

DISCUSSION PAPER SERIES

DP17401

Market Effects of Sponsored Search Auctions

Massimo Motta and Antonio Penta

INDUSTRIAL ORGANIZATION

CEPR

Market Effects of Sponsored Search Auctions

Massimo Motta and Antonio Penta

Discussion Paper DP17401

Published 21 June 2022

Submitted 20 June 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Industrial Organization

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Massimo Motta and Antonio Penta

Market Effects of Sponsored Search Auctions

Abstract

We investigate the market effects of brand search advertising, within a model where two firms simultaneously choose the price of their (differentiated) product and the bids for the advertising auction which is triggered by own and rival's brand keywords search; and where there exist sophisticated/attentive consumers (who look for any available information on their screen) and naive/inattentive consumers (who only look at the top link of their screen), both aware of either brand's characteristics and price. Relative to a benchmark where only organic search exists, in any symmetric equilibrium each firm wins its own brand auction, and advertising has detrimental effects on welfare: (i) the sponsored link crowds out the rival's organic link, thus reducing competition and choice, and leading to price increases; (ii) the payment of the rival's bid (may) raise marginal cost, also contributing to raise market prices. Under extreme asymmetry (there is an incumbent and an unknown new entrant), we do find that the market effect of brand bidding might be beneficial, if the search engine does not list the entrant's link in organic search, and the share of the sophisticated consumers in the economy is large enough for an equilibrium in which the entrant wins the advertising auction on the search for the incumbent's brand to exist.

JEL Classification: N/A

Keywords: online advertising, auctions, oligopoly, Search Engines, brands, Horizontal agreements

Massimo Motta - massimo.motta@upf.edu

ICREA-Universitat Pompeu Fabra and Barcelona School of Economics and CEPR

Antonio Penta - antonio.penta@upf.edu

ICREA-Universitat Pompeu Fabra and Barcelona School of Economics

Acknowledgements

Comments by Francesco Decarolis, Alexandre de Cornière, Bruno Jullien, Michele Polo, Patrick Rey, Anton Sobolev, and participants in the MaCCI Summer Institutes 2021 and 2022, UniBG IO Winter Symposium (2021), MaCCI Annual Conference (2022), BECCLE Competition Policy Conference, and seminars at the Toulouse School of Economics, Joint Vienna Economic Seminar, UPF Internal Micro workshop, Università Bocconi are gratefully acknowledged. Thanks also to Pia Ennuschat for outstanding research assistance and for her comments. Massimo Motta gratefully acknowledges financial aid from the Spanish Agencia Estatal de Investigación (AEI) and FEDER (project ECO2016-76998-P) and from 'Ayudas Fundación BBVA a Equipos de Investigación Científica 2019' (project on "Digital platforms: effects and policy implications"). Antonio Penta acknowledges the financial support of the European Research Council (ERC), ERC Starting Grant #759424. The authors also acknowledge the financial support of the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S). None of the authors is or has been doing consulting work on cases related to this topic.

Market Effects of Sponsored Search Auctions*

PRELIMINARY AND INCOMPLETE

Massimo Motta[†]

Antonio Penta[‡]

June 20, 2022

Abstract

We investigate the market effects of brand search advertising, within a model where two firms simultaneously choose the price of their (differentiated) product and the bids for the advertising auction which is triggered by own and rival’s brand keywords search; and where there exist sophisticated/attentive consumers (who look for any available information on their screen) and naive/inattentive consumers (who only look at the top link of their screen), both aware of either brand’s characteristics and price. Relative to a benchmark where only organic search exists, in any symmetric equilibrium each firm wins its own brand auction, and advertising has detrimental effects on welfare: (i) the sponsored link crowds out the rival’s organic link, thus reducing competition and choice, and leading to price increases; (ii) the payment of the rival’s bid (may) raise marginal cost, also contributing to raise market prices. Under extreme asymmetry (there is an incumbent and an unknown new entrant), we do find that the market effect of brand bidding might be beneficial, if the search engine does not list the entrant’s link in organic search, and the share of the sophisticated consumers in the economy is large enough for an equilibrium in which the entrant wins the advertising auction on the search for the incumbent’s brand to exist.

Keywords: Digital advertising; auctions; oligopoly; search engines; brands; horizontal agreements.

*Comments by Francesco Decarolis, Alexandre de Cornière, Bruno Jullien, Michele Polo, Patrick Rey, Anton Sobolev, and participants in the MaCCI Summer Institutes 2021 and 2022, UniBg IO Winter Symposium (2021), MaCCI Annual Conference (2022), BECCLE Competition Policy Conference, and seminars at the Toulouse School of Economics, Joint Vienna Economic Seminar, UPF Internal Micro workshop, Università Bocconi are gratefully acknowledged. Thanks also to Pia Ennuschat for outstanding research assistance and for her comments. Massimo Motta gratefully acknowledges financial aid from the Spanish Agencia Estatal de Investigación (AEI) and FEDER (project ECO2016-76998-P) and from “Ayudas Fundación BBVA a Equipos de Investigación Científica 2019” (project on “Digital platforms: effects and policy implications”). Antonio Penta acknowledges the financial support of the European Research Council (ERC), ERC Starting Grant #759424. The authors also acknowledge the financial support of the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S). None of the authors is or has been doing consulting work on cases related to this topic.

[†]ICREA, Universitat Pompeu Fabra and BSE, Spain; e-mail: massimo.motta@upf.edu

[‡]ICREA, Universitat Pompeu Fabra and BSE, Spain, and TSE, France; e-mail: antonio.penta@upf.edu

1 Introduction

Online advertising has grown dramatically in the last decade: it has overtaken non-digital advertising, and in 2020 alone it has exceeded US\$ 150 Billion in the US and GBP 14 Billion in the UK.¹ Roughly half of this market consists of *search advertising*, in which sponsored ads appear in response to search queries, for instance on the search engine results pages (e.g., Google, Bing-Yahoo!, etc.), or on selling platforms (e.g., Amazon, AppStore, etc.), and about 40% consists of display advertising, which is showed when a user is viewing content. (The rest of the market is due to online classified advertising.) Within search advertising, there is one particular segment that has special significance, both for its sheer size and for its role in the policy debate. We are alluding to the so called *brand advertising*, which consists of sponsored ads triggered by the search of a brand term, and which represents the main focus of this paper (we shall consider the case of search for generic products, rather than brands, in an extension). For example, when a consumer searches for a firm’s brand or a branded product, the search engine typically auctions off some links on the search engine result page (SERP) to the potential advertisers, which may thus be competing — by bidding in the sponsored search auctions — for the adjudication of a sponsored link that appears on the consumer’s SERP.

One can get a sense of the importance of brand advertising from a variety of sources.² For instance, while neither Google nor Bing publish data about the share of brand advertising, a glance at the list of the top 100 Google searches in the US reveals that many of them are for brand keywords.³ Analysis of searches in the AppStore also points to brand queries being significant, accounting for over 40% of the AppStore traffic.⁴ Furthermore, while all search advertising is characterized for being closer to purchasing decisions than display advertising (whose role is primarily to create brand awareness), the special case of *brand search* is arguably even closer to an actual transaction: when a user enters a brand keyword in her search bar, she likely signals a concrete interest in purchasing that brand. Narayanan and Kalyanam (2015: 393), e.g., report evidence that suggests that brand search is the *last* step in the search process towards a purchase.

Brand advertising is also topical, and potentially controversial, from a policy perspective. Until 2010, bidding on brands had unclear legal status: a 2005 judgment of a French court, for instance, considered the practice of competing for advertisement space on SERP for others’ brands to be illegal, due to a possible infringement of trademark protection. Then, the 2010 *LVMH v. Google* judgment of the Court of Justice of the European Union reversed that earlier judgment. US courts have also decided in a similar direction. Since then, Google itself has encouraged brand advertising, by advising firms to bid on brand searches and by offering instructions on how to optimise these activities.⁵ Marketing practitioners, in turn, have been actively promoting these practices, as any search enquiry for terms such as ‘brand

¹See, respectively, <https://www.pewresearch.org/journalism/chart/sotnm-digital-and-non-digital-advertising-revenue/>, and Competition and Markets Authority, henceforth CMA (2021).

²See also Simonov and Hill (2021).

³See <https://ahrefs.com/blog/top-google-searches/>.

⁴See <https://www.apptweak.com/en/aso-blog/brand-vs-generic-breakdown-in-app-store-search>.

⁵See, e.g., <https://support.google.com/analytics/answer/2531578?hl=en>

advertising’ would show.

However, a recent US judgment may again question the legality of this form of digital advertising. In 2018, the Federal Trade Commission ruled that the agreements between contact lens retailer *1-800 Contacts* and its competitors to refrain from bidding on brand keyword search terms for internet advertisements, violated antitrust law.⁶ However, in June 2021, the U.S. Court of Appeals for the Second Circuit reversed the FTC ruling, on the grounds that there was no evidence of anticompetitive price increases resulting from the restricted competition in online ad auctions.

The unclear (or, at the very least, unstable) legal status of brand advertising reflects an incomplete understanding of its economic relevance and consequences, particularly from the viewpoint of firms’ competition and consumers’ welfare.

Filling this gap is precisely the goal of this paper. Specifically, we aim to explore how the presence of sponsored space on *brand* search result pages, and the auctions that are used to determine the adjudication of such space, affect product market outcomes. A priori, there may be good reasons to believe that sponsored search advertising is procompetitive. For instance, by lowering consumers’ search costs and conveying information on the relevance of the products, it may foster competition and lead to lower prices (e.g., Athey and Ellison (2011), Chen and He (2011)). To the extent that it may give visibility to firms which otherwise would not be considered by consumers, it could again promote competition. However, search advertising may also add to firms’ incremental costs, thereby creating upward pressure on prices. And if a brand search auction was won by the owner of the brand, its sponsored link(s) may have a *crowding out effect* on its competitors, by making the organic links of the rivals less visible, which decreases rather than foster competition.

Our baseline model is the simplest we could devise that distills the fundamental logic of these effects and is still rich enough to account for the possible tensions that we just discussed. We consider two firms, i and j , who face a continuum of consumers who are initially aware of the price and brand of either firm i or firm j . (In an extension, we consider the case where there also exist consumers who are fully informed about both brands.) Depending on the price set by the firm whose brand they know, consumers decide whether to search for that brand term or not. Each SERP always contains exactly two links, which may be both organic (as in the no-ads benchmark), or one sponsored and one organic (as in the sponsored ads treatment). Besides being possibly aware of different brands, consumers are also heterogeneous in their clicking behavior. Specifically, we assume that consumers may be either *naive*, who only consider the link at the top of the page, or *sophisticated*, who consider the links of both firms (if available). In the *no-ads benchmark*, the organic link of the brand searched for will appear first, followed by the rival brand’s link (in an extension, we consider a benchmark where both organic links belong to the brand searched for). In the *sponsored ads treatment*, the top link is for sale, and it will be assigned through a second-price auction, whereas the link in the second position is the organic link of the brand that

⁶1-800 CONTACTS, INC. v. FEDERAL TRADE COMMISSION, UNITED STATES COURT OF APPEALS FOR THE SECOND CIRCUIT, 11 June 2021, Docket No. 18-3848

was searched for.⁷ Thus, compared to the no-ads benchmark, the presence of a sponsored link displaces the second organic link. Whether this effectively results in a crowding out effect depends on the identity of the auction winner: if the firm whose brand was searched wins the auction, then the other link is crowded out; otherwise, the order of the links is inverted compared with the no-ads benchmark, and the winner of the auction may compete for the sophisticated consumers, as well as try to poach the naive ones, who are captive of the sponsored link. Hence, understanding who wins the auction is crucial to determine the overall competitiveness of the market. But this in turn is ultimately related to the optimal bidding strategies, which depend on how firms value the potential “gate-keeping” role of winning the auction on their own brand, versus how the rival firm values the potential to access a pool of consumers who were not initially aware of its brand.

In the sponsored ads case, firms simultaneously set their price and the two bids, respectively on their own brand keyword and on the rival one. In a model with general demand functions and symmetric firms (that is, half of the population is aware of brand 1 and the other half of brand 2; the proportion of naive and sophisticated aware of each brand is the same across brands; and firms have the same production costs) we show that at a symmetric equilibrium with equal prices and bids, each firm wins the auction for its own brand keyword, and that prices are higher than at the benchmark where only organic links are allowed, with overall a negative impact on consumers’ welfare. Intuitively, this is because of the following effects, both exercising upward pressure on prices.

The first effect is due to the *crowding out* that winning one’s own brand auction exercises on the rival’s link: If firm i wins the auction on its own brand, then all consumers who were initially only aware of i , and who searched for its brand, remain captive of that firm. Specifically, even the sophisticated consumers, who would see the organic links of both brands in the no-ads benchmark, now see only two links of firm i (one sponsored, one organic). Thus, they cannot compare offerings, and they are monopolised by i . In other words, instead of having to compete for both types of sophisticated consumers (that is, those aware of i and those aware of j) as in the no-ads benchmark, in the setting with sponsored auctions each firm ends up monopolising all those consumers who are initially aware of its brand.

The second effect is a *marginal cost* effect, which is directly due to the price that each winning firm pays in the auction. Since the payment to the search engine depends on the number of searches and clicks that a firm receives through its sponsored link, and both searches and clicks do in turn depend on the firm’s own pricing decision (consumers decide to search and click based on their information about the price), then the price that firms pay for the auction effectively adds to their marginal costs, causing further upward pressure on prices, everything else constant.

Finally, there is a further negative effect on consumer’s welfare, which is independent of the previous two considerations (i.e., even if prices remained the same as at the benchmark

⁷Note that one could interpret the no-ads benchmark as the result of a policy action – for instance, a law which interprets bidding on brands as infringing intellectual property rights protection; or a non-brand bidding agreement between rival sellers; or of extreme concentration among online advertising intermediaries, which buy sponsored ads on behalf of rival sellers and stop bidding on each other’s brand, in the spirit of Decarolis and Rovigatti (2021).

— which they do not), which we may call *reduced variety* effect: in the sponsored ad case, in particular, the sophisticated consumers only see the brand they had been informed about, which may be less preferred (or more distant) than the rival one, which they would have discovered and maybe purchased from in the no-ads benchmark in which the SERP returns organic links of both firms.

We consider several variations of the baseline model above, and show that the main results that we just described are robust to several changes in the assumptions. For instance, they continue to hold when there exist consumers who are always fully informed about both brands, and whose purchase decisions are entirely unaffected by the existence of online ads; or when SERPs contain more than just two links (so that winning one’s own brand auction does not completely crowd out the rival firm); or allowing for richer forms of heterogeneity in consumers’ behavior.

