

Lifting Growth Barriers for New Firms

Evidence from an Entrepreneurship Training Experiment

with Two Million Online Businesses*

Yizhou Jin[†]

Zhengyun Sun[‡]
(Job Market Paper)

December 1, 2020

[Latest version here](#)

Abstract

Expansion of e-commerce presents new opportunities for SMEs to enter broader markets at lower costs, but the new entrants face barriers to growth after entry. To facilitate the new entrants to overcome these barriers, we implement a training program as a randomized controlled experiment with over two million new sellers on a large e-commerce platform. The training focuses on practical skills specific to online business operations. Treated new sellers with access to the training earn higher revenues. These sellers improve marketing skills and attract more consumers to their online stores. Leveraging detailed consumer-seller matched search and browsing data, we find that consumers have higher purchase probability overall when they encounter new sellers regardless of treatment status. In the cases of purchases, consumers choose treated new sellers over incumbents; moreover, doing so does not lower the quality of their purchases. We use a structural model to characterize consumer demand and recover sellers' underlying quality. Both treated and control new sellers have a higher quality compared to incumbents. The training increases new sellers' likelihood of being encountered by consumers, which improves the matching quality between consumers and sellers. The counterfactual exercise shows that the training leads to higher consumer surplus and sellers' total sales due to market expansion. As the operator of the online marketplace, the platform could earn more profits in both the short and the long run because of the training.

Keywords: e-commerce, platforms, business training, growth barriers

JEL codes: L26, L81, M53, O12, O17

*Sincere thanks to our advisors Jie Bai, Emily Breza and Michael Kremer for their extensive advice and guidance. We thank the product development team at the unnamed e-commerce platform for making the experiment possible. For especially valuable comments and discussion, we thank Long Chen, Shawn Cole, Asim Khwaja, Gabriel Kreindler, David McKenzie, Song Ma, Ariel Pakes, Gautam Rao, Yulu Tang, Yuan Tian, Christopher Woodruff and various seminar participants at the Development Lunch and Seminar at Harvard. We also thank our data contacts and providers for making this project possible. They include the unnamed e-commerce platform and Carolina Harvey. Ability to publish is not contingent on results. Please email carolina_harvey@harvard.edu for questions about the data usage agreements. The trial were registered on the AEA Social Science Registry (AEARCTR-0006725). Human Subjects Approval was obtained from Harvard University (IRB19-0647).

[†]Haas School of Business, UC Berkeley. Email: jyz@berkeley.edu

[‡]Department of Economics, Harvard University. Email: zsun@g.harvard.edu

1 Introduction

As the backbones of many developing economies, small and medium enterprises (SMEs) are heavily influenced by the penetration of digital technology. The accelerating growth of e-commerce is an exemplary case. In China, e-commerce sales grew at an average annual rate of 25 percent for the past five years ([Ministry of Commerce, 2020](#)). In 2019 alone, e-commerce sales grew by 32 percent in India and 35 percent in Mexico ([Lipsman, 2019](#)). The expansion of e-commerce provides particularly exciting new opportunities for SMEs due to the reduced entry costs and expanded market access ([World-Bank, 2016](#); [Lehdonvirta, Kässi, Hjorth, Barnard and Graham, 2019](#)). However, many challenges remain for new entrants to survive and grow after entering these online marketplaces. In particular, these new entrants need to learn about online business operations, which requires skills such as internet marketing and customer management that are different from running offline businesses. Moreover, as [Bai, Chen, Liu and Xu \(2020\)](#) have recently shown, sellers on a major cross-border e-commerce platform need to overcome sizable demand-side search frictions in order to grow.

Lifting growth barriers for promising new entrants could be beneficial for the e-commerce platforms that operate the online marketplaces. As new sellers bring more varieties to the market and increase competitive pressure on incumbents, consumers stand to have a better experience overall. The platform has incentives to be more proactive and to support promising new sellers since these actions might yield long and short-term benefits. In the long-run, a better market environment allows the platforms to attract and retain more consumers and sellers, which coincides with the platform's profit-maximizing goal. In the short-run, the platform could also benefit if sellers earn higher revenues and invest more in marketing on the platform¹. To support the new entrants, one approach that the platform could adopt is to ensure that new sellers master the necessary skills of online business operations so that a knowledge gap does not hinder their growth. In this paper, we study the impacts of one of such efforts by a major e-commerce platform: a large-scale business training program designed to help new sellers overcome growth barriers.

The e-commerce platform's efforts to promote the growth of small businesses with training follow many predecessors' footsteps. However, despite previous efforts, the effects of the training on relevant market participants are still ambiguous. For supported

¹E-commerce platforms could either only operate online marketplaces like eBay and Alibaba or both operate the marketplaces and sell to consumers directly like Amazon and JD.com. In the latter case, the platforms are in direct competition with the sellers they host. However, even in such cases, if the platforms benefit from the insights generated by third-party sellers' activities, they still have incentives to ensure promising sellers have the chance to stand out. For example, third-party sellers' performance could tell the platforms about market demands and current trends. In this case, the short-term motive for the platforms is weaker, but the long-term motive persists. As discussed later on, the platform we collaborate with does not directly sell to consumers.

firms, typical business training that teaches best business practices has mixed impacts on profits and growth (McKenzie and Woodruff, 2014; McKenzie, forthcoming). While management consulting is effective, its high costs make it difficult to scale up (Bloom, Eifert, Mahajan, McKenzie and Roberts, 2013; Bruhn, Karlan and Schoar, 2018). For non-supported firms, scarce evidence shows that spillover could be limited. McKenzie and Puerto (forthcoming) varied the treatment intensity of a training intervention at the market level and found overall market expansion effect without significant spillover on competitors. Apart from the business training, other empirical studies evaluate the spillover effects of firm subsidies (Rotemberg, 2019), credit access (Banerjee and Duflo, 2014), and microfinance (Banerjee, Karlan and Zinman, 2015). However, what is relatively understudied is the impact of interventions to support small businesses on consumers in the markets.

In this paper, we implemented a randomized controlled experiment of a business training intervention with over two million new sellers on a large e-commerce platform to answer the following questions. First, can the training help new sellers on the platform overcome growth barriers? If so, through what channels? Second, how does the training affect consumers' experience on the platform? Third, what are the welfare implications of the training on new sellers, incumbents and consumers?

The e-commerce platform with which we collaborate hosts millions of consumers and sellers. Sellers on the platform are mostly retailers that offer various types of products. We implement the training program at scale, taking advantage of the close to zero marginal dissemination costs online. In contrast to typical business training that teaches generic best business practices, our training program focuses on practical online business operation and marketing skills. We randomly assign access to the training program when new sellers register on the platform. To date, over two million sellers received access to the training. In our study cohorts, 24.9 percent of all the registered new sellers have access to the program, and 24.1 percent of sellers with access took it up within nine months.

To study the impacts of the training on new sellers, we leverage random assignment of the training access and compare the performance of treated and control new sellers. Rich administrative data also allows us to investigate the impacts on sellers' product offerings, marketing and customer service. Next, to evaluate the impacts of the training on consumers, we use rich consumer-seller matched browsing data to recover the sets of sellers that consumers visited when they search for specific products on the platform, and we exploit variations in the search results. Lastly, we use a structural model to characterize consumers' demand and estimate the platform's rules to match consumers and sellers in the search results. With the model, we decompose the welfare impacts of the training on new sellers, incumbents and consumers.

We find that the training changes the experience of new sellers, incumbents and con-

sumers on the platform. First, treated new sellers earn higher revenues. Compared to new sellers in the control group, new sellers with access to the training earn 1.7 percent higher revenues. Using random assignment of the training as the instrument, we find that sellers who participate in the training earn 6.6 percent higher revenues. The revenue gains occur mostly because treated sellers attract more consumers to their sites as they become more engaged in marketing. Specifically, these sellers participate more in pay-per-click ads and promotional events to attract consumers. In addition, treated new sellers improve their customer service quality as they adopt more supplementary services such as the AI assistant to handle customer inquiries. However, we do not find that treated new sellers have significantly higher average purchase probability among visitors or more positive customers' ratings than new sellers in the control group.

As the training helps new sellers accumulate customers, consumers are more likely to encounter these sellers. Overall, consumers have a higher purchase probability when they visit new sellers - treated or control - than when they visit incumbents. This result holds after controlling for consumers, search keyword, and search effort specific effects. To confirm it is the new sellers driving the results, we check whom the consumers choose if they do make purchases and find that they indeed choose treated new sellers over incumbents. In the meantime, we do not find adverse effects on the quality of purchases: consumers are no more likely to request returns or refunds when they purchase from new sellers, while they are as likely to make repeat purchases. Therefore, the training enables promising, higher quality new sellers to interact more often with consumers, which benefits the consumers because the matching quality is higher with these new sellers than with incumbents.

Based on the reduced-form results, we build a structural model to characterize consumer demand and recover the platform's rule to match sellers and consumers. With the model, we use variations in consumers' choice probabilities to recover underlying sellers' quality. Among the set of new sellers and incumbents that consumers visited, both treated and control new sellers have significantly higher underlying quality than incumbents. The difference in consumers' purchase probability when they encounter different types of sellers suggests that the main friction in the market is that high-quality new sellers are not being encountered by consumers often enough. We conduct a counterfactual exercise to evaluate the welfare impacts of the training. We remove training participants' access to the training by lowering the participants' chances to be found by consumers in the matching while holding consumers' search keyword and efforts constant². Doing so causes a 0.1 percent decline in consumer surplus and total sellers' revenues since consumers are less likely to interact with higher quality new sellers. The revenue drop is driven by consumers

²Share of training participants removed is consistent with what we estimated in the reduced form results. We only change the composition of the sellers that the consumers encounter in the counterfactual.

making fewer purchases as they interact with fewer new entrants without the training. In contrast, the degree of revenue re-allocation from new sellers to incumbents is limited in the absence of the training.

Our study relates to several strands of literature. First, empirical works studying growth barriers and firm dynamics for new entrants have recently shifted their focus to demand-side frictions. For offline firms, previous research highlights barriers to growth due to the lack of initial market access ([Atkin, Khandelwal and Osman, 2017](#)), slow customer accumulation ([Foster, Haltiwanger and Syverson, 2016](#); [Piveteau, 2016](#)), and the uncertainty in learning ([Arkolakis, Papageorgiou and Timoshenko, 2018](#); [Berman, Rebeyrol and Vicard, 2019](#)). Overall, our study most closely relates to the work done by [Bai et al. \(2020\)](#). The authors highlight that the demand-side search frictions limit high-quality sellers' growth in a cross-border e-commerce market. Our paper confirms that such frictions are limiting new entrants' growth in the e-commerce market. We experimentally show that the business training that closes sellers' knowledge gap and improves their marketing skills could be an effective strategy not only to lift growth barriers but also to improve consumers' experience on the platform³.

Second, we contribute to an extensive literature on business training intended to help SMEs in the developing world. [McKenzie and Woodruff \(2014\)](#) reviews this literature and finds mixed results on the effectiveness of training for offline firms. [McKenzie \(forthcoming\)](#) summarizes the more recent literature and discusses the difficulties of assessing the impacts of business training programs. Our experiment shows that training is a low-cost way to lift new entrants' growth barriers in online markets, where digital technology helps address challenges to scale and to customize. The design of the online training builds on previous lessons in the literature, incorporating large-scale customization ([Bloom et al., 2013](#); [Bruhn et al., 2018](#)) and rule of thumb style tutorials ([Drexler, Fischer and Schoar, 2014](#)). On specific mechanisms, our finding that better marketing could facilitate the growth of new entrants echoes the findings in [Anderson, Chandy and Zia \(2018\)](#), where the authors show that a business training that teaches marketing skills paves a growth-focused pathway to profits. Marketing and building customer capital is an important mechanism for growth in many settings ([Gourio and Rudanko, 2014](#); [Fitzgerald, Haller and Yedid-Levi, 2016](#)). For broader implications of efforts to support SMEs, some recent empirical studies examine "experimentation at scale" ([Muralidharan and Niehaus, 2017](#)) to evaluate effects on non-treated market participants. For business training, [Calderon, Cunha and Giorgi \(2020\)](#) randomize access to a training at the village level in Mexico and

³A line of literature investigates consumers' search frictions in various online markets and evaluates the platform's design to improve search efficiency. Some examples of empirical works include [Fradkin \(2015\)](#) (Airbnb), [Dinerstein, Einav, Levin and Sundaresan \(2018\)](#) (eBay), [Horton \(2014\)](#) (labor market), [Ursu \(2018\)](#) (Expedia) and [Chen and Wu \(2020\)](#) (AliExpress).

find no spillover effects partially due to the small sample size. Recent work by [McKenzie and Puerto \(forthcoming\)](#) uses a two-stage experimental design where the authors randomize the intensity of a business training intervention at the market level and then randomize individual businesses' access to the training within each market in Kenya. Three years after the training, not only did treated firms earn higher profits, but their success did not come at the costs of their competitors, as the market expanded in terms of sales. Instead of varying treatment intensity across markets, we contribute to this line of work by pointing out a novel channel of welfare gains with unique consumer-seller matched data. In an online market with frictions, business training could promote high-quality sellers in the matching between consumers and sellers, thereby improving both consumers' experience and market efficiency.

