Advertising as Coordination: Evidence from a Field Experiment*

Apostolos Filippas John J. Horton Diego Urraca

November 3, 2023

Abstract

Paid advertising can increase market efficiency by directing buyers to sellers with greater capacity. In a field experiment in a large marketplace, all sellers could advertise, but buyers were randomized into seeing advertising. Contrary to concerns that buyers might infer advertising sellers were adversely selected, treated buyers sought out advertisers. This shift in buyer attention to higher-capacity sellers increased market volume by 2%. Costly advertising was necessary to facilitate this coordination, as mere statements about capacity had become uninformative. A within-seller analysis shows that two years after platform-wide introduction of advertising, advertisers still enjoyed 50% more buyer inquiries while advertising.

JEL Codes: D82, D83, J01, L8, M3

*We gratefully acknowledge generous financial support from the NET Institute (www.netinst.org). Thanks to seminar participants at the NYU Stern Technology Department and Microsoft Research New England. Thanks to Susan Athey, Kyle Bagwell, Michael Ostrovsky, Adam Ozimek, Brad Shapiro, and Robert Somogyi for their helpful comments and suggestions. "Diego Urraca" is a pen name for an employee of the anonymous platform market, and is used to preserve the anonymity of the platform. This work was not subject to prior review and reflects only the opinions of the authors. The latest version of this paper can be found at https://www.apostolos-filippas.com.

1 Introduction

There is a long-standing debate about the role advertising plays in markets (Bagwell, 2007). Critics of advertising contend that advertising can distort consumer preferences, increase costs, concentrate market power, and waste resources. Those more favorably disposed to advertising argue that advertising helps connect buyers and sellers by conveying market-relevant information. For example, advertising can inform potential buyers and sellers about each other, or it can allow advertisers to send a costly signal that may reveal information about product quality (Stigler, 1961; Nelson, 1974; Milgrom and Roberts, 1986). However, these positive conceptions of advertising may be less relevant in a world where market information is almost free, and abundant product information is available online. Woodcock (2017) makes this point in the provocatively and informatively titled law review article "The Obsolescence of Advertising in the Information Age."[1]

In this paper, we consider the role of advertising in the context of a digital marketplace for services. For the first time, sellers were given the opportunity to buy paid advertising. The advertisement was a "badge" with the text "Available Now" that appeared next to a seller in search results. The buyer could also see a notice that the seller paid for this badge. Importantly, these paid advertisements did not give sellers greater visibility to buyers: advertising did not change search rankings, nor the size of the displays in the search results (Edelman, Ostrovsky and Schwarz, 2007; Athey and Ellison, 2011; Decarolis and Rovigatti, 2021). In the experimental phase, all sellers could advertise—but only randomly treated buyers could *see* the advertisements. We, as the experimenters, know which sellers advertised even though untreated buyers did not.

The platform hoped that paid advertising would lead to only relatively more available sellers advertising. Buyers would, in turn, direct their attention to these advertising sellers. The market failure the platform hoped to overcome is that buyers are uncertain about sellers' capacities, which causes them to pursue unavailable sellers (Horton, 2019; Fradkin, 2023). Despite this hope, the platform feared that advertisers would be adversely selected and buyers would learn to simply ignore advertisers, which in turn would cause demand for advertising to disappear.

The goal of the initiative was, in economic terms, to try to make search less "random" (Mortensen and Pissarides, 1994) and more "directed" (Wright, Kircher, Julien and Guerrieri, 2021). Our primary research question is whether the introduction of advertising improved market efficiency in equilibrium, and if so, why. Although our context is one particular digital market, the answers to these questions are informative about the role advertising plays in markets more generally, and the economic problem of strategically missing information that we describe is quite general.

[1] The legal stake was that the FTC should view all advertising as persuasive rather than informative.

Our key results are as follows. We find strong evidence of virtuous selection into advertising, particularly concerning the sellers' capacities to take on more work. For buyers that can see advertising, we find that (1) they seek out advertising sellers, (2) they make more inquiries to sellers in total, (3) they receive more positive responses to their inquiries, and (4) they are more likely to form a contract and transact with a seller. Critically, the increase in transaction probability was not at the expense of non-advertising sellers—there was a *net* increase in sales of about 2%. We next describe our evidence for these results in more detail.

As we observe which sellers chose to advertise, we can compare advertisers to non-advertisers. Counterintuitively, advertising sellers were busier on average—they had more active contracts and were already receiving more buyer inquiries. Despite their greater busyness, advertising sellers were much more likely to respond positively to a buyer's inquiry by submitting a proposal to the project listing. As such, a naive algorithmic approach that directed more buyer attention to less busy sellers would likely exacerbate the problem of buyers pursuing unavailable sellers. "If you want something done, ask a busy person" was empirically true in our case.

Consistent with this virtuous selection into advertising, when we look at impressions—sellers being presented to buyers in an algorithmically ranked order—buyers who could see the advertisements (treatment group) are far more likely to contact an advertising seller. Although we do not know what buyers believed, a parsimonious explanation is that buyers inferred correctly that advertising sellers were more likely to respond to their inquiries. Advertising sellers enjoy an increase in buyer inquiry probability equivalent to about a one-position jump in search rank, in the first search page.

Although buyers showed a strong preference for advertising sellers, non-advertisers were not crowded out by advertisers, at least in aggregate. The reason is that buyers could see advertising sent more recruiting inquiries, presumably because the added information made more sellers appear worthy enough to send an inquiry to. This lack of crowd-out matters because advertising that solely redirects buyer attention from one seller to another might have little welfare import without some offsetting gain to the probability or quality of a match.

One might wonder why paid advertising is needed to reveal the capacity information. The answer lies in the incentives faced by sellers. From the seller's perspective, inquiries and offers are options, and options are valuable even if not pursued. When the cost of receiving an inquiry is low—and it surely is—sellers have little incentive to reveal their capacity truthfully. Empirically, before the experiment, sellers could indicate they were interested in more sales at no cost by reporting high capacity. This "free" advertising had existed for over a decade (Horton, 2019). However, nearly all sellers stated high capacities, and they rarely changed their capacity status, rendering free advertising useless. In short, when advertising was free everyone advertised, and advertising subsequently lost all of its meaning, with the logic of Nelson (1974) still applicable.

The concept of virtuous selection in advertising assumes that sellers exhibit variability in the attributes that buyers care about. This is typically the case in any market, as sellers vary in quality, level of service, and in other attributes. However, in a competitive market, it might be reasonable to ask why variation exists in the vertical attributes of sellers, conditional upon price. In our empirical context, for example, we might wonder why some sellers have lower costs or higher capacities than others, and why this is not reflected in prices.

As an explanation for this variation in capacity, we present a dynamic model of a matching market where the sellers' matching status determines endogenously their expected cost for completing a project. In our model, sellers who are busy are more likely to have higher costs, but buyers search for sellers without knowing these realized costs. In the equilibrium of this market without advertising, welfare is reduced because buyers attempt to trade with "busy" sellers with low capacity. We show that introducing a costly signal creates a separating equilibrium where all sellers with capacities above a threshold choose to advertise, and all sellers below the threshold do not. As a result, buyers direct a higher degree of their search efforts towards sellers with higher capacities, and matching efficiency can increase.

The mechanism through which costly advertising can increase efficiency in our model is by solving a coordination failure. However, a seller with minimal capacity might still advertise if the cost is low enough, as buyer inquiries and offers are valuable.[2] In short, advertising has to be costly enough to create a separating equilibrium but not so costly it leads to a pooling equilibrium where even available sellers do not advertise. The model shows that paid advertising can implement the constrained-efficient equilibrium under plausible conditions. This result requires constant matching efficiency and equilibrium variation only in the probability of acceptance, not price conditional upon a match—a reasonable assumption in a highly competitive market and one borne out in our data. We have no evidence that sellers change their prices depending on whether or not they are advertising.

One concern with our results is that they are the product of a short-term process and that the benefits of advertising would dissipate in the longer run. For example, perhaps buyers only sought-out advertising sellers because the advertising was novel or eye-catching. However, we would expect these effects to die out if buyers learned that, say, sellers were adverseley selected, perhaps on dimensions we are researchers did not understand. We can test this by looking at the effects of advertising over time. Using a a within-seller analysis, we find that two years after platform-wide roll-out, when sellers advertise, they receive 50% more buyer inquiries. This is consistent with the idea that advertising changes the economics of the platform an is not simply a salience effect.

The main contribution of this paper is to show advertising can be a market-coordinating

[2]For example, a custom widget maker who is already very busy might profitably quote an extremely high price to would-be widget buyers—and hence be interested in advertising—even if some less-busy widget maker is the low-cost provider.

mechanism: it can direct buyers to appropriate sellers, resulting in matching formation benefits. An implication is that despite the information-rich context of computer-mediated markets, advertising is not obsolete. This is because the information conveyed by advertising was missing for economic reasons—not technical reasons that digitization alone can solve. Our empirical context is a designed market where the platform can create the possibility of advertising, whereas, in conventional markets, advertising arises organically. But it seems probable advertising in other contexts serves a similar economic function: it helps direct business to appropriate sellers (Bagwell and Ramey, 1994). The reason for this likely generality is the economic problem of buyer uncertainty about seller capacity or suitability is general. So long as not all sellers are equally interested in more sales at a moment in time, advertising can play a role in directing search.

The advertising we study is a kind of sponsored search advertising. As such, our paper is related to a burgeoning but unsettled empirical literature on the effects of sponsored search advertising in online marketplaces. We discuss how our findings can potentially rationalize some of the divergent results. In particular, we argue that whether advertisers are relatively adversely or virtuously selected can determine the effect of advertising on total sales, which is a key point of dispute in the literature.

The rest of the paper is organized as follows. Section 2 describes our study's empirical context and experimental design. Section 3 explores which sellers choose to advertise. We examine the effects of the treatment for buyers in Section 4, and for buyers in Section 5. Section 6 develops a simple model of costly signaling with advertising that can rationalize our results. Section 7 analyzes the data from two-years after the platform introduction of advertising, exploring the long-run effects. Section 8 discusses our results in light of the literature on sponsored search advertising. We conclude in Section 9 with thoughts on future research directions.

2 Empirical context and experimental design

Our study is conducted in a large online labor market (Horton, 2010; Agrawal, Horton, Lacetera and Lyons, 2015; Horton, Kerr and Stanton, 2017). In online labor markets, buyers form contracts with sellers to complete projects remotely. These projects include computer programming, graphic design, data entry, research, and writing. Each market differs in scope and focus, but platforms commonly provide ancillary services, including maintaining project listings, hosting buyer and seller profile pages, arbitrating disputes, certifying seller skills, and maintaining feedback systems (Filippas, Horton and Zeckhauser, 2020). They are broadly similar to a host of other online marketplaces that have arisen in recent years (Einav, Farronato and Levin, 2016). Buyers in this market have extensive sources of information about seller quality—particularly for sellers they are recruiting, as those sellers almost invariably have extensive on-platform experience.

Several features of conventional labor markets also exist in our context. Buyers and sellers are free to enter and exit the market anytime. Buyers post project descriptions, and after buyers and sellers match, they can negotiate over prices, which can either take the form of hourly rates or fixed prices for projects, and they form contracts. More generally, buyers and sellers face substantial search frictions (Horton, 2017, 2019), barriers to entry (Pallais, 2013; Stanton and Thomas, 2016), and information asymmetries (Benson, Sojourner and Umyarov, 2019; Filippas, Horton and Golden, 2022).

2.1 Search and matching in the market

The matching process can be initiated by either the sellers or by the buyers. Sellers can initiate the matching process by searching for and applying for projects. To do so, sellers can view an algorithmically determined ranking of all available projects of interest, access project descriptions and buyer profiles, and apply to projects by writing a cover letter and placing a wage bid. Applications use up "coins," an in-platform currency sold through the platform and costing $0.15 each (Filippas, Horton and Zeckhauser, 2023).

Buyers may also initiate the matching process by inquiring sellers whether they would be interested in applying to the buyer's project. Upon posting a project listing, buyers can view rankings of sellers who are determined algorithmically to be good matches for their projects. Horton (2017) shows that these algorithmic recommendations can be important to match formation. The buyer can further explore the sellers' profiles and, if interested, invite a seller to submit a proposal for the project. We call this buyer invitation to apply an "inquiry." For each project post, buyers may send a fixed number of free inquiries and can purchase the right to send additional inquiries. A seller application following an buyer inquiry uses up no coins. As such, buyer inquiries are valuable to sellers because they both reduce application costs, and because they signal buyer intent that might lead to a paid project.

2.2 Why do buyers not just send a large number of inquiries?

If sending inquiries helps buyers find sellers for their projects, why do buyers not simply send large numbers of inquiries? Of course, some do. But there are several reasons why buyers would economize *ex ante*, even if *ex post* they might wish they had recruited more broadly. One reason could be that the recommended sellers do not fit the buyer's requirements, resulting in only a few sellers worth recruiting. Another reason is that the buyer might decide that searching for additional sellers is not worth the cost. Time and resource constraints can also be a factor, as screening and reviewing candidates is time-consuming. Buyers may therefore wish to avoid being overwhelmed by many applications. Limiting the candidate pool to a manageable size can also help buyers evaluate and compare candidates effectively. Additionally, in a market where sellers can see how many inquiries a buyer sent—as in this one—buyers may be concerned that sending too many inquiries could discourage potential sellers, as it reduces that would-be seller's chances of landing the project, and so they might not apply.

2.3 Experiment

For the experiment, sellers in select technical categories became eligible to advertise. The advertisement simply displayed the phrase "Available Now" on the advertising seller's profile, as well as in all search tiles where the seller appeared. The advertising had no other effect. The price of advertising was fixed to 2 coins per week throughout the experiment. Sellers were notified of the advertising opportunity upon logging into the platform. A total of 243,126 sellers were engaged in the experiment, that is, they became eligible to advertise. Of the 243,126 sellers who were made eligible to advertise, 54,779 sellers (about 22.5%) advertised at least once during the experiment.

Buyers were randomized into a treatment and a control group upon posting a project in the same select technical categories. The only difference between the two groups was that treated buyers could see the seller advertisements. We reproduce the buyer search interface in Appendix A.1.

The experiment began on July 26, 2021 and ended on October 01, 2021. A total of 84,425 buyers were engaged in the experiment. Of that total, 42,474 buyers (50.31%) were allocated to the treatment group, and 41,951 (49.69%) to the control group. The experimental groups were well-balanced across several pre-experimental observables. In Appendix A.2, we report two-sided t-tests for various buyer-level attributes, as well the number of buyers allocated to the control and treatment cells over time. All the evidence is consistent with effective randomization.

Both the possibility of seeing advertising and the possibility of buyer advertising were launched simultaneously—though buyers were added to the experiment at the time of project posting. We observe a high advertising uptake, with the percentage of sellers advertising increasing rapidly at first and then leveling off. By the end of our data, 39.8% of the sellers who applied to at least one project during the experimental period were actively advertising.

Recall that advertising did not affect the sellers' placements in the buyers' search rankings. The rankings were determined algorithmically, and did not take advertising into account throughout the experiment. As such, it is possible that buyers could have little exposure to advertisers, depending on the interplay of search rankings and advertising decisions. However, this was not the case: advertisers made up about 50% of the sellers that buyers were exposed to. We provide more details in Appendix A.3.

The goal of the advertising feature was that, when implemented globally, it would cause buyers to shift their attention to advertising sellers. In a sense, the point of the feature is to "violate" the Stable Unit Treatment Value Assumption (SUTVA). However, this does complicate some interpretations of results. One concern of any field experiment in a single marketplace is that buyers and sellers potentially affect the outcomes of the other, violating the SUTVA. In our scenario, perhaps the treated buyers share information about the advertisement with the control group of buyers. We view this as unlikely as buyers tend not to communicate with each other and are not active in forums (unlike sellers, who often are). Or perhaps treated buyers—by sending more recruiting inquiries—could exhaust sellers' capacities. However, in this case, it would affect both treatment and control buyers. But even if this did occur, this is more of a case of generalizing the experimental results to an equilibrium setting rather than a kind of bias.

2.4 Buyers who could see the advertisements were more likely to contact advertising sellers

Having established that there was substantial seller uptake of advertising, we next turn to the question of whether advertising influenced buyer decision-making. In particular, we focus on whether buyers were more likely to send inquiries to advertising sellers. To do this, we compare buyer inquiries by search position, treatment status, and seller advertising status. We use seller impressions presented to buyers during the experimental period. An impression occurs when a buyer who searches sees a seller on his or her search interface.

