

Sponsored Search Advertisement and Consumer Prices: An Empirical Investigation

**PRELIMINARY DRAFT – PLEASE DO NOT CIRCULATE
EMAIL AUTHORS FOR MOST RECENT VERSION**

Eduardo Schnadower*, Idris Adjerid⁺, Alessandro Acquisti*¹

*Carnegie Mellon University, ⁺University of Notre Dame

Abstract

Targeted advertising is a form of online advertising that relies on knowledge of consumer data to tailor promotional messages. Targeted ads have raised privacy concerns, but the advertising industry has extolled its economic benefits for consumers: via targeted ads, consumers are presented with promotions for products and services they are interested in, reducing their search costs. The effect that targeted advertisement has on consumer welfare, however, depends in part also on the impact targeted ads have on the prices consumers end up paying for advertised products. To date, little is known about the relationship between targeted advertising and consumer prices. We present an ongoing empirical investigation of that relationship, focusing on a specific type of targeted advertising: sponsored search ads. We mine data from both organic and sponsored search results on a popular search engine while searching for a large array of products to compare prices and quality across the results. We present our experimental procedure as well as results from a pilot that highlights preliminary insights on the way targeted ads may impact consumer prices.

1. Introduction

Advertising on the Internet increasingly consists of “targeted” ads. A common and lucrative form of online targeted advertising are “contextual” ads. Contextual ads target consumers based on the content of the page they are visiting, from which they infer consumers’ interests at that moment. A contextual ad, for example, could be an ad for a baby stroller appearing to a consumer visiting a page

¹ eschnado@andrew.cmu.edu, iadjerid@nd.edu, acquisti@andrew.cmu.edu.

with content related to babies or family life. A potentially more invasive version of contextual advertising are ads populated based on the content of a private email that a consumer is currently reading (a common practice by some email clients). While the use of potentially sensitive data for targeting ads has raised privacy concerns since the early days of the commercial Internet (Wang, Lee and Wang, 1998), one of the arguments put forward by proponents of targeted advertising is welfare enhancement (Evans, 2009): targeted ads, it is said, present consumers with advertisements that are more relevant to them, thus reducing their search costs for products, while at the same time helping vendors target specific customers more likely to buy, reducing their advertising costs. In this manuscript, we examine the extent to which contextual targeting affects product purchase options available to consumers. We focus on a particular form of contextual advertising: sponsored search results.

Sponsored search results are results that appear in search engines following a user's search, in addition to so-called "organic" results, and for which advertisers pay a fee based on a cost per click. Typically, vendors compete for sponsored results via second-price auctions. Their bids in such auctions take some user data into account --- such as the search term (which can be used to infer user's interests), location, and device type. While sponsored ads based on search terms may sometimes be innocuous from a privacy perspective, a variety of search terms can raise significant privacy concerns in the context of advertising. For example, a recent Pew report (Fox and Duggan, 2013) finds that over 59% of Americans make medical searches online; these searches include those on depression, suicide, sexually transmitted diseases, rape, and so on. Selling sponsored ads based on these search terms (e.g. to drug companies) would likely be profitable for search engine but may also be perceived as invasive by most consumers. The privacy concerns associated with the sensitivity of some search terms is confounded by the use of location information when generating sponsored ads; prior research

finds that location information can be used to infer characteristics about the user performing the searches (Duckham and Kulik, 2006).

The extent to which online advertising in general, and sponsored search advertising in particular, affects consumer welfare depends, among other things, on the effect that it has on the prices consumer will pay for advertised products. To date, however, little empirical research has investigated the relationship between sponsored search advertising and consumer prices. We mine price and advertising data for a large sample of products sold online using a popular search engine to analyze the relationship between sponsored search advertisements and consumer prices across those products. Although some previous empirical research has investigated how sponsored advertisements affects welfare, that work has tended to be limited in scope: previous research either did not explicitly consider product prices, or was limited to a specific product category or even a single company (see Section 2). By engaging in a large amount of searches that include thousands of products across different vendors, we aim at better understanding the dynamics connecting sponsored and organic search results with prices, vendor quality, and consumer search cost.

We compare prices offered through search-sponsored advertising with prices for identical products that appear in organic search results, while controlling for factors such as vendor quality and search costs. We identify a random sample of 2,000 unique product models across a variety of product categories sold online. For each product model, we perform a search on a popular search engine. We scrape, and then analyze, both organic and sponsored search results, capturing variables such as price, shipping costs, vendor, type of results and order of appearance. By ensuring that the search terms are related to specific product models, we aim at making precise price comparisons between advertised and organic results for identical product models. That is, we make comparisons among different vendors for the same product model. We also control for several aspects that affect positioning in advertisements, such as time of day, day of the week, and location. Finally, we include product model

and vendor fixed effects to capture account for differences in prices between vendors and product types.

The study is ongoing. So far, we have conducted a pilot search for 72 products. All searches were conducted on Google. By introducing automation, we expect to capture all of 2,000 products by the end of June 2018. In this version of the manuscript we report on the results of an analysis applied to the initial set of 72 products. In this preliminary analysis, we already observed trends with respect to prices and sponsored ads. While sponsored results frequently offer better prices than organic search results, that was not regularly the case: for various products, organic search results offered the best prices. As we have not yet analyzed quality data, we refrain from forming any conclusion from the pilot, and offer instead some possible explanations to scrutinize and vet with the full set of data.

The relationship between sponsored searches and product price is, ultimately, likely to depend on consumer and vendor strategies. For example, a quality-minded consumer with high search costs might assume that the first results are always the highest quality ones, whereas a price-minded consumer with a low valuation of time could dig deeper into the results in search for a better price. Different consumers with different preferences in search costs, quality and price may have varying strategies. Vendors, for their part, must consider these possibilities in their bidding strategies, depending on the type of consumer they are targeting. By considering vendor quality and search positions, we aim at not just capturing trends relating sponsored search to prices, but also understand the dynamics behind those trends.

It is important to note that, by focusing on specific product searches, we consider users who are likely considering buying the product being searched, which may not always be the case. Also, as this study only considers search based advertising, we will not consider behavioral based targeting. We also do not address actual consumer behavior, but rather, obtain data of what the consumer is offered. These limitations are mentioned in detail in section 6 and will be addressed in future work.

The rest of the paper is organized as follows. In Section 2 we review previous work in the impact that advertising has on consumer prices. In Section 3 we discuss in detail our empirical approach. In Section 4 we offer details on our data collection procedure. In Section 5 we describe the results of our first pilot and the steps needed to complete the current study. In Section 6 we describe current study limitations and how they may be addressed in future work. Section 7 concludes.

2. Literature Review

2.1. The Impact of Advertising on Consumer Welfare and Price

The impact of advertising on consumer welfare has been a disputed subject among economists. One line of research has argued that advertising benefits consumers by reducing search costs and bringing additional information (Mitra and Lynch, 1996; Akerberg, 2001; Benham, 1972; Nelson; 1974). Under this angle of analysis, intensity of competition is increased, customers are more informed about product, and, as an indirect consequence, prices are lowered thanks to advertising. Another line of work has argued that the opposite can occur, because advertising can also be used to increase market power (Alston, Chalfant and Piggott, 1999; Krishnamurthi and Raj, 1985; Boulding, Lee and Staelin, 1994; Comanor and Wilson, 1979; Nichols, 1985; Dixit and Norman, 1978, Pigou, 1920). Under this alternative scenario, advertisements' main purpose is not to inform, but rather to lure customers away from the competition, which decreases price elasticities of the advertised products. For instance, Pigou argued that inefficient companies who advertise may do so to hinder consumers' comparisons (Pigou, 1920). Conversely, Nelson proposed that low-cost firms can increase sales by both lowering their price and expending more on advertising (Nelson, 1974). Nichols (1985), proposed that whether or not a profit-maximizing amount of advertising improves consumer welfare depends on the market structure and the underlying characteristic of a product. For example, a racquet is the product that is bought, but does not cause utility by itself; its "characteristic," (playing tennis), does, and how advertising affects the utility of playing tennis is what determines whether consumer

welfare is increased by advertising. Others still have recognized that advertising may have simultaneously both positive and negative impacts on consumers (Becker and Murphy, 1993).

2.2. Welfare Implications of Online Ads

How the dynamics underlying traditional advertising translate to *online* advertising is not obvious. Online advertisements are generated dynamically each time a user visits a website, instead of being static (as most offline advertisements). Such advertisements are often generated using vast amounts of consumer data, and tailored towards specific consumers who may be more likely to purchase a product. The implications of such usage of consumer data are nuanced. For instance: on the one hand, targeting may reduce consumer search costs. On the other hand, the same consumer information used in targeting may also be used to infer consumers' reservation prices and be used for first degree price discrimination (Tanner, 2014); in turn, were such pricing strategies to become known to consumers, they may draw significant consumer backlash, as happened in 2000 with Amazon (CNN, 2005).²

Various streams of literature are attempting to understand the economic impact of online advertising, and various theoretical pieces have focused on the relationship between targeting and consumers' economic outcomes. Johnson proposes a model in which consumers benefit from the improved relevancy of ads but are negatively affected by increased volume of advertising (Johnson, 2013). Anand and Shachar suggest that, by using targeted advertising as a signal, targeted ads may convey beneficial information to consumers (Anand and Shachar, 2009). Iyer, Soberman and Villas Boas propose that by identifying comparison shoppers, targeted advertising can be used to persuade them to buy a company's product rather than the competitors' (Iyer, Soberman and Villas-Boas, 2005). These papers collectively suggest that the addition of targeting to traditional advertising brings additional complexities to the impact of ads on prices consumers face and, by extension, their welfare.

