Entry and Advertising on Digital Platforms*

Evidence from a Large E-Commerce Platform

Yizhou Jin and Zhengyun Sun

Dec 31, 2025

Abstract

Millions of businesses join digital platforms each year, yet little is known about their success rates or how platforms manage their entry. Using data from a leading Chinese e-commerce platform, we show that fewer than 1% of new, ID-verified sellers become successful small-to-medium-sized businesses. Analysis of consumer search sessions reveals that consumers prefer entrant products over incumbent ones. Further, increasing entrant exposure on the margin increases overall purchase rate of a search session, indicating that entrants are under-explored compared to social optimum. Finally, we revisit an earlier intervention that increased entrant advertising by 2.3%. We find that it slightly alleviated the under-exploration of entrants, albeit statistically insignificant.

Keywords: digital platform, entry barrier, advertising, consideration set data
JEL codes: L26, M53, D22

*Jin: University of Toronto, yizhou.jin@utoronto.ca; Sun: Uber, zysun92@gmail.com. This paper was previously circulated under the title "AI Training for Online Entrepreneurs." We thank Susan Athey, Jie Bai, Emily Breza, Long Chen, Shawn Cole, Jean-Etienne de Bettignies, Solène Delecourt, Avi Goldfarb, Asim Khwaja, Gabriel Kreindler, Michael Kremer, David McKenzie, Song Ma, Ariel Pakes, Gautam Rao, Yulu Tang, Yuan Tian, Joseph Stiglitz, Christopher Woodruff, Florian Zettelmeyer, and various seminar participants. We also thank the product development team at the unnamed e-commerce platform, and our data contacts and providers for making this project possible. The ability to publish is not contingent on results. Please email carolina_harvey@harvard.edu for questions about the data usage agreements.The trial was registered on the AEA Social Science Registry (AEARCTR-0006725). Human Subjects Approval was obtained from Harvard University (IRB19-0647).

1 Introduction

Millions of businesses join digital platforms each year, which promise vast market access without brick-and-mortar entry costs. However, the large entry volume exacerbates a fundamental information asymmetry problem: platforms do not know the quality of new entrants, leading to a "cold-start" problem (Tadelis 2016). Most studies on this topic focus on the design of reputation and review mechanisms once entrants are explored by consumers, but little is known about how platforms explore entrant quality in practice, or entrants' success rate.[1]

Platforms can explore entrant quality in two main ways. First, they can steer consumers to "sample" new businesses—a bandit exploration strategy. However, it suffers from a free-riding problem as the information generated by one consumer's sampling primarily benefits those who come after (Bolton et al. 2004; Hagiu and Wright 2020a). Moreover, a sophisticated platform can steer consumers toward high-quality established suppliers, further diminishing their incentive to try new products.[2] Platform competition may limit platforms' ability to steer consumers, further diminishing their incentive to promote entry through exploration (Comanor and Wilson 1974). Experiments on e-commerce and service platforms have shown that new suppliers, when exogenously sampled by researchers, saw higher growth and survival rate subsequently (Pallais 2014; Bai et al. 2020).

A second strategy is to rely on voluntary advertising. For example, buying sponsored listings allow entrants to essentially purchase sampling, which is more rewarding for higher quality entrants. In practice, however, established suppliers tend to advertise far more than entrants and, in some cases, see higher returns from advertising (Moshary 2021; Dai et al. 2023; Sahni and Zhang 2024). As many platforms rely on advertising as a primary revenue source,[3] they may also price discriminate all suppliers so that, for an entrant, the cost of advertising exceeds the value of sampling or what is needed to compensate the sampling consumers (Bar-Isaac and Shelegia 2022; Bergemann and Bonatti 2023).

[1] To our knowledge, major platforms rarely disclose the number of new suppliers, their survival rates, or growth trajectories; instead, they primarily share occasional snapshots of the total number of active suppliers. Secondary sources provide some estimates: in 2020, about 1.3 million new sellers joined Amazon Marketplace, 6.8 million joined Taobao and Tmall under Alibaba, and YouTube added 7 million new channels (counting only those with more than five subscribers), according to Marketplace Pulse and Socialblade.

[2] Comanor and Wilson (1974) (chapter 3) formalizes this in a model in which consumers choose between products repeatedly, and the uncertainty about an entrant's product constitutes a disadvantage to the entrant when its products are, on average, expected to be worse than established brands.

[3] For example, Google and Taobao both make over 80% of their revenues from ads in 2022. Since 2017, Amazon Marketplace has seen a 52% annualized increase in ads revenue.

In addition, entrants may suffer information frictions around advertising. To what degree advertising can promote entry is thus an empirical question.

In this study, we obtain administrative data from a major e-commerce platform. We document sellers' entry success rate and develop tests to measure the under-exploration of entrants relative to incumbent sellers. Finally, we revisit a prior experiment that increased entrant advertising frequency and study its effect on entrant exploration.

We find three main results. First, the entry funnel is narrow. In 2019, among the average cohort of potential entrants who completed a rigorous registration process, 51% received traffic from the platform; only 15% became entrants, as defined by having earned revenue. Based on reasonable percentile thresholds defined by the platform, only about 0.5% became SMBs (small- and medium-sized businesses) in the subsequent year. From 2015 to 2019, the platform explored more new sellers but reduced the depth of exploration, reducing the entry funnel, but those that did become SMBs were more likely to persist twelve months after registering.

Second, consumers are more likely to buy from entrants than from incumbent in search sessions. This is inconsistent with substantial bandit exploration so that, on average, an entrant is explored when the expected quality of incumbents are lower. Here, we construct a consideration-set dataset from search session records, each tied to a consumer, a set of search keywords, and a set of products clicked by the consumer. Within consideration sets, we find that consumers are more likely to buy from an entrant than from an incumbent; entrants also see fewer returns and more repeat purchases. Furthermore, comparing nearly identical consideration sets—by conditioning on consumer fixed effects and the interacted fixed effects of the searched keywords, the number of product pages visited, and the search session date—we find that searching an entrant instead of an incumbent is associated with a 3% increase in the consumer's purchase likelihood. Entrants are thus under-explored since producing more information on their quality can expand the market.

Third, to understand how advertising affects entry, we turn to an experiment from a prior study, in which we provided free online training to over two million new sellers on the platform (Jin and Sun 2022). It had precise null effects on nearly all types of seller strategy except advertising (+2.3%), with treated new sellers enjoying a 1.3% increase in traffic. Repeating the consideration-set regressions showed that the quality gap between entrants and incumbents, together with the market expansion effect, are slightly reduced, albeit mostly statistically insignificant. This suggests that enhanced advertising allowed some entrants that are slightly lower quality—yet still superior to the average incumbents—

to reach consumers, mitigating the under-exploration of entrants.

Related Literature This study contributes to the literature on platform entry and advertising. First, we are among the first to document the entry funnel on any major digital platform, and to compare ranking efficiency between entrants and incumbents using search session data. Theoretically, entry has been a central issue in modern platform literature. Earlier work tends to ignore quality uncertainty and focuses on indirect network effects, suggesting that platforms should subsidize entry to attract future consumers—as exemplified by the heavy promotions Uber offered to drivers and riders during its launch. Rochet and Tirole (2003) and Armstrong (2006) discuss how competing platforms may tax suppliers to subsidize consumers when consumers single-home but suppliers multi-home. For example, although Google is free to consumers, the number of advertisers may be lower than the socially optimal level. Recent work tends to recognize that mitigating information asymmetry is a key value proposition of many of today's largest platforms. These platforms play a crucial role in discovering quality and consumer preferences, especially for new products and varieties (Edelman et al. 2007; Athey and Ellison 2011; Tadelis 2016; Hagiu and Wright 2020b). However, advances in data and AI capabilities may further amplify the cold-start problem, raising rather than reducing the barrier to entry (Milgrom and Tadelis 2018). Beyond the issue of entry, the construction and use of our consideration-set data is novel, which allows us to accurately control for unobserved demand heterogeneity, mapping closely to stylized frameworks such as Athey and Ellison (2011).

Second, we provide large-scale experimental evidence that advertising can promote entry and substitute, to a limited degree, for past consumer experience as theorized by Shapiro (1982). Our findings provide an explanation for why prior work on in-platform advertising, which largely focused on its signaling value among established brands, have found mixed results that "may not be true for small and new entities" (Blake et al. 2015).[4] Lastly, our work relates to Li et al. (2020), which studies a reward-for-feedback mechanism on Taobao that "plays a similar signaling role as ads do."

Section 2 describes our setting and the entry funnel. Section 3 provides evidence for the under-exploration of entrants using consideration-set data. Section 4 introduces an experimental variation in entrant advertising and analyze its effect on entrant exploration.

[4] Sahni and Nair (2020a), Long et al. (2022), Sahni and Zhang (2023), and Yu (2024) have found evidence in favor of ads signaling, contrasting with Moshary (2021). They also tended to rely on experiments within established brands or buyer-level randomization that primarily affects those brands.

2 Background: Data and the Entry Funnel

Section 2.1 introduces our e-commerce platform partner and our data. Section 2.2 defines and quantifies the seller *entry funnel* on the platform.

2.1 Platform Business Model and Data

Our research setting is a leading e-commerce platform in China.[5] It hosts millions of active and independent sellers and hundreds of millions of consumers.

Platform business model Most sellers on the platform are third-party retailers. The platform does not directly sell to consumers in almost all sectors, with some notable exceptions such as fresh groceries and financial services. Our analysis excludes these sectors so that the platform only intermediates—and does not compete with—sellers. In this setting, platform revenue primarily comes from in-platform advertising (approximately 80%), with the remaining share made up by sales commissions for large sellers and seller services such as return shipping insurance and advanced data analytics (Jin and Sun 2019).

Almost all traffic to new sellers comes from search: consumers arrive at the platform and search for a desired product; the platform then dissects the search query into keywords and ranks a set of related products and sellers.[6] Sellers can boost their rankings organically or through in-platform advertising. Strategies for boosting organic traffic mainly aim to increase conversion rates and positive reviews: giving discounts or improving product listings as in Lewis (2011), for example.

Sellers can advertise in-platform via sponsored product listings by submitting a daily budget, targeted keywords, and a per-click bid for each keyword and customer segment. In each consumer search session, the platform ranks all bids associated with the search keywords, determines the winning bids, and charges the corresponding sellers. The associated products are then displayed in sponsored positions regardless of their organic search ranking. This is similar to the mechanisms analyzed in Edelman et al. (2007).

Data Our analyses use platform administrative data. The first dataset spans 2015 to 2019, including proportional aggregate statistics on all new sellers on the platform.[7] We describe

[5]Although the platform is based in China, its business model has achieved global success, evidenced by the rapid growth of services such as Temu and Tiktok in North America.

[6]Besides search, some established sellers are matched to consumers via direct platform recommendations. Empirically, however, this is negligible for entrants (Sun et al. 2021).

[7]At the company's request, we do not disclose the raw aggregate levels such as the total number of sellers, which is also not necessary for our funneling analyses.

each statistics in detail below as we introduce the entry funnel.

A second consideration-set dataset is constructed from our prior work Jin and Sun (2022), where we conducted an experiment covering all new sellers on the platform starting May 6, 2019. We describe the data and the experiment in Sections 3.1 and 4.1, respectively.

Compared to prior studies that collaborate with other, often smaller, platforms, a key factor limiting our data access is the rules around algorithmic visibility. Although we observe the actions and outcomes of sellers and buyers in great detail, "behaviors" of the platform's proprietary algorithms—such as product ranking within a search session or the bids in sponsored listings—were classified even among employees. The platform argues that algorithmic transparency would lead to excessive manipulation, or "gaming the system" practices (Petre et al. 2019; Frankel and Kartik 2019; Björkegren et al. 2020).

2.2 The Entry Funnel: Defined and Quantified

Prospective sellers undergo a rigorous verification process to register on the platform. A national ID or a business registry number is required and cross-referenced against the national digital identity database. The platform also collects a small deposit to deter spam. Sellers that successfully registered are considered **potential entrants**. In the following year, those having earned non-zero traffic are considered **explored**; those with non-zero revenue are considered **entrants**. On the intensive margin, we define entrants that are ranked at the top 10% among all sellers by monthly revenue as small-and-medium-sized businesses (**SMBs**).[8] The entry funnel from potential entrants to established brands is illustrated by Figure 1 panel (a).

In 2019, the platform had tens of thousands of potential entrants per *day*. As shown in Table 1, per 10,000 potential entrants, just over half received traffic (were "explored"), and about 15% became entrants. In the first month after registration, only 62 became SMBs, and about two-thirds remained twelve months later. Figure 1 panel (b) redraws the illustrative funnel in panel (a) to reflect the realistic odds.

Over time, the entry funnel appears to have widened at the top: new sellers were more likely to be explored. However, the depth of exploration diminished, as fewer sellers had over 100 unique visitors per month in the year after registering, especially in 2019. Similarly, the lower parts of the funnel grew initially but sharply contracted by 2019. Interest-

[8] This threshold is set by the platform; the corresponding revenue levels are dynamic and proprietary. However, reasonably, they correspond to monthly sales of around one thousand Yuan. More importantly, they highlight the comparison with established sellers: given the large stock of established sellers, we can assume that 10% of them are SMBs.

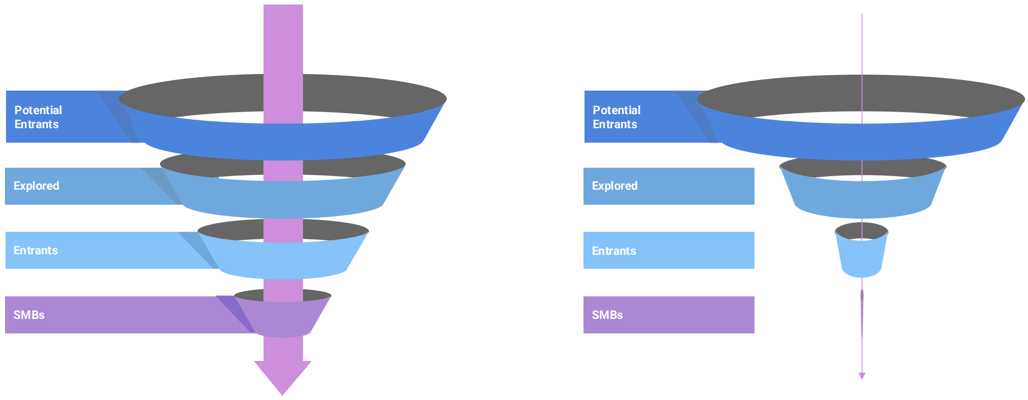

(a) Illustrative Funnel (b) Proportional Funnel

Figure 1: Entry Funnel

Table 1: Entry Funnel from 2015 to 2019

Year	2015	2016	2017	2018	2019
Entry Funnel					
Explored (0+ UV) per 10K P.E.	3246	3917	4536	4552	5069
Explored (100+ UV) per 10K P.E.	2791	3164	3001	2868	2280
Entrants per 10K P.E.	1940	2296	2953	2693	1525
SMB (month 1) per 10K P.E.	87	107	125	124	62
SMB (month 12) per 10K P.E.	59	55	55	51	43
Entry Volume (vs. 2019)					
Total Number of P.E.	1.33x	1.17x	0.86x	0.75x	1x
Number of Explored P.E.	0.85x	0.91x	0.77x	0.67x	1x
Number of Entrants	1.69x	1.76x	1.66x	1.32x	1x
Number of SMBs (month 1)	1.88x	2.04x	1.74x	1.51x	1x
Number of SMBs (month 12)	1.83x	1.49x	1.09x	0.90x	1x

Notes: P.E. standards for "potential entrants" and UV for the average monthly number of "unique visitors." A P.E. is explored if it received more UV than the respective thresholds within a year after its registration. An entrant is a P.E. that received positive revenues during that period. This table is derived based on aggregated administrative data from the platform, with entry volume numbers benchmarked against the 2019 values of the respective rows: for example, in 2015, the total number of P.E. is 1.33x, which means that it is 133% the 2019 total number of P.E.

ingly, although new sellers were least likely to ever become SMBs in 2015 and 2019, those that did become SMBs were more likely to continue thriving at the end of year one. When adjusting for entry volumes, which fell after 2015 but rebounded in 2019, the number of SMBs twelve months after registering bottomed out in 2018.

