Advertising Effects in Equilibrium*

Preliminary and Incomplete·Please do not cite or circulate

Sarah Moshary†

November 27, 2020

Abstract

Platforms face a tradeoff in determining the optimal amount of advertising: while selling advertising space generates revenue, it may also reduce the attractiveness of the platform to consumers, who may in turn exit. As an example, a newspaper must decide how to allocate column inches between advertisements and news stories. Similarly, a search engine must decide how and in what order to present sponsored and organic search results. In order to optimize advertising, these platforms must understand the cost of advertising - that is, how advertising changes consumer behavior. The chief challenge in estimating the effect of advertising is endogeneity. That is, a regression of churn on advertising volumes is confounded by advertisers' targeting policies. This paper exploits an experiment on a large e-commerce platform to estimate the causal effect of advertising on consumer behavior. In the experiment, a random sample of users is shielded from all advertising. The paper documents how advertising affects consumers, sellers, and equilibrium prices on the platform; the results suggest that advertising acts to enhance the salience of sponsored items rather than to signal unobserved quality.

1 Introduction

Platforms that intermediate two-sided markets, such as newspapers, television stations, and search engines, face a tradeoff in selling advertising space: advertising generates revenue, but also potentially affects consumer demand for the platform itself. When inundated by ads, readers may peruse rival newspapers, viewers may change the channel, and customers may query other search engines. Alternatively, advertising may enhance demand if ad content helps consumers identify goods and services that they enjoy.

This second possibility is particularly pertinent for e-commerce platforms that allow sponsored search, where on-site sellers bid for better positions on search result pages (SRPs).[1]

*This project would not have been possible without the help of an anonymous e-commerce platform that provided access to the data. I would also like to thank Sanjog Misra and Steve Tadelis for their feedback.

†University of Chicago Booth School of Business, 5807 S Woodlawn Avenue, Chicago, IL 60637 (email: sarah.moshary@chicagobooth.edu). Financial support for this project was provided by the Robert King Steel Faculty Fellowship.

[1]Lower page ranks correspond to better positions because they are nearer the top of the webpage.

Athey and Ellison (2011) demonstrate that a platform can increase welfare by auctioning SRP positions; this form of advertising reveals information about unobserved firm qualities, as in Nelson (1974b) and Milgrom and Roberts (1986). They highlight that the exact design of the auction matters for the way that sponsored search operates in their model; for example, the common practice of weighting bids by click-through rates coarsens the information revealed by position. A structural empirical literature measures how alternative allocation mechanisms affect platform health, including Yao and Mela (2011), Athey and Nekipelov (2012), Borgers et al. (2013), and Jeziorski and Segal (2015). However, both theory and empirics impose strong assumptions on agent conduct and beliefs: namely, that firms bid optimally given correct expectations about the value of each SRP position and that consumers rationally infer firm quality based on their realized position. Indeed, under alternative assumptions, sponsored search can operate in an entirely different way. As an example, Armstrong, Vickers, and Zhou (2009) develop a model where increasing the prominence of a focal firm coordinates consumer search and alters equilibrium prices. This paper adopts a complementary approach to studying sponsored search that allows me to test these alternative theories; it leverages a field experiment at a large e-commerce platform in the United States to identify the equilibrium effects of sponsored search.

The experiment shields a small fraction (3%) of site traffic on the platform from all sponsored search. To be concrete, on a two sellers platform where both choose to advertise, the researcher would observe firm revenues under dual advertising for 97% of users and under no advertising for 3% of users. This same logic holds for all potential combinations of advertising across the millions of sellers that frequent the platform in reality. Thus, rather than model seller best response functions to predict counterfactual outcomes, the researcher can observe realized seller actions and realized consumer responses. Crucially, the pool of advertising-ineligible users is small relative to site volume, so that the experiment itself is unlikely to alter the sellers' advertising decisions. This final point allows us to interpret the observed advertising levels as equilibrium objects.

This experiment design is atypical in the ad experiment literature, which speaks primarily to the ROI of advertising for a single seller or firm. For example, Ursu (2018) demonstrates that hotel listings with lower page ranks elicit more clicks on Expedia.com, and Sahani and Nair (2019) find that users exposed to a sponsored search campaign on a food delivery app are more likely to order from the promoted restaurant. While instructive about the optimal policy for an individual advertiser, this strand of literature does not speak directly to the question of platform design because it largely abstracts from the general equilibrium effects of advertising. To be clear, general and partial equilibrium effects may differ because changes in own advertising may alter competitors' advertising, entry, and pricing decisions, which in turn may influence consumers behavior and beliefs. For example, ceteris paribus, if a

focal seller successfully bids for the first position on a SRP, then her listings may land an incremental sale. On subsequent searches, however, her competitor might respond with a higher bid, capturing the first position, and "undoing" the focal seller's initial campaign (as in Haan and Moraga-Gonzalez 2011). Alternatively, if the competitor responds slowly or not at all, then incremental sales of the sponsored listing may generate positive reviews, boosting future sales. If ads operate in this fashion, then they may help new sellers overcome the cold start problem, facilitate entry, and toughen competition.

Three features of the experiment examined in this paper are particularly helpful for identifying general equilibrium effects. Of course, the experimental variation also addresses endogeneity concerns arising from the joint determinination of exposures by firm targeting and consumer browsing. Absent exogenous variation, loyal users may see more ads generating a spurious correlation between advertising and sales (Johnson, Lewis, and Nubbemeyer (2017)).

1. **Experiment Duration**. This paper studies five months of data from the experiment, which permits examination of the medium-term effects of platform design. This is longer than typical digital advertising experiments that explore short-term shocks to a single firm's advertising.[2] Consider the simple game outlined in figure 1, where firms A and B compete on an e-commerce platform. Through the lens of the 2×2 game, the standard ad experiment allows the researcher to observe a single column in the payoff matrix (setting aside, for now the distinction between profits and revenues). That is, consider what the researcher observes when seller A experimentally turns on advertising. If seller B did not advertise before the experiment, then this shock facilitates a comparison of payoffs in the first column. Conversely if seller B did advertise before the experiment, then the shock traces out payoffs in the second column. Only if the experiment runs long enough do we have the chance to observe seller B's response to variation in seller A's advertising, permitting observation of payoffs under alternative strategies for seller B. From an experiment design perspective, a key challenge is that is difficult ex ante to know how long is long enough to capture this type of general equilibrium effect. Note also that in a prisoners' dilemma, experiments that randomize only a single firm's advertising will never observe payoffs under (0,0). The experimental variation exploited in this paper shifts all firms' action spaces simultaneously to overcome this challenge.

2. **Number of Players.** One solution to the experiment duration challenge is to randomize advertising by both sellers A and B, as in Dube et al. (2017) or X. Lin et al. (2019). By experimentally varying couponing strategies across two movie theaters, Dube et al.

[2]E.g., Lewis and Reiley 2014 or Johnson, Lewis, and Nubbemeyer (2017); Gordon and Zettelmeyer (2020) report that the median experiment on Facebook runs for 28 days.