² Price discrimination with coupons may still be a more common practice rather than differential pricing, as the former is more widely accepted by consumers than the latter, even though they may be monetarily equivalent (Tanner, 2014).

2.3. Sponsored Searches and Consumer Price

Regarding, specifically, sponsored search advertising, some works have studied positioning of the advertisements and how that is affected by both consumer and vendor strategies. Other works have studied the relationship between ad position and quality. Those works offer important insights regarding the economic implications of sponsored search results, but do not focus on comparing the latter to organic results.

Regarding ad positioning, it is well known that top positions receive more clicks (Ghose and Yang, 2009; Athey and Nekipelov, 2010). However, because clicks do not mean actual purchases, and because many consumers may end up buying a product from different links, this is not always the most profitable strategy (Agarwal Et al, 2011; Ghose and Yang, 2009). The position in which each advertisement appears in a search result page, if it appears at all, is usually determined by a second-price auction mechanism based on the cost per click. Yao and Mela argue that companies that obtain a higher value from each click bid higher. In an experiment, they showed a positive correlation between price and advertiser's bidding (Yao and Mela, 2011).

Another factor that is related to position is quality. By separating between infrequent customers and "expert" consumers, Yao and Mela show that the former are more sensitive to consumer ratings, while the latter are more sensitive to price and slot ranking. They argue that the expert users know that the best quality is usually offered in the first positions, while inexperienced users do not make that assumption. Consumer ratings also affect ad biddings. Athey and Ellison use a theoretical model to show that, when ads are ordered by quality, a higher quality advertiser has a higher marginal benefit for increasing the bid, while lower quality vendors have fewer incentives to compete, and therefore may end up being crowded out (Athey and Ellison, 2011). Price and quality have an important interaction in the positioning and click-through of ads. Animesh, Ramachandran and Viswanathan show that in a low-competition setting, quality-based ads perform better than price-based ads in the

top positions, and this is inverted in the lower positions (Animesh, Ramachandran and Viswanathan, 2010). This is explained by the fact that price-seeking customers have lower search costs than quality-seeking customers, and therefore have a lower marginal search cost. This effect, however, is reduced in a high competition intensity scenario. In our data, we expect to see a lot of interactions between price, quality, and the order of the results. This will help us examine how this might affect consumers with different strategies.

Relatedly, a “Position Paradox” has been proposed by Jerath, et. Al (2011). They argue that, besides position ranking, a website name alone can be enough to convey quality to consumers. Therefore, even if ranked in lower positions, a superior company can receive more clicks than an inferior company in a higher ad position. In this situation, because there are fewer incentives for a high bid by the superior company, the inferior company can try to obtain the highest bid to appeal to consumers who don’t know which company is the superior one (Jareth Et al, 2011). In our study, this could be observed if well-known vendors appeared in lower positions, or not at all, in sponsored search results, while lesser known ones appeared in higher positions.

It is important to note that paid advertising is not the only way in which online companies can make themselves known. Search engines, besides advertisements, show search results based on a “relevancy” algorithm. Websites that appear in those results do not pay for them, and therefore they are called “organic” search results. Yang and Ghose, who performed their analysis on the Google search engine, show that there is a positive interaction in click-through rates between organic and sponsored search results (Yang and Ghose, 2010). That is, when a website appears in both organic and sponsored results, the click-through rates of both links increase. There are well known techniques --- collectively known as “Search Engine Optimization” (SEO) --- that companies apply to their websites to improve their ranking in organic search results and increase traffic (Lee Et al, 2016). Also, keywords of different nature can have different impact on the profitability of both organic and search

results (Ghose and Yang, 2008). Thus, competition strategies might involve a website investing in appearing in both types of results (sponsored and organic), or, conversely, companies that are crowded out of sponsored search results may invest more heavily in SEO.

To our knowledge, most empirical studies on the relationship between sponsored search and consumer welfare have been limited - either because they focus on narrow product categories (such as the software search engine used by Yao and Mela, 2009), or on a limited set of keywords (such as Animesh, Ramachandran and Viswanathan, 2010), or on specific retailers (such as the national retail chain mentioned by Ghose and Yang, 2008 and 2009). In addition, none of the above studies compared prices between organic and sponsored search results. To our knowledge, our study is the first to compare sponsored to organic prices across a wide range of products.

3. Empirical approach.

The objective of our study is to evaluate how (contextually) targeted ads relate to consumer welfare. We focus on sponsored search results, and compare prices associated with both organic and sponsored search results on a popular search engine.

By focusing on sponsored search results, we aim at comparing, across vendors of the exact same product model, whether sellers that appear in sponsored ads sell the product at different prices than what the user could obtain through organic results. In addition to prices, consumer welfare in this context is going to be determined also by vendor quality and search costs. In the following sections, we discuss in detail each of these components.

In this version of the study, we focus on searches for specific, precise product models (as further defined below; in future studies, we plan to apply a similar procedure to broader searches for product types). Because our context of study deals with highly specific searches, it most closely resembles consumers who are performing “transactional” searches --- that is, that they are at least considering obtaining the model in question (a transactional search is defined as a search in which “the intent is

to perform some web-mediated activity”); Broder, 2002).³ The fact alone that advertising for products appears when users search for a specific model is suggestive that vendors have certain expectations regarding the probability that users’ searches are, in fact, transactional.

3.1. Prices

As noted earlier, it is a matter of continued debate whether consumers face higher or lower prices due to, or merely in the presence of, (online) advertising. Different theoretical arguments can lead to a range of predicted effects.

First, when consumers click through an advertisement, that click is more expensive for the advertising merchant than a click on an organic search result. As the click on the advertisement does not guarantee a purchase, the marginal cost of each advertised product increases, which --- under the assumption of competition pushing prices towards vendors’ marginal costs --- may lead to an increase in the price presented to the consumer by advertising merchants.

As we noted earlier, if a company wants to cater to customers who are more priced-focus, it may be willing to sacrifice ranking in results to reduce costs; on the other hand a vendor that is focused on quality may be more incentivized to appear among the first results. It is also possible, of course, that the vendor that offers the best price also offers the best quality. In our context, as we compare vendors for the exact same product, quality is about the service provided by the vendor, not the product itself. In short, different vendors could choose different (bidding and pricing) strategies, attempting to attract consumers that have different preferences in terms of price and quality. In fact, according to the above-mentioned “position paradox,” well-known vendors might choose to sacrifice position in order to reduce costs.

³ Searches could also be informational in some rare cases (i.e., intended to acquire information, such as looking for the product manual). In these cases, differences in prices between sponsored and organic results (and maybe even ads in general) would be less relevant.

Second, while a click on an organic search result does not result in a per-click cost for the linked merchant, appearing at the top of organic results is not easy and can require significant investments by merchants (particularly lesser known and smaller firms). Merchants engage in several manipulations of their sites to increase their probability of showing up on searches in a high position given because appearing first in the search results can impact users' perceptions about a website (Dou et al 2010), and Lee et al. show how these manipulations can help increase website traffic (Lee et al 2016). Vendors therefore have a strong motivation to invest in obtaining high positions in organic search results, which means that higher positioned companies may also face higher costs. However, as these manipulations are done at the website level and do not generate a cost per click, they may be considered akin to fixed costs. Hence, these manipulations may affect prices differently than advertised results. Furthermore, as there may be an interaction between organic and sponsored results (meaning that appearing in one might attract more clicks to the other and vice-versa), it is also possible that the same companies who invest in advertisements could also invest highly in SEO.

Finally, as sellers and prices are shown side by side in sponsored advertisements, price competition is most likely an important element in obtaining clicks among sponsored results, while most organic search results require clicking to actually know the price, meaning that price is less salient in organic results.

The above observations collectively suggest that the relationship between sponsored ads and product prices, relative to product prices in organic results, is not obvious. By mining information from the search results for a wide range of products, we aim at detecting possible different patterns in the prices shown to users and therefore untangle how these different effects play a role.

3.2. Vendor quality

When a consumer has a choice between buying the same exact product from Amazon or from an obscure website, price is not the only factor she may take into consideration. Assuming that the

obscure website has a better price, questions the customer might take into account include: “can I trust them?,” “is the item going to arrive in a certain amount of time and in good condition?,” “are they truly going to send the item that the website says I’m buying?,” and so forth. For an extra level of safety, users might be willing to pay higher prices from better known vendors with better service and guarantees.

As vendors may use position in either sponsored or organic results to convey quality information to the customers, we may see nuanced interactions between price and quality at different positions of the search results. One of the main objectives of this study is to better understand those interactions. To do so, in addition to pricing data, we mine merchant quality information from independent sources to be able to take vendor quality into account. If a user spends time trying to obtain a better price, and this better price is offered by a low quality vendor or scam website, she is likely to be worse off. But if the better price comes from a well-known website, or even an unknown website that accomplishes the order successfully, it means that the time investment paid off.

An important aspect of vendor quality is security. In our pilot, we encountered numerous websites that were most likely fake websites that resemble well-known legitimate stores, as the entire structure was replicated, including the look and feel and third-party vendors, but they always had significantly lower prices. We suspect they intend to collect the customer’s information for illegitimate purposes or charge hidden fees. How pervasive this problem is in organic and sponsored results is something that we analyze in this work.

3.3. Search costs

It takes time for consumers to identify and choose from whom to buy a product. We refer to this time-cost as a “search cost.” Several factors can affect search costs. When the user finds the exact product they want at a price they are willing to pay with the first result they see, the search cost is at a

minimum. However, this cost increases as the user goes through an increasing number of results to find the lowest price in the market or a price compatible with her reservation price.

