

Search Quality and Revenue Cannibalisation by Competing Search Engines*

Greg Taylor[†]

ABSTRACT: Consumers are attracted by high quality search results. Search engines, though, essentially compete against themselves since consumers are induced to substitute away from advertisement links when their organic counterparts are of high quality. I characterize the effect of such revenue cannibalisation upon equilibrium quality when search engines compete for clicks. Cannibalisation provides an incentive for quality degradation, engendering low quality equilibria—even when provision is costless. When consumers exhibit loyalty there is a ceiling above which result quality cannot rise, regardless of what the maximum feasible quality happens to be. Seemingly pro-competitive developments may exert downward pressure on equilibrium quality.

This article was published as as:

TAYLOR, G. (2013): “Search Quality and Revenue Cannibalization by Competing Search Engines,” *Journal of Economics & Management Strategy*, 22(3), 445–467.

1 INTRODUCTION

There has recently been much interest in the possibility that Internet search engines might seek to manipulate the results that they offer. Although these results are typically provided to users for free, their increasing importance as a tool in an information-driven society (and the apparent power thus bestowed upon their

*I express my gratitude to Robin Mason for many useful discussions and comments. Thanks are also due to David Gill, Justin P. Johnson, Alessandro Mennuni, Maksymilian Kwiek, David Myatt, Markus Reisinger, Juuso Välimäki, Alexander White, a coeditor, participants at the ‘Fifth bi-annual conference on The Economics of the Software and Internet Industries’, Toulouse, and seminar participants at the University of Southampton for their comments and suggestions. Two anonymous referees provided exceedingly detailed and generous feedback that has helped to make this paper more readable and also prompted development of a number of the model’s substantive results and extensions. Financial support from the ESRC is gratefully acknowledged. n.b. Owing to journal editorial policies, the spelling of the published article’s title differs from that of this working paper version: Cannibalisation→Cannibalization.

[†]Oxford Internet Institute, University of Oxford, 1 St Giles, Oxford, UK, OX1 3JS. Email: greg.taylor@oii.ox.ac.uk; web: <http://www.greg-taylor.co.uk>.

providers) has prompted scrutiny and, with it, accusations of malpractice.¹ Others, though, have insisted that, since the competition is ‘only a click away’,² search engines will naturally endeavour to provide the best results possible. The lack of a consensus on the incentives facing search engines creates a degree of ambiguity with respect to the appropriate regulatory stance vis-à-vis search engines’ provision of results. In this paper I address the question ‘when does competition suffice to induce search engines to provide high quality search results?’ or, phrased differently, ‘when and why might competing search engines benefit from deliberately degrading their result quality?’

Of particular interest in this matter is the interplay between paid-for sponsored search results (advertisements, or A-links) and free, non-advertisement search results (so-called organic links, or O-links) when there are consumers that search optimally. Intuitively, by providing high quality O-links a search engine attracts consumers to visit its site. This is beneficial if the same consumers, in an attempt to minimise search costs, stay to also click on revenue bearing advertisements rather than continuing their search elsewhere. However, there exists a countervailing effect since search engines face competition for A-link clicks not only from links at rival search engines, but also often from their own organic links. Searchers are presented with a list of organic results that often contains several commercial merchants. For example, in Figure 1 the organic results from a Google search for ‘camera’ include links to Jessops (a British camera retailer) and online retail giant Amazon.co.uk, and there is a danger that consumers’ needs are satisfied by such firms before any A-link is clicked. Surveys, experiments, and empirical studies generally show that consumers are prepared to use organic links to satisfy commercial needs in this way.³ The market is therefore characterised by a kind of revenue cannibalisation that results in a delicate trade-off between the complementary effects of O-links (the incentive to compete for market share) on the one hand, and the desire to have consumers click on advertisements on the other.

I show that this trade-off engenders equilibria in which search engines deliberately degrade their (organic) result quality—even when faced with competition and in the absence of any pecuniary cost for quality provision. This issue becomes particularly acute when one accounts for the importance of consumer loyalty (or

¹This culminated in 2011–2013 with high-profile investigations into Google’s search result generation practices by regulators in both the EU and the US.

²See <http://googleblog.blogspot.com/2008/02/our-secret-sauce.html>, accessed 17th January 2011. See also de Cornière (2011) and Eliaz and Spiegler (2011) for formal models in which competition suffices to induce maximal quality provision.

³See, for example, Agarwal, Hosanagar, and Smith (2011); Jansen, Zhang, and Schultz (2009).

inertia) in the search industry. Under such circumstances competition is initially intense in the sense that chosen quality may be maximal when early technology places a relatively low bound on result quality. However, as one allows the maximum technologically feasible result quality to increase there exists a ceiling above which equilibrium organic link quality will not rise. This implies that fierce competition in the Internet search industry may soften considerably once technology surpasses a particular threshold and that historical performance may not be a reliable barometer for predicting future effort in quality provision.

The possibility of equilibrium quality degradation naturally leads to questions about the prevailing environment's role in inducing such behaviour—and the welfare consequences for users of changes in this environment. Organisational and technical changes within the search industry can have surprising effects. For example, consumers may directly benefit from an improvement in the quality of sponsored links. However, I find that quality degradation is most easily sustained in equilibrium precisely when sponsored link quality is highest. More generally, I show that improvements in sponsored link quality can result in a downward distortion of organic link quality that leaves consumers worse-off overall. Likewise, a reduction in switching costs might be expected to foster competition between search engines. However, since search engines have little incentive to compete for the attention of consumers who will switch before clicking on an advertisement, I find that that high switching costs may be pro-competitive and that making search engines more substitutable in the eyes of consumers can induce them to compete less fiercely in quality. Thus, even accounting for the fact that switching costs impose a direct burden upon consumers, I show that overall consumer utility can fall if such costs are reduced.

Gandal (2001) conducts an empirical study of competition within the Internet

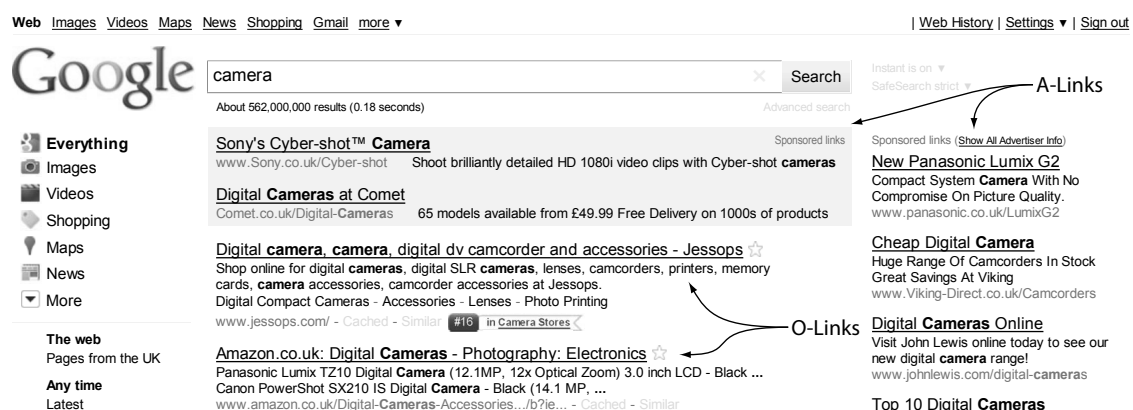


FIGURE 1 O- and A-Links in a Google UK search for “camera”.

search engine market. Two results are of particular interest for the current work: firstly, he finds that consumers are willing to switch between engines in order to find what they are looking for. Secondly, it is shown that the relevance of search results is by far the most important determinant of search engine ranking (as measured by number of searches conducted). This suggests that search engines have a strong incentive to compete on result quality. More recently, Yang and Ghose (2010) have shown that offering organic links increases the volume of sponsored link clicks enjoyed by a search engine, which is consistent with search engines using organic links to compete for user attention. Taken together with the obvious importance of search engines in modern society, these results motivate the model developed below.

This paper is most directly connected to the literature on the industrial organisation of the search engine industry. White (forthcoming) obtains results similar in spirit to those presented here using a monopolist search engine that offers both sponsored results and some quantity of organic links. Additional organic links reduce consumers' search costs and induce more consumers to search, but simultaneously provide competition for those consumers' business in the final goods market—thus reducing consumer prices. This competition amongst sellers reduces the amount that they are willing to pay for advertisements. Unlike White, I provide an explicit model of consumer clicking behaviour and use this to examine an alternative channel for cannibalisation arising as a consequence of consumer substitution between link types, rather than as a result of reductions in advertisers' willingness to pay. In practice, both mechanisms seem likely to be important and including a declining willingness to pay for advertisements in my model serves to further strengthen the cannibalisation effect identified in this paper. de Cornière (2011) and Eliaz and Spiegler (2011) present models of search engines that provide sponsored search results but not organic links. These authors find that competition suffices to induce what amounts to maximal quality provision. By contrast, I find that search engines may wish to degrade the quality of their organic links even under competition—owing to the cannibalisation that these links induce.

de Cornière and Taylor (2012) consider the issue of search engine bias. In their model, a monopoly search engine directs too much traffic towards those horizontally differentiated content-bearing websites that do not provide much competition in the ad market. Allowing the search engine to integrate with a content site need not increase the prevailing level of bias and may, in fact, cause equilibrium bias to fall. In contrast to de Cornière and Taylor (2012), I abstract from the horizontal considerations intrinsic to search bias and focus on the question of whether and when

the presence of competition in the the search industry is sufficient to discipline search engine quality setting behaviour.

