

The Informativeness of On-Line Advertising*

Greg Taylor[†]

Oxford Internet Institute, University of Oxford, Oxford, UK

and

Division of Economics, University of Southampton, Southampton, UK

2nd March 2011

Abstract

Sending general advertisements with inflationary claims may attract additional visitors with whom an advertiser is poorly matched. This is costly when ads are priced per-click because many visitors (clickers) will not purchase. This renders per-click advertising particularly conducive to the transmission of information via ads. The admissibility of information transmission depends not only on advertiser behaviour, but also upon consumers' interpretation of and trust in ads. In less conducive environments, consumers quickly learn to place little stock in the claims they see advertised. This mechanism undermines the ability of advertisers and consumers to communicate under per-impression or per-sale fee structures. Consumers benefit from increased informativeness, but distortions introduced by the market power given to advertisers imply that society may be better-off with no information transmission taking place.

Keywords: obfuscation; advertising; cheap talk; performance-based pricing

JEL classification: L20; D02; D83

This article was published as:

TAYLOR, G. (2011): "The Informativeness of On-Line Advertising," *International Journal of Industrial Organization*, 29(6), 668–677.[‡] doi: 10.1016/j.ijindorg.2011.03.001.

*This and other papers are available at <http://www.greg-taylor.co.uk>. I am indebted to Maksymilian Kwiek, Robin Mason, David Myatt, Joshua Gans (in his editorial capacity), an anonymous referee, and seminar participants at Oxford, Google (Brussels), and Universitat Oberta de Catalunya for very helpful comments and suggestions. Financial support from the ESRC is gratefully acknowledged.

[†]Correspondence address: Oxford Internet Institute, University of Oxford, 1 St. Giles, Oxford, UK, OX1 3JS. Email: greg.taylor@oii.ox.ac.uk.

[‡]Journal pagination is indicated in the margins.

1 INTRODUCTION

p668 I investigate the role played by the prevailing fee structure in determining whether advertisers and consumers are able to communicate and the extent to which information about products can be transmitted via advertisements. At the core of these results is the following simple idea: a per-click advertising fee can serve as a disincentive to sending uninformative advertisements of general appeal by making it costly for an advertiser to attract visits from consumers with whom it is poorly matched. By contrast, since a sale to a poorly matched consumer is as good as one to anyone else, firms that pay for ads on a per-sale basis have an incentive to attract a visit from any consumer that will purchase with positive probability. Likewise, once a consumer is shown an ad, the cost of that impression is sunk and firms that pay per-impression are incentivised to attract any consumer with some positive probability of purchase. This incentive causes advertiser-consumer communication to break down and undermines the existence of fully-informative equilibria. The usefulness of advertisements is then reduced and consumer welfare is harmed. However, for some specifications of search costs, reduced ad information content is beneficial for publishers and society as a whole.

The unique capacity of Internet publishers to monitor and track users' activity has given rise to a number of novel structures for the pricing of advertisement facilities. My focus here will be on three fee structures in particular:

1. Pay-per-click (PPC): Advertisers pay each time a consumer clicks on their ad.
2. Pay-per-impression (PPI) Advertisers are charged each time their ad is shown to a consumer, regardless of whether that consumer takes any further action.
3. Pay-per-sale (PPS): An advertiser must pay for each consumer that clicks on its advertisement and subsequently makes a purchase.

As of 2010, PPI pricing accounted for around 35% of online ad spending, with performance-based pricing (which encompasses PPC and PPS) accounting for a further 61%.¹

Enforcement of honest advertiser behaviour becomes more complicated on the Internet, where a multitude of publishers broadcast advertisements that transcend jurisdictions. Large publishers may intervene to prevent abuse of their advertising resource, but the large and shifting volumes of advertisements handled by major publishers makes perfect enforcement impractical. Small, independent publishers are much less likely to have the resources or inclination to police their advertisers. Moreover, because both the size and types of audience delivered depend upon consumers' interpretation of ads in a non-trivial fashion, it is not a priori clear whether publishers benefit from more informative ads at all.

¹IAB Internet Advertising Revenue Report 2010 First Half-Year Results; available at http://www.iab.net/insights_research/947883/adrevenue-report (accessed 4th Feb 2011).

Arguably, the most effective check on deceptive advertiser behaviour are the consumers themselves. Consumers that view many advertisements each day and are media literate are likely to be intrinsically aware of the appropriate degree of trust associated with the messages contained therein. That is to say, rational consumers cannot be systematically deceived: when the environment is one that is not conducive to honest reporting, consumers are likely to place little stock in the claims that they see advertised. The result is that meaningful communication between advertiser and consumer becomes difficult in such environments. It is this equilibrium intuition that rests at the heart of the present paper.

Although the focus of this paper is on-line advertising, the work is of broader interest insofar as the above fee structures (or their analogues) are used in traditional media advertising. More broadly, other pseudo-advertising environments such as price comparison sites often use a fee structure that occupies one or more of the above categories. For example, the fee structure at *shopper.com* is a combination of per-impression and per-click pricing, and Baye, Morgan, and Scholten (2004) report the efficacy of a per-click fee in enforcing honest reporting (of prices) there. On-line auction platform eBay, whose fee structure most resembles a hybrid fixed fee/PPS arrangement, often hosts listings in which extraneous terms are appended to the title—evidently in the hope of capturing the interest of consumers shopping for substitute goods, and in obvious ignorance of the effect that such behaviour has on the overall usefulness of the search tool. Similar incentive structures also seem likely to be active in other areas of economic activity. For example, a firm that pays its sales force on a salary or commission basis is likely to encourage those salespeople to aggressively pursue leads, whilst a sales team paid per-lead is more likely to be discouraged from wasting time with clients that seem unlikely to buy. These factors may, in turn, influence a consumer's willingness to indulge a salesperson's approach.

Nelson (1974) pioneered the study of the informational role of advertising—notably with the idea that the very existence of an advertisement may be informative. Nelson also acknowledged that advertisements may not always be honest. In the tradition of Nelson, Chen and He (2006) study the implicit information content of ads. They consider bidding in position auctions, which are used to allocate sponsored search advertisements at search engines such as Google. Each firm has some probability of matching with an arbitrary consumer. Since firms that match more often value clicks more highly, consumers are able to make inferences about advertisers' match probabilities based upon the observed ordering of advertisements within a list and knowledge of advertisers' bidding strategies. In contrast to these papers, my focus is on the *explicit* information content of advertising messages, and I show that PPC environments can enable substantial additional information transmission over and above that implicitly revealed by the allocation of ads alone.

The literature has also treated the case in which advertisements are explicitly in-

formative. Anderson and Renault (2006), Grossman and Shapiro (1984), and Meurer and Stahl (1994) are examples of papers of this kind. In contrast to the model below, all of these papers share the assumption that any informational content of an ad must be non-deceptive.

There is a small but growing literature on the efficacy of various advertising fee structures in on-line advertising—with the primary focus having been on vertically differentiated products. Dellarocas and Viswanathan (2008) show that pay for performance fee structures tend to favour low quality firms and yield lower surpluses for all. Sundararajan (2003) shows that performance-based pricing of digital marketing can not screen out low quality advertisers. Agarwal, Athey, and Yang (2009) develop a model to explore the problem of a publisher that must aggregate bids across multiple actions in a pay-per-action environment.

More closely related to this paper, Athey and Ellison (2011) find that pay-per-click position auctions are less conducive to the obfuscation of advertisements than are pay-per-action auctions in a model of sponsored search advertising with firms that have some idiosyncratic probability of satisfying a given consumer. In contrast to Athey and Ellison's model, I model advertisements as cheap talk messages and explicitly examine the effect of fee regime choice on consumers' equilibrium beliefs. The degree of sustainable communication—which depends on both the firms' willingness to talk and the consumers' willingness to listen—grows out of this framework as a major theme of this paper. This framework is useful for understanding when information transmission can be sustained under PPI and PPS fee regimes (Athey and Ellison's model always implies zero information transmission in pay-per-action environments). My paper is also distinct from Athey and Ellison (2011) in modelling PPI regimes (which account for more than a third of on-line ad spending), and considering the effect of pay-per-sale on the pricing of final goods—which I demonstrate to have important welfare implications.

Zhu and Wilbur (2011) model a platform that offers hybrid auctions in which advertisers can choose to pay either per-impression or per-click. Sellers who advertise for the purpose of brand exposure (rather than to directly make sales) have an incentive to exert low effort in attracting clicks and select a PPC payment scheme in order to advertise their brand at little or no cost. The authors show that there exist rational expectations for the publisher that ensure that no advertiser can profit from such behaviour. Unlike Zhu and Wilbur (2011), I make consumer behaviour endogenous and examine the way that consumers respond to explicitly informative ad messages. Moreover, I explore the effect of consumer beliefs on firms' ability to influence their own click through rate. Thus, Zhu and Wilbur (2011) model whether firms do or do not wish to attract clicks, whereas my work is concerned with whether they can or cannot. I show that it is precisely when advertisers wish to inflate their own click-through rate that they are most unable to do so because consumers anticipate this adverse incentive and adjust their response to ad

messages accordingly. This line of reasoning sheds new light on the extent of sustainable advertiser-consumer communication, and also admits a discussion of welfare.