It is generally assumed that consumers look at search results sequentially (Athey and Ellison, 2011). Even if they don't necessarily click on every result, the decision whether to look at a website or not is typically taken in a top-down order. Therefore, customers who look at results further down the list usually incur higher search costs than customers that look only at the first results. However, when the results are side by side, and when they convey price and quality information (even if it is just the vendor's name, which may be enough to inform a customer about quality) without requiring a click, the search costs incurred in comparing are significantly less than when a click is needed (the latter is the case with most organic results, as they don't usually show prices).

But it is not only a matter of what the position of the best offer is in the results. Features that may appear on different websites can increase search costs, limiting the ability of the user to do a thorough comparison of all the options they have. In our pilot, for instance, we observed some discrepancies between the price shown in the ads and the actual price that the website offers once the ad is clicked. This price difference, if not favorable to the user, might cause disappointment and distrust, and is therefore likely to reduce the probability of a user buying the product from that site. This increases the search cost for the consumer, and creates a cost for the vendor in the form of a non-converted click as well as reputation loss. How often this occurs is, therefore, an important matter to consider.

Another issue that increases search costs and which we have observed consists in websites that are not actual sellers, but pretend to be shopping websites. When clicking on the "buy" link on such sites, the user is redirected to a major seller's website. When the user is trying to do comparison shopping (and is expecting to see a different option), ending up looking at results she may have already seen creates a wasted effort and increases search costs. This practice, however, might be convenient

for the vendor, as it is a common technique for improving organic search position (Langville and Meyer, 2011).

Sometimes, in the results, websites appear that do not actually sell the product, but instead only offer reviews, videos, and products that are not what the user is actually searching for, as well as other irrelevant information. It is not always easy to immediately disregard such websites based on the information displayed by a search engine in the results pages, and therefore it is necessary to click on it to be certain. Reviews and videos can have different effects on search costs. If they bring additional information that allows the user to narrow the search, then search costs may be reduced. But if instead the user is already familiar with the product and is only looking for a seller to purchase it from, the impact on search costs may be nuanced. If the review site links to a seller which the consumer has not considered, it might still reduce search cost; but if it does not link to any seller or to a seller the user is already familiar with, search costs will be increased due to the lost time.

In order to consider search costs in our analysis, we capture the order in which each of the results appeared. Once this is known, we can estimate search costs based on the ordering of the results and the number of non-vendor sites that appear in both organic and sponsored results. This allows us to estimate whether the price obtained when doing a more thorough search is offset by the additional search costs.

4. Data collection procedure

In this section we describe the data collection procedure. First, we introduce key terms that will be used in the rest of the manuscript. Next, we describe how product models to be used in the search process are selected. Finally, we describe the data that is being captured in the search process.

4.1. Terminology

To select products to use in the search process, we define a hierarchy of categories:

- **Product Category:** is the most general categorization of products for sale online. These categories are identified by examining major online retailers. Examples of categories are electronics, sports equipment, travel and luggage, among others.
- **Product Subcategory:** Each product category is further divided into several narrower divisions, such as audio and video, camping equipment, and so forth.
- **Product Type:** A product type is a term that refers to a specific kind of product within a general category. For instance, within electronics product types include televisions, headphones, and wearables.
- **Product Model:** A product model is a specific instance of a product type which is sold by one or more vendors and has specific and precise characteristics. If two products are different in size, color, appearance, bundled items, amount of each item included, or other features, they are considered different models, even if otherwise identical. For example, if a box of pencils is sold in a pack of 6 or 12 units, even if it is the exact same pencil, each pack size is considered a different “model.” Also, if a certain kind of stroller is offered in blue, red or black, each color version is of a different “model.”
- **Organic Result:** Results that appear on search engines that are not paid for and appear due to a “relevance” algorithm used by the engine which is not revealed in full to either consumers or sellers.⁴
- **Sponsored Result:** Sellers that wish to show their pages more prominently on the search results pay fees for their offers to be shown as advertisements. These results are typically shown to the consumer as sponsored, and can be of different kinds, as shown in Figure 1.

⁴ Google, for example, states that they use over 200 factors for determining relevance, but the only one they explain is “PageRank”, which is a measure that is based on links from external pages. That means that the higher the number of external pages that link to the website, the higher the probability is of a webpage being shown high on the organic results (Google, 2017a).

Typically, a seller is charged when the consumer clicks on the advertisement, regardless of whether a final purchase is made.

- **Consistent Search:** A “consistent search” means that a search string exists which for a product model generates results in both organic and sponsored results that are, by and large, for the same exact model. For example, a search for “box of 12 pencils, HB, brand X,” even if the search string clearly indicates 12 pencils, may return offers for the 6-pencil packaging a 24-pencil packaging, or even HB pencils of other brands. When that occurs, the result is considered an inconsistent search. On the other hand, a search of “Smart TV brand X model XYZ-123” may return several results with that exact same model of Television. The latter is an example of a consistent search. See Figures 2 and 3 for examples of consistent and inconsistent searches.

4.2. Product selection

For the product selection phase, we leverage the product categories used by three major online retailers: Walmart, Amazon, and Target. We made sample searches for product models across all their categories and determined which product types would be suitable for the study. Suitability was determined based on the following criteria:

- The product type must have products that can be consistently searchable, based on our definition of consistent search.
- The searches must include advertisements. For example, some obscure niche products may generate consistent organic results, but no advertisements. Because the focus of the experiment is to compare organic with sponsored results, without sponsored results, there is no comparison to be made.

Based on that criteria, we selected 865 suitable product types across 232 subcategories and 11 categories (see Table 1 for examples of products in each of the 11 categories).

The next phase of product selection consists of randomly selecting 100 of the 865 usable product types, to generate a representative sample of product models to include in our analysis. For these selected product types, we use a commercial product database⁵ to generate the universe of product models that are currently for sale across online vendors. This online database contains millions of different products across several categories. It includes offer information when the product is being actively sold online, and has varying update frequencies depending on the characteristics of the products. We tested random products on the database and found that when products have at least one active seller and an update date of 30 days or less, about 70% of the products produce consistent searches under our definition. After filtering with those criteria, we sample 20 products per product type, yielding our target sample of 2000 product models for analysis.

4.3. Data capture

Once the product models are selected, we use the Google search engine to perform searches for each product models, and capture the first two pages of results, in the precise order in which they appear. We do not capture search results presented in pages after the first two, as pilot tests showed that, beyond the second page, results tend to include highly irrelevant information. From each result, we capture: date and time of the search, text string used for the search, name of the website, name of the vendor, listing order (that is, position of the result on the search results page), price shown in advertisement (if it was an ad), actual sales price, shipping costs, additional costs (such as oversize fees or others), type of result (based on the Figure 1 taxonomy), and an indicator of whether the result was consistent with the search string.

It is important to note that advertisement space for search results is sold through an automated auction mechanism. For example, vendors buying search advertising space on Google may determine their bidding strategies by using settings based on keywords, physical location, language, time of day,

⁵ Semantics3 Database

day of the week, type of device and whether the site was previously visited by the user (Google, 2017b). To reduce variation, we only consider the English language in desktop computers. To consider location variability, we complete each search from a VPN,⁶ including locations in several cities in the United States, selected randomly from the list of servers owned by a commercial VPN provider. To control for time of day and day of the week, we randomize the order of the product search, so that all product types are searched at varying times. To prevent any impact on search results of clicking on links during the search procedures, cookies are cleared after each search.⁷

5. Pilot results

The main objective of our study is to explore the prices customers are offered when doing searches for a wide range of product models. Ultimately, we plan to collect information on 2000 searches across 11 categories. However, as it is important to test and perfect the procedure, we randomly selected for a pilot study five product types (iced tea makers, baseball bats, walkers, bike helmets and playards) and captured information on 72 randomly selected product models within those types. Table 2 summarizes our search results. The analysis of the results obtained from the pilot is still ongoing. While our empirical approach includes prices, vendor quality, and search costs, the descriptive analysis of the results done so far is currently focused of prices and consistency of results (which affects search costs). We will, however, consider both quality and search costs in the final analysis of both the pilot and the final data. It is important to note that our current results are merely descriptive, as the sample for the pilot was small. However, they still offer preliminary insights that we plan to scrutinize with the full sample.

⁶ A VPN, Virtual Private Network, is a mechanism in which a computer creates an encrypted connection to a server and then access network resources through that server, which masks the real location and identity of the user. We can thus connect to VPN servers located in different cities to show results as if we were searching from varied locations.

⁷ A cookie is a file that is stored in a local computer when visiting a website, and which can be used by the website to identify previous visitors. When cookies are “cleared” all those files are deleted from the computer and therefore the website cannot obtain the information it originally stored.

In the remainder of this section, we will explain first the consistency results, then the price results, and finally the next steps of the experiment.

5.1 Consistency results

As it was already mentioned, a consistent result is a result that offers the exact product model that is being searched for, including color, size, model, amount, and any other characteristics. It also needs to be in stock and available for purchase from the website. If the product model is different, even if varying only by color or size, it is considered inconsistent. If the product model is out of stock or the site does not sell it (for example, it is a review site), the result is considered inconsistent. The results are grouped by the following categories:

- Organic search: As explained before, it is the search result that is based on the engine's relevancy algorithm and is not paid for. All the results that are not organic, are sponsored.
- Featured snippet: These are highlights of an organic result that Google shows as "answers" to what it assumes to be questions people ask in the search engine. As we got one featured snippet in our searches, we are considering it as part of our analysis.
- Sidebar: It is a bar that offers a single product model and a list of vendors and appears on the right side of the organic results.
- Side tile: It appears in a similar position to the sidebar, but instead offers a tile format. In this kind of ads, it is possible to offer results for different models. It uses a grid layout.
- Top bar: A horizontal bar that, similarly to the side tile, can offer different models, but appears on the top on the page and has a single row of product models.
- Top sponsored links: Links that appear right before and look very similar to the organic search results, but have an "ad" indicator.
- Bottom sponsored links: Same as above but appear after the organic search results.