Berman and Katona (forthcoming) are concerned with an alternative mechanism by which search result quality can be degraded. They study the incentives for website owners to engage in search engine optimisation, which improves their search ranking without materially improving their ability to satisfy visitors. Such practices can not only undermine the effectiveness of a search engine, but also divert resources away from substantive content provision. Nevertheless, search engines may become more useful when high quality sites have a relatively high incentive to invest in optimisation. Hagiu and Jullien (2011) consider a monopolist intermediary that can direct consumers to well- or poorly-matched stores. They show that the intermediary will often want to divert search—sending the consumer to a poorly-matched store first in order to induce revenue bearing store visits. The core logic of my paper is similar, but I extend this reasoning into a competitive environment and consider optimal consumer substitution between the different kinds of links that are endemic to the search engine industry.

Pollock (2010) and Telang, Rajan, and Mukhopadhyay (2004) also have theoretical models of the on-line search industry. Telang, Rajan, and Mukhopadhyay (2004) model entry in the Internet search industry and attempt to explain the existence of low quality search engines when consumers pay a price of zero. Pollock (2010) demonstrates a tendency for concentration in the internet search industry and then explores a number of welfare and regulation issues with a monopolist search provider. In both papers, search engine revenues are treated in a reduced form fashion. In this paper, by contrast, I explicitly model consumer link choice, which makes the profitability of a consumer's visit endogenous to the search engine's chosen quality. This introduces a number of new equilibrium considerations which lead to the cannibalisation described above.

Much of the interest in search engines has focused on bidding in the auctions used to sell sponsored search results. In particular, Edelman, Ostrovsky, and Schwarz (2007) and Varian (2007) model the 'position' auction framework that search engines typically use to sell advertisements. Chen, Liu, and Whinston (2007) extend this analysis to consider the optimal design of competing ad auctions, whilst Athey and Ellison (2011) and Chen and He (2011) consider consumers who search optimally through the list of advertisements. Arnold, Darmon, and Penard (2012), Katona and Sarvary (2010), and Xu, Chen, and Whinston (2012) model the impact of organic search results on the bidding incentives for advertisers at a search engine. In these

papers advertisers bid for position in a list of sponsored search results and consumer clicking behaviour is distributed across links following a parametrically prescribed pattern that depends upon the auction's equilibrium outcome. My paper differs from this work by focusing on technology adoption by competing search engines and the ways in which consumers optimally substitute between organic and sponsored links in response to the ensuing search result quality.

On a more general level, a literature on prominence beginning with Armstrong, Vickers, and Zhou (2009) considers the equilibrium effects of 'forcing' consumers to search sellers in a common order. A key result in this literature is that, when all other consumers search in the prescribed order, it is optimal for any individual consumer to do likewise so that prominence is self-reinforcing.

There is also a large literature on the relationship between advertising and other media content—for a summary see Section 10 of Bagwell (2007). Examples include Anderson and Coate (2005), who analyse the provision of programs and advertisements by radio and television broadcasters; Gabszewicz, Laussel, and Sonnac (2001), who study advertising in the newspaper press; and Taylor (2012), who examines content publishers' use of quality investments to hold consumers' attention and thereby appropriate more rents from advertisers. In these papers, much as for search engines, advertisements are provided alongside additional content that is attractive to consumers but does not generate revenue directly. On a related issue, Ellman and Germano (2009) consider the incentive of newspapers to bias their reporting in such a way that readers are more receptive of advertisement messages. The trade-off involved is between increasing the value of an ad, and increasing the readership of the newspaper. The Internet search advertising market differs from these contexts, however, in the fact that, whilst consumers are unlikely to view television ads and programmes as substitutes, consumers can substitute between different kinds of search engine links so that organic search results compete for the same clicks as the revenue generating advertisements appearing alongside them.

2 SIMPLE MODEL

Two search engines, g and m , provide search results to consumers at zero cost. A query at a given search engine returns two results: one organic search result (O-link), and one advertisement (or sponsored-search) result (A-link). Let A_i and O_i respectively denote the A-link and O-link at site i .⁴ A unit mass of homogeneous,

⁴Previous work has focused either on the case in which there is a single list of advertisements only, or on that in which consumer search behaviour follows some specified pattern; using a single link as

risk-neutral consumers have a particular need that they seek to satisfy by searching the Internet. Each time a consumer visits (or re-visits) a search engine they must pay a visit cost, $S > 0$. When the search results are returned the consumer may click on them as he or she pleases, but must pay a further search cost, $s > 0$, for each link that is clicked. If a clicked link matches the consumer's need then the consumer receives an expected surplus, which I normalise to 1. Matches are statistically independent across links and consumers, and I assume that consumers exhibit unit demand so that a satisfied consumer can gain no further utility from searching.

Each search engine implements a proprietary algorithm, inducing a distribution over match probabilities for its respective O-link. Thus search engine i 's choice of 'quality', denoted $p_i \in [0, p^{\max}]$ (where $p^{\max} \in [s, 1]$ is the maximum technologically feasible quality), refers to the expected match probability associated with link O_i .⁵ To emphasise the role of cannibalisation I assume that quality provision is costless—I discuss relaxation of this assumption in section 5.4.

I assume that the same link appears in the A-link slot at both search engines and that search engine i receives an amount b each time A_i is clicked.⁶ To understand the rationale for these assumptions, suppose that each search engine sells its A-link slot by means of a second price auction, and that each advertiser, j , expects to match with a proportion ψ_j of its visitors and makes v_j per match. Knowing that no consumer will click the same link twice (since clicking is costly), it is then a (weakly) dominant strategy for j to multi-home and bid $b_j = v_j \psi_j$ per click for the advertising opportunity at each search engine. Under such circumstances, absent a tie, the winner of both ad auctions will be the same firm. This bidding induces an expected match probability (quality), $q = E(\max \psi_j)$, for the common A-link and a price per click given by the $v_j \psi_j$ of the highest losing bidder. Search engines must take advertiser bids—and hence q —as given. The fixed b assumption can be relaxed and this matter is discussed in Section 5.4. The search engine receives nothing when its organic link is clicked.

The A-link and two O-links point to websites each drawn from separate pools of firms so that there are always three distinct links available to the consumer. Consumers need not observe the realised match probability prior to clicking a link, but are aware of the average match probabilities p_g , p_m , and q —either because

a proxy for each of the lists of O- and A-links allows an explicit and tractable representation of the consumers' optimal click order, whilst preserving the consumers' ability to substitute between the two types of link.

⁵Since consumers and search engines are risk-neutral, only the expected match probability matters.

⁶Qualitatively similar results can be obtained when the search engines offer differing A-links. Intuitively, this strengthens the effect of cannibalisation by making it more attractive to be visited second and thereby weakening the incentive to compete for market share via increases in p .

they search regularly (and learn about quality over a relatively short time-scale), or because the search engine qualities are widely discussed in the press or in society at large.

If $q < s$ then no consumer ever clicks an A-link and search engines, which make zero profits, are indifferent across all choices of p . When $s \leq q < S + s$, consumers click on an A-link if and only if organic link quality is sufficiently high to compensate consumers for the expense of visiting a search engine (which will always be true in equilibrium), but the below analysis otherwise remains essentially unchanged. Hereafter, I assume that $S + s \leq q < 1$, so that consumers are always willing to click on A-links. I am also ruling out the trivial case in which $q = 1$. Briefly, $q = 1$ gives rise to a continuum of equilibria in which search engines choose an arbitrary $p \in [0, 1)$, and consumers click an A-link at an arbitrarily chosen search engine resulting in immediate satisfaction. I assume that half of all consumers visit each search engine first whenever they are indifferent.

To summarise: search engines move first and simultaneously select a quality p_i . Consumers observe p_g, p_m, q, S , and s , and select whether, and in which order to click each link. The game ends when all consumers have had their need met, have exhausted the set of available links, or do not wish to click any further links. Throughout the paper, the solution concept that I use will be that of sub-game perfect Nash equilibrium, focusing on equilibria that are in pure strategies for the search engines.

3 EQUILIBRIUM BEHAVIOUR IN THE SIMPLE MODEL

3.1 Low quality organic results: minimal quality and revenue cannibalisation

The purpose of this subsection is to show how the threat of revenue cannibalisation from organic results can induce deliberate degradation of their quality below the maximum technologically feasible level and to establish a lower bound on sustainable equilibrium quality. Consumers are active participants in this market and in any equilibrium it must be the case that they click links in an order that maximises their expected utility. The complete form of such an optimal click order is formally described in Appendix A, but for now it suffices to note that any optimal click order must involve beginning at the higher quality search engine and clicking the higher quality of the two links there first.

