Sponsored Search in Equilibrium: Evidence from Two Experiments*

Sarah Moshary†

August 17, 2021

Abstract

Advertising affects not only the firm that advertises, but also the platform that hosts the advertisement. Using data from a field experiment at an e-commerce platform, I demonstrate that the effects on the advertising firm and the hosting platform can differ sharply. The experiment blocks all sponsored search advertising for a small fraction of site visitors. Compared to users who are shielded from ads, users who see ads spend significantly more on sponsored listings and significantly less on organic listings. The second effect dominates, revealing that on net, sponsored search reduces total sales on the platform. Using a separate natural experiment, I also find evidence that sponsored search puts upward pressure on prices, which can exacerbate cannibalization. Together, these findings illustrate the cost of advertising from the perspective of the platform.

1 Introduction

Advertising is an important source of revenue for platforms that intermediate two-sided markets, such as television stations, search engines, and e-commerce platforms, but it can also affect consumer demand for the platforms themselves. When inundated by ads, viewers may change the channel, customers may query other search engines, and shoppers may navigate to other e-commerce websites. Alternatively, advertising may enhance demand if ad content helps consumers to identify goods and services that they enjoy. Using data from a large-scale field experiment at an e-commerce platform, this paper shows that paid search advertising shrinks the economic pie. Users who see paid search ads spend more on advertised listings, but less on organic listings.[1] This cannibalization is first order: the net effect is a reduction in the likelihood of making a purchase and in total spending on the platform.

The field experiment blocks all sponsored search advertising for a small fraction (3%) of site traffic on an anonymous e-commerce platform. The experiment allows me to observe

*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 Tom Blake, Jim Dana, Gaston Illanes, Sanjog Misra, Bradley Shapiro, Steve Tadelis, Hal Varian, Caio Waisman, and participants at the NBER Digitization workshop, Columbia IO Seminar, and UCLA Anderson Marketing Seminar for their help and feedback.

†University of Chicago Booth School of Business, 5807 S Woodlawn Avenue, Chicago, IL 60637. Email: sarah.moshary@chicagobooth.edu.

[1]As explained in section 3, the experiment randomizes eligibility for ads. Some users who are eligible for sponsored search ads may not in fact see them, for example, if they do not use the search feature of the platform.

realized seller actions and consumer responses in an equilibrium with sponsored search (the status quo on the platform) and also under an advertising ban. It is unusual to observe outcomes in a setting where advertising is prohibited. For example, if sponsored search is a prisoner's dilemma, then an experiment that randomizes only a single firm's advertising cannot reveal outcomes in a no-advertising environment. The experimental variation exploited in this paper shifts all sellers' action spaces simultaneously to overcome this challenge, shedding light on the payoffs that are relevant to platform design.

Three additional features of the experiment are helpful for identifying the impact of paid search on platform revenues. First, the pool of buyers shielded from sponsored search is small relative to site traffic, so that the experiment itself is unlikely to alter sellers' advertising decisions. Second, the experimental variation addresses endogeneity concerns arising from the joint determination of ad 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). Third, the experiment spans five months, allowing examination of long-term effects beyond what is typical in digital advertising experiments.[2]

The results of the field experiment confirm that sponsored search advertising increases both the prominence and purchasing of sponsored items (items with an active ad campaign during the experiment). However, users who see ads are less likely to complete a purchase on the platform (-0.5%) and they spend less overall on the platform (-1%). These findings highlight an important cost of paid search advertising from the platform's perspective: by decreasing the page rank of sponsored listings, sponsored search displaces organic content, cannibalizing sales of organic listings. An analysis focused only on the sales of sponsored listings would suggest that paid search increases sales by 10%; incorporating losses incurred by organic listings reveals that sponsored search in fact reduces sales. Thus, paid search reduces the commissions that the platform earns on transactions. Paid search may nevertheless be profitable for the platform if the advertising revenue it generates is sufficiently great. I illustrate this through a simple back-of-the-envelope calculation.

The implications for social welfare are more nuanced because they depend on how sponsored search affects equilibrium match quality and search costs. Informative models of advertising, e.g., Nelson (1974b) and Milgrom and Roberts (1986), suggest that advertising may guide consumers to the products that they enjoy most. It is hard, but not impossible, to reconcile these theories with the finding that advertising reduces site-wide transaction volume. One possibility is that advertising eliminates unsatisfactory transactions, where consumers buy products that do not fulfill their needs. However, conditional on an initial purchase, users who see ads are marginally less likely to return to the platform, which does

[2]For example, Gordon, Zettelmeyer, Bhargava, and Chapsky (2019) report that the median experiment on Facebook runs for 28 days.

not support a heightened ability to select high quality products. Further, sponsored search does not change consumer satisfaction as measured by returns and complaints. These findings point against an informative role for advertising, and thus contrast with Sahni and Nair (2019), which shows that disclosing that a listing is sponsored on a food delivery app increases purchasing. Turning to search costs, the evidence is mixed: sponsored search reduces search duration (-1%), but marginally increases clicks. Sponsored search also reduces the use of optional filters: buyers who do not see sponsored ads are more likely to sort their search results in ascending price order (+1%). This finding suggests that general equilibrium effects might include an effect of sponsored search on pricing.

The interplay between advertising and pricing may operate through multiple channels: sellers might pass on the cost of advertising to consumers through higher prices, or alternatively, sellers might adjust prices if advertising alters the elasticity of demand. In the field experiment, users who see sponsored listings buy items with higher prices (+1.6%), which dovetails with their reduced use of the ascending price filter. These patterns are consistent with sponsored search reducing price sensitivity, but they could also be driven by selection (because sponsored search changes the composition of purchasers and the types of items surfaced at the top of SRPs). The field experiment is poorly-suited to identify price effects because sellers set a single listing price for all users. I interpret the prices observed on the platform as prices in an equilibrium with sponsored search because 97% of users are eligible for advertising during the experiment. Thus, to identify price effects, I exploit a natural experiment created by a shift in the platform's advertising policy: in September 2018, 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 suggests that advertising leads to a small increase in transacted prices (on the order of 2%, but statistically insignificant) and a more meaningful increase in the prices for new listings, on the order of 10%. So long as demand is elastic, this increase in prices would tend to reduce platform revenue, amplifying cannibalization.

Relative to the existing ad experiment literature, the design of the field experiment studied here is atypical, as is the focus on spillover and price effects. Most studies of digital advertising speak primarily to the ROI of advertising for a single seller or firm. A few examples include: Lewis and Reiley (2014), which studies the returns to banner ads on Yahoo!; Blake, Nosko, and Tadelis (2015), which estimates the return to paid search advertising for eBay; or Ursu (2018), which demonstrates that hotel listings with lower page ranks elicit more clicks on Expedia.com.[3] Similar to these analyses, I find that advertised items earn higher sales in a sponsored search environment. While instructive about the optimal policy for an individual advertiser, however, this literature does not speak directly to the question

[3]Other work in this vein includes Narayanan and Kalyanam (2015), Dai and Luca (2016), and Gordon et al. (2019).

of platform design because it largely abstracts from the general equilibrium effects of advertising. To be clear, general and partial equilibrium effects may differ for several reasons: changes in own advertising may affect competitors' sales through business-stealing or by increasing category awareness (Shapiro, 2018); own advertising may also affect competitors' advertising, entry, and pricing decisions; these in turn may influence consumer behavior and beliefs. An A/B test that randomizes advertising for a single seller on a platform typically precludes competitor responses. First, if one seller randomizes his ads, then other sellers face a mixed strategy of this seller during the experiment. Competitors may respond to this perceived mixed strategy differently than they would if the experimenting seller instead advertised to all consumers. One solution to this challenge is to randomize advertising by all sellers, as in Dubé, Fang, Fong, and Luo (2017) and X. Lin, Nair, Sahni, and Waisman (2019). However, this solution is infeasible when the set of sellers is large because the number of treatment cells grows exponentially in the number of sellers. 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 to see one seller's ads, then he is eligible to receive other ads, too.