The results for consistency can be seen in Table 3. The most consistent results come from the sidebar, which were consistent 96% of the time. This is not surprising, as the sidebar is designed, specifically, to advertise sellers of a specific model. It is important to note, however, that --- although not present in the pilot --- in some rehearsals prior to capturing data we observed cases in which the sidebar showed a different *model* than what was searched. When that occurs, the whole sidebar is inconsistent.

In general, the results were consistent most of the time for organic results, the top bar, and the side tile. The less consistent results were offered by the top and bottom sponsored links, which were consistent only about 26% of the time. Because our strict definition of “consistent result,” it is possible that these inconsistencies are part of the intended strategy of the advertisers or the search engine: vendors may want to promote competing products to try to lure away comparison shoppers. We did not capture this information in the pilot, but given the results the pilot offered, we are currently capturing more detailed information about inconsistencies in the current process, which would allow us to know when advertisements are from competing brands, or whether they pertain to a complement of a good (for example, if an ad for an ink cartridge appears when searching for a printer).

Although these results are merely descriptive, they have important consequences for search costs. Every time that a consumer clicks on a link and the need is not met, the search cost incurred to finding the desired product model is increased (Athey and Ellison, 2011). Each click is costly for both the consumer and the advertiser, and therefore it is in the interest of both that results are consistent. The inconsistency of sponsored links is something that we are going to look into more deeply when we gather the full set of data, as it has important implications for both sellers and consumers.

5.2 Price results

Our descriptive analysis of the prices in the pilot suggests that several different vendors’ strategies might be at play. We analyze the results from three perspectives: the “cheapest” approach, the

“average” approach, and the “top 3” approach. It is important to note that the prices reported in this preliminary analysis only come from consistent results.

The “cheapest” approach follows the idea of taking the smallest price available in both sponsored and organic search results. Customers who are very price conscious might spend significant time looking for a price they are willing to pay, so they are going to be looking at a large number of potential vendors.

Table 4 shows the summary results. We found that the cheapest offer was in sponsored results 50% of the time, in organic results 23% of the time, and in both 26% of the time. Across our sample, a consumer who is price conscious could save up to 84% (as compared to the cheapest sponsored offered in sponsored search) by going through the organic search results, and 23% would save at least something. It is important to note, however, that quite often more than one website offered the cheapest price. As a result, when the cheapest price appeared in both organic and sponsored results, it was not necessarily the case that it was the same website or websites. Because there were many cases in which the lowest price was offered in both organic and sponsored results, it is possible that many of those cases come from the same website appearing in both. Of the 72 searches, we observed 19 cases (26%) where the lowest price was offered in both. From those 19 cases, 7 (37%) had the same exact websites with the cheapest price appearing in both kinds of results, while 2 cases (11%) had distinct websites offering such a price in organic and sponsored results. The rest of the cases had a mixture of websites that offered the lowest price in both organic and sponsored, and websites with that same lowest price that appeared in only one of the two.

The “average” approach tells us more about the general behavior of prices. Table 5 shows the results. In this approach, we found that organic have a lower average 40% of the time, while sponsored results average is lower 57% of the time. The remaining 3% of the cases had the same average price. Sponsored prices are, in general, lower than organic prices by 5.17%. Under these circumstances, a

consumer who is price conscious, but also has high marginal search costs, might prefer to look just at the sponsored search results, and in average, would save, but in several occasions, would miss out by not looking at organic results. It is also important to observe the bottom section of the table. There are offers that appear only in organic results, offers that appear only in sponsored results, and offers that appear in both. Although on average the results that appear only in organic results are more expensive, we see great variation in price differences.

Finally, with the “top 3” approach we wanted to investigate how a time-conscious customer who only looked at the first three results in both organic and sponsored results would fare. Table 6 shows the results. We see that even in this approach, close to one fifth of the time the prices were lower for organic results.

Even though we have yet to analyze vendor quality, initial results suggest that sponsored results are more likely to have the cheapest and lower average prices. However, a significant percentage of our searches produced the cheapest price among organic search results and organic search results often do provide lower average prices, is of note. This is particularly notable given that sponsored ads are not costless to consumers since they may be privacy invasive and bothersome to consumers.

5.3 Regression Analysis

One of our goals is to investigate how different factors can affect the relationship between prices and types of search results. Table 7 shows the results of several different regression specifications using log price as dependent variable. All regressions were limited to consistent results only and done using product model fixed effects and standard errors clustered by product model. Column 1 shows that, on average, prices are 6% lower in sponsored than in organic results. This coefficient is reduced by half, but remains significant and does not change sign, when adding controls for time of day, city and hourly group (hourly groups were defined by dividing the day in four equal parts), as seen in Column 2. Columns 3 and 4 repeat the same exercise, but separating sponsored results into several

different categories depending of where they show on the page. We observe the same trend in which sponsored prices are lower than organic prices. Finally, Column 5 uses controls for the position (order of appearance), including the page, and interaction terms between position and type of result. The estimates do not change much. Due to space constraints, it is not possible to add all the coefficients to the table, but some of the non-interacted position coefficients were significant and showed a trend of price increasing with position. This is opposite to what was shown in the studies mentioned earlier by Animesh, Ramachandran and Viswanathan (2010); ad Yao and Mela (2011).

Besides average prices, another relevant measure is the probability of finding the cheapest price in a certain set of results. To measure that we employ a logistic regression in which the dependent variable is equal to one if a search result has the lowest price for the specified product model. Again, all the standard errors are clustered by product model. Table 8 shows the results. Column 1 shows the difference in log odds of finding the lowest priced offer between organic and sponsored results. The coefficient is negative, indicating that it is less likely to be found in organic results, but it is not significant. Column 2 shows the results when the sponsored results are separated by different types. Surprisingly, contrasting with what we observed in the previous regression, it seems that in the sidebar there is significantly *less* probability of finding the lowest price, while the opposite happens with the side tile. Column 3 adds controls for position and page number. The position coefficients are not shown in the table, but they show a significant reduction in probability of finding the cheapest price as the position increases. However, this finding may be biased by our small sample size and by several groups of ads having only few positions (for example, it is possible that the right-hand side ads to only have one or two offers). Therefore, we refrain from drawing conclusions from those specific results. We refrain from adding more controls or interactions under the current sample size. Several values of the independent variables have only one result for the dependent variable, which causes the logistic regression to drop a great amount of observations.

Finally, we estimated how different kinds of consumers, based on their search cost preferences, would behave with the results we obtained. As we don't have any behavioral observations, we do so through simulation. We assume a consumer that clicks sequentially in results as they appear on each page. Each result type (right hand side ads, top bar, sponsored links, etc) shown is a "section." The first result shown in a page is always clicked. From the second result on, the user "clicks" on a result if the expected savings from that click exceed a pre-established percentage criterion. If the next click does not exceed that criterion, then the user passes on to the next section, validating the criterion before clicking or otherwise skipping that whole section. One caveat of this simulation is that, for simplicity, we do not consider the fact that many of the sponsored results show price in the ad and therefore it is not necessary to click them to obtain the price. This is also because, quite often, the prices shown in ads are not exact, and we will collect data on those inaccuracies in our final data collection, but it was not collected for the pilot. Once we have data on how accurate those prices are, we will be able to incorporate it into our simulation. We are also assuming that the user will only click in consistent search results, that is, search results for the precise corresponding product model entered in the search term.

For the expected savings, we defined the "savings" from a click as follows: if the link has a lower price than whatever is the current lowest price that has been seen so far, then the percentage savings is measured as $(\text{previous lowest price} - \text{current price}) / (\text{previous lowest price})$. If the price is the same or higher, the savings from that click are zero. This is because we assume that the user will not buy from that vendor if it is more expensive.

Our savings dependent variable is, therefore similar to a censored variable. However, it is not exactly the same. In our case, the real value of savings is zero, even if the percentage difference in price is negative, because when the product is more expensive, the consumer is going choose from a different vendor, and therefore, the savings obtained for clicking that result are zero, making the

negative difference irrelevant. Therefore, we used a two-part model as developed by Olsen and Shafer (2001). The two-part model works as follows: A continuous dependent variable, y , is used to estimate a binary choice model, in which there are two options: $y > 0$ and $y = 0$. So it begins by estimating $P(Y > 0 | X)$, where X are a set of covariates. Once that probability is estimated, a second part, which can be any continuous dependent variable model, is used to estimate $E[Y | Y > 0, X]$. Using both parts, we can use them to estimate $E[Y | X]$.

The results from that model are shown in Tables 9 and 10. In Table 9 we observe the logistic regression results. Column 1 presents the basic results. Columns 2 and 4 add controls for location, day of the week and time of day. Columns 3 and 4 add controls for position. Our logistic regression estimates tell us that an additional click has more probability of generating savings from the topbar and the sidebar than other kinds of result. Although most of the coefficients in the linear regression of the second part are not significant, probably due to the small sample size, the positive coefficients tell us that, conditional on a website generating savings, all sponsored results except for the top sponsored links generate greater savings than organic results. Even though in this regression the position coefficients were not significant, because we need heterogeneity by each position for the simulation to work, we used the estimates from specification 3. Using those estimates and the rules specified above, we estimated how a customer will behave depending on what percentage is the minimum expected savings that they need to click on a result. Table 11 shows the results of the simulation for six different values. The first three rows show the average number of total clicks, clicks on organic, and clicks on sponsored results. The most extreme cases are the one who always clicks on all results, and the one who only clicks on the first result. The others are determined by the minimum expected savings criterion. We see, as expected, that the number of clicks increase as we lower the criterion, as do the chances of finding the overall lowest price. However, the only first result user seems to obtain the lowest price about 24% of the time. On the other hand, by assuming that they

will always buy from the lowest price they find, and if there is more than one offering such price, they will always buy from the last visited, we determined that our simulated most extreme consumer would buy 36% of the time from organic results.