3 Measuring Entrant Under-Exploration: Consideration-Set Analysis

The narrow entry funnel above suggests that, despite the ease of setting up an online store, potential entrants may be under-explored on the platform, making it difficult to access consumers and grow their revenues. This is especially true given the narrowing of the entry funnel over time, which further disadvantages latter potential entrants compared to their predecessors. Our focus henceforth is thus on the exploration of potential entrants.

Note on terminology For expositional simplicity, we slightly abuse the terminology below, and use "entrants" or "new sellers" interchangeably to refer to "potential entrants."

The exploration of entrant is ultimately achieved through consumer search steered by the platform. In this section, we develop an empirical measure of exploration that illustrates the degree to which new sellers are searched on the platform, and the quality of those sellers, relative to incumbents.

Empirically, however, we face two challenges. First, we have little information on the quality of sellers (or their match value with consumers) unless they are sufficiently searched, yet the consumers that search entrants may be different in unobservable ways from those that search incumbents (lower search cost, for example). Second, we do not observe platforms' search ranking or advertising prices as mentioned in Section 2.1, nor are these sufficient to understand exploration without taking into account the equilibrium effect on consumer search — even if the platform ranks new sellers in top positions, they would still be under-explored if consumers do not click on their listings.

To address the first challenge, we rely on our consideration-set data to effectively control for demand heterogeneity. We do this first by comparing the purchase likelihood of new versus incumbent sellers when they are both present in a consideration set. We then look across almost-identical consideration sets—by controlling for interactive fixed effects of the size of, and the date and search keywords associated with the consideration sets—and compare those that contain new sellers with those that do not.

The second challenge pertains to the interpretation of our results. Our analysis should thus be treated as identifying the state of entrant under-exploration as opposed to structurally teasing out the role of platform steering versus consumers' search behaviors.

3.1 Consideration-Set Data

A Platform-Wide Experiment Our consideration-set data comes from a prior experiment featured in Jin and Sun (2022), in which a random subset of all newly registered sellers on the platform were given access to a digital training program. Our research window is roughly a three-month period starting from May 6, 2019, when the experiment started, until August 15, 2019. This dataset is appropriate in the present study of seller entry at large because we can source a representative subset of all new and incumbent sellers. The experiment also caused an exogenous shock to illustrate whether and how advertising can promote entry, which we study in Section 4. Nonetheless, the scale of the intervention and the magnitude of the treatment effect (detailed in Section 4.1) was small enough that it did not qualitatively alter the entry process.

The Dataset We construct consideration-set data based on records from search sessions during the research window. In each session, we observe an anonymized consumer identifier and the keywords parsed from the search query. The *consideration set* consists of all products the consumer had clicked on, each corresponding to one seller. We also observe the product the consumer bought, if any. For a given consumer, we aggregate search sessions with the same keyword combination that occurred on the same calendar day.

On the transaction level, we are able to observe the spending amount. We also observe subsequent requests for returns and refunds issued. By linking consumer and seller IDs, we can further track repeat purchases within the next 30-day window post transaction.

On the seller level, we observe some time-invariant characteristics, including the registration date (which categorizes the seller as new or incumbent), treatment assignment, location based on mailing address, and some basic characteristics of the owner. We also observe a set of time-varying variables describing sellers' actions or states, including the number of product listings, and a composite seller rating across products. To account for heterogeneity across product categories, we develop "sector" dummies based on sellers' product portfolios.[9]

[9] Products are categorized into sectors using the platform's natural-language-processing algorithm based on seller-chosen labels and product descriptions. A seller's sector is based on its top-selling product, which can vary over time. Sellers without listings are put in the "unknown" category.

Sampling Procedures Computationally, we face two challenges that require sub-sampling. First is the sheer volume of data: there were over 10 million visits per day to the new sellers in our experiments alone. Second is the skewness toward incumbent sellers and popular keywords: new sellers appeared in less than 0.1% of consideration sets.

We start with consumers who visited new sellers during the research window, and randomly subsample to form a core consumer sample (CCS). We then pull the calendar date and search query (parsed into keyword combinations) associated with each session conducted by these consumers. We also observe the set of all product pages clicked into, and, if the search session ended in a purchase, the transaction ID. This allows us to construct a consideration set associated with each search session.

To address the skewness of keywords, we filter down to search sessions that contain popular keyword combinations, defined as those that were searched at least three times by consumers in the CCS throughout our research window. Lastly, focusing on the CCS means leaving out a large group of consumers who did not visit new sellers at all during our research window. For each consideration set, we thus pull additional sets conducted by non-CCS consumers, but had the same keyword combinations parsed from the search queries, the same size (number of products visited), and on the same date.[10] This allows us to compare between similar consideration sets using a regression with keyword-size-date interacted fixed effects.

A separate challenge is that many search sessions may not be genuine, especially those that contain only one or two sellers. They make up over 42% of all sessions during our research window yet are only responsible for 1.9% of purchases. We thus restrict our analyses to those with at least three products. This allows us to manage the data volume and to better focus on genuine consumer searches.[11]

Summary Statistics Our dataset contains 6.0 million consideration sets (Table 2), each consisting of 4.6 sellers on average. A total of 65,246 distinct keywords were searched across all major sectors. Despite over-sampling, 82% consideration sets still had no new sellers. These are omitted when analyzing differential seller purchase likelihoods within consideration sets in Section 3.2 ($N = 1,084,228$); we introduce them back in when analyz-

[10]We include up to 50 consideration sets per keyword-size-date combination for computational ease, and we randomly subsample if the number of matching sets exceeds 50.

[11]In particular, we suffer from data contamination from "passive searches" as studied in Ursu et al. (2022), where many consumers were directed to a product page via external links that can be disguised as a search session of size one. In addition, consumers often try different queries before or as they start searching; new consumers also learn about the platform by conducting test searches. These search sessions are typically very small and have close to zero conversion rates.

ing differential outcomes across consideration sets in Section 3.3. There, when conducting our robustness check on returns and repeat purchases, we focus on consideration sets that result in a purchase ($N = 289,258$).

Consumers who visited new sellers are observably different. In particular, they seem to search more sellers regardless of treatment status (Figure 2). This illustrates how demand heterogeneity (such as consumer search cost or purchase intent) can bias quality inference using raw purchase or conversion rates.

Table 2: Summary Statistics: Consideration Set Sample

Statistic	Mean	St. Dev.	Min	Pctl(25)	Pctl(99)	Max
Number of Sellers Browsed	4.6	2.7	3	3	12	55
% of New Sellers	0.57	8.7	0	0	33	100
Purchased (on the same day)	0.180	0.167	0	0	0	1
Purchased from New Sellers	0.001	0.033	0	0	0	1
Pay Amount (USD)	4.8	21.2	0	0	753	2,050
Avg. Seller Price Level (USD)	34.4	86.3	0.0	8.8	582	9,674
Avg. Seller Rating	13.0	2.9	0	11.5	20	20

Notes: $N = 6,007,329$. Unit of analysis is a consideration set. Dollar amount converted to USD based on the average 2019 currency exchange rate of 1 RMB to 0.1448 USD.

Figure 2: Distribution of the Size of Consideration Sets by Seller Type

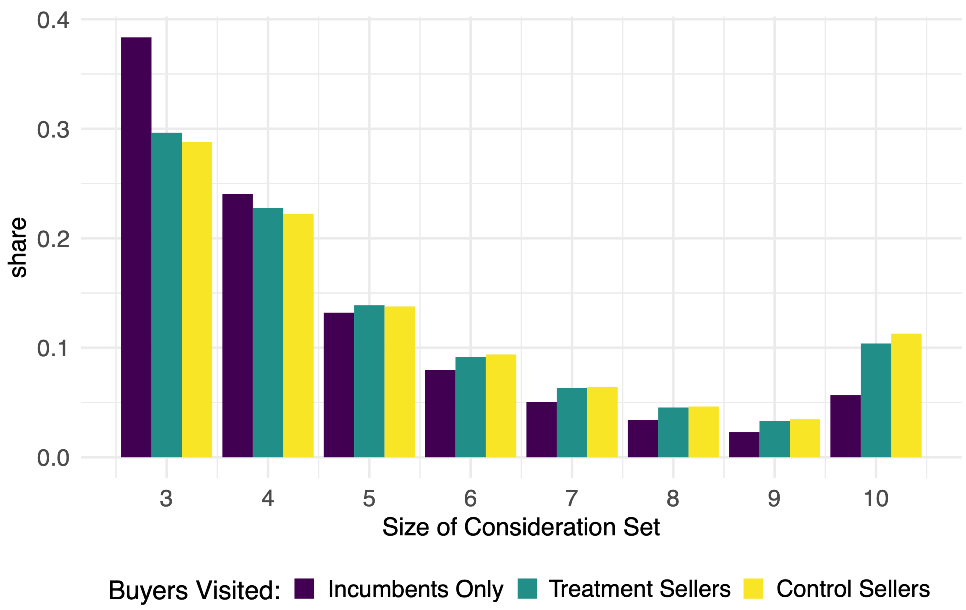

Notes: Density histograms of the size of consideration sets. "Incumbent Only" includes sets that consist only of incumbent sellers. The other groups include sets with at least one control-group or treatment-group seller. We focus on consideration sets containing at least three products for reasons described above.

3.2 Within Consideration-Set Analysis: New vs. Incumbent Sellers

We now conduct a straightforward regression to illustrate that, within a consideration set that contains both new and incumbent sellers, buyers are more likely to purchase from new sellers, leading to a higher average conversion rates. We achieve this by including consideration-set fixed effects while filtering the data to consideration sets that contain products from at least one new and one incumbent sellers.

$$Y_{js} = \beta \mathbf{New}_j + \alpha_s + \varepsilon_{js} \tag{1}$$

The outcome variable is an indicator for a consideration set s, or a search session, resulting in a purchase from seller j. The coefficient of interest, β, measures the differences in purchase likelihood between new sellers, defined as those that registered during our research window, and incumbents. α_s, the consideration-set fixed effects, soak up any unobserved differences across consumers and across each consumer's search sessions.

Table 3: Regression Results: Purchase Conditional on Search

	Has Purchased		Log Spending	
	(1)	(2)	(3)	(4)
Is New Seller	0.013***	0.003***	0.056***	0.010***
	(0.001)	(0.001)	(0.004)	(0.004)
Incumbent Mean	0.17		0.71	
Seller Controls	N	Y	N	Y
Consideration-set Fixed Effects	Y	Y	Y	Y
R^2	0.053	0.073	0.054	0.078
Adj. R^2	0.036	0.057	0.037	0.062

Note: $N = 1,084,228$ consideration sets after filtering down to those that contained at least one new seller. Reports estimates from Equation 1. All variables are described in the paragraphs above, with columns corresponds to dependent variables (labeled at the top) and model specifications (labeled at the bottom). * p<.1, ** p<.05, *** p<.01.

The estimates are summarized in Table 3: surprisingly, new sellers had significantly higher purchasing probability than incumbents, 1.3 percentage point. This is inconsistent with substantial bandit exploration so that, on average, an entrant is explored only when the expected quality of incumbents are even lower.

3.3 Across Consideration Sets Analyses: with vs. without New Sellers

The analysis above compares seller outcome with the same consideration sets, which allow us to control for unobserved demand heterogeneity to the best of our ability. However, doing so filters down to consideration sets that have both new and incumbent sellers. In this section, we remove this filter and compare across *similar* consideration sets that may or may not contain new sellers.

Robustness check: refunds and repeat purchase First, focusing on all consideration sets that resulted in purchases, we can conduct a robustness check of whether the higher purchase probability among new sellers is accompanied by a greater frequency of refunds or a lower frequency of repeat purchases.

Since our goal is to compare similar consideration sets that result in purchases from new versus incumbent sellers, the main independent variable is an indicator of whether a consideration set contains new sellers, $\mathbf{New}_s$. In place of the consideration-set fixed effects (α_s in Equation 1), we include consumer fixed effects $\alpha_i(s)$ and search *keyword-size-date* interacted fixed effects IFE_s, as introduced in Section 3.1. We also average over observable characteristics of all sellers in the consideration set to obtain $\mathbb{X}_s$—number of products offered, average price, consumer star rating, and average order size—and further control for differences across search sessions. The dependent variable is an indicator for return or refund request, or an indicator for the consumer making repeat purchases from the seller from which the original purchase was made.

$$Y_s = \beta \mathbf{New}_s + \mathbb{X}_s \gamma + \alpha_{i(s)} + IFE_s + \varepsilon_s \tag{2}$$

Table 4 report the estimates. Across almost identical search sessions, buying from new sellers is associated with lower likelihood of returns and increased likelihood of repeat purchases. These results simultaneously illustrate the robustness of our prior results, that new sellers offer higher match quality conditional on search, and illustrate how exploring new sellers can be particularly fruitful for repeat purchases.

Market expansion effects Lastly, we relax all data filters and include a majority of consideration sets that neither contains new sellers nor resulted in a purchase. Doing so capitalizes on a novel property of our data: the ability to observe the *potential* market—defined as all consumers that searched at least one seller regardless of whether they made a purchase (from any seller) or not—and thus potential "market expansion effects," i.e. the degree to which marginally increasing the search likelihood of entrants can increase the

Table 4: Regression Results: Post-Purchase Experience

	Dependent variable:			
	Is Returned		Has Repeat Purchased	
	(1)	(2)	(3)	(4)
Has New Seller	-0.006***	-0.004***	0.006***	0.004**
	(0.002)	(0.001)	(0.002)	(0.002)
Incumbent Mean	0.078		0.064	
Seller Controls	N	Y	N	Y
Consumer Fixed Effects	Y	Y	Y	Y
Intent-Effort-Date Fixed Effects	Y	Y	Y	Y
R^2	0.398	0.463	0.301	0.378
Adj. R^2	0.181	0.257	0.050	0.082

Note: $N = 289,259$ consideration sets. We first filter down to those that resulted in a purchase; we then filter down to purchases in which the seller appeared at least one more time in the consumer's future consideration sets. The table reports estimates from Equation 2. All variables are described in the paragraphs above, with columns corresponds to dependent variables (labeled at the top) and model specifications (labeled at the bottom). * p<.1, ** p<.05, *** p<.01.

amount of matches.

The regression specification is similar to Equation 2, with a new sample and variable definitions:

$$Y_s = \beta \mathbf{n}_s + \mathbb{X}_s \gamma + \alpha_{i(s)} + IFE_s + \varepsilon_s \tag{3}$$

We can now use our full sample, including consideration sets that did not result in purchases. The outcome variable is an indicator for the consideration set s resulting in a purchase. The main independent variable is the number of new sellers contained in the search session, $\mathbf{n}_s$. We again include consumer fixed effects, the keyword-size-date interacted fixed effects, and average over observables of sellers in the search session.