Third, we speak to the recent literature that examines the expansion of the digital economy ([Goldfarb and Tucker, 2019](#)) and the line of work on the roles of platforms as regulators of the various markets they host. As technology such as high-speed internet creates business opportunities ([Hjort and Poulsen, 2018](#)), many challenges remain for newcomers. For example, [Couture, Faber, Gu and Liu \(2018\)](#) show the vastly heterogeneous consumption side responses and lack of supply-side reactions as e-commerce penetrates rural retail markets. While the concerns over e-commerce platforms' market power are looming, many recent studies show how e-commerce could foster competition, improve efficiency and boost consumer welfare⁴. Across different domains, peer-to-peer platforms lower transaction costs and reduce search frictions⁵. We add to this discussion by emphasizing the importance of platforms' interventions on reducing frictions and maintaining a more equitable, competitive environment for market participants ([Tadelis, 2016](#); [Hui, Saeedi, Shen and Sundaresan, 2016](#); [Cui, Li and Zhang, 2020](#)). Interventions such as business training help the platform achieve its profit-maximizing goal and improve sellers' and consumers experience.

The rest of the paper is organized as follows. In section 2, we first describe features of the e-commerce platform and then introduce the training intervention and the experimental design. In section 3, we examine the impacts of the training on new sellers. Then, in section 4, we analyze the impacts of the training on consumers. With the reduced-form results, we build a structural model to decompose the welfare implications of the training in section 5. Lastly, in section 6, we conclude.

⁴Some empirical analysis of welfare impacts of e-commerce include [Brynjolfsson, Hu and Smith \(2003\)](#), [Einav, Klenow, Klopock, Levin, Levin and Best \(2017\)](#) and [Jo, Matsumura and Weinstein \(2019\)](#).

⁵Empirical studies situate in different platform markets and show that while frictions still exist, platforms still have the power to improve efficiency using various algorithms and mechanism, see for example [Cohen, Hahn, Hall, Levitt and Metcalfe \(2016\)](#), [Farronato and Fradkin \(2018\)](#) and [Ellison and Ellison \(2018\)](#).

2 Business Training on the E-commerce Platform

2.1 Sellers on the E-commerce Platform

In this paper, we partner with a leading e-commerce platform in China that hosts millions of active sellers and consumers. Total sales on the platform represent a sizable share of all domestic retail sales. Sellers on the platform offer various types of products, where some of the most popular sectors include clothing, cosmetics, home appliances, consumer electronics, and food. The vast majority of sellers on the platform are retailers who source their products from manufacturers or wholesalers. Unlike Amazon, third-party sellers generate a dominant share of sales on the platform. The platform earns revenues from these sellers by offering advertisements, charging commissions, and selling supplementary services. The platform's reliance on the third-party sellers implies that sellers' success aligns with the platform's profit-maximizing goal. Therefore, the platform is motivated to implement policies and programs that bring in more third-party sellers and foster their growth after entry. [Couture et al. \(2018\)](#) investigates one of such efforts where the platform brings e-commerce to rural villages.

Compared to starting an offline business, becoming a registered seller on the platform requires considerably lower monetary and effort costs. To register as a sole proprietor, a potential new seller only needs to complete the authentication with a national ID or a formal business registry but does not need to pay any registration fees. Except for several regulated sectors, the platform does not ask for certificates or charges commission fees⁶. These sole proprietors make up for roughly 97% of all active sellers on the platform. Most active sellers are highly dedicated. Running the e-commerce businesses is a full-time job and the primary source of income for these sellers. According to an online survey with selected new sellers in the training sample, the majority of respondents state that they intend to operate the online store as their full-time jobs. Appendix C discusses more details of the survey.

Despite the easy registration, sellers face growth barriers after entry. First, posting and selling products on the platform come at additional monetary and effort costs⁷. In some

⁶The regulated sectors include food, drug, medical equipment, cigarettes, liquor, infant formulas, and other products subject to public health and safety concerns. Our analysis focus on the C2C ("consumer-to-consumer") sellers. The platform also hosts a small number of "business-to-consumer" (B2C) sellers. These sellers are formally registered, have brands and completed formal applications to operate on the platform. Consumers can access these two types of sellers' sites on the most popular app the platform offers, but B2C sellers have special demarcation for their status and get preferential treatment in search rankings. B2C sellers are also much larger and some of these sellers are internationally recognized brands.

⁷Before posting products on the platform, sellers need to put down small deposits as "consumer protection fees" for potential dispute resolutions. Exact requirements differ by sectors, typically ranging from 0 to 5000 RMB. If sellers decide not to make the deposit, they can still post products on the platform, but their products will get much lower rankings in search results and will not be promoted in other channels

cases, an inexperienced new seller could spend more than 30 minutes to post a product to provide accurate descriptions and pictures that meet the platform's requirements. After posting products, attracting visitors to the sites is the prerequisite to growth for both new sellers and incumbents. The platform uses a sophisticated search and recommendation algorithm that matches consumers to the right sellers to optimize outcomes. Sellers can influence the results by actively engaging in complex strategies to compete for consumers⁸. The most common strategies are purchasing pay-per-click ads, participating in promotions that the platform regularly organizes, and recruiting celebrities to do marketing on social media. In most cases, consumers search for products rather than the sellers, and sellers compete for better rankings in each search session. Such competition in a search session is close to zero-sum⁹. Advertising and promotion will directly influence search rankings, but social media-based marketing operates through a different channel. The sheer number of competitors on the platform and the intense competition between these sellers for consumers attention indicate that marketing is a crucial component of online business operations. According to an online survey with selected new sellers in the training sample, the majority of respondents state that they intend to operate the online store as their full-time jobs. Appendix C discusses more details of the survey.

2.2 Business Training

We collaborate with the platform to implement an online business training program as a randomized controlled experiment. The low dissemination costs for the online intervention make it feasible for the training to reach many sellers. The program was officially launched on May 6, 2019, and is available since then. Over two million new sellers have accessed the training as of June 2020. The training is a standalone program independent from other operations of the platform. In particular, the training's participation and performance do not affect how the platform matches consumers with sellers in the searches.

The platform partners with professional e-commerce service providers to design the training. In contrast to typical business training that teaches generic best business practices, this training focuses on specific challenges of running e-commerce businesses to help new sellers better navigate the platform market. The training materials are organized as sequences of tasks, and each task tackles a specific challenge. In the training, the platform uses administrative data to dynamically match sellers with the most appropriate tasks based on sellers' performance and actions. Each task uses a combination of tutorials, Q&A

⁸Interviews with multiple sellers on the platform suggest marketing spending could account for a significant share of the operating costs. Larger sellers invest even more heavily than small sellers.

⁹Currently over 90% of consumers accessing the platform are from mobile devices rather than from the web. Therefore, competition for ranking is more intense because of the limited space per screen on the mobile device. On the other side, it is hard to define an obvious page break in the search results.

forums and webinars to deliver recommendations.

Participating in the Training New sellers can access the training on the official seller's portal app where the training module appears as a widget on the front page (see figure 1 panel A)¹⁰. The official app is essential for sellers to manage their stores and to communicate with the platform. Therefore, dedicated sellers do not need to invest additional efforts to access the training.

Each task singles out an area of improvement, and sellers can choose which of the tasks to try (see figure 1 panel B). The training tasks are associated with specific performance metrics along with corresponding triggering conditions. For example, a new seller triggers the task "attract more visitors to your store" if the number of visitors she had over the past 30 days is below the 40th percentile among sellers in her sector. An algorithm examines sellers' performance and assigns the most relevant tasks based on their performance and the tasks' triggering conditions.

After taking up the tasks, sellers can access detailed tutorials written by professional e-commerce service providers with whom the platform collaborates. Each task has an associated Q&A forum where sellers can directly reach out to the tutorial's authors. Some service providers also offer live-streamed webinars to communicate with the sellers directly. Tasks have varying completion time based on their difficulties, usually ranging from three days to a month. Reaching the pre-specified targets marks the completion of the tasks. Sellers need to take specific actions or outperform other sellers in their sectors. For example, a seller needs to show a number of visitors above the 60th percentile among sellers in her sector during the past 30 days to complete the "attracting more visitors to your store" task. Sellers earn short-term free accesses to certain supplementary services for each task they complete. These supplementary services support routine online business operations¹¹. The monetary value of the short-term access to the services ranges from \$5 to \$10. Sellers are unaware of the rewards before taking up the first tasks. If a seller fails to complete a task, she can always make additional attempts later¹².

Content of the Training The training tasks cover three major areas of online business operations: basic setup, marketing and customer service¹³. The first type of task focuses on teaching new sellers on how to set up online stores without running into pitfalls or violating the platform's rules. A typical task in this category teaches new sellers on how to post products on the platform. The tutorial of the task contains a step-by-step guide to

¹⁰The app is available on all major operating systems and has a web version.

¹¹As an example, sellers can access a program that allows sellers to print many customized shipping labels with one-click.

¹²The training module does not explicitly state that there is no consequence of not completing the tasks, which might deter some sellers worrying about potential negative consequences. We do not have empirical supports for the direction of selection.

¹³Appendix table 1 provides a list of tasks with explicit contents and classifications.

ensure that sellers follow the platform's rules and help the products get better search rankings¹⁴. Taking up these tasks might increase sellers' likelihood of setting up their online stores and posting products to sell. The second type of task addresses challenges to attract consumers by teaching relevant marketing techniques. Typical tasks in this category teach sellers ways to improve their product titles to get better search rankings, methods to select more suitable keywords used in pay-per-click ads, and techniques to take advantage of hundreds of sales events the platform regularly organizes. Sellers taking up these tasks may be more active in marketing and may improve marketing skills. The last type of task focuses on improving sellers' customer service quality. In this category, typical tasks introduce sellers to many supplementary tools that the platform offers to help sellers better manage their customers. In one task, the tutorial teaches them how to set up an artificial intelligence assistant to answer consumers' inquiries quickly. The adoption of the tools may help improve customer satisfaction and service quality.

The goal of the training is to help new sellers navigate the complex online business environment better. The training emphasizes pushing sellers to stay active in the market and attracting more visitors with better marketing. The focus on customer acquisition echoes earlier findings that demand-side frictions could be the main growth barrier new sellers face in the e-commerce market. While some techniques are relevant only for this platform, many marketing and customer management skills can be easily transplanted when operating other online businesses. Training does not cover more generic business practices such as managing the supply chain, finance and personnel, which are often covered in the related literature (Bloom and Van Reenen, 2007).

For some sellers, the training program helps close their knowledge gap by teaching previously unknown techniques. Others may find that the same information is already available in other sources. The training then functions as a well-structured reminder for new sellers about what should be done at certain stages of their growth trajectory.

2.3 Experimental Design and Implementation

We design and implement the training intervention as a randomized controlled experiment where we randomly offer newly registered sellers access to the training. The access is assigned immediately after the sellers complete the registration. The assignment's timing limits the available baseline information to variables collected during the registration

¹⁴Complexity of the product management system on the platform makes posting products a non-trivial task. Each posting requires sellers to describe the characteristics of the products in great detail. Such information is an essential input to the search algorithms. There are also non-explicit practices that can help promote the product. For example, at least one of the pictures should have a white background to get promotion in non-search channels.

and actions taken on the first day of entry¹⁵. Moreover, because the platform dynamically matches tasks and sellers based on sellers' performance, we cannot randomize the assignment at the task level.

The experiment officially started on May 6, 2019 and has been ongoing since¹⁶. In 2019, about 25,000 new sellers registered every day. From May 6 to October 28, 2019, during the first phase of the formal roll-out of the training intervention, we randomly selected 25% of new sellers to access the training. Later on, we expanded the share of new sellers to access the training to 35% on October 29 and to 90% on December 26. Figure 2 summarizes the timeline of treatment assignment. By June 2020, over two million new sellers have gained access to the training program. For the empirical analysis, we focus on the cohorts of sellers registered between May 6 and August 15, 2019, to track these sellers for a long enough time.

2.4 Training Take-up

Among sellers in the treatment group, 44.6% of the sellers browsed the training during the first month of entry. Within the first nine months after registration, 24.1% of treated sellers took up at least one task and 12.6% completed at least one task. The randomization was successful since none of the sellers in the control group took up any tasks. The adoption rate is typical for online products. However, it is considerably lower than other training programs offline (McKenzie and Woodruff, 2014)¹⁷.

Sellers most actively participate in the training during the first month of entry, partially because the content of the training is most relevant for brand new entrants. Altogether, 49.9% of tasks were taken-up during the first month of entry. During this time, sellers are more likely to pick tasks related to store setup and customer acquisition than those on customer management. Among sellers in the treatment group who eventually posted products, 7% took up tasks before posting. Tasks related to store setup have a higher completion rate at 58.0%, comparing to the average rate at 37.5%.

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 month¹⁸. Conditional on taking up any tasks in the previous month, around 23% of sellers took up more tasks in the following month. Figure 3 shows the share of sellers who

¹⁵Information collected during the registration process include sellers' type (registered as individual or business) and locations. For individual sellers, we also know their gender.

¹⁶The training program went into testing in April 2019, during which about 1% of new sellers received the access.

¹⁷Take-up rates for typical offline training programs are not perfect, usually in the range of 50 to 90%. Compared to the offline setting, the costs of taking up online training are much lower. However, perceived benefits might also be low, especially since many competing training programs exist in the market.

¹⁸Many sellers exit the platform after a month. 89.7% of sellers have visitors during the first month, but only 42.2% do during the third month. The rates are similar for sellers in the treatment and control group.

browsed, took up and completed the tasks in the subsequent months following the entry.