Next, we turn to the analysis of asymmetric equilibria. Specifically, within specific parametric applications, first we show that asymmetric equilibria in which both brand auctions are won by the same firm may exist even in an otherwise perfectly symmetric setting. While the details of the way the crowding out effect displays itself differ from the case of symmetric equilibria, the conclusion that sponsored search have a negative impact on consumers’ welfare holds in this case too.

Finally, we consider two cases of asymmetric environments, which are especially important because a priori they are expected to enhance the procompetitive potential of sponsored ads.

First, we consider the case where an entrant is as efficient as an incumbent, but it is unknown to consumers (and, as a conservative assumption, to the search engine). In this case, the auction for the incumbent’s brand may in principle allow the entrant to gain visibility with consumers and increase competition in the market. Indeed, we find that if the share of sophisticated consumers is high enough, there may exist equilibria in which the entrant wins the auction. In this case, consumers are better off than in the benchmark without auctions: the effect of enhanced competition facilitated by the ads outweighs the upward pricing pressure created by the incremental costs introduced by the payment of the sponsored search ads. Obviously, equilibria in which the incumbent wins the auction unambiguously induce a lower consumer welfare than in the no-ads benchmark.

Second, we consider the case of two firms that are equally known to consumers, and symmetric in all respects, except that one is more efficient than the other. One might expect that the more efficient firm is more likely to win the auction for *both* brand keywords, and that the resulting higher market share of the low-cost firm could benefit consumers. Perhaps surprisingly, it turns out that even when the low-cost firm does win both brand auctions, consumer surplus does not increase relative to the no-ads benchmark. Intuitively, this is because the higher consumer surplus obtained by the consumers who would otherwise just be aware of the high-cost firm is outweighed by the detrimental effect on the consumers who are initially aware of the low-cost firm: due to the crowding out of the high-cost firm’s organic link, and because of the additional marginal cost effect of auctions, these consumers end up paying a higher price, which offset the positive effects that accrue to other types of consumers.

1.1 Related Literature

Theoretical Literature: Given the importance of digital advertising, it is no surprise that there is a recent and growing literature on this topic. In what follows we clarify our relationship with it. Edelman et al. (2007) and Varian (2007) study the properties of position auctions, but they abstract from the ‘objects’ of the auctions being ads, and do not look at their market effects. Athey and Ellison (2011) enrich the setting by introducing consumers, who incur costs when clicking on ads, and act rationally: they decide how many ads to click on and in what order. Ads in their model are efficient:⁸ they help consumers in their search.

Chen and He (2011) are probably the first to incorporate firms’ price decisions in a model with position auctions. In their paper, consumers also search sequentially, and sponsored ads, whose position is determined by a second-price auction, are efficient since they convey information on the relevance of products. As a result, sponsored ads lead to higher output. Note that in their paper firms first choose bids and then prices, so at equilibrium prices are always at their ‘monopoly’ level. Also, the auction does not add to firms’ marginal costs, since consumer search decision is unaffected by firms’ prices.⁹

In Armstrong and Zhou (2011), advertising in a Price Comparison Website (PCW) is also procompetitive: the PCW gives prominence to the firm with lower price, consumers click on it and discover their position on the Hotelling line (further search is costly).

While in previous papers digital advertising tends to be beneficial, in De Cornière (2016) advertising has ambiguous price effects. Firms pay a pay-per-click fee (there is no auction in his model) for targeted search advertising; consumers enter keywords and search sequentially at random through the links that appear. Competition intensifies, because targeted ads reduce search costs (consumers know firms are drawn from better pool). But the advertising fee adds to the firms’ marginal cost.

Although not due to advertising, in Ronayne (2021) internet visibility also translates into higher marginal costs. In his paper, PCWs allow consumers to compare price offerings. However, their procompetitive effects is outweighed by the price-increasing effect due to producers passing on to consumers the cost of the fees imposed by the PCWs.

In Haan and Moraga-González (2011) – which like the previous two papers does not consider position auctions – a firm invests in (traditional) advertising to become more salient, and hence increase the likelihood that consumers search it first. In turn, search costs imply that firms have market power over consumers who visit it. In their paper, the higher search costs the larger the firms’ incentives to invest in advertising, resulting in a positive correlation between search costs, advertising and prices.¹⁰

Unlike the previous literature, our paper is not based on a standard search model. Our

⁸However, the click-weighted auction they analyse in an extension does not necessary result in the right selection of ads, and it may lead to other inefficiencies.

⁹Xu et al. (2012) also study search advertising within a model where firms are asymmetric, consumers discover their preferences upon searching, and the game is a sequential one where firms first bid and then choose prices.

¹⁰Search costs have an ambiguous effect on profits: on the one hand, they increase prices at equilibrium. On the other hand, more advertising creates a dissipation effect. And firms find themselves in a prisoner’s dilemma: if one advertised less than the others, it would be less likely to be visited first by consumers; but each firm has a private incentive to advertise and have a more salient position.

consumers are not searching to learn product characteristics (or their preferences relative to those characteristics) or prices. When they search, they are already aware of one product’s characteristics and price, they are interested in buying it, and they search for that particular brand online in order to purchase it. But in the process, they *might* encounter another brand they were unaware of and consider buying it. Still, there are similarities in the two approaches, and one may also interpret our ‘sophisticated’ consumers as having zero costs of “searching” (or assessing) among the visible links; and our ‘naive’ consumers as having a strong prominence bias (and zero cost of searching/assessing the top link) as well as arbitrarily large costs of “searching” beyond the top link.

Furthermore, the fact that consumers’ search decisions depend in our case on the prices set by the firms imply that the number of keyword auctions triggered by searches, and hence the payments for the sponsored ads, depend on prices and add to marginal costs, thereby tending to push prices upwards. In models where searches and clicks do not depend on prices (e.g., because all consumers always search) the advertising auctions do not affect marginal costs, thereby making it less likely to have anti-competitive effects.

Another possible element of interest is the relationship between placement (or prominence) and prices. In Haan and Moraga-González (2011), Zhou (2010) and Anderson and Renault (2021), firms which are searched first have *lower* prices at equilibrium. The driving force is that a consumer who has walked away from the firm she has visited will have fewer options left, or she signals she did not like the product she had seen, and she is thus less price sensitive.

In Arbatskaya (2007), the prices fall in search order (which is exogenously given, unlike Haan and Moraga-González (2011) where it depends on who advertises most): a consumer who has walked away from a firm she has visited reveals that she has low search costs, thus giving the next firm an incentive to charge a lower price.

In our model if there is a symmetric equilibrium then the question is immaterial: each firm wins its own auction and only its brand will be visible, so at equilibrium each consumer will see only one firm. But whenever asymmetric equilibria exist (whether in a symmetric or asymmetric environment), then the firm which wins the auctions is more prominent and will also have a higher price.

Empirical literature: A number of empirical papers (both in the economics and marketing literature) have studied the effectiveness of brand advertising. Blake et al. (2015) find that in a controlled experiment, when eBay does not bid on its own brand keywords, it loses only about 5% of the traffic.

Coviello et al. (2017) and Golden and Horton (2021) empirically analyse case studies, and find very different results on competitors’ ability to steal a rival brand’s traffic. Simonov et al. (2018) and Simonov and Hill (2021) analyse the same issue within a large-scale quasi-experimental set-up on the Bing search engine, involving (hundreds of) the most searched brands. Among their results closer to the scope of our paper, they find that there is traffic stealing when the focal brand (i.e., the firm which owns the brand) does not advertise. (But this does not imply a high conversion rate, since most links turn out to be “quick-backs”,

namely users go back to the search page immediately after clicking on the rival’s winning link.) Relatedly, the focal brand needs to win the brand auction in order not to lose 50-66% of its traffic,¹¹ thus resulting in a large crowding out effect of organic traffic by paid links. This seems to suggest that obtaining the top position is a necessary condition to win customers.

As for the relationship between sponsored ads in search engines results pages and product prices – the central point of our paper – it has been suggested in the specialised press that the prices of products appearing in Google’s prominent slots are more expensive than elsewhere in the web,¹² which may be consistent with our assumption that not all consumers necessarily make accurate and time-consuming comparisons when shopping online, as well as with our result that when several products are listed, the more expensive get the top positions.

Anecdotal evidence aside, to the best of our knowledge, Sviták et al. (2021) is the only empirical paper which deals with the relationship between sponsored search ads and market prices. They look at the effects of ‘Non-brand bidding agreements’ (NBBAs) in the hotel sector in the Netherlands. These NBBAs were meant to prevent Online Travel Agencies (OTAs) – namely, platforms such as booking.com and Expedia – from bidding in the auction which is triggered when a user is searching for a certain hotel name. They find that NBBAs increase hotel prices, which would show that whatever savings from ads expenses is outweighed by the impact of lower competition between hotels and OTAs. Although obviously related, the object of their study is not the same as ours, since hotels and OTAs are in a vertical relationship, whereas our competing firms are rivals (a NBBA in our setting would amount to the benchmark where a firm cannot bid on a rival’s brand, not to an agreement between a firm and its retailer, like in Sviták et al., 2021).¹³

1.2 Organization of the paper

The paper continues as follows. Section 2 describes the model, solves for the benchmark case and studies the optimal bids. Section 3 analyses a symmetric environment: it compares the symmetric equilibrium with the benchmark within a general model, and then it studies equilibrium solutions (including asymmetric ones, which may also arise) in specific functional form examples. Section 4 analyses a number of robustness checks by relaxing assumptions and extending the model in several directions. Section 5 analyses asymmetric environments, notably one where an incumbent faces an entrant who is completely unknown to consumers, and another where a firm has lower costs than the other. Section 6 briefly discusses possible policy counterfactuals (including the case where whenever a consumer enters a brand keyword in her search bar, an auction for the generic product — rather than the branded product — is triggered) and it concludes the paper.

¹¹In their sample, the owner of the brand almost always wins the top paid position if it participates in the auction.

¹²See e.g., Waters (2013), who also quotes from Ben Edelman: “The PLAs [Product Listing Ads] do a great job at pushing people to things that aren’t the cheapest.”

¹³Another vertical agreement not to bid on brand has been investigated by the European Commission in *Guess*. Closer to our setting is instead the *1 - 800 Contacts v. FTC* case mentioned earlier.

2 The model

We consider two firms, $i \in \{1, 2\}$, interacting in a market in which they compete for consumers as well as for the adjudication of advertisement space on a search engine result page (SERP), which consumers use to access the firms' websites in order to complete their purchases. In the baseline model, consumers are initially aware of either firm 1 or 2, and given its price (which we assume they observe), they decide whether to search for its brandname or not. (As we will explain shortly, we also allow for heterogeneity in consumers clicking behavior). The SERP always contains two links: In the *no advertisement* benchmark, both links are organic. In the *advertisement* case, the first link is sponsored and the second is organic. The search engine platform is monopolistic, and it assigns advertisement space through an auction. So, the sponsored link will belong to the winner of the auction, and the organic link to the firm whose brand was searched. (Note that there are two search results pages, one for consumers who searched one brand and the other for consumers who searched the other. In the ad-case, each will generate a distinct auction for the corresponding sponsored link.) Depending on the composition of the result page that consumers face, which may (or may not) make them become aware of both firms, consumers decide from which firm to buy, if any, and click on the selected firm's link in order to complete their purchase.

In the baseline model, we assume that firms choose their prices and bids (for the advertisement case) simultaneously, and that both firms' links show up in the no-ads benchmark. Variations to these assumptions, as well as alternative benchmarks, such as the case in which consumers search for the generic product rather than for a specific brand, will be considered in later sections.

Firms: Each of the two firms in the industry, $i \in \{1, 2\}$, produces a good at a constant marginal cost, $c_i \geq 0$. We will denote by $j \in \{1, 2\} \setminus \{i\}$ the firm other than i . In the no-ads case, firms simultaneously choose the price of their good, p_i . When sponsored links are present, instead, they also simultaneously submit bids for each of the two possible result pages, one for their own brand and the other for the competitor's. Firms' strategies will thus consist of a triplet (p_i, b_i^i, b_j^i) , where b_i^i and b_j^i denote firm i 's bids in the ad auction for the brand of firms i and j , respectively. As we will discuss shortly, the outcome of the auctions will determine whether consumers have access to both firms, or only one of the two.

Consumers: There is a unit mass of potential consumers with unit demand: given the posted prices, they decide whether to buy the unit of the good, and from which firm, if they become aware of both. The distribution of preferences in the population is summarized by a collection of demand functions that they induce, conditioning on facing either firm as a monopoly, or a duopoly. Specifically, we let $q^{iM}(p_i)$ denote the set of consumers who would be willing to buy from firm i , if it was a monopolist and set price p_i . Similarly, we let $q^{iD}(p_i, p_j)$ denote the set of buyers who would be willing to buy from firm i if they were facing a duopoly in which i 's price is p_i and j 's price is p_j .

With this in mind, it is natural to assume that $q^{iD}(p_i, p_j) \subseteq q^{iM}(p_i)$ for all p_i and p_j (that is, a consumer is willing to buy from i in a duopoly only if, for the same price, she would

be willing to buy from i in a monopoly), and that $q^{iD}(p_i, p_j) \cap q^{jD}(p_i, p_j) = \emptyset$ for all (p_i, p_j) (that is, no consumer buys from both firms). Moreover, it is also natural to assume that the set of consumers who would be willing to buy from either duopolist is no smaller than the set of consumers who would be willing to buy from either monopolist, at the same prices. That is: $q^{iM}(p_i) \cup q^{jM}(p_j) \subseteq q^{iD}(p_i, p_j) \cup q^{jD}(p_i, p_j)$ for all p_i and p_j . The two set inclusions jointly imply the following:

$$q^{iM}(p_i) \cup q^{jM}(p_j) = q^{iD}(p_i, p_j) \cup q^{jD}(p_i, p_j) \text{ for all } (p_i, p_j). \quad (1)$$

With a slight abuse of notation, we will also let $q^{iM}(p_i)$ and $q^{iD}(p_i, p_j)$ denote the *mass* of the corresponding sets, so that as real functions from the set of possible prices they can be interpreted, respectively, as firm i 's *monopolistic* and *duopolistic* demand functions. The set-inclusions above immediately implies that, in terms of demands, we have $q^{iD}(p_i, p_j) \leq q^{iM}(p_i)$ for each i and (p_i, p_j) . Besides these natural restrictions, we also assume that the demand functions satisfy standard smoothness and monotonicity properties. For ease of reference, we list all the maintained assumptions on the demand functions as follows:

Assumption 1 *The demand functions satisfy the following assumptions:*

1. $(q^{iD}(p_i, p_j), q^{iM}(p_i))_{i=1,2}$ are non-increasing in p_i , non-decreasing in p_j , twice differentiable and such that $p_i \cdot q^{iM}(p_i)$ and $p_i \cdot q^{iD}(p_i, p_j)$ are concave in p_i .
2. $q^{iD}(p_i, p_j) \leq q^{iM}(p_i)$ for all i and (p_i, p_j)

These demand functions can be thought of as being derived from quasi-linear utility functions of the form $U_i = f(\theta, \mathbf{y}_i) - p_i$, where θ describes heterogenous consumers, distributed according to some density function, and \mathbf{y}_i is a vector of product characteristics. Under appropriate assumptions, these utility functions will generate smooth demand functions for the monopolistic and duopolistic case, respectively. In such a model, the property in equation (1) would follow directly from the assumption that a consumer with preference parameter θ buys from i if and only if $U_i(p_i, \theta) > \max\{0, U_j(p_j, \theta)\}$ (with the understanding that $U_j(p_j) = 0$ if the consumer only faces i as a monopolist), and if indifferences are broken with equal probability.