Table 1 reports summary statistics on the impressions data, pooling all observations. There were about 3.4M impressions in total. A total of 75,622 unique buyers saw at least one impression. A total of 136,174 unique sellers received at least one impression.

Table 1: Summary statistics for impressions ($n = 3,427,112$)

Variable	Min	25th	Mean	Median	75th	Max	StDev
Position in search (1 = top)	1	8	72.231	25	70	2892	162.380
Seller advertising (AdsOn)	0	0	0.492	0	1	1	0.500
Buyer Inquiry	0	0	0.078	0	0	1	0.268
Seller accepts buyer inquiry	0	0	0.037	0	0	1	0.188
Contract formed	0	0	0.003	0	0	1	0.055

Notes: This table reports summary statistics for impressions of sellers presented to buyers. The reported outcomes are the impression position in the buyer search, whether the seller was advertising when the impression took place, whether the buyer sent an inquiry after the impression, whether the seller responded to the inquiry by applying for the project, and whether a contract was formed.

The number of seller tiles a buyer sees during a search session depends on how extensively the buyer searches. Some searches go quite deep, but most are fairly shallow, as indicated by the maximum "Position in search" value in Table 1. On average, sellers are advertising in about half of the tiles buyers see. Of all the impressions made, buyers make inquiries in about 8%, and about 4% of those inquiries receive a positive response from the seller, i.e., the response rate is about 50%. We can see that the conditional probability that the buyer forms a contract with a seller who has accepted the inquiry is a bit less than 10%.

We are interested in how the seller's position, advertising status, and the buyer's treatment status affect the probability that the buyer sends an inquiry. We begin by exploring this graphically in Figure 1. The x-axis is the seller's position in search (1 = top of page), and the y-axis is the fraction of those sellers that received an inquiry. As expected, buyer inquiry rates are strongly declining in search position. Although these are organic listings, the consumer search pattern of starting at the top and working down is evident, which is why sellers will pay to appear in these top positions in position auctions (Athey and Ellison, 2011). However, recall that advertising sellers were not given additional prominence in our experiment.

Figure 1a slices our data by the sellers' advertising status. In the left facet, we plot the inquiry rates of treatment and control buyers when they encounter non-advertising sellers. Treatment and control buyers respond similarly for every position in the left facet. While it may seem mechanical, as the sellers in this comparison are all non-advertisers, the lack of difference is not necessarily assured: if the treatment prompted buyers to switch from non-advertisers to advertisers generally, we would expect to see a drop in inquiry rates for treated buyers in this panel, since both types of sellers appear mixed together in search results. The absence of this decline provides visual evidence of little crowd-out in the aggregate, which we will substantiate later in the paper. In the right facet, we see the advertising substantially increased the probability that buyers sent inquiries to advertisers. Advertising appears to be as valuable as a higher position in search.

Figure 1b slices our data by the buyers' treatment status. We see that treated buyers have a significantly higher inquiry rate than control buyers. Importantly, comparing the two facets suggests that the ability to view advertisements does not decrease buyer inquiries to non advertising sellers, but rather it increase buyer inquiries to advertising sellers. However, the benefit of advertising makes up for this difference, suggesting that advertising can play a role in directing buyer attention to appropriate sellers.

3 Seller selection into advertising

When we compare the buyer inquiry rates by seller position and seller advertising status, it becomes apparent that buyers who can view advertisements are more interested in advertising sellers. This finding raises the question of why buyers are more interested in these sellers. To answer this question, we investigate the characteristics of sellers who select in to advertising, and explore the reasons why they may be more attractive to buyers.

The sellers who chose to advertise looked quite different from sellers who did not. Table 2 compares advertisers and non-advertisers based on their profile information, work history,

Figure 1: Seller search position and buyer inquiries

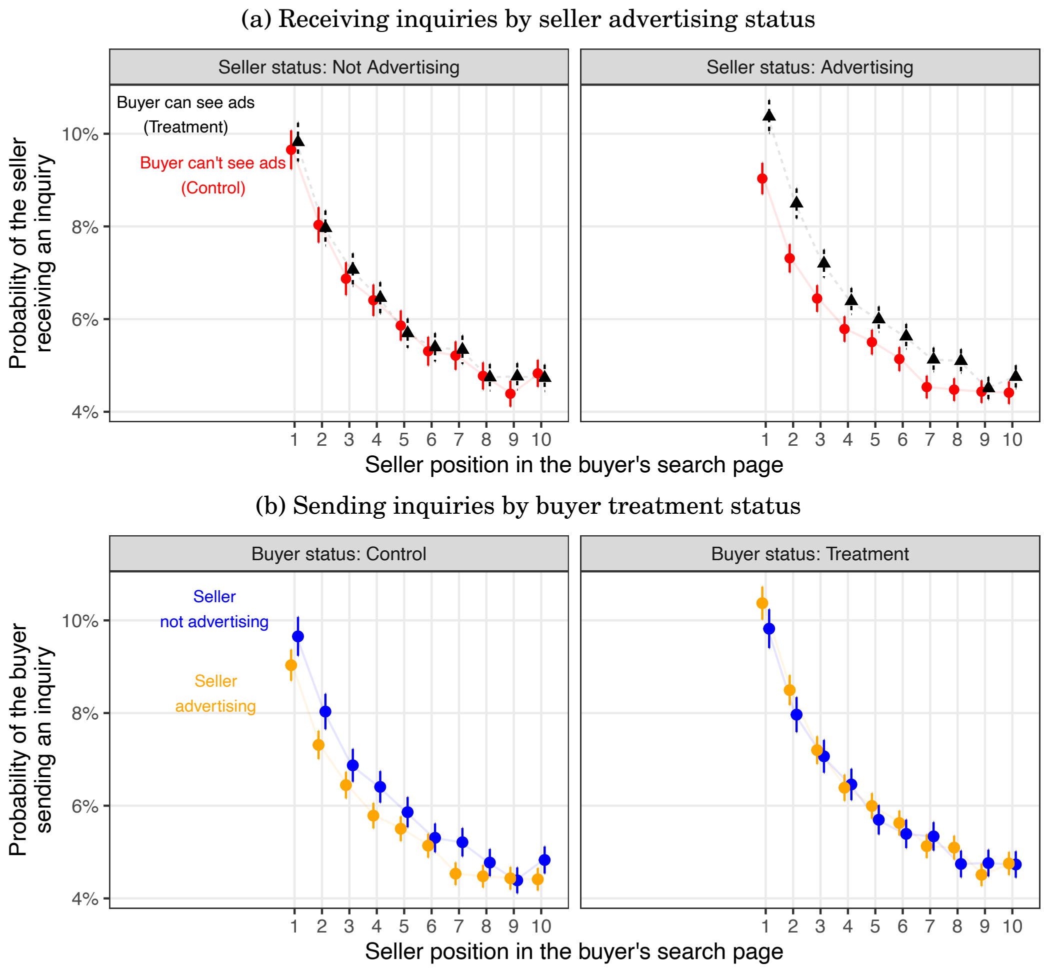

Notes: This figure plots estimates of the probability of a seller impression resulting in a buyer inquiry. The x-axis is the impression position in the buyer's search page (1 = top of page), and the y-axis is the probability that an impression led to a buyer inquiry. The left panel restricts the sample to non-advertising sellers and the right panel to advertising sellers. Estimates for control buyers who could not see the sellers' advertising status are depicted by red circles. Black triangles depict estimates for treated buyers who could see the sellers' advertising status. We report 95% confidence interval for each point estimate.

and pre-experimental and experimental outcomes. For the purposes of this exercise, we define "advertisers" to be the sellers who advertised for at least two full days during the experimental period. To focus on sellers who were active, we restrict our sample to sellers who placed at least one proposal during the experimental period.

Before the start of the experiment, advertising sellers, on average, had more complete platform profiles, had completed more projects successfully, and had received more positive buyer feedback for their past work. They also have slightly lower average hourly rates compared to non-advertisers. The attributes suggest advertising sellers were positively selected from the buyer's perspective.

Table 2: A comparison of the characteristics of advertising and non-advertising sellers

	Non-advertisers (mean)	Advertisers (mean)	Difference (percentage)	p-value
BEFORE THE EXPERIMENT (April 20, 2021 - July 20, 2021)				
Profile information				
stated availability	39.37	39.42	0.15%	0.025
percent completed	87.19	92.61	5.86%	0
contract success rate	0.37	0.47	20.29%	0
feedback score	45.96	53.59	14.22%	0
hourly rate	26.27	25.79	-1.87%	0.003
Seller market activity				
inquiries received	2.95	4.51	34.68%	0
acceptance rate	0.44	0.53	17.5%	0
bids placed	12.33	19.49	36.74%	0
contracts formed	0.3	0.52	42.4%	0
DURING THE EXPERIMENT (July 26, 2021 - October 01, 2021)				
Seller market activity				
inquiries received	4.25	7.23	41.15%	0
acceptance rate	0.47	0.55	13.83%	0
bids placed	18.96	35.76	46.98%	0
contracts formed	0.35	0.75	53.45%	0
Observation counts	46,845	39,088		

Notes: This table reports profile information and project application statistics before and during the experimental period for advertisers and non-advertisers. The sample comprises sellers who applied for at least one project during the experimental period. "Advertisers" are defined as sellers who advertise for at least 48 hours during the same period. For each outcome, we report the mean value for each group, the percentage difference for advertisers, and the p-value of a two-sided test of equal means between the two groups. For each seller, the reported profile attributes are self-reported availability, percentage of platform profile completion, mean buyer-reported contract success rate, and mean feedback score left by previous buyers; the reported contract outcomes are the number of inquires received, and, conditional on receiving at least one inquiry, the acceptance rate, the number of applications made, the number of contracts formed.

Interestingly, advertisers did not differ much from non-advertisers' self-stated capacities to take on more work. As we will show in Section 6, nearly all sellers on the platform report full capacity, and those made eligible to advertise were no exception. Before the start of the experiment, advertisers were already applying for and landing more projects and receiving more inquiries. Conditional on receiving at least one inquiry, their response rate, speed of response, and acceptance rate were also higher than those of non-advertisers. This finding is important because some naive rule, such as "give more impressions to sellers who do not seem busy," would likely have backfired badly, as those who are not busy are comparatively uninterested in more work.

During the experiment, the differences between the two groups remain largely the same in direction and magnitude, except for advertisers' larger bidding and contract formation rates. Advertising sellers continued to have substantially higher positive response rates to seller inquiries. These results suggest that the advertising option separated advertisers and non-advertisers regarding their "true" capacities to take on new projects.

We provide an alternative approach to modeling the sellers' selection into advertising in Appendix A.4. There, we quantify the relative importance of various factors that predict whether a seller will choose to advertise, by reporting the results of a logistic regression where the outcome is an indicator variable for selection into advertising, and the independent variables are pre-experiment seller attributes and outcomes. The positive predictors of advertising are all related to being highly active on the platform: many accepted buyer inquiries, more contracts, a higher acceptance rate of buyer inquiries, and more applications. Furthermore, sellers with higher successful contract completion rates are more likely to select in to advertising. Two factors that predict being less likely to advertise are a higher hourly rate and a greater number of buyers inquiries already received.[3]

4 Effects of seeing advertising: the buyer perspective

We have visual evidence that buyers tend to make inquiries to advertising sellers more often when they are in the treatment and can see advertising. We also have evidence of virtuous selection into advertising, particularly in terms of responding favorably to buyer inquiries. Before exploring buyer behavior and outcomes in more detail, it is useful to consider the economics of the buyer's recruiting problem.

Consider a buyer faces N potential sellers. Let $x_i = 1$ indicate that seller i accepts an inquiry and $x_i = 0$ if the inquiry is not accepted. We can think of these N as the sellers presented in a list of search results. The buyer can contract with a seller if they accept an inquiry. The buyer has utility $u(x_1, x_2, x_3, \ldots x_N)$, or $u(x)$, where $x \in \{0,1\}^N$. Let p_i be the probability that seller i accepts an inquiry if given one, and assume each seller's decision is independent of other sellers' deicions. Suppose they are rank-ordered by p so that p_1 has the highest probability. There are 2^N possible inquiry outcomes that are determined probabilistically. Let X be the collection of possible inquiry outcomes.

The buyer's problem is to decide which sellers to send inquiries to. Assume they all must be sent simultaneously, *ex ante*. This choice is represented by a binary vector r with length N.

[3]The relative price offered by advertisers has been a point of interest to those considering the economic effects of advertising. Schmalensee (1978) shows that firms with lower marginal costs—which might be associated with lower quality—might be most interested in advertising, making advertisers adversely selected. Of course, this would also cause buyers to learn to ignore advertising if being influenced by it lowered their utility. Consumers must believe paying attention to advertising is worthwhile for a separating equilibrium to exist. Outside our context, consumers generally seem to believe that costly advertising indicates product quality—but consumers also state that too much advertising could signal desperation (see Kirmani and Rao (2000) and references therein).

The decision problem is to choose an r that maximizes expected pay-off, subject to the cost of recruiting, $C(r)$, or

$$\arg\max_{r} \sum_{x \in X} u(x) Pr(x|r) - C(r), \tag{1}$$

where $Pr(x|r) = \prod_{i=1}^{N} (r_i p_i)^{x_i}$.

Now suppose one of the applicants, j, advertises, and the only effect is to make the buyer correctly believe that p_j is higher. Qualitatively, how advertising affects the seller's recruiting decision depends on the buyer's utility function u. For example, if all sellers were fungible, that is, $u(x) = |x|_\infty$, then the buyer only needs one seller to accept an inquiry. Then the optimal decision is to send inquiries to the top K sellers, with $1 - (1 - \prod_{i=1}^{K} p_i) > C'(K)$ but $1 - (1 - \prod_{i=1}^{K+1} p_i) < C'(K+1)$. In this scenario, advertising could cause an increase in inquiries sent if it makes seller j marginal, that is, if it "gives" them a p between $C'(K+1)$ and $C(K)$. However, advertising could also decrease the number of inquiries sent: if advertising conveyed that $p_j = 1$, then the buyer would only send one inquiry, which could be less than the counterfactual number in a market without advertising.

If the buyer plans to work with only one seller—as is often the case in our empirical context—then it would *seem* that candidates are substitutes. The reason is that advertising sellers now have higher p's, so recruiting them becomes more attractive, in turn making other sellers relatively less attractive (crowd out). On the other hand, the presence of multiple interested sellers may allow the buyer to obtain the Bertrand competition price rather than the monopolist price (Janssen and Rasmusen, 2002); or if there is some fixed cost to screening, the buyer might need a sufficiently high number of accepted inquiries in the pool. Advertising could then make advertising sellers complements to non-advertising sellers in recruiting.

The point of this decision-theoretic detour is that without more information about the recruiting, screening, and wage determination process, there is no reason to put more structure on the buyer's decision problem. As such, we simply compare outcomes across sellers.

4.1 Constructing a dataset of buyer outcomes and behaviors

We construct a dataset of buyers outcomes during the experimental period. The sample is the first project posted by a buyer. Table 3 summarizes buyer outcomes by treatment status. We report statistics for the treatment and control groups. Although we will report regression results, most of the main effects can be seen by comparing means. We can see that treated buyers sent more inquiries—especially to advertising sellers—received more bids, and, ultimately, formed more contracts. In addition to increases overall, there are also increases on the extensive margin for some of these outcomes. Treated buyers were more likely to send any inquiries at all and more likely to form at least one contract.

Table 3: Summary statistics for buyer outcomes, by experimental group

Variable	AdsVisible (trt)	Min	25th	Mean	Med	75th	Max	StDev
Buyer inquiries sent								
Inquiries sent	1	0	0	3.140	1	4	511	7.509
	0	0	0	3.052	1	4	806	7.268
Inquiries sent (any)	1	0	0	0.568	1	1	1	0.495
	0	0	0	0.559	1	1	1	0.497
Inquiries sent to advertisers	1	0	0	1.485	0	2	249	3.389
	0	0	0	1.397	0	2	277	3.023
Inquiries sent to advertisers (any)	1	0	0	0.440	0	1	1	0.496
	0	0	0	0.425	0	1	1	0.494
Responses								
Buyer inquiries accepted	1	0	0	1.450	0	2	208	2.980
	0	0	0	1.410	0	2	74	2.665
Buyer inquiries accepted (any)	1	0	0	0.470	0	1	1	0.499
	0	0	0	0.459	0	1	1	0.498
Bids								
Bids received	1	0	4	15.556	10	21	255	18.207
	0	0	4	15.503	10	21	198	17.825
Bids received (any)	1	0	1	0.946	1	1	1	0.226
	0	0	1	0.945	1	1	1	0.228
Contracts								
Contracts Formed	1	0	0	0.295	0	1	18	0.525
	0	0	0	0.288	0	1	20	0.537
Contracts Formed (any)	1	0	0	0.274	0	1	1	0.446
	0	0	0	0.265	0	1	1	0.441

Notes: This table reports buyer summary statistics by their treatment assignment (AdsVisible).