Table 5 shows that across similar consideration sets, searching new sellers instead of incumbents increased purchase probability and dollar spent. This implies that marginally increasing search among entrants can lead to market expansion, further confirming the existence of an informational entry barrier in the search process.

Discussion: Entrant Under-Exploration Taken together, our analysis in this section suggest that, in equilibrium, consumers are steered to search entrants when the expected (match) quality of incumbent products is lower than the that of the entrants'. This could happen, for example, in a search session where the consumer searches a set of keywords

Table 5: Regression Results: Market Expansion Effects

	Dependent variable:			
	Has Purchased		Log Spending	
	(1)	(2)	(3)	(4)
Number of New Sellers	0.003***	0.002***	0.011***	0.008***
	(0.001)	(0.001)	(0.003)	(0.003)
Incumbent Mean	0.041		0.17	
Seller Controls	N	Y	N	Y
Consumer Fixed Effects	Y	Y	Y	Y
Intent-Effort-Date Fixed Effects	Y	Y	Y	Y
R^2	0.053	0.073	0.054	0.078
Adj. R^2	0.036	0.057	0.038	0.062

Note: $N = 6,007,329$ consideration sets. The table reports estimates from Equation 3. All variables are described in the paragraphs above, with columns corresponds to dependent variables (labeled at the top) and model specifications (labeled at the bottom). * p<.1, ** p<.05, *** p<.01.

that only matches the product labels from a new seller, despite that seller having no prior transaction or review; or if the incumbent products matching the keywords have received very low ratings historically. Meanwhile, steering consumers to more entrants would expand the market on the margin, implying that some incumbent products were searched over higher-quality ones from entrants. This indicates that entrants on the platform are under-explored compared to social optimum, as producing more information on their quality can expand the market.

Inferring platform conduct is harder without directly observing ranking information. Assuming that consumers search sequentially based on platform ranking, as in canonical frameworks such as Athey and Ellison (2011),[12] our results in this section would indicate that in the average search session, the platform does *not* conduct bandit exploration, by up-ranking entrants with unknown quality for example, to gather information about those entrants' product quality.

[12] This is not without loss of generality. It is possible that the platform conducts bandit exploration by up-ranking new sellers, but consumers skip them in their search, perhaps due to a lack of prior reviews.

4 Can Advertising Promote Entry?

Advertising in the form of sponsored listings can allow new sellers to pay for consumer search and potential reviews, hence overcoming the cold-start problem. For example, there are over one hundred thousand new and incumbent sellers offering sweaters on the platform. Our results above suggest that new sellers offering products with common labels, such as "Christmas" or "cashmere", or "blue", would be unlikely to reach consumers organically. The sellers can, however, bid for those keywords. To the extent that those with particularly compelling and relevant products (and thus high match probability) are willing to bid more, advertising can become a costly signal of match quality.

However, how advertising affects entry is ultimately an empirical question. The costly nature of advertising can exclude liquidity-constraint entrants while compromising its signaling value, especially if the platform fosters an advertising arms race among entrants or with incumbents. Specifically, many regulators have expressed concern that platforms may "sell 'monopoly positions' to sellers by showing buyers alternatives which do not meet their needs" (Crémer et al. 2019).

4.1 A Training Experiment that Increased New-Seller Advertising

In Jin and Sun (2022), we studied a digital training program targeting new sellers on the platform. From May 2019 to June 2020, a subset of all new sellers received access to the program. Among potential entrants, trained ones saw 1.7% higher revenues, almost entirely driven by increased advertising: 2.3% on the extensive margin and 3.0% overall. Traffic from ads increased by 7.6% during the first three-month after registration, which is when the vast majority of training occurred, and 4% over the the full 9-month period.

We refer readers to the original paper for more details, but we highlight three important points. First, in terms of seller strategies, a key result is a null one: despite a broad coverage of business strategies, training had precise zero effect on sellers' product variety, pricing, and ratings.[13]

Second, although the experiment generated statistically significant and economically meaningful impact on entrants, the overall effect on entrant revenues and traffic was

[13]There is a notable exception besides advertising. Trained new sellers were more likely to adopt short-video demos. As studied in Lewis (2011), better product descriptions can reduce asymmetric information on e-commerce platforms and increase matching efficiency. However, new sellers saw no corresponding increase in organic traffic during training, i.e. within the first three months after registration. The magnitude of the increase was also smaller than that of advertising (1.9%).

marginal. This is evident with our aggregate result shown in Table 1, where the entry funnel in 2019 narrowed compared to prior years despite the experimental traffic and revenue increase we observed among treated entrants.

Third, recent studies have questioned the validity of seller-level randomization on digital platforms due to potential SUTVA violations. However, our treatment affected entrants only, whom, despite their large quantity, had little impact on product or advertising equilibrium. Treated new sellers can reasonably be assumed to not have directly competed with their control-group counterparts.

4.2 Are Entrants Still Under-Explored?

Although more advertising raised entrant traffic, it remains unclear what kind of entrants increased advertising and the degree to which that had mitigated the under-exploration of entrants or improved the ability for high-quality ones to reach consumers.

We now replicate the regression analyses in Section 3 but focus on testing the differences across the treatment and control groups. Tables 6 and 7 map to Tables 3 and 4, respectively. Directionally, they suggest that the quality of entrants searched by consumers became slightly worsened with more advertising in the treatment group. The magnitude of the differential, however, is very small and statistically insignificant (except a higher return frequency), meaning that the entrants in the treated group remain substantially superior (in terms of purchase probability and thus match quality) compared to incumbents.

Table 8, which maps to Table 5, tells a similar story: the market expansion effect (of marginally increasing entrant exposure) is still prominent in the treatment group, albeit slightly—and statistically insignificantly—weakened.

Table 6: Regression Results: Purchase Conditional on Search (Treatment - Control T test)

	Has Purchased		Log Spending	
	(1)	(2)	(3)	(4)
Is New Seller	-0.0016	-0.00081	-0.0074	-0.0039
(Treatment - Control T test)	(0.0013)	(0.0013)	(0.0060)	(0.0059)
Seller Controls	N	Y	N	Y
Consideration-set Fixed Effects	Y	Y	Y	Y

Note: $N = 1,084,228$. See notes from Table 3 for more details. * p<.1, ** p<.05, *** p<.01.

Table 7: Regression Results: Post-Purchase Experience (Treatment - Control T test)

	Dependent variable:			
	Is Returned		Has Repeat Purchased	
	(1)	(2)	(3)	(4)
Has New Seller (Treatment - Control T test)	0.0054** (0.0026)	0.0026* (0.0016)	0.0024 (0.0024)	0.0019 (0.0015)
Seller Controls	N	Y	N	Y
Consumer and Keyword-Size-Date Fixed Effects	Y	Y	Y	Y

Note: $N = 289,259$. See notes from Table 4 for more details. * p<.1, ** p<.05, *** p<.01.

Table 8: Regression Results: Market Expansion Effects

	Dependent variable:			
	Has Purchased		Log Spending	
	(1)	(2)	(3)	(4)
Number of New Sellers	-0.00046 (0.00044)	-0.00023 (0.00044)	-0.0025 (0.0020)	-0.0018 (0.0020)
Seller Controls	N	Y	N	Y
Consumer and Keyword-Size-Date Fixed Effects	Y	Y	Y	Y

Note: $N = 6,007,329$. See Table 5 for more details. * p<.1, ** p<.05, *** p<.01.

Taken together, our results in this section suggest that the enhanced advertising likely played a small role in equalizing the differentials between entrants and incumbents: it allowed some entrants that are slightly lower quality—yet still superior to the average incumbents—to reach consumers.

Discussion Ultimately, advertising among entrants was rare regardless of our treatment. Among the new seller cohort, only 7% will have ever conducted advertising in their first year of operation. A 2.3% experimental increase, though statistically significant and economically important as it translates into a 1.3% increase in traffic, is unlikely to address the structural under-exploration of entrants in aggregate. Another way to interpret our result is that, should the platform wish to rely on advertising to generate a socially optimal amount of entrant quality information, a much larger increase in entrant advertising (than what's observed in the experiment) is needed.

5 Conclusions

One of the key promises of digital platforms is their ability to lower entry barriers and decentralize market access. To our knowledge, the present study provides the first documentation of a narrow entry funnel on a dominant digital platform. We also provide evidence for a systematic under-exploration of entrants, in the sense that the provision of more entrant quality information would be socially desirable and expand the market. This simultaneously demonstrates the empirical bite of the classic "cold-start" problem (Tadelis 2016). Lastly, as many digital platforms are ad-supported, we provide experimental evidence that a marginal increase in advertising, while powerful in raising the traffic of individual entrants, had little effect on the structural under-exploration of entrants.

References

Armstrong, Mark (2006). "Competition in two-sided markets". In: *The RAND journal of economics* 37.3, pp. 668–691.

Athey, Susan and Glenn Ellison (2011). "Position auctions with consumer search". In: *The Quarterly Journal of Economics* 126.3, pp. 1213–1270.

Bai, Jie et al. (2020). "Search and Information Frictions on Global E-Commerce Platforms: Evidence from AliExpress". In.

Bar-Isaac, Heski and Sandro Shelegia (2022). *Monetizing steering*. Centre for Economic Policy Research.

Bergemann, Dirk and Alessandro Bonatti (2023). "Data, competition, and digital platforms". In: *arXiv preprint arXiv:2304.07653*.

Björkegren, Daniel et al. (2020). "Manipulation-proof machine learning". In: *arXiv preprint arXiv:2004.03865*.

Blake, Thomas et al. (2015). "Consumer heterogeneity and paid search effectiveness: A large-scale field experiment". In: *Econometrica* 83.1, pp. 155–174.

Bolton, Gary E et al. (2004). "How effective are electronic reputation mechanisms? An experimental investigation". In: *Management science* 50.11, pp. 1587–1602.

Comanor, William S and Thomas A Wilson (1974). *Advertising and market power*. 144. Harvard University Press.

Crémer, J et al. (2019). "Competition Policy for the Digital Era, European Commission". In: *Publications Office of the European Union*.

Dai, Weijia et al. (2023). "Frontiers: Which Firms Gain from Digital Advertising? Evidence from a Field Experiment". In: *Marketing Science* 42.3, pp. 429–439.

Edelman, Benjamin et al. (2007). "Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords". In: *American economic review* 97.1, pp. 242–259.

Frankel, Alex and Navin Kartik (2019). "Muddled information". In: *Journal of Political Economy* 127.4, pp. 1739–1776.

Hagiu, Andrei and Julian Wright (2020a). "Platforms and the exploration of new products". In: *Management Science* 66.4, pp. 1527–1543.

— (2020b). "Platforms and the exploration of new products". In: *Management Science* 66.4, pp. 1527–1543.

Jin, Yizhou and Zhengyun Sun (2019). “Information Acquisition and the Return to Data”. In: *Harvard University Working Paper*.

— (2022). “AI Training for Online Entrepreneurs: An Experiment with Two Million New Sellers on an E-commerce Platform”. In: *Available at SSRN 4334216*.

Lewis, Gregory (2011). “Asymmetric information, adverse selection and online disclosure: The case of eBay motors”. In: *American Economic Review* 101.4, pp. 1535–1546.

Li, Lingfang et al. (2020). “Buying reputation as a signal of quality: Evidence from an online marketplace”. In: *The RAND Journal of Economics* 51.4, pp. 965–988.

Long, Fei et al. (2022). “Designing an online retail marketplace: Leveraging information from sponsored advertising”. In: *Marketing Science* 41.1, pp. 115–138.

Milgrom, Paul R and Steven Tadelis (2018). *How artificial intelligence and machine learning can impact market design*. Tech. rep. National Bureau of Economic Research Cambridge, MA.

Moshary, Sarah (2021). “Advertising Effects in Equilibrium”. In.

Nocke, Volker and Patrick Rey (2021). *Consumer Search and Choice Overload*. Centre for Economic Policy Research.

Pallais, Amanda (2014). “Inefficient hiring in entry-level labor markets”. In: *American Economic Review* 104.11, pp. 3565–3599.

Petre, Caitlin et al. (2019). “Gaming the system: Platform paternalism and the politics of algorithmic visibility”. In: *Social Media+ Society* 5.4, p. 2056305119879995.

Rochet, Jean-Charles and Jean Tirole (2003). “Platform competition in two-sided markets”. In: *Journal of the european economic association* 1.4, pp. 990–1029.

Sahni, Navdeep S and Harikesh S Nair (2020a). “Does Advertising Serve as a Signal? Evidence from a Field Experiment in Mobile Search”. In: *The Review of Economic Studies* 87.3, pp. 1529–1564.

— (2020b). “Sponsorship disclosure and consumer deception: Experimental evidence from native advertising in mobile search”. In: *Marketing Science* 39.1, pp. 5–32.

Sahni, Navdeep S and Charles Zhang (2023). “Are consumers averse to sponsored messages? The role of search advertising in information discovery”. In: *Quantitative Marketing and Economics*, pp. 1–52.

— (2024). “Are consumers averse to sponsored messages? The role of search advertising in information discovery”. In: *Quantitative Marketing and Economics* 22.1, pp. 63–114.

Shapiro, Carl (1982). *Advertising as a Barrier to Entry?* US Federal Trade Commission, Bureau of Economics.

Sun, Tianshu et al. (2021). “The value of personal data in internet commerce: A high-stake field experiment on data regulation policy”. In: *Available at SSRN 3962157*.

Tadelis, Steven (2016). “Reputation and feedback systems in online platform markets”. In: *Annual Review of Economics* 8, pp. 321–340.

Train, Kenneth E (2009). *Discrete choice methods with simulation*. Cambridge University Press.

Ursu, Raluca et al. (2022). “Online advertising as passive search”. In: *NYU Stern School of Business Forthcoming, Columbia Business School Research Paper*.

Yu, Chuan (2024). “The Welfare Effects of Sponsored Product Advertising”. In: *Available at SSRN 4817542*.

A Training Details [For Online Publication]

A.1 Training Take-up Dynamics

Figure A1: Training Take-up Dynamics

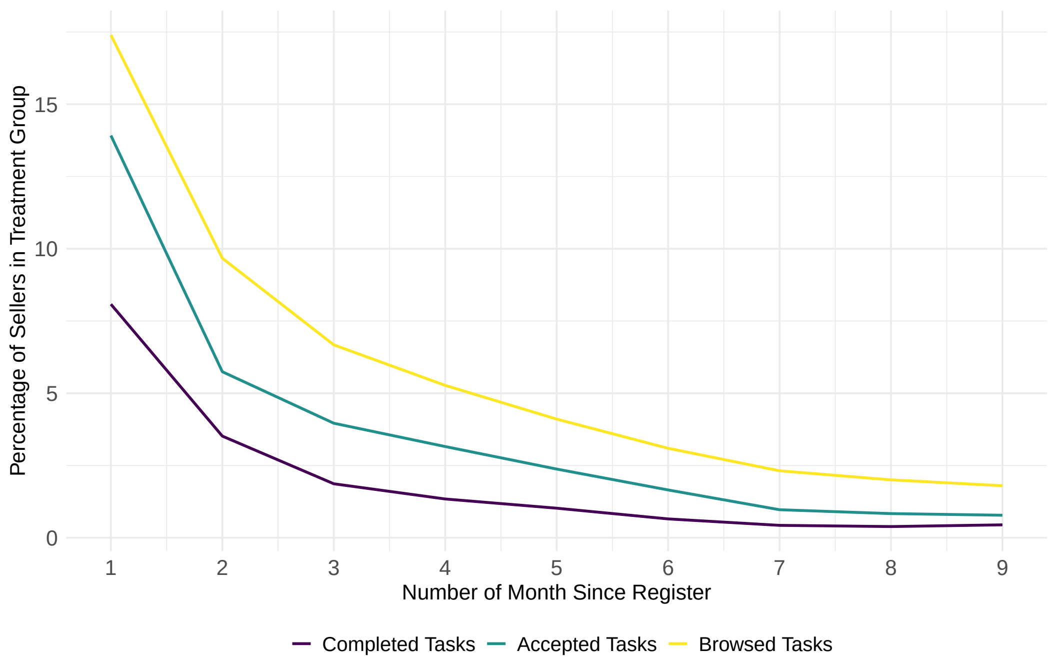

Although sellers have access to the training for at least six months after the registration, the retention rate declined relatively quickly over time. 13.9% of treated sellers took up tasks during the first month, but only 4.0% of sellers continued to do so in the third month.[14] Conditional on taking up any tasks in the previous month, around 23% of sellers took up more tasks in the following month. Figure A1 shows the share of sellers who browsed, took up and completed the tasks in the subsequent months following the entry.