The findings from this field experiment contrast with result in Sahni and Zhang (2020) that enhancing the salience of sponsored listings increases consumer engagement on an anonymous search engine. They interpret increased searches and clicks, particularly in geographies with a high share of new businesses, as evidence of informative advertising that may help resolve the cold-start problem. I find no evidence that small sellers, as measured by earnings before the experiment, benefit disproportionately from sponsored search, and indeed small sellers are less likely to sponsor listings in the setting studied here. This is not to say that advertising results in a prisoner's dilemma; to the contrary, the preponderance of sellers earn higher revenue in the ads-on compared to ads-off environment in the field experiment. A second related paper is Abhishek, Jerath, and Sharma (2019) which finds that increasing the number of positions dedicated to sponsored search does not affect total conversion on Flipkart, a mobile app popular in India. Both Abhishek, Jerath, and Sharma (2019) and Sahni and Zhang (2020) consider changing sponsored search on the margin, so that both treatment and control users experience some degree of paid search advertising. In contrast, this paper considers removing sponsored search entirely. A particular advantage of this field experiment is that ads-off users (who see only organic listings) comprise only a small fraction of website traffic, so that the researcher can interpret seller advertising choices as optimized for an environment with paid search. As an example, the Flipkart experiment reduced the number of paid search slots for a subset of buyers while maintaining the status quo (higher) number of slots for the majority of buyers. Because sellers could not condition their advertising decisions on a consumer's treatment assignment, the experiment cannot capture how

sellers might change their advertising decisions if Flipkart reduced paid search slots for all buyers.

My findings on cannibalization also speak to a theory literature studying the relationship between paid search advertising and transactions. In particular, Athey and Ellison (2011) explore how the design of the position auctions used to allocate space on search result pages affects transaction volumes. In their model, common practices such as weighting bids by click-through rates can reduce equilibrium match quality by coarsening the information revealed by position. Armstrong, Vickers, and Zhou (2009) develop an alternative model where increasing the prominence of a focal firm coordinates consumer search and alters equilibrium prices. This paper can be seen as quantifying the interaction that these papers describe between the advertising and transaction sides of the platform. It also relates to a structural empirical literature, including Yao and Mela (2011), Athey and Nekipelov (2012), Borgers, Cox, Pesendorfer, and Petricek (2013), Jeziorski and Segal (2015), and Choi and Mela (2019) which measure how alternative position allocation mechanisms affect platform outcomes. For tractability, this theory and empirical literature require certain assumptions on agent conduct and beliefs: for example, 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. This paper adopts a complementary approach to studying sponsored search that uses experimental variation and an alternative set of identifying assumptions that I describe in the paper. The implications of the findings extend beyond e-commerce to other settings where there may be a tension between advertising and transactions. As an example, this tension manifests in Wilbur (2016), which argues that television stations ought to favor ads with consumer-friendly content. More broadly, the finding that consumers may dislike ads echoes findings in other contexts, such as Huang, Reiley, and Riabov (2018) that shows consumers listen to Pandora internet radio less often when their programming is interrupted by voice ads.

The rest of the paper proceeds as follows: section 2 details the sponsored search environment on the platform and section 3 the field experiment design. Sections 4 describes the main result that sponsored search increases the sales of sponsored items but reduces overall platform sales. Section 5 presents evidence on how sponsored search affects match quality and search behavior. Section 6 describes the natural experiment and the effect of sponsored search on prices. Section 7 discusses entry, and section 8 concludes.

2 Advertising on the Platform

This paper studies sponsored listings, which are displayed on platform search results pages, interspersed with organic listings. These ads feature items for sale on the platform,

and a click on an ad takes the user to a listing page where they can complete their purchase, as shown in figure 1. Sponsored listings are differentiated from organic listings by a dark outline and a sponsored listing tag. Each listing can appear at most once on a SRP, either as an organic or sponsored listing. Approximately 12.3% of the top 15 positions are occupied by sponsored listings. Sponsored listings may also be placed on the platform's homepage and checkout success pages, although this is far less common.[4]

Figure 1: SRP Example

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 each 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. Organic listings also compete in the auction, but their "bid" is entirely based on their quality ranking; in essence, sponsored listings allow sellers to boost the quality rating of their product for a fee. Some SRP positions are reserved for organic results, although these have dwindled over time. In cases when no sponsored listing is deemed relevant for the search query, only organic results are shown.

During the experiment, approximately 12% 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 the second half of 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

[4]Platform employees report that SRPs account for 80%+ of sponsored ads.

Figure 2: Sponsorship by Seller Size

(a) Any Use of Sponsored Search

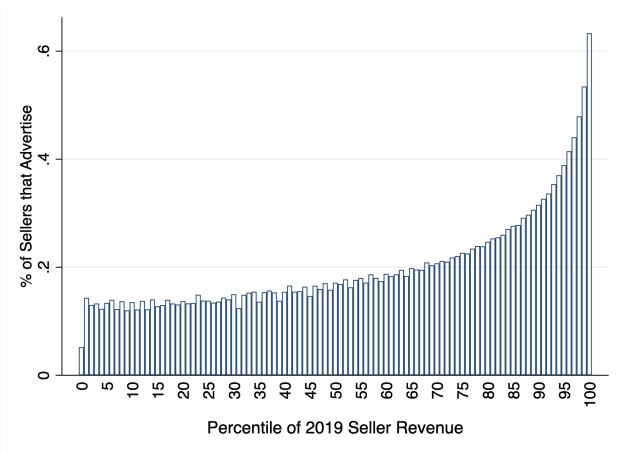

(b) Intensity of Sponsorship

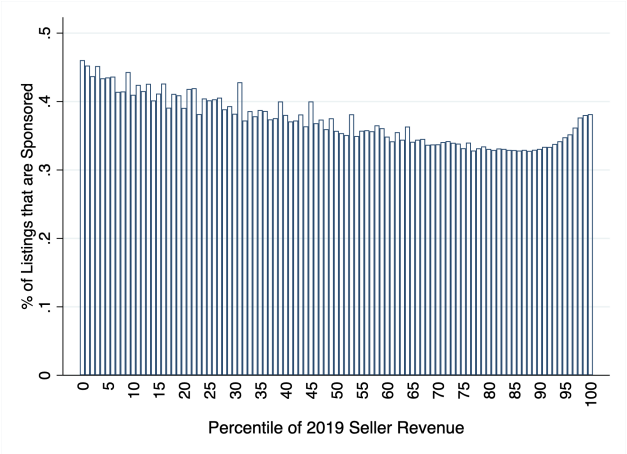

Notes: Advertising=1 if the seller ran at lease one campaign between February 1st and June 30th, 2020. The zero quantile refers to sellers with zero revenue in 2019 (e.g., new sellers). Panel B restricts to sellers that sponsor at least one listing between February 1st and June 30th, 2020.

return to sponsorship for these sellers. Sellers who advertise sponsor 38.6% of their listings on average during the experiment window. Figure 2b shows that the intensity of sponsored search holds a U-shaped relationship with seller size, which is somewhat at odds with the view that firms whose quality is observed (e.g., sellers with an established reputation) might advertise less often.

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 (left panel) and across products (right panel). Both within category and product, sellers tend to set lower prices for advertised items. This finding also cuts against a signaling theory of advertising, wherein high quality firms simultaneously advertise more and charge higher prices (e.g, Milgrom and Roberts, 1986 and Armstrong, Vickers, and Zhou, 2009).