Wilbur and Zhu (2009) consider the problem of click-fraud, that is to say of non-consumers strategically clicking on ads that are priced per-click either to drive rival advertisers out of the market or to inflate publisher revenues. In contrast, I model an increase in an advertiser's unmatched consumer clicks that it either brings upon itself by changing the information content of its advertisement, or else incurs as an endogenous consequence of consumers' inability to extract useful information from the ad. Wilbur and Zhu (2009) do not consider explicitly informative ads, and considerations of consumer-firm communication that are the main focus of the present paper are therefore absent.

Lastly, in a related literature, Ellison and Ellison (2009), Ellison and Wolitzky (2008), and Wilson (2010) discuss firms' general incentive to obfuscate consumer search in order to weaken price competition.

2 MODEL

The idea here is to build a simple model whilst capturing the key stages of on-line product search: (i) visiting a website or submitting a query to a search provider, (ii) finding an advertisement there—typically consisting of a short description—from which an initial assessment of relevance must be formulated, (iii) visiting a relevant link in order to obtain more detailed information, and (iv) (potentially) making a purchase.

There are n firms that produce (at zero cost) a subset of the set of available goods, $\{A, B\}$. In particular, let $\theta_j \in \{\{A\}, \{A, B\}\}$ denote the (privately known) set of goods sold by firm j . Thus, there are single-product firms that stock only a mainstream good, A , and larger multi-product firms that offer, in addition to A , a long-tail or niche good, B .² For brevity, I shall write $\theta_j \in \{A, AB\}$. I assume that a publicly known fraction, γ , of the firms are of type A . In order to sell to the consumers, a firm must occupy an advertising opportunity that is controlled by a monopolist *publisher*. There is one such advertising space, and firms must bid in a second price, sealed-bid auction for the right to become the unique *advertiser* (I later relax the assumption of a single advertising spot). The price reached by the auction determines the advertising fee that the advertiser must pay, with the fee being paid per-impression, per-click or per-sale—depending on the fee structure, $\phi \in \{I, C, S\}$, that is in use.

p670

There are a unit mass of *consumers* identified by their type $t_i \in \{A, B\}$, with these types respectively having frequency δ and $1 - \delta$. I assume that the consumer's search or browsing activity publicly reveals an interest in $\{A, B\}$, but not the specific value of t_i . Each consumer's type represents their ideal product from the set of goods. Writing

²In Section 7, I extend the type space to allow type B single-product firms.

θ for the type of the advertiser, a consumer is perfectly matched with the advertiser when $t_i \in \theta$, and ‘ex-ante unmatched’ if $t_i \notin \theta$. More specifically, this is a search good model in which the consumer learns about his value by visiting firms. Upon arriving at the advertiser and finding $t_i \in \theta$, the consumer learns his value, v_i , which is uniformly distributed on $[0, 1]$, as well as the advertiser’s price, p_j . If i finds that $t_i \notin \theta$ then i and the advertiser are unmatched; however, there is a positive probability $\alpha \in (0, 1)$ that i discovers that A-firms are also able to satisfy him, in which case i again learns his value, v_i , and the firm’s price.³ With probability $(1 - \alpha)$, the consumer finds that an ex-ante unmatched firm offers him zero value. When consumer i visits the advertiser, denote by $\hat{v}_i \in \{0, v_i\}$ his realised value there. When $\hat{v}_i > 0$ I say that i and the advertiser are ex-post matched.

The advertiser must choose a price, p . It can also costlessly transmit an arbitrary, publicly-observable message (advertisement), $m(\theta)$, to the consumers. Visiting the advertiser in order to obtain full product specifications and pricing information is costly—specifically, a cost of $s_i \geq 0$ is incurred by consumer i upon visiting. Denote by $F(\cdot)$ the CDF of s_i , with density $F' = f > 0$ everywhere on its support, $[0, 1/8]$.⁴ The main question addressed in this paper is: to what extent can we expect m to be informative about θ ?

3 PRELIMINARIES

In a perfect Bayesian equilibrium, the consumers use Bayes’ rule to form a posterior belief about θ given m and the equilibrium strategy profile being played. In particular, let $\mu(t|\phi, m)$ be the posterior probability that $t \in \theta$ when the consumer observes ϕ, m . Where no confusion results, I will often write $\mu(m)$ for the set of beliefs. Given $\mu(m)$, the consumer’s optimal strategy is to visit if and only if $E(\hat{v}_i - p | \mu(m), t_i) \geq s_i$, and (conditional on visiting) to purchase when $\hat{v}_i \geq p$.

Nesting the three fee structures together, the advertiser’s profit then has the form

$$(1) \quad \pi^a(\theta, m, p, \phi, \mu(\cdot)) = \lambda^1(\theta, \lambda^2(\cdot))[(1-p)(p - b_S)] - \lambda^2(m, \phi, \mu(\cdot))b_C - b_I,$$

where b_S , b_C , and b_I denote the fee per-sale, per-click and per-impression respectively. I shall refer to $\lambda^1(\cdot)$ as the advertiser’s potential demand. For any given profile of player strategies, this is the mass of ex-post matched consumers. This may not be the same as the number of visitors to the advertiser (clicks)—which I denote by $\lambda^2(\cdot)$. Note that so long as advertisers cannot commit to an advertised price,⁵ the price does not enter

³An alternative interpretation is that α captures the (discounted) value of latent sales that come from exposing a consumer to products that are not of immediate interest but may turn into future sales.

⁴It will turn out that consumers with search costs greater than $1/8$ are never willing to search, and I therefore normalise the population so that such consumers are excluded.

⁵Prices often go unadvertised. Even where prices are advertised, firms have a strong incentive to use high shipping fees and other obfuscation tactics so that the posted price is likely to be a weak indicator of the transfer ultimately made by the consumer. (See Ellison and Ellison, 2009; Ellison and Wolitzky, 2008,

as an argument of $\lambda^2(\cdot)$. The first order condition for profit maximisation implies that the optimal price for the advertiser is the standard monopoly price, $p^* \equiv (1 + b_S)/2$, and consumers will rationally anticipate this as being the price charged. Given this price, the gross surplus accruing to a unit mass of (ex-post) matched consumers can be calculated in the usual fashion by integrating under the demand curve implied by F :

$$(2) \quad GCS = \int_0^{\frac{1-b_S}{2}} \left(1 - q - \frac{1+b_S}{2}\right) dq = \frac{1}{8}(1-b_S)^2.$$

It is immediately apparent that by distorting prices, positive per-sale fees reduce (2).

4 A FULLY-INFORMATIVE EQUILIBRIUM ALWAYS EXISTS UNDER PAY-PER-CLICK

In this section I constrain the publisher to use $\phi = C$ —that is $b_I = b_S = 0$ and $b_C = b^{(n-1)}$, where $b^{(n-1)}$ denotes the second highest bid from the ad auction—and look for equilibria in the resulting game involving the firms and consumers. I begin with the matter of whether or not some information transmission can take place in a PPC regime. As it turns out, full information transmission is always possible. That is to say, the PPC environment can always support a separating equilibrium in which each firm type sends a unique message.

Combining (1) with p^* , and imposing $b_I = b_S = 0$ gives PPC advertiser profits: $\lambda^1(\cdot)/4 - \lambda^2(\cdot)b_C$. Standard second-price auction arguments dictate that it is weakly dominant for firms to bid up to the zero profit level. Thus, the optimal bid is given by

$$(3) \quad b^*(\theta_j, C) = \frac{\lambda^1(\cdot)}{4\lambda^2(\cdot)}.$$

Equilibria in which each firm type sends a distinct signal, $m^*(\theta)$, are essentially equivalent to one another—the most intuitive case is ‘truth-telling’, in which $m^*(\theta) = \theta$. From (2), consumers with $s_i \leq \alpha/8$ always find it optimal to search when $b_S = 0$. Since a truth-telling equilibrium fully reveals a firm’s type, consumers with $s_i > \alpha/8$ click if and only if $t_i \in (m^*)^{-1}(m(\theta))$, where $(m^*)^{-1}(\cdot)$ is the inverse of $m^*(\cdot)$. Thus, when consumers expect honest reporting, sending $m^*(A)$ generates

$$(4) \quad \lambda^2(m^*(A), C, \mu^*(m^*(A))) = \delta + (1 - \delta)F\left(\frac{\alpha}{8}\right)$$

clicks, where $\mu^*(m)$ is the posterior belief that places probability one on the advertiser being of type $(m^*)^{-1}(m)$. A proportion α of the ex-ante unmatched consumers that click realise a positive \hat{v}_i . Thus, single-product advertisers that use m^* have

$$(5) \quad \lambda^1(A, \lambda^2(\cdot)) = \delta + \alpha(1 - \delta)F\left(\frac{\alpha}{8}\right).$$

for a discussion of these and related issues).

Substituting (4) and (5) into (3) yields the optimal single-product firm bid. Note that this bid is decreasing in $F(\alpha/8)$. When $F(\alpha/8)$ is large there are many niche consumers that can search so cheaply that they will visit even when they know the advertiser to be a poor match. The greater is the number of unmatched consumers that will visit a firm, the lower is that firm's conversion rate and hence its value per-click.