[14]Many sellers exit the platform after a month. 89.7% of sellers have visitors during the first month, but only 42.2% still do during the third month. The rates are similar for sellers in the treatment and control group.

A.2 Treatment Effect Dynamics

To explore variations of the treatment effect for sellers at different stages, we estimate the following specification on the balanced sample:

$$Y_{imcs} = \sum_{m=1}^{9} \beta^m Treatment_i M_{im} + \alpha_m + \alpha_c + \alpha_s + \varepsilon_{imcs} \tag{4}$$

where M_{im} is the set of indicators for month $m = 1, ..., 9$. We focus on the set of coefficient $\{\beta^m\}_{m=1}^{9}$ that captures the effect of having access to the training during a particular month m since entering the platform. The specification again controls for month of entry α_m, registered cohort α_c and initial sector affiliation α_s fixed effects. Standard errors are clustered at seller level.

A.3 Heterogeneous Treatment Effects

Is the training particularly effective for certain types of new sellers? We examine the heterogeneous impact of the training by sellers' characteristics in the baseline. Because the content of the training mainly targets basic operations and marketing, we expect sellers with limited previous exposure to e-commerce to benefit more from the training, as it helps to close their knowledge gap. We characterize sellers from the following dimensions: registration type, gender (if registered as individuals), location, whether any products were posted on the first day and whether the store is registered as a B2C store on the first day[15]. Since no sellers took up tasks on the first day of entry, we consider listing products and sellers' registration type on the first day as part of pre-treatment characteristics. We then estimate the following specification:

$$Y_{imcs} = \beta Treatment_i + \gamma Type_i + \delta Treat_i \times Type_i + \alpha_m + \alpha_c + \alpha_s + \varepsilon_{imcs} \tag{5}$$

where as before $Treat_i$ denotes assignment to treatment or control group and $Type_i$ specify whether sellers have characteristics aforementioned. δ captures the heterogeneous treatment effects on different types of sellers. Table A1 summarizes the results on log revenues

[15]To register as a B2C store, potential sellers must obtain formal approval from the platform. The minimum requirements include having a brand name and a formally registered firm. 97% of sellers in the sample are registered as C2C stores. C2C stores can be converted to B2C stores later on. Among all sellers that eventually become B2C stores, 66.7% of them converted later on, and sellers in the treatment group are more like to convert.

by seller types. Overall, estimated β in these specifications have similar magnitudes as estimates using equation **??**, but there are no differential treatment effects by seller types, gender, actions on the first day of entry on revenues. The slightly surprising result is that there is no differential treatment effect for sellers with different preparedness levels. Compared to the rest, sellers who post products on the first day of entry could be better prepared or more experienced. Hence, these sellers might find the basic part of the training less useful, yet we do not find such results[16].

To evaluate the impact of the offline business environment, we group the sellers based on their registered locations[17]. Table A2 presents the results on heterogeneous treatment effect by seller locations on log monthly revenues. γ captures sellers' average performance in different parts of the country compared to those coming from the remaining parts. There are significant variations in average performance for sellers from different parts of the country. Sellers from the southern coastal provinces significantly outperform the rest, while those coming from the less-developed western part lagged. The performance of sellers in different regions is consistent with economic development in the offline world. Training is less helpful for sellers from less developed regions, as these sellers are less likely to take-up the training (see table E2)[18]. Therefore, even the training program offers the same materials to all sellers, sellers coming from less pro-business areas are less likely to take advantage of such knowledge. As a result, the training does not help those lagging behind to catch up, but instead further strengthens the competitive edge of new sellers from more developed regions.

[16]We also do not find training to be more useful for sellers who post products after the first day of entry.

[17]The location information on ID cards for individual sellers may not reflect where the sellers reside at the moment because the location indicates ID card holder's birthplace rather than current residence. The internal migration patterns imply that we are under-counting sellers living in the coastal provinces, as these provinces are major destinations of migration. Similarly, the firm's registered locations might not be the same as where the firms operate, but in this case, the direction of the bias is unclear.

[18]Training participants are significantly less likely to be sellers from the western provinces.

Table A1: Treatment Effect Heterogeneity on Sellers' Basic Types

	Dependent variable: Log Revenues				
	Registration Type		Post Products		B2C Sellers
	Female	Firm	First Day	Later Days	
	(1)	(2)	(3)	(4)	(5)
Treatment	0.016*	0.009	0.014**	0.011*	0.016**
	(0.008)	(0.006)	(0.007)	(0.006)	(0.006)
Seller Type	-0.494***	1.504***	-0.130***	1.265***	3.842***
	(0.006)	(0.009)	(0.008)	(0.008)	(0.050)
Treatment × Seller Type	0.003	0.015	0.013	0.014	0.077
	(0.012)	(0.017)	(0.015)	(0.013)	(0.098)
Observations	6,409,062	6,409,062	6,409,062	6,409,062	6,409,062
R^2	0.138	0.174	0.132	0.163	0.147
Adjusted R^2	0.138	0.174	0.132	0.163	0.147

Notes: Standard errors clustered by seller. All regressions include month, cohort and initial sector fixed effect. Dependent variable is monthly revenues in log scale after adding one to base level. The interaction variables are indicators for whether or not sellers are females, are registered as firms, post products on the very first day of entry or during some later days and lastly whether or not sellers register as B2C sellers. * significant at 10% level, ** significant at 5% level and *** significant at 1% level.

Table A2: Treatment Effect Heterogeneity by Sellers' Registered Location

	Dependent variable: Log Revenues					
	Beijing Vicinity	Resource-Oriented	Northeast	Coastal South	Central	West
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.015**	0.018***	0.019***	0.014*	0.015**	0.019***
	(0.007)	(0.006)	(0.006)	(0.007)	(0.007)	(0.007)
Location	-0.039***	-0.371***	-0.273***	0.470***	-0.128***	-0.642***
	(0.008)	(0.013)	(0.012)	(0.006)	(0.007)	(0.007)
Treatment × Location	0.011	-0.042	-0.048*	0.006	0.008	-0.013
	(0.017)	(0.026)	(0.025)	(0.013)	(0.014)	(0.014)
Observations	6,409,062	6,409,062	6,409,062	6,409,062	6,409,062	6,409,062
R^2	0.132	0.132	0.132	0.138	0.132	0.136
Adjusted R^2	0.132	0.132	0.132	0.138	0.132	0.136

Notes: Standard errors clustered by seller. All regressions include month, entry date and main industry fixed effect. Dependent variable is monthly revenues in log scale. Indicators are sellers registration locations clustered into different regions. Beijing Vicinity includes Beijing, Tianjin, Heibei and Shandong; resource-oriented provinces include Shanxi, Neimenggu, Gansu and Ningxia; northeastern provinces are Heilongjiang, Jilin and Liaoning; coastal southern provinces are Jiangsu, Shanghai, Zhengjiang, Fujian, Guangdong and Hainan; central provinces are Anhui, Jiangxi, Henan, Hubei and Hunan; western provinces are Tibet, Xinjiang, Yunnan, Guangxi, Sichuan, Chongqing, Guizhou, Shaanxi and Qinghai. * significant at 10% level, ** significant at 5% level and *** significant at 1% level.

Figure A2: Quantile Treatment Effect on Traffic Over Time

Figure A3: Quantile Treatment Effect on Revenue Over Time

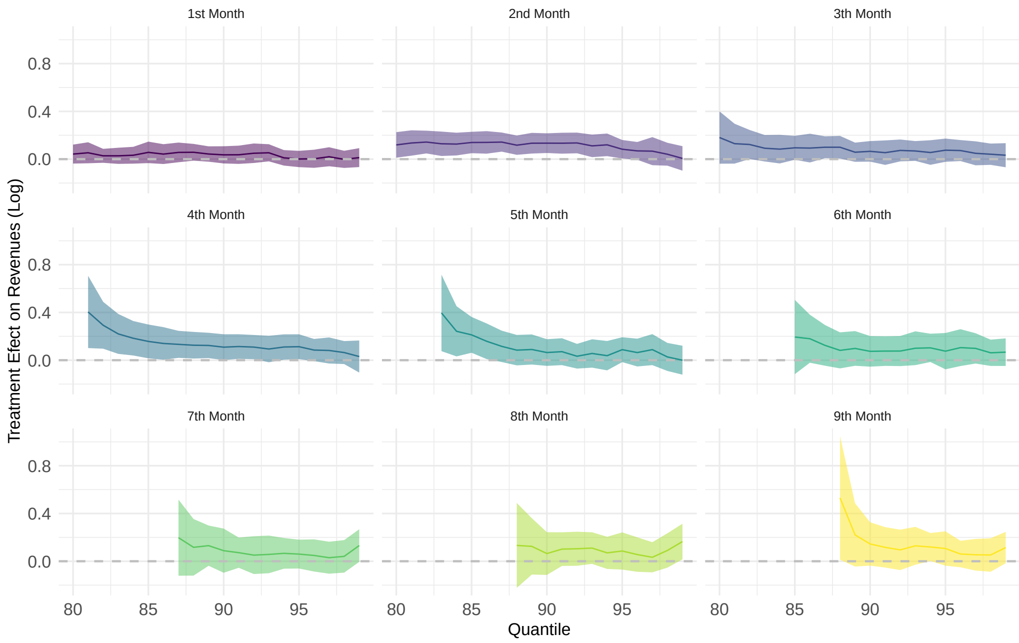

Notes: These figures plot the monthly quantile treatment effect of training. We focus on sellers in the 80th to 100th percentile as only 19% of all sellers have ever earned revenue (Table **??**). Each panel plots the results for a particular month after treatment for our analysis cohort.

B Online Seller Survey[For Online Publication]

We conducted an online survey in August 2019 with sellers to gather some basic demographics information and their opinion about the training. The sampling was stratified by sellers' engagement with the training, and we over-sampled sellers who were more involved with the training, i.e. sellers who took up more tasks. In the end we collected 566 responses. Detailed results are presented in table B1. Since most of the respondents are training participants, they may not form a representative sample of sellers on the platform. These respondents are likely to be more active and have higher sales. Moreover, compared to anecdotal descriptions of typical sellers, these sellers appear to have higher than average ownership of manufacturing factories (32.5%) and offline stores (19.2%).

The survey shows that even among the training participants, sellers differ in terms of their background, experience, education and financial resources. However, while the vast majority of the active new sellers is small and inexperienced, a substantial share of them is reasonably educated and express clear interests to participate in e-commerce. Results from the online survey show that 71.9% have 1 or 2 employees, 74.3% have no or less than one year of experience in e-commerce, and 68.2% have completed at least high school education. About 58.8% of sellers in the sample report that they intend to make running the e-commerce store as their main job and 48.2% have invested more than 10,000 RMB ($1430) into their online businesses. The platform does not have a systematical approach to collecting demographic data from the sellers other than those collected during registration[19].

[19] We could potentially gather more information, such as predicted education, income level and total spending on the platform through the affiliated financial subsidiaries.

Table B1: Summary of Sellers' Survey

Category	Fraction	Category	Fraction
Respondent Chars		**Business Chars**	
Education		*Sources of Supply*	
Primary	2.8%	Own factory	32.5%
Middle School	29.0%	Offline wholesale markets	19.2%
High School	23.7%	Online whotesale markets	21.7%
Some College	28.0%	Distribution/brand subsidary	19.5%
Bachelors	15.4%	Others	7.1%
Master's and Above	0.8%		
Professional Degrees (e.g. MBA)	0.3%	*Number of Employees (inc. owners)*	
		1 - 2 persons	71.9%
Exp in Retail		3 - 5 persons	21.8%
None	36.7%	6 - 10 persons	3.9%
Less than a year	25.6%	>10 persons	2.4%
1 to 3 years	17.2%		
More than 3 years	20.5%	*Total investments*	
		<5k RMB	32.3%
Exp in E-commerce		5k - 10k RMB	19.5%
None	36.3%	10k - 50k RMB	25.2%
Less than a year	38.0%	50k - 100k RMB	9.3%
1 to 3 years	16.5%	100k - 200k RMB	5.0%
More than 3 years	9.2%	>200k RMB	8.7%
Goal			
No specific goal	3.1%		
As part-time job	19.2%		
As main job	58.8%		
Expand offline business online	18.9%		

Notes: Online survey implemented with users assigned to treatment group for the training intervention in August 2018. Separate messages were sent out based on sellers' engagement with the training defined by number of tasks accepted and whether or not sellers have browsed contents of the training. Survey response rates are higher among sellers that were more engaged in the training. All fractions shown adjusted for the sampling and response rate differential.

C Model Appendix[For Online Publication]

Platform exploration is, ultimately, consumer search steered by the platform, or specifically to our setting, by search ranking that may or may not include sponsored listings (advertising). The empirical measure of platform exploration should thus be the degree to which new sellers are sampled, and the quality of the sampled sellers, relative to incumbents.

In this section, we extend Athey and Ellison (2011) and develop an equilibrium model where consumers conduct sequential searches along the platform's search ranking. They differ only in search costs, but incumbent sellers have heterogeneous qualities (match probabilities with consumers) that are ex-ante (before searching) known to the platform but not to consumers or other sellers. Entrants' quality, however, is unknown to the platform, consumers, and all other sellers, creating a "cold-start" problem. We also assume that sellers do not know other sellers' qualities, but do know the total number of ranked sellers and the distribution of qualities.

We assume that bandit exploration takes the form of up-ranking entrants to the top spot with a probability linked to their expected quality, while sponsored listing rewards the highest bidder with that top spot. For simplicity, we further assume that XXXX bidding strategy.

The model generates clear predictions in equilibrium conversion (or match) rates of entrants and incumbents would differ given different exploration intensity and strategy. This echoes our emphasis on entrant outcome as opposed to platform actions alone, it also helps us overcome the lack of ranking and advertising bid data.

Besides the information structure, a main assumption is that consumers have the same tastes and match probabilities to each seller. We thus construct and use our consideration-set data to control for consumer tastes.

C.1 A Model of Platform-Steered Search

Consider a search session in which $J \in \mathbb{Z}[2,\infty)$ incumbent products are matched with the searched keyword, each corresponding to a single seller. Let the quality of each incumbent product j be denoted as $\delta_j \sim F$. Assume that the platform knows incumbent product quality perfectly. The platform then ranks all J products based on a simple sorting algorithm:

$$r_j := \text{rank}^{\text{desc}}\mathbb{E}\left[\delta_j\right] = \text{rank}^{\text{desc}}\delta_j$$

We follow Athey and Ellison (2011) to model consumer search due to the discrete and ordered nature of platform ranking.[20] A continuum of consumers search either until their "need" is met by seller j with probability δ_j—in which case they and the seller each receives a payoff of 1—or until the expected payoff of searching is below their search cost $s_i \sim U[0,1]$. Notice that δ_j equals j's conversion rate.