Figure 3: Listing Price by Sponsorship

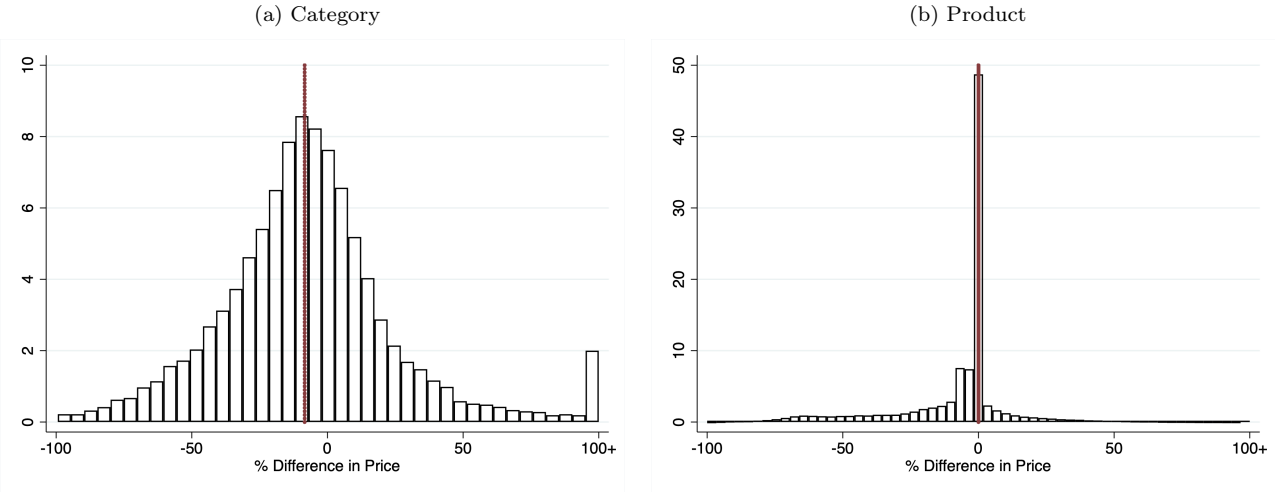

Notes: The percentage difference in price is calculated as $\frac{\bar{p}(sponsored{=}1)-\bar{p}(sponsored{=}0)}{\bar{p}(sponsored{=}0)} \times 100$. The red dotted lined corresponds to the median price difference. Prices are measured exclusive of shipping and insurance charges. Panel A: Each observation is a product category, of which there are more than 17,000. The median price difference across categories is -8.41%. Categories with fewer than 100 sponsored/non-sponsored listings are excluded. Panel B: Each observation is a product (e.g., a unique ISBN or SKU). The median price difference across products is 0%.

3 Field Experiment Design

This paper examines a 2020 experiment run on the platform where sponsored search advertising was blocked for a random sample of 3% of US desktop site traffic. I refer to this treatment as "ads-off." The randomization occurred upon a user's first contact with the site and was maintained throughout the calendar year. All other users 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 user characteristics (including actions and demographics). I call this the "ads-on" condition. This paper analyzes data from February 1 - June 30, 2020, providing an opportunity to examine effects of advertising on purchasing and churn up to a five month horizon.[5] Because the number of site visitors is very large, a random subset of ads-on users is included in the control group (15% of US desktop site traffic). As a randomization check, table 1 provides summary statistics for users included in the experiment. Reassuringly, ads-on and ads-off 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.

[5]During the experiment period, new cookies are continuously randomized into the "ads off" group. Randomization began in January 2020, 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

	Ads-On	Ads-Off	Difference
	(1)	(2)	(3)
Hour of Site Visit	11.8570	11.8568	0.00019
			(0.0011)
Day of Year	107.59	107.57	0.022
			(0.0078)
Share Signed In	0.02380	0.02378	0.000016
			(0.000026)

Notes: This table provides summary statistics for site visitors (cookies) in the test and control groups. Column 3 shows the difference in mean characteristics with standard errors reported below in parentheses.

(Note that while the difference in day of year is statistically significant, the difference is economically tiny.) The likelihood of a prior transaction on the platform is also similar across the test and control groups: ads-on users are a mere 0.00682% more likely to have made a prior purchase than ads-off users (t-stat of 0.03).

I estimate the causal effect of sponsored search on purchasing using a straightforward OLS regression of individual i's spending (y_i) on an indicator for advertising-eligibility ($adson_i$):

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

The coefficient of interest is β, which captures the effect of sponsored search. To preserve sensitive information about the platform's sales, I report effects as percent changes $\left(\frac{\alpha}{\beta}\right)$.

4 Platform Revenue

I first examine the effect of sponsored search on sales of sponsored items, which is the primary focus of the advertising literature. Columns 1-3 in Table 2 present estimates of the relative spending on sponsored items by ads-on vs ads-off users. The point estimates imply that advertising increases spending on sponsored items by approximately 11% and the likelihood of purchase by 6.5%. In short, sponsored listings sell better under sponsored search. The interpretation of this effect differs slightly from most advertising effect estimates in the literature because it captures differences in spending due to own advertising and also rival advertising. In this sense, it is closest to work by Dubé, Fang, Fong, and Luo (2017) and X. Lin, Nair, Sahni, and Waisman (2019), which explicitly study these types of strategic

interactions. Even incorporating interactions, the return to sponsored search for advertised items is large.

Table 2: % Difference in Purchasing

	Sponsored Listings			All Listings		
	Spending		Any Purchase	Spending		Any Purchase
	(1)	(2)	(3)	(4)	(5)	(6)
Sample	Full	Exclude top 0.01%	Full	Full	Exclude top 0.01%	Full
Ads On	0.108***	0.112***	0.065***	-0.005	-0.010*	-0.005***
	(0.009)	(0.008)	(0.002)	(0.011)	(0.005)	(0.001)

Notes: Observations are site visitors from February 1, 2020 - June 30, 2020. Columns 1-3 consider only purchases of listings with an active paid search ad campaign during the experiment window, which can be purchased by ads-off users when they appear as organic listings. Columns 4-6 include all purchases of all listings for sale on the site. Columns 2 and 5 exclude users in the top 0.01% of spending. Heteroskedasticity-robust standard errors in parentheses.

More novel, I next document that this benefit for sponsored listings comes at the expense of overall spending on the platform. That is, I regress a user's total spending on advertising eligibility as in specification (1). Table 2 reports the percent difference in spending between ads-on and ads-off users for the entire sample (column 4) and dropping the top 0.01% of spenders (column 5), where the latter is common practice at the platform to mitigate the influence of outliers. The results indicate that sponsored search reduces total spending by site visitors–including spending both on sponsored *and* organic listings–on the order of 1%, shrinking the economic pie.[6] Column 6 presents estimates where a purchase indicator is the dependent variable; advertising reduces purchase probability by 0.5%, confirming that the decline in spending reflects a true reduction in transaction volume. Appendix table 7 presents analogous regressions with revenue from organic listings as the dependent variable. The results confirm that the decline in overall transactions is due to cannibalization of the sales of organic listings. Figure 4 shows the difference in overall spending between the ads-on and ads-off users for the February cohort for each of their first eight weeks in the experiment. Over this time horizon, the effects of sponsored search is small, negative, and stable.

These estimates show that sponsored search imposes a cost on the platform by reducing the sale of organic listings. This cost is large; it outweighs the 10% increase in sales enjoyed by the advertised items themselves. An experiment aimed only at assessing sponsored items would therefore capture only part of the ad effect relevant to the platform. The implication for e-commerce platforms is that sponsored search poses a tradeoff between transaction

[6]Of course, it is possible that ads-on users increase their spending on other platforms.

Figure 4: Effect over Time for February Cohort

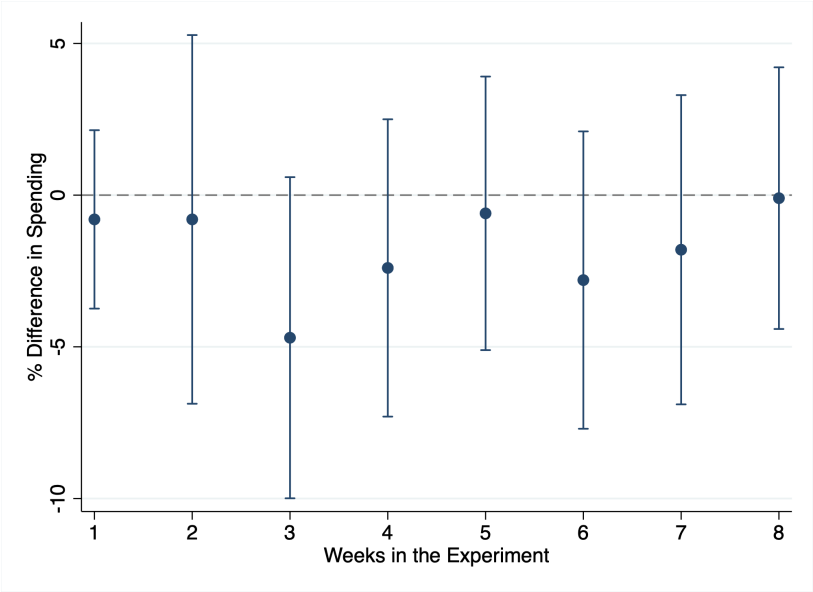

Notes: This plot shows the difference in spending for ads-on vs ads-off users by week since the user joined the experiment along with 95% CIs). The sample consists of users who joined the experiment in February 2020 and is held constant across weeks.

revenue and advertising revenue. While sponsored search is costly, it may still be profitable for the platform so long as advertising revenues outweigh losses in commissions from transactions.