Under μ^* , an advertiser sending $m = m^*(AB)$ receives clicks from all consumers so that $\lambda^2(m^*(AB), C, \mu^*(m^*(AB))) = 1$. If the advertiser is indeed multi-product then its λ s must obey the general property $\lambda^1(AB, \lambda^2(\cdot)) = \lambda^2(\cdot)$, and thus $b(AB, C) = 1/4$.

Note that these optimal bids exhibit the desirable property that a multi-product firm wins the auction whenever such a firm exists, which happens with probability $1 - \gamma^n$.

It is immediately clear that AB -type firms have no incentive to change their messaging behaviour. It remains to be demonstrated that single-product advertisers have no incentive to transmit $m = m^*(AB)$. Honest single-product advertisers attract all of their matched consumers, and continue to do so when sending $m = m^*(AB)$, so that the effect of imitating a multi-product firm is to attract an additional mass equal to $(1 - \delta)[1 - F(\alpha/8)]$ of ex-ante *unmatched* consumers. Thus, the resulting net addition to advertiser profit is

$$\Delta\pi^\alpha = (1 - \delta) \left[1 - F\left(\frac{\alpha}{8}\right) \right] \left(\frac{\alpha}{4} - b^{(n-1)} \right).$$

It follows that a necessary condition for the deviation to be profitable is $\alpha/4 > \min_\theta \{b^*(\theta, C)\}$. This inequality is clearly violated for all α by the multi-product firm's bid of $1/4$. A single-product firm's bid is decreasing in $F(\alpha/8)$, so that the inequality is most easily satisfied when $F(\alpha/8) = 1$. Making this substitution reveals that there is no α such that the inequality is not violated. Thus, conditional on having won the auction, no firm wishes to deviate from the informative reporting strategy and, given this, the bidding strategy derived above is optimal. This leads us to the first main result:

Proposition 1 *A fully informative equilibrium can always be sustained under the pay-per-click regime.*

Equilibrium 1 *Under PPC there exists an equilibrium in which firms bid up to (3). The advertiser transmits a fully-informative message, $m^*(\theta)$, and sets $p = p^*$. Consumers update their beliefs to μ^* , visit if (i) $t_i \in (m^*)^{-1}(m^*(\theta))$, or (ii) $s_i \leq \alpha/8$, and purchase (conditional on having visited) if $\hat{v}_i > p^*$.*

Consumer surplus under this equilibrium can be written as

$$CS = (1 - \gamma^n + \delta\gamma^n) \left(\frac{1}{8} - \int_0^{1/8} sf(s) ds \right) + \gamma^n(1 - \delta) \left[F\left(\frac{\alpha}{8}\right) \frac{\alpha}{8} - \int_0^{\alpha/8} sf(s) ds \right],$$

where the first term is the surplus accruing to consumers that find $t_i \in \theta$, and the second to those with $t_i \notin \theta$.

This equilibrium yields three possible outcomes for the publisher, whose profit is given by $E(b^{(n-1)}\lambda^2)$. Firstly, with probability γ^n , there are only single-product firms. Secondly, with probability $n(1-\gamma)\gamma^{n-1}$ there is precisely one multi-product firm—in which case the per-click cost is set at the single-product firms' bid level, but the number of clicks is that for a multi-product firm. Finally, with the complementary probability there are two or more multi-product firms. Thus,

$$(6) \quad \pi^p = \gamma^n b^*(A, C) \left[\delta + (1-\delta)F\left(\frac{\alpha}{8}\right) \right] + \\ [n(1-\gamma)\gamma^{n-1}] b^*(A, C) + [1 - (\gamma^n + n(1-\gamma)\gamma^{n-1})] \frac{1}{4}.$$

All single-product firms bid up to their value and therefore make zero profit. A multi-product firm makes a positive profit if and only if it is the unique such firm, in which case it pays the single-product firm bid for each of its mass 1 of clicks. Thus, the expected profit for a multi-product firm is $\pi = \gamma^{n-1} [(1/4) - b^*(A, C)]$.

Lastly, total social welfare can be calculated as the sum of consumer and publisher surplus from above along with the expected profits of the advertiser. A more direct method is to note that each visiting consumer with $\hat{v}_i > 0$ generates $3/8$ units of expected surplus, and expends s_i units on search. Thus, total welfare is

$$W = (1 - \gamma^n + \delta\gamma^n) \left(\frac{3}{8} - \int_0^{1/8} sf(s) ds \right) + \gamma^n (1 - \delta) \left[F\left(\frac{\alpha}{8}\right) \frac{3\alpha}{8} - \int_0^{\alpha/8} sf(s) ds \right],$$

Both consumer surplus and social surplus are created (in expectation) whenever a consumer realises $\hat{v}_i > 0$, and lost whenever a consumer pays the search cost; the two expressions are therefore similar.

The advertiser's message to the consumers is costless and unverifiable, making this a game of 'cheap talk'. Multiplicity of equilibria is pervasive in such games, but a common feature is the existence of a pooling (or babbling) equilibrium and such an equilibrium does indeed exist in the pay-per-click environment.⁶ There are, however, compelling grounds for selection of the fully informative equilibrium, namely that it is the unique neologism-proof (see Farrell, 1993) PPC equilibrium. The intuition for this result is as follows: Advertisements are written in natural language, which has an intrinsic meaning outside of the scope of the game. In practice, this allows the transmission of arbitrary messages such as "My type is A , you should believe this statement because no AB -type would wish you to do so." Single-product advertisers make an (expected) loss on unmatched visits, and a profit on matched visits in any PPC equilibrium, and thus have a strong incentive to try to separate themselves in this fashion. It turns out that any PPC

⁶In fact, the babbling PPC equilibrium is equivalent to the babbling equilibrium of the PPI environment (presented in Section 5.2) in the sense that it yields the same surplus for all parties, and hence the same overall welfare.

equilibrium with pooling has such a separating message that is credible in the sense that any AB -type would indeed prefer not to send the message.

Remark 1 *The fully informative equilibrium is the unique neologism-proof equilibrium of the PPC environment.*

It is appropriate to briefly discuss the mechanism responsible for the above results. A firm's willingness to pay for a PPC advertisement is determined by the rate at which it converts clicks into sales. Provided that every advertiser attracts at least some matched consumers, all firms will bid strictly more than the value of an unmatched click. Misleading advertisements that attract poorly matched consumers can therefore not be profitable.

5 PAY-PER-IMPRESSION AND PAY-PER-SALE

5.1 Inadmissibility of informative equilibria under PPI and PPS

In the previous section it was shown that a fully informative equilibrium always exists in the PPC environment. In stark contrast to this result, I show in this section that when the publisher is constrained to use either pay-per-impression or pay-per-sale, the unique non-trivial equilibrium is the babbling equilibrium, in which no information is transmitted.⁷

Proposition 2 *In pay-per-impression and pay-per-sale environments with $\theta \in \{A, AB\}$, there exists no non-trivial equilibrium with informative advertising messages.*

Proof. Let $b_C = 0$ and suppose that the proposition is false. This implies that there exist two equilibrium messages, \hat{m} and \tilde{m} such that $\hat{\mu} = \mu(B|\phi, \hat{m}) > \mu(B|\phi, \tilde{m}) = \tilde{\mu}$, and it must therefore be true that

$$\begin{aligned} \lambda^1(A, \lambda^2(\phi, \hat{\mu})) &= \delta + \alpha(1 - \delta)F\left(\hat{\mu}\frac{1}{8}(1 - b_S)^2 + (1 - \hat{\mu})\alpha\frac{1}{8}(1 - b_S)^2\right) \\ &> \delta + \alpha(1 - \delta)F\left(\tilde{\mu}\frac{1}{8}(1 - b_S)^2 + (1 - \tilde{\mu})\alpha\frac{1}{8}(1 - b_S)^2\right) = \lambda^1(A, \lambda^2(\phi, \tilde{\mu})). \end{aligned}$$

p672 If a single-product advertiser switches from sending \tilde{m} to \hat{m} , then the resulting change in its profits is $\Delta\pi^a = [\lambda^1(A, \lambda^2(\phi, \hat{\mu})) - \lambda^1(A, \lambda^2(\phi, \tilde{\mu}))](1 - p)(p - b_S)$, which is positive for any non-trivial equilibrium. This implies that \tilde{m} is never sent by A -type advertisers, which is not consistent with $\hat{\mu} > \tilde{\mu}$ when \tilde{m} is an equilibrium message. It must then be true that—for any two messages, \hat{m} and \tilde{m} , that are on the equilibrium path— $\mu(B|\phi, \hat{m}) = \mu(B|\phi, \tilde{m})$. This amounts to saying that advertisement messages are completely uninformative. ■

⁷As is typical in search models, there will generally exist trivial equilibria in which the advertiser sets a very high price and consumers do not visit regardless of m .