That consumers prefer to visit high quality search engines first gives search

engines an obvious incentive to increase the quality of their results in order to attract as many users as possible. Indeed, given consumer search behaviour, there is no (cannibalisation) incentive to reduce quality whenever search engines' best links are their advertisements since every consumer will then click on a sponsored link first. This notion—that organic links do not cannibalise sponsored link clicks for O-link qualities that are sufficiently low—can be extended quite naturally to establish that an equilibrium in which organic links are strictly inferior to their sponsored counterparts can only be sustained by a technological constraint on their quality. More precisely,

Lemma 1 *When organic link quality is constrained to be below sponsored link quality (when $p^{\max} < q$), equilibrium quality must be set at the maximum feasible level: $p_g = p_m = p^{\max}$.*

Lemma 2 *When the maximum feasible organic link quality exceeds the sponsored link quality (when $p^{\max} \geq q$), any equilibrium must have $p_g = p_m \geq q$ (search engines set a symmetric organic link quality not less than the sponsored link's quality).*

All proofs are presented in Appendix B. Evidence from, for example, Agarwal, Hosanagar, and Smith (2011) indicates that consumers find organic results to be no less relevant than their sponsored counterparts—even when searching with the intention to buy. My focus for the remainder of the paper will therefore be on the case in which $p^{\max} \geq q$ so that search engines are able to provide organic links whose quality exceeds that of the sponsored results.

Whilst consumers will favour clicking advertisements first for organic links of low quality, high quality organic results induce consumers to substitute away from clicking sponsored links in preference of their superior organic counterparts. When search engines set high organic link qualities there is, therefore, revenue cannibalisation in the sense that some consumers will click on and be satisfied by an organic link without ever clicking on an advertisement. The higher is a search engine's organic link quality, the more consumers will be satisfied in this way and the stronger is the cannibalisation effect. This creates an incentive for search engines to deliberately degrade their quality below the maximum feasible level in order to minimise such cannibalisation. In fact, it transpires that this incentive can suffice to drive equilibrium quality to the floor on admissible quality levels that was established in Lemma 2.

Equilibrium 1 (Low quality cannibalisation equilibrium) *The lowest quality that can ever be sustained in equilibrium has organic link quality set equal to sponsored link quality ($p_g = p_m = q$). Such an equilibrium can be supported whenever advertisements are sufficiently useful—specifically when $q \geq 1/2$.*

There is an essential balancing act in the search engines’ problem: by ‘turning-down’ the quality of its organic links, Google can induce more of its users to click on advertisements. Doing this too much, however, risks causing users to switch to using an alternative search engine as their primary source. Besides establishing that this trade-off can bite under competition, Equilibrium 1 tells us that an important factor governing it is the prevailing sponsored link quality, q . Quality degradation becomes more sustainable as sponsored links improve because any attempt to poach a rival’s users then necessarily involves higher organic link quality and is thus associated with more severe revenue cannibalisation.⁷ Thus, high sponsored link qualities make tacit collusion around low quality (high profit) equilibria easier to sustain by increasing the cost of breaking any implicit arrangement. It follows that, whilst an increase in sponsored link quality may appear to provide for the more efficient allocation of consumers’ attention, it can also facilitate the existence of new, low quality equilibria.

3.2 Higher qualities: switching, sticking, and consumer lock-in

Having considered the lowest possible equilibrium quality, it is interesting to next ask how high equilibrium quality can go. As we will see in this subsection, an important issue in this respect is the relative ease with which consumers can switch search engines mid-query. It is intuitively clear that if the within-site link clicking cost, s , is close to zero then consumers never wish to switch search engines before all links at the first search engine that they visit have been exhausted (because exhausting all of the links at a search engine is then essentially costless). Since A-link clicks then only ever occur at the first-visited search engine, there is an incentive for each search engine to offer a slightly higher quality than its rival in order to capture all of these clicks for itself. This process of one-upmanship is clearly reminiscent of Bertrand price competition and, much like the standard Bertrand case, large portions of search engine profits may be dissipated as search engines compete result quality

⁷Specifically, when a search engine deviates to some $p_i = q + \epsilon$ (ϵ small), the mass of consumers that eventually click its advertisement is $1 - (q + \epsilon)$ so that the deviation becomes less attractive as q becomes large.

up to p^{\max} .^{8,9}

Remark 1 *When within-site link clicking costs, s , approach zero there exists an equilibrium in which search engines set the maximum technologically feasible quality: $p_g = p_m = p^{\max}$.*

More generally, when s is small consumers are effectively locked-in to the first search engine that they visit because the cost of continuing to click links there is low *relative to the cost of switching over to a second search engine*. If, on the other hand, clicking links is quite onerous then the strength of this lock-in effect is attenuated by virtue of the fact that switching search engines (at cost S) becomes relatively less costly vis-à-vis the alternative of staying put to click another link of inferior quality.

More precisely, suppose for concreteness that $q < p_g \leq p_m$. Having incurred the initial visit cost, S , consumers find it optimal to click on m 's organic link first. If g 's organic link is of a high enough quality then a consumer has much to gain by next switching search engines—clicking the O-link at both—before finally clicking on the A-link at the second site he visits. The total expected utility that each consumer gets from ‘switching’ to next click on g 's organic link can be written as

$$(1) \quad U_{\text{switch}} = \underbrace{p_m(1-S-s)}_{\text{Satisfied by } O_m} + \underbrace{(1-p_m)p_g(1-2S-2s)}_{\text{Satisfied by } O_g} + \underbrace{(1-p_m)(1-p_g)q(1-2S-3s)}_{\text{Satisfied by } A_g} + \underbrace{(1-p_m)(1-p_g)(1-q)(-2S-3s)}_{\text{Unsatisfied}}.$$

In contrast, if g 's organic link quality is only marginally larger than that of the sponsored link then the benefit of switching search engines mid-query is small relative to the cost of doing so; consumers then prefer to click both the O-link and A-link at m and conserve on switching cost expenditure. The utility from such ‘sticking’ behaviour is

$$(2) \quad U_{\text{stick}} = \underbrace{p_m(1-S-s)}_{\text{Satisfied by } O_m} + \underbrace{(1-p_m)q(1-S-2s)}_{\text{Satisfied by } A_m} + \underbrace{(1-p_m)(1-q)p_g(1-2S-3s)}_{\text{Satisfied by } O_g} + \underbrace{(1-p_m)(1-q)(1-p_g)(-2S-3s)}_{\text{Unsatisfied}}.$$

⁸Although reminiscent of standard Bertrand competition, the mechanism here differs somewhat since profit dissipation in this model occurs indirectly via consumer substitution between link types (cannibalisation): as qualities are competed upwards an ever smaller fraction of consumers find themselves needing to click on a profit-yielding A-link.

⁹A result analogous to Remark 1 can be established for environments with multiple organic links per-site. Details are provided in the supplementary Appendix.

Setting (1) and (2) equal to one another and simplifying, one finds that the consumer is indifferent between sticking and switching when

$$\frac{p_g}{S+s} = \frac{q}{s}.$$

Clicking a second link at the current search engine costs s , whilst switching to click a link at a different search engine has total cost $S+s$. The above indifference condition therefore says that the utilities per-unit of (search cost) expenditure for next clicking O_g and A_m should be equal if the consumer is to be indifferent between switching and sticking—otherwise, the consumer strictly prefers to click the link with the higher benefit-cost ratio. Writing $\sigma \equiv (S+s)/s$ for the *relative* cost of switching versus sticking, the condition can be rearranged to find the threshold value of p_g that achieves such indifference:

$$(3) \quad p_g = \frac{S+s}{s}q \equiv \sigma q.$$

There is a widespread perception that the Internet makes the cost of switching between sites very low. However, once a consumer is on a site, it is not the absolute size of the switching cost per se that matters, but rather its size relative to the cost of staying put.¹⁰ Thus, the relative cost, σ , of switching versus sticking has an important role to play in driving consumer behaviour. If this cost is high (if $p_g < \sigma q$ in the example above) then consumers face a de facto lock-in and always stick. Conversely, low relative switching costs ($p_g > \sigma q$) induce consumers to switch after their first O-link click.

It follows that, when the relative switching cost is sufficiently high (so that $\sigma q > p^{\max}$), consumers face a de facto lock-in to their first-visited search engine for *every* feasible quality level; much as in Remark 1, search engines then have a persistent and strong incentive to compete to be visited first in order to reap the benefits of having locked-in users who will click A-links.¹¹ When, on the other hand, relative switching costs are low, consumers readily switch between search engines to take advantage of high quality (organic) search results. In this latter case, a search engine facing a high-quality rival has an incentive to respond with *lower* quality organic results in order to capture the A-link clicks of such switchers. Maximal quality provision can not be an equilibrium under such circumstances and quality

¹⁰The other important factor—captured by the presence of p_g and q in (3)—is, of course, the size of the net benefit to switching.