		Seller B: Do Nothing	Seller B: Advertise
Seller A	Do Nothing	$\pi_A(0,0), \pi_B(0,0)$	$\pi_A(0,1), \pi_B(1,0)$
	Advertise	$\pi_A(1,0), \pi_B(0,1)$	$\pi_A(1,1), \pi_B(1,1)$

Figure 1: Seller Payoffs from Advertising

show that incorporating the strategic response of a rival dramatically alters the returns to couponing. Unfortunately, this solution is infeasible when the set of sellers is large because the number of treatment cells grows exponentially in the number of sellers.

3. **Level of Randomization**. If the experiment randomizes at a more granular level than seller B's advertising decision, as in Johnson, Lewis, and Nubbemeyer (2017), then during the experiment, seller B is playing against a mixed strategy of seller A. Seller B may respond to this perceived mixed strategy differently than she would to A advertising as a pure strategy. Of course, if the mixed strategy is of inherent interest, then this experiment design is a feature rather than a flaw. In the platform advertising experiment examined here, when a seller chooses to promote a listing, she is not playing against an artificial mixed strategy–if a user is eligible for advertising by firm A, then he is eligible for advertising by all sellers.

The results show that sponsored search increases both the salience and purchasing of sponsored items; sponsored items are 15% more likely to hold the top position on a SRP and they earn 12% higher revenues compared to a purely organic search environment. Sponsored search also leads to shorter search duration (15% faster). These patterns are consistent with a signaling role for advertising that reduces search costs. However, sponsored search also reduces the overall likelihood that a user makes a purchase on the platform (-0.5%) and reduces average overall spending on the platform by 1%. Conditional on an initial purchase, advertising-eligible users are less likely to return to the platform, which does not support a heightened ability to select high quality products. Rather, these patterns suggest that sponsored search works to enhance the prominence of the sponsored items, as in Armstrong, Vickers, and Zhou (2009). In their model, consumers engage in costly search to discover product prices and match quality. If a prominent firm exists, they search it first. Consequently, the prominent firm faces more elastic demand and sets lower prices (relatively to an equilibrium without advertising). In contrast, less-prominent rival firms set higher prices, ultimately leading to a decline in transactions.

I next consider the interplay between sponsored search and price, which is a second point of divergence between the competing views of advertising. A fist finding is that sponsored listings tend to be less expensive than organic listings, again contravening the intuition in Milgrom and Roberts (1986) where price and advertising act as complementary signaling tools. However, recovering the causal effect of sponsored search on prices requires more legwork; the 2020 experiment is poorly-suited to study price effects because sellers set a single listing price for all users. Because 97% of users are eligible for advertising, the prices observed on the site correspond to prices in a sponsored search equilibrium. To identify price effects, I exploit a 2018 shift in the platform's advertising policy; in September, the platform dramatically increased the number of SRP slots available for sponsored listings. A difference-in-differences estimator using categories exempted from sponsored search as controls reveals that advertising leads to a small decline in transacted prices (on the order of 1%, but statistically insignficant) and an increase in listing prices.

Taken together, these results lend support to the view that sponsored search acts primarily to enhance the salience of promoted items. As pointed out by Armstrong, Vickers, and Zhou (2009), advertising may therefore benefit firms that do and do not advertise because the latter benefit from higher prices in the marketplace. Indeed, a striking finding from the experiment is that more sellers benefit from advertising than are harmed by it, even though most sellers do not run campaigns. These findings also speak to optimal platform design by highlighting the tradeoff between advertising revenue and transaction volumes.

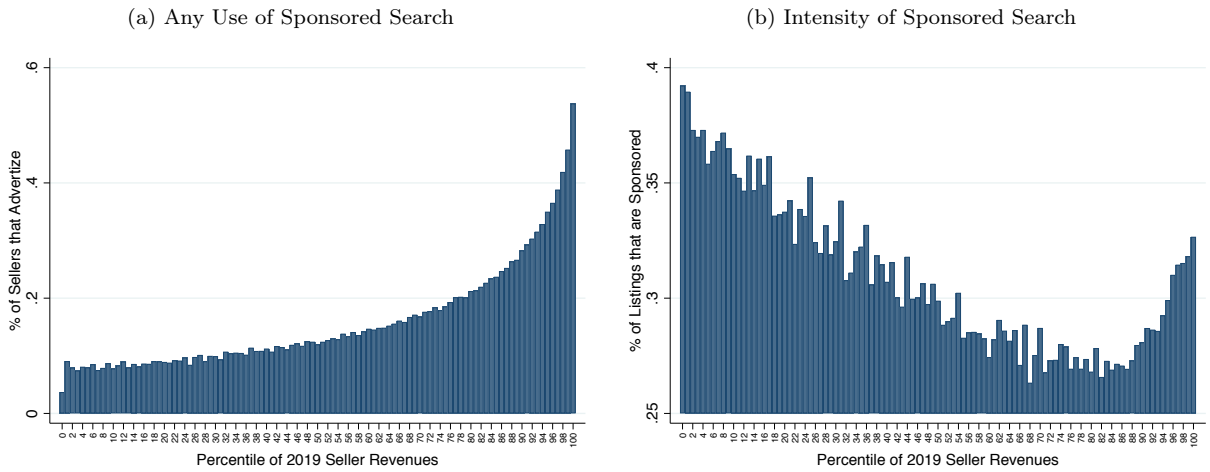
The rest of the paper proceeds as follows: section 2 describes the sponsored search environment on the platform and section 3 the experimental design. Sections 4 and 5 present results on how sponsored search affects buyers and sellers, respectively. Section 6 estimates the effect of sponsored search on prices. Section 7 concludes.

2 Advertising on the Platform

This paper focuses on sponsored listings, which are displayed on platform search results pages (SRPs), interspersed with organic listings. These ads feature items for sale on the platform, and a click on an ad takes the user to a product page where they can complete their purchase (as shown in appendix figure **??**). Sponsored listings are differentiated from the organic listings by a dark outline and a sponsored listings tag. Approximately 4.86% of items sold on the platform are sponsored at some point.[3] Sponsored listings may also be placed on the platform's homepage and checkout success pages, although this is far less common (SRPs account for over 80% of sponsored advertising).

[3]Based on data from February 2019.

Figure 2: Sponsorship by Seller Size



Notes: Advertising=1 if the seller ran at least one campaign between February 1st and June 30th 2020. The zero quantile refers to sellers with zero revenue in 2019 (e.g., new sellers). Among sellers that do advertise, they sponsor 32.3% of their listings on average. Notes for panel B: Sample is sellers that sponsor at least one listing between February 1st and June 30th 2020. The zero quantile refers to sellers with zero revenue in 2019 (e.g., new sellers). Among sellers that do advertise, they sponsor 32.3% of their listings on average.