5.2 Babbling under pay-per-impression

Given Proposition 2, any equilibrium admissible in a PPI environment must involve advertiser babbling. I therefore look for an equilibrium of the PPI regime in which all θ transmit $m(\theta) = \tilde{m}$. Advertiser profits at the monopoly price are given by $\pi^a = (\lambda^1(\cdot)/4) - b_I$, which implies an optimal bid of

$$(7) \quad b^*(\theta_j, I) = \frac{\lambda^1(\cdot)}{4}.$$

In a babbling equilibrium, both advertiser types share the same λ^2 , which includes a strictly positive mass of B -type consumers. It follows that AB -type advertisers convert more of their visitors and $b^*(AB, I) > b^*(A, I)$. Consumer i with $t_i = B$ expects, then, to receive

$$U_i = (1 - \gamma^n) \frac{1}{8} + \gamma^n \frac{\alpha}{8} - s_i \equiv s^* - s_i$$

from visiting, and $\lambda^2(I, \mu(\tilde{m})) = \delta + (1 - \delta)F(s^*)$. As always, multi-product advertisers have $\lambda^1 = \lambda^2$. Single product advertisers have $\lambda^1(A, \lambda^2(\cdot)) = \delta + \alpha(1 - \delta)F(s^*)$.

Equilibrium 2 *When the publisher is constrained to use a PPI fee structure there exists a babbling equilibrium in which firms bid up to $\lambda^1(\cdot)/4$. The successful firm transmits a meaningless message, \tilde{m} , and sets $p = p^*$. Consumers update their beliefs, visit if (i) $t_i = A$, or (ii) $s_i \leq s^*$, and purchase (conditional on having visited) if $\hat{v}_i > p^*$.*

Note that, even when the advertising message is completely uninformative, there is belief updating in equilibrium. In particular, $\mu(B|I, \tilde{m}) = 1 - \gamma^n$, which is greater than the corresponding prior probability $(1 - \gamma)$ because multi-product firms always outbid their single-product rivals. This result—that advertisements have an intrinsic informational content by virtue of certain types' greater willingness to pay for them—exists in a similar spirit to the intuition of Nelson (1974) and Chen and He (2006).

Consumer surplus under this equilibrium can be written as

$$CS = \delta \left(\frac{1}{8} - \int_0^{1/8} sf(s) ds \right) + (1 - \delta) \left(F(s^*)s^* - \int_0^{s^*} sf(s) ds \right).$$

Publisher profit under PPI is simply $b^{(n-1)}$ which, in expectation, is calculated as

$$\pi^p = [1 - \gamma^n - n(1 - \gamma)\gamma^{n-1}] \frac{\delta + (1 - \delta)F(s^*)}{4} + [\gamma^n + n(1 - \gamma)\gamma^{n-1}] \frac{\delta + \alpha(1 - \delta)F(s^*)}{4},$$

where $\gamma^n + n(1 - \gamma)\gamma^{n-1}$ is the probability of there being one or fewer multi-product firms. As in Equilibrium 1, single product firms here bid up to their value for the advertisement slot, and make zero profits. A multi-product advertiser makes positive profits if and only

if it is the unique such firm, so that profits for such firms are given by

$$(8) \quad \pi = \gamma^{n-1} \left[\frac{\delta + (1-\delta)F(s^*)}{4} - \frac{\delta + \alpha(1-\delta)F(s^*)}{4} \right] = \gamma^{n-1} \left[\frac{(1-\alpha)(1-\delta)F(s^*)}{4} \right].$$

Lastly, as before, social surplus is generated and dissipated under the same circumstances as consumer surplus so that one can obtain an expression for total welfare by substituting the welfare generated per unit mass of matched consumers (3/8) in place of gross consumer surplus in the CS equation:

$$W = \delta \left(\frac{3}{8} - \int_0^{1/8} sf(s)ds \right) + (1-\delta) \left(3F(s^*)s^* - \int_0^{s^*} sf(s)ds \right).$$

5.3 *Equilibrium under pay-per-sale*

Substituting p^* into (1), and imposing $b_C = b_I = 0$ reveals that advertiser PPS profits are positive for all $b_S < 1$ —so that the optimal bid for all firms is 1. Such behaviour engenders a trivial equilibrium in which $p^* = 1$, consumers do not search, and consumer surplus, publisher profit, and firm profit are all equal to zero. More fundamentally, single- and multi-product firms each share a common value per-sale (namely, $p^* - b_S$), and it is therefore impossible for the publisher to ex ante discriminate between the two on the basis of a per-sale fee. The result is likely to be sub-optimal choice of advertiser in the *short-run*. In the *long-run*, the publisher may be able to infer firm types by monitoring metrics such as the firm's conversion rate, and thereby select *AB*-type advertisers systematically.

Suppose that, instead of using an auction to allocate its advertisement, the publisher sets a take it or leave it price, b_S , for the slot and selects a firm to fill it. Given that advertisers must transmit meaningless messages in a PPS equilibrium, a *B*-type consumer expects $U_i = (1 - b_S)^2 s^{**} - s_i$ from visiting, where $s^{**} = s^*$ in the long-run, and $s^{**} = (1 - \gamma)/8 + \gamma\alpha/8 (< s^*)$ in the short-run. It follows that the number of visitors attracted to the advertiser is $\lambda^2(\cdot) = \delta F((1 - b_S)^2/8) + (1 - \delta)F((1 - b_S)^2 s^{**})$, and that $\lambda^1(A, \lambda^2(\cdot)) = \delta F((1 - b_S)^2/8) + \alpha(1 - \delta)F((1 - b_S)^2 s^{**})$, whilst $\lambda^1(AB, \lambda^2(\cdot)) = \lambda^2(\cdot)$ as usual.

Using p^* , publisher profit is then given by

$$\pi^P = E(\lambda^1(\cdot))(1 - b_S)b_S/2.$$

This function is bounded in b_S , and is positive for all $0 < b_S < 1$ so that there exists at least one b_S yielding (positive) maximised publisher profits. Let \mathcal{B} denote the set of all such b_S .⁸ For example, when F is given by the uniform distribution on $[0, 1/8]$, $\mathcal{B} = \{1/4\}$ and the publisher induces $p^* = 5/8$. Given \mathcal{B} , it is possible to characterise equilibrium in

⁸If the publisher cannot offer consumers a credible commitment to a choice of b_S —say, because consumers do not directly observe b_S before making their visit decision—then the best that it can do is to choose the b_S that maximises $(1 - b_S)b_S/2$, so that $\mathcal{B} = \{1/2\}$. Consumers will then rationally expect $p = 3/4$.

the pay-per-sale game.

Equilibrium 3 *When the publisher uses a PPS fee structure, any non-trivial equilibrium of the PPS game has $b_S \in \mathcal{B}$. The chosen advertiser transmits a meaningless message, \tilde{m} , and sets $p = p^* = (1 + b_S)/2$. Consumers visit if (i) $t_i = A$ and $s_i \leq (1 - b_S)^2/8$, or (ii) $s_i \leq (1 - b_S)^2 s^{**}$, and purchase (conditional on having visited) if $\hat{v}_i > p^*$.*

Consumer surplus and social welfare under Equilibrium 3 are equal to

p673

$$(9) \quad CS = \delta \left[F \left(\frac{1}{8} (1 - b_S)^2 \right) \frac{1}{8} (1 - b_S)^2 - \int_0^{(1-b_S)^2/8} sf(s) ds \right] + \\ (1 - \delta) \left[F \left((1 - b_S)^2 s^{**} \right) (1 - b_S)^2 s^{**} - \int_0^{(1-b_S)^2 s^{**}} sf(s) ds \right],$$

$$W = \delta \left[F \left(\frac{1}{8} (1 - b_S)^2 \right) \frac{1}{8} (3 - 2b_S - (b_S)^2) - \int_0^{(1-b_S)^2/8} sf(s) ds \right] + \\ (1 - \delta) \left[F \left((1 - b_S)^2 s^{**} \right) (3 - 2b_S - (b_S)^2) s^{**} - \int_0^{(1-b_S)^2 s^{**}} sf(s) ds \right].$$

The latter is simply the consumer surplus expression with the gross consumer surplus replaced by the expected social value of a match, which is $[3 - 2b_S - (b_S)^2]/8$.

Firm profits under Equilibrium 3 are $\pi = \Pr(\text{chosen as advertiser}) \cdot \lambda^1(\theta, \lambda^2(\cdot))(1 - b_S)^2/4$. In the short-run, each firm is chosen with equal probability so that $\Pr(\text{chosen as advertiser}) = 1/n$. In the long-run, the publisher systematically selects multi-product firms, leaving such firms with $\Pr(\text{chosen as advertiser}) = (1 - \gamma^n)/n(1 - \gamma)$, whilst single-product firms have $\Pr(\text{chosen as advertiser}) = \gamma^n/n$. Note that, in contrast to PPC and PPI fee structures, single product firms make positive profits.

6 WELFARE

As n becomes large, or γ small, an AB -type advertiser is realised in Equilibria 1 and 2 with probability close to 1. Moreover, as α or δ approach 1, all firms are functionally similar to AB -types—satisfying all visitors. These four cases are equivalent from a welfare perspective, and the limiting surpluses are identical for the two equilibria in question since both induce the same product market price. The resulting surplus values are presented in Table 1 for consumers (CS), the publisher (π^P), firms (π), and society overall (W).

Remark 2 *As $n \rightarrow \infty$, $\alpha \rightarrow 1$, $\delta \rightarrow 1$, or $\gamma \rightarrow 0$, Equilibria 1 (PPC) and 2 (PPI) yield equal expected payoffs for all.*

Table 1: Expected payoffs from Equilibria 1 (PPC) and 2 (PPI) with large n , large α , large δ , or small γ .