Columns 4 to 6 of table 1 summarize characteristics of treated sellers by whether or not they participate in the training. Compared to sellers in the treatment group who did not participate, sellers who took up tasks are slightly more likely to be registered as firms but are less likely to be female¹⁹. Sellers from more economically developed southern coastal provinces are more likely to participate and take up more tasks conditional on participation²⁰. Sellers from the southern provinces contribute more to total sales on the platform than sellers from anywhere else. We hypothesize that sellers who post products on the very first day of entry are better prepared. The early-movers are more likely to become training participants, but conditional on participation, they do not take up more tasks.

3 Effect of the Training on Sellers

In this section, we discuss the impacts of the training on alleviating new sellers' growth barriers. We analyze the overall impacts of accessing the training on sellers' performance and their strategies.

3.1 Data: the New Seller Panel

Our main data source is the administrative data that the platform collects on sellers' performance, strategies, characteristics, and participation in the training²¹. As mentioned, we focus on the cohorts of new sellers registered between May 6 and August 15, 2019. On average, 22,230 new sellers registered on the platform each day during the sample period.

We require newly registered sellers to log in to the official seller's portal app at least once within the first seven days of registration and have completed the entire registration process to be included in the final sample²². Since sellers will not find out whether they have access to the training program before the first login to the Seller's Portal, the login requirement does not induce selection. However, the training could potentially affect the frequency of subsequent logins²³. The full sample consists of 712,118 sellers, out of which

¹⁹As mentioned, sellers can either register with a national ID card (as an individual) or with a formal business registry (as a firm).

²⁰The coastal provinces are Guangdong, Jiangsu, Zhejiang, Shanghai and Fujian.

²¹Similar to other studies (Zhang, Dai, Dong, Wu, Guo and Liu, 2019) on online businesses, our data only includes activities that are observable online. Specifically, we do not have costs or other offline, supply-side information. Therefore our measure of performance would be revenues rather than profits. Moreover, we do not have access to sellers' outcomes on other platforms if they are multi-homing. Multi-homing is more common for larger, more established sellers than small sellers. In our sample, this possibility might not be a significant concern for small new sellers. Survey evidence previously mentioned suggests that the vast majority of the sellers have a minimal offline presence and operate solely online.

²²Sellers would have their basic information and locations recorded if they completed the registration.

²³We empirically test the impacts of the treatment on login but do not find any significant differences.

177,026 (24.8%) sellers are randomly assigned to access the training. We check the balance of the treatment assignment using sellers' characteristics collected during the registration, i.e. their types and locations (column 1 to 3 in table 1). All the characteristics are balanced across treatment and control except for sellers' registration type, where treated sellers are more likely to be firms but the difference is small²⁴. Among sellers registered as individuals, 45.3% are females. Gender distribution is balanced, so are the locations at both province and city level.

We construct a balanced monthly panel with all the sellers in the final sample. The panel spans the subsequent nine months following the registration, where we define each month as a 30-days period relative to the date of entry. The entry day is also the day of the treatment assignment. We collect sellers' performance measures on revenues, numbers of visitors and conversion rates. Conversion rate is the share of visitors who make purchases, which is the most commonly used metric for measuring efficiency. For example, the platform uses the conversion rate to evaluate its search and recommendation algorithm's efficiency. We also collect sellers' quality measures, including their customer ratings (on the accuracy of the product description, customer service and logistic), the likelihood for consumers to request refunds or returns and the frequency of platform's rules violations. On sellers' strategies, we observe their product offerings, pricing level, marketing and customer service. We do not observe actual spending on marketing. Instead, we observe some proxies for sellers' engagement, such as the number of products participating in pay-per-click ads. For sellers with treatment access, we obtain their entire history of interactions with the training program on tasks take-up and completion²⁵.

In the final sample, 40.6% of registered sellers have never posted any products to sell. For sellers who posted products or earned revenues, we obtain their affiliated sectors²⁶. For sellers without product postings and sellers who exit the market, we replace missing values with zero for outcome measures such as the number of visitors and revenues²⁷. Conversion

²⁴26.4% of sellers in the control group are registered as firms. The share of firms in the treatment group is 26.8%.

²⁵As described in section 2.2, the assignment of tasks is individualized with temporal variations. We do not keep track of the tasks that were assigned to sellers daily.

²⁶Sellers' sectors are determined by the products they posted and sold. Therefore, sellers will not have a sector affiliation if they do not post anything. Moreover, a significant share of sellers is labeled as selling second-hand products, which are treated differently in the search results. We group sellers without sectors, sellers who sell second-hand products and sellers selling unclassified products together. Since sellers could change their sectors, we use the first sectors that the sellers identify with as their affiliated sectors. The results do not change if we use sellers' most frequently affiliated sectors.

²⁷Platform automatically remove a registered seller from the platform if the seller does not have any active product posting over the past four weeks. Notice that there are no requirements on the number of visitors attracted or revenues made. When removed, the seller's site is inaccessible, and the platform stops collecting data. Sellers have the option to re-open their store, at which point the platform will start to collect the information again under the same ID. We do not exclude sellers without product posting in the main analysis because encouraging sellers to post products is an important part of the training. Hence the decision to post products could be affected by the training. We do not find a significant difference in the likelihood of posting products

rates are undefined if the sellers attract no visitors. Similarly, the quality measures such as customer ratings are undefined if the sellers do not have any orders. We leave these variables as missing. Distribution of the outcome variables such as the number of visitors and revenues is extremely skewed. For the main analysis, we convert these variables to log scale²⁸.

3.2 Impacts of the Training

Overall Impacts We first evaluate the overall intent-to-treat (ITT) effect with the new seller panel. For seller i during (relative) month m belonging to an entry-date cohort c with affiliated sector s , we run the following specification:

$$Y_{imcs} = \beta Treatment_i + \alpha_m + \alpha_c + \alpha_s + \epsilon_{imcs} \quad (1)$$

$Treatment_i$ is an indicator for having access to the training, α_m , α_c and α_s are month, cohort and sector fixed effects. Standard errors are clustered by seller.

Table 2 presents the estimated results on sellers' performance. Access to the training leads to a 5% increase in the likelihood of earning revenues (column 1) as well as a 1.7% increase in revenues earned (column 2)²⁹. Since revenues are zero for 60.5% of seller-month pairs, the unconditional average revenue is close to zero. Restricting the sample to seller-month pairs with positive revenues (column 3), sellers in the treatment groups earn 2.6% higher revenues. Treated sellers earn higher revenues because they attract more visitors to their sites. Treated sellers attracted 1.3% more visitors to their sites (column 4), and they have 0.8% more consumers making purchases (column 6). Conditional on having some visitors, treated sellers attract 2.4% more visitors (column 5) and consequently have 1.8% more consumers making purchases (column 7). However, we do not find a significant improvement in treated sellers' conversion rates, as shown in column 8. We use consumer-side data to explore the conversion in section 4. Figure A1 presents the quantile treatment effect on log revenues separately for each month. The variations of treatment effects by different quantiles are small. The treatment effect is slightly larger for sellers in the middle of the revenue distribution, especially for sellers on the edges of earning revenues. Since

in the subsequent months, but we cannot rule out the possible treatment effect.

²⁸59.5% of seller-month observations have no visitors. To avoid dropping most of the sample, we add one to the number of visitors, revenues and other performance measures before taking logs. We also use inverse hyperbolic sine transformation for the main outcomes and reach a similar conclusions.

²⁹Table A1 presents the treatment effect on revenues in levels with different winsorization thresholds. To address the large concern due to large number of zeros, we also use inverse hyperbolic transformation for the revenues (column 4), the result is similar to using log with plus one. Besides the specification with raw revenues, all the estimated effects are positive, but not all significant. The results are very sensitive to extreme values at the top of the distribution. Because of the training content, we do not expect that the training could have meaningful impacts on these top sellers.

the impacts of training on seller posting products are limited (table A3 column 1), we restrict the sample to sellers that have ever posted product in the first nine months and find similar results as in the full sample (table A3). In this subsample, treated new seller earn 2.4% higher revenues than control sellers (column 3).

Table A7 summarizes the impacts of training on sellers' observed quality metrics. We do not find treated sellers to have significantly higher customer ratings than sellers in the control group. For all three types of ratings, namely accuracy of product descriptions, customer service and logistics, treated sellers obtain slightly higher scores than control sellers, but the difference is not significant. These two groups of sellers also have a similar percentage of refunds and complaints, the share of positive reviews, and the frequency of violating the platform's rules. The point estimates suggest that sellers with training access weakly outperform control sellers for most quality metrics.

Focusing on sellers who participate in the training, we use the random assignment of the training access as an instrument for actual participation and an indicator for taking up any tasks during the sample period as the first stage variable. Column 1 of table A5 shows the first stage specification as equation 1. On average, 25.7% of treated sellers took up some tasks. The rest of the table presents the two stage least square estimates on sellers' performance. For sellers taking up the tasks, they earn 6.6% higher revenues and attracted 5.2% more visitors to their sites. We do not find significant impacts on conversion rates. Comparing sellers who took up some tasks to those who did not in the treatment group, it is obvious that training participants significantly outperform non-participants along with all performance measures (table A6).

We analyze temporal variations of the treatment effects in ?? and conclude that the temporal variations are small. The experimental design limits the baseline heterogeneity we could capture to basic types, locations and actions on the first day. We discuss the heterogeneous treatment effects in details in appendix A.1. In a nutshell, we do not find a significant difference by sellers' registration type or level of preparedness³⁰. Instead, the offline business environment could play a role. Sellers from less-developed regions are less likely to participate in the training, and the training is less useful for these sellers.

The magnitude of the treatment effect on revenues is positive but small. Over the nine months, treated sellers earned \$1.8 million higher total revenues. Assuming the treatment effect is of a similar magnitude for all cohorts of new sellers, all two million treated sellers combined could earn about \$4.7 million higher revenues. Higher revenues could be a result of market expansion and business stealing. We discuss these two possibilities in section 4 and section 5.

³⁰We measure the level of preparedness by whether or not sellers post products on the first day of registration.

Sellers' Strategies Next, we discuss how the new sellers' strategies change relates to the observed increase in revenues and the number of visitors. We focus on observable strategies, including pricing, product offerings, marketing and customer service. We find treated sellers change their marketing strategies and slightly improve customer service quality, but do not behave differently on other dimensions.

Table 3 presents estimated coefficients β in specification 1 on the treatment indicator, where each cell corresponds to a separate regression on a specific outcome variable. Although many training tasks focus on technical and administrative barriers sellers may face when setting up their online stores, the training has limited impacts on incentivizing market participation in terms of posting products or putting down security deposits. The platform strongly recommends sellers to put down small deposits to cover potential consumers' losses³¹. The results suggest that closing the knowledge gap alone is not enough since many offline constraints are still limiting. For example, sellers need to find sources of supplies and have available funds to cover operation costs. Eventually, 65.0% of sellers in the treatment and 64.8% of sellers in the control group posted products, but the difference is not significant. Moreover, treated sellers are not accelerating the speed of posting products among the subset of sellers who posted products after the first day of entry (table A2 columns 1 and 2). Similarly, treated sellers are no more likely to put down security deposits (table 3). As expected, sellers do not behave differently in terms of the number of products offered, the likelihood of moving into different sectors or setting different prices. These strategies are not covered by the training and are more affected by the offline environment.

Treated sellers are more likely to follow the platform's recommendations to adopt free supplementary tools that help improve the quality of customer service³². Specifically, we find that treated sellers have a slightly shorter average response time when consumers make inquiries and have higher conversion rates among consumers who made inquiries (table 3 section on customer service). These results are driven by treated sellers' higher likelihood of adopting the AI-backed customer assistant to help answer consumers' basic questions³³. Although treated new sellers do not have significantly higher customer ratings of their service quality, improved customer service quality could still contribute to treated new sellers' revenue increases as indicated by higher purchase probabilities among consumers who make inquiries.

Training helps improve treated sellers' marketing skills (table 3 section on marketing).

³¹Sellers can still post products if they do not put down deposits. However, these sellers will get much lower rankings in search results.

³²The supplementary tools that the training promoted focus on improving customer service quality and are not the same as those given out as rewards for completing the training tasks

³³Currently we do not have data on sellers' actual subscriptions to the AI assistants and other supplementary service.

Treated sellers have more products participating in pay-per-click ads where sellers bid for better search rankings with specific search keywords, and they have a higher share of visitors coming from the paid channels³⁴. In addition to advertising, treated sellers are also more likely to participate in the limited-time promotional events that the platform regularly organizes³⁵. The products on sale get preferential treatment in the search rankings and additional exposure since consumers can find these products from other channels besides the main search and recommendation program. More active participation in marketing could involve higher costs for the treated sellers. Because we do not observe actual spending on marketing, we are unable to claim that the training definitively increases sellers' profits.

Marketing is an indispensable part of online business operations, and attracting visitors is the key to success. However, marketing capacity and quality might not perfectly correlate (Hu and Ma, 2020). The training intervention that either improves sellers' marketing skills or raises sellers' awareness of marketing helps new sellers accumulate more consumers³⁶. As a result, incumbents and non-treated new sellers could have fewer visitors. Such reallocation has ambiguous implications for consumers. The ambiguity is a common concern for typical training interventions that promote specific groups because of potential adverse selection. In such cases, consumers might interact with lower quality firms more often due to the interventions. We assess the impacts of the training on consumers in section 4.