To illustrate this point, I conduct a back-of-the-envelope calculation that highlights the role of ad rates (γ) and commissions (τ), both of which are percentage fees levied on transactions. Absent sponsored search, the platform's profits are $\pi(0) = \tau \cdot R(0)$, where $R(0)$ is total transaction revenue. If the platform allows paid search, then its revenue becomes $\pi(1) = (\tau + \gamma \cdot \delta) \cdot R(1)$ where δ is the fraction of purchases attributed to sponsored listings (this fraction is relevant because sellers only pay the ad rate γ on sales attributed to the ad). The advertising experiment examined in this paper pegs $R(1)/R(0)$ at approximately 0.99, so that sponsored search increases platform profits so long as $\gamma/\tau \geq 0.01/\delta$. To give a sense for magnitudes, the commission rate τ ranges from 8%-15% on Amazon.com; 20% on GrubHub.com; 3%-13% on eBay.com; 3%-20% on Walmart.com; and 5% on Etsy.com.[7] Figure 5 plots the break-even share of purchases attributed to sponsored search (δ) as a function of the ad rates (γ) in the neighborhood of the commissions charged by these large platform.

[7]Fee information is available at `https://get.grubhub.com/grubhub-profit-calculator` (GrubHub); `https://www.ebay.com/help/selling/fees-credits-invoices/selling-fees?id=4364` (eBay); `https://www.etsy.com/legal/fees;https://sellerhelp.walmart.com/seller/s/guide?article=000006011` (etsy); and `https://sell.amazon.com/pricing.html` (Amazon).

Even at the high end of commissions, paid search increases platform profits under fairly modest ad fees ($\gamma \approx 5\%$) and ad attribution rates ($\delta \approx 3\%$); it is therefore unsurprising that the platforms listed above offer sponsored search ads.

To give a sense for what drives cannibalization, I next investigate how paid search affects the ordering of listings on SRPs. For each listing with an active campaign, I calculate the difference in the likelihood that it ranks first on the SRP for ads-on vs ads-off users and then average across listings. Figure 6 presents these calculations: sponsored search meaningfully increases the likelihood that an listing that a seller has selected for sponsorship surfaces in each of the top five positions. On average, sponsored listings occupy 12.3% of the top 15 positions for ads-on users.

Figure 5: Break-Even Ad Analysis

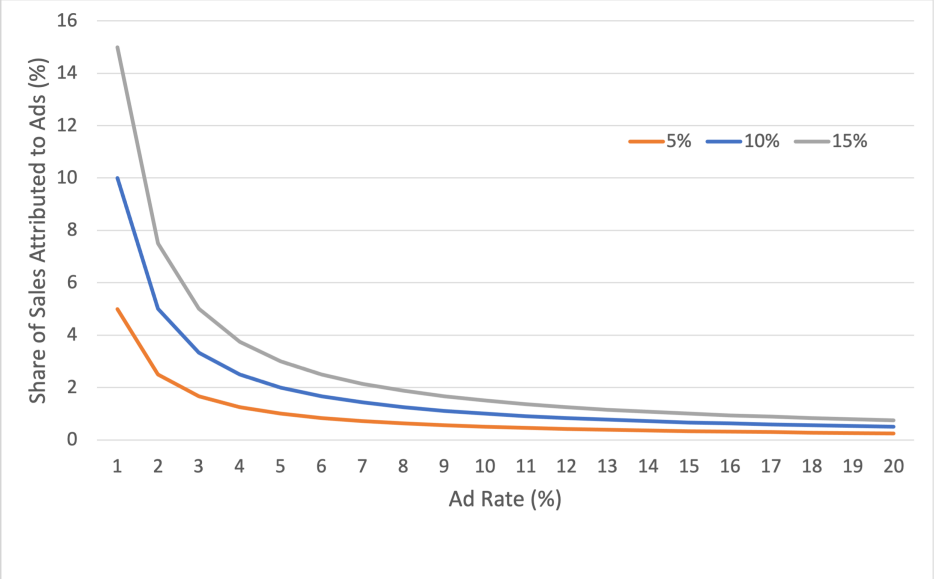

Notes: This figure plots the share of platform revenue that must be attributed to sponsored search advertising for sponsored search to increase platform profits as a function of the ad rate.

Figure 6: Effect of Sponsorship on SRPS

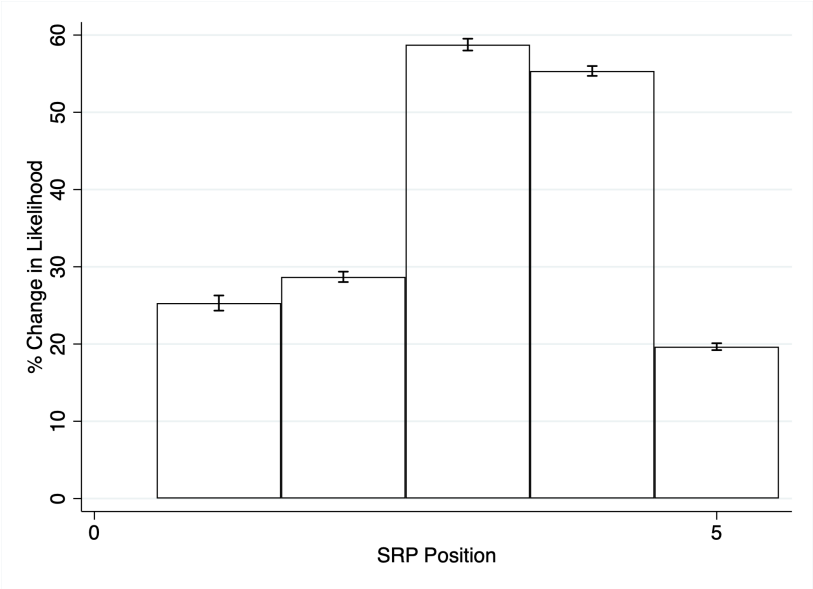

Notes: This figure shows the increased likelihood of surfacing among the top 5 search positions for sponsored items relative to a world without sponsored search. The sample includes all searches conducted by users in the experiment during February 2020 on the US desktop site. Percent changes are calculated from an item-level i regression of the count of SRP appearances among test and control users: $napperances_{ig} = \alpha + \beta \cdot adson_g + \epsilon_{ig}$. I plot the ratio $\frac{\hat{\beta}}{\hat{\alpha}}$ and the associated 95% confidence intervals.

One potential challenge in interpreting these results is that the field experiment randomizes sponsored search eligibility for cookies rather than users, as is typical in e-commerce settings (e.g., Blake, Moshary, Tadelis, and Sweeney (Forthcoming)). It is possible that a single individual might be randomized into different experiences on their laptop and desktop computers. Even on a single device, a user might be assigned to a different 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). Fortunately, more than 80% of users are associated with only a single cookie during the experiment, as shown in figure 7, which suggests that fragmentation is limited in this context.

While this histogram is heartening, it may understate fragmentation because I do not observe all cookies associated with each user. 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 across all cookies. Table 3 reports results, which are similar to those estimated using the full sample. The estimates indicate a negative effect of sponsored search on overall sales and the likelihood of a transaction. Thus, fragmentation does not appear to be a first order concern in this setting.