CS	π^P	π	W
$1/8 - E(s_i)$	$1/4$	0	$3/8 - E(s_i)$

It is not instructive to repeat this exercise for the PPS case, since the welfare values then depend upon the chosen b_S . One can, however, consider limiting welfare values under idealised circumstances (i.e. the highest value for each group that could be obtained if b_S and F were exogenously set to maximise that group's surplus):

Proposition 3 *As $n \rightarrow \infty$, $\alpha \rightarrow 1$, $\delta \rightarrow 1$, or $\gamma \rightarrow 0$, Equilibrium 3 (PPS) is worse for publishers, consumers, and society more generally than either Equilibrium 1 (PPC) or Equilibrium 2 (PPI). For large n , PPS is no better for firms than either alternative fee structure.*

Proof. Under long-run PPS, consumer surplus and social welfare are maximised when $b_S = 0$ —in which case they are identical to those in Equilibrium 2 for any F . Since $b_S > 0$ will hold in any non-trivial equilibrium (and consumers surplus and social welfare are lower under short-run PPS) consumers and society more generally must be worse off under PPS than either PPC or PPI when at least one of n, α, δ , or $(1 - \gamma)$ is large. The publisher can make at most $1/8$ under PPS (by setting $b_S = 1/2$, and having $E(\lambda^1(\cdot)) \approx 1$), which is strictly less than the value of $1/4$ in Table 1. For large n firm expected profits under PPS are zero. ■

More generally, consumer surplus from Equilibria 1, 2 and 3 can, for any parameter configuration, be ranked in that order. For a given consumer-firm pair, Equilibria 1 and 2 yield the same surplus conditional on the consumer visiting (since the price is the same in both cases). However, the PPI equilibrium induces excessive visiting in the case that the advertiser is type A , and too few visits in the case that the advertiser is of type AB . By contrast, consumers in an informative equilibrium visit precisely when it is privately optimal to do so—namely when $s_i \leq E(\hat{v}_i | \mu^*(m^*(\theta)), t_i) = E(\hat{v}_i | \theta, t_i)$.

Equation 9 is no greater than the surplus under Equilibrium 2 or, by the above discussion, Equilibrium 1 when $b_S = 0$. Now, increasing b_S affects consumer surplus through two channels: (i) the surplus derived by inframarginal consumers from each unit consumed decreases, (ii) marginal consumers (those with s_i closest to $(1 - b_S)^2 s^{**}$) are induced to abstain from searching. The first channel has a non-positive effect on consumer surplus. The second channel must also have a non-positive effect on consumer surplus since participation by marginal consumers under the original b_S implies that their expected surplus from visiting was non-negative. It follows that Equilibrium 3 provides lower consumer surplus than does Equilibrium 1 or 2—strictly so in any non-trivial case.

Remark 3 *Consumer surplus under Equilibrium 1 (PPC) is no lower than that under Equilibrium 2 (PPI) which, in turn, is at least as great as in Equilibrium 3 (PPS).*

An analogous argument can be extended to establish that social welfare is (weakly) lower under Equilibrium 3 than Equilibrium 2. In particular, note that this is clearly true for $b_S = 0$, and that participation by marginal consumers implies a loss of surplus when such consumers are excluded by an increase in b_S .

Remark 4 *Welfare is lower under Equilibrium 3 (PPS) than under Equilibrium 2 (PPI).*

Remark 4 does not extend to a ranking of Equilibrium 1. In order to extract profit from the advertising resource that it controls, the publisher must endow the advertiser with the power to price above marginal cost. The result is a wedge driven between consumer surplus and social welfare. Thus, whilst consumer visit decisions are privately optimal under an informative equilibrium, they may not be socially optimal; in particular, consumers will tend to visit too infrequently—not internalising the profits that their visits generate for the advertiser and publisher. When the distribution of search costs is such that a reduction in information transmission induces additional visits from those consumers for whom s_i lies between the private and social benefit of visiting, an increase in overall surplus can result. As an example, social welfare under Equilibrium 2 is greater than that under Equilibrium 1 for $\alpha = 3/4$, $\delta = \gamma = 1/2$, $n = 2$, and F given by the inverse quadratic distribution, $F(s) = 192s^2 - 1024s^3$. Long-run social welfare is identical for $b_S = 0$ under Equilibria 2 and 3, so that the same argument holds for a comparison of pay-per-click and pay-per-sale when b_S is low. Since welfare is decreasing in b_S , and since welfare is zero when $b_S = 1$, there exists a b_S above which pay-per-click offers higher welfare than does pay-per-sale.

Remark 5 *Improved information transmission can be harmful to social welfare.*

The publisher faces a trade-off between wishing to increase the number of clicks (by reducing information transmission in the event that there is a single-product advertiser) on the one hand, and wishing to increase consumer confidence in ad messages and the value of each click on the other. These conflicting incentives give rise to a degree of ambiguity in the optimal choice of fee structure. In the case that F is uniform, consumer confidence considerations always dominate and we have the following result, whose proof may be found in the appendix:

Proposition 4 *For F uniform on $[0, 1/8]$, Equilibrium 1 (PPC) always yields higher publisher profits than does either Equilibrium 2 (PPI) or Equilibrium 3 (PPS).*

p674

For more general F the publisher may wish to restrict information transmission. For example, Equilibrium 2 yields higher publisher profit than does Equilibrium 1 if $\delta = 1/2$,

$\alpha = \gamma = 4/5$, $n = 2$, and F is given by the inverse quadratic distribution, $F(s) = 192s^2 - 1024s^3$. Some intuition can be gained by noting that publisher profit in Equilibrium 2 is $E(b^{(n-1)}) = E(\underline{\lambda}^1)/4$ (where $\underline{\lambda}^k$ denotes the λ^k that would have been induced had the second highest bidder won the auction), whilst that in Equilibrium 1 is $E(b^{(n-1)}\lambda^2) = E(\underline{\lambda}^1\lambda^2/4\underline{\lambda}^2)$ (both equilibria induce the same distribution of advertiser types). Since $\lambda^2 \geq \underline{\lambda}^2$, a sufficient condition for Equilibrium 1 to be more profitable is that it generates a larger $E(\underline{\lambda}^1)$. A move to informative advertisements reduces the number of B -type visits vis-à-vis Equilibrium 2 when the advertiser is of type A , and increases the number of such visits when the advertiser's type is AB . Thus, PPC yields higher expected profits for the publisher when it is relatively likely that there are at least two AB -type firms (when γ is not too large and n is not too small). We should, then, expect PPC to be favoured in 'thick' markets with a relatively large supply of potential advertisers.

A per-sale fee distorts p^* , which both undermines the publisher's ability to extract advertiser profit and reduces consumers' propensity to click—thus creating a double hurdle for PPS to be more profitable than either PPC or PPI. Proposition 5 establishes an upper bound on the relative profitability of PPS when PPS profits attain their theoretical maximum and F is allowed to be general.

Proposition 5 *If niche (B-type) consumers are a minority (if $\delta \geq 1/2$) then Equilibria 1 (PPC) and 2 (PPI) always yield higher publisher profits than does Equilibrium 3 (PPS).*

Proof. Since $(1 - b_S)b_S/2$ can not be greater than $1/8$ and $E(\lambda^1(\cdot)) < 1$, publisher profit under PPS must be strictly less than $1/8$. Profits from Equilibrium 1 are minimised when $\gamma = 1$, and are therefore sure to be greater than $1/8$ when

$$\frac{1}{4} \left[\delta + \alpha(1 - \delta)F\left(\frac{\alpha}{8}\right) \right] \geq \frac{1}{8} \Leftrightarrow \delta \geq \frac{1 - 2\alpha F\left(\frac{\alpha}{8}\right)}{2 - 2\alpha F\left(\frac{\alpha}{8}\right)}.$$

PPI is always more profitable than PPS if the number of AB -type firms is not equal to 1 since it then induces a higher click-through rate and extracts all of the firms profits (which are larger than those from PPS since PPI does not distort p^*). When there is precisely one AB -type firm PPI profits are greater than $1/8$ if

$$\frac{1}{4} \left[\delta + \alpha(1 - \delta)F(s^*) \right] \geq \frac{1}{8} \Leftrightarrow \delta \geq \frac{1 - 2\alpha F(s^*)}{2 - 2\alpha F(s^*)}.$$

■

Proposition 5 compares the theoretical upper bound for PPS profits with worst-case scenario PPI and PPC profits. In practice, expected PPS profits will typically be substantially below their theoretical upper bound and publishers will not benefit from implementing a PPS regime except under the most extremal parameter configurations—preferring instead to choose the better of PPC or PPI. This may explain why some early experimenters such as Google have subsequently abandoned their per-sale pricing scheme.

7 EXTENSIONS

7.1 *Symmetric Type Space*

I now extend the type space for firms so that $\theta_j \in \{\{A\}, \{B\}, \{A, B\}\}$. With no loss in generality, let $\delta \geq 1/2$. I assume that the fraction of firms that are type A , B , and AB is $\gamma/2$, $\gamma/2$ and $(1 - \gamma)$ respectively. It is fairly straightforward to verify that Proposition 1 continues to hold under such circumstances. In particular, (3) continues to determine firm bids. Since $\lambda^1/\lambda^2 \geq \alpha$, it follows that $\alpha/4 \leq \min_{\theta} \{b^*(\theta, C)\}$. Thus, no firm can profitably deviate from a fully informative equilibrium.