4.2 Buyer outcomes

As our outcomes are counts, we estimate a Poisson regression

$$\mathbb{E}[y_j \mid \text{ADSVISIBLE}_j] = \exp(\beta_0 + \beta_1 \text{ADSVISIBLE}_j),$$

where y_j is the buyer outcome of interest, ADVISIBILE_j indicates whether buyer j was assigned to the treatment group. With our binary treatment indicator, the β_1 coefficient is approximately the percentage change from control to treatment. Table 4 reports the results, using the first project posted by the buyer during the experimental period as the sample.

The outcome in Column (1) is the count of buyer inquiries sent. Treated buyers sent more inquiries overall, with an increase of about 3% more. In Column (2), the outcome is inquiries sent to advertising sellers. There is nearly a 6% increase in inquiries sent to advertising sellers. This reflects the pattern in Figure 1, borne out in buyer-level estimates.

Table 4: Effects of treatment on buyer outcomes, Poisson regression

	Dependent variable:				
	Buyer inquiries sent	Inquiries to advertisers	Inquiries to non-advertisers	Proposals from inquiries	Contracts formed
	(1)	(2)	(3)	(4)	(5)
ADSVISIBILE (trt)	0.028^{*}	0.061^{***}	0.0001	0.028^{**}	0.027^{***}
	(0.016)	(0.015)	(0.020)	(0.014)	(0.008)
Control	1.116^{***}	0.334^{***}	0.504^{***}	0.344^{***}	-1.246^{***}
	(0.012)	(0.011)	(0.014)	(0.009)	(0.006)
Observations	84,425	84,425	84,425	84,425	84,425

Notes: This table reports Poisson regression estimates where the depent variables are buyer outcomes, and the independent variable is an indicator for treatment status. The sample comprises each buyer's first project post during the experimental period. Robust standard erorrs are reported in parentheses. Significance indicators: $p \le 0.1 : \ddagger$, $p \le 0.05 : *$, $p \le 0.01 : **$, and $p \le .001 : ***$.

One might wonder if the treatment caused buyers to substitute away from non-advertising sellers. To test for this, in Column (3), the outcome is inquiries sent to *non* advertising sellers. Surprisingly, we get a precise zero, with no net decline in inquiries going to non-advertisers.

The net effect of this increase in buyer inquiries is that treated buyers had more sellers to choose from. In Column (4), the outcome is the count of sellers who responded positively to the inquiry. We can see that there was an increase. Interestingly, the magnitude of the effect is similar to what we saw in Column (1).

Did these additional proposals matter in forming contracts? In Column (5), the outcome is the number of contracts formed. We can see that treated buyers formed substantially more contracts, with an effect size of about 2.7% more.

The notion that advertising can expand total industry sales has limited prior empirical support, with own and rival advertising tending to cancel out (Lambin, 1976). In the sponsored search context, the evidence ranges from positive (Sahni and Nair, 2019), to neutral (Abhishek, Jerath and Sharma, 2022), or even negative (Moshary, 2021). Blake, Nosko and Tadelis (2015) found that brand keyword ads took traffic from their own organic listings, but Golden and Horton (2021) found no evidence of advertising "business stealing" from even close rivals.

4.3 Model-free evidence on the shift towards advertisers

Table 4 showed clear evidence that treated buyers sent more inquiries to advertising sellers but no clear evidence of a shift away from non-advertisers. This is surprising, as we might expect that buyers who send few inquiries may shift their attention to advertisers, at the expense of non-advertisers. To build confidence in this result, we take a model-free, graphical approach.

Figure 2 shows the distribution of buyer inquiries by treatment status and seller advertising status. The plot shows the empirical cumulative distribution functions for inquires sent to advertising sellers (left panel) and non-advertising sellers (right panel). We can see clear evidence that the treatment worked to increase inquiries to advertising sellers. However, there is no discernible shift for non-advertising sellers.

Figure 2: Buyer inquiry quantiles, by seller advertising decision and buyer treatment status

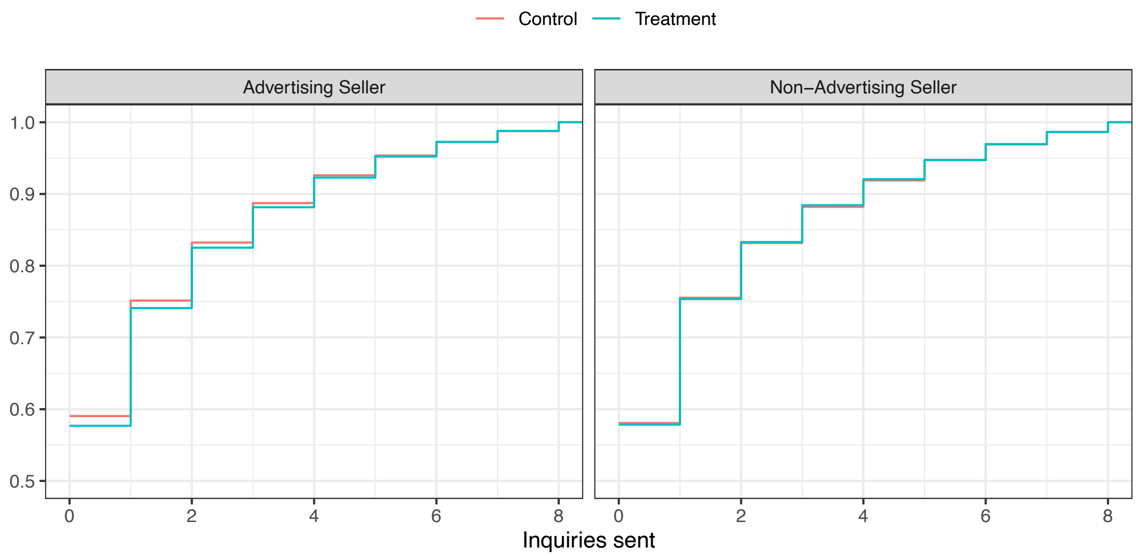

Notes: This figure shows the empirical cumulative distribution functions of the inquiries sent by buyers to advertisers and non-advertisers. In each panel, we split the data by the buyer's treatment status.

4.4 Discussion of buyer results

It is possible that the increase in inquiries for seller profiles that added a badge was due to a novelty effect. However, this explanation is unlikely, given two key pieces of evidence. First, in the aggregate, sellers did not turn off the badge over time, suggesting that the badge continued to be effective in attracting inquiries. Second, when we expanded the sample to include subsequent projects by buyers, the treatment effect of the badge got stronger. In Appendix A.5, we expand the analysis to include subsequent project openings posted by buyers. This suggests that the increase in inquiries was not just a temporary novelty effect but instead had a lasting impact on buyer behavior.

We do not know what buyers believed when they encountered the advertising badge. Our analysis is limited to examining the resulting behaviors after the badge was seen or not seen. It is possible that some buyers assumed the badge was platform-inferred, others believed it was paid advertisement, and some ignored it altogether. Buyers interested in how the badge was determined could easily discover this information. Whatever beliefs buyers started with, presumably, there was learning about the true nature of the badge over time. The fact that the effects of the badge got stronger over time (Appendix A.5) suggests the effects of the badge were not transitory.

5 Effects of advertising: the seller's perspective

Although randomization took place at the level of the buyer, we can also examine the effects of advertising from the sellers' perspective. The seller's advertising decision depends on the revenue increase minus the cost. Although we observe the cost of advertising, we do not observe the seller's costs of actually performing the work, so we do not know their margins. However, we can characterize what margins would make advertising rational.

5.1 Advertising sellers are more likely to receive a buyer inquiry and to form a contract when an impression is with a treated buyer

We first examine the effects of advertising for sellers at the impression level. We construct a sample where each observation is a seller impression in a buyer's search interface during the experimental period. Our estimation strategy is to regress each outcome of interest on indicators for the treatment status of the buyer who posted the project and viewed the impression, the advertising status of the seller, and the interaction of the two terms. Our specification is:

$$y_{ip} = \beta_0 + \beta_1 \text{ADSVISIBILE}_p + \beta_2 \text{ADSON}_{ip} + \beta_3(\text{ADSVISIBLE}_p \times \text{ADSON}_{ip}) + \epsilon_{ip},$$

where p is a project post, i is a seller, y_{ip} is an outcome for the buyer-seller interaction for project p, ADSVISIBLE_p indicates whether buyer who posted the project p was assigned to the treatment group, ADSON_{ip} indicates whether the seller was advertising during that interaction, and and ϵ_{ip} is an error term. Crucially, because advertising is randomly visible due to our randomized assignment, the estimates for the coefficients β_1 and β_3 can be interpreted causally. The estimate for β_1 measures the additional effects of a buyer being able to see ads, even when a particular seller is not advertising. We expect this estimate to be equal to zero unless there are spillovers from advertising: a buyer who sees ads can become more or less active even with a seller who is not advertising. The estimate for β_3 measures the effect of advertising visibility on the outcomes of *advertisers* who appeared in the treated buyers' search interfaces.

We report the estimated effects in Table 5. The coefficient $\hat{\beta}_1$ on ADSVISIBLE is close to zero for all dependent variables. This suggests no strong spillovers or crowd-out of non-advertising sellers by advertising sellers: For buyers in the treatment, encountering a seller without ADSON did not make it less likely for them to contact that seller. This is consistent with what we learned from our buyer-focused analysis showing no reduction in inquiries sent to non-advertisers.

Table 5: The effects of advertising for sellers at the impression level.

	Dependent variable:		
	Buyer sent inquiry	Seller applied after inquiry	Contract formed
	(1)	(2)	(3)
CONSTANT	0.080***	0.032***	0.003***
	(0.001)	(0.0005)	(0.0001)
ADSVISIBLE	0.003‡	−0.00001	−0.0001
	(0.002)	(0.001)	(0.0001)
ADSON	−0.010***	0.009***	0.0005***
	(0.001)	(0.001)	(0.0001)
ADSON×ADSVISIBLE	0.004**	0.003***	0.0002‡
	(0.001)	(0.001)	(0.0001)
Observations	3,427,112	3,427,112	3,427,112
R^2	0.0003	0.001	0.00003

Notes: This table reports regressions where the independent variables are each buyers's treatment status, an indicator for whether a seller viewed in search was renting the badge, and an interaction term between the two. The dependent variables are indicators for whether (1) the buyer sent an inquiry to the seller, (2) the seller applied for the project after receiving the inquiry, and (3) whether the buyer and the seller formed a contract. Observations are on the project post/buyer impression level, and we cluster standard errors on the project post and seller level. Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Viewing the advertisement causally increased the probability of an advertiser receiving an inquiry substantially. We can see this in the coefficient on ($\text{ADSVISIBLE}_p \times \text{ADSON}_{ip}$)). The estimated effect is about 4.83% higher than the estimated base probability of receiving an inquiry.

In Column (2), we see advertising sellers were more likely to respond positively following a buyer inquiry, that is, to apply for the project. Moving further downstream in Column (3), we can see they were more likely to form a contract. Although the effect is somewhat imprecise, the advertiser was about 7.72% more likely to form a contract with the buyer following the impression. Given this, the effect of advertising for a given seller is mechanically heterogeneous in levels but could be proportional.

5.2 Total effects of advertising from the seller perspective

In deciding whether or not to advertise, what matters for sellers is not the per-impression effect of advertising but the net effect over some period of time. We now switch our analysis to the seller-period level rather than the seller-impression level. We construct a two-period panel: before and after advertising was introduced to the market. We will use that panel for a difference-in-differences analysis of the effect of advertising on seller outcomes.

Table 6 reports summary statistics for the panel. Note that rather than simply having an advertising on/off indicator as an independent variable, we have advertising days as our key dependent variable. This is useful as advertisers can decide how many days to advertise.

Table 6: Summary statistics for seller difference-in-differences panel

Variable	Min	25th	Mean	Median	75th	Max	StDev
Advertising adopter	0.00	0.00	0.45	0	1.00	1.00	0.50
Contracts formed	0.00	0.00	0.47	0	0.00	139.00	1.32
Amount of advertising (days)	0.00	0.00	19.99	0	43.99	66.08	26.34
Buyer inquiries received	0.00	0.00	4.63	1	4.00	805.00	13.07

Notes: This table reports summary statistics for seller impressions in the buyer search interface.

We look at the number of buyer inquiries and contracts in each period. We estimate the following specification:

$$y_{it} = \alpha_i + \beta_0 \textsc{ExpPeriod}_t + \beta_1(\textsc{AdDays}_i \times \textsc{ExpPeriod}_t) + \epsilon_i,$$

where α_i is a seller-specific fixed effect, ExpPeriod is an indicator for the post-period, and AdDays measures the number of days the seller advertises in the post-period. We cluster standard errors at the individual seller level. Table 7 reports the results.

One complication is that sellers differ radically in how many impressions they receive, despite the advertising having a fixed cost not commensurate with the number of impressions. This is because of the algorithmic nature of the search rankings, which attempts to use historical data on buyer interest to rank sellers. As we would expect the effects of advertising to be proportional to impressions, this would tend to make effects necessarily multiplicative rather than having some constant effect in levels. This would argue for a log or Poisson transformation of the outcome, but because there are many zeros, the inclusion of seller-specific fixed effects means the estimation would necessarily have to drop observations. But we also expect many zero-to-one effects making this dropping unwise, so simply using a linear model seems like the best option.

Our primary interest is the interaction between the period indicator and the number of advertising days in the table's second row. In Column (1), we can see that each day of advertising lead to approximately 0.04 more inquiries per advertising day.

The coefficient on ExpPeriod captures this mean of the difference between the two periods. The experimental period lasted 66 days, and the pre-period lasted 91 days, but our sample is made up of sellers who applied for at least one project during the experimental period; as looking for new projects is episodic, we expect the pre-period to be characterized by less activity

Table 7: The effects of advertising for sellers at the seller level

	Dependent variable:			
	Buyer inquiries received	Log rate bid	Contracts formed	
	(1)	(2)	(3)	(4)
ExpPeriod	1.154***	0.038***	0.050***	−0.036***
	(0.030)	(0.002)	(0.005)	(0.006)
Log proposals made				0.456***
				(0.006)
ExpPeriod×AdDays	0.040***	0.00001	0.004***	0.003***
	(0.002)	(0.0001)	(0.0002)	(0.0002)
Seller fixed effects	✓	✓	✓	✓
Observations	171,866	119,052	171,866	151,024
R^2	0.872	0.951	0.749	0.781

Notes: This table reports regressions where the dependent variables are an indicator for the experimental period, and the same variable interacted with the number of days each seller advertised during the experimental period. The dependent variables are indicators for whether (i) the number of inquiries the seller received from buyers, (ii) whether the seller received any inquiries from buyers, (iii) the number of project contracts the seller formed, and (iv) whether the seller formed any contracts. We include a seller fixed effect, and cluster standard errors on the seller level. Significance indicators: $p \le 0.1 : \ddagger$, $p \le 0.05 : *$, $p \le 0.01 : **$, and $p \le .001 : ***$.

compared to the experimental period.

In Column (2), the outcome is the average log wage bid. There is no evidence that advertising sellers increased or lowered their listed rates when advertising. Note, however, that this is a within-seller effect and that advertisers had lower rates in levels. As such, not finding a price reduction is not evidence against Benham (1972) or the arguments made in Bagwell and Ramey (1994).

Do the extra inquiries from buyers result in more contracts for sellers? In Column (3), the outcome is the number of contracts formed. We find that advertisers form around 0.004 more contracts per day compared to non-advertisers. However, the estimate in Column (3) may be too high, as sellers might have also increased their project-finding efforts by applying to more projects. This channel could be strongly linked to the decision to advertise. Although our seller fixed effect helps control for this, sellers could still choose to advertise and increase their search intensity simultaneously.

To deal with this issue, in Column (4), we include a control for the number of proposals sent during the period to proxy for search intensity. This is technically an improper regression—we are controlling for a downstream outcome—but it is likely still informative about whether advertising was really driving increased business for sellers. Including this reduces the treatment effect by about 1/4, suggesting that the buyer inquiry channel to invitations was important but somewhat overstated by our difference-in-differences approach.

In terms of the returns to advertising, the results imply that if a buyer inquiry is worth at least a $1 and a contract at least $10, then advertising has a positive return on investment at the 2 coins per week price (see Section 2 for more details on coins).