¹¹Note that Remark 1's example of $s \rightarrow 0$ is simply a special case of this more general lock-in phenomenon since $s \rightarrow 0$ implies $\sigma q \rightarrow \infty$.

will be competed down to the point at which attracting additional consumers to switch and click an A-link is no longer possible. Putting this together:

Equilibrium 2a (Maximal quality equilibrium) *There exists an equilibrium in which search engines set the maximum feasible quality ($p_g = p_m = p^{\max}$) if and only if relative switching costs, σ or sponsored link quality, q , are sufficiently high—specifically when $\sigma q \geq p^{\max}$.*

Equilibrium 2b (Reduced quality switching equilibrium) *If relative switching costs, σ , or sponsored link quality, q , are sufficiently low (specifically if $\sigma q < p^{\max}$) there exists an equilibrium in which both search engines set quality σq . This is then the highest sustainable equilibrium quality.*

Search engines might be tempted to increase their quality to attract more first-time visitors in the hope that these consumers will stick-around to click a revenue bearing advertisement. Low switching costs, however, ensure that many of the consumers attracted by such an increase in quality will ultimately switch away from the deviating site before clicking on a sponsored link—making the putative deviation unprofitable. Equilibrium 2b thus reveals an additional channel by which search engines are incentivised to limit their provision of quality. In particular, low switching costs make consumers relatively fickle and reduce the pay-off to competing for their attention, thereby weakening quality competition. In this sense, Equilibria 2a and 2b establish that large relative switching costs can be pro-competitive by engendering ex ante competition to lock-in consumers who will stay on-site to click sponsored results.

We can examine the consequences of an increase in switching costs for consumers by substituting the equilibrium quality level given in (3) into consumers' utility (Equations 1 or 2) and differentiating with respect to the visit cost, S . It is easily verified that $\partial U / \partial S > 0$ so that making switching more costly always results in higher utility in Equilibrium 2b.¹² Thus, even after accounting for the fact that higher visit costs impose a direct welfare burden on consumers, an increase in S causes utility to increase in Equilibrium 2b by inducing stronger competition and higher equilibrium quality levels.

Corollary 1 *Increases in switching costs can induce higher equilibrium quality and make consumers better-off.*

¹²In particular $\partial U / \partial S = 2(1-q)(q-s)[(1-q)s - qS] / s^2 \geq 0 \Leftrightarrow \sigma q \leq 1$, which always holds in Equilibrium 2b.

Importantly, consumers are shielded from the usual negative (i.e. anti-competitive) consequences of high switching costs since the ensuing ‘lock-in’ is transitive—holding only within a single search query. Between queries, consumers can easily change to another search provider since the cost, S , of beginning a new search is the same at each search engine. Indeed, as we will see in the next section, it is this combination of strong within-query and weak between-query lock-in that ensures that the very highest quality of results can persist in equilibrium.

In summary, the simple model has produced equilibria which serve to highlight three distinct forces at work in the search market;¹³ namely, revenue cannibalisation, competition for market share, and intra-query switching. Which of these equilibria can be supported depends upon the size of switching costs and the quality of sponsored search results.

At any time there exists precisely one equilibrium in which result quality is relatively high. Indeed, quality in this equilibrium is forced to its maximum feasible level by strong ex ante competition for user visits when switching costs are high (Equilibrium 2a). If switching sites is relatively easy, however, then consumers are too fickle to merit strong competition for their attention and even the highest admissible equilibrium quality is then limited to be below the maximum feasible level (Equilibrium 2b).¹⁴

In addition to this ‘good’ equilibrium, high sponsored link qualities can foster the existence of a low quality equilibrium (Equilibrium 1) by giving search engines much to lose from any increase in quality that attracts more users at the cost of inducing consumers to click more organic links. Search engines then have an incentive to tacitly collude in quality degradation in order to avoid collectively incurring high degrees of revenue cannibalisation.

4 CONSUMERS WITH SEARCH ENGINE ‘LOYALTY’

Although the existence of equilibria with degraded quality has been established, the simple model above has engendered a stark equilibrium in which search engines compete quality up to its maximum feasible level in spite of the consequent revenue cannibalisation. This is ultimately a product of the discontinuous demand induced by the assumption that an arbitrarily small quality advantage is sufficient to capture the entire market. In practice, whilst the official Google position (for example) is

¹³Lemma 2, along with the reasoning developed in this subsection, implies that Equilibria 1, 2a, and 2b are exhaustive under pure search engine strategies.

¹⁴Note that if the cost, S , of visiting a search engine is close to zero then $\sigma q \rightarrow q$; it follows from Equilibrium 2b and Lemma 2 that q is then the unique equilibrium quality level.

that they are ‘one click away from losing every customer’,¹⁵ many consumers exhibit at least some degree of loyalty (or inertia) towards their preferred search engine. Demand should therefore be expected to react smoothly to small changes in quality. Indeed, search engines actively promote technologies—such as browser toolbars—that significantly lower the cost of accessing the default search engine and thereby induce such entrenchment. One question that naturally arises is therefore whether and when the very highest quality of O-links can be sustained if there is some degree of continuity to the demand faced by each search engine, and that is the concern of this section.

Moreover, the low quality equilibrium (Equilibrium 1) presented above arises as a consequence of a kind of tacit collusion between search engines—each search engine being prepared to degrade its quality only if its rival does likewise. In light of this, it is interesting to ask whether cannibalisation can ever make quality degradation always individually rational (in the sense that each search engine wishes to degrade its quality regardless of its rival’s strategy). It is to this end that I invoke a standard Hotelling linear city model in which consumers have heterogeneous visit costs distributed along a segment of the real line.

More concretely, suppose that the consumers are now uniformly distributed along a line of unit length with g located at point 0, and m at point 1. A consumer located at point x must pay a cost $S_g(x) = xS$ to visit g and $S_m(x) = (1 - x)S$ to visit m , where S now functions as a transport cost. As in the previous section, I maintain the assumption that $1 > q > S + s$, which says that every consumer would, in principle, be willing to click either search engine’s sponsored link. The model is otherwise as described in section 2. To begin, I focus on the limiting case with small clicking costs, s . This rules out potentially complicated switching behaviour (which, in general, will now be location dependent).

Given the structure of visit costs, consumers now determine which search engine to visit first in consideration not only of link qualities, but also their own personal search engine preference (i.e. their location, x). In particular, for some x^* the best response for a consumer is to visit g first whenever $x \leq x^*$ (otherwise begin at m), and once there to click on the organic link first if its quality exceeds that of its sponsored counterpart, and vice-versa. As in section 3, consumers stop clicking if their need is satisfied, or if the quality of the O-link at the second site is too small to merit a visit. An immediate question, then, is whether it is ever optimal for search engine i to choose some link quality $p_i < S$ —thus inducing some far away consumers to click

¹⁵See <http://www.google.com/competition/competitiveweb.html>, accessed 13th February 2012.

only the two links at its rival. In fact, it is possible to show that sponsored link quality constitutes a floor on equilibrium organic link quality.

Lemma 3 *Offering an organic link whose quality is below that of the sponsored link (setting $p_i < q$) is a dominated strategy for the search engines.*

The logic here is essentially the same as in Lemma 2: for very low organic link qualities search engines can increase their quality without incurring any cannibalisation (since consumers will always click an A-link first), and this is therefore a costless way to attract new users. Given Lemma 3, one can use the utility functions implied by $p_g, p_m \geq q$ to identify the location of the consumer indifferent between first visiting g and m .¹⁶ This location, x^* , satisfies

$$(4) \quad \frac{x^*}{1-x^*} = \frac{p_g + (1-p_g)q}{p_m + (1-p_m)q},$$

which simply says that the relative market shares of the two search engines are given by the relative probabilities with which they offer a link that satisfies a given consumer—higher quality search engines enjoying a larger share of the market. Thus g 's market share, x^* , is continuously increasing in its organic link quality—a benefit that must be weighed against the fact that the proportion of visitors who are satisfied by an organic link without clicking on an ad is also increasing in p_g .

A clear implication of Lemma 3, is that there can be no equilibrium with search engines setting quality below q . I now turn my attention to the converse possibility of equilibria with organic link qualities at least as high as their sponsored counterpart. Firstly, much as in the simple model, the prospect of revenue cannibalisation can foster an equilibrium in which search engines jointly reduce their quality to the minimum admissible level.

Equilibrium 3 *There exists a \underline{q} such that there is an equilibrium with organic link quality set equal to sponsored link quality ($p_g = p_m = q$) whenever $q \geq \underline{q}$ (when sponsored link quality is sufficiently high).*

This result is broadly analogous to Equilibrium 1 and is sustained by the same considerations—in particular, high sponsored link qualities again facilitate the existence of this kind of low-quality equilibrium. More interesting is what now happens for higher equilibrium qualities:

¹⁶That is, x^* is the solution to $p_g(1-Sx) + (1-p_g)q(1-Sx) + (1-p_g)(1-q)p_m(1-S) + (1-p_g)(1-q)(1-p_m)(-S) = p_y(1-S(1-x)) + (1-p_y)q(1-S(1-x)) + (1-p_y)(1-q)p_g(1-S) + (1-p_y)(1-q)(1-p_y)(-S)$.




FIGURE 2 Values of \bar{q} and \underline{q} for various technological quality ceilings, p^{\max} , and sponsored link qualities, q .