Remark 6 *A fully informative equilibrium can always be sustained under the PPC regime in the symmetric type space model.*

What of PPI and PPS? Any signal that is sent only by AB -types cannot be supported in equilibrium since both single-product firm types would prefer to deviate: transmitting $m(AB)$ and receiving a superset of the visitors obtained in the putative equilibrium. Thus, a fully-informative equilibrium cannot be sustained under PPI or PPS.

Remark 7 *A fully informative equilibrium can never be sustained under the PPI or PPS regimes in the symmetric type space model.*

If AB - and A -type advertisers pool together then it is no longer clear that B -type advertisers wish to imitate such firms as doing so may deter high-cost B -type consumers from visiting. Thus, it seems that some information transmission might be possible, even when ads are priced per-impression or per-sale. Under certain conditions, this indeed turns out to be the case. Lemmas 1 and 2 formalise this intuition.⁹

Lemma 1 *When $\theta_j \in \{\{A\}, \{B\}, \{A, B\}\}$ and $\delta \geq 1/2$, there exists a partial pooling equilibrium with $m(A) = m(AB) \neq m(B)$ under a PPI fee regime if and only if*

$$(10) \quad \frac{1 - F(s^{PPIpp})}{1 - F(\alpha/8)} \geq \alpha \frac{\delta}{1 - \delta}, \text{ where } s^{PPIpp} \equiv \left(\frac{1 - \gamma^n + \alpha(1 - \frac{1}{2^n})\gamma^n}{1 - \gamma^n + (1 - \frac{1}{2^n})\gamma^n} \right) \frac{1}{8}.$$

Lemma 2 *When $\theta_j \in \{\{A\}, \{B\}, \{A, B\}\}$ and $\delta \geq 1/2$, there exists a partial pooling equilibrium with $m(A) = m(AB) \neq m(B)$ under a PPS fee regime if and only if*

$$(11) \quad \frac{F(\frac{1}{8}(1 - b_S)^2) - F((1 - b_S)^2 s^{PPSpp})}{F(\frac{1}{8}(1 - b_S)^2) - F(\frac{\alpha}{8}(1 - b_S)^2)} \geq \alpha \frac{\delta}{1 - \delta},$$

where b_S is a publisher-optimal per-sale price, $s^{PPSpp} = s^{PPIpp}$ in the long-run and

$$(12) \quad s^{PPSpp} = \left(\frac{1 - \gamma}{1 - \gamma + \frac{\gamma}{2}} \right) \frac{1}{8} + \left(\frac{\frac{\gamma}{2}}{1 - \gamma + \frac{\gamma}{2}} \right) \frac{\alpha}{8}$$

⁹Analogous equilibria can be constructed with $m(B) = m(AB) \neq m(A)$ (or, equivalently, $\delta < 1/2$).

in the short-run.

Proofs for Lemmas 1 and 2 are contained within the Appendix. A higher γ is more conducive to the existence of partial pooling equilibria: as γ becomes low, consumers that observe $m(A \cup AB)$ assign a high-likelihood to the advertiser being multi-product and sending such a signal becomes tempting for B -types. Likewise, a relatively equal market share is necessary for the admission of a partial pooling equilibrium: B -types have little to gain from honesty when their ‘turf’ constitutes a relatively small fraction of the market.

Type B advertisers do not wish to pool with multi-product firms when the positive signal that this sends to A -types is offset by a negative signal to B -types. Building on this intuition, Lemmas 1 and 2 serve to emphasize the importance of consumers’ interpretation of advertising messages in enabling advertiser-consumer communication.

7.2 Multiple advertising opportunities

p675

I now consider the case in which there are $l > 1$ advertising opportunities (or slots) to be allocated, retaining the assumption that $\theta_j \in \{\{A\}, \{B\}, \{A, B\}\}$. For convenience, identify firms $j \in \{1, \dots, l\}$ with the slot that they occupy in equilibrium.

An equilibrium concept is required for the pricing of ads when $l > 1$. A good choice is to require that the allocation of slots be a Walrasian allocation and the price for each slot a Walrasian price—that is, given the vector of prices, no firm should wish to change its allocated slot. Varian (2007) shows that the set of such core outcomes corresponds to the set of Nash equilibria of the full-information generalised second price auction when consumers consider ads in a fixed order. Walrasian pricing also intuitively accommodates non-ordered ads: in such case it makes sense to think of ads being allocated by a uniform price auction, whose equilibria are also Walrasian. More generally, the Walrasian requirement represents a minimum condition for the stability of the resulting allocation when ads are allocated by some non-auction mechanism such as prices posted by the publisher.

Any equilibrium consumer search rule must have consumers click in an order that maximises expected utility, which amounts to a requirement that consumers click ads with a higher expected match probability first. Of interest is the manner in which consumers break ties between advertisers offering the same expected utility. One possibility is that consumers choose between such advertisers at random. When advertisers are expected to be ordered randomly, a *random search* rule makes all advertising slots equivalent from the advertiser’s perspective and can therefore be thought of as modelling the case in which there is no obvious ordering or list structure to the set of advertisements. In contexts such as search engine advertising ads have an obvious prominence structure. One way to capture *ordered search* is to suppose that consumers break ties

by clicking more prominent ads first. For example, when consumers expect advertisers to transmit informative messages ordered search has consumers first read the entire list from top to bottom in search of a matched advertiser, visiting the first such firm that is encountered. If the bottom of the list is reached without encountering a firm that appears to be matched, the consumer (who now knows the structure of the entire list) reverts to clicking unmatched-links either randomly or in a top-down fashion. A third alternative is *hierarchical search*: having the consumer consider the most prominent ad first before subsequently choosing amongst other ads at random. This is the kind of behaviour one might expect, for example, if all consumers view the ad on a web site's homepage first, but then surf idiosyncratic paths through the site's sub-pages, viewing different ads en route.

Random, ordered, and hierarchical search are all examples of what I will refer to as non-discriminatory (ND) search rules. Formally, given messages (m_j, \mathbf{m}_{-j}) , and the types of j 's advertiser rivals, θ_{-j} , the consumers' search rule induces a probability of a type t consumer clicking on j 's ad,¹⁰ which I denote by $\psi_j(m_j, \mathbf{m}_{-j}, \theta_{-j}; t)$. Write $\mu_j(t|m_j)$ for the consumers' posterior belief that $t \in \theta_j$ when the advertiser in slot j sends m_j .¹¹

Definition 1 A search rule is said to be non-discriminatory (ND) if $\psi_j(m_j, \mathbf{m}_{-j}, \theta_{-j}; t) = \psi_j(m'_j, \mathbf{m}_{-j}, \theta_{-j}; t) \forall \mathbf{m}_{-j}, \theta_{-j}$ whenever $\mu_j(t|m_j) = \mu_j(t|m'_j) = 1$.

In words, an ND search rule demands that a consumer who is certain that a firm will offer a match does not discriminate against that firm on the basis of the precise message it sends. In an informative ND equilibrium consumers are not attracted by claims to sell or not sell irrelevant items. Since consumers can always click at random, an ND rule always exists.

The consumers' search rule can influence firms' incentives for honest messaging. For example, if $l \geq 3$ and type t consumers break ties with the preference order over messages $m(t) > m(AB)$, selecting amongst firms transmitting identical messages at random then it is easy to find parameters such that all AB -type advertisers find it optimal to deviate to a single-product firm's message in any putative fully-informative equilibrium—even under PPC pricing. However, when one restricts attention to ND search behaviours, the incentive properties observed in the single-slot model are preserved. The first result of this section is that we can always find an informative equilibrium under PPC by hav-

¹⁰This probability accounts for both the consumer's (potentially random) click order and the possibility that the consumer is satisfied by some firm clicked before j . When a consumer visits an ex ante unmatched firm he will update his beliefs about the utility of future clicks. Eventually, the consumer may wish to drop out altogether, and ψ also accounts for the probability of his doing so before reaching j .

¹¹In some instances, when informative messages are expected, the consumer's out of equilibrium posterior may not be well defined (e.g, if consumers expect AB -type firms to be in more prominent slots but observe a firm transmitting $m(AB)$ in a less prominent slot than one known to contain an A -type). In such cases I assume that beliefs are updated based on the equilibrium messaging strategies alone as this is the rule that admits the greatest scope for breakdowns in information transmission. This assumption is made for concreteness and is not necessary to establish the existence of informative equilibria under PPC.

ing consumers use an ND search rule. The proofs for this section can be found in the appendix.

Proposition 6 *There exist equilibria of the PPC environment in which all links clicked with positive probability are posted by advertisers transmitting fully-informative messages. Any consumer-optimal non-discriminatory search rule can support such an equilibrium.*

In the case of PPI/PPS it is possible to construct informative equilibria for some search rules and parameter configurations when $l > 1$. For example, suppose that $l = 2$, $\delta = 1/2$, that consumers expect truthful reporting, and that consumers of type t use a search rule that breaks ties with the preference order over messages $m(t) > m(AB)$, selecting amongst firms with identical messages at random. Regardless of the configuration of firm types that occupy the two slots, no firm then has an incentive deviate from the informative messaging strategy under PPI or PPS for any ad price that leaves advertisers with non-negative profits. Note that $m(t) > m(AB)$ violates Definition 1, and is therefore not ND. In fact, ND search rules can never support a fully-informative equilibrium under PPI or PPS.