5.3 Spillovers effects from other advertising sellers within the consideration set

Sellers appear in buyer search along with other sellers, typically in groups of 10, because of pagination in the interface. From the perspective of a given seller—call them the focal seller—there is variation in the number of advertising sellers that appear on each page. Are the number of other advertisers on the page good or bad from the focal seller's perspective? Or put differently, do other advertisers "crowd out" and "crowd in" buyer attention on the social seller?

The crowd-out effect is that other advertising sellers may draw attention away from the focal seller, removing the potency of his or her ad. The crowd-in effect is that advertising by other sellers in the consideration set may draw more attention to the focal seller—perhaps mitigating any concern that advertising sellers are "desperate." Note that this kind of informational spill-over is different from advertising by one brand potentially benefits some other brand in the same industry, as in Shapiro (2018), which is presumably more due to awareness.

For each buyer-seller impression, we can regress whether the buyer made an inquiry to the seller on a seller-specific fixed effect, a position-specific fixed effect, and all the interactions of the buyer's treatment status, the seller's advertising status, and the fraction of *other* sellers in the search results that were also advertising. As these kinds of triple difference regressions are difficult to interpret, we plot the implied effects in Figure 3.

$$
\begin{aligned}
y_{jp} = \alpha_p + \textsc{Position}_p + \beta_1 \textsc{AdsVisibile}_p + \beta_2 \textsc{AdsOn}_{ip} + \beta_4 \textsc{FracAdvertising} + {} \\
\beta_3(\textsc{AdsVisible}_p \times \textsc{AdsOn}_{ip}) + {} \\
\beta_5(\textsc{AdsVisibile}_p \times \textsc{FracAdvertising}) + {} \\
\beta_6(\textsc{AdsOn}_{ip} \times \textsc{FracAdvertising}) + {} \\
\beta_7(\textsc{AdsOn}_{ip} \times \textsc{FracAdvertising} \times \textsc{AdsVisible}) + {} \\
\epsilon_{ip},
\end{aligned}
$$

The left facet shows effects when buyers cannot see advertising because the buyer is in the control group. We can see that when $\textsc{FracAdvertising} = 0\%$, an advertising focal seller is considerably less likely to receive an inquiry. An advertising seller with few other advertising sellers is perceived as "bad" from the control buyer's standpoint, even though the buyer does

Figure 3: The effect of seller advertising on receiving buyer inquiries, conditional on the number of other advertisers in search results.

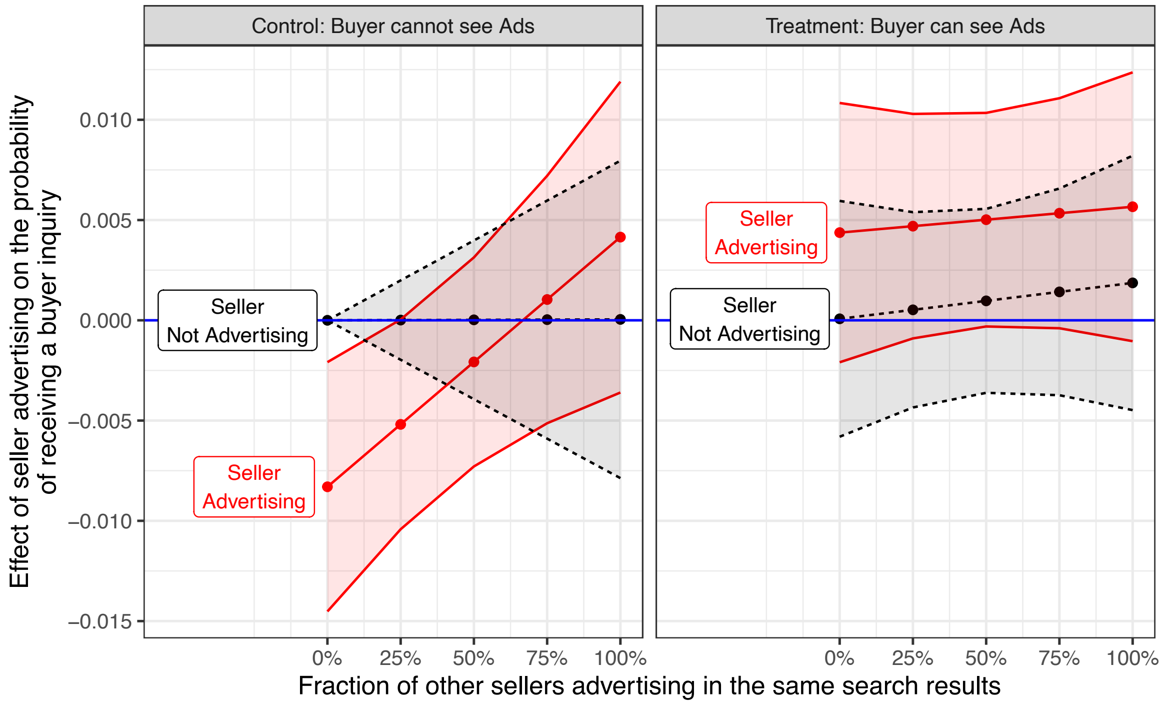

Notes: This figure reports the effects of seller advertising on buyer inquiries, conditional upon how many other sellers in search results were also advertising at the same time and conditional on whether the buyer could see advertising.

not know the focal seller is advertising. However, as the fraction increases, this lower inquiry rate for advertising sellers goes away, as we would expect: At FRACADVERTISING = 100%, the focal advertising seller is not selected, as every other seller is also advertising.

In the right facet, results are shown when buyers are in the treatment and can see advertising. Advertising sellers now enjoy greater buyer interest than non-advertising sellers. Interestingly, this benefit does not diminish much as the number of advertising sellers in the buyer's consideration set increases. Furthermore, there is no evidence that non-advertising sellers are crowded out: Their inquiry rate does not decline with a higher fraction of advertising sellers.

One might think that at FRACADVERTISING = 100%, a non-advertising seller would be strongly disadvantaged generally. One might also think that a seller advertising when FRACADVERTISING = 100% would get no relative benefit. That does not seem to be the case. This is further evidence that advertising mainly increased the number of buyer inquiries rather than crowded-out inquiries to non-advertisers.

6 How paid advertising can increase efficiency

A key result of the experiment is that advertising shifted buyer attention to more available sellers. These sellers were more likely to transact, leading to an overall increase in matches formed. In this section, we examine further the conditions under which paid advertising can increase market efficiency. We first show empirically that costless advertising fails to increase market efficiency. We then develop a simple model that helps us to examine under which conditions costly advertising can increase market efficiency.

6.1 The near uselessness of free advertising

Buyer inquiries are common on the platform: more than half of the buyers send at least one inquiry after they post a project. Because inquiries are limited and searching for sellers is costly, buyers want to send inquiries to sellers who are likely to accept them. As such, we would expect buyers to try to infer seller capacity to take on new projects.

In the earlier days of the platform, buyers had little to work with. But later on, the platform introduced a signaling mechanism that allowed sellers to declare their "availability"—their capacity to take on new projects—on their profiles. The availability signaling feature allowed sellers to put one of three messages on their profiles about their availability: (1) "Less than 30 hrs/week," (2) "More than 30 hrs/week," and (3) "As Needed - Open to Offers." Sellers could change their availability at any point in time. Horton (2019) analyzes the impact of this feature and finds that it was effective, in that sellers signaling high capacity received more buyer inquiries, rejected fewer inquiries, and were more likely to form a contract.

Although the self-reported seller availability feature appeared promising initially, its effectiveness deteriorated. Figure 4 shows how sellers used that feature before the experiment. Panel (a) reports the distribution of sellers' choices using a cross-section that spans half a year before the commencement of the experiment. Only about 4.2% of the sellers reported limited capacities. Instead, sellers overwhelmingly reported that they had high capacities: 88.6% reported they were available full-time, and 7.3% reported they were available as needed.[4]

Panel (b) in Figure 4 shows the inquiry acceptance rates conditional on seller availability status. Only about 36.2% of buyer inquiries were accepted by the sellers. Somewhat surprisingly, sellers who were ambiguous about their availability status had very similar acceptance rates (37.8%) to sellers who indicated that they were highly available (37.9%). Furthermore, sellers indicating that they were highly available received about 52.1% of buyer inquiries and the rest of the inquiries reached sellers who signaled other availability levels.

Figure 5 plots the percentage of sellers who changed their self-reported availability status each month, using data spanning more than two years before the experiment. For each month,

[4] Horton (2019) finds that 22% of sellers choose "As needed," 33% choose "30 hours or less," and 45% choose "open to offers."

Figure 4: Statistics on sellers' self-reported availability

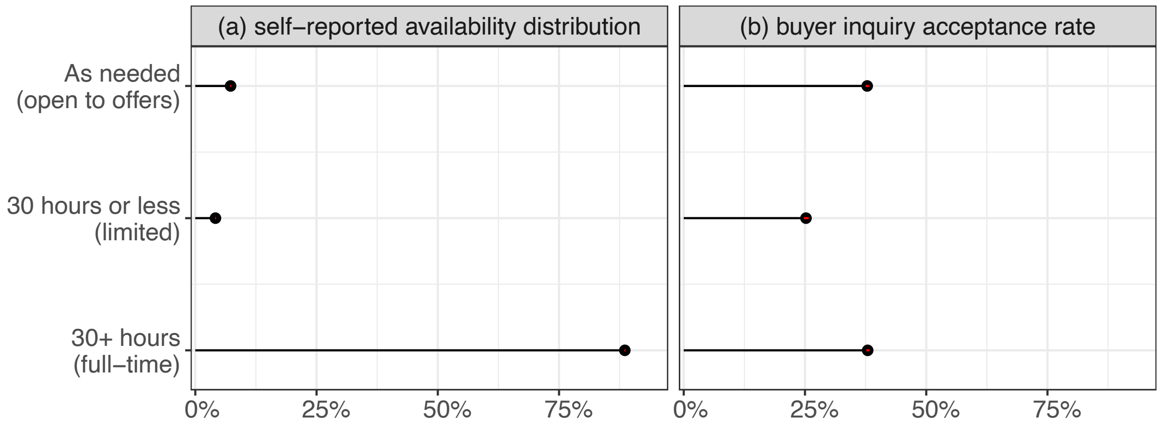

Notes: This figure reports summary statistics for the sellers' self-reported availability. Panel (a) reports point estimates of the distribution of the seller availability choices. Panel (b) reports point estimates of the probability that sellers accept a seller inquiry, conditional on their self-reported availability. A 95% confidence interval is plotted for each point estimate—the confidence intervals are not visible in this plot because they are narrow.

the sample of sellers is restricted to active sellers, defined as those sellers who applied for at least one project during that or the previous month. Sellers did not change their availability status often: the percentage of sellers changing their availability status decreased over time from about 4.5% to about 0.6%.

Figure 5: Sellers changing their self-reported availability over time

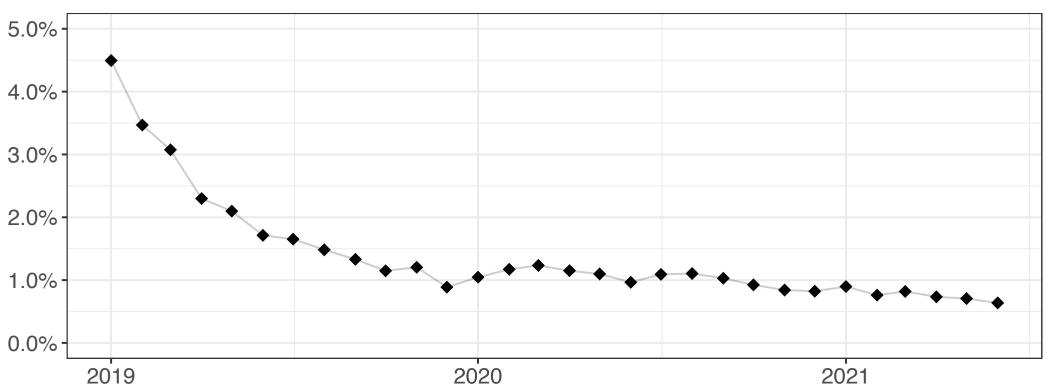

Notes: This figure plots the percentage of active sellers who changed their self-report availability at least once each month. The period covers about 2.5 years before the commencement of the experiment. The sample comprises active sellers, defined as those sellers who applied for at least one project during the current or the previous month.

Taken together, the evidence above suggests that (i) sellers were overstating their capacity to take on new projects, (ii) buyers were aware of this misreporting and responded by "mixing" their inquiries to sellers with lower self-reported availability statuses as well, and (iii) the usefulness of the feature likely decreased over time. In other words, simply allowing sellers to advertise "for free" did not work. It is worth noting that even if the platform attempted to incentivize the sellers to keep their self-reported availability statuses up to date—e.g., by de-prioritizing sellers who haven't changed their statuses recently in the buyers' search rankings—then sellers would likely have continued overstating their capacities.

6.2 Why do prices not clear the market?

In our setting, what differs among sellers is not some fixed attribute (as quality does in Nelson (1974) and Milgrom and Roberts (1986)), but rather their capacity to take on more work. A reasonable question is then why, in a highly competitive market, does this variation in capacity exist. Or put differently, why do prices alone not clear the market? For example, one could imagine slack sellers posting lower prices to get more work, and constrained sellers posting higher prices. The fact that they do not—or cannot do so sufficiently to make advertising unnecessary—-suggests that matching frictions matter in explaining the market.

In a market with frictions, it would seem that advertising could help direct search to more available sellers. However, if even "busy" sellers value buyer inquiries, advertising would have to be costly enough to deter them from advertising, but not so costly that "available" sellers also choose not to advertise. Furthermore, the benefit of advertising is endogenous, as it depends on how buyers change their search behavior, which in turn depends on which sellers select into advertising. These equilibrium considerations are sufficiently complex that even if our empirical results suggest advertising works to coordinate the market in practice, a natural question is whether it works in theory. Furthermore, the notion of "works" is under-specified, as it is really a question of welfare, which is hard to answer without a formal model.

We construct a model of a large matching market with endogenous seller capacity. Our main modeling goals are to understand when it is possible to have an equilibrium in which such advertisements are credible and what are the efficiency implications of such an equilibrium. Critically, we do not assume that sellers differ in capacity but rather micro-found it with a matching process that leaves some sellers with slack capacity and buyers unable to direct their search to these sellers.

6.3 Model

A unit mass of buyers and a unit mass of sellers interact in a matching market. Sellers produce and sell a homogenous good that buyers value at $v > 0$. At each point, a seller is characterized by her state $s \in \{a, b\}$, where a means the seller is "available" and b means she is "busy." The value of the state is private information and affects the seller's output cost.

Buyers and sellers potentially meet in several markets. If the seller can not signal her availability, only one market exists. But if a successful signaling technology is available, there could be separate markets for busy and available sellers. If a market has x buyers and y

sellers, then $m(x,y)$ "meetings" occur. We make the following standard assumptions about the matching function m:

Assumption 1. The matching function $m(x,y)$ is continuously differentiable, quasiconcave, increasing, homogenous of degree 1 (constant returns to scale), with $m(x,0) = m(0,y) = 0$ and $m(1,1) \le 1$.

When a seller in state s and a buyer meet, the seller draws a cost $C_s \ge 0$ for completing the project from a distribution with cdf F_s. If the value to the buyer exceeds the cost to the seller, $C_s \le v$, then a match takes place and produces a surplus $v - C_s$. A fraction $\alpha \in (0,1)$ of the surplus goes to the seller, and the rest to the buyer. We assume that a busy seller is more likely to have a higher cost of completing the project:

Assumption 2 (Stochastic dominance)**.** For any cost $c \in \mathbb{R}_+$, $F_a(c) = \Pr\left(C_a \le c\right) \ge \Pr\left(C_b \le c\right) = F_b(c)$, where the inequality is strict for $c = v$.

The matching process affects how sellers transition between the two states. Time is discrete. A seller who is busy in period t and gets matched in period t continues being busy in period $t+1$; otherwise, the seller becomes available in period $t+1$. A seller available at t who gets matched in that period becomes busy in $t+1$; otherwise, the seller remains available in $t+1$. A new unit mass of myopic buyers enters the market at every period.

6.4 Payoffs and the law of motion

Payoffs. A buyer and a seller in state s who meet each other match if $C_s \le v$. Let p_s be the probability of a match conditional on a meeting. Clearly, $p_s = F_s(v)$, with $p_a > p_b$. Let w_s be the expected value of the surplus conditional on the meeting, equal to $w_s = \mathbb{E}\left[\max\{v - C_s, 0\}\right]$. We will see later that all objects in our model depend on the distributions F_s only through the values of p_s and w_s. For example, the expected payoff to the seller of type s from a meeting is αw_s, and the expected payoff to the buyer is $(1-\alpha)w_s$. The following lemma demonstrates that we can treat the values (p_a, p_b, w_a, w_b) as model fundamentals, with the only restrictions on them being that $p_a > p_b$ and $w_a/w_b \in [1, +\infty)$.