Equilibrium 4 *There exists a \bar{q} such that an equilibrium with organic link quality strictly greater than sponsored link quality can be sustained if and only if $q \leq \bar{q}$ (if sponsored link quality is sufficiently low). The unique such equilibrium has search engines set their quality equal to*

$$(5) \quad \frac{1-3q}{3-3q}$$

(which lies between q and $1/3$ for every $q \leq \bar{q}$) when doing so is technically feasible, and equal to the maximum feasible quality, p^{\max} , otherwise.

The form of \underline{q} and \bar{q} is summarised in Figure 2.¹⁷ The lower solid line shows \underline{q} , the locus of points above which there is no profitable deviation from Equilibrium 3. The upper solid line is the path, \bar{q} , below which there is no profitable deviation from Equilibrium 4.

A search engine that increases its organic link quality will induce consumers to substitute away from clicking on its sponsored links resulting in revenue cannibalisation. If the maximum feasible quality is relatively low then a search engine may find that it can increase its quality until the technological constraint binds without being much troubled by such cannibalisation. This is because, when O-link quality is constrained to be low, most organic links will necessarily fail to satisfy

¹⁷The analytic expressions for \underline{q} and \bar{q} can be found in the proofs of Equilibria 3 and 4 respectively.

consumers (who, therefore, ultimately click an advertisement anyway). However, as the march of technological progress admits ever higher result qualities—leaving ever more consumers satisfied by organic links—there comes a point at which further quality improvements sacrifice more consumer clicks to cannibalisation than they attract from the rival search provider. Search engines then do not wish to implement additional (feasible) improvements in their algorithm regardless of their rival’s strategy and it is always individually rational for search engines to degrade their result quality.¹⁸

Having historically observed seemingly vigorous quality improvements, a casual observer of a maturing search market might be tempted to conclude that competition alone is sufficient to ensure maximal quality provision. Equilibrium 4, however, cautions that one should not complacently count on this pattern being repeated indefinitely in the face of improving search technology since search engines’ ability to increase quality will eventually outstrip their willingness to do so. In terms of Figure 2, if the maximum feasible link quality (and sponsored result quality) are non-decreasing over time then the industry must eventually cross into either (i) the (dotted) region in which the quality ceiling in (5) binds, or (ii) the (hatched) region where no high quality equilibrium can be supported at all.

It has already been seen that an increase in sponsored result quality may admit new equilibria such as Equilibria 1 and 3 by facilitating collusion between search engines. In addition, note that the ‘high quality’ Equilibrium 4 has organic link quality that is (weakly) decreasing in sponsored link quality. Indeed, substituting the equilibrium quality (5) into consumers’ utility and differentiating reveals that

$$\frac{\partial U}{\partial q} = -\frac{4}{9(1-q)^2} < 0,$$

so that equilibrium consumer utility is *decreasing* (and total search cost expenditure increasing) in sponsored link quality once the quality ceiling bites: Although consumers directly benefit from an increase in sponsored result relevance, it transpires that this benefit is more than offset by the reduced incentive for organic link quality provision. Consumers may thus indirectly pay a heavy price for improvements in sponsored link quality, and claims that such improvements are made with consumers’ welfare in mind might be viewed with suspicion.

Corollary 2 *Improvements in sponsored link quality can decrease equilibrium qual-*

¹⁸In particular, note that profit for a search engine that sets $p > q$ is $\pi = (1-p)x^*b$. This is concave with $p \rightarrow 1 \Rightarrow \pi \rightarrow 0$ so that search engines will always eventually wish to degrade their quality.

ity and make consumers worse-off.

Why should equilibrium organic link quality decrease when sponsored link quality increases? The answer lies in the fact that as sponsored links become more useful consumers are more often satisfied by them and thus become less dependent on organic links to fulfil their needs. This implies that a search engine's market share is less sensitive to its organic link quality when sponsored link quality is high.¹⁹ The result is that search engines have less to gain by increasing their quality (but no less to loose) and thus implement lower qualities in equilibrium.

5 DISCUSSION

In this section I briefly describe the results from a number of extensions to the model. Further details are available in the supplementary appendix.

5.1 *Acquiring loyal consumers*

One can extend the model of Section 4 to allow for asymmetric distributions of consumer loyalty by supposing that consumers are distributed along the Hotelling line according to the generalised linear distribution with density $f(x) = 2ax + 1 - a$. Thus, $a = 0$ describes to the uniform case considered above, whilst $a > 0$ ($a < 0$) corresponds to a skewed distribution with most consumers loyal to m (to g). In the long-run, one might expect the degree of loyalty to a search engine to be determined endogenously in the sense that consumers decide which search engine to become loyal to (e.g. which toolbar to install) after considering the relative quality of the two. To capture this, I let $a = (p_m - p_g)/(p_m + p_g)$, so that more consumers become loyal to g whenever its quality increases relative to that of its rival. Relative to the case with loyalty exogenously fixed (given in Equilibrium 4), this kind of endogenous loyalty induces higher search engine qualities and can support equilibrium qualities above q for a strictly larger range of q . This is intuitive: endogenous loyalty gives search engines an additional marginal incentive to increase result quality in the short-run because doing so convinces consumers to adopt search-engine specific investments and thus become more loyal in the long-run. Despite this, it remains the case that cannibalisation considerations ultimately cause search engines to distort their quality downwards and imply a binding ceiling on equilibrium quality.

¹⁹This much can clearly be seen in (4), which reacts more slowly to changes in p_g when q is large.

Remark 2 *When search engines can endogenously induce consumer loyalty, cannibalisation imposes a ceiling on the maximum admissible equilibrium result quality independent of any technical feasibility considerations.*

5.2 Commercial versus informational searches

It is natural to think of consumers as being heterogeneous in dimensions besides their affinity to a particular search engine. For instance, a search for “car insurance”, or “London hotel” betrays a patently commercial intent, whereas a search for “Greg Taylor” is (alas!) much less likely to be conducted for commercial ends. One might expect a sponsored result to have a higher match probability for a commercial searcher than for an informational one. For their part, search engines cannot possibly hope to provide an individually tailored algorithm for each of the infinite variety of possible keywords. One can model this by supposing that each consumer has an idiosyncratic sponsored link quality, but that search engines must provide the same organic quality to all.

Writing $\omega(\cdot)$ for the density of q in the population, and recalling that the location, x^* , of the consumer who is indifferent between visiting g or m first (characterised in (4)) depends upon the A-link quality, one can write g ’s profits as

$$\pi_g = \left[(1 - p_g) \underbrace{\int_0^{p_g} \omega(q) \cdot x^*(q) dq}_{\text{Mass of consumers who visit } g \text{ first and have } p_g \geq q.} + \underbrace{\int_{p_g}^1 \omega(q) \cdot x^*(q) dq}_{\text{Mass of consumers who visit } g \text{ first and have } p_g < q.} \right] b.$$

It is clear that this profit function evaluates to zero at $p_g = 1$ so that an interior solution must prevail for any distribution of q —that is to say, cannibalisation must eventually impose a competitive ceiling on equilibrium result quality. In fact, if one takes the relatively simple case in which ω is uniform in the consumer population, the unique symmetric equilibrium quality level turns out to be $p_g = p_m = 0.245$.

Remark 3 *When consumers vary with respect to the commercial value of their search, cannibalisation imposes a ceiling on the maximum admissible equilibrium result quality independent of any technical feasibility considerations.*

5.3 ‘Loyalty’ with non-trivial clicking costs

It is interesting to extend the loyalty model of Section 4 to allow non-trivial link clicking costs, $s \gg 0$. As in Section 3.2, the presence of a positive link clicking costs means that some consumers may wish to switch away from a search engine before

exhausting its supply of links. Under such circumstances, making search engines more substitutable (i.e. reducing transport cost, S) causes equilibrium quality to fall: Whilst a search engine has a strong incentive to compete for the attention of sticking consumers, it would prefer any switching consumers to visit its rival first—leaving it to subsequently capture the sponsored link clicks. Thus, whilst the presence of sticking consumers exerts upward pressure on organic result quality, switchers have the opposite effect. All else equal, making search engines more substitutable induces more consumers to switch mid-query and thus decreases the mass of consumers over whom the search engines wish to compete. This is basically the same intuition described in Section 3.2. As before, it is the cost of switching relative to sticking that turns out to be important.

Remark 4 *In any equilibrium in which some consumers switch, equilibrium quality is increasing in the degree of differentiation and decreasing in the cost of clicking links.*

When, on the other hand, search costs are such that all consumers wish to stick, the trade-off faced by a search engine is the familiar one: reducing quality mitigates cannibalisation but also drives some users to instead use a rival search service. What is new here is that high qualities are especially attractive to consumers when the link clicking cost is high because they reduce the average number of clicks needed to attain satisfaction. A quality advantage therefore has an amplified effect on market share when $s > 0$, and this provides an additional marginal incentive to increase quality—resulting in qualities above the level described in Equilibrium 4. This notwithstanding, cannibalisation continues to impose a ceiling on the maximum sustainable equilibrium quality.