Proposition 7 *If consumers use a non-discriminatory search rule then there is no fully-informative equilibrium of either the PPI or PPS environment.*

8 CONCLUSION

Charging advertisers on a per-click basis provides a disincentive to attracting poorly matched consumers, and can therefore encourage the transmission of informative advertisements. Conversely, advertisers who pay for advertisements on a per-impression basis, or pay only in the event of a sale find attracting visits from poorly matched consumers to be profitable provided that those consumers will buy with some positive probability—however small that probability is. Such fee structures are therefore much less conducive to the equilibrium transmission of information. A key message of this work is that the informative capacity of advertisements depends as much upon consumers' interpretation of (and trust in) ads as on the ad messages themselves. Crucially, it is precisely when advertisers wish to inflate their own click-through rate that they are most unable to do so because consumers anticipate this adverse incentive and adjust their response to ad messages accordingly.

Consumers benefit from improved transparency but the effect of improved information transmission on publisher profits and social welfare is ambiguous. When there are many advertisers competing for the advertising resource, the auction mechanism does a good enough job of selecting advertiser to mitigate the effects of reduced information transmission. Under such circumstances, pay-per-impression and pay-per-click are

welfare equivalent, whilst pay-per-sale delivers consistently lower welfare owing to the distortionary effects of a marginal tax on firm sales.

APPENDIX A OMITTED PROOFS

Proof of Proposition 4. (i) Setting $F[x] = 8x$, and differentiating the difference between publisher profits under Equilibrium 1 and Equilibrium 2 yields

$$\frac{\partial(\pi^P|_{\phi=C} - \pi^P|_{\phi=I})}{\partial\alpha} = \frac{(1-\alpha)\gamma^n(1-\delta)}{4} \left(-2 - \underbrace{\frac{-2\gamma^{1+n} - n(1-\gamma)\left(2\gamma^n - \frac{\delta(1+\alpha+(1-\alpha)\delta)}{(\alpha+\delta-\alpha\delta)^2}\right)}{\gamma}}_{\text{term A}} \right),$$

It is easily verified (by differentiating) that term A is minimised by $\gamma = 1$. Making this substitution reveals term A to be bigger than -2 , so that the difference in profits decreases in α . From Table 1 this difference is zero for $\alpha = 1$; it is therefore positive $\forall \alpha < 1$.

(ii) Setting $F[x] = 8x$, noting that $\mathcal{B} = \{1/4\}$ for uniform search costs, and simplifying:

$$\pi^P|_{\phi=C} - \pi^P|_{\phi=S} = \frac{1}{512} \left[101 - (1-\alpha)\gamma^n [2(37+64\alpha) + 27(1-\alpha)\gamma^n] (1-\delta) - \frac{128n(1-\alpha)\alpha(1-\gamma)\gamma^{-1+n}(1-\delta)}{\alpha+\delta-\alpha\delta} \right].$$

Both the second and third terms within the large square brackets are decreasing in δ (in absolute terms) so that $\pi^P|_{\phi=C} - \pi^P|_{\phi=S}$ is minimised by $\delta = 0$ for any (α, γ, n) . Making the $\delta = 0$ substitution reduces the problem to one of showing that

$$(13) \quad 101 \geq (1-\alpha)\gamma^{n-1} (128n(1-\gamma) + \gamma(2(37+64\alpha) + 27(1-\alpha)\gamma^n)).$$

Taking a first order condition reveals that the value of α that maximises the right hand side of (13) is $\alpha^* = 1 - [64n(1-\gamma) + 101\gamma] / [\gamma(128 - 27\gamma^n)]$ (the second order condition reveals that the right hand side of (13) is concave in α). Now, note that $\partial\alpha^*/\partial\gamma = n [8192 - 27(64 + 64n(1-\gamma) + 101\gamma)\gamma^n] / [\gamma^2(128 - 27\gamma^n)^2]$, which has the same sign as $(8192 - 27(64 + 64n(1-\gamma) + 101\gamma)\gamma^n)$. Moreover, $\partial(8192 - 27(64 + 64n(1-\gamma) + 101\gamma)\gamma^n)/\partial\gamma = -27(1+n)(64n(1-\gamma) + 101\gamma)\gamma^{n-1} \leq 0$. This implies that $\partial\alpha^*/\partial\gamma$ can be negative for some $\gamma \in [0, 1)$ only if it is negative at $\gamma = 1$. However, $\partial\alpha^*/\partial\gamma|_{\gamma=1} = 37n/101 > 0 \Rightarrow \partial\alpha^*/\partial\gamma > 0$. Thus, $\max_{\gamma,n} \alpha^* = \alpha^*|_{\gamma=1} = 0$. Concavity then implies that the right hand side of (13) is maximised by $\alpha = 0$.

$$(14) \quad \lim_{\alpha \rightarrow 0} \lim_{\delta \rightarrow 0} (\pi^P|_{\phi=C} - \pi^P|_{\phi=S}) = \frac{1}{512} [101 - \gamma^{n-1} (128n(1-\gamma) + \gamma(74 + 27\gamma^n))].$$

The problem is reduced to one of showing that the square bracket term is non-negative.

$$(15) \quad \frac{\partial}{\partial \gamma} \gamma^{n-1} (128n(1-\gamma) + \gamma(74 + 27\gamma^n)) = 2n\gamma^{n-2} \underbrace{(64n(1-\gamma) - 64 + \gamma(37 + 27\gamma^n))}_{\xi}$$

$$\frac{\partial \xi}{\partial \gamma} = 37 - 64n + 27(1+n)\gamma^n, \quad \frac{\partial \xi}{\partial \gamma \partial n} = 27\gamma^n(1 + (1+n)\ln[\gamma]) - 64 < 0.$$

Thus, $\partial \xi / \partial \gamma$ is maximised by $n = 2$, in which case it is equal to $81\gamma^2 - 91 < 0$. This implies that ξ is always decreasing in γ . Substituting $\gamma = 1$ into ξ results in $\xi = 0$ so that (15) is non-negative; (14) is, then, minimised by $\gamma = 1$. Substituting this into (14) reveals that $\pi^p|_{\phi=C} - \pi^p|_{\phi=S} \geq 0$. ■

p677

Proof of Lemma 1. Consider the messages $m(A) = m(AB) \equiv m(A \cup AB) \neq m(B)$. The signal $m(B)$ is, then, fully informative and induces

$$\lambda^2(m(B), \cdot) = \delta F\left(\frac{\alpha}{8}\right) + (1-\delta); \quad \lambda^1(B, \lambda^2(m(B), \cdot)) = \alpha \delta F\left(\frac{\alpha}{8}\right) + (1-\delta).$$

Now, the equilibrium bids induce a probability, $z(\theta, \gamma, n)$, of a type θ firm being the winner of the auction. After observing $m(A \cup AB)$, the consumers use Bayes' rule to update their beliefs. A given B -type consumer then expects utility of

$$\left(\frac{z(AB, \gamma, n)}{z(AB, \gamma, n) + z(A, \gamma, n)} \right) \frac{1}{8} + \left(\frac{z(A, \gamma, n)}{z(AB, \gamma, n) + z(A, \gamma, n)} \right) \frac{\alpha}{8} - s_i \equiv s^{PPIpp} - s_i$$

from visiting, whilst A -types expect $1/8 - s_i$. This implies that $\lambda^2(m(A \cup AB), \cdot) = \lambda^1(AB, \lambda^2(\cdot)) = \delta + (1-\delta)F(s^{PPIpp})$, and $\lambda^1(A, \lambda^2(\cdot)) = \delta + \alpha(1-\delta)F(s^{PPIpp})$. Firm bids are given by (7). Thus, for $\delta \geq 1/2$, $b(AB, I) > b(A, I) > b(B, I)$. It immediately follows that $z(AB, \gamma, n) = 1 - \gamma^n$, and $z(A, \gamma, n) = (1 - 2^{-n})\gamma^n$. Firm profits are increasing in λ^1 and neither AB - nor A -type firms wish to deviate to $m(B)$ so long as doing so reduces their λ^1 , which is true when

$$\min \left\{ \frac{\delta}{1-\delta}, \frac{\delta}{\alpha(1-\delta)} \right\} \geq \frac{1 - F(s^{PPIpp})}{1 - F(\alpha/8)}$$

respectively. Satisfaction of these conditions is guaranteed since $\delta \geq 1/2$. The analogous condition for a type B firm not wishing to transmit $m(A \cup AB)$ is (10). Thus, when (10) holds, conditional on having won the auction, no firm wishes to deviate from the equilibrium reporting strategy and the bids implied by (7) are also optimal. ■

Proof of Lemma 2. Consider the messages $m(A) = m(AB) \equiv m(A \cup AB) \neq m(B)$. The signal $m(B)$ is, then, fully informative and induces

$$\lambda^2(m(B), \cdot) = \delta F\left(\frac{\alpha}{8}(1 - b_S)^2\right) + (1-\delta)F\left(\frac{1}{8}(1 - b_S)^2\right),$$

$$\lambda^1(B, \lambda^2(m(B), \cdot)) = \alpha \delta F\left(\frac{\alpha}{8}(1-b_S)^2\right) + (1-\delta)F\left(\frac{1}{8}(1-b_S)^2\right).$$