Lemma 1. *For any ratio $w_a/w_b \in [1, +\infty)$ and probabilities p_a, p_b satisfying $p_a > p_b$, there exist a pair of distributions F_a and F_b such that $F_a(v) = p_a$, $F_b(v) = p_b$, and F_b dominates F_a in the sense of Assumption 2.*

Proof. This proof and all other omitted proofs can be found in Appendix B. □

The law of motion. Let $A_t \ge 0$ and $B_t \ge 0$ be the measures of all available and busy sellers in period t, with $A_t + B_t = 1$. Suppose that all buyers and sellers meet in one market. Then the number of matches M_t in period t is given by $M_t = m(1,1)B_t p_b + m(1,1)A_t p_a$ because, with random matching, we expect the fraction of meetings with type s to be proportional to their mass. As the number of matches is equal to the number of sellers getting matched, and all matched sellers become busy in the next period, the law of motion for the mass of busy sellers becomes

$$B_{t+1} = M_t = m(1,1)B_t p_b + m(1,1)A_t p_a. \tag{2}$$

Now suppose there are two distinct markets for available sellers and busy sellers. In that case, we also need to distinguish between the buyers shopping in the market for available sellers, with a mass of R_t^a, and those buyers shopping for the busy sellers, with a mass of R_t^b. The total number of matches formed is then $M_t = m(R_t^a, A_t)p_a + m(R_t^b, B_t)p_b$, and the law of motion for the number of busy sellers is

$$B_{t+1} = m(R_t^a, A_t)p_a + m(R_t^b, B_t)p_b. \tag{3}$$

6.5 Pooling ("no advertising") equilibrium

We study the existence of an equilibrium in which sellers cannot credibly signal their availability. We call this the pooling equilibrium. We restrict attention to the economy's steady state, in which the number of busy and available sellers remains fixed over time. With only one market, the agents do not make any choices, making the characterization of the pooling equilibrium straightforward.

Definition 1. A *stationary pooling equilibrium* is a collection $(R^a, R^b, A, B) \in \mathbb{R}_+^4$ such that (1) $A + B = 1$, (2) $R^b = B$, and (3) $B = m(R^a, A)p_a + m(R^b, B)p_b$.

The definition of the stationary pooling equilibrium is written as if there are two separate markets, with A sellers and R^a buyers shopping in the market for available sellers, and B sellers and R^b buyers shopping in the market for busy sellers. This choice will become convenient when we study the advertising equilibrium in which there will be two separate markets. In deriving the law of motion (2) for the economy with one market, we noted that the total number of meetings $m(1,1)$ would be split between available and busy sellers with weights A and B, respectively. Constant returns to scale implies $Am(1,1) = m(A,A)$ and $Bm(1,1) = m(B,B)$, meaning that we can think of the matching process with one market as if it is actually taking place in two separate markets, where the number of buyers and the sellers are equal to each other. We now establish the existence and uniqueness of the stationary pooling equilibrium.

Proposition 1. *There exists a unique stationary pooling equilibrium.*

6.6 Advertising equilibrium

We now study the existence of an "advertising" equilibrium, in which the two types of sellers are able to credibly signal their availability through costly advertising. In such an equilibrium, buyers and sellers must decide which market they want to transact in. While sellers are free to sell in any market, we will construct an equilibrium in which they choose the market that matches their state.

If a seller of type s enters the available market, the probability that she meets a buyer is $m(R^a,A)/A$—this is the total number of meetings divided by the number of sellers in that market. The expected surplus conditional on a meeting is w_s, which is determined by the seller's type and not by the market she is in. The seller receives a fraction α of the surplus. Finally, the seller must pay π to advertise her availability. Therefore, the payoff to selling in the "available" market is $U_s(a) = \alpha w_s m(R^a,A)/A - \pi$, and the payoff from transacting in the "busy" market is $U_s(b) = \alpha w_s m(R^b,B)/B$. Note that, as long as $m(R^b,B)$ is positive, sellers are guaranteed a positive payoff from participating in the busy market, meaning that the "participation constraint" is slack.

For both available and busy sellers to reveal their types, the following two incentive compatibility conditions should hold:

$$\alpha w_a m(R^a,A)/A - \pi \geq \alpha w_a m(R^b,B)/B \tag{4}$$

$$\alpha w_b m(R^b,B)/B \geq \alpha w_b m(R^a,A)/A - \pi \tag{5}$$

These constraints can be equivalently rewritten as two constraints on the price that can support the advertising equilibrium

$$\alpha w_a \left(\frac{m(R^a,A)}{A} - \frac{m(R^b,B)}{B} \right) \geq \pi \geq \alpha w_b \left(\frac{m(R^a,A)}{A} - \frac{m(R^b,B)}{B} \right). \tag{6}$$

If the allocation (R^a,R^b,A,B) admits prices satisfying (6), then that allocation is incentive-compatible for the sellers. Notice that the non-emptiness of the set of advertising prices (6) is equivalent to the condition

$$\frac{m(R^a,A)}{A} \geq \frac{m(R^b,B)}{B} \tag{7}$$

We summarize this in Proposition 2:

Proposition 2. *For an advertising equilibrium to exist, a necessary condition is that sellers have to enjoy higher meeting rates when they advertise.*

We are now ready to define an advertising equilibrium and prove its existence.

Definition 2. A *stationary advertising equilibrium* is a collection $(R^a, R^b, A, B, \pi) \in \mathbb{R}_+^5$ such that (i) $A + B = 1$, (ii) $m(R^a, A)w_a/R^a = m(R^b, B)w_b/R^b$, (iii) $B = m(R^a, A)p_a + m(R^b, B)p_b$, and (iv) π satisfies condition (6).

In the definition above, condition (ii) is the indifference condition for the buyers. Since all buyers are identical, they should be indifferent between the two markets in equilibrium. Interestingly, since $w_a \geq w_b$, this implies that buyers must have a *lower* meeting rate in the "available" market, but they are compensated by a higher matching rate, conditional upon a meeting.

Proposition 3. *A stationary advertising equilibrium exists and in such equilibrium $B \geq R^b$, where the inequality is strict if $w_a > w_b$.*

6.7 Welfare properties

We have characterized the two types of equilibria that a market can have—one where sellers can advertise their availability and one where they cannot. Next, we study the welfare properties of the two equilibria.

Social welfare at a point (R^b, B) that lies on the frontier of stationary allocations $B = p_a m(1 - R^b, 1 - B) + p_b m(R^b, B)$ is

$$W(R^b, B) = w_a m(1 - R^b, 1 - B) + w_b m(R^b, B) \tag{8}$$

A stationary allocation is *constrained-efficient* if it maximizes Equation 8 subject to the stationarity constraint.

It turns out that characterizing the welfare-maximizing allocation and comparing it to the two types of equilibria that we study is challenging at the level of generality that we have maintained so far. Meetings in the available market produce higher surplus w_a but are, in a sense, more expensive because they are more likely to lead to a match ($p_a > p_b$) and exhaust the endogenous "budget" of busy sellers B the economy can sustain in a stationary environment. Our model places few restrictions on (p_a, p_b, w_a, w_b) (see Lemma 1) and the form of the meeting function m, and, as a result, no strong welfare statements can be obtained. Proposition 4 below presents a welfare claim that holds under our general conditions, while Proposition 5 demonstrates that stronger results require additional assumptions on the model. We then show how additional structure placed on the model can strengthen the welfare results.

Proposition 4. *At the point of the stationary allocations frontier that defines the pooling equilibrium ($B(R^b) = R^b$), the derivatives of total matches B and aggregate welfare $W(R^b, B(R^b))$ with respect to the number of buyers in the busy market R^b are negative.*

Proposition 4 shows that as we *move* along the stationary frontier from the pooling equilibrium toward an advertising equilibrium, both the welfare and the total matches have to increase. Intuitively, the major reason for this result is that the pooling equilibrium is characterized by the "proportionality" condition $B = R^b$ and the stationarity condition $B = p_a m(1-R^b, 1-B) + p_b m(R^b, B)$, both of which ignore the difference in surplus w_a vs. w_b that different types of meetings produce. In the advertising equilibrium, buyers take into account the difference in expected surpluses, and they "move" away from the pooling equilibrium in the "right" direction. However, the advertising equilibrium, once reached, is not necessarily efficient. In fact, we have examples where the advertising equilibrium over- or undershoots the welfare-maximizing allocation.

Proposition 5. *The welfare in the advertising equilibrium is not necessarily higher than that in the pooling equilibrium.*

How do we gain traction? There are two additional assumptions that, together, make the model very tractable. The first assumption puts additional structure on the relationship between the match probabilities (p_a, p_b) and the expected surplus from a meeting (w_a, w_b).

Definition 3. A *"no price dispersion" market* is one in which $\frac{p_a}{p_b} = \frac{w_a}{w_b}$, or that expected pay-off to a seller conditional upon a match is the same regardless of the seller type.

Recall that $\mathbb{E}[\max\{v - C_s, 0\}] = \mathbb{E}[v - C_s \mid v - C_s \geq 0] p_s$. The type of the seller—or the level of her costs C_s—affects the social surplus through both the probability with which she matches *and* the size of the pie conditional on the match. In a "no price dispersion" market, the value $\mathbb{E}[v - C_s \; v - C_s \geq 0]$ does not depend on the type of the seller. This is true in any class of models where the distribution of C_s conditional on being below v is the same for both types; one simple example is C_s taking two values ("high" and "low") where only the "low" one leads to match. This assumption aligns the weights that the welfare function in Equation 8 puts on meetings in different markets with the "prices" those meetings have in the stationary allocation equation, leading to the following result.

Proposition 6. *Under "no price dispersion," maximizing welfare (Equation 8) is synonymous with maximizing the number of matches B, or the equilibrium number of busy sellers.*

Proof. Social welfare is maximized by solving

$$\begin{aligned} &\max_{R^b, B} w_a m(1-R^b, 1-B) + w_b m(R^b, B) \\ &\quad \text{s.t.} \quad B = p_a m(1-R^b, 1-B) + p_b m(R^b, B). \end{aligned}$$

Multiplying the objective function by a constant p_b / w_b and using the fact that $\frac{p_a}{p_b} = \frac{w_a}{w_b}$, we can transform the objective function into the expression for B. □

Proposition 6 makes the characterization of the socially-efficient outcome simple. Intuitively, this is because when realized surplus is the same for all matches, the social planner simply wants to maximize the number of people who are working. The following restriction on the meeting function m, in turn, simplifies the characterization of the advertising equilibrium.

Assumption 3. The elasticity of the meeting function with respect to the number of buyers in a market is constant: $\frac{\partial m(x,y)}{\partial x}\frac{x}{m(x,y)} \equiv \eta$.

Assumption 3 effectively restricts the meeting function to be Cobb-Douglas and does not hold in the case of the more general CES matching function, as the elasticity depends on the particular values of x and y.[5]

In Proposition 7, we combine Definition 3 and Assumption 3 to show that advertising equilibrium is constrained-efficient. It is an equilibrium that maximizes the total number of busy sellers, B, which is also the number of matches formed each period.

Proposition 7. *With no price dispersion and constant matching efficiency, the advertising equilibrium is constrained-efficient.*

Proof. Proposition 6 shows that the social planner would like to maximize the total number of busy sellers B. To show this is B-maximizing equilibrium is the same as the advertising equilibrium, consider the first order condition for the B-maximization problem:

$$\frac{d}{dR^b}[p_a m(1-R^b,1-B)+p_b m(R^b,B)] = 0$$

$$-\eta p_a \frac{m(1-R^b,1-B)}{1-R^b} + \eta p_b \frac{m(R^b,B)}{R^b} = 0$$

$$w_a \frac{m(1-R^b,1-B)}{1-R^b} = w_b \frac{m(R^b,B)}{R^b},$$

which is the buyer indifference condition for the stationary advertising equilibrium from Definition 2. ☐

Figure 6 illustrates the situation. The y-axis is the number of "busy" sellers, B, and the x-axis is the share of buyers recruiting in the "busy" sub-market, R^b. The heavy dark line depicts the set of stationary allocations $B = m(R^a,A)p_a + m(R^b,B)p_b$. These are combinations of buyers and sellers in the "busy" market such that these shares would remain unchanged, given how matching works.

The 45-degree line is consistent with the pooling condition $R^b = B$ when meetings are random. The curved dashed line indicates the buyer indifference condition $m(R^a,A)w_a/R^a =$

[5] Empirically, there is work estimating matching function parameters with Cobb-Douglas and the more general CES matching function (see Bernstein, Richter and Throckmorton (2022) for an overview), along with the non-parametric approaches Lange and Papageorgiou (2020). Every example we are aware of is from conventional labor markets. The literature is unsettled on how consequential the constant elasticity assumption is in practice.

$m(R^b,B)w_b/R^b$, which is met when pursuing busy and available sellers offer the same expected pay-off for buyers. Where these two curves intersect the curve of stationary equilibria, we have the advertising and pooling equilibria.

Figure 6: Illustration of the busy seller equilibria

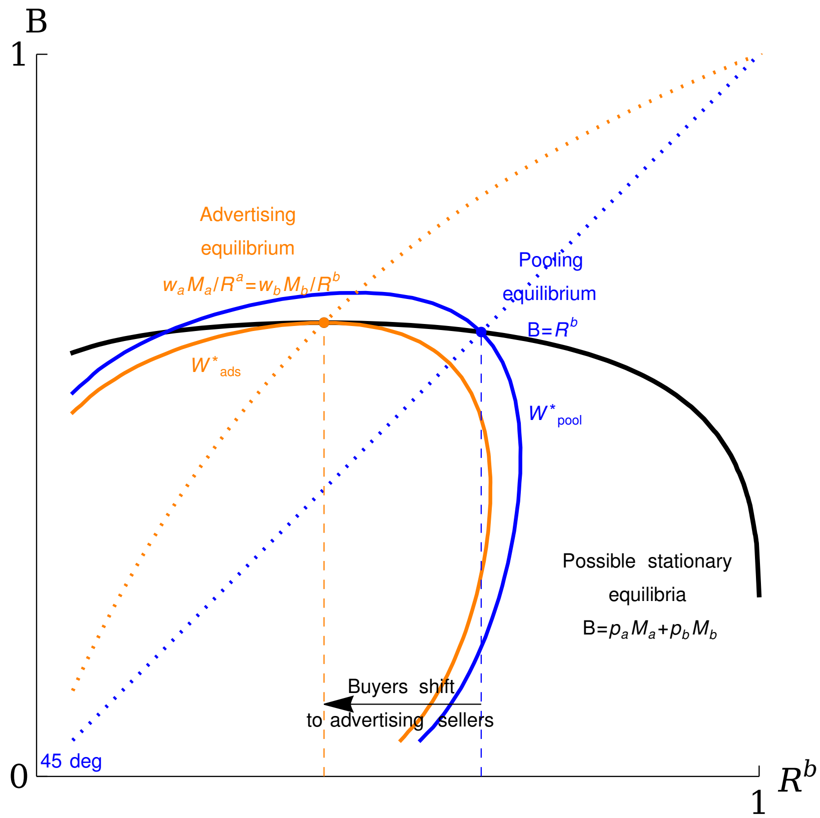

Notes: The heavy dark line indicates stationary equilibria: $B = m(R^a,A)p_a + m(R^b,B)p_b$. The buyer indifference condition between the busy and available seller markets is indicated by the curved line from the origin, $m(R^a,A)w_a/R^a = m(R^b,B)w_b/R^b$. Where it intersects the stationary equilibria frontier is the advertising equilibrium. The random matching condition is a 45-degree line from the origin, $R^b = B$, and where it intersects the stationary equilibria frontier is the pooling (no advertising) equilibria. The iso-welfare curves for these two equilibria are indicated by W^*_{pool} and W^*_{ads}. Note that the advertising equilibria iso-welfare curve is tangent to the frontier. This is the social welfare-maximizing equilibrium.

Note that in Figure 6, from the pooling condition, R^b buyers would prefer to pursue available sellers if they could condition on buyer status, as the indifference point for buyers is with a lower equilibrium fraction of buyers pursuing busy sellers, i.e., a lower R^b. The arrow indicates this reduction in recruiting busy sellers that occurs with advertising.

