac{\partial E\pi}{\partial \zeta} \frac{d\zeta}{d\gamma} \\ &= (r-m) \frac{S}{N} z^{N-1} \\ &\quad + \left((r-m) \gamma \frac{S}{N} (N-1) z^{N-2} \right) \left(\frac{-z^{N-1} S (r-m) (N-1)}{(N-1)^2 z^{N-2} S (r-m) \gamma + N (N-1) z^{-N-1} \varepsilon} \right) \\ &= (r-m) \frac{S}{N} z^{N-1} \left(1 - \left(\frac{z^{N-2} S (r-m) \gamma}{z^{N-2} S (r-m) \gamma + \frac{N\varepsilon}{N-1} z^{-N-1}} \right) \right) > 0\end{aligned}$$

To establish the impact on the comparison site's profits, notice that its equilibrium profits are

$$U^* = (1 - \zeta^N) c$$

Since ζ and c are independent of λ , it then immediately follows that $dU^*/d\lambda = 0$. Since we have already established that an increase in ε increases ζ and reduces the optimal c , we have

$$\frac{dU^*}{d\varepsilon} < 0$$

Finally, since we have already shown that an increase in γ reduces ζ and increases the optimal c , we have

$$\frac{dU^*}{d\gamma} > 0$$

This completes the proof.

Proof of Proposition 8

It is sufficient to prove parts (a) and (b) of the proposition to establish the result.

Proof of (a): Fix $\varepsilon = 0$ and suppose that the price comparison site chooses $T = 0$ and $c = (r - m) \gamma$. We claim that it is an equilibrium for firms to advertise with probability one and charge a price of r . To see this, note that if a firm conforms to the putative strategy and advertises a price of r on the site, it earns

$$\begin{aligned} \pi &= (r - m) \lambda L + ((r - m) \gamma - c) \frac{S}{N} \\ &= (r - m) \lambda L \end{aligned}$$

since $c = (r - m) \gamma$, $\alpha = 1$, and $T = 0$. If the firm deviates and does not advertise, it earns.

$$\begin{aligned} \pi_0 &= (r - m) \lambda L + (r - m) (1 - \alpha)^{N-1} \frac{\gamma S}{N} \\ &= (r - m) \lambda L \end{aligned}$$

since $\alpha = 1$. Thus, each firm is indifferent between advertising or not, so $\alpha = 1$ is a (weak) best response. Next, notice that a firm cannot profitably deviate by lowering its price. This is obvious when the firm does not advertise and, when it does advertise, it earns negative profits per shopper if it chooses any price $p < r$. Therefore, pricing at r , and hence the asserted strategies, are an equilibrium response to $T = 0$ and $c = (r - m) \gamma$. Finally, notice that with this fee structure the price comparison site captures all of the contestable surplus, so there is no contract yielding it higher profits. Thus, when $\varepsilon = 0$, $c^*(0) = (r - m) \gamma$ and $U^*(0) = S (r - m) \gamma$.

Fix $\varepsilon > 0$. Equation (13) implies

$$\lim_{\varepsilon \rightarrow 0} \frac{\varepsilon}{\zeta^{N-1}} = \lim_{\zeta \rightarrow 0} \frac{S(r-m)\gamma(N-1)}{((N-1)\zeta^{-N} + 1)} = 0$$

since ζ is monotonically decreasing in ε (by Proposition 7) and, in fact, equation (13) implies $\zeta \rightarrow 0$ as $\varepsilon \rightarrow 0$. Using the expression for c^* in Proposition 1,

$$\begin{aligned} \lim_{\varepsilon \rightarrow 0} c^*(\varepsilon) &= (r-m)\gamma \frac{N-1}{N} - \lim_{\varepsilon \rightarrow 0} \frac{\varepsilon}{\zeta^{N-1} S} \\ &= (r-m)\gamma \frac{N-1}{N} \end{aligned}$$

and hence the comparison site's profits converge to

$$\lim_{\varepsilon \rightarrow 0} U^*(\varepsilon) = (1 - \zeta^N) S c = S(r-m)\gamma \frac{N-1}{N}$$

Proof of (b): First, note that as $\varepsilon \rightarrow 0$, $n^* \rightarrow N$. Taking limits of equations (3) and (4) as $\varepsilon \rightarrow 0$ thus yields $\lim_{\varepsilon \rightarrow 0} c^*(\varepsilon) = v$ and $\lim_{\varepsilon \rightarrow 0} U^*(\varepsilon) = \left(1 - (1 - \alpha)^N\right) v$. Next, notice that when $\varepsilon = 0$, the same expressions obtain.