Consumers share a common prior $\delta_j \sim U[0,1]$ and know the platform's ranking rule. They search sequentially based on r_j because the platform ranks truthful and have superior information on δ_j, hence the term "*platform-mediated search*."[21] The only source of uncertainty is the stochasticity of whether a product of quality δ will meet a given consumer's need. But in aggregate, market equilibrium can be characterized by quality and traffic cascades as seller ranking r worsens (i.e., discretely increases in level):

Proposition 1 (Quality and Traffic Cascade) *The expected quality of the r-ranked product is $\theta(r,J)$, and an r-ranked product will have expected traffic $\phi(r)$, where*

$$\theta(r,J) = \frac{J-r+1}{J+r} \tag{6}$$

$$\phi(r,J) = \prod_{r'=1}^{r-1}\left(\frac{2r'-1}{J+r'}\right) \tag{7}$$

Figure C1: Illustration of Proposition 1: the Quality and Traffic Cascade

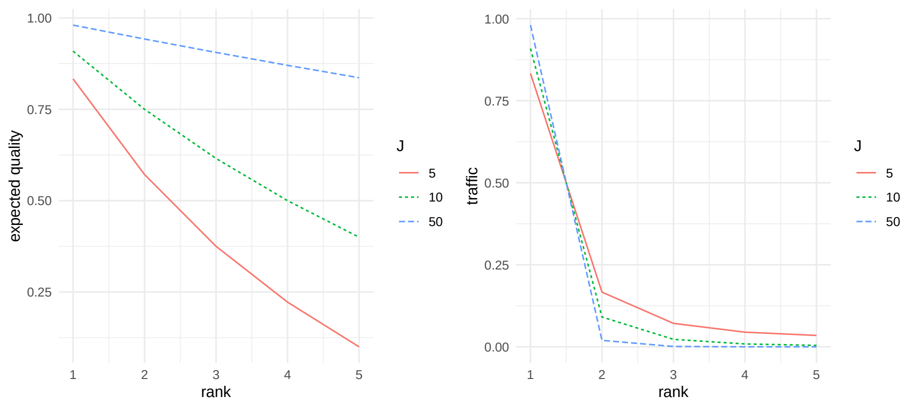

Notes: Each line and color represents a particular value of J. The x-axis represents the rank r, the y-axis plots the equilibrium quality $\theta(r,J)$ and traffic $\phi(r,J)$ as outlined in Equations (2) and (3).

[20] In Appendix E, we instead use a BLP-type econometric model. Model estimates based on our consideration-set data reach similar conclusions as Section 6 as we find superior conversion rates among entrants, particularly trained entrants, as well as an overall market expansion effect of training.

[21] For a detailed analysis of the optimal platform ranking rule, see Nocke and Rey (2021), where our top-down quality sorting rule and the induced search mediation is one of multiple equilibria that can arise.

Discussion: pricing and differentiation The model implicitly assumes that sellers differ by a one-dimensional price-inclusive quality measure, δ. This represents individual search sessions well, in which rankings are based on the likelihood of satisfying a consumer's specific search query and need regardless of whether sellers are horizontally or vertically differentiated. For example, in a search session for "red sweaters," a very well-made (in the vertical sense) green-sweater entrants to advertise, then it would have weakened ad screening despite potentially raising traffic and revenues for those green-sweater entrants.

Given platform-mediated search and the ad ranking rule, a seller's pricing is deterministic given its cost structure and ad price a. It can thus be profiled out.[22] In this sense, interpreting δ as a price-inclusive quality measure is without loss of generality.

Lastly, we assume that $q \ll 1$: entrants have much lower expected quality than incumbents. This can result from learning-by-doing, exiting of low-quality incumbents, or overly optimistic entrant beliefs. Further, if high-quality entrants must buy ads and pay for sampling, this entry barrier further pressures them to raise prices and thus have lower δ.

C.2 Cold-Start Entry and Bandit Exploration

Now assume that an entrant $j = J+1$ is also matched to the keyword. However, due to the "cold-start" problem (entrant quality is private information not only to buyers but also to the platform), the risk-neutral platform relies on a common belief about entrant quality, which is drawn from a Beta distribution $Beta[q, 2-q]$. The expected quality of the entrant is thus $\mu = q/2$ and the distribution is the same as that of incumbent quality if $q = 1$.

The benefit of exploration is dynamic, so we assume that there are two periods in total, where the platform conducts bandit exploration in the first period, updates its prior on entrant quality, and ranks the entrant according to the new expected quality in period 2. Since belief updating is discrete in this case, we assume that there are N total consumers, so that the posterior mean of entrant quality is simply $\frac{\text{number of buyers}+q}{\text{traffic}+2}$.

Proposition 2 (Bandit Exploration) *If the platform ranks the entrant as follows in the first period,*

$$r_{J+1} = S + (1-S) \cdot \text{rank}^{\text{desc}}[\mu] \quad (8)$$
$$\textbf{where } S \sim Bernoulli(\kappa)$$

[22] Argument for if δ is endogenous, then everyone will choose $\delta/2$.

Then the welfare-maximizing platform would set any κ when $\mu > \theta(1,J)$. κ is weakly decreasing in J and in μ but weakly increasing in N.

To see this, XXX.

Corollary 1 (Bandit Exploration and Conversion Rate) *Entrants' conversion rate is weakly decreasing in κ.*

$$r_{J+1} = S + (1-S) \cdot \text{rank}^{\text{desc}}[\mu] \qquad (9)$$
$$\textbf{where } S \sim Bernoulli(\kappa)$$

Then the welfare-maximizing platform would set any κ when $\mu > \theta(1,J)$. κ is weakly decreasing in J and in μ but weakly increasing in N.

C.3 Advertising

The platform holds a second-price auction for the top-ranking spot. For tractability, we assume that incumbents submit bids without considering the entrant. This is most appropriate in situations in which $q \ll \frac{1}{2}$???, where the average entrant. This way, the bid strategy for each incumbent is known. We also avoid modeling incumbents' consideration of tie-break rules between advertising and exploration.

What we model is the simultaneous presence of rent-seeking from incumbents and information signaling among entrants.

Proposition 2 (Ads as Screening Device and Entry Barrier) *If the platform ranks the entrant as follows*

$$r_j(A_j) = A_j + (1-A_j)\left\{\text{rank}^{\text{desc}}\left[\mathbb{E}[\delta_j|b_j]\right]\right\} \qquad (10)$$
$$\text{where} \quad A_j = \mathbf{1}\left[b_j > b_{j'}\right] \quad \forall j' \neq j$$
$$\text{and} \quad \mathbb{E}[\delta_j|A_j] = \delta_j \quad \forall j < J+1$$

Then the equilibrium bidding strategy for each incumbent is.

$$b_{j=1,\dots,J}(\delta) = \frac{1}{2}\delta\left[\phi(1,J) - \phi(\theta^{-1}(\delta,J),J)\right]$$

The entrant will bid $b(\theta(1,J))$ if and only if $\delta_{J+1} > \theta(1,J)$.

Corollary 2 (Advertising and Conversion Rate) *Introducing the position auction reduces the conversion rate among incumbents. Advertised entrants have higher conversion rates than incumbents.*

Looking at the bid function, it is clear that the top-ranked product may or may not be ranked first, depending on how fast the traffic differential increases relative to how fast δ reduces.

Corollary 3 (Entrant Bidding and Conversion Rate) *The entrant's expected conversion rate is weakly higher than incumbents', conditional on winning the bid, but its bidding likelihood weakly reduces with κ.*

Entrants know that the expected winning bid B and will bid if and only if $\delta_{J+1} \geq B$

Corollary 4 (Non-Advertised Conversion Rate Comparison) *Conditional on not advertising, gap between the entrant's and the incumbents' expected conversion rate weakly decreases with κ. It is positive only if $\mu > \frac{1}{2}$, and as $\mu \to \frac{1}{2}$, $\kappa \to 0$.*

Advertising weakly increases unsponsored incumbent conversion rates, κ weakly reduces entrant conversion rates.

To see this, Equation 11 specifies the payoff of advertising for the entrant. Here, the equilibrium becomes "muddled" due to potential advertising insertion: the r-ranked product may be the r-th or the $(r+1)$-th highest in quality. Denote the new equilibrium cascade functions as $\theta^e(r, J+1)$ for expected quality, and $\phi^e(r, J+1)$ for traffic.[23]

$$\left[\phi^e(1, J+1) - \phi^e\left(\lceil \theta^{e,-1}\left(\frac{q\omega}{2}, J+1\right)\rceil, J+1\right)\right] \cdot \delta - a \tag{11}$$

The square bracketed term expresses the traffic differential between the two ads states and is positive by definition of ϕ^e, which may most accurately describe a situation with probabilistic native advertising (Sahni and Nair 2020b).[24] The equilibrium cutoff ω is thus positive as long as $a > 0$; the platform, on the other hand, can always set a arbitrarily close

[23] $\theta^{e,-1}(\cdot)$ is the inverse quality cascade function and $\lceil \cdot \rceil$ is the ceiling function. See Appendix C.4 for analytical formulas for these given a tractable assumption.

[24] If consumers can skip the advertised product as indicated by Ursu et al. (2022), we can still prove that the square-bracketed term in Equation 11 remains positive. Specifically, the second term, the expected traffic assigned to an unadvertised entrant (whose expected quality is $\frac{\mu\omega}{2}$), is less than $\phi\left(\lceil \theta^{-1}(\frac{\mu\omega}{2}, J+1)\rceil, J+1\right)$ due to the lack of information precision on entrant quality. It is thus lower than $\frac{\mu\omega}{2}$, leaving the bracketed term positive still since $\phi^e(1, J+1) = \frac{\mu}{2} + \frac{\mu\omega}{2}$.

or equal to 1 and satisfy its incentive compatibility.[25]

High-quality entrants must suffer the signaling cost a that incumbents of the same quality do not, which turns the "cold-start" problem into a concrete entry barrier. The platform thus faces a trade off in setting a higher a: while it screens out higher-quality entrants with lower chance of muddling the full-information equilibrium, setting $a = 1$ (assuming $J \ll \infty$ and $\theta(1,J) < 1$) would also inefficiently prevent entry.

Corollary 1 (Ad Screening and Conversion Rate) *The entrant's expected conversion rate is increasing in a. It is higher than that of the average incumbent only if $a \gg 0$.*

Corollary 1 highlights a comparative static between equilibrium conversion rates and the intensity of advertising screening. The expected entrant conversion rate is given by

$$\begin{aligned}&\int_{\delta'} \phi^e(r_{J+1}(\mathbf{1}[\delta' > \omega]), J+1)\mathrm{d}G(\delta') \\ =&\frac{q\left[\phi^e(1,J+1)(1+\omega)+\phi^e\left(\left\lceil\theta^{e,-1}\left(\frac{q\omega}{2},J+1\right)\right\rceil,J+1\right)\omega\right]}{2\left[\phi^e(1,J+1)(1-\omega)+\phi^e\left(\left\lceil\theta^{e,-1}\left(\frac{q\omega}{2},J+1\right)\right\rceil,J+1\right)\omega\right]}\end{aligned} \tag{12}$$

Recall that ω is determined by setting Equation 11 to zero and solving for δ. When $a = 0$, $\omega = 0$, the entrant's expected conversion rate is $\frac{q}{2} \ll \frac{1}{2}$ while that of incumbents must be larger than $\frac{1}{2}$ due to platform ranking. In addition, as the platform sets higher a, there is a first-order increase in ω, triggering a secondary decrease in the square bracketed term (in Equation 11) as the expected quality of an unadvertised entrant becomes slightly higher. Figure C2 in Appendix C.4 provides a simulated visualization of the relationship between a and ω with an added tractability assumption.

C.4 Discussion

What does it really look like?

This section provides a tractable characterization of θ^e and ϕ^e as outlined in Section **??**.2.

[25]We note the theoretical and empirical limitation of advertising as a screening device, documented by Blake et al. (2015), Frankel and Kartik (2019), and Moshary (2021), etc. Whenever advertising becomes high-stake enough, its ability to signal δ will suffer owing to other dimensions of private information (gaming, budget constraints) or strategic interactions. This concern is limited in our setting, as we focus on the extensive margin of advertising among entrants that rarely advertises.

Proposition A1 *If the platform ranks the entrant as follows*

$$r_{J+1}(A_{J+1}) = A_{J+1} + (1 - A_{J+1}) \cdot (J+1) \tag{13}$$

Then the unique BNE can be characterized by the expected quality cascade function θ^e*,*

$$\theta^e(r, J+1) = \frac{\mu}{2}\left[(1-\omega^2)\frac{J-r+1}{J+r-4} + \omega\frac{J-r+1}{J+r}\right] r > 2$$

and the expected traffic cascade function $\phi^e(r)$*.*

$$\phi^e(r, J+1) = \frac{\mu}{2}\left((1-\omega^2)\prod_{r'=1}^{r-1}\frac{2r'-5}{J+r'-4} + \omega\prod_{r'=1}^{r-1}\frac{2r'-1}{J+r'}\right) \forall r > 3$$

The entrant advertises if and only if $\delta > \omega \in \mathbb{R}[0,1]$*, so that*

$$\omega = \frac{a}{\left[\frac{\mu(1+\omega)}{2} - \phi^e(J+1, J+1))\right]}$$

We note three points regarding the tractability assumption above, which amounts to a complete downranking of unadvertised entrants. First, this assumption has almost no bite when J is large or when q is small. In these cases, traffic ϕ decreases so fast wrt. r that down-ranking an unadvertised entrant from, say the 2nd page of search results (as in Proposition 2) to the last (as in Proposition A1) makes very little difference in terms of its realized traffic. Second, some degrees of downranking is warranted when consumers are risk- or loss-averse, which is realistic but assumed away for simplicity in the main model. Lastly, notice that the formula is different but remains tractable for $\theta^e(r, J+1)$ and $r \leq 2$ or $\phi^e(r, J+1)$ and $r \leq 3$ since they require special calculation for the expected quality and traffic of the top-ranked product.

Figure C2 presents the equilibrium advertising cutoff ω, the platform revenue, and consumer search as a function of ads pricing a (x-axis), product variety J, and $\mu = \frac{q}{2}$ which denotes the expected entrant quality when incumbent quality is distributed $U[0,1]$. Figure C3 presents a bigger set of equilibrium quantities assuming socially optimal a. Figure C4 plots the equilibrium quantities that optimize different agents' payoffs ("opt-type"), setting $J = 5$ and $\mu = \frac{1}{2}$. This is meant to demonstrate incentive misalignments across the platform, entrants, incumbents, and consumers.