Examples might include a variety of vertical product differentiation models in the spirit of Mussa and Rosen (1978) and Shaked and Sutton (1982) – where, e.g., $U_i(p_i, \theta) = \theta u_i - p_i$, u_i represents the quality of the good and θ is either the consumer's taste for quality or her income – as well as various applications of the Hotelling model, where $U_i = v_i - t \cdot (|\theta - l_i|)^\alpha - p_i$ where v_i is the baseline utility for the consumers, l_i and θ denote, respectively, the location of firm i and of the consumer, and t is the 'transport cost'. In the applications throughout the paper, we will use the latter model with $v = 1$, firms located at $l_1 = 0$ and $l_2 = 1$, $\alpha = 1$ and $\alpha = 2$, and θ distributed uniformly either on the segment $[0, 1]$ or along the real line.

Consumers' Types: At the outset, consumers may *either* be aware of brand 1 or brand 2, and may either be *naive* or *sophisticated*. Hence, the set of consumer *types* is

$T = \{1, 2\} \times \{n, s\}$, with generic element $t = (i, n)$ or $t = (i, s)$ to denote respectively the naive and sophisticated types who are aware of firm i . We use the term ‘ i -types’ to refer to types $t \in \{(i, n), (i, s)\}$, the term ‘naive types’ to refer to types $t \in \{(1, n), (2, n)\}$, and the term ‘sophisticated types’ to refer to types $t \in \{(1, s), (2, s)\}$. As we discuss next, type-1 and type-2 consumers differ in their decision to search, while *naive* and *sophisticated* types differ in their clicking behavior. For each $i \in \{1, 2\}$, we let $s_i \in [0, 1]$ denote the share of consumers who are ex ante aware of firm i ’s brand name ($s_1 + s_2 = 1$), and let η_i denote the fraction of type- i consumers who are *sophisticated*. We shall assume throughout the paper that a consumer aware of brand i is also aware of its product characteristics and price, and that all such types are distributed independently of preferences. Hence, the demand functions discussed above will be maintained with the same proportions within each of the four groups.

Decision to Search: Differently from the approach followed by the economic literature on search, where consumers search in order to discover prices or characteristics of the products on offer, we assume here that consumers search in order to carry out a purchase. They already have information on a certain brand, they are willing to buy it, and they introduce that brand keyword in the search engine in order to find the webpage where they can make their purchase. Arguably, this is consistent with actual users’ brand search activity, which is the last stage of their information acquisition process, and just precedes their purchase decision.

Type- i consumers’ decision to search will generally depend on the price set by the firm they are aware of, p_i . For the time being, we assume that type- i consumers know the price p_i , and decide to search if and only if they are willing to buy from i at that price p_i . Hence, given p_i , the total number of type i -consumers who search i ’s brand is equal to $S_i(p_i) = s_i \cdot q^{iM}(p_i)$ (In Section 3.2 we discuss the case where all type- i consumers search).¹⁴ We shall refer to the function $S_i(p_i)$ as the ‘search function’. Under the maintained Assumption 1, $S_i(p_i)$ is twice differentiable, *weakly* decreasing in p_i , and such that $p_i \cdot S_i(p_i)$ is concave.

Search Results and Clicks: We assume that each search results page contains two links. If no advertisement is sold by the search engine, both links are ‘organic’, otherwise the first link is sponsored and the second organic. In the *no-ads* case, both firms’ links show up on the result page, the first being that of the firm whose brand was searched. (In Section 3.2 we consider the alternative benchmark where only the organic links for the brand that was searched appear on the results page.) In the *advertisement* case, the sponsored link corresponds to the firm that has won the auction, and the organic link belongs to the firm whose brand was searched. Figure 1 illustrates consumers’ screens and shows which links appear in the different cases.

Naive and *sophisticated* consumers differ in their clicking behavior. Specifically, we assume that *naive* consumers always and exclusively click on the *first link at the top of their*

¹⁴Note that we could define the volume of i -type consumers who search as a function of p_i without necessarily assuming that i -type consumers know p_i . For instance, they may each receive a noisy signal of what p_i is, and decide to search if and only if some statistics of their interim beliefs over the true price is low enough. As long as their interim beliefs over p_i are first-order stochastically ranked according to the true price set by the firm (for instance, if the individual consumer’s noisy signal took the form $\hat{p}_i = p_i + \epsilon$, where ϵ is a zero-mean noise), the expected number of i -type consumers who search would be a function of p_i (and, of course, on the parameters of the noise distribution) even though none of the consumers necessarily know the price p_i .

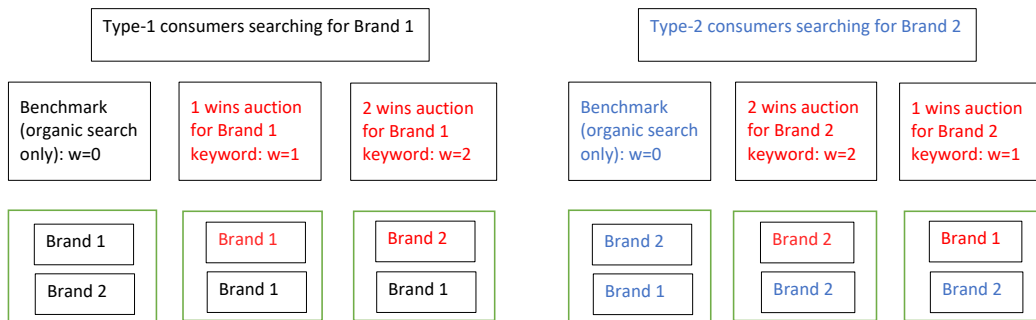


Figure 1: Illustration of the consumer's screen

Illustration of the screen appearing to type-1 consumers (left panel, first three columns) and respectively type-2 consumers (right panel, last three columns) when they search for brand 1 and respectively brand 2, under (i) organic search, (ii) firm i winning own brand auction ($w = i$), and (iii), firm j winning the auction for brand i keyword ($w = j$), for $i, j = 1, 2$ and $i \neq j$.

search result page. This link would correspond to the winner of the auction in the case with sponsored link, or the first organic link in the benchmark with no sponsored links. *Sophisticated* consumers instead always click on the link of the firm they were *not* initially aware of, if available, in order to compare offers; they click on the link of the firm they are aware of, if it is the only available one (in the other cases, they may or may not also click on the other link), or if – after having clicked the other brand's link – they decide that they prefer the brand they had originally been informed about. This implies that, both in the no-ads benchmark and if j wins the auction, then (i, s) consumers first click on the link which belongs to firm j . If instead i wins the auction, so that no link to j is available, then we assume that that they click on the sponsored link, which is the most prominent. (Variations of this baseline model are considered in Section 4, where we vary the assumptions on consumers' clicking behavior, on the number of links displayed, and on consumers' information.)

Note that one could reinterpret the naive consumers as paying attention exclusively to the top link on their search results page, and the sophisticated consumers as paying attention to both links. Also, the assumption that naive consumers only click on one link implies that type- i naive consumers who happen to click on j 's link but for some reason decide not to buy product j (see below), do not go back to i . This can be seen as an extreme form of rational inattention or of choice overload, where naive consumers' attention cost increases sharply after the first (more prominent) link. Alternatively, one may interpret sophisticated consumers as having zero cost of searching/assessing the visible links, and the naive consumers as having a prominence bias and zero search/assessment costs for the top link, but very large costs for searching/assessing any other link.

It goes without saying that these types, with their starkness, are not meant to be taken literally. Rather, they represent stylized versions of consumers with different levels of contestability, which may arise in the real world from different sources (not just cognitive heterogeneity, as suggested by the language of the model, but from a variety of reasons, such as idiosyncratic preference shocks, heterogenous search costs or biases for the familiar brand, etc.). The continuity of the model with respect to the η parameter means that any combina-

tion of less stark behavioral types that might exist in the real world, with different levels of contestability, can be suitably captured by some convex combination of these two archetypes, that represent in a way the ‘extremal points’ on the spectrum of possible consumers’ contestability (with our *sophisticated* types being maximally contestable, and the *naive* not contestable at all).

Formally, we let $w \in \{0, 1, 2\}$ denote the *winner of the auction*, where $w = 0$ refers to the *no-ads* case, and for each type $t \in \{1, 2\} \times \{n, s\}$, each firm $i \in \{1, 2\}$ and each $w \in \{0, 1, 2\}$, we let $CTR_t^{i|w}(p_i, p_j)$ denote the probability (or *click-through rate (CTR)*) that a type- t clicks on i ’s link, given that the ‘winner’ of the auction is w , and given prices p_i and p_j .

From Clicks to Sales, and Conversion Rates: Upon clicking on a link, consumers learn the brand, location and the price set by the firm which controls it (if they didn’t know them before). Then, given the information they have, they decide whether to purchase or not. We assume that sales can only occur through clicks, but of course consumers may click on i ’s link and still not purchase from it (as discussed, for instance, (i, s) consumers may click on j ’s link to learn their price, but then decide to go back to and buy product i .) We thus represent consumers’ purchasing behavior by means of ‘click-to-sales’ (CTS) rates: for each type $t \in \{1, 2\} \times \{n, s\}$ and each firm $i \in \{1, 2\}$, we let $CTS_t^i : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow [0, 1]$ denote the probability that a type- t consumer purchases from firm i , given that she has clicked on its link, as a function of the two prices.

It is convenient to also introduce the overall *conversion rate*, that is the probability that a consumer of type $t \in \{1, 2\} \times \{n, s\}$ ends up buying from $i \in \{1, 2\}$, given that her search led to a result page in which the sponsored link (if any) was posted by $w \in \{0, 1, 2\}$. Since conversion rates are derived from the CTRs and CTSs, they can be thought of as functions $x_t^{i|w} : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow [0, 1]$, as follows:

$$x_t^{i|w}(p_1, p_2) = CTR_t^{i|w}(p_1, p_2) \cdot CTS_t^i(p_1, p_2). \quad (2)$$

The conversion rates of firm i are effectively driven by consumers’ demand, as “filtered” by the CTRs. For instance, consider the case where the (i, s) types see both links on their search results page (that is, if $w = 0$ or $w = j$). Whether organic or sponsored, the sophisticated types compare both firms’ offerings before deciding, and hence the firms compete for these consumers. As a result, conversion rates for firm i in this case are:

$$x_{(i,s)}^{i|0}(p_i, p_j) = x_{(i,s)}^{i|j}(p_i, p_j) = \frac{q^{iD}(p_i, p_j)}{q^{iM}(p_i)} \leq 1. \quad (3)$$

To understand the conversion rates for firm j , instead, it is useful to think of the demands in terms of *sets* of consumers who would be willing to buy from either firm, as we explained at the beginning of Section 2. In this case, among the set of all consumers in $q^{iM}(p_i, p_j)$ who have searched for i , it will be only those in $q^{iM}(p_i, p_j) \cap q^{jD}(p_i, p_j)$ who would buy from j . It is easy to see that, since $q^{iD}(p_i, p_j) \cap q^{jD}(p_i, p_j) = \emptyset$ for all (p_i, p_j) , and as long as condition (1) holds, $q^{iM}(p_i, p_j) \cap q^{jD}(p_i, p_j) = q^{iM}(p_i, p_j) \setminus q^{iD}(p_i, p_j)$. Hence, the *mass* of buyers for

firm j in this situation is proportional to $(q^{iM}(p_i, p_j) - q^{iD}(p_i, p_j))$. Thus, conditioning on the mass of consumers who have searched for i , which is $q^{iM}(p_i, p_j)$, the conversion rate for firm j in this case is:

$$x_{(i,s)}^{j|0}(p_i, p_j) = x_{(i,s)}^{j|j}(p_i, p_j) = \frac{(q^{iM}(p_i) - q^{iD}(p_i, p_j))}{q^{iM}(p_i, p_j)} \leq 1. \quad (4)$$

Consider now the case where type- i consumers, whether sophisticated or naive, can see only brand i because $w = i$ and both a sponsored and an organic link to i appear on type- i 's page. Firm i 's conversion rate will be equal to 1, since all type- i consumers will buy (if they have been searching for product i , they will also be willing to buy it):

$$x_{(i,s)}^{i|i}(p_i, p_j) = x_{(i,n)}^{i|i}(p_i, p_j) = \frac{q^{iM}(p_i)}{q^{iM}(p_i)} = 1, \quad (5)$$

Correspondingly, firm j 's conversion rate will be nil:

$$x_{(i,s)}^{j|i}(p_i, p_j) = x_{(i,n)}^{j|i}(p_i, p_j) = 0. \quad (6)$$

When type- i naive consumers are searching and auctions do not take place ($w = 0$), brand i appears as the top link in organic search and such consumers will always click there. Since they have no surprise (they were willing to buy i at price p_i and this information is confirmed when clicking), they will always buy i . Since they never go beyond the first click, they will never click on j . As a result, the conversion rates for firm i and j will be 1 and 0 respectively:

$$x_{(i,n)}^{i|0}(p_i, p_j) = \frac{q^{iM}(p_i)}{q^{iM}(p_i)} = 1, \quad x_{(i,n)}^{j|0}(p_i, p_j) = 0. \quad (7)$$