4 Effect of the Training on Consumers

As discussed in the previous section, the entrepreneur training causes treated new sellers to attract more consumers to their sites and improves customer service quality. As a result, consumers' experience on the platform could be affected when consumers interact with different types of sellers. In this section, we evaluate the impacts of the training on consumers by answering the following questions. First, when a consumer visits more treated or control new sellers during a search session, is she more likely to find what she needs and make a purchase from some sellers in the set? Second, when a consumer visits both new sellers and incumbents in a search session, from whom is she more likely to purchase? Changes in consumers' search experience could affect overall purchase probability (mar-

³⁴The usual paid channels include pay-for-clicks ads in search results, advertising spots in the AI-powered recommendations, headline figures and social media campaign.

³⁵The classic sales events the platform regularly offers are product specific limited-time discounts. Sellers select the set of products to participate and submit applications to the platform to be included.

³⁶Since we do not observe sellers' operating or marketing costs including their spending on the platform, we are unable to identify if sellers' investment in the marketing ended up yielding higher profits. However, the bottom line is that training should have no direct impacts on sellers' available financial resources to invest in marketing, even though sellers might shift investments from other channels to marketing.

ket expansion) and choices within a set of visited sellers (market allocation). Empirically, we use the detailed consumer-seller matched browsing data and exploit variation in sellers' composition that consumers visit, given their interests and search efforts. The set of sellers that a consumer visits in a search session is determined by the platform's search algorithm and her own browsing behaviors. While both are affected by consumers' characteristics and past behaviors, the search algorithm has some random assignment procedure when matching new sellers³⁷. For consumers, explicitly picking out new sellers from the search results is nearly impossible without visiting sellers' sites.

4.1 Training and Market Expansion

Sample Construction To evaluate how interacting with new sellers affect consumers' subsequent experience, we identify consumers with the same interests (searching the same query) and the same search efforts (visiting the same number of sellers) on the same day. These consumers ended up visiting different sets of sellers. To be precise, we construct a consumer-search session sample using administrative data from the platform in the following steps (summarized in A3). First, we draw a random sample of new sellers from the experimental sample³⁸. Next, we identify a set of consumers who have visited the sampled sellers' sites between August 1 and December 31, 2019, and obtain the search keywords they used to find these sellers. Then, we find another set of consumers who searched the same keywords on the same day and visited the same number of sellers as the previous group of consumers did but only visited incumbents³⁹. As mentioned, sellers that each consumer visited are not randomly selected because both search rankings and consumers' browsing behaviors are endogenous. However, without the platform's explicit promotion in the search outcomes, it is almost impossible for consumers to specifically look for new sellers when searching for products⁴⁰. Therefore, our empirical strategy exploits

³⁷We do not have access to the actual search algorithm, and the algorithm is too complex and too dynamic to be summarized as simple rules. Discussions with internal engineers suggest that the algorithm does have random components when matching for new products and new sellers. However, there are no explicit rules on how such matching associates with consumers' characteristics.

³⁸We have to take a sub-sample from the full sample mainly due to computational reason. All sellers in the experimental sample combined could attract more than 10 million visitors on a typical day.

³⁹We limit the number of matched consumers per search query-search effort-date set at 50 for the computational reason. Such sampling implies that we over-sampled sellers with fewer visitors and keywords that are less popular. We also require that consumers visit at least three sellers because consumers who visit only one or two sellers might have different mindsets. On the one hand, those consumers might not be serious about purchasing because of their limited search efforts. On the other hand, those consumers might be looking for very specific sellers, especially those they have purchased previously, as the algorithm tends to promote these sellers. We are unable to separate these two possibilities, and these motivations could result in opposite purchasing behaviors.

⁴⁰A consumer can only determine whether the seller is newly registered if they visit the seller's site and click on a specific section about the store's necessary information. In the search results, besides prices, product titles, and category-specific labels, the displayed information also include the cumulative number of orders and sellers' locations. The cumulative number of orders could correlate with sellers' age, but the correlation

variations in the platform's search algorithm, which determines the pool of sellers that consumers could access. For these two groups of consumers who visited some new sellers or only visited incumbents, we obtain their search and purchase history for a month before and after the event. We aggregate the final sample to the consumer-search session level. A search session is a search query-search efforts-date combination. Our primary outcome measure for each consumer-search session is whether the consumer places an order with any sellers in the set and her total spending. Therefore, each observation corresponds to a specific consideration set that the consumer uses to make a purchase decision. We obtain sellers' pricing, the number of products offered, and their customers' ratings as the main control and average these measures to consumer-search session level. For the consumers, we obtain their search and purchase history around the time of the session, which allows us to summarize consumers' characteristics and preferences.

The final sample consists of 1,381,273 consideration sets (consumer-search sessions) spanning 153 days in the second half of 2019. Table A10 summarizes the main variables. The complete sample consists of 515,748 consumers. On average, each consumer appears in 2.68 search sessions, where 44.2% of consumers appear in only one search session. These consumers searched 13,593 distinct keywords, spanning most of the popular sectors on the platform. In 18.8% of the search sessions, consumers placed an order from some sellers they visited on the same day. Average consumers visited 4.89 sellers per search session. In 61.1% of the sessions, consumers visited three or four sellers.

Interacting with new sellers is rare because new sellers only account for a relatively small share of all the sellers on the platform⁴¹. In our sample, consumers visited treated new sellers in 6.5% of search sessions and control new sellers in 9.2% of sessions⁴². New sellers that appear in consumers' consideration sets are larger and more successful than average new sellers in the experimental sample, especially since a significant proportion of the new sellers have none or very few visitors. Therefore, the consumers' results do not speak to the average new sellers, but only new sellers in the top end of the distribution whom the consumers can encounter when searching on the platform.

The characteristics of consumers who visited new sellers are different from those who only visited incumbents. In these search sessions, consumers visit more sellers. Consumers also spend more money and search more intensively in the week before the search

is not perfect. Consumers can search specifically for new stores by including relevant keywords in the search query, but the search results do not always respect that intention. Moreover, we control the search keyword in the primary analysis.

⁴¹In 3.2, we analyze the number of visitors *new sellers* attracted and compare between treated and control *new sellers*. In this section, we compare new sellers to *incumbents*.

⁴²As described in the sample construct, the probability of interacting with new sellers are calculated based on the sample. It may not be the same as the probabilities on the entire platform. We over-sampled sellers attracting fewer visitors and less popular keywords. The actual probability of having some new sellers in the consideration sets over all the search sessions should be lower than what we have here in this sample.

event. For search sessions involving new sellers, consumers visited three or four sellers in only 51% of these sessions (figure A4). Conditional on search effort, keyword and date, consumers visiting some new sellers spend 38% more than consumers only visiting the incumbents. We address the selection on the consumer side in the empirical analysis with consumer and search session fixed effects and a rich set of controls for consumers' characteristics at the time.

Empirical Strategy When a consumer visits more treated or control new sellers in a search session, is she more likely to find what she needs and make a purchase? We answer this question using the consumer-search session sample with the following specification:

$$Y_{is} = \beta^t \mathbf{T}_{is} + \beta^c \mathbf{C}_{is} + \mathbf{X}_{is} \gamma + \alpha_s + \alpha_i + \epsilon_{is} \quad (2)$$

for consumer i and search session s . The main variables of interests are \mathbf{T}_{is} and \mathbf{C}_{is} , defined as the share of treated or control sellers in the consideration set for consumer i in search session s . Since neither the platform nor consumers know if new sellers have access to the training, the comparison between treatment and control new sellers is not subject to selection bias. We add the fixed effect α_s to capture variations across search sessions. Controlling for the search query addresses heterogeneous responses when consumers search for different types of products. For example, consumers tend to explore more options when searching for horizontally differentiated products such as clothing than when they search for vertically differentiated products such as home supplies. Controlling for the number of sellers that consumers visited in a search session helps alleviate two competing concerns. On the one hand, consumers visiting more sellers are more dedicated as they invest more search efforts, which increases their likelihood of purchase. On the other hand, consumers browsing more sellers might also be less satisfied with their previous matches, forcing them to search more intensively and lower their overall probability of purchase. Therefore, the consideration set's size serves as a proxy for the searches' intensity and quality. Similarly, we control for the date because the platform organizes many promotional events year-round, and these events could have differential impacts on new sellers and incumbents. We control consumer time-invariant characteristics with α_i to address consumers' idiosyncratic variations in their pickiness, purchasing power, experience and familiarity with the platform. In addition, we add control variables \mathbf{X}_{is} that includes sellers' average pricing level, the number of products offered and sellers' ratings in the consideration sets⁴³ as well as consumer i 's total spending and search intensity in the previous week. This specification does not control for consumers' search query specific preferences, e.g. a consumer might be unusually picky when choosing printing paper even though her

⁴³Ratings are determined by the cumulative number of positive reviews sellers get. Hence they reflect sellers' size more than quality. The price level is measured in the log scale.

fellow shoppers view printing paper as a homogeneous product. Such idiosyncratic preferences could affect consumers' search behaviors, but it is unclear how such taste difference would bias the way consumers interact with new sellers.

Results Table 4 presents the results of specification 2 on the purchase probability and log order size. On the day of the visit, compared to visiting a set of sellers with incumbents only, a consumer is 11% more likely to purchase if her consideration set only consists of new sellers (column 1). Conditional on consumers' recent behaviors and sellers' characteristics, consumers are 1.9% more likely to make a purchase with one more treated or control new seller in the set (table A12 column 1). Incorporating the possibility that consumers might place an order in later days, we reach a similar conclusion. Having more new sellers in the consideration sets significantly increases the likelihood for consumers to make purchases by 1% in the next three days (column 2 in table 4) and 1.1% in the following week (column 3). Having a search set with only new sellers increases total spending by at least 6.5% (column 4 in table 4), while having one extra new sellers in the consideration increases total spending for the specific search session by 1.1% (column 4 in table A12). Higher purchase probability indicates that consumers benefit from better matching quality in the searches. Higher intensity of interactions with new sellers therefore leads to market expansion.

Comparing control new sellers to treated new sellers who appear in consumers' consideration sets, we find that for the most part interacting with these two types of new sellers leads to a similar increase in consumers' purchase probability. We test the difference in the estimated coefficients β^c and β^t and find that the difference between these estimated coefficients is small and insignificant except for log total spending (column 4 table 4). If we use the number of treated new sellers as the main explanatory variable rather than the share of the new sellers in the consideration sets, there is no significant difference between estimated β^c and β^t (table A12). Since the training increases new sellers' likelihood of appearing in the consideration sets and improves their customer service quality, there is both selection and treatment effect. Overall we do not find adverse selection because of the training since interacting with treated new sellers is associated with higher matching quality similar to visiting non-treated new sellers. More importantly, consumers' search efficiency rises as they interact with both types of new sellers compared to only incumbents in the consideration sets. The gap in purchase probability implies that there could be significant frictions that hinder new sellers' growth.

4.2 Training and Market Reallocation

In this section, we ask, when a consumer visits both new sellers and incumbents in a search session, from whom she is more likely to purchase? We construct a consumer-seller matched pair sample to analyze consumers' choices and the resulted market allocation.

Data and Sample Construction To construct the sample, we restrict the attention to consideration sets where consumers have visited at least one new seller and have made purchases from some sellers in the sets. We construct a consumer-seller pair level sample where each pair is associated with a specific consideration set hence a search session defined by search query-search effort-date combination. For each consumer-seller pair, we use whether or not a consumer makes a purchase from the specific seller within a given period of time as well as the size of the order as the main outcomes. When a consumer makes a purchase, we obtain quality measures on whether or not the consumer requests returns or refunds and whether or not she makes repeat purchases from the same seller in the following month. To control for the impacts of sellers' strategies and characteristics, we again collect data on sellers' pricing level, the number of products offered and the ratings⁴⁴.

The final sample consists of 300,273 consumer-seller pairs corresponding to 42,004 consumers' 61,280 consideration sets (defined per search session - consumer) spanning 153 days. Table A11 presents summary statistics of the main variables used. These search sessions are the same as those we analyzed in section 4.1, except in this sample, we obtain more detailed consumer-seller level interactions associated with these sellers. 98,631 sellers appear in the sample, with 3,687 of them belong to the treatment group and 7,329 sellers in the control group. As discussed earlier, these new sellers are highly selected as they are much larger and more active than average new sellers. For consumers making purchases on the same day, average sellers have 18.3% chances to be selected. The average order size is \$31.5, and the median order size is \$20.4⁴⁵. For 9.1% of purchases, consumers request refunds or returns. Consumers place repeat orders in the following month in 4.9% of the cases. We use quantity weighted average prices at the seller level as a proxy for seller's pricing level⁴⁶.

⁴⁴We also restrict the sample to the set of consumers who appear in at least two search sessions. Since we control for consumer fixed effect in the main specification, including consumers with one search session does not change the results.

⁴⁵Throughout the paper, we assume 1USD equals 7RMB. These statistics do not correspond exactly to the average of pay amount in table A11 because we exclude the observations where consumers did not place any orders with the sellers.