Figure 7: Histogram of Cookies per User

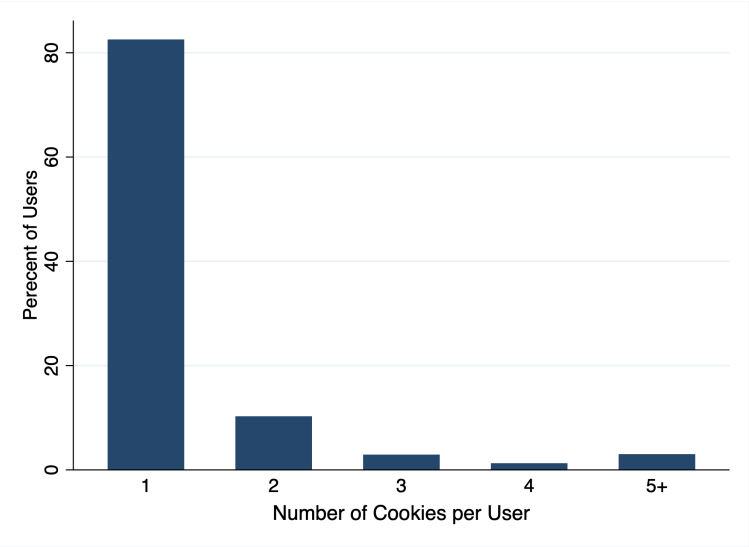

Notes: This histogram shows the distribution of the number of cookies identified with each user in the experiment. This linkage is established when a user logs into their account on the platform. Cookies that are not associated with any user are excluded.

Table 3: Effect of Sponsored Search on Experienced Users

	All Spending		Any Purchase
	(1)	(2)	(3)
Sample	Full	Exclude top 0.01%	Full
Ads On	-0.033∗	-0.009	-0.004∗
	(0.013)	(0.008)	(0.002)

Notes: Observations are site visitors from February 1, 2020 - June 30, 2020 in the experiment who were identified with a user id before the experiment (before Jan 1, 2020). Heteroskedasticity-robust standard errors in parentheses. Column 2 excludes users who were among the top 0.01% of spenders in the experiment.

5 Match Quality and Search Costs

That sponsored search leads to a decline in transactions does not necessarily imply that consumers dislike paid search ads. One possibility is that sponsored search steers consumers away from unsatisfying products that they would regret purchasing. This in turn could provide a benefit to the platform and to sellers by reducing the administrative and logistical costs of returns and disputes over faulty products. If advertising steers consumers to high quality products, then we might expect them to purchase repeatedly, as in Milgrom and Roberts (1986). Table 4 presents results to the contrary: ads-on purchasers are marginally less likely to complete a second purchase (-0.02%) although the difference is not statistically

Table 4: Effect of Sponsored Search on Complaints

	Purchase Twice		Any Complaint		Complaints per Transaction	Complaint on 1st Transaction
	(1)	(2)	(3)	(4)	(5)	(6)
Sample	Full	At least one purchase	Full	Conditional on Purchase	Conditional on Purchase	Conditional on Purchase
Ads On	-0.006***	-0.002	-0.011***	-0.006*	-0.002	-0.005
	(0.001)	(0.001)	(0.003)	(0.003)	(0.004)	(0.004)

Notes: Sample includes complaints for purchases on the platform between February 1, 2020 - June 30, 2020 by site visitors in the experiment as of March 2021. A buyer complaint comprises any of the following scenarios: the buyer initiates a product return; the buyer leaves negative or neutral comments; the buyer does not in fact receive the product. Heteroskedasticity-robust standard errors in parentheses.

significant. Admittedly, many of the items for sale on the platform are durable goods that we would not expect consumers to purchase twice within a few months. However, to the extent that a consumer learns about platform reliability through a transaction, a positive initial experience might spur future purchases of other goods, an effect that I do not detect here.

To provide further evidence on quality, I leverage a second measure of consumer satisfaction: complaints to the platform. A complaint comprises any of the following scenarios: the buyer initiates a product return; the buyer leaves negative or neutral comments; or the buyer does not in fact receive the product. Table 4 presents results of regressing different measures of complaints on sponsored search eligibility. The third column shows results for the entire sample; ads-on users are less likely to initiate a complaint (approximately -1%), which is unsurprising because they are also less likely to transact on the platform. The coefficient on the ads-on indicator halves when I condition on at least one purchase (column 4), but it is still difficult to interpret because conditional on purchasing, ads-off consumers engage in more transactions. Column 5 therefore normalizes complaints by the number of total transactions, and the estimated effect attenuates and loses statistical significance. Column 6 considers the likelihood of a complaint on the first transaction, and the point estimate halves again. The 95% CI precludes a 1% reduction in the likelihood of a buyer complaint on his or her first transaction. Taken together, the estimates in the last two columns of table 4 suggest that sponsored search does not meaningfully reduce the likelihood of a disappointing purchase.

A third avenue through which sponsored search can affect consumer welfare is via search costs. In particular, sponsored search is hypothesized to benefit consumers 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 affects consumer search. Ads-on users do indeed search for shorter periods of time (-1.08%, column 1). They click on roughly the same number of items (difference of 0.3%, not statistically significant), but conduct 3% fewer searches (significant at the 1% level). In contrast with the purchasing results described above, these patterns do not provide a clear conclusion as to whether advertising facilitates or inhibits search. However, they do provide a clear example where purchasing and clicks move in the opposite direction (purchasing falls, but clicks marginally increase), which cautions researchers against proxying purchasing with upper funnel metrics in similar settings.

Table 5: Effect of Sponsored Search on Search Behavior

% Difference between Ads-On and Ads-Off Users		
	All (1)	Conditional on Purchase (2)
Search Duration	-1.09% (0.026)	0.02% (0.18)
# Searches	-3.09% (0.34)	0.52% (0.37)
# of Items Clicked	0.30% (0.36)	1.42% (0.46)
Page Rank of Initial Click on a SRP	-6.61 (0.07)	-3.61% (0.73)
Page Rank of Purchase	_	6.11% (0.033)
Out-of-Order Search	6.47% (0.13)	4.56% (0.43)

Notes: This table reports search behaviors for ad-eligible vs -ineligible users. Effects are reported as percent differences with standard errors in parentheses. The sample is search sessions for all site visitors that participated in the experiment during February 2020. Search duration measures session length in milliseconds. The number of searches and items clicked are measured per user across all sessions.

6 Price Effects of Advertising

This section investigates the interplay between advertising and pricing. As summarized by Bagwell (2007), the economics literature advances a range of hypotheses about this relationship: sellers may pass along advertising costs to consumers through higher prices; 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 toughen price competition if it lowers barriers to entry (e.g., by mitigating the cold start problem for new sellers).

A first finding from the 2020 field experiment is that ads-on users purchase items that are 1.61% more expensive than ads-off users. This difference is driven by differences in the transacted price of both sponsored and organic items (+1.605% vs +1.613%, with t-statistics of 6.10 and 1.67, respectively). Second, ads-on users are 1.09% (t-stat of 53.9) less likely to search based on price using the ascending price filter (that is, altering the default search ordering so that the least expensive items appear first). These two patterns are consistent with advertising increasing the elasticity of demand. However, because ads-on users see different listings at the top of SRPs, this is not definitive evidence that they are relatively price inelastic. For example, it is possible that ads-on users do not filter on price because sponsored items are less expensive.

A Natural Experiment in Sponsored Search

In the 2020 field experiment, sellers are unable to set separate prices for ads-on and ads-off users. Because 97% of site traffic sees sponsored listings, I interpret observed prices as equilibrium prices in an "ads-on" environment.[8] 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 2018, the platform reserved most 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 paid search volumes in a five month window.[9] Ceteris paribus, if sponsored search intensifies (softens) price competition, then we would expect prices to fall (rise) following this marked increase in advertising volumes.