Finally, when j wins the auction ($w = j$), naive consumers of type i will click on j 's link, and they will never go back to i . Hence, firm i will not be able to sell to type- i naive consumers, and its conversion rate will be 0. As for j 's conversion rate, it would depend on additional assumptions on naive consumers' behaviour. In the worst case from the viewpoint of firm j , these naive consumers will act as if they were in a duopoly, comparing j 's price with that of firm i , whose memory may inhibit purchasing from j even if i 's link has been displaced by j 's at the top of their search result page. In this case, j 's conversion rates would be effectively the same as if they were sophisticated, and hence, as explained for the case of equation (4), the conversion rate for firm j would be equal to $\frac{(q^{iM}(p_i) - q^{iD}(p_i, p_j))}{q^{iM}(p_i, p_j)}$. In the best case scenario, instead, naive types may 'forget' about i , and buy from j whenever it gives them positive utility. In this case, it would be naive buyers in the set $q^{iM}(p_i, p_j) \cap q^{jM}(p_i, p_j)$ who buy from j . Again, since $q^{iD}(p_i, p_j) \cap q^{jD}(p_i, p_j) = \emptyset$ for all (p_i, p_j) , and under condition (1), we have that $q^{iM}(p_i, p_j) \cap q^{jM}(p_i, p_j) = (q^{iM}(p_i, p_j) \setminus q^{iD}(p_i, p_j)) \cup (q^{jM}(p_i, p_j) \setminus q^{jD}(p_i, p_j))$. Hence, the conversion rate for firm j in this case would be $\frac{(q^{iM}(p_i, p_j) - q^{iD}(p_i, p_j)) + (q^{jM}(p_i, p_j) - q^{jD}(p_i, p_j))}{q^{iM}(p_i, p_j)}$. Overall, the conversion rate over (i, n) types are:

$$x_{(i,n)}^{i|j}(p_i, p_j) = 0; \quad (8)$$

$$x_{(i,n)}^{j|j}(p_i, p_j) \in \left[\frac{(q^{iM}(p_i, p_j) - q^{iD}(p_i, p_j))}{q^{iM}(p_i, p_j)}, \frac{(q^{iM}(p_i, p_j) - q^{iD}(p_i, p_j)) + (q^{jM}(p_i, p_j) - q^{jD}(p_i, p_j))}{q^{iM}(p_i, p_j)} \right]. \quad (9)$$

Figure 2 illustrates the search function and some conversion rates for the Hotelling model where firms are located at l_1 and l_2 and consumers have utility function $U_i = 1 - p_i - t | \theta - l_i |$, with $i = 1, 2$, where consumers' location θ is uniformly distributed on the real line, t is a disutility (or transport cost) parameter and p_i is product i 's price. Consumers who are indifferent between buying from 1 and 2 are denoted by θ_{12} ; θ'_{10} and θ_{10} denote consumers who are indifferent between buying from i and not buying at all. Note that although type-1 and type-2 consumers have identical preferences (there is no correlation between consumers' location and their brand awareness), conditional on a search being made, there is a fundamental asymmetry between the two firms, due to the self-selection of consumers' search decisions. That is because type-1 consumers who search are closer to l_1 , and hence even at equal prices firm 2 obtains a lower share of type-1 searchers than firm 1 (in the figure, firm 2's demand is given by $(\theta_{10} - \theta_{12})$, which is lower than firm 1's demand, which is $(\theta_{12} - \theta'_{10})$). Thus, a *familiarity bias*, that favors the known brand, may arise due to a pure selection effect, independent of other possible behavioral sources.




Figure 2: Illustration: Hotelling model

Illustration of the search function $S_1(p_1)$ and conversion rates $x_{(1,s)}^{k|w}$ and $x_{(1,n)}^{k|w}$ in the Hotelling model with consumers uniformly distributed on the real line and firms located at l_1 and l_2 . Indifferent consumers between 1 and 2 are denoted by θ_{12} , between buying i and not buying by θ'_{10} and θ_{10} . (In this example, $s_1 = 1$).

2.1 The no ads benchmark

We begin our analysis by first considering the case where there are no sponsored ads. This is a useful benchmark from a theoretical perspective, but it may also be relevant from a policy viewpoint, as a possible counterfactual to the current situation where competitive bidding on

brand keywords is allowed.

With *no sponsored links*, the first link that each user finds on the search result page is the link of the brand she has searched for, followed by the link of the other firm. Hence, firm i 's link will be at the top of i -types' result page, and at the bottom for j -types. Naive consumers only click on the top link (i.e., the link of the firm the searched for), and hence they do not learn about the other firm nor about its price; sophisticated consumers instead become aware of both firms' prices and compare them in order to make their purchasing decision. As a consequence, firm i 's demand in this case is:

$$Q_i^0(p_1, p_2) = S_i(p_i) \left[\eta_i \cdot x_{(i,s)}^{i|0}(p_1, p_2) + (1 - \eta_i) \cdot x_{(i,n)}^{i|0}(p_1, p_2) \right] + \eta_j S_j(p_j) \cdot x_{(j,s)}^{i|0}(p_1, p_2).$$

By replacing the conversion rates derived above, this can be written as:

$$Q_i^0(p_1, p_2) = (1 - \eta_i)q^{iM}(p_i) + \eta_i q^{iD}(p_i, p_j) + \eta_j [q^{jM}(p_j) - q^{jD}(p_i, p_j)]. \quad (10)$$

The profit function in turn is

$$\pi_i^0(p_1, p_2) = Q_i^0(p_1, p_2) (p_i - c_i). \quad (11)$$

Since each firm i is a monopolist over the *naive* consumers of type- i , competition between firms is limited to the sophisticated consumers of both types, who end up seeing both links on their search result page. Hence, the optimal pricing will trade off the ability to attract sophisticated consumers, with the revenues lost on the 'captive' naive consumers of type i . Moreover, under the maintained assumptions, $\pi_i^0(p_1, p_2)$ is also differentiable in p_i and concave (strictly so if $(q_i^M)_{i=1,2}$ and $(q_i^D)_{i=1,2}$ are *strictly* decreasing in p_i , or if $c_i > 0$), and hence the Nash equilibria (\bar{p}_1, \bar{p}_2) in the no-ads benchmark are fully characterized by the two firms' first-order conditions at the equilibrium prices, $\frac{\partial \pi_i^0}{\partial p_i}(\bar{p}_i, \bar{p}_j) = 0$. Formally: (\bar{p}_1, \bar{p}_2) is a Nash equilibrium of the *no-ads benchmark* if and only if

$$Q_i^0(\bar{p}_i, \bar{p}_j) + \frac{\partial Q_i^0}{\partial p_i}(\bar{p}_i, \bar{p}_j) (\bar{p}_i - c_i) = 0 \text{ for each } i = 1, 2. \quad (12)$$

Existence of a Nash equilibrium comes from standard results, since firms' payoff functions are continuous and concave in their own price, and since under standard assumptions the range of relevant prices can be compactified. This of course does not imply uniqueness, since multiple pairs of prices may satisfy Condition (12). However, under the maintained assumptions, and if $(q_i^M)_{i=1,2}$ and $(q_i^D)_{i=1,2}$ are *strictly* decreasing in p_i or if $c_i > 0$, the following condition is sufficient for uniqueness: for all i and (p_i, p_j) ,

$$\left| 2 \frac{\partial Q_i^0}{\partial p_i}(p_i, p_j) + \frac{\partial^2 Q_i^0}{\partial^2 p_i}(p_i, p_j) (p_i - c_i) \right| > \left| \frac{\partial Q_i^0}{\partial p_j}(p_i, p_j) + \frac{\partial^2 Q_i^0}{\partial p_i \partial p_j}(p_i, p_j) (p_i - c_i) \right|, \quad (13)$$

In words, this condition requires the concavity in own action of each firm's profit function, to be stronger than the cross derivative with respect to both prices. Intuitively, this limits

the strength of the strategic externalities, i.e. how strongly a firm's optimal price depends on the price set by the other firm (cf. Moulin, 1984, and Ollar and Penta, 2017).¹⁵

2.2 Bidding on brands

If there are sponsored links, we need to distinguish the four cases in which firm i wins no auction, both, or only one of them (which could be either on its own brand keyword, or the competitor's). We first consider i 's profits realized via search result pages for i 's own keyword, in the case in which firm i loses or wins the corresponding auction.

If firm i *loses* this auction, then the demand it faces from these searches is:

$$S_i(p_i) \left[\eta_i \cdot x_{(i,s)}^{i|j}(p_i, p_j) + (1 - \eta_i) \cdot x_{(i,n)}^{i|j}(p_i, p_j) \right] = S_i(p_i) \cdot \eta_i \cdot x_{(i,s)}^{i|j}(p_i, p_j).$$

This is the sum of sales made through clicks of sophisticated consumers who searched for i (who still see its organic link after the sponsored link by firm j), and those made to the naive consumers of type i who searched for i but clicked on j , whose sponsored linked showed at the top of their result page (this second term is zero under our assumptions).

If firm i *wins* this auction, then both naive and sophisticated consumers of type- i will only see i 's link, and hence i 's demand from these searches is:

$$S_i(p_i) \left[\eta_i \cdot x_{(i,s)}^{i|i}(p_i, p_j) + (1 - \eta_i) \cdot x_{(i,n)}^{i|i}(p_i, p_j) \right] = S_i(p_i).$$

It follows that the *net value of winning* the auction on its *own* brand is:

$$VPI_i^i = S_i(p_i) \left[\eta_i \left(1 - x_{(i,s)}^{i|j}(p_i, p_j) \right) + (1 - \eta_i) \right] \max \{ (p_i - c_i), 0 \}. \quad (14)$$

Since in this case all consumers only see i 's link, all the $S_i(p_i)$ consumers who searched for i would click on its sponsored link, and hence the net *value-per-click* from winning the auction on its own brand is:

$$VPC_i^i = \left[\eta_i \left(1 - x_{(i,s)}^{i|j}(p_i, p_j) \right) + (1 - \eta_i) \right] \max \{ (p_i - c_i), 0 \}. \quad (15)$$

We turn next to i 's potential profits that could be realized via search result pages for j 's keyword. If firm i *loses* that auction, then it would make zero profits on this page: if the sponsored link is won by j , the only remaining organic link would still be j 's, and hence neither the naive nor the sophisticated j -type consumers would become aware of i , and hence they would not purchase from it. If instead firm i *wins* the auction for the sponsored link on j 's brand search results, then firm i would face the following demand:

$$S_j(p_j) \left[\eta_j \cdot x_{(j,s)}^{i|i}(p_i, p_j) + (1 - \eta_j) \cdot x_{(j,n)}^{i|i}(p_i, p_j) \right].$$

¹⁵Besides ensuring uniqueness of the Nash equilibrium, in this context Condition (13) also guarantees that it is globally stable with respect to the *Cournot tatonnement* dynamics (cf. Moulin, 1984).

Hence, the net value of winning the auction on the *opponent's brand* is:

$$VPI_i^j = S_j(p_j) \left[\eta_j \cdot x_{(j,s)}^{i|i} (p_i, p_j) + (1 - \eta_j) \cdot x_{(j,n)}^{i|i} (p_i, p_j) \right] \max \{ (p_i - c_i), 0 \}. \quad (16)$$

Now, all the $S_j(p_j)$ consumers who searched for j would click on i 's sponsored link (the sophisticated will click to discover i 's price; the naive because its link is at the top of their page). Hence, the *value-per-click* for the opponent's brand auction is:

$$VPC_i^j = \left[\eta_j \cdot x_{(j,s)}^{i|i} (p_i, p_j) + (1 - \eta_j) \cdot x_{(j,n)}^{i|i} (p_i, p_j) \right] \max \{ (p_i - c_i), 0 \}. \quad (17)$$

Note that, other things being equal, the higher the *profit margin*, the higher the value of winning either auction, and hence the bids that a firm will want to place. This observation will play an important role in determining the outcome of the auctions, and will tend to result in the upward pricing pressure which will emerge in our model.

The Auction We assume that a single ad slot is sold for each brand keyword, via a sealed-bid *second price auction*. Hence, similar to the above, for each brand/keyword $k = 1, 2$, firms 1 and 2 simultaneously submit bids (b_1^k, b_2^k) , the highest bidder wins the auction and obtains the slot, and pays a *price-per-click (PPC)* equal to the second highest bid (in case of ties, we assume that they are broken in favor of firm k , i.e. the one whose brand was searched).¹⁶ Hence, for each $k \in \{1, 2\}$, taking (p_1, p_2) as given, the optimal (undominated) bid, for each $k \in \{1, 2\}$, is $b_i^k = VPC_i^k$. Ignoring the irrelevant case in which $p_i < c_i$, we obtain the following *optimal bids* for the two auctions:

$$b_i^i = \left[\eta_i \left(1 - x_{(i,s)}^{i|i} (p_i, p_j) \right) + (1 - \eta_i) \right] (p_i - c_i), \text{ and} \quad (18)$$

$$b_i^j = \left[\eta_j \cdot x_{(j,s)}^{i|i} (p_i, p_j) + (1 - \eta_j) \cdot x_{(j,n)}^{i|i} (p_i, p_j) \right] (p_i - c_i). \quad (19)$$

The next result, which follows directly from these conditions, states an important property of the auction outcomes when firms are fully symmetric:

Lemma 1 *Under our maintained assumptions, if $p_i - c_i = p_j - c_j$ and $\eta_i = \eta_j = \eta$, then each firm wins the auction for its own brand keyword.*

Proof. First, substituting the symmetry conditions $(p_i - c_j) = (p_j - c_j)$, and $\eta_i = \eta_j = \eta$ into equations (18-19), we have that $b_i^i \geq b_i^j$ if and only if $1 \geq \eta \left(x_{(i,s)}^{i|i} + x_{(i,s)}^{j|j} \right) + (1 - \eta) \cdot x_{(i,n)}^{j|j}$. Then, note that $x_{(i,s)}^{i|i} + x_{(i,s)}^{j|j} = 1$. That is, all type (i, s) will buy when j wins the auction: since they did search for i , they are willing to buy from it (it still gives them non-negative surplus), but now they see both brands. Thus, they may buy from j instead, but abstaining

¹⁶For the single-slot case, this is in fact the mechanism that is actually used by all major search engines. That is because the GSP auction (which is adopted, for instance, by Google, Taobao, Bing; cf. Varian (2007) and Edelman et al. (2007)) coincides with the baseline second-price auction when a single slot is sold, and is typically implemented in the PPC format that we consider. An alternative payment system that is sometimes adopted is the *pay-per-impression (PPI)* (or pay-per-view). In our baseline model the payment of the auction is triggered whenever there is a click on the sponsored link, and given our assumptions, all those who search will also click. This implies that PPI and PPC are identical.

from the purchase could only decrease their utility. Next, recall that $x_{(i,n)}^{j|j} \leq 1$, since not all type- i naive consumers will necessarily buy j even if this is their only purchase chance (for instance, because j may give them negative net surplus). Hence, $b_i^i \geq b_j^i$. Given the tie-breaking rule that favors the owner of the brand, the result follows. ■

As a final remark, note that in real sponsored search auctions the search engine typically determines the winner of the auction by adjusting the monetary bid by a *quality score*. This means that firm j would have to bid significantly higher than i to win the ad slot for i -s brand. We have just showed that, in the symmetric equilibria of symmetric environments (which will be characterized in the next Section), each firm wins the auction for its own brand even in the absence of quality scores.