Although we cannot draw robust conclusions from this sample, we observed some anecdotal cases which are potentially telling of the data we may find during the actual study. There were cases in which a lower priced major vendor did not show in sponsored results but did show in organic, and vice versa. We saw some vendors with lower prices that appeared in sponsored results but only after clicking on “view all sellers and prices” - which means that they were considerably less salient. In addition, we observed some significantly discounted offers that only appeared in organic results. All these examples show the degree of richness of the data that we are capturing, and what we might expect to observe from our analysis.

5.5 Next steps

The results from the pilot suggest that there may be many different dynamics at play between the different categories of sponsored and organic search results. To better understand the implications of these results, we still need to gather the information regarding the vendors that offered the product models, so we can assess their quality. Quality is possibly a significant factor driving the results, and therefore it is essential for a deeper analysis. Also, as the data we gather includes the order in which the results appear, we need to take that into account in order to consider how that affects the search costs that consumers face. However, the “top 3” approach suggests that the search cost to find a cheaper offer in the organic results might not be necessarily high.

At this moment we are capturing the full sample of data. To collect the final data, we select 2,000 products from a commercial product database.⁸ Selecting products from that database is a lengthy process, as it contains millions of products and downloading the necessary information in order to be

⁸ The database is called Sematinc3

able to do the selection takes time. Simultaneously, we are currently collecting data from those products already selected. However, we have not collected yet enough data to do any significant analysis.

It is important to note that the pilot results were obtained manually. In the pilot, the process was the following:

1. Each research assistant would clear cookies, and connect to a VPN using a random US location.
2. The assistant would perform the search for a product they had assigned, and capture manually information that includes: date and time, the location from which the VPN is connected, name of the website, name of the vendor (if it is a third-party vendor), type of result (sidebar, top bar, sponsored links, organic, Etc.), price, shipping cost, other costs, amount that would be needed to obtain free shipping (if there is one), and whether the result is consistent or not, according to the definition given earlier. If it was inconsistent, they would write a text commenting on the nature of the inconsistency.
3. After capturing all the information for that single product, the assistant would delete cookies, disconnect from the VPN and reconnect from another location.

The pilot involved a significant amount of manual capturing of information, and it was a lengthy, slow process. In order to make it feasible to capture 2,000 products, we developed a web scraper that captures most of the information needed by the research assistant. However, due to the broad range of websites that appear in the results, and the different page structures, it is not possible to completely automate the collection of information. Below is a description of what the web scraper does:

1. It connects to the VPN, using a random US based server for each product.

2. It performs the search, downloads the HTML files⁹ of the search results and each link appearing on those results.
3. It generates screenshots¹⁰ from the websites that are linked in the search results of the first two results pages.
4. It creates a spreadsheet that has all of the information that the research assistant captured manually in the pilot with the exception of price, shipping costs, other costs, amount required to free shipping, vendor (for websites that have third party vendors) and consistency of the result. In addition, it adds links to the original search result page and to the locally generated screenshot.

Under our current process, after running the scraper for our batch of products, a research assistant opens the generated spreadsheet and for each row on the spreadsheet:

1. Clicks on one of the generated screenshots (this has zero loading time).
2. Tries to use the information on the screenshot to fill out the missing information.
3. If the screenshot does not have the missing information, clicks on the link to the original result and tries to obtain it from there (loading times and number of clicks required to obtain the information may vary).

In addition, we have codified the most common inconsistencies found in the pilot, trained the RAs on how to spot these inconsistencies, and ask them to indicate why a result was not a match for the model searched. Table 12 shows the most common inconsistencies. The RAs still have the option to include a free text description of inconsistent results and the protocol is flexible to their inclusion as data collection proceeds.

⁹ HTML files are the way in which the content of the webpages is stored. Each file corresponds to a single page as seen in the browser.

¹⁰ As downloaded HTML files do not always render correctly when opening them locally, we generate screenshots of how the webpage looks like at the moment the scraper visits it to make it easier for the RA to see.

The process is ongoing. We expect the team of research assistants to be able to capture around 200 products per week, which would allow us to start the early data analysis by the start of July 2018.

Currently our analysis does not evaluate whether differences in vendor quality exists between sponsored and organic results. This limits our ability to draw robust conclusions about whether consumers are, in fact, better off from the price trends we observe. We will conduct this additional analysis once we have quality data on the vendors in our analysis. Data on quality will be obtained from the following sources:

- Better business bureau: An organization that gives letter grades to businesses based on their practices and consumer complaints.
- www.resellerratings.com and www.sitejabber.com: two community driven online review sites that give grades based on the 0-5 star scale.
- Third party vendor reviews from Amazon, Walmart, Ebay, Sears and Target, which depending on the website may use a 0-100 scale or a 0-5 star scale.

Although reviews and qualifications are not necessarily a perfect measure of a vendor's quality, it has been shown that online reviews affect the customer's perception of quality of an online store (Ulz, et al, 2012; Xu and Kim, 2008). Therefore, we use reviews as a proxy of the vendor's true quality.

Once obtained data on quality, we will be able to incorporate it in our models. Specifically, due to the complex relationship between quality and price, it may be necessary to estimate a simultaneous equations model. Additionally, in our simulations, when quality is observed, it may be desirable to pay a higher price for the same product if it comes from a higher quality vendor. We will take that into consideration. In addition, we will also collect vendor characteristics such as firm size and age, as these may also be related to customer's perceptions of quality.

In the current set of data, probably because of the small size of the sample, location was not a significant factor. However, with the full sample, we will be able to observe whether there is a

significant influence of location in the pricing of sponsored search advertising. Previous studies have suggested evidence of location based price discrimination (Mikianis, et. al, 2012). As in our design we will only perform the searches once for each model once, we will not be able to detect price discrimination per se, but we will observe how location affects the relative prices and qualities of sponsored and organic search results.

6. Limitations and Future Work

Our analysis focuses on contextually targeted ads, and the search terms used imply that the consumer may have a very specific idea of what they intend to buy (assuming the search is transactional). While this assumption is realistic for certain users and products, other models of user behaviors are possible – such as models where users search for broad categories of products (rather than specific models) or, in fact, are exposed to behaviorally targeted advertising for products they had not even searched. As we mention further below, in future work we will expand our analysis both to broad product type searches, as well as to behaviorally targeted (rather than contextually targeted) ads, as a key potential benefit of targeted advertising online is to provide consumers with new products not known to them, but are welfare enhancing for them.

Relatedly, another limitation of this project is that we are confining ourselves to search based advertising, while in real life users are going to see advertisements not only on search pages, but also on web pages that take into account their web browsing behavior to suggest products or websites.

Also, due to the requirement of “consistent search,” some products that have inconsistent results (even if the search terms are specified) had to be excluded. Some examples are clothing, school supplies, food, furniture, jewelry, among others. Despite the exclusion of these products, the range of products available for inclusion in the study is still very broad and should still yield generalizable insights.

Finally, we are currently only taking into account what the user sees, but not what they actually do with what they see. That is, we are not considering consumer behavior.

We plan to address these future limitations in further studies currently in the design phase. First, as noted, we plan to relax the specific product assumption by doing a similar analysis with searches for product *types* (as defined above) instead of product *models*. This will also allow us to include previously excluded products. One of the main obstacles in the design of this experiment is going to be that when searching for product types, while advertised results offer a range of models (see Figure 4), many links in organic results lead to product or category listings (see Figure 5) that do not offer specific products but require further clicking and searching. This analysis, however, will allow us to determine how well the full diversity of available products is represented in the advertisements, and therefore how likely is a consumer to find what he is looking for in the ads (even when they don't know the exact model they want).

A second study is going to incorporate behavioral advertising, that is, advertisements that appear on different websites that take into account the user's browsing patterns, instead of considering only search-based ads. In this study, we aim to determine whether the information brought by behavioral ads offers any benefits to consumers instead of price, quality or novelty (whether the behavioral ad introduced the consumer to welfare enhancing products they had previously been unaware of).

We also plan to consider conducting an experiment with human subjects in which actual consumer behavior is observed. Through such experiment we intend to capture how the users interact with the different kinds of ads, organic search results, and what they end up actually buying, as compared to what they could have bought (based on the previous experiments).

7. Conclusions

While the benefits that online targeted advertising provides to sellers have been widely studied, how targeted ads affect customers is a matter that still requires attention. As more and more personal

information is collected and used to try and sell products to consumers, whether they are benefited or damaged by these practices is certainly an issue that deserves attention, as it could have important implications for policy makers. In this manuscript, we detailed a study aimed at analyzing the impact that sponsored search advertising has on consumer prices. As prices are not the only aspect that affects consumer welfare in online shopping, we also detailed how we are going to consider search costs and vendor quality, which are important aspects of the online shopping experience.

This study, however, has certain limitations because of the assumption of specific model search, the consistent search requirement, only search based advertisements, and not considering actual consumer behavior. We will attempt to address these limitations in future work.