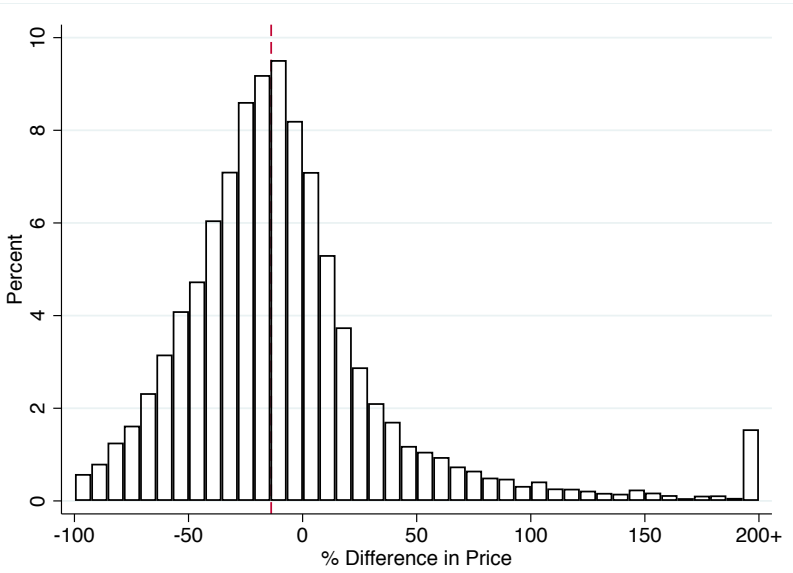
Sellers can promote their listings in most but not all categories. Notable exceptions include automobiles and real estate, which I leverage to identify the price effects of sponsored search in section 6. Positions are allocated through pay-per-sale first price scoring auctions. In the auction, bids are weighted by quality scores that are unknown to participants but are linked to click-through rates. Sellers submit at most one bid for each listing, where they bid a fraction of the item's final sale price. They only pay the platform when a buyer clicks on the sponsored listing and purchases the sponsored item within 30 days of the click. Some SRP positions are reserved for organic results, although these have dwindled over time. In instance when no sponsored listing is deemed relevant for the search query, only organic results are shown.

During the experiment period, approximately 9% of sellers choose to promote at least one listing. Figure 2a breaks down the use of advertising by seller size, which is measured by centile of sales in 2019 (new sellers correspond to the zero centile). The likelihood of sponsoring at least one listing increases monotonically in seller size, which is consistent with large sellers investing more in managing their presence on the platform and/or a larger return to sponsorship for these sellers. Sellers who advertise sponsor 32.3% of their listings on average during the experiment window. Figure 2b show that the intensity of search holds a U relationship with seller size. This pattern is somewhat at odds with Athey and Ellison, 2011, who suggest that firms with high observable quality (e.g., sellers with an established reputation) might be less willing to bid on sponsored search positions.

To give a sense for the types of listings that are sponsored, figure 3 plots the difference

in the average price of sponsored and non-sponsored listings across categories. In the modal category, sellers charge lower prices for the items that they sponsor. This finding is a first piece of evidence against a signaling theory of advertising, where high quality firms advertise more and charge higher prices (e.g, Milgrom and Roberts, 1986 and Armstrong, Vickers, and Zhou, 2009).

Figure 3: Listing Price by Sponsorship



Notes: Each observation is a product category (there are 17,000 plus categories). The percentage difference in price is calculated as $\frac{\bar{p}(sponsored=1)-\bar{p}(sponsored=0)}{\bar{p}(sponsored=0)} \times 100$. Categories with fewer than 100 listings are excluded. Prices are measured exclusive of shipping and insurance charges. Including these shifts the median price difference from -12.46% to -13.78%.

3 Experiment Design

This paper examines a 2020 experiment run on the platform where 3% of US desktop site traffic was randomly chosen to become ineligible for sponsored advertising. The randomization occurred upon a cookie's first contact with the site and was maintained throughout the calendar year. All other cookies were eligible to see sponsored listings, but their exposure depended upon their actions on the site and the demand for advertising, which may be a function of cookie characteristics (including actions and demographics). The paper analyzes data from February 1 - June 30, 2020, providing an opportunity to examine effects of advertising on purchasing and churn up to a three month horizon.[4] Because the number of site visitors is very large, a random subset of advertising-eligible users is included in the control group (15% of US Desktop site traffic). As a randomization check, table 1 provides summary

[4]During the experiment period, new cookies are continuously randomized into the "ads off" group. Randomization began in January 2018, but involved preliminary trouble-shooting. February 1 was selected as the start date based on conversations with the ads team that this constituted the beginning of the full-fledged experiment.

Table 1: Randomization Check

2020 Sponsored Listings Experiment			
	Ads On	Ads Off	Difference
	(1)	(2)	(3)
Hour of Site Visit	11.93	11.93	-0.001
			(0.001)
Days Since Experiment Start	75.59	75.57	0.022
			(0.008)
Share Signed In	0.036	0.037	-0.0004
			(0.000006)

Notes: Standard errors are reported in parentheses in column 3.

statistics for users included in the experiment. Unfortunately, observable user characteristics are few (unless the user signs into their site account, an action that itself might be affected by advertising). Reassuringly, treatment and control users appear similar in terms of the hour of their initial site visit, the date of their visit, and the likelihood that they are signed in to their account on the platform before entering the experiment, which proxies for experience.

I estimate the causal effect of sponsored search on consumer behavior using a straightforward OLS regression of an outcome y_i, such as individual i's spending, on an indicator for treatment, $adson_i$:

$$y_i = \alpha + \beta \cdot adson_i + \epsilon_i. \tag{1}$$

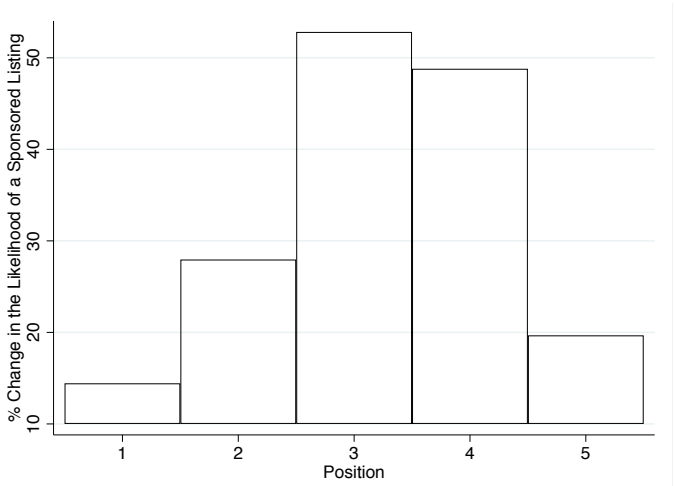
Before presenting results on spending, I investigate how sponsored search affects the ordering of listings on SRPs. One possibility is that sponsored search does not alter the order of listings at all–that the winning bidder for the first position might otherwise hold the highest organic search score.[5] To understand whether and how advertising affects SRPs, I calculate the percentage difference in the likelihood that a sponsored listing surfaces in the first through fifth SRP position when a user is/is not eligible for sponsored search. Figure 4 presents these calculations: sponsored search meaningfully increases the likelihood that an item that *would be sponsored* surfaces in each of these top positions. This variation is key to testing advertising models where sponsorship increases prominence (and therefore the likelihood that a consumer considers the promoted item). The next section investigates how this change in SRPs affects consumer behavior.

4 Effect on Consumer Behavior

If sponsored search informs consumers about promoted products, whether by signaling quality or changing consideration sets, then advertising-eligible users ought to increase spend-

[5]Even in this extreme case, test and control users would not see the exact same SRP. Those eligible for ads would also see the "sponsored" tag marking advertised listings.