3 Symmetric environments

In this Section we study an economy with perfectly symmetrical firms, i.e. such that $c_1 = c_2$; $s_1 = s_2 = 1/2$; and $\eta_1 = \eta_2 = \eta$. First we characterize the symmetric (pure) Nash equilibria in both bids and prices (Section 3.1); then we compare them with those of the no-ads benchmark, and we show that prices are always higher with sponsored search auctions (Section 3.2). In Section 3.3 we discuss some illustrative examples that are based on parametric models, and show that despite the full symmetry between the firms, asymmetric equilibria in which one firm wins both auctions may also exist.

3.1 Symmetric Equilibria in bids and prices

In a symmetric equilibrium, both firms choose the same price, $\tilde{p}_1 = \tilde{p}_2$, and place the same bids, $(\tilde{b}_1^1, \tilde{b}_1^2) = (\tilde{b}_2^2, \tilde{b}_2^1)$ on their own and on the opponents' brand auction. Then, Lemma 1 directly implies that, in any such equilibrium, each firm i wins the auction on its own brand, and hence all the i -types who searched only see i 's link and remain unaware of firm j in equilibrium. This means that each firm i ends up being a monopolist over consumers of type- i (both naive and sophisticated), and there is no competition between firms in the product market. Moreover, for each of the $S_i(\tilde{p}_i)$ consumers that search its brand and click on its link, the firm pays an amount equal to the losing bid \tilde{b}_j^i placed by its competitor. Thus, firm i 's overall profits in a symmetric equilibrium are:

$$\tilde{\pi}_i(\tilde{p}_1, \tilde{p}_2) = S_i(\tilde{p}_i)(\tilde{p}_i - c_i) - \tilde{b}_j^i \cdot S_i(\tilde{p}_i). \quad (20)$$

It follows that *necessary conditions* for any symmetric equilibrium, $(\tilde{p}_i, \tilde{b}_i^i, \tilde{b}_i^j)$, are that it satisfies both conditions (18), (19), and the F.O.C. $\frac{\partial \tilde{\pi}_i}{\partial p_i}(\tilde{p}_i, \tilde{p}_j) = 0$, that is:

$$S_i(\tilde{p}_i) + \frac{\partial S_i(\tilde{p}_i)}{\partial p_i}(\tilde{p}_i - c_i) - \tilde{b}_j^i \cdot \frac{\partial S_i(\tilde{p}_i)}{\partial p_i} = 0. \quad (21)$$

Note that, under the maintained assumptions, the profit function in (20) is concave. However, these conditions are *not* sufficient for an equilibrium. In equilibrium, firms should also have

no incentive to unilaterally change their price and bids so as to change the outcome of the auctions (e.g., so as to win both, lose both, or only win the auction for the opponent's brand), thus moving to a different profit function (we will illustrate this point in Section 3.3). If such profitable deviations exist, then there are no (pure) symmetric equilibria.

3.2 Comparing auction and benchmark equilibria

Before proceeding to a comparison with the equilibrium outcome in the no-ads benchmark, it is useful to compare the equilibrium profit functions in (20) with their counterpart from the no ads benchmark in Section 2.1 (equations (10)-(11)). Two main differences are immediately apparent, with distinct effects on the resulting equilibrium prices. First, the presence of the advertisement cost in the advertisement case provides a further, separate channel through which optimal pricing may be affected. This is clearly seen from inspecting the necessary condition (21), from which it is clear that the more consumers' search is responsive to i 's price, the higher $\partial S_i(\tilde{p}_i)/\partial p_i$, and hence the stronger the marginal cost effect of advertisement. Second, while in the no-ads benchmark each firm i is a monopolist over the *naive* consumers of type- i , but competes with firm j for the sophisticated consumers of both types, in the symmetric equilibria with sponsored ad auctions each firm i ends up being a monopolist over *all* consumers of type- i , and there is no competition between firms in the product market.

Formally, imposing the conditions for symmetric environments into equations (10)-(11), and rearranging terms, the equilibrium profits in the no-ads benchmark can be written as

$$\bar{\pi}_i(p_1, p_2) = \hat{\pi}_i(p_i) + \eta(p_i - c_i) [q^{iD}(p_i, p_j) + q^{jM}(p_j) - q^{jD}(p_i, p_j)], \quad (22)$$

where $\hat{\pi}_i(p_i) = (1 - \eta)(p_i - c_i)S_i(p_i)$. Similarly, rearranging terms in equation (20), the profits under the ad auctions can be written as:

$$\tilde{\pi}_i(p_1) = \hat{\pi}_i(p_i) + \eta(p_i - c_i) [q^{iM}(p_i)] - \tilde{b}_j^i \cdot S_i(p_i). \quad (23)$$

Next, take the FOCs at the benchmark and denote as (\bar{p}_i, \bar{p}_j) the price pair which solves $\frac{\partial \bar{\pi}_i(\bar{p}_1, \bar{p}_2)}{\partial p_i} = 0$. When evaluated at the benchmark equilibrium prices, the price under the ads auctions will (weakly) increase if $\frac{\partial \tilde{\pi}_i(\bar{p}_1)}{\partial p_i} \geq \frac{\partial \bar{\pi}_i(\bar{p}_1, \bar{p}_2)}{\partial p_i}$, or:

$$\eta \left((\bar{p}_i - c_i) \left[\frac{\partial q^{iM}(\bar{p}_i)}{\partial p_i} - \frac{\partial q^{iD}(\bar{p}_i, \bar{p}_j)}{\partial p_i} + \frac{\partial q^{jD}(\bar{p}_i, \bar{p}_j)}{\partial p_i} \right] \right) - \tilde{b}_j^i \cdot \frac{\partial S_i(\bar{p}_i)}{\partial p_i} \geq 0, \quad (24)$$

where we have used $q^{iM}(\bar{p}_i) = q^{jM}(\bar{p}_j)$, $q^{iD}(\bar{p}_i, \bar{p}_j) = q^{jD}(\bar{p}_i, \bar{p}_j)$, and $\frac{\partial q^{jM}(p_j)}{\partial p_i} = 0$.

The first effect we alluded to, that is that the auctions may cause an *extra marginal cost*, due to the payment of the rival's bid, is reflected by the last term in this condition, in that $\frac{\partial S_i(\bar{p}_i)}{\partial p_i} \leq 0$. It is straightforward to see that, as long as $S_i(\cdot)$ is responsive to p_i , this term would always translate into a higher price, other things being equal. The strength of this effect obviously grows with the price elasticity of search.

Abstracting from this, the second effect instead is captured by the term in square brackets

in (24): since $\frac{\partial S_i(\bar{p}_i)}{\partial p_i} \leq 0$ a sufficient condition for (24) to hold is that:

$$\frac{\partial q^{iM}(\bar{p}_i)}{\partial p_i} - \frac{\partial q^{iD}(\bar{p}_i, \bar{p}_j)}{\partial p_i} + \frac{\partial q^{jD}(\bar{p}_i, \bar{p}_j)}{\partial p_i} \geq 0. \quad (25)$$

Some standard algebraic manipulation allows to write this condition as follows:

$$-\epsilon_M + \frac{q^{iD}}{q^{iM}} \epsilon_D + \frac{q^{jD}}{q^{iM}} \gamma_D \geq 0,$$

where ϵ_M , ϵ_D , and γ_D are respectively (the absolute values of) the own price elasticities of monopoly and duopoly demand functions, and the cross-elasticity of the demand function. Since at the symmetric equilibrium $q^{iD} = q^{jD}$, this can be rewritten as:

$$\epsilon_M \leq \frac{q^D}{q^M} (\epsilon_D + \gamma_D). \quad (26)$$

If (25) holds, then ad auctions increase prices, regardless of the marginal cost effect we discussed above, and even if consumers' search decisions were completely inelastic.

The following “*trade integration v. autarky*” interpretation may help gain some intuition about this condition, which effectively characterizes the sign of the comparison between the equilibrium prices in an *integrated* duopoly, and in two *segregated* monopolistic markets. Consider a simple international trade example with two symmetric countries, 1 and 2, each with a domestic firm, where consumers have a demand that exhibits no ‘love for variety’, so $q_i^{iM}(p) = q_i^{iD}(p, p) + q_i^{jD}(p, p)$, and where there are no transport or other trade costs. Under an autarkic scenario, where each firm sells only in the domestic market, profits are:

$$\tilde{\pi}_i^A(p_1) = (p_i - c_i) [q_i^{iM}(p_i)]; \quad (27)$$

whereas in a trade integration scenario, they would be:

$$\bar{\pi}_i^T(p_1, p_2) = (p_i - c_i) [q_i^{iD}(p_i, p_j) + q_j^{iD}(p_i, p_j)],$$

which, by using symmetry and the ‘love for variety’ condition, can be rewritten as:

$$\bar{\pi}_i^T(p_1, p_2) = (p_i - c_i) [q_i^{iD}(p_i, p_j) + q_j^{jM}(p_j) - q_j^{jD}(p_i, p_j)]. \quad (28)$$

(Note that the relevant portions of profit functions (27) and (28) correspond to the functions (23) and (22) respectively). By using the same procedure as above, it is straightforward to see that condition (26) characterizes the conditions under which moving from an international trade to an autarky equilibrium would lead to higher prices.

Finally, note that even if prices were equal, there would be a further welfare effect on consumers. Namely, at the no-ads benchmark, at least some consumers (namely, the sophisticated types) have access to both firms, and hence they may buy a product which is closer to their location or ideal variety. In contrast, at the sponsored ads auction equilibrium, they may either have to buy a product which gives them lower net utility, or not buy at all. *Ceteris*

paribus, this reduces the consumer surplus at the auction equilibrium. This can be neatly illustrated with a simple Hotelling example: Suppose that firms 1 and 2 are located at l_1 and l_2 and they set the same price. When firm 1 is a monopolist, consumers located close to l_2 may not buy, or they may buy but incur a high disutility (or transport) cost. If those consumers could also buy from 2, even if $p_2 = p_1$ their surplus would be greater.

The discussion in this subsection can be summarised as follows.

Remark 1 *In a symmetric equilibrium with sponsored ads auctions, type- i consumers will only see brand i and will buy at monopoly price from firm i , whose marginal costs are given by $c_i + \tilde{b}_j^i \cdot \frac{\partial S_i(p_i)}{\partial p_i}$. A sufficient condition for this price to be higher than at the benchmark is given by (26), which states that this monopoly price is higher than the benchmark duopoly price that would prevail when the firms are competing to attract sophisticated consumers of both types. Consumers are also worse off relative to the benchmark equilibrium because of less choice, and because the payment of the auction may represent an additional marginal cost. The latter effect is present whenever consumers' search decisions are responsive to firms' prices.*

The driving force behind these results is the fact that the sponsored links *crowd out* the organic ones: first, absent the ads, type- i sophisticated consumers might be able to see and buy product j , whereas with the sponsored ad auction firm j 's organic link disappears, leaving them captive to firm i ; second, absent the ads, the naive consumers would click (and buy) from an organic link, not from a sponsored link, which adds to the firms' marginal costs thereby affecting prices.¹⁷

Alternative benchmark and price elasticity of the search function. In the baseline model, we assumed that absent the sponsored links, when type- i consumers search for brand i , the search engine results page (SERP) returns organic links for both firms i and j . As a result, if in the advertising case firm i wins the auction on its own brand (which is the case in any symmetric equilibrium, cf. Lemma 1), j 's link is effectively crowded out by the sponsored link, contributing to the anti-competitive effects that we discussed in the previous section. If one changed the benchmark so that SERPs only return links of the firm whose brand was searched, this *crowding out* effect obviously disappears. Note, however, that the other effect that we discussed above – namely, the *marginal cost* effect due to the auction payment – would still remain. Hence, as long as consumers' search decision is responsive to prices, the presence of brand advertising would induce upward pressure on prices.

If, on the other hand, the search function is constant in prices, it follows directly from the analysis in Section 3.2 that the *marginal cost effect* disappears: advertising is a fixed cost, and it does not affect pricing decisions. Hence, in that case, only the *crowding out* effect remains: Under the auction, the equilibrium price would be the monopoly price, $\tilde{p}_i = p^M$; at the benchmark instead it will be a weighted average of the monopoly price p^M and the price

¹⁷Whether sophisticated consumers click on the organic or the sponsored link when they see both links from the same firm does not affect the results. See Section 4.

prevailing at a duopolistic equilibrium where firms compete for the sophisticated consumers of both types.

Obviously, if (i) the benchmark only returns links of the firm whose brand was searched *and* (ii) search decisions are completely inelastic to prices, then the presence of advertisement has no effect on prices.

3.3 An Hotelling Application

To further illustrate the points made in Sections 3.1 and 3.2, we consider next a parametric example, based on a symmetric version of the Hotelling model, with $c_1 = c_2 = 0$, $\eta_1 = \eta_2 = \eta$, and $s_1 = s_2 = 1/2$. Consumers have utility $U_i = 1 - p_i - t | \theta - l_i |$, where θ , which describes their location, is uniformly distributed on the real line; l_i is the location of firm $i = 1, 2$, for $l_1 = 0$, $l_2 = 1$; and $t \in (4/21, 1/2)$ is the transport (or disutility) cost.¹⁸ The consumers indifferent between buying i and not buying are $\theta_{i0}^L = l_i - (1 - p_i)/t$ and $\theta_{i0} = l_i + (1 - p_i)/t$. The indifferent consumer between buying 1 and 2 is $\theta_{12} = (t - p_1 + p_2)/(2t)$. Using these expressions, the conversion rates for 1 are:

$$x_{(1,s)}^{1|2} = \frac{\theta_{12} - \theta_{10}^L}{\theta_{10} - \theta_{10}^L} = \frac{2 - 3p_1 + p_2 + t}{4(1 - p_1)}; \quad x_{(2,s)}^{1|1} = x_{(2,n)}^{1|1} = \frac{\theta_{12} - \theta_{20}^L}{\theta_{20} - \theta_{20}^L} = \frac{2 - p_1 - p_2 - t}{4(1 - p_2)}. \quad (29)$$

And symmetrically for firm 2. Next, replace these values into (18) and (19) to obtain the optimal bids:

$$b_i^i = \frac{p_i [2(2 - \eta) - p_i(4 - 3\eta) - \eta(p_j + t)]}{4(1 - p_i)}; \quad b_j^i = \frac{p_j(2 - \eta) [2 - t - p_i - p_j]}{4(1 - p_i)}. \quad (30)$$

Note that $\partial b_i^i/\eta < 0$ and $\partial b_j^i/\eta < 0$: the higher the share of sophisticated consumers in the economy, the lower the bids. This is because winning naive consumers is more valuable: if a firm loses them, they will not buy from it (not so for the consumers of type (i, s)).