Remark 5 *Any equilibrium in which no consumers switch has a maximum admissible quality imposed by cannibalisation independent of any technical feasibility considerations.*

5.4 Costly quality provision

Throughout this paper I have assumed that quality can be provided at no cost in order to emphasise the role of the cannibalization effect. I have also assumed a constant, exogenous value for b . If, instead, $b_i = b(p_i)$ then, intuitively, competition between advertisers and the firms behind organic links may cause $b' < 0$.²⁰ A moderate, non-decreasing cost of quality provision $C(p_i)$, or a moderately negative b' would serve to

²⁰White (forthcoming) provides a formal model of such a phenomenon.

further discourage upward quality deviations.²¹ In the simple model, these factors therefore leave the results qualitatively unchanged. The effects of such a modification on the model when consumers exhibit search engine loyalty depend on the precise form of the $C(\cdot)$ and $b(\cdot)$ functions, but intuitively act in the same direction as cannibalisation in applying downward pressure on result quality—ensuring that such a modification to the model cannot increase the maximum sustainable equilibrium quality, but may decrease it.

5.5 Market structure

In the results above, competition for visits prompts search engines to cannibalise their revenues from A-link clicks. When this competition is taken away so is the incentive to provide O-links of a high quality so that only the cannibalisation effect remains. A monopolist will therefore not wish to offer an organic link whose quality exceeds its sponsored counterpart.²²

More generally, in an n -search engine oligopoly, results analogous to Equilibria 2a and 2b can easily be obtained using the same intuition developed in Section 3.2. More interesting is the effect upon Equilibrium 1 of having $n \geq 2$ search engines. Supposing that indifferent consumers visit each search engine first with probability $1/n$ reveals that

Remark 6 *The lowest quality that can be sustained in equilibrium has organic link quality set equal to sponsored link quality ($p_i = q \forall i$). Such an equilibrium can be supported whenever advertisements are sufficiently useful or consumers are sufficiently inclined to click on ads—specifically when $q \geq (n - 1)/n$.*

It is immediately clear that the condition for existence of the ‘low quality’ collusive equilibrium becomes less demanding as the number of search engines in the industry is reduced so that less competitive industries are more susceptible to this kind of systematic quality reduction. Search engines may therefore have an incentive to consolidate since this can create new equilibria with higher total industry profits (but lower organic result quality). Reduced competition can thus spill-over into the quality of search services enjoyed by consumers. This may prove to be an important consideration in evaluating merger proposals if the consumer’s search experience is part of the regulator’s objective.

²¹Excessive quality costs may result in equilibrium non-existence—see Pollock (2010).

²²If $S + s > q$ for some proportion of the consumer population then the monopolist may still wish to set a $p > q$ in order to induce those consumers to search—see White (forthcoming) for an analysis of this issue.

6 CONCLUSION

In this paper I have examined equilibrium behaviour in a simple model of the Internet search market and considered the question of when competition will suffice to induce the provision of high quality search results. I find that if search engines compete on result quality and consumers select search engines according to link relevance then search engines always provide positive quality organic links to attract consumers to their site. It is in this fashion that search engines ‘cannibalise’ their own revenue streams since the organic links that they provide compete for clicks with their own advertisements. The revenue cannibalisation effect engenders equilibria with degraded result quality, even when quality provision is without pecuniary cost. This result is particularly stark when one accounts for the loyalty or inertia that consumers exhibit towards their preferred/default search engine—in which case there exists a ceiling above which equilibrium quality cannot climb, regardless of improvements in search technology. These findings imply that, even where competition has been fierce in the past, unbridled expansion of search result quality may not be sustainable into the indefinite future.

Increasing the quality of sponsored search results helps to satisfy consumers who find those links useful, but can also facilitate collusion between search engines—which may respond with a reduction in organic link quality. This can leave consumers worse-off overall. In contrast, switching costs may be pro-competitive by inducing consumers to spend more time clicking sponsored links on a current site rather than switching to an alternative, and thereby creating a strong incentive to compete *ex ante* for consumer visits.

APPENDIX A OPTIMAL SEARCH STRATEGY IN THE SIMPLE MODEL

First, let

$$e \equiv \begin{cases} g & \text{if } p_g > p_m \text{ and } p_g \geq s \\ y & \text{if } p_g < p_m \text{ and } p_m \geq s \\ g \text{ w.p. } 1/2, y \text{ w.p. } 1/2 & \text{if } p_g = p_m \text{ or } \max\{p_g, p_m\} < s, \end{cases}$$

and $-e \equiv \{g, y\} - \{e\}$. Thus e is the search engine that a consumer prefers to visit first.

Strategy 1 Suppose that $q \geq S + s$. Any consumer best response strategy maps the link qualities $\{p_g, p_m, q\}$ into a click order (a_1, a_2, a_3) in the following manner. Begin

by clicking a_1 thus:

$$a_1 = \begin{cases} O_e & \text{if } p_e > q \\ A_e & \text{if } p_e < q \\ A_e \text{ w.p. } \lambda, O_e \text{ w.p. } 1 - \lambda & \text{if } p_e = q, \end{cases}$$

where a fraction λ of consumers break indifference by clicking on an A-link. If the consumer's need was met by a_1 then stop clicking (i.e. $a_2 = a_3 = \emptyset$), otherwise click a_2 :

$$a_2 = \begin{cases} A_e & \text{if } a_1 = O_e \text{ and } p_{-e} < \sigma q \\ O_{-e} & \text{if } a_1 = O_e \text{ and } p_{-e} > \sigma q \\ A_e \text{ w.p. } \lambda, O_{-e} \text{ w.p. } 1 - \lambda & \text{if } a_1 = O_e \text{ and } p_{-e} = \sigma q \\ O_e & \text{if } a_1 = A_e \text{ and } p_e \geq s \\ \emptyset & \text{if } p_e < s, \end{cases}$$

where $\sigma = (S + s)/s$ is the relative cost of switching versus sticking—see Section 3.2. If the consumer's need was met by a_1 or a_2 then stop clicking (i.e. $a_3 = \emptyset$), otherwise click a_3 :

$$a_3 = \begin{cases} O_{-e} & \text{if } a_2 \neq O_{-e}, \text{ and } p_{-e} \geq S + s \\ A_{-e} & \text{if } a_2 = O_{-e} \\ \emptyset & \text{if } p_{-e} < S + s. \end{cases}$$

APPENDIX B OMITTED PROOFS

Proof of Lemma 1. Consumers must use a strategy of the form given in Strategy 1. If $p_i \leq p_{-i} < p^{\max} \leq q$ then i 's A-link is clicked with some probability less than 1 and i has a profitable deviation in setting $p_i \in (p_{-i}, p^{\max})$. If $p_i < p_{-i} = p^{\max}$ then all A-link clicks occur at $-i$ and i makes zero profits. It is, however, possible for i to make positive profits by setting $p_i = p_{-i} = p^{\max}$, in which case its profits are $[\lambda + (1 - \lambda)(1 - p^{\max})]b/2 > 0$ when $p^{\max} = q$ and $b/2 > 0$ when $p^{\max} < q$. ■

Proof of Lemma 2. I begin by establishing the symmetry result: In equilibrium, the consumers' strategies must have the form of Strategy 1 in Appendix A. Suppose, for concreteness, that $p_i > p_{-i}$. If $p_i \leq q$ then $-i$ makes zero profits, but can make positive profits by setting some $p'_{-i} \geq \max\{s, p_i\}$. When $p_i > q$ and $p_{-i} < \sigma q$, i receives a profit of $(1 - p_i)b$ —which is decreasing in p_i —and therefore prefers to reduce p_i slightly. If $p_{-i} = \sigma q$ then i 's profits are $\lambda(1 - p_i)b$, so that i again prefers to decrease p_i . Finally, if $p_i > p_{-i} > \sigma q$ then the consumer uses click order $\{O_i, O_{-i}, A_{-i}\}$ and i 's profits are zero. There is thus a profitable deviation for i which has it set $p'_i \in (\sigma q, p_{-i})$. Since all

equilibria are symmetric in search engine strategies when $p^{\max} > q$, $\min\{p_g, p_m\} < q$ implies $p_g = p_m = p < q$. Some $i \in \{g, y\}$ must, then have its A-link clicked with probability less than 1. It follows that i has a profitable deviation $p'_i \in (p, q)$. ■

Proof of Equilibrium 1. Consumers must use a strategy of the form given in Strategy 1. When $p_g = p_m = q$, search engine expected profits are $\pi_i = (\lambda + (1 - \lambda)(1 - p_i))b/2$. Consider a deviation by search engine i to $p'_i < p_{-i} = q$. Consumers now click in the order $\{A_{-i}, O_{-i}, O_i\}$. i 's profits are zero. Suppose, instead, that i sets $p'_i > p_{-i} = q$. Consumers click $\{O_i, A_i, O_{-i}\}$, which gives a profit for i of $\pi'_i = (1 - p'_i)b$. Since this is decreasing in p'_i , it suffices to consider the limiting case with $p'_i = q$. For the deviation to be non-profitable, π_i must be greater than π'_i , which gives $q \geq 1/(1 + \lambda)$ when q has been substituted in place of p_i and p'_i . Since the consumer is indifferent about click order when $p_g = p_m = q$, any λ constitute a best response so that the proposed strategies form an equilibrium for $q > 1/2$. ■

Proof of Equilibrium 2a. (1: sufficiency of $p^{\max} \leq \sigma q$): Let $p^{\max} \leq \sigma q$. The consumers' strategies must have the form of Strategy 1 in Appendix A. $p_i > p^{\max}$ is not possible. Deviating to $p'_i < p^{\max} (< \min\{p_{-i}, \sigma q\})$ implies that consumers never click A_i , and i 's profits are zero.