Consumers update their beliefs using Bayes' rule. When $\delta \geq 1/2$ and $m(A) = m(AB) \neq m(B)$, the publisher's preference ordering over advertisers is $AB > A > B$. Thus, the expected surplus from visiting for a B -type consumer who observes $m(A \cup AB)$ can be written as $U_i = (1-b_S)^2 s^{PPSpp} - s_i$, where s^{PPSpp} is equal to s^{PPIpp} in the long-run, and is given by (12) in the short-run. The implied consumer behaviour is

$$\begin{aligned} \lambda^1(AB, \lambda^2(\cdot)) &= \lambda^2(m(A \cup AB), \cdot) = \delta F\left(\frac{1}{8}(1-b_S)^2\right) + (1-\delta)F\left((1-b_S)^2 s^{PPSpp}\right), \\ \lambda^1(A, \lambda^2(\cdot)) &= \delta F\left(\frac{1}{8}(1-b_S)^2\right) + \alpha(1-\delta)F\left((1-b_S)^2 s^{PPSpp}\right). \end{aligned}$$

Advertiser profit is given by $\lambda^1(\cdot)(1-p^*)(p^* - b_S)$, and firms therefore wish to deviate in messaging strategy if and only if there is an alternative signal resulting in a higher $\lambda^1(\cdot)$. Thus, neither AB -type nor A -type firms wish to deviate to $m(B)$ provided that

$$\min\left\{\frac{\delta}{1-\delta}, \frac{\delta}{\alpha(1-\delta)}\right\} \geq \frac{F\left(\frac{1}{8}(1-b_S)^2\right) - F\left((1-b_S)^2 s^{PPSpp}\right)}{F\left(\frac{1}{8}(1-b_S)^2\right) - F\left(\frac{\alpha}{8}(1-b_S)^2\right)},$$

Satisfaction of this condition is guaranteed since $\delta \geq 1/2$. A type B firm prefers not to transmit $m(A \cup AB)$ so long as (11) is satisfied. Thus, when (11) holds, no advertiser wishes to deviate from the equilibrium reporting strategy. ■

Proof of Proposition 6. Suppose that consumers expect informative messages and fix a Walrasian allocation and (per-click) price vector. The Diamond paradox applies so that optimal firm pricing is given by p^* . For firm j , changing from $m(A)$ to $m(AB)$ yields a change in click volume of

$$(16) \quad \Delta \lambda^2 = \delta \left[\overbrace{\psi_j\left(m(AB), \mathbf{m}_{-j}^*, \boldsymbol{\theta}_{-j}; A\right) - \psi_j\left(m(A), \mathbf{m}_{-j}^*, \boldsymbol{\theta}_{-j}; A\right)}^{\text{Change in A-type clicks = 0}} \right] \\ + (1-\delta) \left[\underbrace{\psi_j\left(m(AB), \mathbf{m}_{-j}^*, \boldsymbol{\theta}_{-j}; B\right) - \psi_j\left(m(A), \mathbf{m}_{-j}^*, \boldsymbol{\theta}_{-j}; B\right)}_{\text{Change in B-type clicks} \geq 0} \right] \geq 0,$$

where the zero value for the first bracket follows from Definition 1, and the inequality for the second from the requirement that the search rule be optimal (that consumers click on ads promising $t \in \theta$ first). Since AB -type firms have $\lambda^1 = \lambda^2$, sending $m(A)$ (or, by symmetry, $m(B)$) is not optimal for any per-click price below $1/4$. Since A -type firms (or, by symmetry, B -types) only increase the mass of unmatched consumers by deviating to $m(AB)$, doing so is not optimal when the per-click price is above $\alpha/4$. A similar exercise

obtains for A -types that deviate to $m(B)$ (or, by symmetry, vice-versa):

$$\Delta\lambda^2 = \delta \overbrace{\left[\psi_j \left(m(B), \mathbf{m}_{-j}^*, \boldsymbol{\theta}_{-j}; A \right) - \psi_j \left(m(A), \mathbf{m}_{-j}^*, \boldsymbol{\theta}_{-j}; A \right) \right]}^{\text{Change in } A\text{-type clicks} \leq 0} + (1 - \delta) \underbrace{\left[\psi_j \left(m(B), \mathbf{m}_{-j}^*, \boldsymbol{\theta}_{-j}; B \right) - \psi_j \left(m(A), \mathbf{m}_{-j}^*, \boldsymbol{\theta}_{-j}; B \right) \right]}_{\text{Change in } B\text{-type clicks} \geq 0}.$$

The firm thus trades some matched consumers for unmatched consumers, which is never optimal when the per-click price is above $\alpha/4$.

Since all clicks are worth at least $\alpha/4$ to every firm, every Walrasian price vector has an analogue in which all clicks cost at least $\alpha/4$ (such a price can be implemented by auction with a reserve price equal to $\alpha/4$). It remains to show that there exists a Walrasian allocation for all non-discriminatory search rules with informative ads. Such an allocation can be found as follows: (i) If $\cup_j \theta_j = \{A\}$ then (since all firms are of the same single-product type) the slots may be allocated arbitrarily. Writing $\psi_j^t \equiv \psi_j \left(m^*(\theta_j), \mathbf{m}_{-j}^*, \boldsymbol{\theta}_{-j}; t \right)$, a Walrasian price can be found by setting the price per-click for advertiser/slot j equal to

$$\frac{\delta \psi_j^A}{\delta \psi_j^A + (1 - \delta) \psi_j^B} \frac{1}{4} + \frac{(1 - \delta) \psi_j^B}{\delta \psi_j^A + (1 - \delta) \psi_j^B} \frac{\alpha}{4},$$

or $\alpha/4$ if $\delta \psi_j^A + (1 - \delta) \psi_j^B = 0$ (a symmetric price can be set if all firms are of type B). (ii) If $\cup_j \theta_j = \{A, B\}$ then a Walrasian allocation must have slots allocated such that each of A and B are offered by some advertiser (otherwise a firm offering the missing good can replace an advertiser and enjoy a per-click value of $1/4$). A Walrasian price can then be found by having each advertiser pay a price per-click of $1/4$. ■

Proof of Proposition 7. Suppose that the consumer search strategy is ND, and that consumers expect honest reporting. For an A -type advertiser, deviating to $m(AB)$ yields the change in click-throughs given in (16), where the inequality is strict with positive probability (e.g., for the case in which all advertisers are of type A). Such a deviation increases the firm's number of expected ex-post matches, and hence its profits. ■

REFERENCES

- AGARWAL, N., S. ATHEY, AND D. YANG (2009): "Skewed Bidding in Pay Per Action Auctions for Online Advertising," *American Economic Review Papers and Proceedings*, 99(2), 441–447.
- ANDERSON, S. P., AND R. RENAULT (2006): "Advertising Content," *American Economic Review*, 96(1), 93–113.

- ATHEY, S., AND G. ELLISON (2011): "Position Auctions with Consumer Search," *Quarterly Journal of Economics*, 126(3), 1213–1270.
- BAYE, M. R., J. MORGAN, AND P. SCHOLTEN (2004): "Price Dispersion in the Small and in the Large: Evidence from an Internet Price Comparison Site," *Journal of Industrial Economics*, 52(4), 463–496.
- CHEN, Y., AND C. HE (2006): "Paid Placement: Advertising and Search on the Internet," Working Paper.
- DELLAROCAS, C., AND S. VISWANATHAN (2008): "The Holy Grail of Advertising? Allocative Efficiency and Revenue Implications of "Pay-per-Action" Advertising in Environments with Quality Uncertainty," Working Paper.
- ELLISON, G., AND S. F. ELLISON (2009): "Search, Obfuscation, and Price Elasticities on the Internet," *Econometrica*, 77(2), 427–452.
- ELLISON, G., AND A. WOLITZKY (2008): "A Search Cost Model of Obfuscation," Working Paper.
- FARRELL, J. (1993): "Meaning and Credibility in Cheap-Talk Games," *Games and Economic Behaviour*, 5(4), 514–31.
- GROSSMAN, G. M., AND C. SHAPIRO (1984): "Informative Advertising with Differentiated Products," *Review of Economic Studies*, 51(1), 63–81.
- MEURER, M., AND D. O. STAHL (1994): "Informative Advertising and Product Match," *International Journal of Industrial Organization*, 12(1), 1–19.
- NELSON, P. (1974): "Advertising as Information," *Journal of Political Economy*, 82(4), 729–754.
- SUNDARARAJAN, A. (2003): "Pricing Digital Marketing: Information, Risk Sharing and Performance," Working Paper.
- VARIAN, H. R. (2007): "Position auctions," *International Journal of Industrial Organization*, 25(6), 1163–1178.
- WILBUR, K. C., AND Y. ZHU (2009): "Click Fraud," *Marketing Science*, 28(2), 293–308.
- WILSON, C. M. (2010): "Ordered search and equilibrium obfuscation," *International Journal of Industrial Organization*, 28(5).
- ZHU, Y., AND K. C. WILBUR (2011): "Hybrid Advertising Auctions," *Marketing Science*, 30(2), 249–273.