At a symmetric equilibrium $p_i = p_j = p$, and hence (as shown by Lemma 1 for the general model), we have $b_i^i \geq b_j^j$ for any $\eta \in [0, 1]$ – holding with equality for $\eta = 1$ – so that each firm i wins the auction on its own brand and makes profits $\tilde{\pi}_i(p_1, p_2) = p_i S_i(p_i) - \tilde{b}_j^i \cdot S_i(p_i)$. To find the equilibrium, we need to solve the system of the two equations $\frac{\partial \tilde{\pi}_i(p_1, p_2)}{\partial p_i} |_{(p_i=p_j)} = 0$ and $b_j^i(p_i, p_j) = \frac{p_i(2-\eta)[2-t-p_i-p_j]}{4(1-p_i)}$ (cf. conditions (18), (19), and (21) in Section 3.1). The solution of these necessary conditions is given by:

$$\tilde{p}_i = \tilde{p}_j = \frac{8 + 2t - \eta t + 2\eta - \sqrt{[2(4 + \eta) + (2 - \eta)t]^2 - 32(2 + \eta)}}{4(2 + \eta)}. \quad (31)$$

Hence, if this profile of prices and bids is actually an equilibrium (recall, from the discussion in Section 3.1, that these conditions are only necessary), then it is also the *unique*

¹⁸The restriction $t < 1/2$ ensures that in the standard model with all consumers fully informed duopoly prices are lower than monopoly prices (when the market is served by just one of the two firms); $t > 4/21$ is needed for an interior equilibrium of the duopoly model with type-1 consumers to exist. At the equilibrium, which is asymmetric, $p_1^d < p_2^d$ (firm 1's price is lower because p_1 determines the set of consumers searching, and hence the potential demand of firm 1), and if t was too low then firm 1 would get all the demand (that is, the indifferent consumer between 1 and 2 would lie to the right of l_2).

symmetric equilibrium.

Existence of the symmetric equilibrium. To ensure that these strategies are global optima, and hence that a symmetric equilibrium exists, we also have to check that given $(\tilde{p}_j, \tilde{b}_j^j, \tilde{b}_j^i)$ that satisfy (30) and (31), firm i does not have an incentive to set a triplet (p_i, b_i^i, b_i^j) that changes the outcome of the auctions. It can be shown that such deviations may actually be profitable for sufficiently high values of η (details available from the authors), and hence symmetric equilibria do not exist, when the fraction of sophisticated consumers is especially high. Figure 3 illustrates the region of existence of the symmetric equilibrium where each firm wins its own brand name keyword auction and loses the rival's (WL, WL): the equilibrium exists only if there is a sufficiently high share of naive consumers in the population (that is, below the curve), with the relevant threshold being increasing in the transportation costs.

Intuitively, if $\eta \rightarrow 0$, naive consumers represent most of the population and b_j^j is much higher than b_j^i . Therefore, winning the rival brand's auction would be very costly (precisely because b_j^j is high), and losing its own auction would imply losing most sales (naive consumers do not “see” the second link). Hence, there is no incentive to deviate. Consider instead the case where $\eta \rightarrow 1$. At the candidate equilibrium, $b_i^i = b_j^i$, so for firm i winning its own auction is relatively costly. Firm i could then deviate by losing its own auction and undercutting the rival's price, by setting $p_i' < p_j$. This would allow it to save incremental costs (it does not have to pay b_j^j) while still winning most of type- i consumers thanks to the lower price p_i' . (Recall that sophisticated consumers of type i see both links when j wins brand i 's auction). Also, the lower price expands the total size of the market, since the set $S_i(p_i)$ of searchers increases when p_i falls.




Figure 3: Equilibria in a Symmetric Environment

The symmetric equilibrium (WL, WL) where each firm wins own brand auction exists in the orange region; the asymmetric equilibrium (WW, LL) (and symmetrically, (LL, WW)) where a firm wins both auctions exists in the blue region; In the white region there are no pure-strategy equilibria.

Comparison with the no-ads benchmark. There are two ways to check that this price is higher than at the no-ads benchmark. One is to compute the latter, which turns out to be $\bar{p}_i = \bar{p}_j = 1/2$, and verify that $\tilde{p}_i > \bar{p}_i$. The second method is to apply directly conditions (24) and (25). One can then check that the sufficient condition (25) holds with equality in this model, and hence the term in square brackets in equation (24) is zero.¹⁹ Prices, however, are still strictly higher in the sponsored search auction case, because of the marginal cost effect due to the auction payments. Condition (24) thus holds with a strict inequality because $-\tilde{b}_j^i \cdot \frac{\partial q_{(i,s)}^{i,M}(\bar{p}_i)}{\partial p_i} > 0$.

Asymmetric equilibria. Interestingly, despite the symmetry of the model, there may also exist asymmetric equilibria in which one of the two firms wins both auctions. If, say, firm 1 wins both auctions, then the firms' profits will be:

$$\pi_1 = \frac{(1-p_1)(p_1-b_2^1)}{t} + \frac{(1-p_2)}{t} \left(\frac{p_1(2-p_1-p_2-t)(2-\eta)}{4(1-p_2)} - b_2^2 \right); \quad \pi_2 = \frac{p_2\eta(2+t+p_1-3p_2)}{4t}.$$

After taking the FOCs and replacing the optimal bids given by (30), one can find the candidate equilibrium prices, bids, and profits $(\hat{p}_1, b_1^1(\hat{p}_1, \hat{p}_2), b_1^2(\hat{p}_1, \hat{p}_2))$ (details available from the authors). Similar to the above, here too one needs to check that there are no profitable deviations – in particular, firm 2 may find it profitable to outbid firm 1 on its own (i.e., brand 2) keyword auction. Figure 3 illustrates the region of existence of the asymmetric equilibrium (either (WW, LL) or (LL, WW)) where one firm wins both auctions: the equilibrium exists only if there is a sufficient share of sophisticated consumers in the population.

Note that also in such an asymmetric equilibrium (when it exists), the net effect of sponsored search auctions on brand keywords is detrimental to the consumers. Intuitively, for 1 to win both auctions, firm 1 needs to have higher prices than firm 2, and indeed $\hat{p}_1 > \bar{p} > \hat{p}_2$. But by winning its own auction firm 1 crowds out the organic link of firm 2, resulting in both naive and sophisticated type-1 consumers to be monopolised by firm 1. This effect outweighs the beneficial effect due to firm 2's prices being lower than at the benchmark.

To understand the existence conditions, note that when most consumers are sophisticated, firm 2 sells nothing to type-1 consumers, but it does sell to type-(2, s) consumers, who still “see” brand 2, and hence it does make positive profits. To win its own auction, firm 2 should considerably raise its bid (at the candidate equilibrium, $\hat{p}_1 > \hat{p}_2$ and $\hat{b}_1^2 > \hat{b}_2^2$), but this is expensive. As the share of naive consumers rises, this deviation may become profitable. To see why, consider the case where most consumers are naive. Then at the candidate equilibrium firm 2's profits are very small (in the limit, if $\eta \rightarrow 0$, $\bar{\pi}_2 \rightarrow 0$). By raising the bid on its own

¹⁹That is, in this version of the Hotelling model, the ‘integrated duopoly’ price is the same as the ‘autarkic monopoly’ price. While this makes the example rather special and intuitively less desirable, it has the advantage that the result that sponsored ad auctions are detrimental to consumers does not depend on the benchmark chosen for the no-ads case. The reason is that the price at the no-ads benchmark when type- i consumers see organic links of both brands is the same as in an alternative benchmark, in which they only see organic links of brand i . In a variation of the Hotelling model, with consumers positioned on the unit interval, quadratic transportation costs, and firms located at the extremes, the ‘integrated duopoly’ would have lower prices than the ‘autarkic monopoly’, and hence sponsored links would induce higher prices due to both the *crowding out* and the *marginal cost* effect in the baseline benchmark, and only the latter effect in the alternative benchmark.

auction, b_2^2 , firm 2 would obtain all of the type-2 naive consumers, and hence this deviation would be profitable.

4 Extensions and variations

In this section we discuss several variations and extensions of the baseline model, and show that our main results on the price effects of brand auctions are robust to a variety of changes to the model.

4.1 Letting sophisticated consumers favor organic links

Some users may want to avoid clicking on sponsored links when organic links providing similar information are available. (This may be because of aversion to anything related to publicity, or because they fear that sponsored links may provide them with biased information relative to organic links, or again because they do not like it that, when clicking, a payment to the search engine is triggered.) In our context, suppose that sophisticated consumers behave in this way. Then, when firm i wins the auction for brand i , and both the sponsored and the organic link to website i are provided, the sophisticated consumers would click the latter link. As a consequence, firm i would only have to pay for the clicks of the naive consumers, while still managing to crowd out the rival brand link from the search engine results page. This increases the value-per-click that i gets from winning the auction on its own brand.²⁰ Moreover, since nothing changes for j 's bid on i 's brand (type- i sophisticated consumers would be interested in comparing offerings, and the only way to do so when they search for i would be to click on j 's sponsored link, if available), this behavior would further favor firm i on its own brand auction (and, thus, further impair its position on j 's brand). But nothing would change at the equilibrium, since the losing bid is unchanged.

4.2 Introducing perfectly informed consumers

In addition to our naive and sophisticated consumers there may also be consumers with identical preferences but who are perfectly informed, or who have gathered information on the products on offer without using the search engine. Then, their demands will not be affected by the existence of the auction and who wins it, since they have information on both firms no matter what the SERP shows. It follows that the bids made by each firm will not change: VPC_i^i and VPC_I^j are the same as in the base model, and so will be b_i^i and b_i^j . Thus, in a symmetric equilibrium we will still have $b_i^i \geq b_j^j$, and hence each firm wins the auction on its own brand.

Formally, letting μ denote the proportion of fully informed consumers, we shall have $(1 - \mu) \cdot \eta$ sophisticated consumers and $(1 - \mu)(1 - \eta)$ naive ones. At the no-ads benchmark,

²⁰Formally, letting β_i^i denote the optimal bid in this case, we would have $\beta_i^i = \frac{VPI_i^i}{(1-\eta)S_i(p_i)} > \frac{VPI_i^i}{S_i(p_i)} = b_i^i$, where b_i^i denotes the optimal bid obtained in the previous section.

firm i 's profits can be written as:

$$\bar{\pi}'_i(p_1, p_2) = \hat{\pi}'_i(p_i, p_j) + \eta(1 - \mu)(p_i - c_i) [q^{i,D}(p_i, p_j) + q^{j,M}(p_j) - q^{j,D}(p_i, p_j)],$$

where now $\hat{\pi}'_i(p_i, p_j) = \mu(p_i - c_i)q^{iD}(p_i, p_j) + (1 - \eta)(1 - \mu)(p_i - c_i)S_i(p_i)$. Similarly, the profits under the ad auctions are:

$$\tilde{\pi}'_i(p_1) = \hat{\pi}'_i(p_i, p_j) + \eta(1 - \mu)(p_i - c_i) [q^{iM}(p_i)] - (1 - \mu)q^{iM}(p_i)\tilde{b}_j^i.$$

The difference between the two profit functions therefore is unaffected, compared to the case in the baseline model (cf. equations (22)-(23)), except for the presence of a multiplicative term $(1 - \mu)$. Hence, following the same steps as in the previous section, it is easy to see that the conditions for the sign of the price comparisons are unaffected (so, for instance, condition (25) is still sufficient for the price to increase with brand auctions). What does, change, however, are the conditions for equilibrium existence, which become more permissive as the proportion of fully informed consumers gets larger. The reason is that, as more perfectly informed consumers are present, the incentives to deviate in a way to win both auctions become weaker, and hence the candidate equilibrium strategies are immune to ‘global’ deviations for a broader range of parameters. In terms of Figure 3, for instance, the locus of points where the equilibrium (WL, WL) exists shifts up as the share of fully informed consumers moves from $\mu = 0$ (which corresponds to the figure) to $\mu = 1$. Further, it turns out that once μ is sufficiently high, there is no region of the parameter space where no pure strategy equilibria exist, and instead there will be a region where both the symmetric and the asymmetric equilibria exist.

4.3 More than two links and consumer types

In the baseline model we also assumed that the SERP only contains two links, resulting in the sponsored link crowding out an organic link, and hence preventing sophisticated consumers from seeing the rival’s link any longer, when the owner of the brand they searched wins the auction. Consider now an alternative model where there are three links on the page, just one slot being auctioned off for ads, and three categories of consumers, denoted by C_n , with $n = \{1, 2, 3\}$ being the number of links, starting from the top, that consumers of category C_n pay attention to. Thus, C_1 consumers behave like the naive in the baseline model, the C_2 consumers like the sophisticated, and the C_3 will always be fully informed, like those of the previous section. Hence, by denoting their respective shares in the population as $(1 - \mu)(1 - \eta)$ for the C_1 , $(1 - \mu)\eta$ for the C_2 , and μ for the C_3 -types, the model is identical to the one we discussed in the previous section 4.2, which as explained gives rise to the same results as in the main model. Therefore, under this reinterpretation, even if the sponsored link does not fully displace an organic link, the qualitative message from the baseline model would be unchanged, the only differences being in the existence conditions for the equilibria, which under this variation become more permissive.

4.4 Sequential game: prices then bids

We now modify the assumption that firms simultaneously choose prices and bids. In this Section, we study the Hotelling model on the real line (see Section 3.3) under the assumption that firms first choose prices and then bids.²¹ This variation does not affect the no-ads benchmark, as there is no bidding stage. As for the bidding stage when sponsored link are present, the conversion rates and optimal bids remain the same as in the baseline model, given the first period prices (cf. (29) and (30)). We focus next on the symmetric equilibria of the sequential game, in which firms choose price in the first period, anticipating the impact that this will have on the second period bids and auction outcomes.