⁴⁶Almost all the sellers on the platform offer multiple products. However, in the search sessions, consumers access a particular seller from the product page. However, due to data limitation and the pricing strategies' complexity, we do not observe real-time prices that the consumers observe. Moreover, since we aggregate the outcomes to the seller level and consumers could browse and purchase multiple products from the sellers they visit, it is unclear how to aggregate prices without observing what products the consumers visit. Therefore, we use seller level, quantity weighted pricing as a proxy for the seller's pricing level. The current algorithm encourages sellers to design their pricing strategies to target a specific group of consumers with comparable purchasing power. Therefore, the variations in prices across products in a store could be more limited than across seller variations.

Empirical Strategy To test the impacts on allocation, we use the following specification:

$$Y_{ijs} = \beta^t \mathbf{T}_j + \beta^c \mathbf{C}_j + \mathbb{X}_j \gamma + \alpha_{is} + \epsilon_{ijs} \quad (3)$$

The specification includes consideration set α_{is} fixed effects to address consideration set level heterogeneity. Consideration set is a set of sellers a consumer eventually chooses from in a search session. In this way, we control for consumer - search session specific idiosyncratic variations and only evaluate consumers' choices between sellers in the sets. The outcomes of interests are \mathbf{T}_j and \mathbf{C}_j , indicators for whether or not the seller j belongs to the treatment or control group. To test the differential impacts of interacting with treated and control new sellers, we compare the coefficients β^t and β^c . As before, we include a set of seller level controls \mathbb{X}_j on seller j 's pricing, the number of products offered and its ratings.

Results The estimated results are summarized in table 5. By restricting to consideration sets that consumers purchased from, we control for the possibility of market expansion and evaluate the market allocation between new sellers and incumbents. Consumers are more likely to choose new sellers, especially new sellers in the treatment group, over incumbents that appear in their consideration sets. Specifically, on the day of the visit, consumers are 5.9% more likely to choose a treated new seller than an incumbent in the same consideration set (column 2). Comparing treated new sellers with control new sellers, consumers' purchase probability with the former is significantly higher by about 4.2% (column 2). The differential impacts of interacting with treated or control new sellers are similar if we use purchase within a week or amount spending as the outcomes. In columns 2, 4 and 6, we include seller level controls. The coefficients are quantitatively similar. Results here suggest that consumers' higher purchase probability with new sellers is not driven by sellers' charging different prices.

Using the consumer-seller matched sample, we show that training benefits the consumers by improving their matching quality and confirm that the improvements on matching are because of consumers interacting with new sellers. The results shut down the potential negative selection induced by the training. As discussed in section 3.2, training helps treated new sellers to attract more sellers to their sites. Therefore, from the consumers' perspective, more consumers visit treated new sellers and treated new sellers are more likely to appear in consumers' consideration sets. Consequently, if the consumers do not visit more sellers or search more intensively, control new sellers and incumbents might be less likely to appear in the consideration sets. Hence, by changing the compositions of sellers that consumers visited, the training could cause market expansion and market reallocation. Such a shift in the allocation of consumers' visit likelihood could benefit the consumers but may come at the costs of control new sellers and incumbents. Current ev-

idence suggests that new sellers have higher quality than incumbents so that reallocation improves overall efficiency. We quantify and decompose such impacts in section 5.

Besides purchase probability, table 6 shows how post-purchase experience may differ when consumers purchase from new sellers instead of the incumbents. We restrict the sample to consumer-seller pairs where consumers make purchases on the day of visit and use specification 3 to evaluate the consequences on the likelihood of return, refund and repeat purchase⁴⁷. Overall, placing an order with new sellers does not significantly negatively impact consumers' likelihood to request returns or refunds (columns 1 and 2). Moreover, there is no difference in the likelihood of making repeated purchases from new sellers versus incumbents (column 3). These results show that while consumers are more likely to purchase from new sellers, their experience with the new sellers is no worse than their experience after purchasing with incumbents. Hence, higher purchase probability does not come at the cost of lowering purchase quality.

5 Decomposition of the Impacts of the Training

Motivated by the reduced-form evidence, we use a structural model to characterize growth barriers new sellers face. Our model focuses on consumers' purchase decisions given their consideration sets and uses variations in consumers' choices to identify the primary sources of frictions: the mismatch between sellers' true quality and the total number of visitors they manage to obtain. The model uses a flexible function to characterize the platform's matching rule based on observable sellers' performance, namely their lagged number of visitors and conversion rates⁴⁸.

To evaluate the welfare implications of the training, our model emphasizes how the training changes the sellers' likelihood to appear in consumers' consideration sets that consequently changes matching quality and welfare. Because of the training, treated new sellers adjust their strategies accordingly to attract more visitors to their sites. Treated new sellers capture more attention from the consumers, which implies that the incumbents and control new sellers will be less likely to enter consumers' consideration sets. That is, training induces changes to the composition of sellers in consumers' consideration sets. We expect the negative spillover on the non-treated new sellers to be limited because the vast majority of sellers on the platform are incumbents. If reallocation occurs randomly among all the sellers, then by chance, most of the market reallocation will come from the

⁴⁷We could not add seller level controls because the model cannot be identified due to a large number of fixed effects and control variables. Hence the control variables only include consumer baseline characteristics.

⁴⁸We do not directly model supply-side responses because empirical changes in strategies that could affect consumer demand such as price adjustments and product introductions are rare. Instead, most of the actions that the sellers take concentrate on marketing, which is captured by the number of consumers they attracted in previous periods.

incumbents. In the model, we take consumers' search and browsing behaviors as given and only consider the market reallocation among different types of sellers⁴⁹.

5.1 Model Setup

Consumer Demand As mentioned, we do not explicitly model consumers' search process and how they arrive at the observed consideration sets. Instead, we take these consideration sets as given and consider the conditional purchase decisions. Specifically, a consumer $i \in I$ searches a query and generates a consideration set K_i . Each set K_i consists of a group of sellers $j \in K_i$ that could be either new sellers (treated or control) or incumbents. Consideration sets K_i could have different sizes, which we do not model. Consumer i solves the following maximization problem to choose from which seller $j \in K_i$ she wants to purchase:

$$\max_{j \in K_i} U_{ij} = V_{ij} + \epsilon_{ij} = x_j \beta - \alpha p_j + \xi_j + \epsilon_{ij} \quad (4)$$

p_j is the price level seller j charges, and x_j is the set of strategies and characteristics seller j adopts that might affect consumers' purchase decisions. ξ_j is the unobserved seller j 's underlying quality and is our main object of interest. Sellers with higher quality ξ_j yield higher utility for all consumers visiting their sites. The main source of friction the model captures comes from the mismatch between sellers' quality ξ_j and their likelihood to appear in consumers' consideration sets. We do not explicitly model the treatment effect. However, since we use post-treatment data to estimate ξ_j , it captures both sellers' underlying quality and effect of the treatment. ϵ_{ij} is the I.I.D. consumer-seller idiosyncratic preference that reflects unobservable components affecting consumers' decisions.

We assume that ϵ_{ij} follows a type I extreme value distribution and consumers' outside option of not purchasing from any seller have zero utility. We get the following familiar logit formulation

$$P_{ij} = \frac{\exp(x_j \beta - \alpha p_j + \xi_j)}{1 + \sum_{k \in K_i} \exp(x_k \beta - \alpha p_k + \xi_k)} \quad (5)$$

where 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

⁴⁹On the one side, treated new sellers are more likely to appear in consumers' search sets, resulting in market reallocation from incumbents and control new sellers to treated new sellers. By closing the knowledge gap, training helps to ensure that new sellers can participate in the competition for consumers' attention. On the other hand, consumers may change their search behaviors in response to the changing composition of the sellers they visit. Consumers need to spend fewer efforts to search if the quality of the matches improves, which is welfare improving for the consumers but could limit fellow sellers' chances to appear in consumers' consideration sets. Alternatively, matching with higher quality sellers in the search sessions may induce consumers to do more searches in the future, raising the likelihood of purchases from other sellers. This channel could potentially benefit all the sellers if the market for consumers' attention expanded. Compared to the market reallocation and the direct impacts of changes in the consideration sets on consumers' purchases, changes in consumers' search behaviors are second-order.

specific price coefficient α_s .

Endogenous Strategies One major 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 estimated coefficients. 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

$$\begin{bmatrix} p_{jt} \\ \text{prod}_{jt} \end{bmatrix} = \mathbf{Z}\beta^{fs} + \xi_j^{fs} + \xi_t^{fs} + \xi_s^{fs} + \epsilon_{jt} \quad (6)$$

ξ_j^{fs} , ξ_t^{fs} and ξ_s^{fs} are seller, time and sector fixed effect. The set of instruments Z are variables that capture the stringency of the platform's rule enforcement. These instruments include frequency of different types of rule violations and shares of sellers identified as frequent rule violators in the corresponding sectors⁵⁰. 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 that sellers obey relevant state regulations and maintain a well-functioning market. In the cases of rule violations, the platform could downgrade sellers in the search rankings, remove access to sellers' products or even their sites and, in some cases, call for legal solutions. Such punishment could have significant impacts on sellers' business operations. The platform frequently adjusts the design and enforcement of the regulations as the business environment fluctuates. We exploit changes in the strictness of rule enforcement at the sector-month level. When the platform strengthens the rule enforcement, sellers could be more cautious about posting more products, charging extreme prices or engaging in unruly promotions. On the other side, when the platform enforces stricter rules, rule-obeying sellers could benefit as the platform regulates their unscrupulous competitors' behaviors, allowing the rule-obeying sellers to increase their market shares.

Matching The rule for matching sellers and consumers is the most important device that the platform has to improve consumers' experience and to support promising sellers. We simplify the complex matching rules used in the search and recommendation algorithm by highlighting the reliance on the previous period's conversion rates and the number of visitors. In the matching, 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 impacts of the training on attracting consumers. Sellers' strategies, including their participation in the training, do not directly factor in the matching process, but they can affect the lagged results. We capture the evolution of the number of visitors over time with the following

⁵⁰We use number of visitors each seller attracts as the weights when we aggregate these instruments to sector level. Rule violations come in different level of seriousness and we separately calculate them by type.

model:

$$\mathcal{T}_{jt} = f(\mathcal{T}_{jt-1}, C_{jt-1}, C_{jt-2}) \quad (7)$$

\mathcal{T}_{jt} is the current period total number of visitors for the seller j and \mathcal{T}_{jt-1} is the previous period's number of visitors⁵¹. C_{jt-1} and C_{jt-2} are the conversion rates in the previous two periods.

5.2 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 basic model.

Consumer Demand We use the consumer-seller pair sample to estimate the demand parameters in particular sellers' ξ_j . Appendix B.1 describe the detailed sample construction process. Due to computation constraints, we sample a subset of sellers for actual estimation. The final sample is a seller-consumer matched pair dataset, and we explore seller level variations. x_j includes the number of products sellers offer as well as sellers' ratings and p_j is the average price level that sellers charge⁵². 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 .

We jointly estimate consumer demand and sellers' strategy using the instruments described above. The final set of instruments include 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 B.2.

Matching To estimate the matching rule, we use the the new seller panel. 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:

$$T_{jts} = f(T_{j,t-1}, C_{j,t-k}) + g(I_{j,t-k}(t-k=0)) + h(I_{j,t-k}^c(C_{j,t-k}=0)) + \xi_k^{trf} + \xi_t^{trf} + e_{jts} \quad (8)$$

$C_{j,t-k}$ are lagged conversion rates in past periods $k = 1, 2$, $I_{j,t-k}(t-k=0)$ is an indicator for the initial two periods⁵³ and $I_{j,t-k}^c(C_{j,t-k}=0)$ are indicators for lagged conversion rates being exactly zero. The specification also includes product category, calendar time

⁵¹We again convert the number of visitors to log scale to improve the fit because of the skewness of the distribution.

⁵²As 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 offer as well as any sales or promotions sellers offer.

⁵³Sellers have no past history on number of visitors and conversion rate in the initial periods hence are subject to different matching rules.

and relative month fixed effect. In the baseline specification we start with linear functions for $f(\cdot)$, $g(\cdot)$ and $h(\cdot)$. In this setup, all the right-hand-side variables are determined in the previous periods.

5.3 Estimation Results

Following the procedure described in section 5.2, we estimate the baseline model and quantify the welfare implications of the training program with counterfactual exercises.

Panel A of figure A5 plots the distribution of the estimated sector fixed effect ξ_s . Panel B of figure A5 plots the distribution of the price elasticity for sellers in the sample. The average price elasticity is -0.22. The elasticity we estimate here is much smaller than the typical elasticities observed in the literature (Broda and Weinstein, 2006). The difference occurs because consumers are choosing between products in their consideration sets, rather than choosing among all the products offered in the market. When constructing the consideration sets, consumers already restrict their choices to a narrower range of prices. In our sample, 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. Therefore the estimated price elasticity with the consideration set is lower.

The main parameters of interest are ξ_j for $j \in J$. ξ_j captures the sellers' post-treatment underlying quality, incorporating selection and treatment effect. The final sample contains 52,241 sellers, out of which 8.33% are new sellers. 28% of sellers in the sample have some purchase records. Table 7 presents the distribution of ξ_j for different subset of sellers. On average, new sellers have higher ξ_j than incumbents, both among sellers with and without purchase records. The difference is summarized in table 8. Among sellers with purchase records, the estimated quality of new sellers is 17.9% higher than the incumbents, and among those without purchase records, new sellers' estimated quality ξ_j is 11.2% higher. As shown in the distribution of ξ_j (figure 4), the difference is not driven by a small set of extremely high-quality new sellers whose ξ_j land on the right tail of the distribution. Instead, the results are driven by median new sellers having higher underlying quality than the median incumbents. The sample of new sellers is a highly selected subset from all new sellers. 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.