Acknowledgments

This research was in part supported via a grant from the Alfred P. Sloan Foundation. Acquisti also gratefully acknowledges support from the Carnegie Corporation of New York via an Andrew Carnegie Fellowship and from the National Science Foundation (Award # 1514192). For a list of Acquisti's additional grants and funding sources, please visit www.heinz.cmu.edu/~acquisti/cv.htm. Eduardo Schñadower also acknowledges the National Science Foundation for the travel grant given to present at WEIS.

References

- Akerberg, D. A. (2001). Empirically distinguishing informative and prestige effects of advertising. *RAND Journal of Economics*, 316-333.
- Agarwal, A., Hosanagar, K., & Smith, M. D. (2011). Location, location, location: An analysis of profitability of position in online advertising markets. *Journal of marketing research*, 48(6), 1057-1073.
- Alston, J. M., Chalfant, J. A., & Piggott, N. E. (1999). *Advertising and consumer welfare*. Centre for International Economic Studies.
- Anand, B. N., & Shachar, R. (2009). Targeted advertising as a signal. *QME*, 7(3), 237-266.

- Animesh, A., Ramachandran, V., & Viswanathan, S. (2010). Research note—Quality uncertainty and the performance of online sponsored search markets: An empirical investigation. *Information Systems Research*, 21(1), 190-201.
- Animesh, A., Viswanathan, S., & Agarwal, R. (2011). Competing “creatively” in sponsored search markets: The effect of rank, differentiation strategy, and competition on performance. *Information Systems Research*, 22(1), 153-169.
- Athey, S., & Ellison, G. (2011). Position auctions with consumer search. *The Quarterly Journal of Economics*, 126(3), 1213-1270.
- Athey, S., & Nekipelov, D. (2010, May). A structural model of sponsored search advertising auctions. In *Sixth ad auctions workshop* (Vol. 15).
- Becker, G. S., & Murphy, K. M. (1993). A simple theory of advertising as a good or bad. *The Quarterly Journal of Economics*, 108(4), 941-964.
- Benham, L. (1972). The effect of advertising on the price of eyeglasses. *The Journal of Law and Economics*, 15(2), 337-352.
- Boulding, W., Lee, E., & Staelin, R. (1994). Mastering the mix: Do advertising, promotion, and sales force activities lead to differentiation?. *Journal of marketing research*, 159-172.
- Broder, A. (2002, September). A taxonomy of web search. In *ACM Sigir forum* (Vol. 36, No. 2, pp. 3-10). ACM.
- CNN (2005) "Web sites change prices based on customers' habits." CNN. Accessed February 14, 2018. <http://edition.cnn.com/2005/LAW/06/24/ramasastry.website.prices/>
- Comanor, W. S., & Wilson, T. A. (1979). The effect of advertising on competition: A survey. *Journal of economic literature*, 17(2), 453-476.
- Dixit, A., & Norman, V. (1978). Advertising and welfare. *The Bell Journal of Economics*, 1-17.
- Dou, W., Lim, K. H., Su, C., Zhou, N., & Cui, N. (2010). Brand positioning strategy using search engine marketing. *Mis Quarterly*, 261-279.
- Duckham, M., & Kulik, L. (2006). Location privacy and location-aware computing. *Dynamic & mobile GIS: investigating change in space and time*, 3, 35-51.
- Evans, D. S. (2009). The online advertising industry: Economics, evolution, and privacy. *Journal of Economic Perspectives*, 23(3), 37-60.
- Farahat, A., & Bailey, M. C. (2012, April). How effective is targeted advertising?. In *Proceedings of the 21st international conference on World Wide Web* (pp. 111-120). ACM.
- Fox, S., & Duggan, M. (2013). Health Online 2013. Retrieved from <http://www.pewinternet.org/2013/01/15/health-online-2013/>
- Ghose, A., & Yang, S. (2008, August). Comparing performance metrics in organic search with sponsored search advertising. In *Proceedings of the 2nd International Workshop on Data Mining and Audience Intelligence for Advertising* (pp. 18-26). ACM.
- Ghose, A., & Yang, S. (2009). An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Science*, 55(10), 1605-1622.
- Google. (2017). How Google Search Works - Search Console Help. Retrieved November 30, 2017, from <https://support.google.com/webmasters/answer/70897?hl=en>
- Google. (2017). Targeting your ads - AdWords Help. Retrieved August 30, 2017, from <https://support.google.com/adwords/answer/1704368?hl=en>
- Iyer, G., Soberman, D., & Villas-Boas, J. M. (2005). The targeting of advertising. *Marketing Science*, 24(3), 461-476.

- Jerath, K., Ma, L., Park, Y. H., & Srinivasan, K. (2011). A “position paradox” in sponsored search auctions. *Marketing Science*, 30(4), 612-627.
- Johnson, J. P. (2013). Targeted advertising and advertising avoidance. *The RAND Journal of Economics*, 44(1), 128-144.
- Krishnamurthi, L., & Raj, S. P. (1985). The effect of advertising on consumer price sensitivity. *Journal of Marketing Research*, 119-129.
- Langville, A. N., & Meyer, C. D. (2011). *Google's PageRank and beyond: The science of search engine rankings*. Princeton University Press.
- Lee, S., Jang, W., Lee, E., & Oh, S. G. (2016). Search engine optimization: A case study using the bibliographies of LG Science Land in Korea. *Library Hi Tech*, 34(2), 197-206.
- Mikians, J., Gyarmati, L., Erramilli, V., & Laoutaris, N. (2012, October). Detecting price and search discrimination on the internet. In *Proceedings of the 11th ACM Workshop on Hot Topics in Networks* (pp. 79-84). acm.
- Mitra, A., & Lynch, J. G. (1996). Advertising effects on consumer welfare: prices paid and liking for brands selected. *Marketing Letters*, 7(1), 19-29.
- Nichols, L. M. (1985). Advertising and economic welfare. *The American Economic Review*, 75(1), 213-218.
- Nelson, P. (1974). Advertising as information. *Journal of political economy*, 82(4), 729-754.
- Olsen, M. K., & Schafer, J. L. (2001). A two-part random-effects model for semicontinuous longitudinal data. *Journal of the American Statistical Association*, 96(454), 730-745.
- Pigou, A. C. (1932). *The economics of welfare, 1920*. *McMillan & Co., London*.
- Tanner, A. (2014, June 27). Different Customers, Different Prices, Thanks To Big Data. Retrieved February 14, 2018, from <https://www.forbes.com/sites/adamtanner/2014/03/26/different-customers-different-prices-thanks-to-big-data/#17d6c7cc5730>
- Utz, S., Kerkhof, P., & Van Den Bos, J. (2012). Consumers rule: How consumer reviews influence perceived trustworthiness of online stores. *Electronic Commerce Research and Applications*, 11(1), 49-58.
- Wang, H., Lee, M. K., & Wang, C. (1998). Consumer privacy concerns about Internet marketing. *Communications of the ACM*, 41(3), 63-70.
- Xu, Y. C., & Kim, H. W. (2008). Order effect and vendor inspection in online comparison shopping. *Journal of Retailing*, 84(4), 477-486.
- Yang, S., & Ghose, A. (2010). Analyzing the relationship between organic and sponsored search advertising: Positive, negative, or zero interdependence?. *Marketing Science*, 29(4), 602-623.
- Yao, S., & Mela, C. F. (2011). A dynamic model of sponsored search advertising. *Marketing Science*, 30(3), 447-468.



Graco Aire3 Click Connect Stroller - Pierce



All Shopping Videos Images News More Settings Tools

Page 2 of about 10,300 results (0.63 seconds)

[A]

Shop for Graco Aire3 Click Conne... on Google

Sponsored ⓘ

Graco Aire 3 Stroller Click ... \$127.99 Target Free shipping	Graco Aire3 Click Connect ... \$179.99 Target ★★★★★ (8)	Graco Aire3 Click Connect ... \$129.99 BabyDiapersMart Free shipping	Graco FastAction Fold Sport ... \$230.99 Target ★★★★★ (128)	Graco Breeze Click Connect ... \$113.99 Target Free shipping

[B]

Graco Aire3 Click Connect Ultra-lightweight Fold...

4.5 ★★★★★ 174 user reviews



Shop now

Sponsored ⓘ

Pierce ▾

\$127.99 · Target	Free shipping
\$179.99 · Toys R Us	Free shipping
\$179.99 · Bed Bath & Beyond	Free shipping
\$179.99 · Macy's	Free shipping
\$179.99 · buybuy BABY	Free shipping

View all sellers and prices

Reviews

Pros: Grows With Child · Smooth Ride · Lightweight · Easy To Maneuver

View all reviews

Product details

Brand: Graco

Seating capacity: Single

View more details

[C]

Graco® Official Site - Free Shipping Over \$50 - gracobaby.com

Ad www.gracobaby.com/Official ▾

Innovating In Child Safety For 60+ Years. Shop Graco Official Site Online Today!

Safe Seat Engineered · 60+ Years In Business · 24 hr. Customer Support

Types: Car Seats, Strollers, Playards, Highchairs, Swings

Travel Systems

Take Your Infant Or Toddler From Stroller To Car With Just 1 Click!

Jogger Strollers

Shop 3-Wheel Designs For Superior Suspension And Maneuverability.

[D]

Graco Aire 3 Click Connect Stroller - Pierce Reviews - Influenster

https://www.influenster.com › Family › Baby Gear › Graco Strollers ▾

★★★★★ Rating: 5 - 1 review

Aug 5, 2016 - Before you buy Graco Aire 3 Click Connect Stroller - Pierce, check out 1 Influenster reviews. Morenike K. said "Listen up. If you live in a city like ...