(2: necessity of $p^{\max} \leq \sigma q$): Let $p^{\max} > \sigma q$. If $p_g = p_m = p^{\max}$ then all consumers switch and $\pi_i = \frac{1}{2}(1 - p^{\max})^2 b$. i , has a profitable deviation in setting $p'_i \in (\sigma q, p_{-i})$, which yields to i profits of $(1 - p^{\max})(1 - p'_i)b$. ■

Proof of Equilibrium 2b. In equilibrium, the consumers' strategies must have the form of Strategy 1 in Appendix A. With $p_g = p_m = \sigma q > q$, the expected profit for i is $\pi_i = \frac{1}{2}[\lambda(1 - p_i) + (1 - \lambda)(1 - p_{-i})(1 - p_i)]b$. A deviation by i that has it set $p'_i < p_{-i} = \sigma q$ leaves it with a profit of zero since the consumer never clicks on A_i if $p_i < \min\{\sigma q, p_{-i}\}$. Suppose instead that i deviates with $p'_i > p_{-i} = \sigma q$. The pay-off for i becomes $\pi'_i = \lambda(1 - p'_i)b$. Since this is decreasing in p'_i , it suffices to consider the limiting case of $p'_i = p_i = \sigma q$. The deviation is not profitable so long as $\pi_i \geq \pi'_i$. Substituting σq for p'_i , p_i , and p_{-i} in this expression and rearranging yields $\lambda \leq (1 - \sigma q)/(2 - \sigma q)$. Given that the consumer is, by definition, indifferent over all λ when $p_g = p_m = \sigma q$, satisfaction of this condition is consistent with equilibrium.

From Lemma 2, $p_g = p_m$ in equilibrium. Suppose that $p_g = p_m = p > \sigma q$, inducing profit of $\frac{1}{2}(1 - p)^2 b$. Each engine can increase its profits by deviating to $p' \in (\sigma q, p)$, yielding profit $(1 - p)(1 - p')b$. Thus $p > \sigma q$ can't be an equilibrium. ■

Proof of Lemma 3. Suppose that $p_g < q$. Since s is small, consumers click on A_g only if they visit g first. Denote by x^* the mass of such consumers. If p_g is a best response, then it must be the case that $x^* > 0$, since g can always induce nearby

consumers to visit it first by setting a p_g close enough to p_m .

Consider, then, the interior case with $0 < x^* < 1$. Since $p_g < q$, rational consumers always click on A_g before O_g . Thus g 's profits are given by x^*b . A consumer that clicks A_g first must either use the click order $\{A_g, O_g, O_m\}$ or else use $\{A_g, O_g, \emptyset\}$, which implies utility functions $U(A_g, O_g, O_m) = q(1-Sx) + (1-q)p_g(1-Sx) + (1-q)(1-p_g)p_m(1-S) - (1-q)(1-p_g)(1-p_m)S$, and $U(A_g, O_g, \emptyset) = q(1-Sx) + (1-q)p_g(1-Sx) - (1-q)(1-p_g)Sx$ respectively. Any rational consumer who does not find it optimal to use either of the above two click orders must visit m first. The possible equilibrium click orders in use by such consumers are $\{A_m, O_m, O_g\}$; $\{A_m, O_m, \emptyset\}$; $\{O_m, A_m, O_g\}$; and $\{O_m, A_m, \emptyset\}$, which have utility functions defined analogously to those above.

Now, given that $S_g(x)$ is continuously increasing in x , the set of x for which consumers click A_g first must be $[0, x^*]$ and there must exist a marginal consumer at x^* who is just indifferent between the click orders in use by those consumers at $x^* + \epsilon$ and $x^* - \epsilon$, ϵ small. That is to say, $\max\{U(A_g, O_g, O_m), U(A_g, O_g, \emptyset)\}$ must be equal to the maximum of $U(A_m, O_m, O_g)$, $U(A_m, O_m, \emptyset)$, $U(O_m, A_m, O_g)$, and $U(O_m, A_m, \emptyset)$ at $x = x^*$. Consider a small increase in p_g to $p'_g \in (p_g, q)$. It is straightforward to verify that

$$\min \left\{ \frac{\partial U(A_g, O_g, O_m)}{\partial p_g}, \frac{\partial U(A_g, O_g, \emptyset)}{\partial p_g} \right\} \geq \max \left\{ \frac{\partial U(A_m, O_m, O_g)}{\partial p_g}, \frac{\partial U(A_m, O_m, \emptyset)}{\partial p_g}, \frac{\partial U(O_m, A_m, O_g)}{\partial p_g}, \frac{\partial U(O_m, A_m, \emptyset)}{\partial p_g} \right\}.$$

Thus, the increase in p_g causes the marginal consumer to strictly prefer some click order that has him click A_g first: the mass of consumers clicking A_g , (and hence g 's profits) is thus increased. By the continuity of p_g , there exists such a p'_g for all $p_g < q$.

■

Proof of Equilibrium 3. I show that there is no profitable deviation for g , and appeal to symmetry to complete the argument for m . Given that consumers do not switch when s is small, and given that they click on links at the first site in declining order of quality, g 's profits can be written as

$$(6) \quad \pi_g = x^*(1 - p_g)b$$

when $p_g > q$,

$$(7) \quad \pi_g = x^*b$$

when $p_g \in [S, q)$, and $\pi_g = x^* [(1 - \lambda)(1 - q) + \lambda]b$ when $p_g = q$, where a fraction λ of consumers break indifference in favour of A-links.

It is immediately apparent that any equilibrium with some $p_i = q$ must have $\lambda = 1$ —otherwise i could do better by reducing its quality. With $\lambda = 1$, profit with $p_g = q$ collapses to (7). Since, by Lemma 3, profit from any $p_g < q$ is increasing in p_g it suffices to consider deviations to $p_g > q$. The quasi-reaction function that returns the optimal $p_g > q$ is found from the first order condition for (6):

$$(8) \quad p_g^* = \frac{2q^2 - 2q - p_m(1 - q)^2 + \sqrt{(1 - q)^2(p_m^2(1 - q)^2 + q(1 - q) + p_m(1 + q - 2q^2))}}{(1 - q)^2}.$$

Substituting $p_m = q$ into (8) gives the p_g that maximises (6):

$$(9) \quad p_g^*(q) = \frac{4q^2 - 3q - q^3 + Z}{(q - 1)^2},$$

where $Z = \sqrt{(q - 1)^2 q (2 + 3q - 4q^2 + q^3)}$. For any $q \geq (1/3)(3 - \sqrt{6}) \approx 0.1835$, $p_g^*(q) \leq q$. Moreover, twice differentiating (6) with respect to p_g yields the following:

$$\left. \frac{\partial^2 \pi_g}{(\partial p_g)^2} \right|_{p_m=q} = \frac{2q(2 + q - 7q^2 + 5q^3 - q^4)}{(p_g(q - 1) + (q - 3)q)^3} b,$$

which is negative for $q \in (0, 1)$ so that profits are concave—this implies that profits are decreasing above q when $p_g^*(q) \leq q$. Note also that, for any p_m , (6) is less than (7) when both are evaluated at $p_g = q$. This implies that when $q \geq (1/3)(3 - \sqrt{6})$ deviations to some $p_g > q$ are not profitable.