6.8 The model predictions versus the data

In this section, we connect some of our theoretical results to our empirical findings. The model is, of course, highly stylized compared to the complexity of the empirical context. In reality, the number of buyers and sellers are both endogenous in the long run, capacity is not a binary state but a matter of degree, buyers send multiple inquiries, and so on. As such, the main focus is on comparing the model's welfare predictions.

Perhaps the biggest strength of our setting is that it allows for a test of Proposition 4. We were able to establish that, as one starts "moving away" from the equilibrium without advertising to the one with it, the total number of buyers searching in the 'busy' market should decrease, and the number of matches should increase. Both findings are strongly supported by the empirical findings of Section 4. Our increase in total transactions could also be consistent with the stronger assumptions of Proposition 8. Furthermore, although Proposition 5 shows that advertising does not necessarily increase welfare, the results of Section 4 are also consistent with the prediction that introducing advertising increases equilibrium matches.

The equilibrium notion in the model required buyers to be indifferent between seeking out available and busy sellers. Sellers, on the other hand, should enjoy higher "meeting" rates for the separating equilibrium to exist (see Proposition 2). A reasonable proxy for a "meeting" is a buyer inquiry, as it may or may not lead to a "match" (a contract being formed). Consistent with our model, advertising sellers enjoyed more buyer inquiries (see Table 2). However, the buyers hardly had an identical meeting rate in both markets, as evidenced by our finding of a higher number of sellers' proposals in response to buyers' inquiries (Table 4). One explanation for this discrepancy is that a stationary equilibrium did not generate our data. Another possibility is that there is differentiation among sellers that our "available" versus "busy" characterization is not complete. As we saw, without the ability to see advertising, buyers were slightly more likely to seek out non-advertisers.

7 Long-run evidence of a sustained equilibrium

We have so far focused on the time period following the introduction of advertising to the market. One concern with our approach is that the effectiveness of advertising may dissipate in the long-run. Next, we present evidence that advertising sustained its efficiency in the long-run.

We collected additional data on buyer inquires and seller advertising two years after the experiment took place. To stay close to the design of the original experiment, we restrict our sample to sellers who received at least one inquiry during July, 2023, and we we collect data for these sellers for the month of August, 2023. It is worth noting that during this time period (i) sellers from all platform categories were eligible to advertise, and (ii) seller advertising choices were not taken into account when determining seller rankings.

Our empirical approach relies in utilizing within-seller variation in advertising choices to estimate the effect of advertising on the number of inquiries received by the sellers. Specifically, we estimate the following regression:

$$y_{it} = \beta_0 + \beta_1 \text{ADVERTISING}_{it} + f_i + \tau_t + \epsilon_{it},$$

where y_{it} is the number of inquiries received by seller i on day t, ADVERTISING_{it} is a dummy variable indicating whether seller i advertised on day t, f_i is a fixed effect for seller i, and τ_t is a fixed effect for day t. The coefficient β_1 captures the effect of advertising on the number of inquiries. The main identifying assumption is that a seller who engages in advertising doesn't do so when her profile becomes more attractive or visible to the buyers. We cluster standard errors at the individual seller level. Table 8 reports the results.

Table 8: The effects of advertising on the sellers two years after the experiment

	Dependent variable:	
	Buyer inquiries received	
	(1)	(2)
CONSTANT	0.069^{***}	
	(0.001)	
ADVERTISING_{it}	0.096^{***}	0.036^{***}
	(0.003)	(0.001)
Daily fixed effects	×	✓
Seller fixed effects	×	✓
Observations	2,118,075	2,118,075
R^2	0.009	0.493

Notes: This table reports regressions where the dependent variable is the number of buyer inquiries seller i received on day t and the independent variable is an indicator for advertising activity by i on day t. The sample consists of all sellers who received at least one buyer inquiry in July 2023, and we use August 2023 for the analysis. Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

The first column of Table 8 reports the results of an OLS regression without fixed effects, and the second column reports the results with fixed effects. The latter estimates that one day of advertising increases the number of buyer inquiries by about 0.036 per day, which is more than a 50% increase relative to the average number of inquiries received by sellers on days when they didn't advertise. The point estimate is remarkably close to the one in the original experiment (see Table 7). Two years after its introduction, advertising appears to continue to lead to more buyer inquiries for the advertising sellers. We take it as evidence that the virtuous equilibrium we identified in Section 3 is sustained in the long run.

8 Implications for sponsored search advertising

Our focus is on advertising as a solution to a market failure, which is not exclusive to digital markets. However, our context also relates closely with the sponsored search context, which is of great practical and policy importance. A key welfare consideration with sponsored search advertising is the effect it might have on the volume of economic transactions and the competitiveness of markets. For example, if advertising simply favors market incumbents and results in higher prices, we might take a dimmer view of the industry from a policy perspective. But if it instead enables greater competition and serves a low deadweight loss for platforms to raise revenue, we might be more enthusiastic.

In sponsored search advertising, buyers are given a more prominent position, and to the extent that would-be buyers have limited attention and sponsored ads are more likely to be seen, crowd-out of non-advertising sellers is nearly guaranteed. Or, as in our empirical case, if advertising renders sellers more salient—but does not give them more "real estate"—non-advertisers could still be crowded out. The welfare question is whether net sales decline if advertising re-directs buyer attention to advertisers who are less suitable on some dimension.

There are three papers particularly closely related to our own: Sahni and Nair (2019), Moshary (2021), and Hui and Liu (2022). Each paper varies exposure or salience to advertising at the buyer level. Sahni and Nair (2019) and Moshary (2021) have sharply contrasting results: the former finds advertising increases overall sales, while Moshary (2021) finds that it decreases overall sales. The proposed mechanism in Moshary (2021) is that buyers do not "like" the platform's advertisements, at least relative to its organic listings. This seems likely, given that the design seemingly mechanically ensures advertisers are relatively adversely selected compared to organic sellers: "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."

Sahni and Nair (2019) randomize buyers to being able to see *whether* sellers have advertised without changing the collection of available sellers (as do we). Moshary (2021) randomizes buyers to be able to see advertising sellers, with treated buyers having search results with advertisers removed if those advertisers would not have won the auction organically. In short, the consideration set changes in the context of Moshary (2021) but it does not in the context of Sahni and Nair (2019).

Hui and Liu (2022) work in the same empirical context as Moshary (2021), but their design includes variation in the salience of the sponsored listing and, critically, in other sources of information about seller quality. Hui and Liu show that consumers on this platform generally do not like advertisements *but* that perceptions of quality strongly mediate this effect: advertising sellers gain far more business from advertising when their listing has a badge indicating they are higher-quality sellers. This echoes results in Barach, Golden and Horton (2020) that show that platform-provided quality signals strongly influence buyers' selection.

We believe we offer a parsimonious explanation for the main differences in findings: what matters is whether advertisers are relatively positively or virtuously selected from the consumers' perspective. When sellers are adversely selected, as in Moshary (2021) and advertisers displace organic listings, efficiency decreases. When sellers are adversely selected but there is no displacement, buyers simply ignore advertisers. But when sellers are virtuously selected relative to organic listings, giving advertisers greater prominence is likely a free lunch from the platform's perspective.

If advertisers are adversely selected, the platform faces a trade-off between ad commissions and revenue from organic transactions, as the ads crowd out more preferred sellers (Choi and Mela, 2019; Balseiro and Désir, Forthcoming). Suppose advertisers are adversely selected but not given any special prominence. In that case, it seems difficult for this kind of advertising to work in equilibrium: buyers would learn to ignore advertisers, and sellers would prefer to pool with non-advertisers. However, if sellers are positively selected, the platform faces no trade-off. This might at first seem like something out of the platform's control. However, numerous policy choices available to the platform can influence selection into advertising to ensure that the "good are rich" (and spend money on advertising) rather than hope that "the rich are good."

What makes a seller virtuously selected is presumably context-dependent. In our setting, the seller's capacity to take on more work is likely a key consideration and a seller's private information. Advertising sellers were highly virtuously selected on that dimension, and buyers acted accordingly. One challenge is that advertising sellers can displace organic sellers; there is a horse race between the quality of organic results to a search query and the quality of the targeting tools. If the platform is very good at targeting—if the organic results are excellent—then it is harder for advertisers to be virtuously selected.

9 Conclusion

We show that advertising can help to overcome a market failure by serving as a signal. To serve this function, advertising had to be costly, consistent with Nelson (1974) and Milgrom and Roberts (1986). Signaling was about capacity in our case, but it is easy to imagine this could vary based on the context. The common economic problem is buyer uncertainty about seller suitability—which could be price, capacity, quality, or some other vertical attribute.

It is interesting to note that other economic institutions have evolved to solve the capacity problem, although we are aware of no cases where the signal cost is determined centrally. In many of these scenarios, the cost of signaling a willingness to "trade" is more of a hassle cost or a technical cost, and the solution is imperfect. In the conventional advertising context, the cost of ads was partially physical—literal paper and ink—and the price to advertise was determined by whatever firm had the next higher value for the advertising spot. But then creatives could spend money on the ads—as much as they like—to create a stronger signal. In our setting, the platform can pick a price for advertising that maximizes the informational content. We do not explore the question of the optimal quantity of advertising to maximize information, but it poses an interesting market design problem.

References

Abhishek, Vibhanshu, Kinshuk Jerath, and Siddharth Sharma, "The impact of sponsored listings on online marketplaces: Insights from a field experiment," feb 2022. Available at SSRN: `https://ssrn.com/abstract=3013468` or `http://dx.doi.org/10.2139/ssrn.3013468`.

Agrawal, Ajay, John J Horton, Nicola Lacetera, and Elizabeth Lyons, "Digitization and the contract labor market: A research agenda," in "Economic analysis of the digital economy," University of Chicago Press, 2015, pp. 219–250.

Athey, Susan and Glenn Ellison, "Position auctions with consumer search," *The Quarterly Journal of Economics*, 08 2011, *126* (3), 1213–1270.

Bagwell, Kyle, "The economic analysis of advertising," *Handbook of industrial organization*, 2007, *3*, 1701–1844.

_ **and Garey Ramey**, "Coordination economies, advertising, and search behavior in retail markets," *The American Economic Review*, 1994, pp. 498–517.

Balseiro, Santiago R. and Antoine Désir, "Incentive-compatible assortment optimization for sponsored products," *Management Science*, Forthcoming, *0.*

Barach, Moshe A., Joseph M. Golden, and John J. Horton, "Steering in online markets: The role of platform incentives and credibility," *Management Science*, 2020, *66* (9), 4047–4070.

Benham, Lee, "The effect of advertising on the price of eyeglasses," *The Journal of Law and Economics*, 1972, *15* (2), 337–352.

Benson, Alan, Aaron Sojourner, and Akhmed Umyarov, "Can reputation discipline the gig economy? Experimental evidence from an online labor market," *Management Science*, 2019.

Bernstein, Joshua, Alexander W. Richter, and Nathaniel A. Throckmorton, "The matching function and nonlinear business cycles," Technical Report 2201, FRB of Dallas Working Paper February 2022.

Blake, Thomas, Chris Nosko, and Steven Tadelis, "Consumer heterogeneity and paid search effectiveness: A large-scale field experiment," *Econometrica*, 2015, *83* (1), 155–174.

Choi, Hana and Carl F. Mela, "Monetizing online marketplaces," *Marketing Science*, 2019, *38* (6), 948–972.

Decarolis, Francesco and Gabriele Rovigatti, "From mad men to maths men: Concentration and buyer power in online advertising," *American Economic Review*, October 2021, *111* (10), 3299–3327.

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz, "Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords," *American economic review*, 2007, *97* (1), 242–259.

Einav, Liran, Chiara Farronato, and Jonathan Levin, "Peer-to-peer markets," *Annual Review of Economics*, 2016, *8*, 615–635.

Filippas, Apostolos, John Horton, and Richard Zeckhauser, "The surprisingly low cost of free goods: Stimulating consumption with in-kind transfers," *Working paper*, 2023.

_ , John J Horton, and Joseph Golden, "Reputation inflation," *Marketing Science*, 2022.

_ , _ , and Richard J Zeckhauser, "Owning, using, and renting: Some simple economics of the "sharing economy"," *Management Science*, 2020, *66* (9), 4152–4172.

Fradkin, Andrey, "Search, matching, and the role of digital marketplace design in enabling trade: Evidence from Airbnb," 2023. Working Paper.

Golden, Joseph and John Joseph Horton, "The effects of search advertising on competitors: An experiment before a merger," *Management Science*, 2021, *67* (1), 342–362.

Horton, John J, "Online labor markets," *Internet and Network Economics*, 2010, pp. 515–522.

_ , "The effects of algorithmic labor market recommendations: Evidence from a field experiment," *Journal of Labor Economics*, 2017, *35* (2), 345–385.

_ , "Buyer uncertainty about seller capacity: Causes, consequences, and a partial solution," *Management Science*, 2019, *65* (8), 3518–3540.

_ , William R Kerr, and Christopher Stanton, "Digital labor markets and global talent flows," Technical Report, National Bureau of Economic Research 2017.

Hui, Xiang and Meng Liu, "Quality certificates alleviate consumer aversion to sponsored search advertising," 2022. CESifo Working Paper No. 9886, Available at SSRN: `https://ssrn.com/abstract=4196282` or `http://dx.doi.org/10.2139/ssrn.4196282`.

Janssen, Maarten and Eric Rasmusen, "Bertrand competition under uncertainty," *The Journal of Industrial Economics*, 2002, *50* (1), 11–21.

Kirmani, Amna and Akshay R Rao, "No pain, no gain: A critical review of the literature on signaling unobservable product quality," *Journal of marketing*, 2000, *64* (2), 66–79.

Lambin, Jean-Jacques, *Advertising, competition and market conduct in oligopoly over time; an econometric investigation in western European countries*, Amsterdam (Netherlands) North Holland Pub., 1976.

Lange, Fabian and Theodore Papageorgiou, "Beyond Cobb-Douglas: Flexibly estimating matching functions with unobserved matching efficiency," Technical Report 26972, National Bureau of Economic Research April 2020.

Milgrom, Paul and John Roberts, "Price and advertising signals of product quality," *Journal of Political Economy*, 1986, *94* (4), 796–821.

Mortensen, Dale T and Christopher A Pissarides, "Job creation and job destruction in the theory of unemployment," *The review of economic studies*, 1994, *61* (3), 397–415.

Moshary, Sarah, "Sponsored search in equilibrium: Evidence from two experiments," *Available at SSRN 3903602*, 2021.

Nelson, Phillip, "Advertising as information," *Journal of Political Economy*, 1974, *82* (4), 729–754.

Pallais, Amanda, "Inefficient hiring in entry-level labor markets," *American Economic Review*, 2013.

Prada-Sarmiento, Juan David, "A note on concavity, homogeneity and non-Increasing returns to scale," *Working paper*, 2010.

Sahni, Navdeep S and Harikesh S Nair, "Does advertising serve as a signal? Evidence from a field experiment in mobile search," *The Review of Economic Studies*, 10 2019, *87* (3), 1529–1564.

Schmalensee, Richard, "A model of advertising and product quality," *Journal of Political Economy*, 1978, *86* (3), 485–503.

Shapiro, Bradley T., "Positive spillovers and free riding in advertising of prescription pharmaceuticals: The case of antidepressants," *Journal of Political Economy*, 2018, *126* (1), 381–437.

Stanton, Christopher T and Catherine Thomas, "Landing the first job: The value of intermediaries in online hiring," *The Review of Economic Studies*, 2016, *83* (2), 810–854.

Stigler, George J, "The economics of information," *Journal of Political Economy*, 1961, *69* (3), 213–225.

Woodcock, Ramsi A, "The obsolescence of advertising in the information age," *Yale Law Journal*, 2017, *127*, 2270.

Wright, Randall, Philipp Kircher, Benoît Julien, and Veronica Guerrieri, "Directed search and competitive search equilibrium: A guided tour," *Journal of Economic Literature*, 2021, *59* (1), 90–148.

A Additional experimental details and results

A.1 Buyer view of the advertising information

Figure 7 shows an example of a buyer's view of sellers during the experiment. It depicts the case where the seller on the top has chosen to advertise, and the seller in the bottom has chosen not to advertise. Upon hovering over the advertising "badge," the buyer could see the text: "This seller is promoting that they're open to more work." Advertising was visible only while the buyer was searching for sellers to send inquiries to, and the ability to view the advertising information was the sole difference between treated and control buyers.