Of course, prices might change over time for reasons unrelated to the platform's paid

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

[9] 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

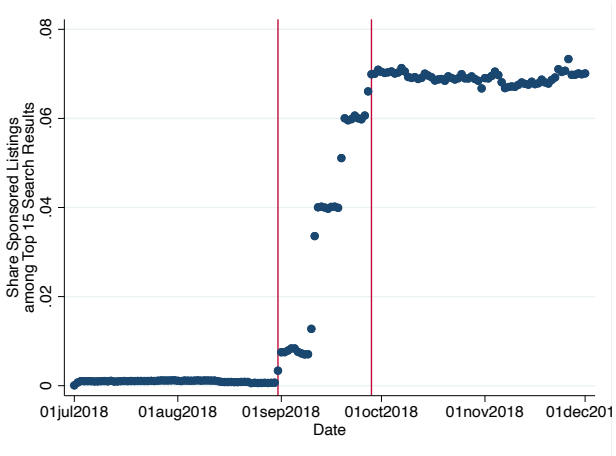

Notes: This plot shows the percent of top 15 SRP positions allocated to sponsored listings from July1-Nov 30, 2018 across all site visitors. The dates delineated by dashed red lines span August 31-September 28, 2018.

search policy. I therefore 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 platform-wide trends and shocks that might influence prices. The estimating equation specifies the log price in category c on date t as a function of category fixed effect Ω_c, date fixed effects Λ_t, and interactions between indicators for the ramp-up period, $Interim_t$, the post period, $Post_t$, and an indicator that a category allows sponsored search, $Eligible_c$:

$$logp_{ct} = \delta \cdot Post_t \times Eligible_c + \gamma \cdot Interim_t \times Eligible_c + \Omega_c + \Lambda_t + \epsilon_{ct} \tag{2}$$

Column 1 in table 6 reports estimates of the key parameter δ, the effect of sponsored search on transacted prices. The point estimate suggests that sponsored search increases transacted prices by 2.6%, but I cannot reject the null hypothesis of no effect. In column 2, I estimate specification (2) at the product level, which permits inclusion of product fixed effects. Including these fixed effects eliminates variation in prices due to inventory changes, but also restricts attention to products that can be categorized (such as a textbooks that can be linked to an ISBN). The estimate of δ barely moves. Figure 2(a) shows the relative change in residual transacted prices for categories that do/do not allow sponsored search; it is hard to discern a price increase.

I next consider the effect of sponsored search on listing prices, with a focus on new and revised listings. I focus on these new and revised listings because they require the seller to actively set prices. Estimates are presented in table 6 column 3. They suggest that the

Table 6: Price Effects of Sponsored Search

	Transactions		Listings
	(1)	(2)	(3)
Post × Eligible Category	0.026	0.025	0.139*
	(0.025)	(0.016)	(0.064)
Interim × Eligible Category	0.006	-0.006	0.031
	(0.025)	(0.025)	(0.047)
FE	Date & Category	Date & Product	Date & Category

Notes: Data from February - November 2018. Regressions are weighted by volume. The outcome variable is log price.

prices for new and revised listings increased following the expansion of sponsored search, on the order of 10%.[10] Figure 9(b) presents the graphical analogue of this regression; with the exception of one week in July, there is virtually no difference in the movement of prices in categories that allow/do not allow sponsored search until October. Starting in October, a gradual increase in listing prices emerges. Of course, a causal interpretation of these patterns requires stronger assumptions than the 2020 field experiment analysis. As a robustness check, I therefore replicate figure 2 using data from 2017. I find no increase in listing prices, as shown in figure 12, although there is the number of listings is increasing over this period. Thus, these findings should be interpreted with caution. That said, the observed increase in listing prices dovetails with the comparative static in Armstrong, Vickers, and Zhou (2009), where most firms increase prices in response to the advent of advertising.

7 Entry

This section considers a second general equilibrium effect: whether sponsored search encourages entry. One potential benefit of advertising is that it allows new firms and products to gain a foothold in the market. This potential to solve the "cold-start" problem is highlighted by Sahni and Zhang (2020) as a key mechanism by which sponsored search can increase consumer welfare. Just as with pricing decisions, however, identifying the effect of advertising on entry poses an empirical challenge that is not amenable to the design of the 2020 field experiment wherein sellers make a single entry decision that does not vary depending on the user's treatment status. Nonetheless, I can test whether sponsored search primarily benefits new and small sellers, who have yet to establish a reputation on the platform, by comparing seller revenue among ads-on and ads-off users in the 2020 experiment. From the perspective of each seller, the 2020 experiment randomly shields a 3% share of their market from advertising.

[10] New listing data is too sparse to allow for the inclusion of product fixed effects within the difference-in-difference.

Figure 9: Prices in 2018
Ineligible vs Eligible Categories

(a) Transacted Prices

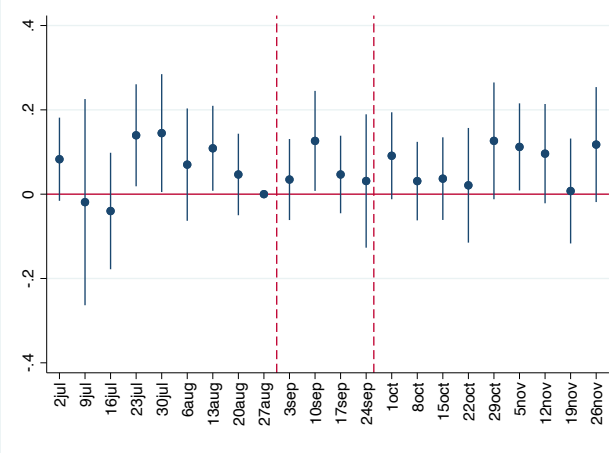

(b) Listing Prices

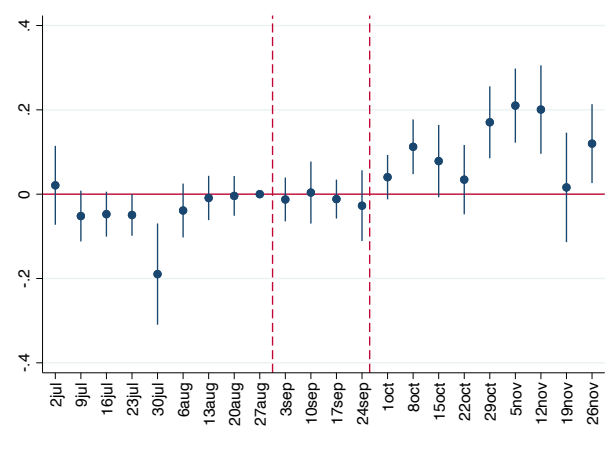

Notes: The dates delineated by dashed red lines span August 31-September 28, 2018. Left plot: this plot shows relative transaction prices in categories that do/do not allow sponsored search advertising. Right plot: this plot shows the relative prices of new and modified listings in categories that do/do not allow sponsored search.

Figure 10: Share of Sellers with Higher Revenue under Advertising
by Quantile of Sales in 2019

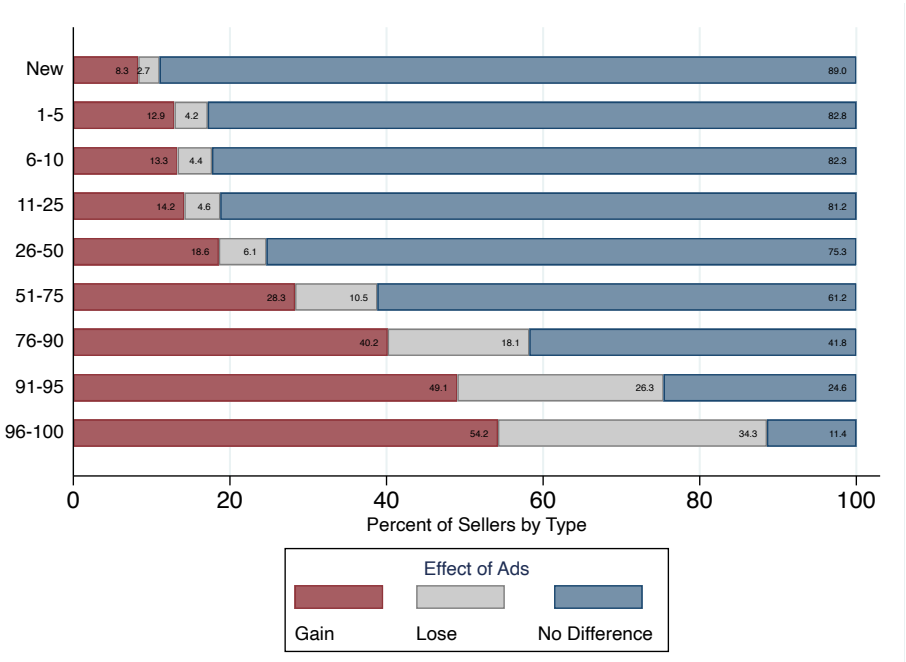

Notes: The sample is all sellers with at least one listing between February 1 - June 30, 2020 on the platform's US site. Sellers are grouped by quantile of sales July-December 2019. "New" denotes sellers who had no transactions in the latter half of 2019 (e.g., sellers that entered the site in 2020).