Figure 4: Effect of Sponsorship on SRPs



Notes: Data from the February 2020 US Desktop site. Percent changes are calculated based on a SRP-level regression of an indicator for whether the nth ranked listing was sponsored on whether the SRP browser was ads-eligible: $\text{sponsored}_s = \alpha + \beta \cdot \text{ads on}_s + \epsilon_s$. I then report the ratio of $\frac{\hat{\beta}}{\hat{\alpha}}$.

ing on sponsored items. On the other hand, if sponsorship merely annoys or distracts consumers, then purchasing of sponsored items should fall. Table 2 examines the relative spending on sponsored items by ad-eligible vs -ineligible users. The point estimates imply that advertising increases spending on sponsored items by approximately 9.8% and the likelihood of purchase by 7.13%. Sponsored listings do better under sponsored search.

I next document the effect of sponsored search on overall consumer spending. Table 3 reports the percent difference in spending between ads-eligible and ineligible users (the ratio of coefficients $\frac{\beta}{\alpha}$ from equation 1). Column 1 reports estimates using the entire sample, while column 2 reports estimates dropping users in the top 0.01% of spending. The results indicate that sponsored search reduces spending by site visitors on the order of 0.98%, shrinking the economic pie.[6] However, a decline in spending does not necessarily imply that consumers dislike sponsored search; one possibility is that sponsored search steers consumers to less expensive items as in Stigler (1961) or Armstrong, Vickers, and Zhou, 2009, which could quantity but decreasing total spending. A regression of a purchase indicator on ad-eligibility

[6]The average proportion of positions allocated to sponsored listings for ads-eligible users: 12.3% of the top 15 positions are occupied by sponsored listings.

Table 2: % Difference in Purchase of Sponsored Items

	Spending	Spending	Purchase Probability
	(1)	(2)	(3)
Ads On	0.09754***	0.10259***	0.07125***
	(0.00767)	(0.00505)	(0.00160)

Notes: Sample includes all users in the experiment. Heteroskedasticity-robust standard errors in parentheses.

Table 3: Effect of Sponsored Search on Purchasing

	% Change in				
	Spending	Spending	Purchase Once	Purchase Twice	Purchase Twice Conditional on Initial Purchase
	(1)	(2)	(3)	(4)	(5)
Ads On	-0.00447	-0.00986*	-0.00467***	-0.00617***	-0.00151
	(0.0114)	(0.00463)	(0.00101)	(0.00142)	(0.001000)

Notes: Heteroskedasticity-robust standard errors in parentheses. Sample in column 2 drops the top 0.01% of spenders. Probabilities are measured out of 1.

reveals a 0.47% decline in the likelihood of purchase (column 3). Further, if advertising steers consumers to high quality products, then we might expect them to be more likely to repeat purchase as in Milgrom and Roberts, 1986. Admittedly, many of the items for sale on the platform are durable goods that we would not expect consumers to purchase twice within a short window. However, the platform's reputation mechanism could play a similar role to repeat purchasing, allowing an initial sale to generate more future sales for higher quality firms. To the contrary, ad-eligible purchasers are marginally less likely to complete a second purchase (-0.015%, column 5), although the difference is neither statistically significant. Appendix **??** presents additional evidence that effects persist over time for a fixed cohort and across cohorts in the experiment.

One challenge in interpreting these results is that the experiment randomizes sponsored search eligibility for cookies rather than users, as is typical in e-commerce settings (e.g., Blake et al. (2019)). This design means that a single individual might be ineligible for advertising on their laptop, but receive advertising on their mobile phone. Even on a single device, a user might be assigned to a new experience if they clear their cache or switch browsers. Thus, while the experiment is valid for estimating cookie-level effects of advertising, it may overstate or understate the user-level effect of advertising, (see T. Lin and Misra 2019 for a longer discussion of fragmentation). To address this concern, I repeat the main analysis for users who were logged in to the site before randomization, and for whom I can therefore track effects on all spending, both on- and off-cookie. Table 4 reports results. The results in columns 2 and 3 are remarkably similar to the cookie-level results, indicating small negative effects on overall sales and purchasing.

Table 4: Effect of Advertising Eligibility and Exposure on Experienced Users

	% Change in		
	(1)	(2)	(3)
	Spending	Spending	Purchase Probability
Ads On	-0.0322*	-0.0110	-0.00405*
	(0.0132)	(0.00877)	(0.00173)

Notes: Heteroskedasticity-robust standard errors in parentheses. Sample in column 2 drops the top 0.01% of spenders. Sample is cookies that were logged in before January 1, 2020.

Effect on Search Behavior

Consumer browsing data permits a second test of advertising theory. One way that sponsored search is hypothesized to benefit consumers is by expediting the search process (e.g., Athey and Ellison, 2011 or Armstrong, Vickers, and Zhou, 2009). If search costs are high, this benefit could swamp the losses from the reduction in transactions documented above. Table 5 presents results on how advertising eligibility affects consumer search. Search duration is indeed 15% shorter in the sponsored search environment (column 1). However, users also click on 4.3% more listings and are 8.08% more likely to engage in out-of-order search, where a consumer clicks first on a higher ranked item before returning to click on a lower-ranked item second. In other words, consumers are less likely to engage in top-down search in a sponsored search environment. This pattern suggests that sponsored search erodes consumer beliefs that better positions imply higher quality. However, it is possible that consumers engage in "out-of-order" search precisely because they believe that sponsored listings offer higher quality. As described in section 2, the platform commingles sponsored and organic results, where positions are based both on relevance scores and bids (which are zero for organic listings), so that some organic results may appear before the first sponsored listing. To understand the out-of-order search behavior in the ads-eligible group, I calculate the relative likelihood of consumers skipping over sponsored listings to explore higher ranked organic items and vice versa:

$$\frac{Pr\{\text{out-of-order, sponsored first, organic second}\}}{Pr\{\text{out-of-order, organic first, sponsored second}\}} = 1.36.$$

In words, out-of-order search occurs more commonly when consumers are skipping over organic listings to examine sponsored listings. This finding is consistent with consumers inferring unobserved quality from sponsorship.

5 Effect on Seller Revenues

I next investigate how sponsored search affects sellers on the platform. From the perspec-

Table 5: % Difference in Search Behavior
Ad Eligible vs Ineligible Users

	All	Conditional on Purchase
	(1)	(2)
Duration	-15.44%	-0.64%
# of Items Clicked	4.33%	4.52%
Out-of-Order Search	8.08%	8.33%
Page Rank of Purchase	_	5.04%

Notes: Sample comprises all users in the experiment during February 2020.

tive of each seller, the 2020 experiment randomly shielded a 3% share of their market from advertising. The experiment therefore allows the researcher to observe each seller's revenues with and without advertising. Intuitively, new and small sellers have the most to gain from sponsored search; they have yet to establish a reputation on the platform, so their quality may be less observable. Absent sponsored search, these sellers might be relegated to higher search positions. Alternatively, it is possible that large, established sellers leverage sponsored search more effectively. They may employ more sophisticated bidding strategies and/or have higher quality scores, inducing an advertising scale economy. Advertising may even constitutes a prisoner's dilemma for sellers. This section estimates the effect of advertising on revenues earned by different types of sellers in order to distinguish these possibilities.

To understand how sponsored search changes the distribution of sales across all seller types, table 6 reports quantile effects of advertising on revenue across sellers The mean effect is small, negative and statistically insignificant The effect on low quantiles (1st and 5th) is also negative, but the effect on high quantiles (90th, 95th, 99th) is positive and large. The quantile effects reveal that sponsored search increases the dispersion in revenues, but more sellers gain than lose.