5.4 Counterfactual: Welfare 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 sample sellers from the sellers' pool associated with the search keyword according to certain 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 to appear 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 B.3.

Baseline: Predicted Number of Visitors 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 5.3. The current logit-specification allows us to calculate consumer surplus as

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

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

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

Sellers will not earn any revenues if they do not appear in the consideration sets. Table 9 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 pool⁵⁴.

Welfare Impact of the Training Since the training intervention increases treated new sellers' likelihood to appear in consumers' consideration sets, we estimate the impacts of the

⁵⁴The sample we used in this estimate is not the same as the sample used in 4 because of the model specification. See appendix B.1 for more details on how the sample is constructed.

training by limiting some new sellers' appearance in the search results. Without the training, treated new sellers' likelihood to appear in the sampling pool should be similar to that of the control sellers. Assuming that only training participants are affected by the training, we evaluate the welfare of the training by randomly removing 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. As a result, more consumers will end up visiting non-participants, control new sellers and incumbents instead of training participants, which is what we expect in the absence of the training. Restricting the effect of the training on training participants allows us to account for selection into training participation. The average difference between the estimated quality of training participants and that of non-participants is small. Appendix B.3 discuss the details of the counterfactual exercise.

As presented in the first row of table 9, treated new sellers' revenue share drops by 7.7% in the absence of training. Total sellers' revenues decrease by 0.05%, and consumer surplus decreases by 0.07% as a result of lowering the likelihood to visit higher quality training participants. Even though welfare loss is small in percentage terms, the absolute magnitude of welfare loss could be substantial because of the volume of total transactions on the platform. To decompose the source of revenue growth induced by the training, we compare the changes in the market shares by different types of consideration sets and sellers in these sets. Compared to the baseline, revenue generated from consideration sets with treated new sellers significantly drops when we limit treated new sellers' presence. The drop suggests that most of the revenue growth induced by the training is driven by market expansion as the consumers are more likely to purchase when treated new sellers appear in the consideration sets. For market reallocation, we compare the market share of different types of sellers if they appear in the consideration sets with at least one new seller. The market share of the treated new sellers drops in these sets, while control new sellers capture a higher market share.

6 Conclusions

In this paper, we study how a business training intervention can be an effective way to lift growth barriers new entrants face in a competitive e-commerce platform where sellers face demand-side frictions. Leveraging the experimentally randomized access to the training and the unique consumer-seller matched search and browsing data, we show that the business training helps new sellers increase their presence in consumers' consideration sets and earn higher revenues. The resulting changes in the composition of consumers' consideration sets are beneficial for the consumers as they enjoy higher matching efficiency without lowering the purchases' quality. Using a structural model where we highlight the mis-

match between the number of consumers sellers acquire and their quality, we show that the training increases consumers' welfare and total revenues by limiting the extent of misallocation. The improved matching quality and new sellers' higher service quality could lower consumers' search costs as well, which we do not account for in the current analysis. Enhancing matching quality to improve consumers' experience and promising sellers' growth potential is consistent with the platform's long-term profit-maximizing goal. In the short-run, market expansion and sellers' increased engagement with online marketing also contributes to the platform's profits.

Our findings provide one of the first sets of direct empirical evidence on the consumer welfare implications of an intervention that supports the subset of firms. Consumers do not experience adverse selection in this context because the training reduces market frictions and improves matching quality for the higher quality new sellers. As the market operators, the platforms could play critical roles in lifting growth barriers with proper interventions. Doing so is in alignment with the platforms' incentives as profit-maximizing firms in both the short and long run. Although large e-commerce platforms' looming market power should not be overlooked, these platforms indeed create enormous opportunities for SMEs. They have incentives and the capacity to take more active roles to foster the efficiency and equality in online marketplaces they host.

References

- Anderson, Stephen J., Rajesh Chandy, and Bilal Zia**, "Pathways to Profits: The Impact of Marketing vs. Finance Skills on Business Performance," 2018, 64 (12), 5559–5583.
- Arkolakis, Costas, Theodore Papageorgiou, and Olga A Timoshenko**, "Firm learning and growth," *Review of Economic Dynamics*, 2018, 27, 146–168.
- Atkin, David, Amit K Khandelwal, and Adam Osman**, "Exporting and firm performance: Evidence from a randomized experiment," *The Quarterly Journal of Economics*, 2017, 132 (2), 551–615.
- Bai, Jie, Maggie Chen, Jin Liu, and Daniel Yi Xu**, "Search and Information Frictions on Global E-Commerce Platforms: Evidence from AliExpress," 2020.
- Banerjee, Abhijit, Dean Karlan, and Jonathan Zinman**, "Six randomized evaluations of microcredit: Introduction and further steps," *American Economic Journal: Applied Economics*, 2015, 7 (1), 1–21.
- Banerjee, Abhijit V and Esther Duflo**, "Do firms want to borrow more? Testing credit constraints using a directed lending program," *Review of Economic Studies*, 2014, 81 (2), 572–607.
- Berman, Nicolas, Vincent Rebeyrol, and Vincent Vicard**, "Demand learning and firm dynamics: evidence from exporters," *Review of Economics and Statistics*, 2019, 101 (1), 91–106.
- Bloom, Nicholas and John Van Reenen**, "Measuring and explaining management practices across firms and countries," *The Quarterly Journal of Economics*, 2007, 122 (4), 1351–1408.
- , **Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts**, "Does management matter? Evidence from India," *The Quarterly Journal of Economics*, 2013, 128 (1), 1–51.

- Broda, Christian and David E Weinstein**, "Globalization and the Gains from Variety," *The Quarterly Journal of Economics*, 2006, 121 (2), 541–585.
- Bruhn, Miriam, Dean Karlan, and Antoinette Schoar**, "The impact of consulting services on small and medium enterprises: Evidence from a randomized trial in Mexico," *Journal of Political Economy*, 2018, 126 (2), 635–687.
- Brynjolfsson, Erik, Yu Hu, and Michael D Smith**, "Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers," *Management Science*, 2003, 49 (11), 1580–1596.
- Calderon, Gabriela, Jesse M Cunha, and Giacomo De Giorgi**, "Business Literacy and Development: Evidence from a Randomized Controlled Trial in Rural Mexico," *Economic Development and Cultural Change*, 2020, 68 (2), 507–540.
- Chen, Maggie X and Min Wu**, "The value of reputation in trade: Evidence from Alibaba," *Review of Economics and Statistics*, 2020, pp. 1–45.
- Cohen, Peter, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe**, "Using big data to estimate consumer surplus: The case of uber," Technical Report, National Bureau of Economic Research 2016.
- Couture, Victor, Benjamin Faber, Yizhen Gu, and Lizhi Liu**, "Connecting the countryside via E-commerce: evidence from China," Technical Report, National Bureau of Economic Research 2018.
- Cui, Ruomeng, Jun Li, and Dennis J Zhang**, "Reducing discrimination with reviews in the sharing economy: Evidence from field experiments on Airbnb," *Management Science*, 2020, 66 (3), 1071–1094.
- Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan**, "Consumer price search and platform design in internet commerce," *American Economic Review*, 2018, 108 (7), 1820–59.
- Drexler, Alejandro, Greg Fischer, and Antoinette Schoar**, "Keeping it simple: Financial literacy and rules of thumb," *American Economic Journal: Applied Economics*, 2014, 6 (2), 1–31.
- Einav, Liran, Peter J Klenow, Benjamin Klopock, Jonathan D Levin, Larry Levin, and Wayne Best**, "Assessing the gains from e-commerce," *Unpublished manuscript*, 2017.
- Ellison, Glenn and Sara Fisher Ellison**, "Match quality, search, and the Internet market for used books," Technical Report, National Bureau of Economic Research 2018.
- Farronato, Chiara and Andrey Fradkin**, "The welfare effects of peer entry in the accommodation market: The case of airbnb," Technical Report, National Bureau of Economic Research 2018.
- Fitzgerald, Doireann, Stefanie Haller, and Yaniv Yedid-Levi**, "How exporters grow," Technical Report, National Bureau of Economic Research 2016.
- Foster, Lucia, John Haltiwanger, and Chad Syverson**, "The slow growth of new plants: Learning about demand?," *Economica*, 2016, 83 (329), 91–129.
- Fradkin, Andrey**, "Search frictions and the design of online marketplaces," *Working Paper*, 2015.
- Goldfarb, Avi and Catherine Tucker**, "Digital economics," *Journal of Economic Literature*, 2019, 57 (1), 3–43.
- Gourio, Francois and Leena Rudanko**, "Customer capital," *Review of Economic Studies*, 2014, 81 (3), 1102–1136.
- Hjort, J and J Poulsen**, "The arrival of fast internet and skilled job creation in Africa," *American Economic Review*, 2018.

- Horton, John Joseph**, “Misdirected search effort in a matching market: causes, consequences and a partial solution,” in “Proceedings of the fifteenth ACM conference on Economics and computation” 2014, pp. 357–357.
- Hu, Allen and Song Ma**, “Persuading Investors: A Video-Based Study,” Technical Report, Working paper 2020.
- Hui, Xiang, Maryam Saeedi, Zeqian Shen, and Neel Sundaresan**, “Reputation and Regulations: Evidence from eBay,” *Management Science*, 2016, 62 (12), 3604–3616.
- Jo, Yoon J, Misaki Matsumura, and David E Weinstein**, “The Impact of E-Commerce on Relative Prices and Consumer Welfare,” Technical Report, National Bureau of Economic Research 2019.
- Lehdonvirta, Vili, Otto Kässi, Isis Hjorth, Helena Barnard, and Mark Graham**, “The global platform economy: A new offshoring institution enabling emerging-economy microproviders,” *Journal of Management*, 2019, 45 (2), 567–599.
- Lipsman, Andrew**, “Global Ecommerce 2019: Ecommerce Continues Strong Gains Amid Global Economic Uncertainty. eMarketer.,” Technical Report, <https://www.emarketer.com/content/global-ecommerce-2019> 2019.
- McKenzie, David**, “Small business training to improve management practices in developing countries: Reassessing the evidence for “training doesn’t work”,” *Oxford Review of Economic Policy*, forthcoming.
- **and Christopher Woodruff**, “What are we learning from business training and entrepreneurship evaluations around the developing world?,” *The World Bank Research Observer*, 2014, 29 (1), 48–82.
 - **and Susana Puerto**, “Growing Markets through Business Training for Female Entrepreneurs: A Market-Level Randomized Experiment in Kenya,” *American Economic Journal: Applied Economics*, forthcoming.
- Muralidharan, Karthik and Paul Niehaus**, “Experimentation at scale,” *Journal of Economic Perspectives*, 2017, 31 (4), 103–24.
- of Commerce, PRC Ministry**, “E-commerce in China,” 2020.
- Piveteau, Paul**, “An empirical dynamic model of trade with consumer accumulation,” Technical Report, Working Paper 2016.
- Rotemberg, Martin**, “Equilibrium effects of firm subsidies,” *American Economic Review*, 2019, 109 (10), 3475–3513.
- Tadelis, Steven**, “Reputation and feedback systems in online platform markets,” *Annual Review of Economics*, 2016, 8, 321–340.
- Train, Kenneth E**, *Discrete choice methods with simulation*, Cambridge University Press, 2009.
- Ursu, Raluca M**, “The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions,” *Marketing Science*, 2018, 37 (4), 530–552.
- World-Bank**, *World Development Report 2016: Digital Dividends*, World Bank Publications, 2016.
- Zhang, Dennis J, Hengchen Dai, Lingxiu Dong, Qian Wu, Lifan Guo, and Xiaofei Liu**, “The value of pop-up stores on retailing platforms: Evidence from a field experiment with alibaba,” *Management Science*, 2019, 65 (11), 5142–5151.

7 Tables and Figures

Table 1: Baseline Summary Statistics: Treatment Assignment and Participation

	Full Sample			Sellers in Treatment Group		
	Treatment (1)	Control (2)	(1) - (2) Difference (3)	Participants (4)	Non Participants (5)	(4) - (5) Difference (6)
Is Firm	0.268 (0.443)	0.264 (0.441)	0.004*** (3.33)	0.275 (0.446)	0.266 (0.442)	0.008*** (3.18)
Female Owner (among individual sellers)	0.455 (0.497)	0.453 (0.497)	0.002 (1.27)	0.442 (0.497)	0.458 (0.498)	-0.16*** (4.61)
Region: Coastal South	0.435 (0.496)	0.434 (0.496)	0.001 (0.28)	0.518 (0.500)	0.413 (0.492)	0.105*** (3.573)
Region: West	0.118 (0.323)	0.118 (0.323)	0.0001 (0.34)	0.089 (0.284)	0.126 (0.332)	-0.038*** (21.41)
List Products on Day One	0.213 (0.409)	0.212 (0.409)	0.0005 (0.44)	0.245 (0.430)	0.204 (0.403)	0.04*** (16.04)
Number of Listed Products	1.539 (2.166)	1.535 (2.164)	0.004 (0.65)	2.307 (1.988)	1.342 (2.167)	0.965*** (80.81)
Traffic	2.813 (2.289)	2.807 (2.283)	0.006 (0.91)	4.028 (2.259)	2.501 (2.191)	1.527*** (115.35)
Conversion Rate	0.051 (0.161)	0.051 (0.162)	0.0002 (0.52)	0.055 (0.120)	0.050 (0.171)	0.006*** (7.34)
Revenues	2.145 (3.312)	2.134 (3.304)	0.01 (1.144)	3.842 (3.726)	1.708 (3.048)	2.134*** (100.66)
Observations	177,026	535,092	712,118	36,189	140,837	177,026

Notes: Columns 1, 2, 4, and 5 present means and standard deviations (in parentheses). Column 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. Column 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. * significant at 10% level, ** significant at 5% level and *** significant at 1% level.