Figure 11: Effect of Sponsored Search by Seller Type

(a) % Change in Revenue

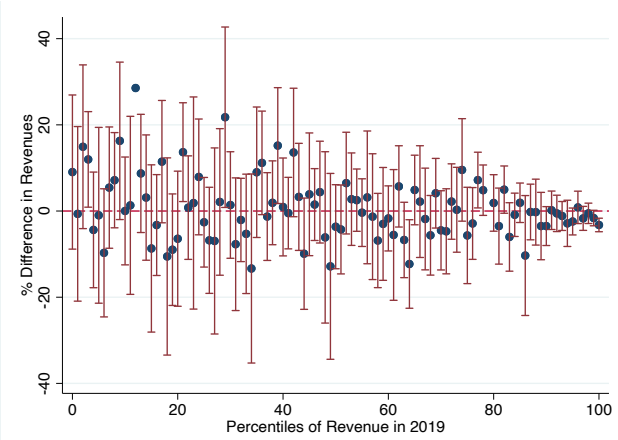

(b) % Change in Revenue per Sponsored Listing

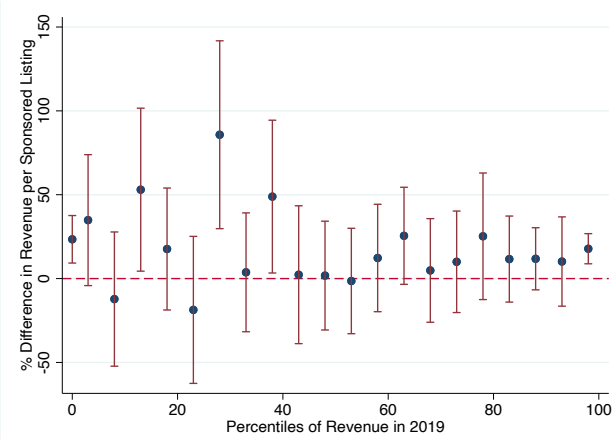

Notes for figure a: For each percentile of seller size by revenue in 2019, this figure plots the average percent difference in sales between ads-on and ads-off users in the 2020 experiment. Specifically, I run the following regression centile-by-centile, where s denotes the individual seller and g denotes the consumer group (test or control): $revenues_{sg} = \alpha + \beta \cdot adson_g + \epsilon_{sg}$. This gives a set of parameters $\{\alpha_i, \beta_i\}_{i=0}^{100}$ where i denotes the seller centile in 2019 sales. I then plot $\frac{\hat{\beta}_i}{\hat{\alpha}_i} \times 100$ along with 95% confidence intervals. The zero quantile refers to sellers with zero revenue between July-December 2019, e.g., new sellers. *Figure B*: For each percentile of seller revenue, I plot the percent change in revenue per sponsored listing, pooling across vintiles.

To measure seller size, I link sellers with their transactions and listings in Q3 and Q4 of 2019, the year before the experiment. I bin sellers by their percentile of sales 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 10). New sellers, defined as those with no sales in the second half of 2019, are plotted separately as the 0th quantile. Over three times as many new sellers earn higher revenues in the paid search environment compared to an ad ban. However, the modal large seller also earns higher revenue under sponsored search–in fact, across all seller types, more sellers benefit from advertising than are harmed by it.

To give a sense for magnitudes, Figure 6 plots the percent difference in sales among ads-on and ads-off users for each seller segment. The confidence intervals are wide, and there is no clear correlation with seller size. In principle, effects might differ for small and large sellers for at least two reasons: first, as shown in Figure 2, the use of sponsored search varies with seller size; and second, the effect of sponsorship may vary with size if larger sellers are better at designing campaigns (e.g., selecting which items to sponsor) or if reputational spillovers matter more for small sellers. To disentangle these possibilities, I examine the

relationship between revenues *per sponsored listing* and the user's ad eligibility. Across the board, estimates are large and positive, including for new sellers and the very largest sellers. Seller effects are further explored in Appendix B, which includes regression tables showing the effect of paid search on revenues. Overall, these estimates show that sponsored search can indeed boost the sales of new and small sellers, but that benefits are not unique to this group.

8 Conclusion

Sponsored search is an important component of the digital economy, surpassing $50 billion in revenue in 2019.[11] Naturally, much attention in industry and academia has focused on quantifying the returns to paid search from the perspective of the advertising firm to understand whether this money is well-spent. These ROI analyses often hold fixed the actions of other economic agents. This paper analyzes the return to paid search advertising for the platform with particular attention to general equilibrium effects, including spillovers to other products and price adjustments. I first show that paid search cannibalizes organic sales, so that the net effect is a reduction in total transaction volume on the platform. I establish this finding using data from a large-scale field experiment at an anonymous e-commerce platform that blocks all paid search advertising for 3% of site visitors. Over the course of five months in 2020, these users spend more and are more likely to complete a purchase than their counterparts who see ads. One immediate implication is that advertising incurs an opportunity cost in commissions for e-commerce platforms.

The finding that advertising shrinks the economic pie contravenes a central prediction of models of informative advertising, wherein ads guide consumers to high quality products. Further, the results of the experiment do not support the hypothesis that sponsored search increases equilibrium match quality, as ads-on users are equally likely to engage in a return or complaint as users who do not see ads. Nor do the results of the experiment suggest a marked decline in search costs. Indeed, while ads-on users spend less time on the platform, they click on more listings, highlighting that clicks can be a poor proxy for purchase in e-commerce settings.

Because the experiment shields a mere 3% of site traffic from paid search, I interpret observed pricing decisions during the experiment as equilibrium prices under paid search. To understand how advertising affects pricing decisions, I exploit the ramp up of sponsored search on the platform in 2018 that affected some categories but not others. While the difference-in-differences analysis requires stronger identification assumptions, the results suggest that advertising exerts upward pressure on prices, which would tend to amplify the

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

negative effect of sponsored search on transactions and revenues. These findings are also consistent with observed differences in purchasing between ads-on and ads-off users in the field experiment: ads-on users buy more expensive items and are less likely to filter items by price.

Just as advertising affects prices, it may also affect seller entry decisions. While I cannot get at these directly, the 2020 field experiment does reveal that small sellers do not benefit disproportionately from an ads-on environment. Further, a large entry effect is hard to reconcile with descriptive evidence that new sellers seldom participate in ad auctions on the platform. Low participation rates might indicate a fixed cost of designing a campaign, which the platform could alleviate by offering seller tools. However, the results of the field and natural experiments analyzed in this paper suggest that platforms must weigh the cost in forgone transaction revenues in such an expansion of sponsored search.

References

Abhishek, Vibhanshu, Kinshuk Jerath, and Siddhartha Sharma (2019). "Advertising on Online Marketplaces: Information Asymmetry and the Relevance of Sponsored Listings". In: *Working Paper*.

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). "The Economic Analysis of Advertising". In: *Handbook of Industrial Organization* 3, pp. 1701–1844.

Blake, Tom, Sarah Moshary, et al. (Forthcoming). "Price Salience ad Product Choice". In: *Marketing Science*.

Blake, Tom, Chris Nosko, and Steven Tadelis (2015). "Consumer Heterogeneity and Paid Search Effectiveness: A Large-Sale Field Experiment". In: *Econometrica* 83.1, pp. 155–174.

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

Choi, Hana and Carl Mela (2019). "Monetizing Online Marketplaces". In: *Marketing Science* 38.6, pp. 913–1084.

Dai, Daisy and Michael Luca (2016). "Effectiveness of Paid Search Advertising: Experimental Evidence". In: *Working Paper*.

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

Gordon, Brett et al. (2019). "A Comparison of Approaches to Advertising Measurement: Evidence from Big Field Experiments at Facebook". In: *Marketing Science* 38.2.

Huang, Jason, David H. Reiley, and Nickolai M. Riabov (2018). "Measuring Consumer Sensitivity to Audio Advertising: A Field Experiment on Pandora Internet Radio". In: *Working Paper*.

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.

Narayanan, Sridhar and Kirthi Kalyanam (2015). "Position Effects in Search Advertising and their Moderators: A Regression Discontinuity Approach". In: *Marketing Science* 34.3, pp. 388–407.