Table 6: Distribution of Seller Revenues with/without Advertising

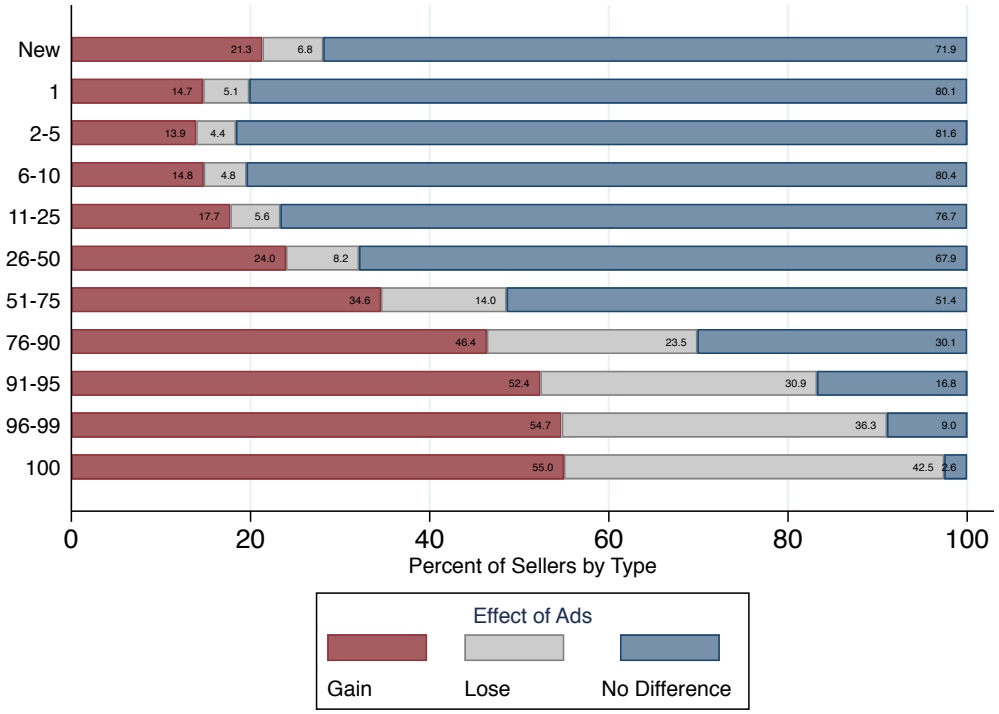
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ads On	-0.221	-255.000	-28.842	0.000	0.252	17.000	46.822	224.200
	(0.400)	(0.838)	(0.096)	(.)	(0.008)	(0.037)	(0.091)	(0.694)
Quantile	Mean	1	5	10	75	90	95	99

Notes: The 25th and 50th quantiles are zero. Seller revenues are measured in dollars.

To get a sense for the types of sellers that benefit from sponsored search, I link sellers with their transactions and listings in 2019, the year before the experiment. I bin sellers by their percentile of sales in 2019, and then plot the share of sellers in each bucket that earn more/less/equal revenue per user in the ads-on compared to ads-off conditions (Figure 5).

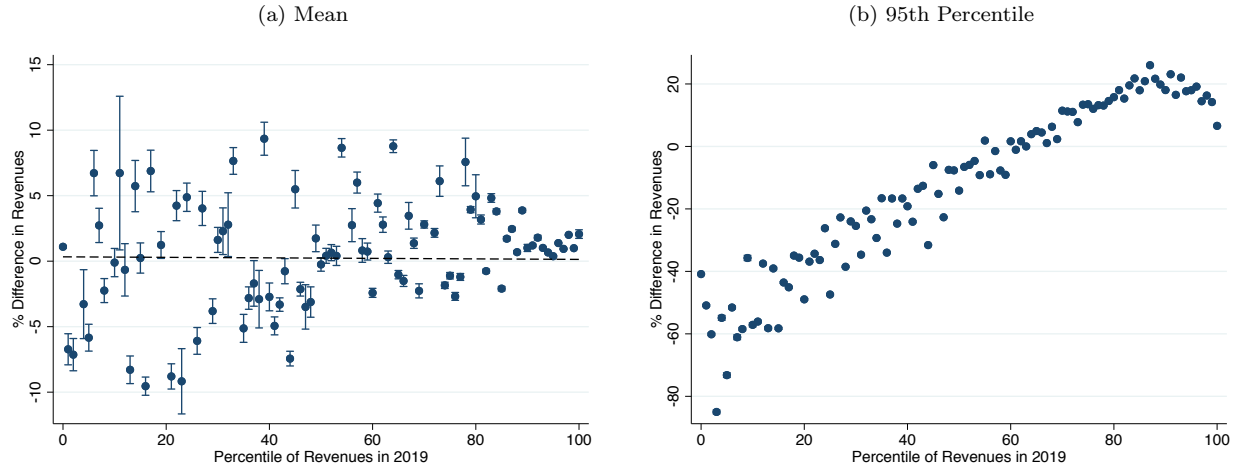
New sellers, defined as those with no sales in 2019, are plotted separately as the 0th quantile. Over three-times as many new sellers earn higher than lower revenues in a sponsored search environment. However, the modal large seller also earns higher revenue under sponsored search.

Figure 5: Share of Sellers with Higher Revenue under Advertising by Quantile of Sales in 2019



Notes: Sample is all sellers with at least one listing 02/01/2020-06/30/2020 on the platform's US site. Sellers are grouped by quantile of sales July-December 2019. "New" denotes sellers had no transactions in the latter half of 2019, e.g. sellers that entered the site in 2020.

Figure 6: Effect of Advertising by Seller Type



Notes: For each percentile of seller sales, this figure plots the average percent difference in sales between ads on and ads off users. Specifically, I run the following regression centile-by-centile: $\text{revenue}_{sg} = \alpha + \beta \cdot \text{ads on}_g + \epsilon_{sg}$. I then plot $\frac{\hat{\beta}}{\hat{\alpha}} \times 100$ along with 95% confidence intervals. The zero quantile refers to sellers with zero revenue in 2019 (e.g., new sellers).

Figure 6 plots the percent difference in sales among eligible/ineligible users across seller types. The plot gives a sense for the magnitudes of gains and losses. Both new sellers and large sellers (above 85th percentile in revenue) earn higher revenues under sponsored search, on the order of 1%, while the effects for small- and medium-sized sellers are mixed. The second panel plots the percent change in the 95th quantile of sales for each group; the quantile effects increase strongly in seller size.