Table 2: Overall Treatment Effects on Sellers' Performance

	<i>Dependent variable:</i>							
	Log Revenues	Any Revenues	Log Revenues	Log # Visitors		Log # Buyers		Conversion
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	0.017*** (0.006)	0.002** (0.001)	0.026** (0.010)	0.013** (0.006)	0.024*** (0.008)	0.008*** (0.003)	0.018*** (0.006)	0.0001 (0.0002)
Dep Var Mean	1.39	0.19	7.14	1.73	4.29	0.57	1.4	0.04
Sample	Full	Full	Earn Revenues	Full	Have Visitors	Full	Have Visitors	Have Visitors
Observations	6,409,062	6,409,062	1,253,284	6,409,062	2,593,762	6,409,062	2,593,762	2,593,762
R ²	0.132	0.152	0.081	0.207	0.111	0.105	0.070	0.043
Adjusted R ²	0.132	0.152	0.081	0.207	0.111	0.105	0.070	0.043

Notes: Dependent variables are monthly outcome for all sellers in the new seller sample. Number of visitors, number of buyers and revenues (total payments received) are monthly total in log after adding one to the level. Any revenues is an indicator for earning positive revenues during the month. Conversion is the conversion rate defined as the share of visitors that make a purchase. Column 3 restricts the sample to sellers that earn non-zero revenue during the month. Column 5, 7 and 8 include only sellers that have visitors during the month. All regressions include cohort, relative month and main sector fixed effect as described in equation 1. Dependent variable means are calculated with sellers in the control group. Standard errors are clustered by seller. * significant at 10% level, ** significant at 5% level and *** significant at 1% level.

Table 3: Treatment Effect on Sellers' Strategies

Variable	Treatment	Dep Var Mean	Variable	Treatment	Dep Var Mean
Market Participation			Marketing		
Post Products	0.001 (0.001)	0.34	Paid Ads (Product Counts)	0.002** (0.001)	0.07
Paid Deposits	0.001 (0.001)	0.31	Paid Ads (Traffic Share)	0.001* (0.0004)	0.03
			Promotion	0.0002* (0.0001)	0.0008
Service			Pricing		
Active Time (min)	-0.078 (0.128)	19.17	Avg. Price Per Buyer	0.005 (0.006)	4.39
Reply Time (sec)	-67.725* (34.581)	23341	Avg. Price Per Product	0.0003 (0.001)	3.83
Conversion Rate	0.0014*** (0.0005)	0.1			

Notes: Table presents estimated coefficients β on treatment assignment dummy with specification 1. Standard errors clustered by seller. All regressions include month, entry date and main industry fixed effect. Post products and paid deposits are indicators for having any active product postings or having put down some deposits during the month. Active time is total number of minutes that sellers' account is active and can answer customer inquiries. Reply time is number of seconds average customers waited to hear responses from sellers when making inquiries. Conversion rates are measured as share of consumers making purchases among those who made inquiries. Paid ads (product counts) is number of products participating in paid-for-clicks ads. Paid ads (traffic share) is number of consumers visiting sellers' sites from paid channels in log scale (including through paid for clicks ads and other channels). Promotion is number of times sellers participate in the limited time promotional events that the platform regularly organizes. Average price per products and average price per buyers are seller-level prices measured in log scale. Sellers do not have a price measure if they have zero orders. * significant at 10% level, ** significant at 5% level and *** significant at 1% level.

Table 4: Consumers' Interaction with New Sellers and Purchases

	<i>Dependent variable: Purchase</i>				
		Purchase		Log Spending	Log Order Size
	Same Day	In 3 Days	In a Week	Same Day	
	(1)	(2)	(3)	(4)	(5)
% Treated Seller	0.004*** (0.001)	0.004*** (0.002)	0.005*** (0.002)	0.014** (0.007)	0.017** (0.007)
% Control Seller	0.005*** (0.001)	0.004*** (0.001)	0.005*** (0.001)	0.028*** (0.006)	0.030*** (0.006)
Incumbent Mean	0.18	0.21	0.22	0.76	0.78
Treatment - Control	-0.0074 (0.0018)	0.00004 (0.0019)	-0.0065 (0.002)	-0.014* (0.0081)	-0.013 (0.0086)
Observations	1,381,273	1,381,273	1,381,273	1,381,273	1,381,273
R ²	0.680	0.668	0.657	0.698	0.691

Notes: All regressions include search keywords-date-size of consideration set fixed effects, consumer fixed effects and control for average sellers' price level, ratings and number of products offered as well as consumers' baseline characteristics following equation 2. Purchase are dummies for consumers purchasing from some sellers in the consideration set on the day of visit, within 3 days and within a week of visit. Spending is total payments made and order size is the size of order before applying discounts. The latter is the main performance metrics for sellers on the platform. For 10% of the cases, consumers placed orders but did not complete the payments. The bottom rows present t-tests for listed coefficients with standard errors. * significant at 10% level, ** significant at 5% level and *** significant at 1% level.

Table 5: Purchase and Spending within Consideration Sets

	<i>Dependent variable:</i>					
	Purchase		Purchase		Log Spending	
	Same Day		In a Week		Same Day	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated Seller	0.016*** (0.003)	0.010*** (0.003)	0.010*** (0.003)	0.014*** (0.004)	0.074*** (0.015)	0.050*** (0.016)
Control Seller	0.009*** (0.003)	0.003 (0.003)	0.007** (0.003)	0.010*** (0.003)	0.039*** (0.012)	0.010 (0.013)
Incumbent Mean	0.17	0.17	0.21	0.21	0.71	0.71
Treatment - Control	0.0063 (0.0041)	0.0071* (0.0041)	0.0026 (0.0045)	0.0041 (0.0044)	0.035* (0.0019)	0.039** (0.0019)
Seller Controls	No	Yes	No	Yes	No	Yes
Observations	300,273	300,273	300,273	300,273	300,273	300,273
R ²	0.119	0.126	0.076	0.084	0.145	0.149

Notes: Sample restricted to set of consumers appearing in at least 2 sets, sets with at least one new seller where consumers made a purchase within a week. Column 2, 4 and 6 include seller's ratings, price level and number of products listed. Purchase are dummies for consumers purchasing from the specific sellers in the consideration set on the day of visit or within a week of visit. Spending is total payments made. All regressions include consideration set fixed effects. * significant at 10% level, ** significant at 5% level and *** significant at 1% level.

Table 6: Effect on Quality of Purchase

	<i>Dependent variable:</i>		
	Refund Same Order (1)	Return (2)	Repeat Purchase In a Month (3)
Treated Seller	-0.005 (0.018)	0.0001 (0.012)	0.020 (0.016)
Control Seller	-0.027* (0.015)	-0.015 (0.009)	0.016 (0.013)
Incumbent Mean	0.072	0.025	0.051
Treatment - Control	0.022 (0.021)	0.015 (0.014)	0.0049 (0.02)
Control	Yes	Yes	Yes
Observations	54,936	54,936	54,936
R ²	0.950	0.939	0.943

Notes: Sample restricted to set of consumers appearing in at least 2 sets, sets with at least one new seller and consumer-seller pairs where consumers actually placed orders. Control variables include consumer's recent spending and searching, consumers' ratings (a proxy for their cumulative experience on the platform) as well as seller's ratings, price level and number of products listed. All regressions include consideration set fixed effects. * significant at 10% level, ** significant at 5% level and *** significant at 1% level.

Table 7: Estimated Sellers' Type ξ

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
All Sellers	52,241	0.141	0.314	-2.828	0.100	0.100	5.024
By Purchase Status							
No Purchase	37,617	0.098	0.051	-2.828	0.100	0.100	0.101
Has Purchase	14,624	0.252	0.573	-1.965	0.100	0.100	5.024
By Sellers' Type							
Control	3,366	0.136	0.296	-1.445	0.100	0.100	4.939
Treatment	988	0.150	0.353	-0.853	0.100	0.100	3.668
Incumbent	47,887	0.142	0.314	-2.828	0.100	0.100	5.024
By Type Among Sellers with Purchase							
Control	788	0.258	0.593	-1.370	0.100	0.100	4.939
Treatment	252	0.296	0.679	-0.853	0.100	0.100	3.668
Incumbent	13,584	0.251	0.570	-1.965	0.100	0.100	5.024

Notes: Distribution of estimated ξ_j on a sub-sample with 3000 new sellers and the associated incumbents. See appendix B.1 for details. Sellers are grouped based on whether or not at least one consumers have made purchase from these sellers.

Table 8: Sellers' Characteristics and Estimated Type ξ

	<i>Dependent variable:</i>		
	ξ_j		
	Full	Purchase	No Purchase
	(1)	(2)	(3)
Control	0.015 (0.010)	0.142*** (0.032)	0.021*** (0.003)
Treatment	0.017* (0.019)	0.109*** (0.055)	0.013** (0.005)
Constant	0.096*** (0.003)	0.615*** (0.007)	-0.110*** (0.001)
Observations	52,241	14,624	37,617

Notes: Distribution of estimated ξ_j on a sub-sample with 3000 new sellers and the associated incumbents. See appendix B.1 for details. Sellers are grouped based on whether or not at least one consumer has made purchase from these sellers.

Table 9: Welfare and Market Share Decomposition

Revenue	CS	Market Share		Market Expansion		Market Allocation	
		Treatment	Control	Treatment	Control	Treatment	Control
Restrict Access to Training Participants							
-0.05	-0.07	-7.71	0.17	-10.87	0.09	-4.94	3.17
(0.05)	(0.03)	(0.63)	(0.11)	(0.44)	(0.16)	(0.52)	(0.34)

Notes: Welfare and market share calculated with a random sample of 60,000 consumers. See B.3 for details. 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. To capture the impacts on market expansion, we calculate the share of revenues coming from consideration sets with at least one treated new seller, at least one control new seller or only incumbents and calculate market shares of these sets under different matching rules. Market allocation is calculated as share of revenues from treated new sellers, control new sellers and incumbents that appear in the consideration sets that include some new sellers. Results presented here are percentage difference in market share from the baseline.

Figure 1: Screenshot of Training Module and Tasks on Seller's Portal App



Figure 2: Experimental Sample and Treatment Assignment Over Time

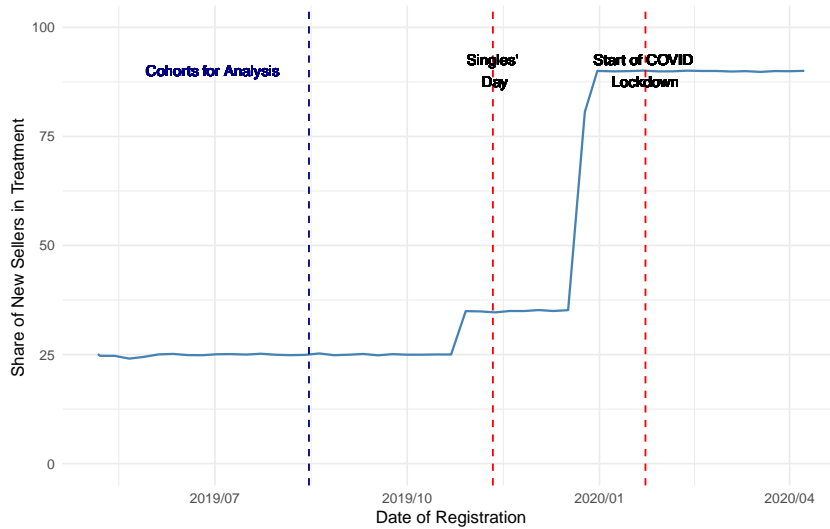


Figure 3: Seller Retention: Training Participation





Figure 4: Distribution of Estimated ξ_j



A Additional Empirical Results: Impacts on New Sellers

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

$$Y_{imcs} = \sum_{m=1}^9 \beta^m Treatment_i M_{im} + \alpha_m + \alpha_c + \alpha_s + \epsilon_{imcs} \quad (9)$$

where M_{im} is the set of indicators for month $m = 1, \dots, 9$. 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.