Figure C2: Equilibrium Quantities as Functions of μ, J, and a

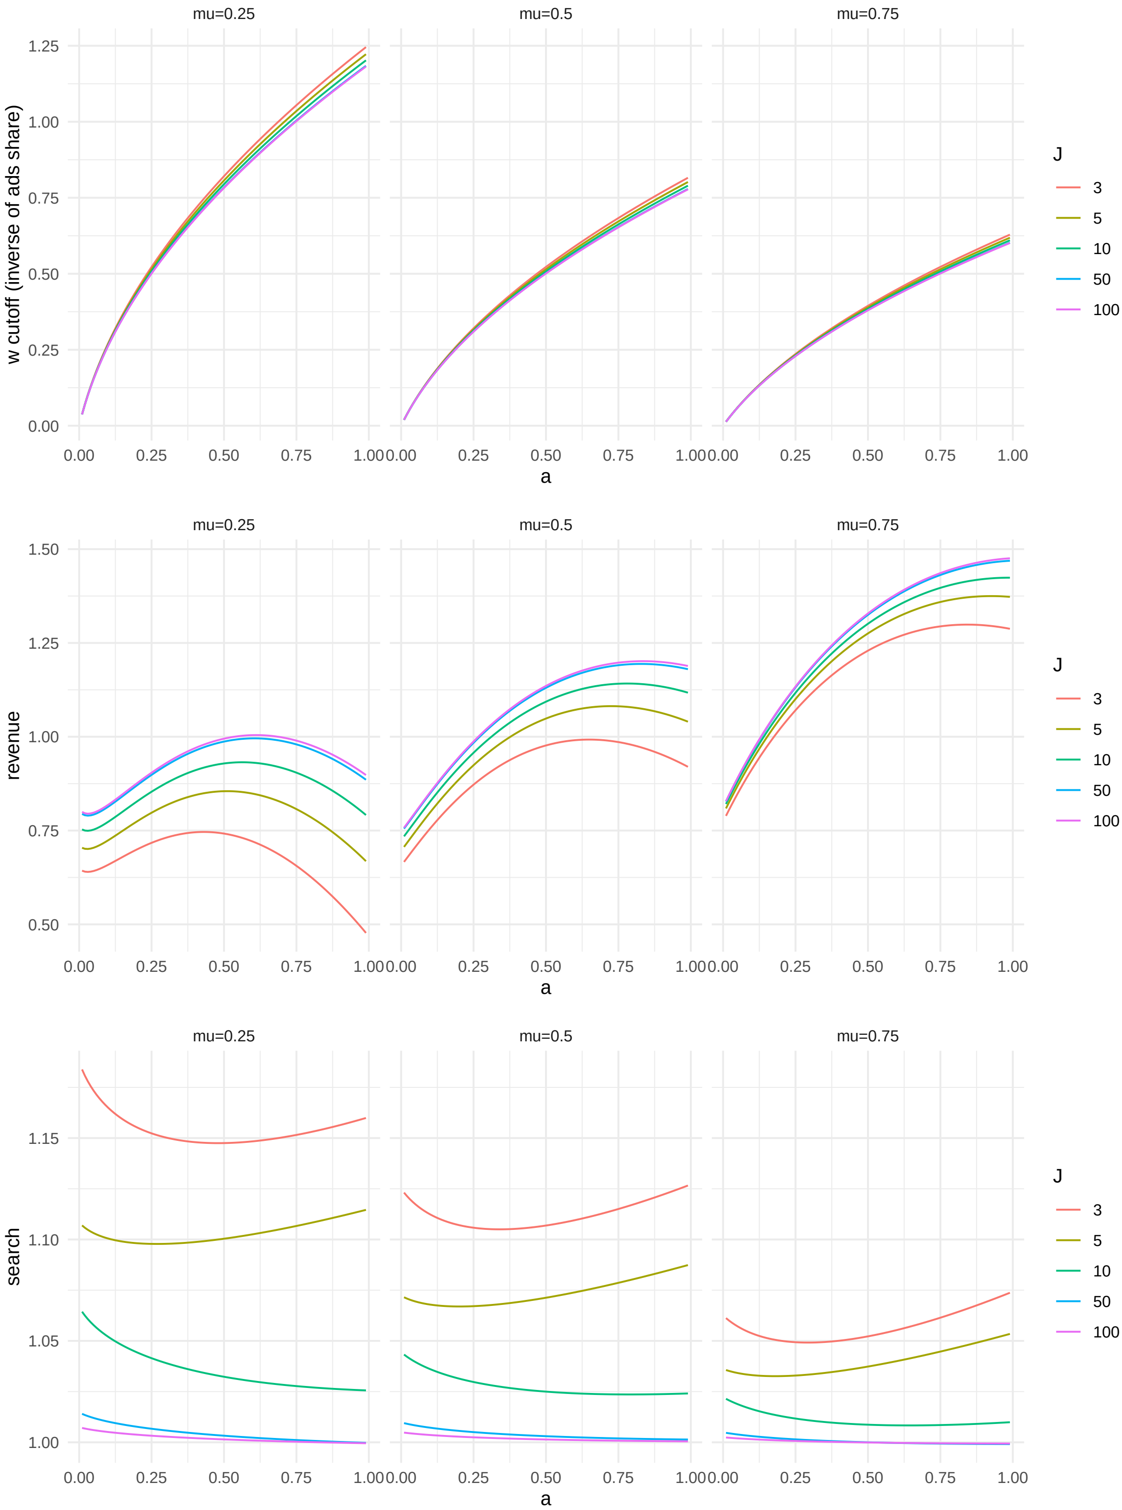

Figure C3: Equilibrium Quantities as Functions of μ and J

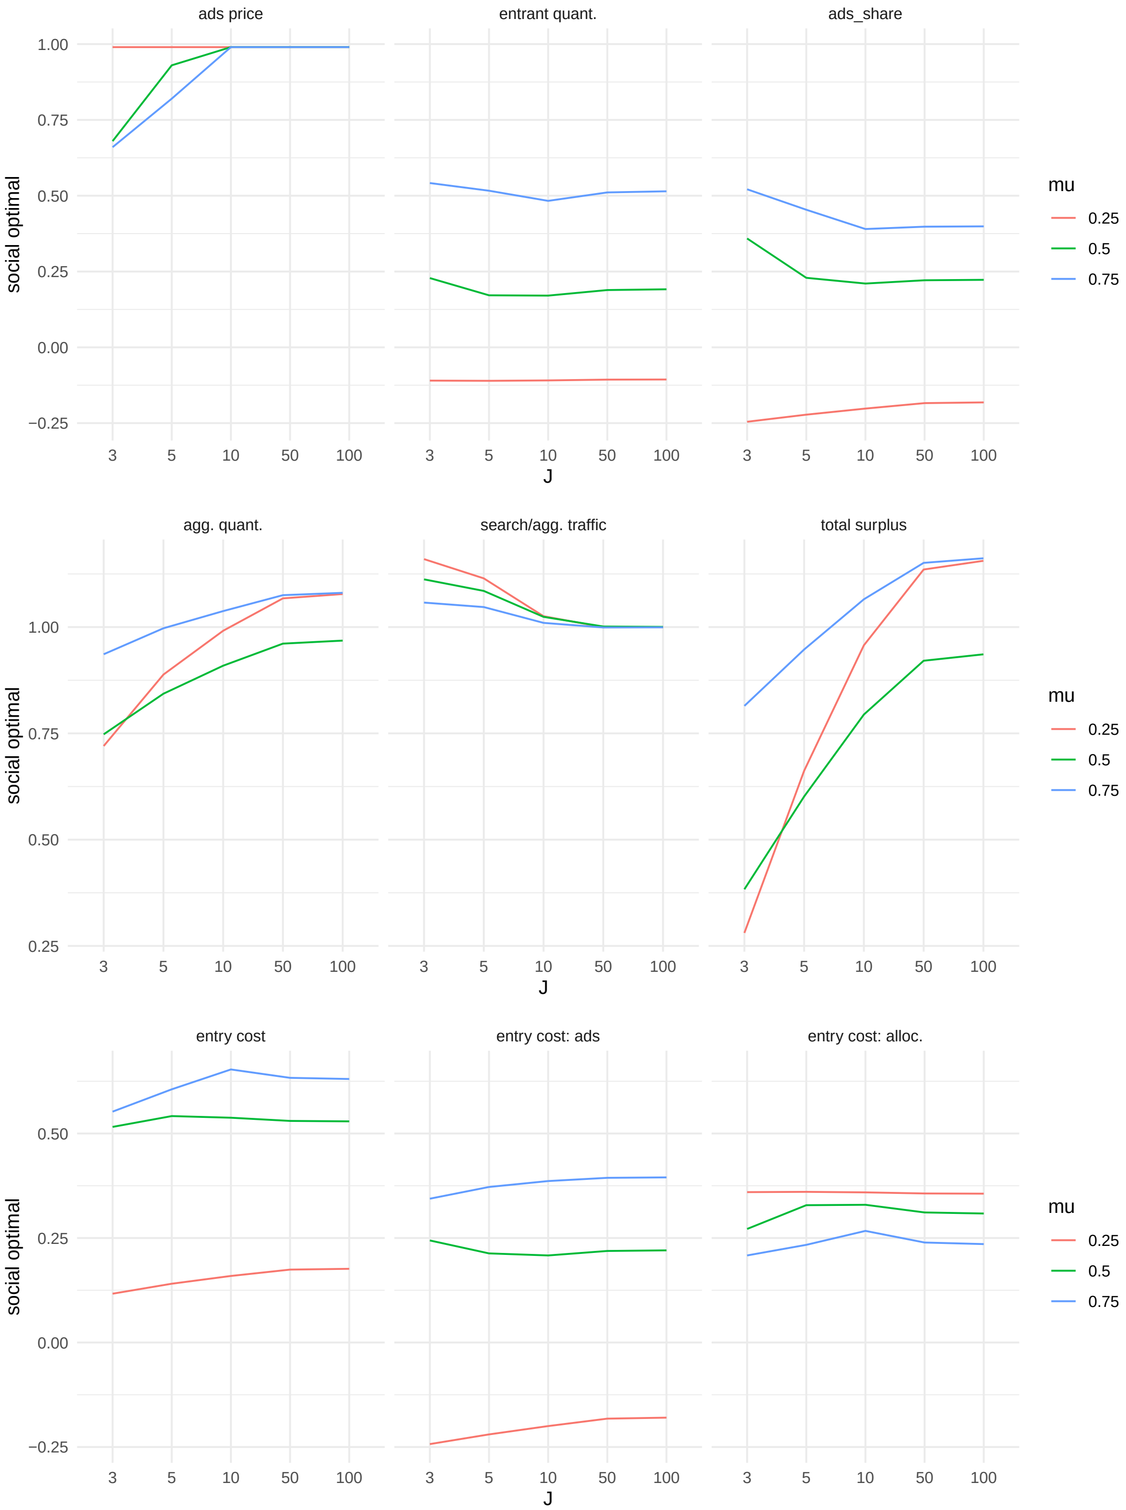

Figure C4: Equilibrium Quantities that Maximize Different Agents' Payoffs

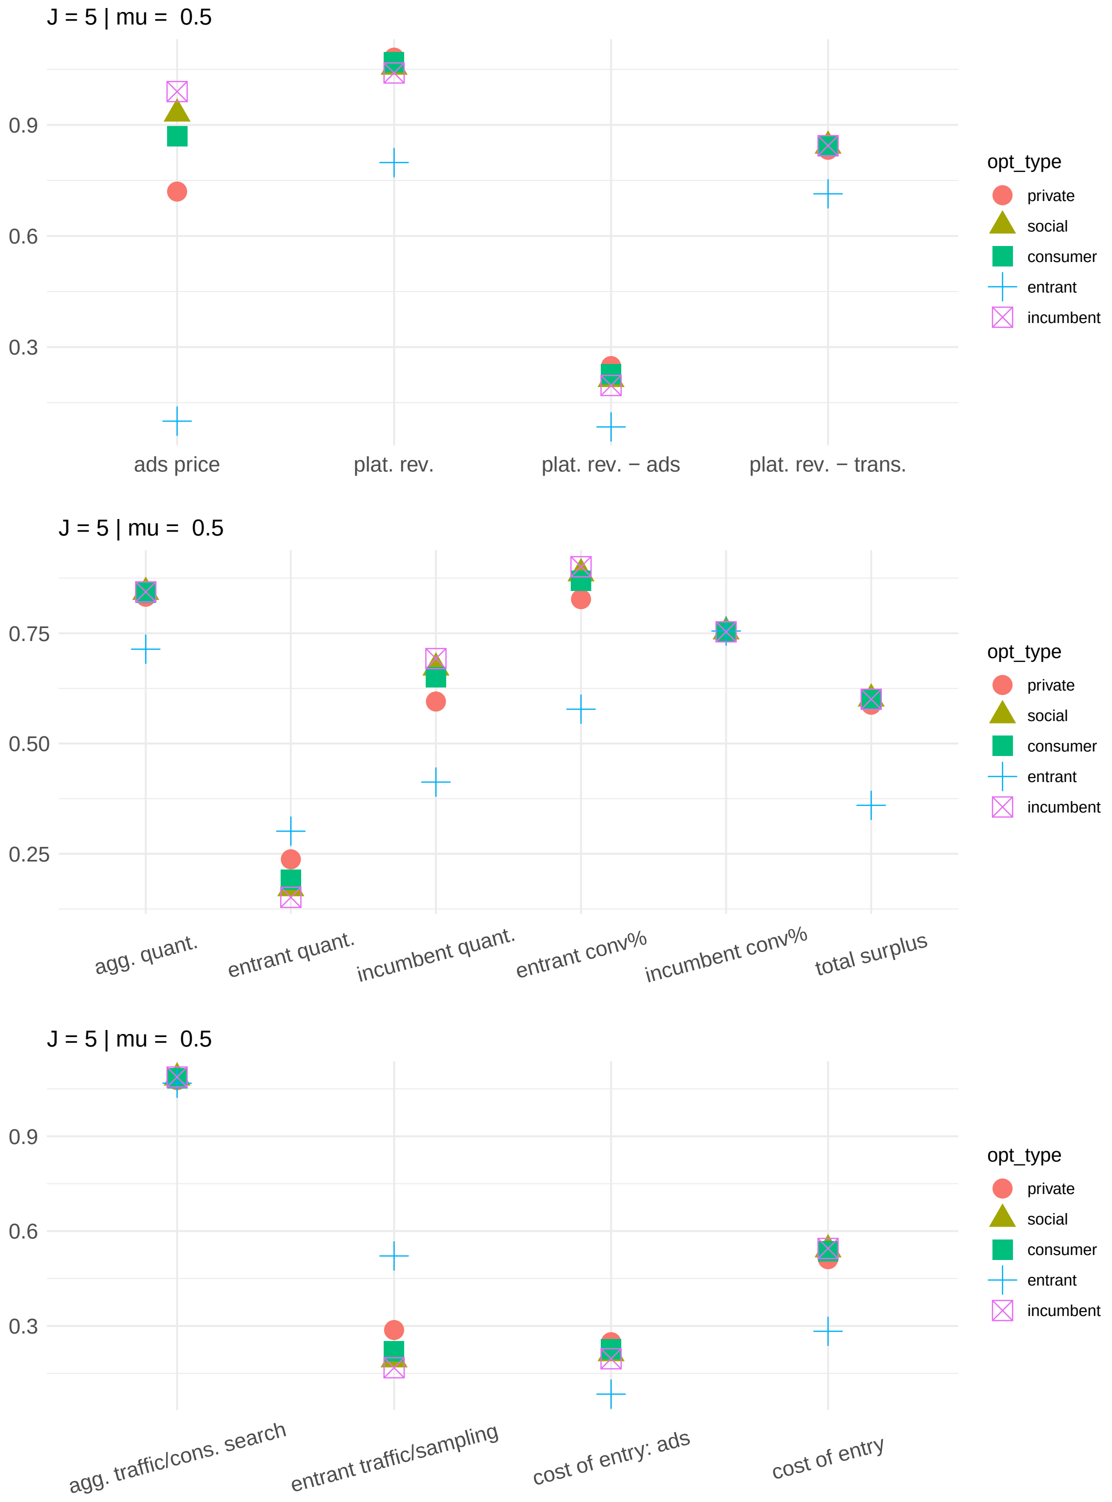

D Additional Tables and Figures [For Online Publication]