Two factors may drive these heterogeneous effects: differential use of sponsored search, as shown in figure 2a, and differential effects *per* listing. For example, larger sellers may be better at designing campaigns (e.g., selecting which items to sponsor). Alternatively, potential reputational spillovers of sponsored search might matter more for small sellers. To disentangle these possibilities, I examine the relationship between per-listing revenues and advertising eligibility for sponsored items. The baseline specification is

$$\text{revenue per sponsored listing}_{sg} = \alpha_0 + \alpha_1 \cdot \text{ads on}_g + \Gamma_s + \epsilon_{slg} \tag{2}$$

where s denotes seller, g denotes user group (advertising eligible/ineligible), and the sample comprises sellers that sponsor listings. Estimates of the regression parameters pooled across all sellers are presented in table 7. On average, sellers earn 12.4% higher revenues on their sponsored listings in a sponsored search environment ($\frac{\alpha_1}{\alpha_0} \times 100$, column 1). I then extend the specification to allow for differential effects across listings types l, which measures how sponsored search affects standard listings (listings that are not sponsored), :

Table 7: Effect of Advertising on Sales of Sponsored & Non-Sponsored Items

	% Change in Revenue per Item		
	(1)	(2)	(3)
Sample	Promoted	All	All
Ads On	12.42	0.805	-2.53
	(1.615)	(3.531)	(1.75)
Ads On × Sponsored		21.61	25.31
		(4.40)	(3.22)
Sponsored		40.159	59.38
		(18.025)	(5.32)
Ads On × Sister Sponsored			40.60
			(39.07)
Seller FE	Yes	Yes	Yes

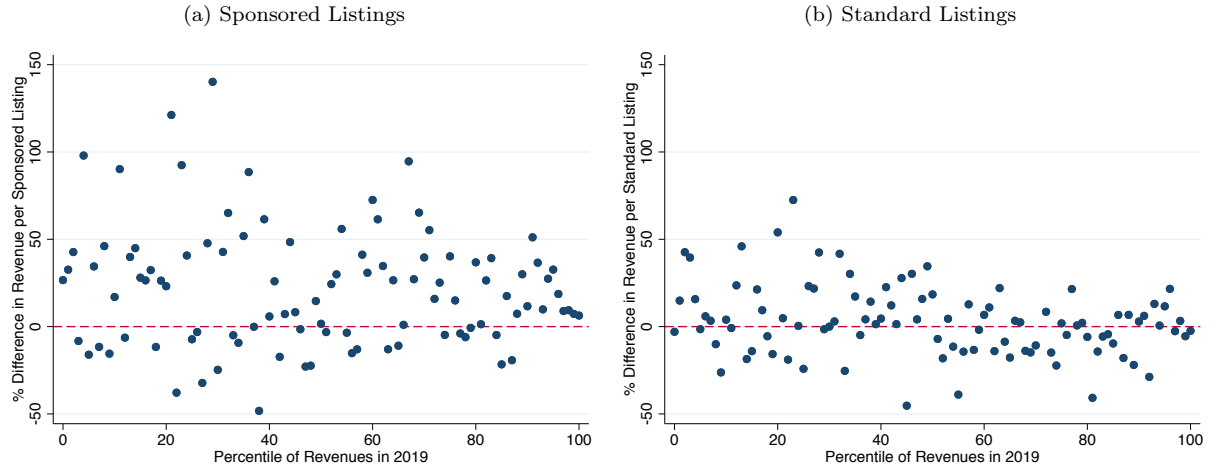
Notes: Standard errors are clustered by seller. Sample in columns (2) & (3) includes all sellers with a listing in the US site during the experiment; in column 1, it is restricted to sellers with at least one sponsored listing. To be clear, each observation seller × listing-type combination (where listing type is promoted or not promoted). If a seller does not have any listings of a particular type, then that combination is excluded. All regressions include a constant. This table reports the ratio of the coefficient over the constant so that magnitudes are interpretable as percent changes.

$$\text{revenue per listing}_{sgl} = \alpha_0 + \alpha_1 \cdot \text{ads on}_g + \alpha_2 \cdot \text{ads on}_g \times \text{sponsored}_l + \alpha_3 \times \text{sponsored}_l + \Gamma_s + \epsilon_{slg}. \tag{3}$$

The estimates in table 7 column 2 reveal that sponsored items earn 40% higher revenue even among ads-ineligible users ($\frac{\alpha_3}{\alpha_0} \times 100$). This difference reflects both seller decisions about which listings to sponsor, but also potential spillover effects. As an example, sponsorship might generate positive reviews that increase sales to all users.

Figure 7 plots these effects separately across seller types. Both new sellers and large sellers see an increase in per-listing revenue for sponsored items, and a small negative effect for standard listings. New sellers enjoy a particularly large increase in per-revenue sales from sponsorship, which is consistent with a role for quality signaling.

Figure 7: Effect of Sponsored Search on Listing Revenues



Notes: Standard error bars are excluded for clarity but included in appendix table **??**.

Finally, I examine whether the benefits of sponsorship extend beyond the focal item to other listings offered by the same seller. If sponsorship signals seller quality rather than listing quality, then sponsoring a subset of listings could boost the sales of other, standard listings. I extend specification 2 to include standard listings and add a covariate to indicate whether a standard listing is posted by a seller who sponsors *other* listings, sister sponsored$_s$:

$$\begin{aligned}\text{revenue per listing}_{sgl} = \alpha_0 &+ \alpha_1 \cdot \text{ads on}_g + \alpha_2 \cdot \text{ads on}_g \times \text{sponsored}_l \\ &+ \alpha_3 \cdot \text{ads on}_g \times \text{sponsored}_s + \alpha_4 \cdot \text{sponsored}_l \\ &+ \alpha_5 \cdot \text{sister sponsored}_s + \Gamma_s + \epsilon_{slg}.\end{aligned} \tag{4}$$

The coefficient α_3 captures the spillover effect of advertising on standard listings sold by sellers who advertise. Again, to protect proprietary data and for ease of interpretation, I present the ratio of each coefficient to the constant, so that the magnitudes report in table 7 can be interpreted as percent changes. The point estimate for the halo effect is large but imprecisely estimated.

Taken together, the seller-side analysis reveals three facts: first, sponsored search increases dispersion in seller revenues; second, more sellers benefit from sponsored search than are harmed from it; and third, for high quantiles, effects are increasing in seller size. This final pattern is hard to reconcile with a signaling model, as ostensibly, the platform can accurately measure seller quality after thousands of transactions.

6 Price Effects of Advertising

This section investigates the interplay between advertising and pricing. As described by Bagwell (2007), an extensive literature in economics has generated a wide array of hypotheses: advertising may increase the elasticity of demand–either by attracting relatively inelastic consumers or by increasing (perceived) product differentiation–putting upward pressure on prices; advertising could decrease elasticities, particularly for the focal firm, by coordinating search activity (Armstrong, Vickers, and Zhou (2009)); or advertising might increase price competition if it lowers barriers to entry, e.g., by mitigating the cold start problem for new sellers. A first finding is that ads eligible users purchase items that are 1.47% more expensive than ineligible users, although the difference is not statistically significant (t-stat of 1.30). This difference is driven primarily by differences in the transacted price of sponsored items; among ads-eligible users, the average purchase price of sponsored items is higher than for standard items (1.65% vs 1.21%, with t-statistics of 5.48 and 1.37, respectively).