Suppose $q < (1/3)(3 - \sqrt{6})$. If $p^{\max} \in (q, p_g^*(q))$, concavity of (6) implies that the best *feasible* $p_g > q$ is $p_g = p^{\max}$. Deviation to such a p_g is non profitable when

$$\frac{p^{\max} + q - p^{\max} q}{p^{\max} + 3q - p^{\max} q - q^2} (1 - p^{\max}) b \leq \frac{2q - q^2}{4q - 2q^2} b = \frac{b}{2},$$

i.e. when q is no less than

$$(10) \quad \frac{1}{2} \left[1 + 3p^{\max} - 2(p^{\max})^2 - \sqrt{1 + 2p^{\max} + 13(p^{\max})^2 - 12(p^{\max})^3 + 4(p^{\max})^4} \right].$$

Therefore, when $p^{\max} < p_g^*(q)$, q is given by (10). If $p^{\max} \geq p_g^*(q)$, substituting (9) into (6) yields g 's deviation profits thus: $\pi_g = (2q - 3q^2 + q^3 - Z)(3q^2 - q^3 - q - 1 + Z)b/Z(1 - q)^2$, which is decreasing in q over the relevant range. The values of $q \in [0, 1]$ that

equate this with compliance profits ($b/2$) are 0.08356 and 0.616981. The second root can be ignored since it is greater than 0.1835. Thus, for $p^{\max} \geq p_g^*(q)$, deviation is non profitable for any q above $\underline{q} = 0.08356$. Figure 2 summarises—a $p_g = p_m = q$ equilibrium can be admitted anywhere above the \underline{q} line. ■

Proof of Equilibrium 4. I show that there is no profitable deviation for g , and appeal to symmetry to complete the argument for m . Begin by assuming that p^{\max} is a non-binding constraint on the search engine's quality choice. Taking the derivative of (6) with respect to p_g , yields the quasi-reaction function (8), which is valid for $p_g > q$. The corresponding function for m is symmetric. Differentiating g 's quasi-reaction function with respect to p_m , and substituting $p_m = 1$ yields $\partial p_g^* / \partial p_m|_{p_m=1} = 0.06066$. Now, taking the second derivative of p_g^* gives

$$\frac{\partial^2 p_g^*}{(\partial p_m)^2} = - \frac{(1-q)^4}{4[(1-q)^2((p_m)^2(1-q)^2 + q(1+q) + py(1+q-2q^2))]^{3/2}} < 0.$$

Since the second derivative is negative, and the first derivative is positive at $p_m = 1$, the first derivative of p_g^* must be positive for all $p_m \in [0, 1]$. That the quasi-reaction functions are symmetric, increasing and concave implies that there can be at most two points of intersection, and that both of these must have $p_m = p_g$.

Imposing $p_m = p_g$ for symmetry and solving the quasi-reaction function gives (5) and $p_g = p_m = q/(q-1)$. The second solution is non-positive. Since these are the only two points of intersection of the two $p > q$ quasi-reaction functions, the only possible equilibrium behaviour (for non-binding p^{\max}) in which $p_g, p_m > q$ is given by (5).

By construction, when p_m plays according to (5), no $p_g > q$ can yield a higher profit for g than will compliance with (5). Moreover, by Lemma 3, any deviation to $p'_g < q$ is less profitable than some $p''_g \in (p'_g, q)$. It suffices, then, to show that the limit of (7) as $p_g \rightarrow q$ can not be higher than the profit from compliance with the proposed equilibrium.

Substituting (5) into (6) gives profits for compliance with the candidate equilibrium thus:

$$(11) \quad \pi_g = \pi_m = \frac{1}{3-3q} b.$$

Substituting $p_g = q$, and $p_m = (1-3q)/(3-3q)$ into (7) gives an expression for the maximal deviation profits:

$$(12) \quad \pi_g = \frac{3q^2 - 6q}{3q^2 - 6q - 1} b.$$

Equating (11) and (12) yields the cubic $-9q^3 + 24q^2 - 12q + 1 = 0$. This has two roots in the interval $[0, 1]$, namely $\{0.1042, 0.5228\}$.

For $0 < q \leq 0.1042$ profits from compliance exceed those from deviation; for $0.1042 < q < 0.5228$ deviation appears strictly profitable, and for $0.5228 \leq q$ compliance is again optimal. However, for $q \geq 1/3(3 - \sqrt{6}) \approx 0.1835$, (5) demands a $p_g \leq q$. Thus, (5) constitutes a valid equilibrium strategy only when $0 < q \leq 0.1042$ —for $p^{\max} \geq (1 - 3q)/(3 - 3q)$, $\bar{q} = 0.1042$. Since (6) is concave and (7) always exceeds (6) for p_g sufficiently close to q , there can be no equilibrium with $p_g, p_m > q$ when (5) is less than q .

The above implies that whenever $p^{\max} \geq (1 - 3q)/(3 - 3q)$, $p_g = p_m = (1 - 3q)/(3 - 3q)$ is the unique equilibrium with $p_g, p_m > q$. When $p^{\max} < (1 - 3q)/(3 - 3q)$, such an equilibrium is no longer possible. That the unconstrained quasi-reaction function given in (8) is increasing, and that (6) is concave jointly imply that any intersection in the interval $[0, 1]$ of the quasi-reaction functions constrained by $p^{\max} < (1 - 3q)/(3 - 3q)$ must occur at $p_g = p_m = p^{\max}$. Profits when both search engines play in this fashion are $\pi_i = (1 - p^{\max})b/2$. Deviating to q when $p_{-i} = p^{\max}$ yields a profit for i of $\pi_i = b(2 - q)q/[p^{\max}(1 - q) + (3 - q)q]$. Thus, such a deviation is non-profitable when

$$(13) \quad q \leq \frac{1 + 4p^{\max} - (p^{\max})^2 - \sqrt{1 + 4p^{\max} + 14(p^{\max})^2 - 4(p^{\max})^3 + (p^{\max})^4}}{2(1 + p^{\max})},$$

and (13) characterises the value of \bar{q} when $p^{\max} < (1 - 3q)/(3 - 3q)$. ■

REFERENCES

- AGARWAL, A., K. HOSANAGAR, AND M. D. SMITH (2011): “Sponsored Search: Do Organic Results help or hurt the Performance and under what conditions?,” in *ICIS 2011 Proceedings*, Shanghai, China.
- ANDERSON, S. P., AND S. COATE (2005): “Market Provision of Broadcasting: A Welfare Analysis,” *Review of Economic Studies*, 72(4), 947–972.
- ARMSTRONG, M., J. VICKERS, AND J. ZHOU (2009): “Prominence and Consumer Search,” *RAND Journal of Economics*, 40(2), 209–233.
- ARNOLD, M., E. DARMON, AND T. PENARD (2012): “To Sponsor or Not to Sponsor: Sponsored Search Auctions with Organic Links,” Working Paper.
- ATHEY, S., AND G. ELLISON (2011): “Position Auctions with Consumer Search,” *Quarterly Journal of Economics*, 126(3), 1213–1270.

- BAGWELL, K. (2007): "The Economic Analysis of Advertising," in *Handbook of Industrial Organization*, ed. by M. Armstrong, and R. Porter, vol. 3, pp. 1701–1844. Elsevier, Amsterdam, The Netherlands.
- BERMAN, R., AND Z. KATONA (forthcoming): "The Role of Search Engine Optimization in Search Rankings," *Marketing Science*.
- CHEN, J., D. LIU, AND A. B. WHINSTON (2007): "Competing Keyword Auctions," Working Paper.
- CHEN, Y., AND C. HE (2011): "Paid Placement: Advertising and Search on the Internet," *Economic Journal*, 121(556), F309–F328.
- DE CORNIÈRE, A. (2011): "Search Advertising," Working Paper.
- DE CORNIÈRE, A., AND G. TAYLOR (2012): "Integration and Search Engine Bias," Working Paper.
- EDELMAN, B., M. OSTROVSKY, AND M. SCHWARZ (2007): "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97(1), 242–259.
- ELIAZ, K., AND R. SPIEGLER (2011): "A Simple Model of Search Engine Pricing," *Economic Journal*, 121(556), F329–F339.
- ELLMAN, M., AND F. GERMANO (2009): "What do the Papers Sell? A Model of Advertising and Media Bias," *Economic Journal*, 119(537), 680–704.
- GABSZEWICZ, J. J., D. LAUSSEL, AND N. SONNAC (2001): "Press Advertising and the Ascent of the Pensée Unique," *European Economic Review*, 45(4–6), 641–651.
- GANDAL, N. (2001): "The Dynamics of Competition in the Internet Search Engine Market," *International Journal of Industrial Organization*, 19(7), 1103–1117.
- HAGIU, A., AND B. JULLIEN (2011): "Why Do Intermediaries Divert Search?," *RAND Journal of Economics*, 42(2), 337–362.
- JANSEN, B. J., M. ZHANG, AND C. D. SCHULTZ (2009): "Brand and its Effect on User Perceptions of Search Engine Performance," *Journal of the American Society for Information Science and Technology*, 60(8), 1572–1595.
- KATONA, Z., AND M. SARVARY (2010): "The race for Sponsored Links: Bidding Patterns for Search Advertising," *Marketing Science*, 29(2), 199–215.

- POLLOCK, R. (2010): “Is Google the Next Microsoft: Competition, Welfare and Regulation in Online Search,” *Review of Network Economics*, 9(4), 1–29.
- TAYLOR, G. (2012): “Attention Retention: Targeted Advertising and the Ex Post Role of Media Content,” Working Paper.
- TELANG, R., U. RAJAN, AND T. MUKHOPADHYAY (2004): “The Market Structure for Internet Search Engines,” *Journal of Management Information Systems*, 21(2), 137–160.
- VARIAN, H. R. (2007): “Position auctions,” *International Journal of Industrial Organization*, 25(6), 1163–1178.
- WHITE, A. (forthcoming): “Search Engines: Left Side Quality versus Right Side Profits,” *International Journal of Industrial Organization*.
- XU, L., J. CHEN, AND A. WHINSTON (2012): “Effects of the Presence of Organic Listing in Search Advertising,” *Information Systems Research*, 23(4), 1284–1302.
- YANG, S., AND A. GHOSE (2010): “Analyzing the Relationship Between Organic and Sponsored Search Advertising: Positive, Negative, or Zero Interdependence?,” *Marketing Science*, 29(4), 602–623.