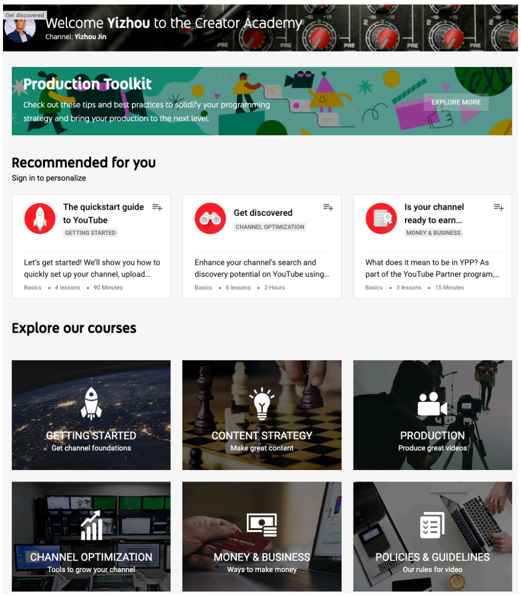

(a) Youtube Creator Academy

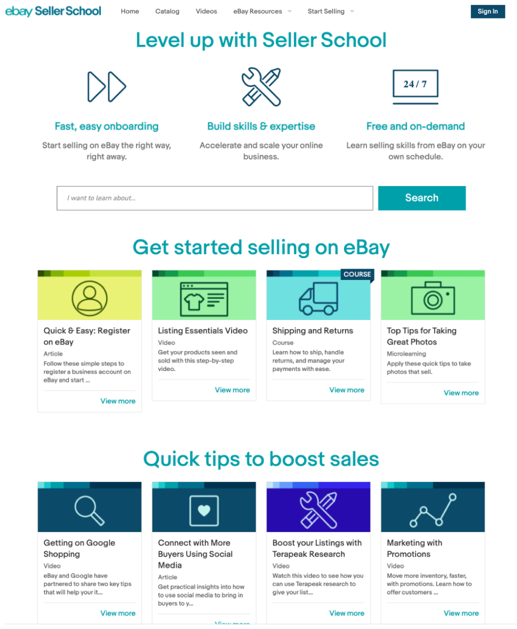

(b) Ebay Seller School

Figure E1: Examples of Entrepreneur Training Programs on Digital Platforms

Table E1: Example of Tasks and Their Classifications

Task	**Indicator**	**Area of Focus**	**Function**	**Type**
Acquire customers' reviews	reviews	ratings	knowledge	outcome
Acquire free traffic	visitors from search channel	marketing	knowledge	outcome
Choose proper promotion products	payment received	basic	knowledge	outcome
Complete pricing and inventory information	payment received	basic	knowledge	outcome
Complete an order	payment received	basic	knowledge	outcome
Expand base of followers	followers	customers	knowledge	outcome
Improve "add to shopping cart"	add to cart	marketing	knowledge	outcome
Improve buyer review section	reviews	basic	reminder	outcome
Improve conversion rate: inqury	conversion	service	knowledge	outcome
Improve conversion rate: make payment	conversion	marketing	knowledge	outcome
Improve fans' engagement	followers	customers	reminder	outcome
Improve payments from returning customers	payments received	customers	knowledge	outcome
Engage with customers via microblog	followers' activities	customers	reminder	action
Improve ratings on customer service	ratings	ratings	knowledge	outcome
Improve ratings on product quality	ratings	ratings	knowledge	outcome
Decorate store front-page on app	decoration	basic	reminder	action
Improve tag/bookmark rates	bookmarked	marketing	knowledge	outcome
Improve ratings on delivery	ratings	service	knowledge	outcome
Optimize products' titles	traffic	marketing	knowledge	outcome
Participate in official sales events	sign-up	marketing	knowledge	action
Participate in product position auctions	paid ads	marketing	knowledge	action
Pay security deposits	deposits	basic	reminder	action
Post products on store page	number of products	basic	reminder	action
Set up automated position auction campaign	paid ads	marketing	knowledge	action
Setup bonus after purchase	bonus	basic	reminder	action
Setup free return and refund	return policy	basic	reminder	action
Setup free trial / offer free samples	free trial	basic	reminder	action
Setup store coupons and discount	coupons	basic	reminder	action
Shorten average time to delivery	delivery time	service	knowledge	outcome
Shorten response time to customer inquires	response time	service	knowledge	outcome
Subscribe to storefront customization	"hip-store" subscription	basic	reminder	action
Upload videos for product descriptions	videos	basic	reminder	action
Upsell to increase per-consumer spending	avg. order size	marketing	knowledge	outcome

Notes: Listed tasks are a subset of all tasks offered to the sellers. Over time service providers also created more tasks and the platform invested in streamlining and regularizing the tasks offered. Tasks are ordered in sequence of priorities. Each task is triggered by a particular indicator. For the outcome based tasks, comparisons are made with other sellers in the same industry. Tasks are classified based on main area of focus, the functions they served and how they are evaluated.

Table E2: Balance Check and Selective Participation
(at time of registration and treatment assignment)

	Full Sample			Sellers in Treatment Group		
	Treatment	Control	(1) - (2) Difference	Participants	Non Participants	(4) - (5) Difference
	(1)	(2)	(3)	(4)	(5)	(6)
Is Firm	0.268	0.264	0.004***	0.275	0.266	0.008***
	(0.443)	(0.441)	(3.33)	(0.446)	(0.442)	(3.18)
Female Owner	0.455	0.453	0.002	0.442	0.458	-0.16***
(among individual sellers)	(0.497)	(0.497)	(1.27)	(0.497)	(0.498)	(4.61)
Region: Coastal South	0.435	0.434	0.001	0.518	0.413	0.105***
	(0.496)	(0.496)	(0.28)	(0.500)	(0.492)	(3.573)
Region: West	0.118	0.118	0.0001	0.089	0.126	-0.038***
	(0.323)	(0.323)	(0.34)	(0.284)	(0.332)	(21.41)
List Products on the 1st Day	0.213	0.212	0.0005	0.245	0.204	0.04***
	(0.409)	(0.409)	(0.44)	(0.430)	(0.403)	(16.04)
# Listed Products	1.539	1.535	0.004	2.307	1.342	0.965***
	(2.166)	(2.164)	(0.65)	(1.988)	(2.167)	(80.81)
Traffic	-	-	0.006	-	-	1.527***
	(2.289)	(2.283)	(0.91)	(2.259)	(2.191)	(115.35)
Conversion Rate	0.051	0.051	0.0002	0.055	0.050	0.006***
	(0.161)	(0.162)	(0.52)	(0.120)	(0.171)	(7.34)
	(3.312)	(3.304)	(1.144)	(3.726)	(3.048)	(100.66)
Revenues	-	-	0.01	-	-	2.134***
	(3.312)	(3.304)	(1.144)	(3.726)	(3.048)	(100.66)
Observations	177,026	535,092	712,118	36,189	140,837	177,026

Notes: The timing of the assignment limits the available baseline information to variables collected during the registration and actions taken on the first day of entry, most of which is listed here. Columns 1, 2, 4, and 5 present means and standard deviations (in parentheses). Columns 3 shows the difference in means across the treatment and control group (training participants and non-participants) in the full sample with the corresponding t-statistics in parentheses. Columns 6 shows the comparison between training participants and non-participants. For the second row (female owner) the value in parentheses shows the estimate with a further sample restriction that includes only those not registered as firms. Participation is defined as having taken up any tasks during the nine-month period. Traffic, conversion rate, revenues and number of product posted are for the first month outcomes.

Table E3: Intent-to-Treat Effects on Revenues (Robustness Check on Revenue Definition)

	Dependent variable:					
	Monthly Revenues					
	Indicator	Log	I.H.S.	Raw / Winsorized		
				99th	99.5th	99.9th
	(1)	(3)	(4)	(5)	(6)	(7)
Treatment	0.002**	0.017***	0.018***	25.436**	26.175	85.084
	(0.001)	(0.006)	(0.007)	(12.478)	(22.308)	(54.321)
Dep Var Mean	0.19	-	-	-	-	-
Observations	6,409,062	6,409,062	6,409,062	6,409,062	6,409,062	6,409,062
R^2	0.152	0.132	0.134	0.027	0.024	0.019
Adjusted R^2	0.152	0.132	0.134	0.027	0.024	0.019

Notes: Dependent variables are total revenues in the seller sample. All regressions include cohort, initial sector and relative month fixed effect. In column 4, the revenue is transformed with inverse hyperbolic sine. We calculate thresholds for winsorization by the relative month since registration. Dependent variable means calculated with sellers in the control group.* significant at 10% level, ** significant at 5% level and *** significant at 1% level.

Table E4: ATT Results with Alternative Take-up Definition: the Number of Tasks Taken-up

	Dependent variable:				
	Any Revenues	Revenues	# Visitors	# Buyers	Conversion Rate
	(1)	(2)	(3)	(4)	(5)
# Tasks	0.028***	0.284***	0.208***	0.139***	0.002
	(0.011)	(0.088)	(0.077)	(0.043)	(0.002)
Dep Var Mean	0.19	-	-	0.57	0.04
Observations	6,409,062	6,409,062	6,409,062	6,409,062	2,593,762
R^2	0.155	0.136	0.209	0.109	0.049

Notes: The first stage specification is the following:

$$Tasks_{it} = \sum_{k=1}^{9} \gamma_k Treatment_i \mathbf{I}_{t=k} + \alpha_t + \alpha_s + \alpha_c + \varepsilon_{it}$$

where $Tasks_{it}$ is number of tasks taken-up for seller i during month i. $\mathbf{I}_{t=k}$ is an indicator for being in the kth month since registration. All specifications are two-stage least square results using treatment assignment interacting with dummies for relative month as the instrument. Number of visitors, number of buyers and revenues (total payments received) are monthly total in log after adding one to the level. All regressions include cohort, relative month and initial industry fixed effect. Standard errors clustered at seller level. Dependent variable means calculated with sellers in the control group. * significant at 10% level, ** significant at 5% level and *** significant at 1% level.

Table E5: OLS Results on Main Outcomes

	Dependent variable:							
	Log Revenues	Any Revenues	Revenues	Log # Visitors		Log # Buyers		Conversion
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Take-up Tasks	1.847***	0.223**	0.771***	1.735**	0.703***	0.825***	0.686***	0.005***
	(0.017)	(0.002)	(0.019)	(0.014)	(0.015)	(0.009)	(0.012)	(0.0004)
Dep Var Mean	-	0.19	-	-	-	0.57	1.4	0.04
Sample	Full	Full	Earn Revenues	Full	Have Visitors	Full	Have Visitors	Have Visitors
Observations	1,593,234	1,593,234	314,376	1,593,234	646,894	1,593,234	646,894	646,894
R^2	0.193	0.204	0.108	0.274	0.128	0.160	0.098	0.048
Adjusted R^2	0.193	0.204	0.1087	0.274	0.128	0.159	0.097	0.047

Notes: Sample restricted to sellers with access to training. Main explanatory variable is an indicator of take up of at least during the sample period. Traffic (number of visitors), number of buyers and revenues (total payments received) are monthly total in log after adding one to the level. All regressions include cohort, relative month and initial industry fixed effect. Standard errors clustered at seller level. Dependent variable means calculated with sellers in the control group. * significant at 10% level, ** significant at 5% level and *** significant at 1% level.

Table E6: Treatment Effect on Sellers' Ratings and Refunds

Variable	Treatment	Dep Var Mean	Variable	Treatment	Dep Var Mean
Ratings			**Refunds and Reviews**		
Products	0.004	1.25	% Refund (Amount)	-0.0004	0.21
	(0.004)			(0.001)	
Service	0.004	1.26	% Complaints	-0.003	0.05
	(0.004)			(0.006)	
Logistics	0.005	1.26	Rule Violations	0.0002	0.22
	(0.004)			(0.001)	
			% Good Reviews	-0.00001	0.99
				(0.0002)	

Notes: Table presents estimated coefficients β on treatment assignment dummy with specification **??**. Standard errors clustered by seller. All regressions include month, entry date and main industry fixed effect. Ratings are customer ratings variables that the platform calculates and assignes to sellers based on customers' reviews and ratings. The ratings scale between 0 to 5, on the dimensions of accuracy of product descriptions, quality of customer service and logistics. % refunds calculated as total refunds requested over total payments made. % complaints defined as number of complaints over total number of orders. Rule violations is the frequency that sellers violate the platform's rules, see more details in appendix E, descriptions of the instruments. Share of good reviews is the share of good reviews out of all the reviews that sellers get. Vast majority of the reviews are positive. * significant at 10% level, ** significant at 5% level and *** significant at 1% level.

Figure E2: Training Participation: Tasks Take-up and Completion

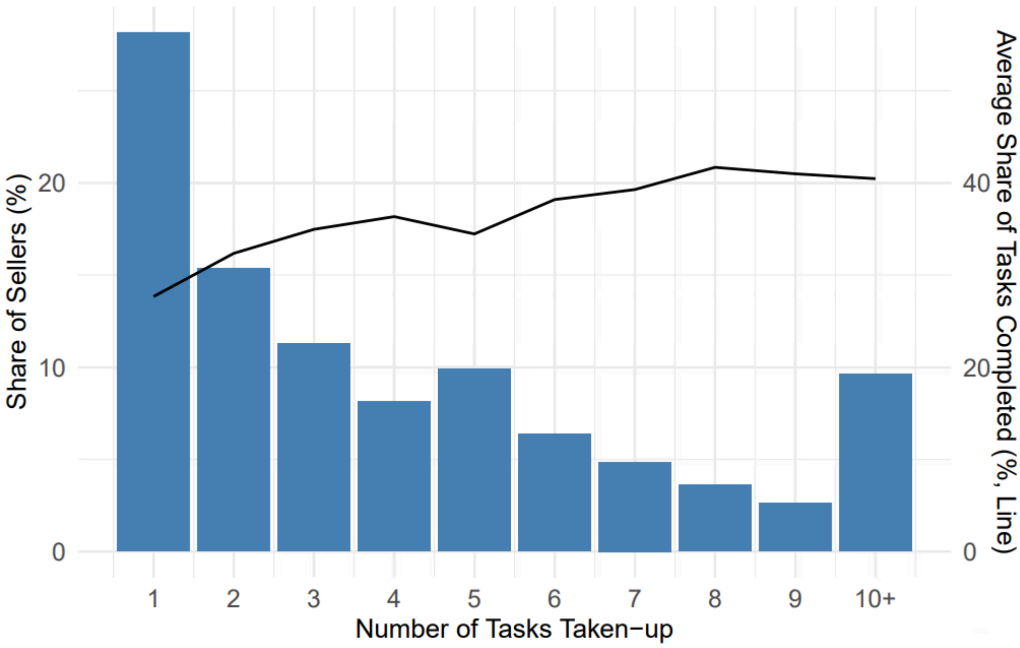

Figure E3: Procedures of Constructing the Consumer Sample

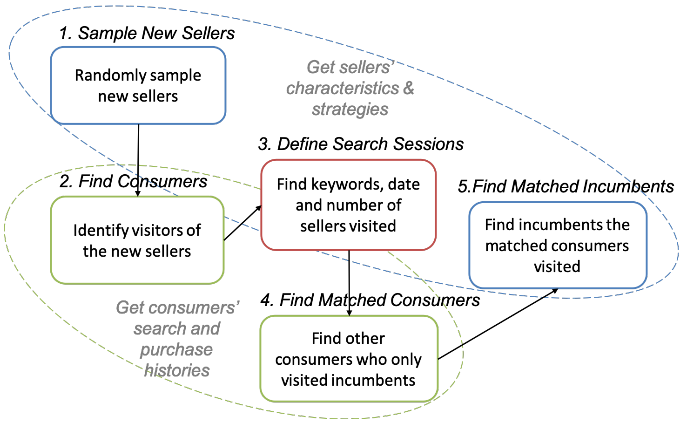

E Structural Estimation[For Online Publication]

This section quantifies the welfare gain by augmenting our model, first introduced in Section **??**, and taking it to the consideration-set data.