Figure 7: An example of a buyer interface during the experiment

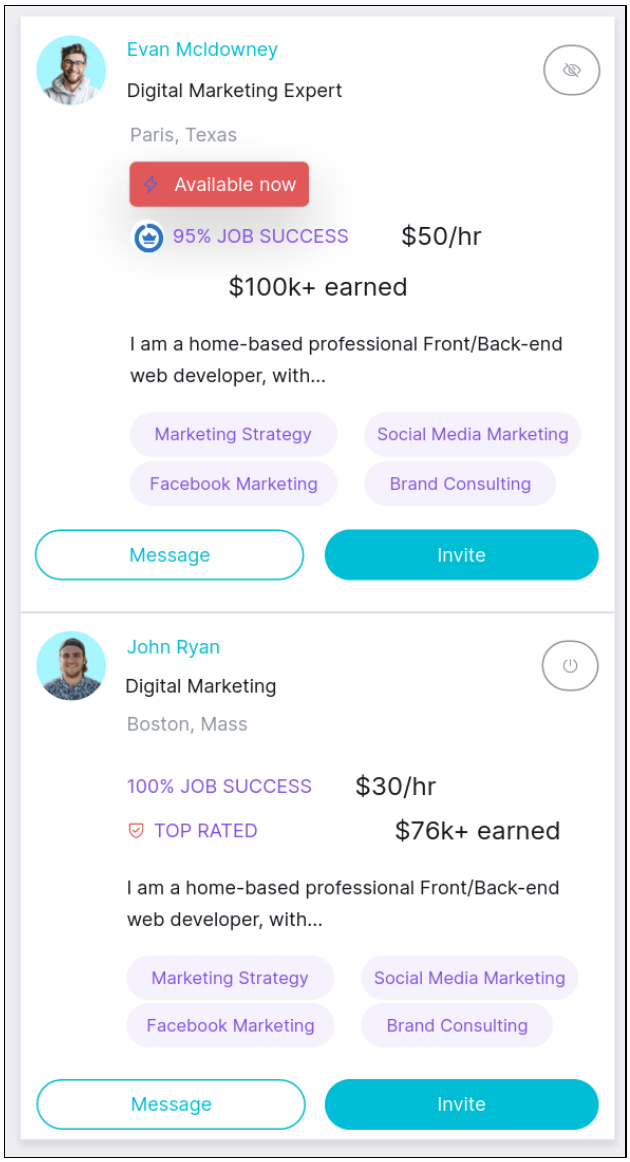

A.2 Internal validity

The experimental groups were well-balanced across several pre-experimental observables. To assess whether the randomized assignment was performed correctly, we test for systematic differences in observable pre-treatment outcomes between buyers assigned to the control and the treatment groups. Table 9 reports two-sided t-tests for various buyer project-specific outcomes. Figure 8 plots the number of buyers allocated to the experimental cells over time.

Table 9: Balance test table for allocated buyers

	Control mean $\overline{X}_{CTL}$	Treatment mean $\overline{X}_{T}$	p-value
Project-specific outcomes			
number of project posts	2.36	2.31	0.269
number of inquiries sent to sellers	4.04	4.71	0.337
number of seller applications received	20.91	21.04	0.694
number of contract offers extended	0.67	0.67	0.858
Observation counts	41,951	42,474	0.072

Notes: This table reports averages and p-values of two-sided t-tests for various pre-treatment outcomes, for buyers assigned to the control and treatment group. The reported outcomes are (i) the number of projects posted, (ii) the number of inquiries sent to sellers per post, (iii) the number of seller applications received per post, (iv) the number of contract offers extended per post.

Figure 8: Buyers allocated to the control and treatment groups over time

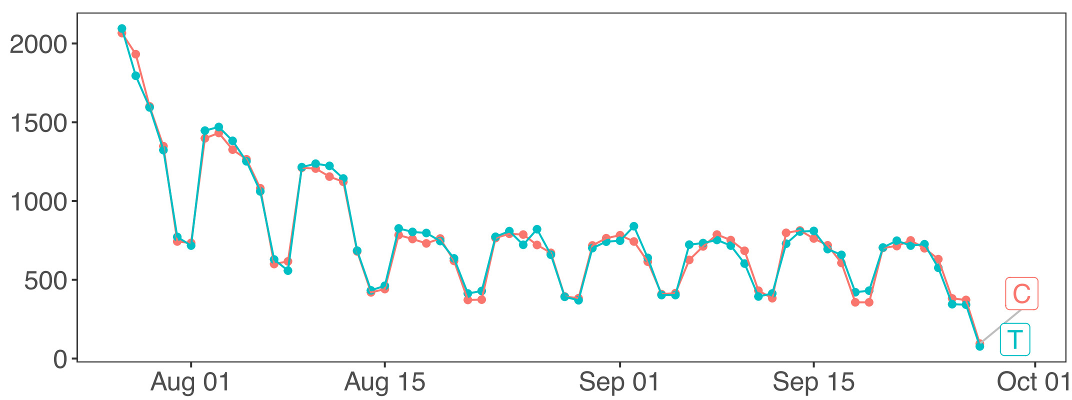

Notes: This figure plots the number of buyers allocated to the control and treatment groups each day of the allocation period. The allocation period began on July 26, 2021 and ended on September 27, 2021.

A.3 Advertising uptake and exposure

Figure 9 provides more details on the sellers' uptake of advertising over time, and the buyers' exposure to advertising over time. Figure 9a displays the number of sellers advertising each day during the experiment. We also examines dynamic measures of seller capacity and compare them to advertising uptake. Specifically, we look at whether sellers who advertised were actively applying for projects organically. To do so, we define a seller as actively applying for projects at a given time if they had applied for at least one project in the previous seven days. We normalize this number by the total number of sellers who applied to at least one project during the experimental period. Our findings indicate that many advertisers also actively applied to projects without being recruited via an inquiry.

Figure 9: Details on advertising uptake and exposure

(a) Sellers' uptake of advertising

(b) Distribution of the percentage of advertising sellers in buyer search

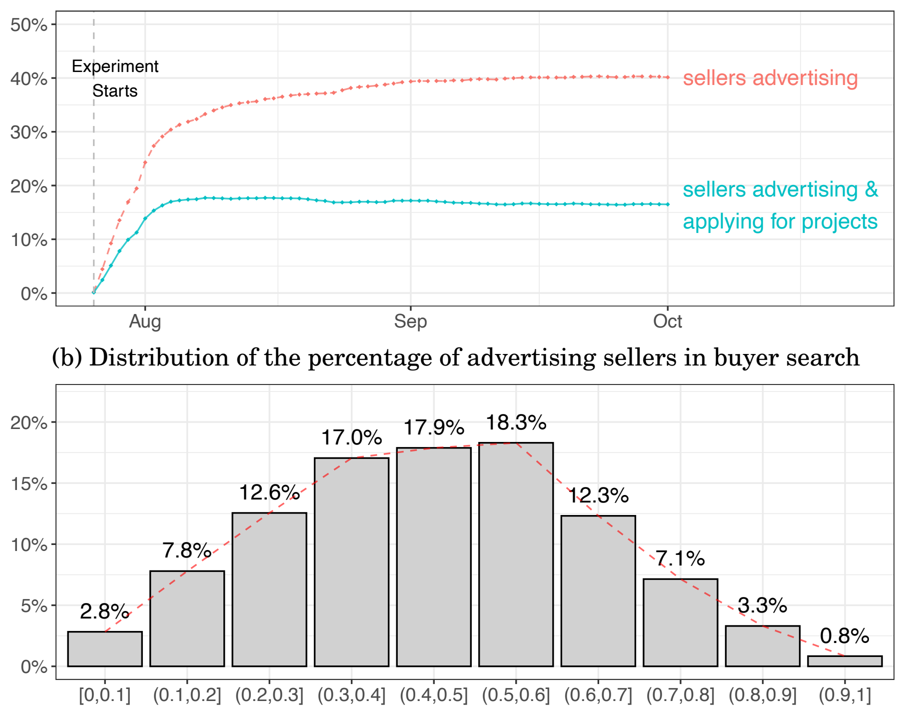

Notes: Panel (a) plots the sellers' advertising uptake over time. It plots the number of sellers advertising, and both advertising and actively applying for projects, for each day during the experimental period. We define a seller as "active" on a given day if she applied for projects within the previous seven-day window. Panel (b) depicts the distribution of the percentage of advertisers in buyer search, using data for the last week of the experiment. Panel (b) uses data only for treated buyers.

Advertisers were displayed prominently in buyer search, with about 49.6% of all impressions and 52.6% of first-page impressions coming from advertisers. Nevertheless, each buyer's experience may have differed, even if advertisers were commonplace. Figure 9b shows the distribution of the percentage of advertisers seen by buyers using data from the last week of the experiment. We can see that exposure to advertising was widespread and that seeing no advertisers was a rare event for buyers.

A.4 Statistically modeling seller selection into advertising

We quantify the relative importance of various factors that predict whether a seller will choose to advertise. To do this, we report the result of a logistic regression where the outcome is an indicator variable for selection into advertising, and the independent variables are pre-experiment seller attributes and outcomes. We standardize the independent variables (mean 0 and 1 standard deviation) and use them as predictors in the logistic regression. Standardization allows for comparing the relative importance of the predictors in the model.

Figure 10 reports those coefficients, ordered from largest to smallest. The coefficients are the effects on log odds; above each effect, we report the implied percentage change in advertising probability from a one standard deviation increase in that measure. The baseline advertising adoption level is 45% of sellers.

Figure 10: Relative importance of factors predicting seller selection into advertising

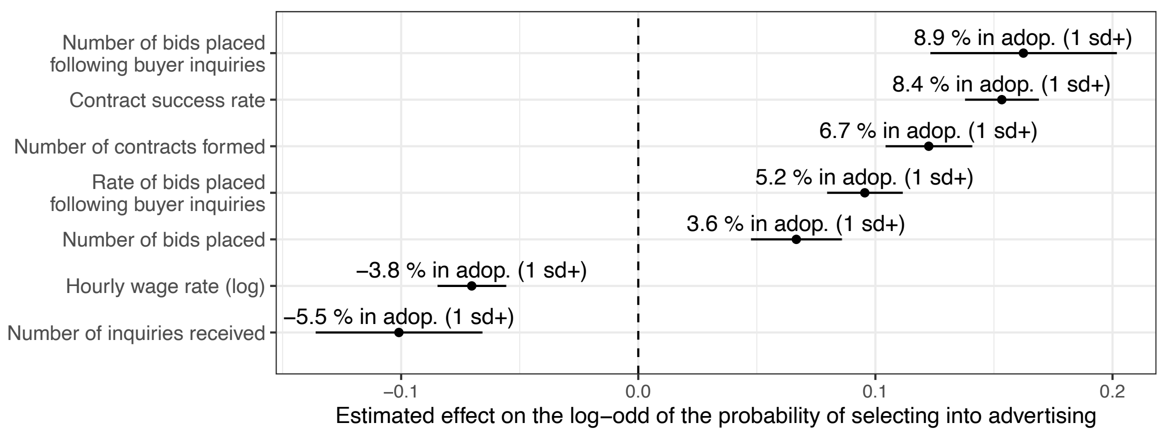

Notes: This figure reports estimates of the effects of pre-experiment seller attributes on the probability that they select in to advertising. The estimates are obtained through a logistic regression where the outcome is a binary indicator for the seller selecting into advertising, and the independent variables are standardized attributes. We report a 95% confidence interval around each estimate and the implied percentage change in the probability of selecting in to advertising from one standard deviation increase in the corresponding attribute.

The positive predictors of advertising are being highly active: many accepted buyer inquiries ("Number of bids placed following buyer inquiries"), more contracts ("Number of contracts formed"), a higher acceptance rate of buyer inquiries ("Rate of bids placed following buyer inquiries"), and the total number of bids (including when not solicited by buyers). We can also see that sellers with higher reputation scores—as indicated by their successful contract completion rate ("Contract success rate")—are more likely to select in to advertising. In contrast, those with higher wage asks are less likely to select in to advertising. Regarding magnitudes, the coefficient on the contract success score is 0.15, which implies that a seller with a 1 SD higher score has a 8.4% higher probability of adopting advertising than the baseline, assuming log-odds are linear in the predictors.

The two factors that predict being less likely to advertise are (a) a higher hourly rate and (2) a greater number of buyers inquiries already received. With all of the caveats needed for this cross-sectional analysis, a simple interpretation is that buyers advertising were interested in more work, as evinced by a relatively lower number of seller inquiries and a lower wage ask.

The relative price offered by advertisers has been a point of interest to those considering the economic effects of advertising. Schmalensee (1978) shows that firms with lower marginal costs—which might be associated with lower quality—might be most interested in advertising, making advertisers adversely selected. Of course, this would also cause buyers to learn to ignore advertising if being influenced by it lowered their utility. Consumers must believe paying attention to advertising is worthwhile for a separating equilibrium to exist. Outside our context, consumers generally seem to believe that costly advertising indicates product quality—but consumers also state that too much advertising could signal desperation (see Kirmani and Rao (2000) and references therein).

A.5 Expanding the sample and looking longer term

The Poisson regression analysis (see Table 4) used a sample comprising only the first project posted during the experimental period for each buyer allocated to the experiment. One might worry that there is a novelty effect that wears off with greater experience.

To expand our analysis, we can exploit that some buyers can post many project posts during the experiment. This gives us three possible ways of analyzing the experiment: (1) "Pooled" uses the entire sample of project-post-level outcomes and clusters at the buyer level, (2) "Averaged" uses buyer-level averages for each outcome, and (3) "First Project" restricts the sample to each buyer's first project-post-level outcomes during the experiment. Our preferred specification is "Pooled" given that it captures the most behavior, but the other two will allow us to examine how the effects of the treatment varied over time.

For many of our outcomes, it makes sense to think of the outcome as having an extensive margin—was there any of this behavior or outcome? Accordingly, we examine the effects of the treatment in two different ways. First, we consider the extensive margin effect of the treatment, applying the indicator variable transformation on each outcome. This allows us to study the fractions of buyers that had non-zero outcomes during the experimental period. Second, we consider the intensive margin of the treatment. To deal with extreme values, we winsorize the distributions of the non-binary outcomes at the 99% level.

We regress each outcome on indicators for the treatment, i.e.,

$$y_j = \beta_0 + \beta_1 \text{ADSVISIBLE}_j + \epsilon,$$

where y_j is the buyer outcome of interest, ADVISIBILE_j indicates whether buyer j was assigned to the treatment group, and ϵ is an error term. We report the estimated effects as percentage changes over the control group outcome in Figure 11, by plotting the least squares estimate $\hat{\beta}_1/\hat{\beta}_0$ for each of the three active treatment groups, along with a 95% confidence interval around each point estimate.

Treated buyers were more likely to send at least one inquiry: compared to a baseline of 53.09% for the control group, the increase was 1.33 percentage points (2.51%). Similarly, treated buyers sent 0.1 (3.56%) more inquiries per post. Interestingly, treated buyers redirected their recruiting efforts substantially towards advertising sellers. Buyers who were able to see the advertising information were 1.86 percentage points (4.59%) more likely to send an inquiry to an advertiser, and sent on average 0.09 more inquiries (7.08%) to advertising sellers.

Treated buyers subsequently received better responses to their inquiries. In particular, more of the treated buyers' inquiries were responded to, as well as responded to fast (defined as getting a response within 48 hours), both on the extensive and on the intensive margin. Importantly, treated buyers were 1.33 percentage points (3.04%) more likely to receive a seller application following an inquiry, and received 0.06 (4.29%) more such applications. While seller applications following an buyer inquiry constitute only a small fraction of the applications buyers receive, they are important: treated buyers were 0.83 percentage points (2.63%) more likely to make at least one hire.

The above results suggest that the advertisement had an immediate, positive effect on buyers. Almost all estimated treatment effects in the "first" sample seem smaller than in the "all" sample. This suggests that the experiment's effects are unlikely to be transitory, as they seem to increase in magnitude over time. It is also worth noting that advertising did not affect the sellers' rankings on the platform, sellers may not have been adopted immediately, and advertising was likely priced sub-optimally low during its roll-out; we explore these factors in the rest of the paper.