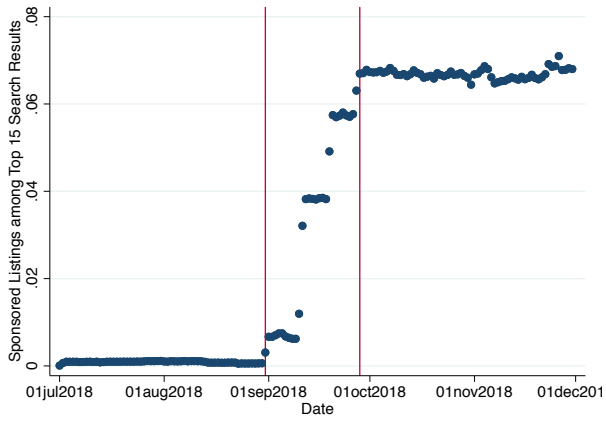
These purchase price patterns fly in the face of the prediction in Armstrong, Vickers, and Zhou (2009) that the prominent firm charges lower price in equilibrium. However, these price comparisons do not incorporate potential general equilibrium effects because sellers are unable to set separate prices for ads-on and -off users in the experiment. Because 97% of site traffic sees sponsored listings, I interpret observed prices as equilibrium prices in an "ads on" environment.[7] To identify the price effects of sponsored search this paper therefore exploits a natural experiment: the ramp up of sponsored search in September 2018. Before September, the platform reserved a certain number of SRP positions for organic listings. Many of these positions became available for sponsored listings over the course of September. Figure 8 shows how this policy change increased the volume of sponsored search during this period[8] Ceteris paribus, if sponsored search intensifies (softens) price competition, then we would expect prices to fall (rise) following this marked increase in advertising volumes.

[7] It is possible that a user in the ads-ineligible group who is aware of the A/B test could interpret prices as signals of quality. To rule out this type of SUTVA violation would require testing on different platforms.

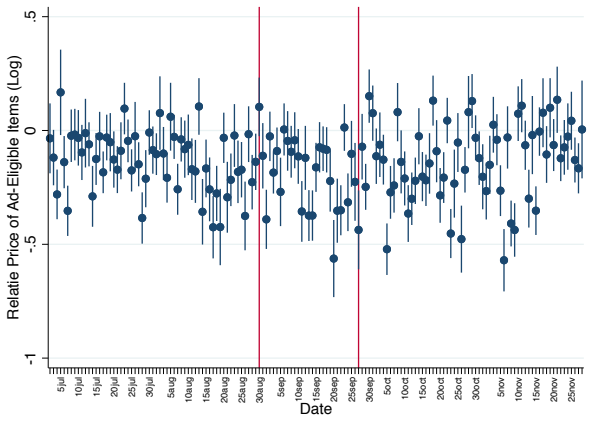
[8] Conversations with the ads team at the platform suggest that this date was not chosen to coordinate with any other major policy change.

Figure 8: Sponsored Search in 2018

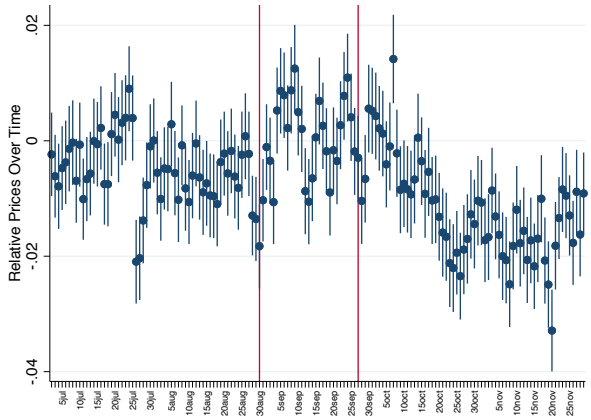
(a) Sponsored Search Volumes



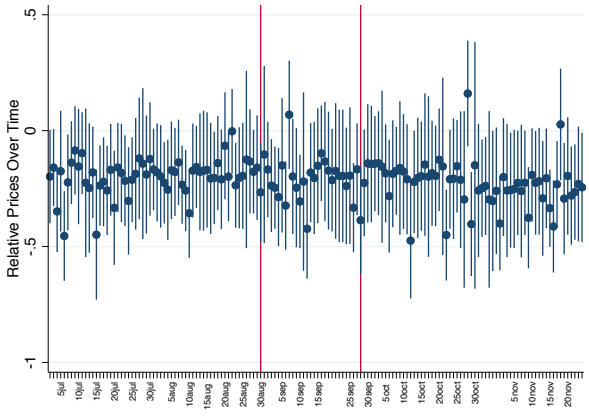
(b) Transcated Prices - Eligible Categories



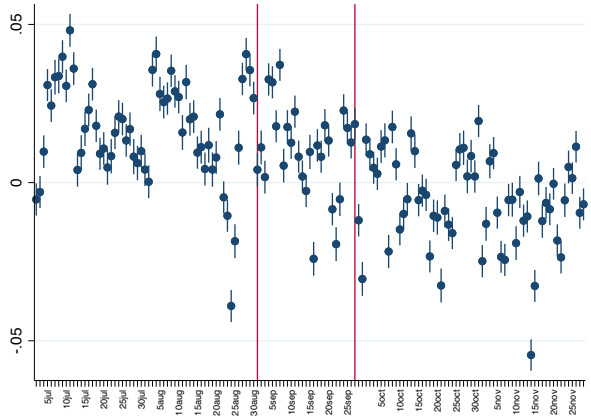
(c) Transacted Prices
Fixed Assortment - Eligible Categories



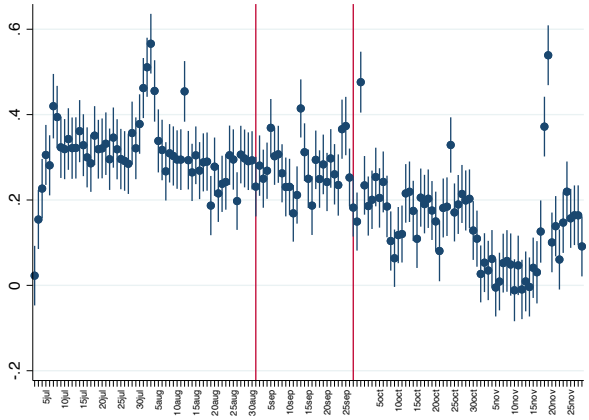
(d) Transacted Prices
Fixed Assortment - Ineligible Categories



(e) Listing Prices
Eligible Categories



(f) Listing Prices
Ineligible Categories



Notes: The dates delineated by the red lines span August 31st through September 28th. Only eligible listing formats are included.

Of course, prices might change over time for reasons unrelated to sponsored search policy. To mitigate this concern, I exploit listings and sales in three categories where the platform does not offer sponsored search: travel, real estate, and miscellaneous (e.g., career development self-help guides). While these control categories differ from the treatment categories in price levels, identification requires only that these categories capture broad trends on the platform. The estimating equation is run at the category (c) by day (t) level, where I include indicators for the ramp-up period, $Interim_t$, and the post period, $Post_t$ (dates after September 28, 2018), and their interactions with $Eligible_c$, an indicator that takes a value of one for all categories that permit sponsored search. I also include category fixed effects Ω_c.

$$log\bar{p}_{ct} = \alpha + \beta \cdot Post_t + \delta \cdot Post_t \times Eligible_c + \gamma \cdot Interim_t + \lambda \cdot Interim_t \times Eligible_c + \Omega_c + \epsilon_{ct} \quad (5)$$

The first two columns in table 8 reports results for transacted prices, and indicate no change following the increase in sponsored search. This pattern is also borne out in figure 8 (b). In columns 3 and 4, I estimate 5 using raw transaction prices, which allows me to include product fixed effects. This exercise restricts to products that can be categorized (e.g., a textbook that can be linked to an ISBN or a mobile phone that can be linked to an SKU), but it reduces any variance induced by inventory changes on the platform. The point estimate of δ is negative and economically meaningful, but statistically insignificant. In contrast, in columns 5 & 6, I focus on the prices for new or revised listings, and the reverse trend emerges; listing prices tend to increase following the change in SRPS. This dovetails with the comparative static in Armstrong, Vickers, and Zhou (2009), where most firms increase prices in response to the advent of advertising.