E.1 A Structural Econometric Model

Consumer Demand In order to quantify consumer welfare, we augment our economic model in Section **??** to fit our consumer choice data in consideration sets. Each consumer i arrives at the platform with a predetermined consideration set K_i that consists of a group of seller-product pairs $j \in K_i$, each provides utility U_{ij}, which consists of δ_{ij}, consumers' valuation of the product, and noise ε. Compared to the simple model, we incorporate consumer heterogeneity but do not allow the degree of information asymmetry to vary based on seller type, both informed by our reduced-form analyses above.

$$U_{ij} = \delta_{ij} + \varepsilon_{ij} = x_j\beta - \alpha p_j + \xi_j + \varepsilon_{ij} \tag{14}$$

p_j is the price level of j, and x_j a set of observable characteristics. ξ_j is the residual, unobserved product value that is our main object of interest. We interpret ξ_j as *product quality* in the vertical differentiation sense. Since the platform observe and utilize p_j and x_j, the main source of friction our model captures comes from a mismatch between sellers' unobserved quality ξ_j and their likelihood of appearing in consumers' consideration sets. ε_{ij} is the i.i.d. consumer-seller idiosyncratic preferences that follows a type I extreme value distribution. Lastly, we assume consumers' outside option of not purchasing from any seller has zero utility, which gives way to the following familiar logit formulation:

$$P_{ij} = \frac{\exp(x_j\beta - \alpha p_j + \xi_j)}{1 + \sum_{k \in K_i} \exp(x_k\beta - \alpha p_k + \xi_k)} \tag{15}$$

P_{ij} is consumer i's probability of purchasing from seller j. We later enrich the baseline model by adding consumer side heterogeneity, sector-specific fixed effect ξ_s and sector-specific price coefficient α_s.

Matching The algorithm that matches sellers and consumers is the most important means the platform has to improve consumer experience and support promising sellers. But it can also strongly contribute to the entry barrier that we have uncovered if it relies heavily on past consumer choice data. We simplify the complex matching rules used in the search and recommendation algorithm by focusing on the reliance on the previous period's conversion rates and the number of visitors. Here, conversion rate directly reflects seller-specific consumer demands affected by sellers' underlying quality ξ_j. The last period's total number of visitors summarizes sellers' characteristics, especially their marketing skills and the impact of the training on attracting consumers. In contrast, sellers' actions and training participation do not directly enter the matching function. We capture the evolution of the number of visitors over time with the following model:

$$\mathscr{T}_{jt} = f(\mathscr{T}_{jt-1}, C_{jt-1}, C_{jt-2}) \tag{16}$$

$\mathscr{T}_{jt}$ is the current-period traffic (number of visitors), which is modeled as a function of its lag $\mathscr{T}_{jt-1}$,[26] and conversion rates C over the previous two periods.

E.2 Identification and Estimation

We estimated the model using simulated maximum likelihood following Train (2009). To better fit the empirical setting, we make the following changes to the baseline model.

Consumer Demand We use the consumer-seller pair sample to estimate the demand parameters in particular sellers' ξ_j. Due to computation constraints, we sample a subset of sellers for actual estimation. The final sample is a seller-consumer matched pair dataset, where we explore seller-level variations. x_j includes the number of products sellers offer and sellers' ratings. p_j is the average price level that sellers charge.[27] To account for sector-level heterogeneity, we add a sector-specific intercept ξ_s in the baseline model and enrich the baseline model by estimating the sector-specific price coefficient α_s.

Endogenous Strategies One concern with the baseline model is that pricing level p_j and strategies such as the number of products offered in x_j could correlate with ε_{ij}, which would bias the estimates. To address this concern, we use a set of instruments to jointly determine the number of products posted prod_{jt} and pricing level p_{jt} with

$$\begin{bmatrix} p_{jt} \\ \text{prod}_{jt} \end{bmatrix} = \mathbf{Z}\beta^{fs} + \xi_j^{fs} + \xi_t^{fs} + \xi_s^{fs} + \varepsilon_{jt} \tag{17}$$

[26]We use log scale as the traffic distribution is very skewed.

[27]As mentioned before, ratings capture sellers' size more than quality. The price level seller charges is the quantity-weighted prices of all products sold by the sellers during the day. Such weighted prices reflect the relative popularity of products sellers' offerings as well as any sales or promotions.

ξ_j^{fs}, ξ_t^{fs} and ξ_s^{fs} are seller, date and sector fixed effect. The instruments Z capture the stringency of the platform's rule enforcement. These instruments include frequencies of rule violations and shares of sellers identified as frequent rule violators in the corresponding sectors.[28] The most common rule violations include infringing on intellectual property rights, selling counterfeits and providing false or misleading product information. The platform enforces comprehensive rules to ensure implementation of relevant government regulations and to maintain a well-functioning market. In the cases of rule violations, the platform might downgrade sellers in the search rankings, remove access to sellers' products or even their sites and, in some cases, pursue legal action. Such punishment could have significant impact on sellers' business operations. The platform frequently adjusts the rules in response to changes in the business environment. We exploit changes in the stringency of rule enforcement at the sector-month level. When the platform strengthens the rule enforcement, sellers could be more cautious about posting products, charging extreme prices or engaging in unruly promotions. Conversely, the rule-obeying sellers are better positioned when the platform imposes stricter regulations on their unscrupulous competitors.

We jointly estimate consumer demand and seller strategy using the instruments described above. The final set of instruments includes the average frequency of rule violations and the share of sellers labeled as frequent rule violators in sector s over the previous 30-day period. In the baseline model, we do not explicitly account for temporal variations. More detailed estimation procedures are described in **??**.

Matching We use the the new seller panel to estimate the matching function. We use the following specification to distinguish new entrants with no previous history and sellers with zero conversion rates in the previous periods from the rest:

$$T_{jts} = f(T_{j,t-1}, C_{j,t-k}) + g(I_{j,t-k}(t-k=0)) + h(I^c_{j,t-k}(C_{j,t-k}=0)) + \xi_k^{trf} + \xi_t^{trf} + e_{jts} \tag{18}$$

$C_{j,t-k}$ are lagged conversion rates in past periods $k = 1,2$, $I_{jt-k}(t-k=0)$ is an indicator for the initial two periods, and $I^c_{j,t-k}(C_{j,t-k}=0)$ are indicators for lagged conversion rates being exactly zero.[29] The specification also includes product category, calendar time and relative month fixed effect. In the baseline specification we start with linear functions for

[28] We use the number of visitors each seller attracts as the weights when we aggregate these instruments to sector level. Rule violations come in different degrees of severity which we calculate separately.

[29] Sellers have no past history of number of visitors and conversion rate in the initial periods and hence are subject to different matching rules.

$f(\cdot)$, $g(\cdot)$ and $h(\cdot)$. In this setup, all the right-hand-side variables are determined in the previous periods.

E.3 Estimation Results

The main parameters of interest are ξ_j for seller $j \in J$, which we detail in Table D2. These seller fixed effects can be interpreted as a quality measure for their products and services. Focusing on those with at least one purchase record (28% of all sellers), new sellers have significantly higher ξ_j than incumbents. Among sellers with purchase records, the estimated quality of new sellers is 23.1% higher than incumbents. This means that the average product offered by new sellers provide 65-cent higher consumer gain. These results confirm what we found in the reduced-form analysis: new sellers have higher underlying quality, allowing them to out-compete incumbents in the same consideration sets. To unpack the welfare implications of the training intervention, we turn to counterfactual scenarios in the matching.

There is substantial cross-sector ("s") heterogeneity. Panel A of Figure D1 plots the distribution of the estimated sector fixed effect ξ_s. Panel B of figure D1 plots the distribution of the price elasticity for sellers in the sample. The average price elasticity is -0.22.[30]

E.4 Counterfactual: Welfare Impact of the Training

We analyze the welfare consequences of the training by considering how changes in the likelihood for different types of sellers to appear in consumers' consideration sets induced by the training affect consumer surplus and sellers' revenues.

To conduct the counterfactual analysis, we randomly sample a subset of sellers along with their associated search sessions. The potential pool of sellers a consumer searching a particular keyword could choose from consists of all sellers visited by any consumers searching that keyword in the full sample. For each seller, we obtain their strategies x_j and estimated quality ξ_j. Based on the estimated results, we calculate consumers' utility when they visit a particular seller j, V_j. In the baseline model, V_j is the same across consumers. To construct the consideration sets, we hold constant consumer-search session pairs and

[30]The elasticity we estimate here is much smaller than the typical elasticity observed in the literature broda2006. The difference occurs because consumers are choosing between products in their consideration sets, rather than choosing among all the products offered on the market. When constructing the consideration sets, consumers already restrict their choices to a narrower range of prices. The average standard deviation of prices among all sellers that some consumers visited is 3.14 times higher than the average standard deviation of prices among sellers in the consumer-specific consideration sets.

sample sellers from the sellers' pool associated with the search keyword according to sampling weights. We match the number of sellers sampled with the observed number of sellers each sampled consumer visited in the corresponding search sessions. Therefore in the counterfactual, we only vary the composition of the consideration sets and hold everything else constant. Since training increases treated new sellers' likelihood of appearing in consumers' consideration sets, we evaluate the welfare of the training by restricting training participants' probability of being sampled, as described below. More details about construction of the counterfactual consideration sets are described in **??**.

Baseline In the baseline version, sellers' sampling weights are given by their predicted numbers of visitors as determined by the empirically estimated matching rule described in section E.3. The current logit specification allows us to calculate consumer surplus as

$$CS = \frac{1}{\alpha} \sum_i \log[\sum_{j \in K_i} \exp(V_j)]$$

Sellers' revenues are given by the probability of being chosen and the price level:

$$R_j = \sum_i \frac{\exp(V_j)}{1 + \sum_{k \in K_i} \exp(V_k)} p_j$$

Sellers will not earn any revenues if they do not appear in the consideration sets. Table D3 summarizes the results of welfare decomposition. In the baseline, new sellers only capture 6% of total revenues even though they represent 8.3% of sellers in the pool[31].

Welfare Impact of the Training Program In this exercise, we use our structural model to compute the welfare impact of training. In particular, we calculate a counterfactual scenario in which the distribution of unobserved quality among new sellers in the treatment group–and hence their likelihood to appear in consideration sets—is equalized with those in the control group.[32]

Overall, out of all products in consideration sets, only 4.3% are from new sellers. Those from trained new sellers make up only 0.25%. However, removing the training program from such a small portion of the market is enough to reduce total revenue reduces by 0.05%, with consumer welfare dropping by 0.07%. In particular, treated sellers suffer a

[31]The sample we used in this estimate is not the same as the sample used in **??** because of the model specification.

[32]We randomly remove a subset of training participants from the sampling pool that consumers could choose from, until the likelihood of treated new sellers' appearance in the consideration sets matches those of control new sellers.

7.7% drop in market share, obtaining a smaller share of a smaller pie.[33] Both control-group and incumbent sellers see moderate increase in market share.

Table D1: Prediction Precision: Traffic

Lagged Traffic Degree	Lagged Conv. Rate Deg	FE	R^2	RMSE	MAE
1	1	Y	0.63	1.63	1.18
1	1	N	0.6	1.7	1.23
2	1	N	0.6	1.69	1.22
2	1	N	0.61	1.68	1.21
3	2	N	0.6	1.69	1.22

Notes: Table shows measure of prediction's precision with different specifications on current period traffic. Precision calculated with on the out-of-sample data.

Table D2: Sellers' Characteristics and Estimated Type ξ_j

Statistic	N	Mean	St. Dev.	Min	Max
All Sellers	52,241	0.141	0.314	−2.828	5.024
By Purchase Status					
No Purchase	37,617	0.098	0.051	−2.828	0.101
Has Purchase	14,624	0.252	0.573	−1.965	5.024
By Sellers' Type					
Control	3,366	0.136	0.296	−1.445	4.939
Treatment	988	0.150	0.353	−0.853	3.668
Incumbent	47,887	0.142	0.314	−2.828	5.024
By Type Among Sellers with Purchase					
Control	788	0.258	0.593	−1.370	4.939
Treatment	252	0.296	0.679	−0.853	3.668
Incumbent	13,584	0.251	0.570	−1.965	5.024

Notes: Distribution of estimated unobserved seller characteristics, ξ_j, on a sub-sample of new sellers and incumbents conditional on making at least one sale. See Appendix E for more details.

[33] Percentage point are hidden at data provider's request.

Table D3: Welfare and Market Share Decomposition

Revenue	CS	Market Share		
		Treatment	Control	Incumbent
-0.05	-0.07	-7.71	0.17	0.13
(0.05)	(0.03)	(0.63)	(0.11)	(0.02)

Notes: Welfare and market share calculated with a random sample of 60,000 consumers. "CS" stands for consumer welfare. Results presented in the table are percentage difference in comparison to baseline level where traffic assignment is determined by predicted traffic only. Bootstrapped standard errors are in parentheses. Revenues and consumer surpluses are changes from baseline level revenues and consumer surpluses respectively. See Appendix E for more details.

Figure D1: Cross-Sector Differences

(a) ξ_s (b) α_s

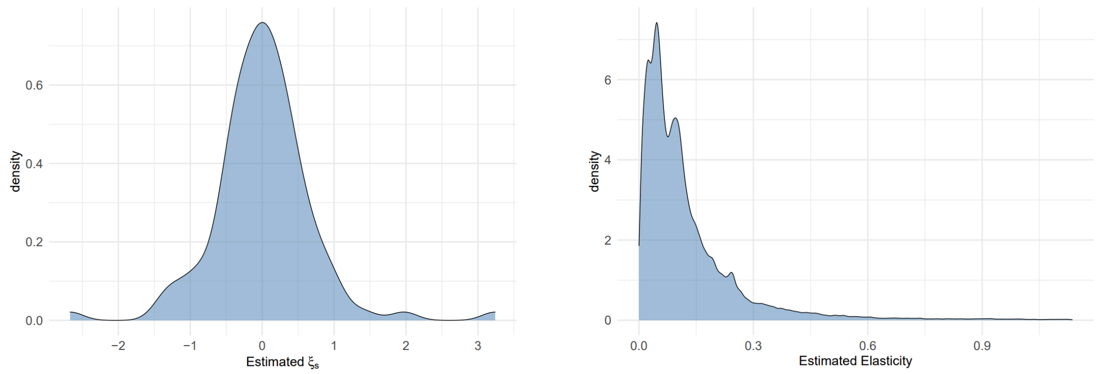

Notes: These figures plot the distribution of two key latent parameters from our model estimation that represent heterogeneity across sectors. ξ_s measures sector-level residual in unobserved seller quality, while α_s measures the sector-specific price elasticity.