References

- [1] Armstrong, Mark (2006): “Competition in Two-Sided Markets,” *RAND Journal of Economics*, 37(3), 668-691.
- [2] Athey, Susan and Glenn Ellison (forthcoming): “Position Auctions with Consumer Search,” *Quarterly Journal of Economics*.
- [3] Baye, Michael R. and John Morgan (2001): “Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets,” *American Economic Review*, 91(3), 454-474.
- [4] Baye, Michael R. Baye and John Morgan (2002), “Information Gatekeepers and Price Discrimination on the Internet,” *Economics Letters*, 76(1), 47-51.
- [5] Baye, Michael R., John Morgan, and Patrick Scholten (2004a): “Price Dispersion in the Small and Large: Evidence from an Internet Price Comparison Site,” *Journal of Industrial Economics* 52(4), 463-469.
- [6] Baye, Michael R., John Morgan, and Patrick Scholten (2004b): “Temporal Price Dispersion: Evidence from an Online Consumer Electronics Market,” *Journal of Interactive Marketing* 18(4), 101-115.
- [7] Baye, Michael R., John Morgan, and Patrick Scholten (2006): “Information, Search, and Price Dispersion,” in *Handbook of Information Systems* Vol. 1 (Terrence Hendershott, ed.), Amsterdam: Elsevier, 323-375.
- [8] Brynjolfsson, Erik and Michael D. Smith (2000), “Frictionless Commerce? A Comparison of Internet and Conventional Retailers,” *Management Science* 46(4), 563-585.
- [9] Brown, Jennifer, and John Morgan (2009): “How much is a Dollar Worth? Tipping versus Equilibrium Coexistence on Competing Online Auction Sites,” *Journal of Political Economy*, 117(4), 668-700.
- [10] Caillaud, Bernard and Bruno Jullien (2003): “Chicken and Egg: Competition among Intermediation Service Providers.” *RAND Journal of Economics*, 34(2), 521-552.

- [11] Chevalier, Judith and Austan Goolsbee (2003): "Measuring Prices and Price Competition Online: Amazon.com and BarnesandNoble.com," *Quantitative Marketing and Economics* 1, 203-222.
- [12] Edelman, Benjamin, Michael Ostrowsky and Michael Schwarz (2007): "Internet Advertising and the Generalized Second-Price Auctions: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97(1), 242-259.
- [13] Edelman, Benjamin and Michael Schwarz (2006): "Optimal Auction Design in a Multi-Unit Environment: The Case of Sponsored Search Auctions," Harvard University Working Paper.
- [14] Ellison, Glen and Sara Fisher Ellison (2009), "Search, Obfuscation, and Price Elasticities on the Internet," *Econometrica*, 77(2), 427-452.
- [15] Hossain, Tanjim and John Morgan (2009): "The Quest for QWERTY," *American Economic Review*, 99 (2), 435-440.
- [16] Hossain, Tanjim, Dylan Minor and John Morgan (forthcoming): "Competing Matchmakers: An Experimental Analysis," *Management Science*
- [17] Liebowitz, Stanley and Stephen Margolis (1994): "Network Externality: An Uncommon Tragedy," *Journal of Economic Perspectives*, 8(2), 133-150.
- [18] Nahm, Jae (2003), "The Gatekeeper's Optimal Fee Structure when Sellers can Price Discriminate," *Economics Letters*, 80 (1), 9-14.
- [19] Narasimhan, Chakravarthi (1988): "Competitive Promotional Strategies." *Journal of Business*, 61, pp. 427-449.
- [20] Rochet, Jean-Charles and Jean Tirole (2003): "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association*, 1(4), 990-1029.
- [21] Rosenthal, Robert W. (1980): "A Model in Which an Increase in the Number of Sellers Leads to a Higher Price," *Econometrica* 48(6), 1575-1580.
- [22] Tellis, Gerard J., Yiding Yin, and Rakesh Niraj (2009): "Does Quality Win: Network Effects versus Quality in High Tech Markets," *Journal of Marketing Research*, 46(2), 135-149.

- [23] Varian, Hal (1980): “A Model of Sales,” *American Economic Review* 70, 651-659.
- [24] Varian, Hal (2007): “Position Auctions,” *International Journal of Industrial Organization* 25(6), 1163-1178.
- [25] Weyl, E. Glen (2010): “A Price Theory of Multi-Sided Platforms,” *American Economic Review* 100(4), 1642-1672.

Figure 1:

Representative Ads Displayed on Search Engine, Social Networking and Email Platforms

Bing (Search: *How to play a guitar*)

Sponsored sites

[Learn to play Guitar](#)

Affordable **Guitar** Lessons-Pro Teachers For All Ages - Begin today!
[TakeLessons.com/Guitar](http://www.TakeLessons.com/Guitar)

[Learn To Play Guitar Now](#)

Free 77 Page Guide Reveals **How To Learn & Play Guitar** Like a Pro Now
www.GuitarChordsMadeSimple.com

[Easy Guitar Lesson Videos](#)

24 Free Video **Guitar** Lessons - Fast, Fun, Step-by-Step Results
Video.GuitarTricks.com/Lessons

[Best 10 Guitar Lessons](#)

We Exposed 10 Top **Guitar** Courses! of 2011 - Which Ranks #1?
www.OnlineLessonReviews.com/Guitar

[acoustic guitar lessons](#)

In-Depth Reviews of **Guitar** Lessons sites: Smart or Scam? We Tell.
www.reviewsnest.com/guitarlessons

LinkedIn (Social Network)

Ads by LinkedIn Members	
Create Rapid eLearning Easily create interactive eLearning. Games, Simulations & more. Free Trial.	PhD Education Leadership Online degree prepares you to lead in educational settings.
	Find A Great Job Now Find Out Who's Hiring - Jobs Paying Over \$100,000
	

Ads

[Ridiculous HDTV Sales](#)

Every Day People Save 90% OFF - The Best and Newest Televisions!
www.quibids.com

[Healthcare & Hospital TV](#)

Great Deals on Healthcare Grade Hospital LCD TV's . Shop Now!
mdmhealthcare.com/HospitalTV's

[Panasonic Full HD 3D](#)

Experience Panasonic's Full HD 3D Technology At The Official Site.
Panasonic.com/3D

Gmail (Subject Line: *Where did you buy your TV?*)