Figure 11: Effects on buyer outcomes from being able to see advertising

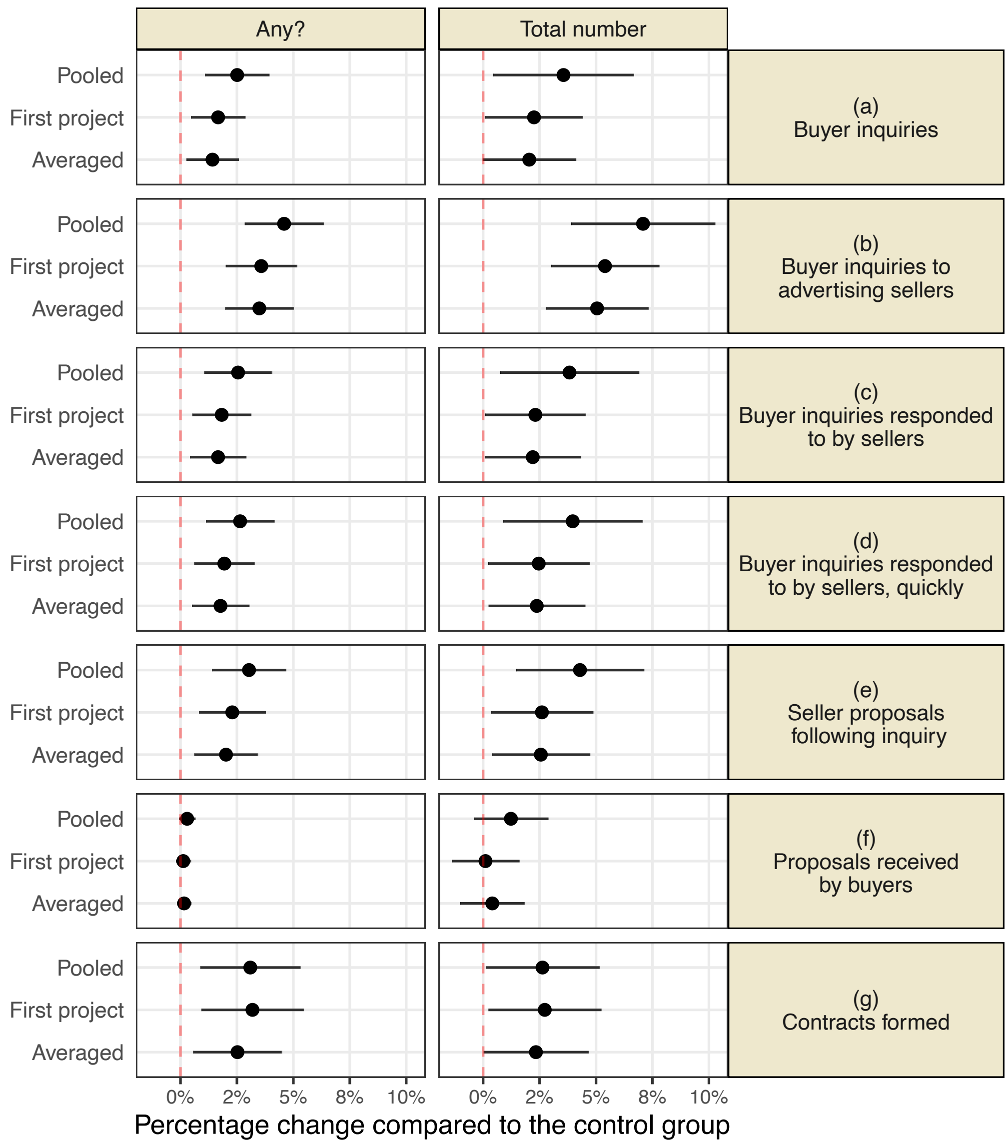

Notes: This figure plots estimates of the effects of the treatment on buyer outcomes, using cross-sectional data. Each panel reports point estimates as the percentage change in the treatment group over the control group, along with a 95% confidence interval. Panels on the left examine the extensive margin effect, with the dependent variable being the indicator variable transformation of each outcome. Panels on the right examine the intensive margin effect, with the dependent variable being the "raw" outcome winsorized at the 99% level. Estimates are computed for three different samples: (i) "all" uses the entire sample and clusters standard errors on the buyer level, (ii) "avg" uses buyer-level averages for each outcome, and (iii) "first" only uses each buyer's first project post during the experiment. See Section A.5 for details on the sample construction and the estimation strategy, and for a discussion of the results.

B Proofs of our theoretical results

B.1 Proof of Proposition 1

The stationary pooling equilibrium could be obtained in closed form. Using conditions (1) and (2), condition (3) can be written as

$$B^\star = m(1-B^\star, 1-B^\star)p_a + m(B^\star, B^\star)p_b$$

By constant returns to scale, the above becomes $B^\star = m(1,1)p_a(1-B^\star) + m(1,1)p_b B^\star$, which implies

$$B_{pool} = \frac{m(1,1)p_a}{1+m(1,1)(p_a-p_b)}, \qquad A_{pool} = \frac{1-m(1,1)p_b}{1+m(1,1)(p_a-p_b)} \tag{A1}$$

B.2 Proof of Proposition 3

Our proof proceeds by demonstrating that there always exists a collection (R^a, R^b, A, B) that satisfies conditions (1)-(3) in the definition of a stationary advertising equilibrium (Definition 2). We then show that those conditions guarantee the existence of the price π that supports the separation of the types. Since $R^a = 1-R^b$ and $A = 1-B$, we can work with the pair (R^b, B).

Step 1. In this step, we demonstrate that the stationarity condition (3) can be rewritten as a concave function $B(R^b)$ with $B(0) > 0$ and $B(1) < 1$.

To show this, fix $R^b \in [0,1)$, and consider the function $f(B) = m(1-R^b, 1-B)p_a + m(R^b, B)p_b - B$. We have $f(0) = m(1-R^b, 1)p_a > 0$ and $f(1) = m(R^b, 1)p_b - 1 < m(1,1)p_b - 1 \le p_b - 1 < 0$. Since $f(B)$ is continuous, a solution $B(R^b)$ to $f(B) = 0$ exists for any value of $R^b \in [0,1)$. Note, in particular, that for $R^b = 0$ we have $f(0) > 0$ and hence $B(0) > 0$.

This solution $B(R^b)$ is unique because $f(B)$ is a concave function. If $B_1 < B_2$ are two distinct solutions, then there exists $\lambda \in (0,1)$ such that $B_1 = \lambda \times 0 + (1-\lambda) \times B_2$. By concavity of f, we get

$$0 = f(B_1) = f(\lambda \times 0 + (1-\lambda) \times B_2) \ge \lambda f(0) + (1-\lambda) f(B_2) = \lambda f(0) > 0,$$

which is a contradiction.

The concavity of $f(B)$ follows from two observations. First, under the assumptions that we make, $m(x,y)$ is a concave function (Prada-Sarmiento, 2010). Second, it is simple to show that if $m(x,y)$ is concave, then so is $m(1-x, 1-y)$. Finally, a sum of concave functions is concave.

One remark is in order about the solution of the equation $f(B) = 0$ when $R^b = 1$. In that case, we are solving $m(1,B)p_b - B = 0$. Under our assumptions, $B = 0$ is a solution, and any solution is less than 1. In the uniqueness argument above, we assumed that $f(0) > 0$, which doesn't hold when $R^b = 1$. This development can generate one additional solution (more than one would still be ruled out by concavity unless $m(1,B)p_b - B$ is 0 everywhere). However, the stationary equilibrium with ($R^b = 1, B = 0$) is of little economic interest because, in such an equilibrium, all the buyers shop in the market for busy sellers while all the sellers are available. If another solution exists, then we use it instead, but if not—our results are not affected by this corner case.

Step 2. In this step, we demonstrate that the indifference condition (2) can be expressed as an increasing function $B(R^b)$ that satisfies $B(0) = 0$, $B(1) = 1$, and $(w_a > w_b) \Longrightarrow (B(R^b) > R^b$ for $R^b \in (0,1))$.

The collection of all points (R^b, B) which make the buyers indifferent between the two markets is given by

$$w_b \frac{m(R^b, B)}{R^b} = w_a \frac{m(1-R^b, 1-B)}{1-R^b} \tag{A2}$$

Using constant returns to scale, this can be written as

$$w_b m\left(1, \frac{B}{R^b}\right) = w_a m\left(1, \frac{1-B}{1-R^b}\right) \tag{A3}$$

For any $R^b \in (0,1)$, there is a unique value of B that satisfies this equation. To see that, consider $f(B) = w_a m\left(1, \frac{1-B}{1-R^b}\right) - w_b m\left(1, \frac{B}{R^b}\right)$. We have $f(R^b) = w_a - w_b \geq 0$. If $w_a = w_b$, it is easy to see that $B = R^b$ is the only solution. In the interesting case of $w_a > w_b$, note that $f(1) = -w_b m(1, 1/R^b) < 0$. The solution exists by continuity of f and it is unique by the monotonicity of $m(1,x)$. Note that, unless $w_a = w_b$, the solution satisfies $B > R^b$ for $R^b \in (0,1)$. It is trivial to see that the solution $B(R^b)$ is an increasing function of R^b.

We now show that $\lim_{x\to 1} B(x) = 1$. First of all, note that $B(R^b)$ is bounded above by 1: for any $R^b \in (0,1)$, $B = 1$ is too 'large' to satisfy equation (A3). Since $B(R^b)$ is an increasing function, then by the monotone convergence theorem and the continuity of $B(R^b)$, a limit at $R^b \to 1$ exists. Now suppose that $\lim_{x\to 1} B(x) < 1$. Then

$$\lim_{x\to 1} w_a m\left(1, \frac{1-B}{1-x}\right) \geq w_a m(1,1),$$

since $m(1,\infty)$ is either infinity, in case $m(1,x)$ is bounded, or at least exceeds $m(1,1)$. Then, by taking the limits of both sides in equation (A3), we get

$$w_b m\left(1, \lim_{x\to 1} B(x)\right) = w_a m\left(1, \infty\right) \geq w_a m(1,1).$$

This is a contradiction, as, generally, $w_b < w_a$ and $m\left(1, \lim_{x\to 1} B(x)\right) < m(1,1)$. Instead, if $\lim_{x\to 1} B(x) = 1$, then both the numerator and the denominator of $(1-B)/(1-R^b)$ are going to 0 as $R^b \to 1$. That allows the equation to be satisfied, provided that the convergence to 0 happens at a particular rate:

$$w_b m(1,1) = w_a m\left(1, \lim_{x\to 1} \frac{1-B(x)}{1-x}\right) \tag{A4}$$

$\lim_{x\to 1} \frac{1-B(x)}{1-x} = L$ where L satisfies

$$m(1,L) = \frac{w_b}{w_a} m(1,1) \tag{A5}$$

The limit of $B(x)$ as $x \to 0$ depends on the functional form of $m(x,y)$. If $m(x,y)$ is unbounded, then $\lim_{x\to 0} B(x) = 0$, as we will show momentarily. However, if the function is bounded above, then it is no longer the case. One example is $m(x,y) = \min(x,y)$, where $B(x)$ takes the form $B(x) = 1-(w_b/w_a)(1-x)$. However, if $m(1,1/x) \to \infty$ as $x \to 0$, then $B(x)$ has to approach 0 for equation (A3) to hold.

Step 3. There is a pair (R^b, B) that satisfies both the indifference condition (2) and the stationarity condition (3).

Let $B = f_1(R^b)$ be the function that describes the indifference condition and $B = f_2(R^b)$ be the function that describes the stationarity condition. We established that $f_1(0) = 0$, while $f_2(0) > 0$. Similarly, we showed that $f_1(1) = 1$ and $f_2(1) < 1$. Since both functions are continuous, there exists a value R^b where the two cross.

Step 4. If (R^b, B) is a point described in Step 3, then there exists a price of advertising π that fulfills the separation condition (7).

We established that all the points that satisfy the indifference condition are such that $B \geq R^b$. That implies $A = 1-B \leq 1-R^b = R^a$ and, hence, $R^a/A \geq R^b/B$. The separation condition (7) is satisfied when

$$\frac{m(R^a, A)}{A} = m(R^a/A, 1) \geq m(R^b/B, 1) = \frac{m(R^b, B)}{B},$$

which holds since $m(x,1)$ is an increasing function.

B.3 Proof of Proposition 4

Let $B(R^b)$ be the frontier of all stationary allocations. Total welfare at any point R^b can then be written as

$$W(R^b, B(R^b)) = w_b m(R^b, B(R^b)) + w_a m(1-R^b, 1-B(R^b)) \tag{A6}$$

We take the derivative of the welfare with respect to R^b and evaluate it at the point where $R^b = B$, which characterizes the pooling equilibrium. That derivative is negative, as we now show. The fact that the welfare-maximizing level R^b lies to the left of the point $R^b = B$ then follows from the concavity of $W(R^b)$, which we also demonstrate.

The derivative of interest has three components:

$$\frac{d}{dR^b}W(R^b,B(R^b)) = \left.\frac{\partial W(R^b,B)}{\partial R^b}\right|_{R^b=B} + \left.\frac{\partial W(R^b,B)}{\partial B}\right|_{R^b=B} \times \left.\frac{dB}{dR^b}\right|_{R^b=B} \tag{A7}$$

Let us evaluate them all.

$$\frac{\partial W(R^b,B)}{\partial R^b} = w_b\frac{\partial m(R^b,B)}{\partial R^b} + w_a\frac{\partial m(1-R^b,1-B)}{\partial R^b} \tag{A8}$$

$$\left.\frac{\partial W(R^b,B)}{\partial R^b}\right|_{R^b=B} = \frac{\partial m}{\partial R^b}(1,1)(w_b - w_a) < 0 \tag{A9}$$

Here we used the fact that partial derivatives of a function that is homogenous of degree 1 are homogenous of degree 0, i.e., $\frac{\partial m}{\partial R^b}(1,1) = \frac{\partial m}{\partial R^b}(x,x)$ for any x.

An identical exercise shows that

$$\left.\frac{\partial W(R^b,B)}{\partial B}\right|_{R^b=B} = \frac{\partial m}{\partial B}(1,1)(w_b - w_a) < 0 \tag{A10}$$

The last part is of special interest because this derivative tells us the change in total matches as we move away from the pooling equilibrium towards an advertising equilibrium where R^b is lower. Since $B(R^b)$ is defined implicitly as the solution of $B = m(1-R^b,1-B)p_a + m(R^b,B)p_b$, we use the inverse function theorem to obtain $B'(R^b) = \frac{dB}{dR^b}$.

$$B'(R^b) = p_b m_1(R^b,B) + p_b m_2(R^b,B)B'(R^b) - p_a m_1(1-R^b,1-B) - p_a m_2(1-R^b,1-B)B'(R^b)$$

Plugging in $B = R^b$ and using the homogeneity of m (and its partial derivatives m_1 and m_2) we get

$$B'(R^b)|_{R^b=B} = -\frac{(p_a - p_b)\frac{\partial m}{\partial R^b}(1,1)}{1+(p_a - p_b)\frac{\partial m}{\partial B}(1,1)} < 0 \tag{A11}$$

The equation above proves that introducing advertising should locally increase the number of matches.

The signs of our three derivatives are not enough to determine the sign of the overall expression A7. We evaluate that expression to find

$$\left.\frac{\mathrm{d}W}{\mathrm{d}R^b}\right|_{R^b=B} = -\frac{\frac{\partial m}{\partial R^b}(1,1)(w_a - w_b)}{1+\frac{\partial m}{\partial B}(1,1)(p_a - p_b)} < 0 \tag{A12}$$

B.4 Proof of Proposition 5

It is proof by example.

Let $m(x,y) = \min\{x,y\}$, $p_a = 1/2$, $p_b = 2/5$, $w_b = 1$, $w_a = 5/4$. Then the pooling equilibrium has $R^b = B = 5/11$ and total welfare is 25/22. The advertising equilibrium has $R^b = 1/4$, $B = 2/5$, and $W = 1$. Note that, while $m(x,y) = \min\{x,y\}$ doesn't fully satisfy the assumptions we placed on the meeting function, its CES approximation does. We find that $m(x,y) = ((1-\alpha)x^\rho + \alpha y^\rho)^{1/\rho}$ with $\alpha = -1/2$ and $\rho = -10$ gets very close to the Leontief example and also produces higher welfare in the pooling equilibrium.

Interestingly, if we set $\rho = 0$ and get a Cobb-Douglas meeting function in the example above, the pooling equilibrium and welfare do not change. However, the advertising equilibrium and the welfare with that allocation do change to $R^b \approx 0.35$, $B \approx 0.457$, and $W \approx 1.1425 > 25/22 = 1.136$. Ultimately, there appears to be a connection between the elasticity of substitution between buyers and sellers in the matching function and the efficiency of the advertising equilibrium.

Figure 12 provides an example economy where welfare is identical in the pooling and the advertising equilibria. The point where welfare is maximized is where the frontier of stationary allocation is tangent to iso-welfare curves. As the diagram plot shows, the welfare-maximizing allocation lies somewhere between pooling and advertising equilibria. For this example, although moving from the pooling equilibrium to the advertising equilibrium is welfare-improving, the actual advertising equilibria offers the same welfare.

Figure 12: Welfare indifference curves and the conditions defining the advertising and the pooling equilibria.