Figure 1: Basic schematics of a results page. [A] Shows sponsored results tiles, which can have a different product model in each tile and can appear in the top, bottom or right side of the page. [B] Shows a sponsored result sidebar, which shows different vendors for the same exact product model. [C] Shows a sponsored inline result, which looks very similar to organic (unpaid) results but has a green "Ad" indicator, and can appear either before or after organic search results in each page. [D] Shows an example of an organic search result. Organic results can lead to vendors, product listings videos, reviews, among others, and so it does not always lead to a page in which the product can be bought.



staedtler noris hb pencils pack of 50



All Shopping Images Videos News More Settings Tools

About 369,000 results (0.59 seconds)

Shop for staedtler noris hb pencils... on Google

Sponsored ⓘ

				
Staedtler Noris School Pencils ... \$17.99 Bonanza	Staedtler Noris 122-HB ... \$8.45 Amazon.com	Sharpened Pencils ... \$16.99 Nordisco	Staedtler Noris Norris School ... \$11.82 Bonanza	Staedtler Noris 120 Hb Pencil ... \$10.38 eBay Free shipping

Amazon.com : Staedtler Noris HB Pencil (Tub of 50) : Office Products

<https://www.amazon.com/Staedtler-Noris-HB-Pencil-Tub/dp/B003HIG6OQ> ▼

In the unmistakably iconic black and yellow design these Staedtler pencils are easy to sharpen and erase. Ideal for everyday use at home, school or in the office these **Staedtler HB Pencils** are great for sketching, note taking and much more. This great value **pack of 50** us perfect for bulk buying for office or student use.

Amazon.com : STAEDTLER NORIS SCHOOL PENCILS HB [Box of 36 ...

<https://www.amazon.com/STAEDTLER-NORIS-SCHOOL-PENCILS.../B00M6CPTO2> ▼

★★★★★ Rating: 3.5 - 13 reviews

Amazon.com : **STAEDTLER NORIS SCHOOL PENCILS HB** [Box of 36] : Office Products. ... **Staedtler Noris School Pencils - Pack of 36 HB** Grade Sharpened; The School Favourite for Decades. The **Staedtler Noris** Get a **\$50** Amazon.com Gift Card instantly upon approval for the Amazon Rewards Visa Card Apply now ...

Figure 2: An example of an inconsistent result. Notice how the advertisements and organic results offer different presentations with differing amounts of the pencil, even if the search string clearly states that a pack of 50 is what is being looked for.

XBR X900E - Sony

<https://www.sony.com/electronics/televisions/xbr-x900e-series> ▼

★★★★★ Rating: 4.5 - 220 reviews - \$899.99 to \$2,799.99

4K Ultra HD. ... Smart TV (Android TV™) ... Discover the next level of 4K HDR entertainment with the power of the 4K HDR Processor X1™ and our evolved X-tended Dynamic Range PRO technology for exceptionally high contrast, detail, and clarity.

[XBR X900E Series Reviews ...](#) · [Full Specifications](#) · [XBR49X900E](#) · [Contrast](#)

Sony XBR-X900E series Specs - CNET

<https://www.cnet.com/products/sony-xbr65x900e/specs/> ▼

View full Sony XBR-X900E series specs on CNET. ... Remote Control. Type. remote control. Remote Control Model. Sony RMF-TX200U. Video. Refresh rate. 960 Hz. Connections. Type ... Auto On/Off. Yes. Header. Brand. Sony. Product Line. Sony XBR. Model. 65X900E. Series. BRAVIA X900E Series. TV System. Additional ...

Amazon.com: Sony 4K Ultra HD Smart TV XBR65X900E: 65 Inch LED ...

<https://www.amazon.com/Sony-XBR65X900E-65-Inch-Ultra-Smart-.../B01MZF81NS> ▼

★★★★★ Rating: 4.2 - 253 reviews

The X900E's narrow, exquisitely designed aluminum frame keeps you focused on the screen, while cables stay cleverly hidden at the back and the front. -Sony 4K Ultra HD Smart LED TV (2017 Model XBR-65X900E) -4K HDR Processor X1 -Dimensions without stand: 57" x 32.8" x 2.4" | with stand: 57" x 35.3" x 10.3"

Sony X900E Review (XBR49X900E, XBR55X900E, XBR65X900E ...

www.rtings.com › [TV](#) › [Reviews](#) › [Sony](#) ▼

★★★★★ Rating: 8.1/10 - Review by Daniel O'Keefe

Mar 17, 2017 - The Sony X900E is a great 4k TV that offers some of the best picture quality found in an LED TV. HDR content looks particularly good on this TV since it gets very bright, and it handles motion exceptionally well. Its only real downside is the degradation of the image when viewed at an angle. This TV is ...

Sony 65" Class (64.5" Diag.) - LED - 2160p - Smart - 4K ... - Best Buy

<https://www.bestbuy.com> › [TV & Home Theater](#) › [TVs](#) › [4K Ultra HD TVs](#) ▼

★★★★★ Rating: 4.8 - 1,616 reviews - \$1,499.99 - In stock

Model: XBR65X900E; SKU: 5748207. Rating 4.8 out of 5 stars. ... This Sony HDR Ultra TV is compatible with PlayStation Vue, letting you live stream events as they happen. 4K Ultra HD TV Buying Guide ... Sanus - Premium Series Advanced Tilt TV Wall Mount For Most 42" - 90" TVs - Extends 5.75" - Black. \$89.99. On Sale.

Sony XBR-X900E-Series 65"-Class HDR UHD Smart XBR-65X900E

https://www.bhphotovideo.com/c/.../sony_xbr_65x900e_x900e_series_65_4k.html ▼

★★★★★ Rating: 4.9 - 11 reviews - \$1,498.00

Buy Sony XBR-X900E-Series 65"-Class HDR UHD Smart LED TV features UHD 3840 x 2160 LED Panel, HDR10-Compatible. Review Sony Televisions, TVs & Entertainment. ... Activate An Eligible Sony Bravia 2017 Android TV. Redeem on Google Play Movies And TV App. Offer ends: DEC 31 '18. Share To Win \$1000 ...

Sony BRAVIA X900E Series XBR 65X900E - 65" LED Smart TV - 4...

4.7 ★★★★★ 2,951 user reviews



Shop now

Sponsored ⓘ

65" ▼

\$1,499.99 · Best Buy

Store pickup

\$1,498.00 · B&H Photo-Video-Audio

Free shipping, no tax

\$1,449.00 · EbuyUSA

Free shipping, no tax

\$1,498.00 · Dell

Free shipping

\$1,499.91 · PC Richard & Son

Free shipping

View all sellers and prices

Reviews

Pros: Attractive Design · Easy To Use · Easy To Set Up

[View all reviews](#)

Product details

Release Date: January 2017

Feature: High Definition, Smart TV

Weight: 48.5 lbs

Screen size: 65 in

Brand: Sony

Display type: LED

HDTV format: 4K

[View more details](#)

Similar Televisions For Gaming



Figure 3: An example of a consistent search. Notice that both organic and sponsored search results show sellers for the same exact product model.

strollers 

All Shopping Images Maps News More Settings Tools

About 44,100,000 results (0.60 seconds)

Shop for strollers on Google

Sponsored 

 <p>Uppababy Cruz 2017 Stroller ... \$549.99 Bed Bath & Bey... ★★★★★ (16)</p>	 <p>Infant Bugaboo 'Buffalo' Stroller \$119.95 Nordstrom ★★★★★ (72)</p>	 <p>Infant Mima Xari Black Chassis ... \$1,424.00 Nordstrom Free shipping</p>	 <p>Evenflo Pro Series Pivot ... \$279.99 Target Free shipping</p>	 <p>UPPAbaby 2017 Vista Double ... \$1,039.97 Albee Baby ★★★★★ (120)</p>
---	---	---	---	--

Strollers : Target

<https://www.target.com/c/strollers-baby/-/N-5xtk7> ▼

Items 1 - 24 of 384 - Graco strollers & gear* ... Car seat & stroller toys. ... Baby Trend Snap-N-Go EX Universal Infant Car Seat Carrier already viewed.

Chicco : Strollers · Graco : Strollers · Lightweight Strollers · Britax : Strollers

Baby Strollers - Infant & Toddler - Babies"R"Us - Toys"R"Us.com

<https://www.toysrus.com/products/baby-strollers.jsp> ▼

Babies"R"Us is home to an extensive inventory of baby strollers that keep baby comfortable and secure as you move through the day together. ... To give you the ultimate in peace of mind, select from the most popular infant and toddler stroller models from leading manufacturers ...

Figure 4: Results when searching for a product type instead of a product model. Notice the variation now shown in the advertised results.

Shop by category



Baby Gifting



Full-size strollers



Jogging strollers



Travel systems



Lightweight strollers



Double, triple & quad strollers



Car seat carriers



Premium strollers



Stroller accessories



Car seat & stroller toys



Top stroller registry picks



Stroller deals

shop by brand



Figure 5: Sample landing page from organic search results when searching for “strollers”. Notice that no specific product is offered and more clicking is required.

Category Name	Some (non-exhaustive) Examples of articles
Baby	Strollers, Swings, Cribs, Car Seats
Sports Equipment	Exercise machines, sport-specific gear, exercise gadgets
Outdoors	Camping tents, flasks, sleeping bags, climbing gear
Personal care, beauty and health	Walking aids, hair clippers, curling irons, thermometers, cosmetics
Kitchen and dining	Cookware, silverware, china, coffee makers, kitchen furniture
Home, garden and pets	Patio heaters, coat racks, mailbox covers, post lights
Tools and home improvement	Locks, toolboxes, drills, door bells, light bulbs
Appliances	Fans, humidifiers, refrigerators, electric skillet
Travel and luggage	Backpacks, luggage, travel totes
Electronics	Cameras, TV's, tablet computers, hard drives, car GPS
Office	Paper binder, hand trucks, file boxes, electric sharpeners

Table 1: The 11 categories we found with consistently searchable products.