Nelson, Philip (1974b). "Advertising as Information". In: *Journal of Political Economy* 82.3, pp. 729–754.

Sahni, 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*.

Sahni, Navdeep and Clark Zhang (2020). "Search Advertising and Information Discovery: Are Consumers Averse to Sponsored Messages?" In: *Working Paper*.

Shapiro, Bradley (2018). "Positive spillovers and free riding in advertising of prescription pharmaceuticals: The case of antidepressants". In: *The Journal of Political Economy* 126.1, pp. 381–437.

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

Wilbur, Kenneth (2016). "Advertising Content and Television Advertising Avoidance". In: *Journal of Media Economics* 29.2, pp. 51–72.

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

A Effect on Organic Revenue

The table below shows the effect of sponsored search on the sales of organic items, the analogue of table 2.

Table 7: % Difference in the Purchasing of Organic Items

		Spending	Any Purchase
	(1)	(2)	(3)
	Full	Exclude top 0.01%	Full
Ads On	-0.028*	-0.036***	-0.017***
	(0.013)	(0.005)	(0.001)

Notes: Observations are site visitors from February 1, 2020 - June 30, 2020. Organic items are those without an active ad campaign during the experiment window. Columns 2 excludes users in the top 0.01% of spending. Heteroskedasticity-robust standard errors in parentheses.

B Effects on Seller Revenues

This appendix examines the relationship between per-listing revenues and advertising for sellers on the platform. The baseline specification is:

$$\text{revenue per sponsored listing}_{sg} = \alpha_0 + \alpha_1 \cdot adson_g + \Gamma_s + \epsilon_{sg} \tag{3}$$

where s denotes seller, g denotes user group (ads-on or ads-off), Γ_s are seller fixed effects, and the sample comprises sellers that sponsor listings. Essentially, this framework compares purchasing rates between ads-on and ads-off users within seller. Estimates of the regression parameters are presented in table 8. On average, sellers earn 17% higher revenues on their sponsored listings in a sponsored search environment ($\frac{\hat{\alpha}_1}{\hat{\alpha}_0} \times 100$, column 1). I then extend the specification to allow for differential effects across listings types $l \in \{sponsored, standard\}$ to measure how sponsored search affects standard listings (listings that are not sponsored and so surface as organic listings):

$$\begin{aligned}\text{revenue per listing}_{slg} = \alpha_0 + \alpha_1 \cdot adson_g + \alpha_2 \cdot adson_g \times sponsored_l \\ + \alpha_3 \times sponsored_l + \Gamma_s + \epsilon_{sg}.\end{aligned} \tag{4}$$

The estimates in table 8 column 2 hints at positive selection into sponsorship: sponsored items earn 47% higher revenue even among ad-ineligible users ($\frac{\hat{\alpha}_3}{\hat{\alpha}_0} \times 100$). This difference could reflect both seller decisions about which listings to sponsor, but also potential spillovers. As an example, sponsorship might generate positive reviews that increase sales to all users.

Figure 11 (b) plots these effects separately across seller types, but the confidence intervals are wide and it is hard to draw a conclusion about differential effects.

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 listing is posted by a seller who sponsors at least one listing, $sponsoringseller_s$:

$$\begin{aligned} \text{revenue per listing}_{slg} = \alpha_0 + \alpha_1 \cdot adson_g + \alpha_2 \cdot adson_g \times sponsored_l \\ + \alpha_3 \times sponsored_{sl} + \alpha_4 \cdot adson_g \times sponsoringSeller_s \\ + \Gamma_s + \epsilon_{sg} \end{aligned} \tag{5}$$

The coefficient α_4 captures the spillover effect of advertising on standard listings sold by sellers who advertise other listings. 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 8 can be interpreted as percent changes. The point estimate for the spillover effect is economically and statistically insignificant.

Table 8: Effect of Sponsored Search on Revenues

	Revenue per Listing		
	(1)	(2)	(3)
Sample	Promoted	All	All
Ads On	17.000***	4.427	3.404
	(2.378)	(3.155)	(2.883)
Ads On × Sponsored		16.414*	7.904
		(8.274)	(18.591)
Sponsored		46.807***	51.317***
		(10.393)	(9.482)
Ads On × Sponsoring Seller			0.096
			(0.172)
Seller FE	X	X	X

Notes: Standard errors are clustered by seller. Column 1 restricts to sellers with at least one sponsored listing. Columns 2 & 3 include all sellers with a listing in the US site during the experiment. Each observation is a seller × sponsorship combination (sponsorship = 1 or 0). This table reports the ratio of the coefficient over the constant so that magnitudes are interpretable as percent changes.

C Placebo Event Study

This appendix section replicates the event study plots in sections 6 and 7 for 2017, when there was minimal paid search advertising on the platform. Figure 12 (a) shows the evolution of listing prices for new and revised listings, which are essentially flat. However, as shown in figure (b), the relative number of listings is increasing over the entire window. To the extent that this increase is driven by an omitted variable that also differentially affects eligible categories in 2018 and after the introduction of sponsored search, that omitted variable would bias the estimated price effects.

Figure 12: Placebo Event Study in 2017:
Categories that Do/Do Not Allow Sponsored Search

(a) Prices for New & Revised Listings

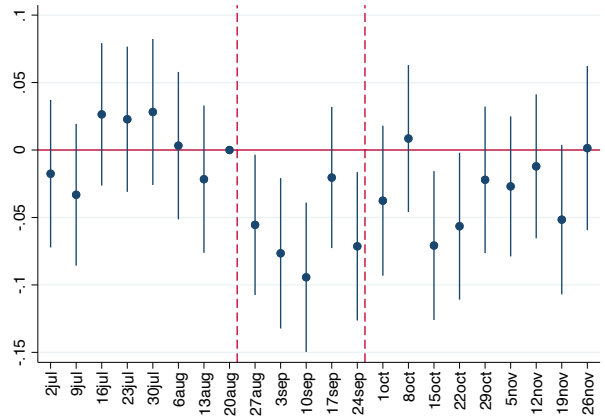

(b) Log Number of Listings

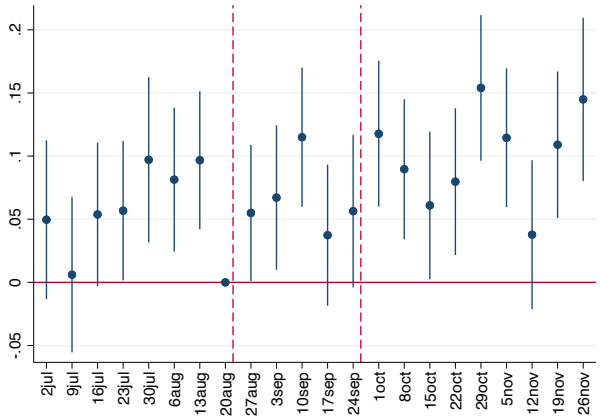

Notes: Figure a shows the relative prices for new and revised listings in 2017 in categories that are eligible vs ineligible for advertising under the 2018 platform rules. Prices are measured in logs. Figure b shows the relative number of new listings, which are also measured in logs, for a balanced panel of categories (where there is at least one new listing per day).

D Out-of-Order Search

One clear difference in the observed search behavior of ads-on and ads-off users is their choice of where to click. Ads-on users initially click on listings nearer the top of the results page, but conditional on making a purchase, they select items that are further from the top of the SRP. This pattern is consistent with sponsored search attracting attention, but losing it more often than counterfactual organic listings. I probe this further by considering out-of-order search, where a consumer clicks first on a higher-ranked item before returning to click on a lower-ranked item second; ads-on users are more likely to click out-of-order. In other words, ads-on users are less likely to engage in top-down search. This pattern

suggests that sponsored search erodes the consumer belief 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-on group, I calculate the relative likelihood of consumers skipping over sponsored listings to explore higher-ranked organic items and vice versa:

$$\frac{Pr\{\text{out-of-order, sponsored first, organic second}\}}{Pr\{\text{out-of-order, organic first, sponsored second}\}} = 1.42$$

In words, out-of-order search occurs more commonly when consumers skip over organic listings to examine sponsored listings. This finding is consistent with consumers inferring unobserved quality from sponsorship. It is also consistent with advertising increasing salience, perhaps via the "sponsored" tag that delineates sponsored listings.