In this case, firm i 's profits will be given by $\pi_i^{WL} = (p_i - b_j^i(p_i, p_j)) \cdot S_i(p_i)$, with $b_j^i(p_i, p_j) = \frac{(2-\eta)(2-t-p_i-p_j)}{4(1-p_i)}$ and $S_i(p_i) = (1 - p_i)/t$.²² (The index WL refers to a firm winning its own brand auction and losing the rival's.) By taking the FOCs and solving, we find the candidate equilibrium prices and by substitution the candidate equilibrium bids:

$$p_i = p_j = p^{WL} = \frac{4}{6 + \eta}; \quad b_i^i = \frac{8 - t\eta(6 + t) - 2\eta^2}{12 + 8\eta + \eta^2}; \quad b_i^j = \frac{(2 - \eta)[2(2 + \eta) - t(6 + \eta)]}{12 + 8\eta + \eta^2}.$$

Note that the in the auction configuration the candidate equilibrium price is higher than at the benchmark for any η , which confirms the results obtained in the main model.

To verify that this is indeed an equilibrium, first note that at these prices profits are $\pi^{WL} = \frac{2\eta(2+\eta)+t(12-4\eta-\eta^2)}{t(6+\eta)^2}$. Then, as usual, we need to check that – given $p_j = p^{WL}$ – firm i has no incentive to change its price so as to either win both auctions, or to win none. The latter deviation turns out to be more binding. To look for it, consider that $\pi_i^{LL} = p_i \cdot \eta \cdot (2+t-3p_i+p^{WL})/(4t)$, – note that by deviating the firm loses both bids and hence does not pay additional costs, and that it can sell only to the type- i sophisticated consumers – leading to the optimal deviation $p_i' = \frac{t(6+\eta)+2(8+\eta)}{6(6+\eta)} < p^{WL}$. By substitution, one can find the optimal deviation profits and compare them with the candidate equilibrium profits. Figure 4 shows the region where the symmetric equilibrium exists (light-colored area), and where the asymmetric one with one firm winning both auctions exists (blue-shaded area). The results are qualitatively similar to the case of simultaneous moves (in particular, consumer surplus is always higher at the benchmark than with the sponsored ads).

²¹It seems natural to assume this order of moves, since firms are more likely to modify their bids than their price in the short-run, and it is difficult to think that firms' prices are contingent on the result of the bids. However, a few papers have posited the inverse order of moves. See e.g., Chen and He (2011), Xu et al. (2012), and Anderson and Renault (2021).

²²When computed in a neighbourhood of the equilibrium, $\partial b_i^i/\partial p_i > 0$, that is, an increase in price makes it more likely to win – ceteris paribus – the own brand auction. The sign of $\partial b_i^j/\partial p_i$ is a priori ambiguous, whereas $\partial b_j^j/\partial p_i < 0$.




Figure 4: Equilibria in sequential game

Sequential game where first set first prices and then bids. The symmetric equilibrium where each firm wins its own brand auction exists in the orange region; the equilibria where one firm wins both auctions exist in the blue region.

4.5 Product Search

The model can also be adapted to the case in which consumers search for *products* instead of *brands*. For instance, a type- i consumer who is willing to buy i does not enter ‘brand i ’ in her search bar, but the generic product (say, ‘tennis shoes’, or ‘tablet’, etc.). Consider a symmetric setting. In the no-ads benchmark, both firms’ links will show up as organic, in first or in second position with equal probability. The naive consumers will only ‘see’ the first link, which may be from the firm they already know about, or from the other. As with brand search, they will purchase the brand appearing in this first link if this gives them non-negative surplus. The sophisticated consumers instead will see both links, and hence always learn about both brands. As usual, they will buy the brand which gives them higher surplus. In the sponsored case, instead, since both types enter the same query, there will be a single auction; the first link on the SERP (the sponsored one) will be taken by the winner of the auction, and the other (the organic one) will be assigned to either firm with equal probability. This way, all the naive consumers, and half of the sophisticated ones, will only ‘see’ the link of the winner of the auction, with the losing firm being only accessible to *half* of the sophisticated consumers (partial *crowding out*). Purchase decisions will then take place as usual.

Compared with the case of brand search, the main difference in this setting is that symmetric equilibria will cease to exist: at any symmetric profile, each firm would win with probability $1/2$, but an arbitrarily small increase of their bids will make the probability of winning jump to one, taking the firm to a higher profit function. Thus, all equilibria will be asymmetric, and they only exist as long as the fraction of sophisticated consumers is high enough. The winner of the auction places a higher bid, commanded by the higher price that it will charge, optimally, given that its winning position would make it a monopolist for *all* the

naive consumers. The loser of the auction, instead, will only access half of the sophisticated, and always as a duopolist, since these consumers will also have access to the sponsored link of the winner of the auction. Thus, confined to a smaller market and with no one to exercise monopolistic power on, the optimal price will be lower for this firm, in turn commanding a lower bid. For such an equilibrium to exist, however, it must be that the market that remains for the losing firm is not too small, or it would be profitable to try to win the auction. This explains the lower bound on the share of sophisticated consumers required for the existence of these asymmetric equilibria, as showed by Figure 5.




Figure 5: Equilibrium in product search

Case where consumers search for the generic product (rather than for the specific brand they know). An asymmetric equilibrium where one of the firms wins the product auction exists in the blue region; no equilibrium in pure strategies exists in the white region.

In such asymmetric equilibria, it will always be the case that the winning firm charges a higher price than in the no-ads benchmark, due to both a *marginal cost* and to a *crowding out* effect. The losing firm, in contrast, may charge a higher or lower price than in the no-ads benchmark, depending on the degree of substitutability between the two products. In the Hotelling model, for instance, for high transportation costs t the equilibrium price of the winning firm will be so high that the best response of the losing firm (in the duopolistic market consisting of half of the sophisticated consumers) is higher than the equilibrium price in the no-ads benchmark, in which firms split the market evenly; for low transportation costs, instead, the threat of losses on the inframarginal units for the winning firm is high enough that it sets a lower price (though still higher than the no-ads benchmark), putting extra competitive pressure on the losing firm that is left with no captive consumers, which may thus best respond by charging a price that is lower than in the no-ads benchmark. The effect on consumers' welfare therefore is always negative, *except* for half of the sophisticated consumers who have access to the organic link of the losing firm: for some of those consumers – specifically, for those who actually end up buying from the losing firm – the welfare effect of the sponsored link is ambiguous, and will depend on the effect that it has on the losing firm (it may be positive if the two products are highly substitutable, not otherwise).

Another possible comparison is between the equilibrium outcomes of product search and brand search. In the region where the product search (asymmetric) equilibrium exists, under brand search there may exist a symmetric equilibrium or an asymmetric one.

In the former case, the firm who loses the auction under product search always sets lower prices than at the brand search equilibrium, whereas the firm which wins the auction under product search may set higher or lower prices than at the brand search equilibrium in different areas of the region where both the product and the brand search symmetric equilibria exist. However, consumer surplus is always higher under the product search equilibrium (also due to sophisticated consumers ‘seeing’ both brands — and hence having the possibility to minimise transportation costs — with 50% probability).

In the latter case, the ranking of prices for both firms depends on parameter values. But again, it turns out that consumer surplus is higher at the product search equilibrium.

In conclusion, if brand search was replaced by product search, there would be a gain in overall consumer surplus. We shall return to this result when discussing policy interventions in Section 6.

In any case it is worth stressing that considering product search instead of brand search mainly entails changes in the conditions for existence, but does not change the main qualitative punchline of our analysis, namely that sponsored search auctions would be detrimental to consumers relative to a benchmark where only organic search results exist.

5 Asymmetries

We have so far analysed the case where firms are symmetric, but arguably it is when firms are asymmetric that the procompetitive potential of sponsored ads is the highest. For instance, an entrant may be as efficient as an incumbent, but it is unknown to consumers and hence they cannot buy from it. In this situation, *if* the entrant managed to win the sponsored search auction for the incumbent’s brand, then it would be able to sell to at least some of the consumers who searched, and the ensuing market competition may reduce market prices.²³ Another example could be one in which one firm is more efficient than the other. One might then expect that, holding everything else constant, the more efficient firm is more likely to win the auction for *both* brand keywords, and hence that the increased market share of the low cost firm could potentially outweigh the (still present) anti-competitive effects seen in the symmetric case.

In this Section we study two applications which fit these examples. In Section 5.1 we consider the case where a new firm appears that is completely unknown to both the consumers and the search engine; in Section 5.2 we consider the case with asymmetric costs.

²³In Prat and Valletti (2021) two platforms sell advertising through auctions, and two firms, an Incumbent and an Entrant, compete in a product market. By posting an ad on both platforms, the Incumbent (already known by consumers) could foreclose the Entrant, who is unknown absent advertising. When the two platforms are independent, it is too costly for the Incumbent to foreclose the competitor. But if they merge (and coordinate their “selling of attention” to users), it is profitable: the reduction of product market competition leads to higher seller profits which are partly extracted by the platforms. Our setting is of course very different, but the common point is that even if a priori an entrant may be able to use ads to become visible to consumers, an incumbent may succeed in obtaining the sponsored ads available and hence keeping its market power.

5.1 An incumbent and a new entrant

There is an incumbent, firm 1, which is known to all consumers in the population; and an entrant, firm 2, which is completely unknown to them: $s_1 = 1$; $s_2 = 0$. Other than that, the model coincides with the Hotelling model with consumers uniformly distributed on the real line and linear disutility costs analysed in Section 3.3, with $c_1 = c_2 = 0$ and $\eta_1 = \eta_2 = \eta$. In order to avoid stacking the odds against a beneficial effect of the auctions, let us assume the search engine is *not omniscient*, that is, the organic search results *do not* reveal the existence of firm 2 when no consumer knows it. (Apart from making it more likely to find a procompetitive effect of auctions, this is not necessarily unrealistic: if users are unaware of a firm, its website will attract little traffic and search algorithms will not rank it highly, making it less likely to appear in a high position in the SERP.) Hence, absent sponsored links, it would return just organic links to the incumbent firm 1. In this case, the benchmark equilibrium is the straightforward monopoly equilibrium, with $\bar{p}_1 = p_1^M = \operatorname{argmax}\{p_1 \cdot 2(1 - p_1)/t\} = 1/2$.

By using the conversion rates derived in section 3.3, and applying the expressions for the optimal bid functions given by (18) and (19), we have:

$$b_1^1 = p_1 \cdot \frac{4 - p_1(4 - 3\eta) - \eta(2 + p_2 + t)}{4(1 - p_1)}; \quad b_2^1 = p_2 \cdot \frac{(2 - \eta)(2 - p_1 - p_2 - t)}{4(1 - p_1)}. \quad (32)$$

Note there is only one auction, since nobody will search for brand 2. Hence, there are only two possible cases, which we analyze next: either the incumbent wins the auction, or the entrant does.

5.1.1 The incumbent wins

In any equilibrium in which the incumbent wins the auction, its profits are $\pi_1(p_1) = \frac{2(1-p_1)(p_1-b_2^{1,h})}{2}$, and so it must be $p_1 = \operatorname{argmax}\{\pi_1(p_1)\} = \frac{1+b_2^{1,h}(p_1,p_2)}{2}$, from which we can already observe that in any such equilibrium, if it exists, the price will be higher than at the benchmark, where $\bar{p}_1 = 1/2$. Given the best replies in (32), $p_1 = \frac{1+b_2^{1,h}(p_1,p_2)}{2}$ implicitly defines a schedule $p_1 = f(p_2)$. Additionally, the candidate equilibrium must be such that $b_1^1(p_1, p_2) \geq b_2^1(p_1, p_2)$, or else it would not be the incumbent which wins the auction. There may be a continuum of prices p_2 which satisfy these conditions, leading to a multiplicity of equilibria in which the incumbent wins the auction.

Let us focus on the equilibrium where the entrant makes the highest possible bid. This is the most relevant case to consider, since it makes it the hardest for the incumbent to win the auction. Given any p_1 , the price at which the entrant's optimal bid is the highest is: $p_2^h = \operatorname{argmax}\{b_2^1\} = (2 - p_1 - t)/2$, and $b_2^{1,h}(p_1) = \frac{(2-p_1-t)^2(2-\eta)}{16(1-p_1)}$. By replacing and solving $p_1 = \frac{1+b_2^{1,h}(p_1)}{2}$, we obtain the candidate equilibrium price:

$$\tilde{p}_1 = \frac{28 - 2t - 2\eta + t\eta - 4\sqrt{-2 + 5t(2 - \eta) - 2t^2(2 - \eta) + 3\eta}}{34 - \eta}.$$

By using \tilde{p}_1 and p_2^h we can then check that it is indeed the case that $\tilde{b}_1^1 > b_2^{1,h}$ for any admissible value of t and η .

However, to ensure that this is indeed an equilibrium, we also have to check that neither firm 1 nor firm 2 have an incentive to deviate by respectively losing and winning the auction. It turns out that the former deviation, where given (\tilde{p}_2, b_2^1) firm 1 deviates and sets a pair (p_1, b_1^1) such that it loses the auction, is more restrictive. Figure 6 illustrates the parameter region where this condition holds, and hence where an equilibrium in which the incumbent wins the auction exists.²⁴ This is the case whenever the transportation cost t is sufficiently high or if there are enough naive consumers in the population. As explained above, since the price is always higher than at the benchmark and consumers' choice does not increase, consumer surplus is unambiguously lower than at the benchmark.




Figure 6: Equilibria with an Incumbent and an Entrant

The equilibrium (L, W) where the entrant wins the auction on the incumbent brand keyword exists in the blue region. The equilibrium (W, L) where the incumbent wins the auction exists in the orange region.

5.1.2 The entrant wins

If the entrant wins the auction for 1's brand keyword, firms' profits will be:

$$\pi_1 = \frac{\eta p_1 (2 - 3p_1 + p_2 + t)}{2t}; \quad \pi_2 = \frac{p_2 (2 - \eta) (2 - p_1 - p_2 - t)}{2t} - \frac{2(1 - p_1)}{t} b_1^1,$$

and the candidate equilibrium prices are $\tilde{p}_1 = (6 + t)/13$ and $\tilde{p}_2 = (10 - 7t)/13$. Note that firm 2's optimal price does not depend on b_1^1 , since its marginal profit does not depend on it (this is a case where the payment of the bid does not add to marginal costs). Further, $\tilde{p}_1 < \tilde{p}_2$, because firm 1 has an incentive to keep the price low so as to expand the market size (that is, the set of consumers who are willing to buy and thus search).

²⁴In a small region where both t and η are small, there are no real number prices where the equilibrium where the entrant makes its highest possible bid exists. However, in that region there are other pairs (p_1, p_2) at which the incumbent wins the auction

By replacing price expressions into the optimal bids, (18) and (19), we obtain:

$$\tilde{b}_1^1 = \frac{(6+t)[28-18\eta-t(4+3\eta)]}{52(7-t)}; \quad \tilde{b}_2^1 = \frac{(2-\eta)(10-7t)^2}{52(7-t)},$$

and a necessary condition for the equilibrium to exist is $\tilde{b}_2^1 > \tilde{b}_1^1$, which amounts to $\eta > \frac{-16+142t+51t^2}{4+88t+23t^2} \equiv \tilde{\eta}(t)$. In other words, it is only when there is a large enough share of sophisticated consumers (who are less “valuable” to the incumbent) that the entrant can win the auction.