Table 8: Price Effects of Advertising

	Transactions				Listings	
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.0298**		-0.0132***		-0.0205***	
	(0.0108)		(0.0025)		(0.0061)	
Interim	-0.0396*		0.0100***		0.0114*	
	(0.0175)		(0.0011)		(0.0054)	
Eligible Category × Post		-0.0022		-0.0236		0.1121*
		(0.0144)		(0.0302)		(0.0463)
Eligible Category × Interim		-0.0497*		0.0006		-0.0678*
		(0.0237)		(0.0265)		(0.0286)
FE Type	Category	Category, Date	Product	Product, Date	Category	Category, Date

Notes: Data from February - November 2018. Thirty-three categories in columns 1 and 2. Regressions are weighted by cell size.

7 Conclusion

Sponsored search is an important component of the digital economy, surpassing $50 billion in 2019.[9] In a seminal paper, Ostrovsky and Schwarz (2016) demonstrate the potential gains from leveraging economic theory to design sponsored search, focusing on optimal reserve prices in auctions that many platforms use to allocate search positions. However, economic theory offers conflicting views on how and why sponsored search affects the interactions between consumers and firms on e-commerce platforms. One vein of literature dating back to Nelson (1974b) makes the case that advertising can be informative, allowing consumers to identify high quality products, thereby improving welfare. A second vein argues that advertising mainly enhances salience, allowing firms to alter consumer consideration sets. To understand how sponsored search works in practice, this paper leverages a large-scale field experiment on an e-commerce platform operating in the US in 2020. The experiment shielded 3% of site from all sponsored search, allowing the researcher to study the equilibrium effect of advertising. The findings bear out two predictions common to both perspectives on advertising: consumers are more likely to search sponsored items and they are also more likely to purchase sponsored items. However, the results show that sponsored search reduces consumers spending by 0.9%, which contravenes a central prediction of models of informative advertising; it suggests that sponsored search works primarily by affecting firm prominence.

[9]Nicole Perrin. Nov 18, 2019. eMarketer.com. "Google Maintains Wide Lead in Net US Search Ad Revenues."

Further evidence to support this view is the revealed distribution of benefits across sellers; the largest sellers on the platform appear to benefit most from the sponsored search environment. These are precisely the sellers whose latent types ought to be revealed by thousands of customer reviews.

These results speak to the equilibrium effects of advertising in ways that typical advertising experiment cannot. The existing literature chiefly analyzes the partial equilibrium effect of advertising on own- or rival-sales that holds the actions of other advertisers as fixed. In reality, however, when one firm increases its advertising, other firms react. This experiment captures these general equilibrium effects by eliminating all advertising for a small portion of site traffic. That said, there is still scope for understanding other GE effects related to seller price and entry decisions. The challenge in this experiment is that sellers can set only a single price to users who are/are not eligible for advertising. Because the experiment shields only 3% of site traffic from advertisements, I interpret observed prices as those consistent with an equilibrium where advertising is permitted. To understand how sponsored search affects equilibrium prices, I exploit the ramp-up of sponsored search on the platform in 2018. The results, while imprecise, suggest that transacted prices fall, but listing prices rise – a finding that dovetails with Armstrong, Vickers, and Zhou (2009). They describe an equilibrium where advertising creates a prominent firm that consumers search first; this firm reduces its price, while less-prominent rivals raise their prices. Just as sponsored search can affect prices, it may also affect entry decisions. I leave investigation of these effects to future work.

References

Armstrong, Mark, John Vickers, and Jidong Zhou (2009). "Prominence and Consumer Search". In: *RAND Journal of Economics* 40.2, pp. 209–233.

Athey, Susan and Glenn Ellison (2011). "Position Auctions with Consumer Search". In: *Quarterly Journal of Economics* 126.3, pp. 1213–1270.

Athey, Susan and Denis Nekipelov (2012). "A Structural Model of Sponsored Search Advertising Auctions". In: *Working Paper*.

Bagwell, Kyle (2007). "Handbook of Industrial Organization". In: ed. by Mark Armstrong and Robert Porter. Vol. 3. Elsevier. Chap. The Economic Analysis of Advertising, pp. 1701–1844.

Blake, Tom et al. (2019). "Price Salience ad Product Choice". In: *Working Paper*.

Borgers, Tilman et al. (2013). "Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence". In: *American Economic Journal: Microeconomics* 5.4, pp. 163–187.

Dube, JP et al. (2017). "Competitive Price Targeting with Smartphone Coupons". In: *Marketing Science* 36.6.

Gordon, Brett and Florian Zettelmeyer (2020). "Close Enough? A Large-scale Exploration of Non-experimental Approaches to Advertising Measurement". In: *WP*.

Haan, Marco A. and Jose Luis Moraga-Gonzalez (2011). "Advertising for Attention in a Consumer Search Model". In: *The Economic Journal* 121.552, pp. 552–579.

Jeziorski, Przemyslaw and Ilya Segal (2015). "What Makes Them Click: Empirical Analysis of Consumer Demand for Search Advertising". In: *American Economic Journal: Microeconomics* 7.3, pp. 24–53.

Johnson, Garrett A., Randall A. Lewis, and Elmar I. Nubbemeyer (2017). "Ghost Ads: A Revolution in Measuring Ad Effectiveness". In: *Journal of Marketing Research* 54.6, pp. 867–884.

Lewis, Randall A. and David H. Reiley (2014). "Online Ads and Offline Sales: Measuring the Effects of Retail Advertising via a Controlled Experiment on Yahoo!" In: *Quantitative Marketing and Economics* 12, pp. 235–266.

Lin, Tesary and Sanjog Misra (2019). "The Identity Fragmentation Bias". In: *Working Paper*.

Lin, Xiliang et al. (2019). "Parallel Experimentation in a Competitive Advertising Marketplace". In: *Working Paper*.

Milgrom, Paul and John Roberts (1986). "Price and Advertising Signals of Product Quality". In: *Journal of Political Economy* 94.4, pp. 796–821.

Nelson (1974b). "Advertising as Information". In: *The Journal of Political Economy* 82, pp. 729–54.

Ostrovsky, Michael and Michael Schwarz (2016). "Reserve Prices in Internet Advertising Auctions". In: *Working Paper*.

Sahani, Navdeep and Harikesh Nair (2019). "Does Advertising Serve as a Signal? Evidence from a Field Experiment in Mobile Search". In: *The Review of Economic Studies*.

Stigler, R. George J. (1961). "he Economics of Information". In: *The Journal of Political Economy* 69.3, pp. 213–225.

Ursu, Raluca (2018). "The Power of Rankings: Quantifying the Effects of Rankings in Online Consumer Search and Purchase Decisions". In: *Marketing Scieence* 37.4, pp. 530–552.

Yao, Song and Carl Mela (2011). "A Dynamic Model fo Sponsored Search Advertising". In: *Marketing Science* 30.3, pp. 447–468.