Product category	Product type	Number of models	Consistent organic results	Consistent sponsored results
Appliances	Ice tea maker	16	147	155
Baby	Playard	14	101	135
Personal care, beauty and health	Walker	12	103	149
Outdoors	Bike helmet	13	109	113
Sports equipment	Baseball bat	17	110	168
Total		72	570	720

Table 2: Summary of the search results.

Summary	Total	Bottom sponsored link	Featured snippet	Organic results	Sidebar	Side tile	Top sponsored results	Top bar
Consistent	1290	62	1	570	153	189	48	267
Inconsistent	1154	186	0	467	6	103	145	247
Total	2444	248	1	1037	159	292	193	514
%Consistent	52.78%	25.00%	100.00%	54.97%	96.23%	64.73%	24.87%	51.95%

Table 3: Consistency of results across all product searches.

Cheapest price approach	
Percentage of time the cheapest offer was in	
Organic	24%
Sponsored	50%
Both	26%
Maximum differences between cheapest offers in organic and sponsored	
When organic was cheaper	84%
When sponsored was cheaper	60%
Average differences between cheapest offers in organic and sponsored	
When organic was cheaper	22%
When sponsored was cheaper	23%
Average difference in general (negative indicates sponsored is lower)	-5.56%
When the cheapest prices appeared in both:	
Times the vendors were the same in organic and sponsored	37%
Times all vendors in organic and sponsored were different	11%

Table 4: Summary information about the cheapest price found in organic and sponsored search results. Total number of products searched: 72

Average price approach			
Percentage of time average price was lower for			
Organic			40%
Sponsored			57%
Both			3%
Maximum differences between average offers in organic and sponsored			
When organic was lower			29%
When sponsored was lower			44%
Average differences between average offers in organic and sponsored			
When organic was cheaper			7%
When sponsored was cheaper			12%
Average difference in general (negative indicates sponsored is lower)			-5.32%
Average price differences between sites that appear			
	Average	Minimum	Maximum
Only in organic vs only in sponsored	13%	-45%	111%
In both vs only in sponsored	-0.6%	-54%	49%
Only in organic vs In Both	14.70%	-11%	121%

Note: negative indicates the element on the left is lower.

Table 5: Summary information for the average prices found in organic and sponsored search results. Total number of products searched: 72

Top 3 approach	
Percentage of time minimum price was lower for	
Organic	19%
Sponsored	39%
Both	44%
Maximum differences between lowest offers in organic and sponsored	
When organic was lower	32%
When sponsored was lower	64%
Average differences between lowest offers in organic and sponsored	
When organic was cheaper	8%
When sponsored was cheaper	23%
Average difference in general (negative indicates sponsored is lower)	-7.05%

Table 6: Summary results of the top 3 results in organic vs sponsored. Total number of products searched: 72

	(1)	(2)	(3)	(4)	(5)
	logprice	logprice	logprice	logprice	logprice
Sponsored result	-0.0647*** (0.0173)	-0.0336*** (0.0111)			
Organic result	0 (.)	0 (.)			
Featured Snippet			-0.156*** (0.0130)	-0.153*** (0.0499)	-0.0988*** (0.0220)
Organic			0 (.)	0 (.)	0 (.)
Sidebar			-0.0495* (0.0280)	-0.0214 (0.0140)	-0.0340 (0.0363)
Topbar			-0.0632** (0.0265)	-0.0366** (0.0145)	-0.0452 (0.0445)
Bottom sponsored link			-0.0503 (0.0322)	-0.0256 (0.0240)	-0.0285 (0.0490)
Sidetile			-0.0895*** (0.0197)	-0.0475*** (0.0134)	-0.0947*** (0.0355)
Top sponsored link			-0.0369 (0.0293)	-0.0148 (0.0185)	0.0485 (0.0402)
Page number=1					0 (.)
Page number=2					0.0208 (0.0147)
Controls for position and interaction terms for link type * position					*
Controls for vendor, day of the week, city and hourly group		*		*	
Constant	5.190*** (0.00610)	6.802*** (0.0501)	5.189*** (0.00748)	6.801*** (0.0501)	5.129*** (0.0204)
Observations	1277	1276	1277	1276	1277
Adjusted R ²	0.937	0.978	0.937	0.978	0.939

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Regressions of the log price on the type of link. All of the regressions include product model fixed effects and robust standard errors clustered by product model.

	(1)	(2)	(3)
	1=lowest price	1=lowest price	1=lowest price
Sponsored Result	0 (.)		
Organic Result	-0.216 (0.196)		
Organic		0 (.)	0 (.)
Sidebar		-0.595* (0.359)	-0.675** (0.335)
Topbar		0.208 (0.245)	0.181 (0.251)
Bottom sponsored link		0.167 (0.435)	-0.0283 (0.413)
Sidetail		0.705*** (0.219)	0.732*** (0.240)
top sponsored link		0.0911 (0.497)	-0.151 (0.489)
Controls for position and interaction between position and link type			*
Page number=1			0 (.)
Page number=2			-0.243* (0.137)
Constant	-1.485*** (0.178)	-1.701*** (0.264)	-1.286*** (0.288)
Observations	1277	1276	1259

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Regressions for the probability of finding the cheapest result in a link of specified type.

	(1)	(2)	(3)	(4)
	% Savings	% Savings	% Savings	% Savings
logit				
Featured Snippet	0 (.)	0 (.)	0 (.)	0 (.)
Organic	0 (.)	0 (.)	0 (.)	0 (.)
Sidebar	1.333*** (0.372)	1.391*** (0.388)	1.218*** (0.380)	1.224*** (0.388)
Topbar	1.305*** (0.335)	1.407*** (0.383)	1.197*** (0.337)	1.253*** (0.383)
bottom sponsored link	0.164 (0.640)	0.327 (0.700)	-0.0430 (0.637)	0.0912 (0.686)
sidetile	-0.0984 (0.418)	-0.0638 (0.448)	0.232 (0.512)	0.385 (0.564)
top sponsored link	-0.707 (1.050)	-0.598 (1.058)	-0.784 (1.091)	-0.606 (1.052)
Page number=1			0 (.)	0 (.)
Page number=2			-0.805** (0.370)	-0.881** (0.384)
Controls for city, day of the week and hourly group		*		*
Controls for position			*	*
Constant	-3.143*** (0.220)	-3.141*** (0.670)	-2.868*** (0.493)	-3.051*** (0.802)
Observations	1213	1194	1196	1177

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Logistic regression part of the two-part model used for the simulations.

	(1)	(2)	(3)	(4)
	% Savings	% Savings	% Savings	% Savings
regress				
Organic	0 (.)	0 (.)	0 (.)	0 (.)
Sidebar	0.00989 (0.0392)	0.0298 (0.0480)	-0.00858 (0.0472)	0.0599 (0.0606)
Topbar	0.0414 (0.0385)	0.0555 (0.0386)	0.0336 (0.0491)	0.0812 (0.0504)
bottom sponsored link	0.0296 (0.0996)	0.0588 (0.0976)	0.0318 (0.0986)	0.0582 (0.0950)
sidetile	0.0165 (0.0539)	0.0236 (0.0689)	0.0368 (0.0720)	0.0274 (0.0952)
top sponsored link	-0.0794*** (0.0290)	-0.0329 (0.0454)	-0.110** (0.0516)	0.00729 (0.0773)
Page number=1			0 (.)	0 (.)
Page number=2			-0.0408 (0.0579)	0.000383 (0.0854)
Controls for city, day of the week and hourly group		*		*
Controls for position			*	*
Constant	0.127*** (0.0290)	0.214 (0.138)	0.155*** (0.0553)	0.301** (0.142)
Observations	1213	1194	1196	1177

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Linear regression part of the two-part model.

Expected savings criterion	Clicks all	0.05%	0.10%	0.50%	1.00%	Clicks only the first
Avg results clicked per product	16.86	15.71	15.57	7.78	3.72	1.00
Avg organic clicked per product	7.72	6.92	6.92	2.50	0.19	0.07
Avg sponsored clicked per product	9.14	8.79	8.65	5.28	3.53	0.93
% of time lowest price found	100.00%	97.22%	97.22%	69.44%	47.22%	23.61%
% of time bought from organic	36.11%	31.94%	31.94%	34.72%	9.72%	6.94%

Table 11: Simulation results.

Inconsistency	Description
Different model	At least one of the characteristics of the product (color, quantity, model number, technical specs, etc.) is different to what the search term was for.
Product listing	Instead of a landing page showing the product searched, it shows a list of products.
Unrelated	The website contains information unrelated to the product being searched.
Same brand	When it is a different model or a product listing but it is of the same brand as the one being searched.
Competing brand	When it is a different model or a product listing and it shows a competing brand.
Main page	When instead of landing on a product page, the link lands on the main page of a website. (For example: http://www.mystore.com/)
Out of stock/unavailable	When the product appears in the page but it is shown that it is either out of stock or no longer available.
Price aggregator	When the website shows prices for the product sold in other websites, instead of selling it directly.
Review/Video	When the landing page shows a review or video of the product and does not offer to sell it.
Not selling	When the website is not selling the product, or it is an unrelated website that does not sell whatever they offer.
Suspicious website	When it is reasonable to suspect that the website is a scam, phishing or otherwise malicious website.
Foreign website	When the website offers the product in a currency different from US Dollars.

Table 12: Most common types of inconsistencies that appeared during the pilot.