For sufficiency, we also need to check that, given firm 2’s pair $(\tilde{p}_2, \tilde{b}_2^1)$, firm 1 does not want to deviate and set (p'_1, b_1) so as to outbid firm 2. This deviation would give it profits $\pi'_1(p'_1) = \frac{2(1-p'_1)(p'_1-\tilde{b}_2^1)}{t}$. The optimal deviation price is $p_1^{t*} = \max(p_1^t, \operatorname{argmax}\{\pi'_1(p'_1)\})$. The final step is to substitute it into the deviation profits and compare them with the candidate equilibrium profits. Figure 6 also draws the locus of the points where firm 1’s deviation profits equal the candidate equilibrium profits, and the region where the equilibrium (L, W) where the entrant wins the incumbent’s brand auction exists: It exists as long as t is sufficiently high and/or there are enough sophisticated consumers in the population.

When the equilibrium with entry exists, the effect on consumer surplus is a priori ambiguous, since there are different effects at work. On the one hand, we have $\tilde{p}_1 < 1/2 < \tilde{p}_2$, the latter inequality implying that some consumers will buy at a higher price than at the benchmark. On the other hand, consumers located closer to l_2 will save transport costs relative to the benchmark (where 1 monopolises the market). Further, most sophisticated consumers will buy from 1, and therefore their situation will be improved with respect to the benchmark. It turns out that the second effect dominates, and overall consumer surplus (as well as total surplus) is higher when the entrant wins the auction than at the benchmark.

Remark 2 *To conclude, brand search ad auctions may in some circumstances allow a new entrant to acquire a visibility that otherwise it would not have, and obtain a share of the market. When this happens, the sponsored ad auction is procompetitive. For this procompetitive effect to take place, however, it is necessary that (i) at the benchmark, the search engine is as ‘ignorant’ about the entrant as consumers are, that is, it does not list the entrant’s brand in its organic results; and (ii) there is a large enough proportion of sophisticated consumers in the economy, as showed by the light-shaded region in Figure 6. If, instead, the incumbent wins the auction, then the effect of brand auctions is always detrimental to the consumers.*

5.2 A firm is more efficient than the other

One may think the auction is procompetitive also when one firm is more efficient than the other: arguably, if it wins the sponsored links triggered by consumers aware of its rival firm, the increased visibility may allow the efficient firm to win market share. To explore this case, we rely again on the Hotelling model on the real line and we assume that $c_1 = 0 < c_2 \leq 1/2$. Otherwise, firms are symmetric: $\eta_1 = \eta_2 = \eta$ and $s_1 = s_2 = 1/2$. To avoid unnecessarily

complex expressions, we set $t = 1/3$.²⁵

To look for an equilibrium where firm 1 wins both auctions, we can use the expressions of the bids (30), of the conversion rates (29), and of the associated profit functions when 1 wins both auctions (3.3), all derived in Section 3.3, and with c_2 which can differ from zero and $t = 1/3$. By taking the FOCs, replacing the optimal bids and solving the system, we obtain candidate equilibrium prices (\hat{p}_1, \hat{p}_2) (details available from the authors). It can be checked that these prices satisfy the necessary second order conditions for a local optimum, and by replacing them into the usual bidding functions and in the profit expressions we obtain our relevant candidate equilibrium variables, $(\hat{p}_i, \hat{b}_i^i, \hat{b}_i^j)$. As usual, at this point, we need to check that firm 1 has no incentive to unilaterally deviate by setting a triplet at which it loses brand 2's auction, and firm 2 has no incentive to unilaterally deviate by setting a triplet at which it wins its own auction.²⁶




Figure 7: Equilibria with Cost Asymmetry

The equilibrium where the more efficient firm wins both brand keywords auctions is in the orange region. In the green region there additionally exists an equilibrium such that the inefficient firm wins both auctions.

The latter deviation turns out to be more binding, and the light-shaded region in Figure 7 illustrates the combination of parameters (c_2, η) where the equilibrium with 1 winning both auctions exists. Intuitively, the larger the cost asymmetry (that is, the higher c_2) the bigger the profit margin commanded by firm 1 relative to firm 2, and the easier for 1 to win both auctions; also the higher the share of the sophisticated in the economy (η), the bigger the incentive for firm 1 to gain visibility with type-2 consumers, since they will compare offerings. Letting t vary, and repeating the same exercise, the region where the equilibrium in which 1 wins both auctions exists shrinks as t increases. Intuitively, the higher the transportation

²⁵When a firm has higher costs than the other, we need to have sufficient differentiation for both of them to be able to sell in a duopolistic environment. It can be showed that this amounts to assuming that $t > (4+5c_2)/21$, and so $t = 1/3$ ensures that both firms are able to sell in a duopoly even as $c_2 = 1/2$.

²⁶It is straightforward that firm 1 losing its own auction is never profitable as it would imply losing monopoly profits on type-1 consumers, and firm 2 deviating so as to win also the auction on brand 1 is never profitable because too expensive.

costs, the less fierce the competition between firms, and hence the lower the opportunity for firm 1 to win type-2 consumers.

Finally, we study the price effects of the auctions, when this equilibrium exists. First of all, it turns out that relative to the benchmark prices, the price of firm 1 increases ($\hat{p}_1 > \bar{p}_1 = 1/2$) and that of firm 2 decreases ($\hat{p}_2 < \bar{p}_2 = (1+c_2)/2$). Intuitively, firm 1 now monopolises type-1 consumers, and the payment of the bid on its own auction increases its marginal cost (the payment of b_2^2 instead is a fixed cost), whereas firm 2 now can only sell to sophisticated consumers (all naive buy from 1) and will therefore have to decrease the price in order to have their demand. The overall effect on consumer surplus is therefore a priori ambiguous.²⁷ Calculating the consumer surplus, however, shows that even when the brand search auctions allow the efficient firm to win market share, this is never beneficial for consumers with respect to the benchmark case.

Figure 7 also shows that there is a small region of parameter values where there is an equilibrium where it is the inefficient firm which wins both auctions.²⁸ Although it might seem surprising at first, recall that when firms are perfectly symmetric, and most consumers are sophisticated, there is an asymmetric equilibrium in which either firm can win both auctions. By continuity, if the inefficient firm is only marginally less efficient than the rival, the equilibrium where it wins both auctions still exists. As soon as the cost differences start to be relevant, however, such an equilibrium disappears.

6 Summary, policy counterfactuals and concluding remarks

In this paper we have investigated the market effects of online advertising triggered by search of brand terms. We have made use of a simple model where two firms simultaneously choose the price of their (differentiated) product and the bids for the sponsored ads auction on own and rival's brand keywords, and where sophisticated/attentive consumers (who consider any available information on their screen) and naive/inattentive consumers (who only consider the top link on their screen) are aware of either brand's characteristics and price. These consumers search in order to make the online purchase of the brand they are aware of, provided it gives them non-negative utility.

Relative to a benchmark where only organic search exists, we have showed that in any symmetric equilibrium each firm wins its own brand name auction, and that advertising lowers welfare, due to the following mechanisms: firstly, the sponsored link crowds out the rival's organic link, thus reducing competition and leading to price increases and reduced choice for consumers; secondly, the payment of the rival's bid (we consider a second price auction) raises marginal cost, also contributing to raise market prices. This result is robust to a number of extensions, including the presence of any proportion of fully informed consumers who do not need to resort to search engines to execute their purchase; of more available links so that paid links never completely crowd out organic links; and an alternative timing where

²⁷One has also to consider another negative effect on type-1 consumers: whereas at the benchmark the sophisticated buy from both firms, under the auction they will buy just from 1, implying – other things being equal – a higher disutility for those located closer to 2.

²⁸There also exist equilibria where each firm wins its own brand auction but they are not analysed here.

price decisions are taken before decisions on bids.

We also considered specific functional forms applications where asymmetric equilibria may arise, both in symmetric environments and in asymmetric ones. Among the cases analysed, we have found that the market effect of brand name bidding might be beneficial only in a very asymmetric market where a new entrant is completely unknown to consumers (so all consumers search for the incumbent and nobody searches for the entrant); the search engine does not list the entrant’s link in organic search; and the share of the sophisticated consumers in the economy is large enough (condition for an equilibrium in which the entrant wins the advertising auction on the search for the incumbent’s brand to exist).

Overall, therefore, we find that brand search advertising may well have welfare detrimental effects relative to a benchmark where only organic links exist. This gives some support to the view that agreements not to bid on brand keywords should not necessarily be considered anti-competitive, a view which is arguably in line with the recent *1-800 Contacts* judgment.

One may however wonder whether organic search is the only possible counterfactual to brand bidding. A possible policy may consist of prohibiting brand bidding but allowing for generic product bidding. Suppose for instance that whenever consumers entered a search for “brand X running shoes”, or “brand Y running shoes”, advertisers could just bid on the keyword “running shoes” (and of course could not make the bid conditional on whether a specific brand is or not included in the search query). The analysis carried out in Section 4.5 suggests that consumer welfare is lower under brand bidding than under product search bidding. Therefore, whether the alternative of brand auctions is just organic search results or product search auctions, a policy prohibiting brand bidding would be beneficial.

Another possible counterfactual to brand bidding competition is one where a firm cannot place bids in an auction for a rival’s brand keyword. This is arguably in line with some early judgments, which found that bidding on others’ brandnames would be a violation of intellectual property rights. The owner of the brand would still have an incentive to bid, because the sponsored link would crowd out the organic link of the rival. As a result, each firm would sell only to its ‘captive’ consumers (that is, those aware of its brand and price) and a monopolistic situation would arise. As in our base model, if condition (26) holds, the result stated in Remark 1 follows, and this restricted version of brand bidding also harms consumers. Of course, in the context of our model — where only one bidder would be allowed to bid — the platform would be expected to introduce a reservation price in the auction. Since the reservation price represents an incremental cost, whether this case results in higher or lower prices than under brand bidding competition will trivially depend on the magnitude of the reservation price relative to the losing bid in the competition scenario.

Our model was the simplest we could devise that allows to identify the main effects of brand search advertising. Being very stylized, there are a number of dimensions in which it could be extended so as to possibly give rise to new insights. In particular, we have considered s_i and s_j , the shares of the population aware of brands i and j , as exogenous. In a richer model, before competing in market prices and for the adjudication of search ads, firms may invest in display advertising, which would endogenously determine s_i and s_j . This would also formalize the complementarity between display and search advertising

which has been emphasized in recent policy reports on digital advertising. Or the search engine, whose role is completely passive here, may be modeled as an active player, which could for instance endogenously determine the number of sponsored v. organic links that it intends to have on the search engine results page. Or again, the model could be adapted to consider intra-brand (rather than inter-brand) bidding competition: some recent antitrust cases concerned situations where a producer of goods or services was competing downstream with other retailers of its own brand, and our formalization might shed insights on the effects of those cases too.

7 References

- Anderson, S. P. and R. Renault (2021). Search Direction: Position Externalities and Position Auction Bias. CEPR DP 16724.
- Arbatskaya, M. (2007). Ordered search, *Rand Journal of Economics*, vol. 38, pp. 119-27.
- Armstrong, M. and Zhou, J. (2011). Paying for prominence. *The Economic Journal*, 121(556), 368-395.
- Athey, S., and G. Ellison (2011). Position Auctions with Consumer Search. *Quarterly Journal of Economics* 126 (3): 1213-70.
- Blake, T., Nosko, C., and Tadelis, S. (2015). Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment. *Econometrica*, 83(1), 155-174.
- Chen Y., and C. He (2011). Paid-Placement: Advertising and Search on the Internet. *Economic Journal*, 121: F309-F328.
- Coviello L, Gneezy U., and Goette L. (2017). A large-scale field experiment to evaluate the effectiveness of paid search advertising. CESifo Working Paper Series No. 6684.
- Decarolis F. and G. Rovigatti (2021). From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising. *American Economic Review*. 111(10): 3299-3327.
- De Cornière, A. (2016). Search Advertising. *American Economic Journal: Microeconomics*. 8(3): 156-88.
- Desai, P.S., Shin, W. and Staelin R. (2014). The Company That You Keep: When to Buy a Competitor's Keyword. *Marketing Science*, 33(4): 485-508.
- Edelman, B., M. Ostrovski and M. Schwartz (2007). Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review*. 97(1): 242-259.
- Golden, J. and J.J. Horton (2021). The effects of search advertising on competitors: an experiment before a merger. *Management Science*, 67(1): 342-362.

- Haan M.A. and Moraga-Gonzalez, J.L. (2011). Advertising for attention in a consumer search model. *The Economic Journal*, 121(552), 552-579.
- Moulin, H. (1984). Dominance Solvability and Cournot Stability. *Mathematical Social Sciences*, 7, 1998.
- Mussa, M. and S. Rosen (1978). Monopoly and product quality. *Journal of Economic Theory*, 18, 2: 301-317.
- Narayanan S. and K. Kalyanam (2015). Position Effects in Search Advertising and their Moderators: A Regression Discontinuity Approach. *Marketing Science* 34(3): 388-407.
- Ollar, M. and A. Penta (2017). Full Implementation and Belief Restrictions. *American Economic Review*, August, 2017.
- Ronayne David (2021). Price Comparison Websites. *International Economic Review*, 62(3): 1081-1110.
- Shaked, A. and J. Sutton (1982). Relaxing Price Competition Through Product Differentiation. *The Review of Economic Studies*, 49(1): 3-13.
- Simonov A., and S. Hill (2021). Competitive Advertising on Brand Search: Traffic Stealing and Click Quality, *Marketing Science*, 40 (5).
- Simonov A., C. Nosko and J. Rao (2018). Competition and Crowd-out for Brand Keywords in Sponsored Search, *Marketing Science* 37 (2): 200-215.
- Varian, H. (2007). Position Auctions. *International Journal of Industrial Organization*, 25(6): 1163-1178.
- Sviták, J., J. Tichem and S. Haasbeek (2021). Price effects of search advertising restrictions. *International Journal of Industrial Organization*, 77, 102736.
- Waters, R. (2013). Google criticised as product listing adverts push up prices, *Financial Times*, November 24.
- 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.
- Zhou, J. (2010). Ordered search in differentiated markets, *International Journal of Industrial Organization*, vol. 28, pp. 253-62.