The treatment effect on performance is relatively short-lived. Figure A2 presents the estimated coefficients $\{\beta^m\}_{m=1}^9$ in specification (9) on monthly revenues, traffic, conversion rate and number of consumers making purchases. The treatment effect of the training access is positive and significant during the 2nd, 3rd and 4th months on traffic and revenues, but remain insignificant for conversion rate. However, the magnitude of estimated coefficients does not differ significantly from month to month. IV results with actual task take up as the first stage variable follow similar temporal patterns, where the impacts are the strongest during the second and third month on traffic and revenues. The pattern of treatment effect on revenues is consistent with the timing of sellers' participation in the training: most of the training participants took up tasks during the initial months (as described in section 2.4), and there could be a lag before actions induced by training become effective

A.1 Heterogeneity

Is the training particularly effective for certain types of new sellers? We examine the heterogeneous impacts 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), locations, whether any products were posted on the first day and whether the store is registered as a B2C store on the first day⁵⁵. Since no sellers took up tasks on the first day of entry, we consider listing products and

⁵⁵To 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.

sellers' registration type on the first day as part of pre-treatment characteristics. We then estimate the following specification

$$Y_{imcs} = \beta Treatment_i + \gamma Type_i + \delta Treat_i \times Type_i + \alpha_m + \alpha_c + \alpha_s + \epsilon_{imcs} \quad (10)$$

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 A8 summarizes the results on log revenues by sellers' types. Overall, estimated β in these specifications have similar magnitudes as estimates using equation 1, but there are no differential treatment effects by sellers' 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 results⁵⁶.

To evaluate the impacts of the offline business environment, we group the sellers based on their registered locations⁵⁷. Table A9 presents the results on heterogeneous treatment effect by sellers' 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 1)⁵⁸. 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.

⁵⁶We also do not find training to be more useful for sellers who post products after the first day of entry.

⁵⁷The 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.

⁵⁸Training participants are significantly less likely to be sellers from the western provinces.

B Details on Structural Estimation

B.1 Data and Sampling

The sample we use for the structural estimation of consumer demand is an adapted version of the sample used in 4.2. The population is the consumer-seller pair sample where each consumer i seller j pair belongs to a search query - efforts - date combination s . The difference from the sample used in 4.2 is that we also include consideration sets that contain incumbents only, as opposed to just new sellers. Because of the computational constraint and setup of the model, we cannot estimate the model on the full sample. Instead, we use the following approach to construct the sample:

1. Randomly sample 3,000 new sellers (treated and control) from the pool of new sellers that appear in the full consumer-seller pair sample.
2. For each new seller, obtain all the incumbents that appear in the same consideration sets s as these new sellers do, include these sellers to the estimation sample.
3. For all the incumbents who appear in the same sets with the new sellers, obtain all the consideration sets these sellers appear in and other incumbents that appear in the same sets as they do, add these sellers to the estimation sample.
4. Iterate the previous step until all the new sellers and incumbents in the estimation sample appear in at least two consideration sets.

We require all the sellers in the estimation sample to appear in at least two consideration sets because otherwise, the ξ_j would not be identified. As a result, our final sample consists of sellers who have higher traffic shares because they are more likely to appear in multiple consideration sets. We use traffic to refer to the number of visitors a seller obtains in a 30-day period. The final sample consists of 52,241 sellers and 1,312,967 observations in 323,584 consideration sets. 3,366 sellers are in the control group, and 988 sellers belong to the treatment group. We have more new sellers than what we originally sampled because additional new sellers are incorporated into the estimation sample in the iteration process. 18.1% of the observations are selected as part of the estimation sample.

The estimation sample is at the seller level, even though in the search process, consumers access specific sellers' sites by searching specific products rather than the front page of the sites. Sellers almost always offer multiple products and sometimes could span different sectors. Consumers could purchase multiple products from the sellers they visited, and in particular, they could purchase products other than the ones that direct them to sellers' sites. Since we do not observe consumers' browsing history in the stores, we cannot fully capture such a process. Instead, since we are predominantly interested in the

seller level characteristics, we aggregate all the purchase and browsing behaviors to the consumer-seller level, rather than to the consumer-seller-product level. The number of products offered and the ratings are measured at the seller level. Pricing is the quantity weighted prices of all products that the sellers offer.

The instruments we used for sellers' pricing and the number of products offered are variables that capture the platform's rule enforcement intensity. The variables are constructed as averaging over all the sellers belonging to the same sector during the past 30 days period weighted by each seller's traffic. We subtract seller j 's own weighted measure from the weighted mean. These instruments include frequencies of identified and enforced rule violations and share of sellers identified as selling fake or counterfeit products or as boosting sales with fake orders. We standardized these variables to have mean zero and unit standard deviations for estimation.

B.2 Estimation Details

We modify the baseline model as described in 5.2. In the actual estimate, we use 5 to approximate the probability of consumer i choosing a specific seller j in the consideration set, which gives the probability parameter in a Bernoulli distribution. We use the simulated maximum likelihood to identify the parameters of interests by matching with realized purchases on the day of visits. In the baseline model, we include the sector specific fixed effect ξ_s but keep the price coefficient α constant for all sectors. We enrich the model with sector specific coefficient α_s and the estimated results are similar. To account for endogeneity of product offerings and pricing, we use instruments mentioned above. We model price level and log number of products offered following multinormal distribution where the respective means are determined by equation 6. We jointly estimate the first stage for strategies with the consideration set based demand.

To estimate the matching rule, we use the new seller sample described in 3 where we re-define month relative to the time when sellers first post the products. As described in 5.2, we include dummies for having no visitors or zero conversion rate in the previous periods. The traffic measure is converted to log scale. We use a flexible polynomials for $f(\mathcal{T}_{jt-1}, C_{jt-1}, C_{jt-2})$ and test for the changes in R^2 when adding higher order terms of $\mathcal{T}_{jt-1}, C_{jt-1}$. We also add relative month, calendar date and initial sector fixed effects. Adding higher order terms of lagged traffic and conversion does not significantly improves precision of predicted traffic. Excluding the fixed effects will reduce R^2 by construction but the impacts on predicted traffic are small. The most important predictor is the lagged traffic and the relationship between current period traffic and previous period traffic is close to linear. To test the precision of the prediction, we use a cross validation method and calculate the average residuals on the training sample. Table A13 shows measures of

prediction precision on current period traffic with different specifications. For the actual estimate to generate predicted traffic, we use estimated coefficients on \mathcal{T}_{jt-1} and C_{jt-1} and C_{jt-2} with linear specification without fixed effects.

B.3 Counterfactual Details

To run the counterfactual analysis, we randomly sample 60,000 consumer - search keyword - date combinations and obtain their corresponding number of sellers visited. We construct the potential pool of sellers that a specific consumer searching a particular keyword could sample from as **all** the sellers who were visited by any consumers searching those keywords on that date. For each seller in the pool, we calculate their predicted traffic \hat{T}_j using the estimated $f(\cdot)$ as described in the previous section using lagged traffic and lagged conversion rates. The predicted traffic is a good approximation of the actual traffic these sellers acquire. Ideally, we should use query-specific predicted traffic as the sampling weights, but such data is not currently available. The sampling weights for seller j in a specific query-date pool is given by $w_{js}^0 = \frac{\hat{T}_j}{\sum_{k \in S} \hat{T}_k}$. We use w_{js}^0 as the baseline sampling weights. With the estimated sampling weights for all the sellers in the pool, we randomly sample sellers from a consumer-set specific pool where the number of sellers is the same as the number of sellers the consumer visits during that particular search session. Therefore, the only part that is changed in the counterfactual analysis is the composition of sellers in consumers' consideration sets. In contrast, the size of the sets and which keywords consumers searched are all kept constant.

For the counterfactual analyses, we adjust sellers' sampling weights to consider the welfare impacts of the current training. To quantify the impacts of the training, we assume that treated new sellers should have the same behaviors as the control new sellers. Among all the sellers that appear in the sampling pooling, 5.9% of them are control new sellers and these sellers account for 3.2% of observations in the sampling pool. Currently, 1.74% of sellers in the pool are treated new sellers, and they make up for 1.07% of the appearance. Without the training, treated new sellers should make up for a similar share of the observations in the sampling pool as the control new sellers do, in which case they should make up for 0.94% of the observations, which is a 12.2% drop from their current shares. If we further assume that only training participants are subject to the training's influence, and since that training participants make up for 51.16% of the treated new sellers, their appearance in the sampling pool should be dropped by 24.66%. Therefore, in the counterfactual analysis, we randomly drop 24.66% of training participants' appearance from the sampling pool and recalculate other sellers' sampling weights after removing consumers' access to these sellers. With the new sampling pool and the updated weights, we reconstruct consumers' consideration sets and calculate the impacts on consumer surplus and

sellers' revenues following the specifications in section 5.4. To decompose the changes in sellers' revenues, we compute the market share of three types of consideration sets and different types of sellers' market shares if they appear in the consideration sets with some new sellers under each counterfactual scenario. We then calculate the changes in market shares from the baseline market allocation.

C Online Sellers' Survey

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 A15. 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 are small and inexperienced, a substantial share of them are 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 collect demographic data of from the sellers other than those collected during the registration⁵⁹.

⁵⁹We could potentially gather more information such as predicted education, income level and total spending on the platform through the affiliated financial subsidiaries.

D Additional Tables and Figures

Table A1: Treatment Effects on Monthly Revenues

Indicator	<i>Dependent variable:</i>						
	Raw	Monthly Revenues		Winsorized			
		Log	I.H.S.	99th	99.5th	99.9th	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treatment	0.002** (0.001)	-939.245 (1,447.148)	0.017*** (0.006)	0.018*** (0.007)	25.436** (12.478)	26.175 (22.308)	85.084 (54.321)
Dep Var Mean	0.19	7018.16	1.39	1.39	1322.79	2019.61	3472.93
Observations	6,409,062	6,409,062	6,409,062	6,409,062	6,409,062	6,409,062	6,409,062
R ²	0.152	0.0001	0.132	0.134	0.027	0.024	0.019
Adjusted R ²	0.152	0.00004	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 A2: Treatment and Speed of Listing Products

<i>Dependent variable:</i>					
Number of Days Passed Before Listing First Items					
	OLS	OLS	IV	IV	OLS
	(1)	(2)	(3)	(4)	(5)
Treatment	0.057 (0.144)	0.067 (0.207)			
Took-up Tasks (Fitted)			0.835 (2.121)	0.664 (2.046)	
Took-up Tasks					7.180*** (0.597)
Sample	All	Late	All	Late	Late, Treatment
Dep Var Mean	16.05	23.99	16.05	23.99	24.03
Observations	476,292	318,792	476,292	318,792	79,477
R ²	0.001	0.002	0.001	0.002	0.004
Adjusted R ²	0.001	0.001	0.001	0.002	0.003

Notes: Sample restricted to sellers who have posted at least one product during the sample period. Dependent variable is number of days passed since registration before sellers posted the first product. Columns 2, 4, and 5 further restricted sample to sellers who posted products on the second day or later. Column 5 again restricts the sample to sellers who have listed products on the second day or later and are assigned to the treatment group. For instrumental regressions in column 3 and 4, the instrument is being assigned to the treatment group. All regressions include cohort fixed effect. Dependent variable means calculated with sellers in the control group. * significant at 10% level, ** significant at 5% level and *** significant at 1% level.

Table A3: Main Results Among Sellers Ever Posted Products

	<i>Dependent variable:</i>					
	Post Products (1)	Any Revenues (2)	Revenues (3)	# Visitors (4)	# Buyers (5)	Conversion Rate (6)
Treatment	0.001 (0.001)	0.002* (0.001)	0.024*** (0.010)	0.018** (0.008)	0.012** (0.005)	0.0001 (0.0002)
Dep Var Mean	0.59	0.32	2.32	2.82	0.95	0.04
Observations	6,409,062	3,802,509	3,802,509	3,802,509	3,802,509	2,593,762
R ²	0.380	0.100	0.081	0.118	0.064	0.045
Adjusted R ²	0.380	0.100	0.081	0.118	0.064	0.045

Notes: The outcome variable for column 1 is an indicator for seller ever posting a product using the full sample. In column 2 to 5 the sample is restricted to sellers that ever post any products in 9-month. In column 6 the sample is further restricted to sellers who have visitors during the month. 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 A4: IV Results on Main Outcomes

	<i>Dependent variable:</i>					
	Take-up Training (1)	Any Revenues (2)	Revenues (3)	# Visitors (4)	# Buyers (5)	Conversion Rate
Treatment	0.257*** (0.001)					
Take-up Training		0.006**	0.066*** (0.024)	0.052** (0.021)	0.032*** (0.012)	0.0001 (0.001)
Dep Var Mean	0	0.19	1.39	1.73	0.57	0.04
Observations	6,409,062	6,409,062	6,409,062	6,409,062	6,409,062	2,593,762
R ²	0.232	0.153	0.133	0.208	0.107	0.049
Adjusted R ²	0.232	0.153	0.133	0.208	0.106	0.048

Notes: The first stage variable is an indicator for whether or not sellers have taken up at least one task during the nine-month period. Column 1 presents the estimated on first stage outcome with treatment assignment as the instrumental variable. All specifications are two-stage least square results using treatment assignment as the instrument. 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 A5: IV Results on Main Outcomes

	<i>Dependent variable:</i>					
	Take-up Training	Any Revenues	Revenues	# Visitors	# Buyers	Conversion Rate
	(1)	(2)	(3)	(4)	(5)	
Treatment	0.257*** (0.001)					
Take-up Training		0.006**	0.066*** (0.024)	0.052** (0.021)	0.032*** (0.012)	0.0001 (0.001)
Dep Var Mean	0	0.19	1.39	1.73	0.57	0.04
Observations	6,409,062	6,409,062	6,409,062	6,409,062	6,409,062	2,593,762
R ²	0.232	0.153	0.133	0.208	0.107	0.049
Adjusted R ²	0.232	0.153	0.133	0.208	0.106	0.048

Notes: The first stage variable is an indicator for whether or not sellers have taken up at least one task during the nine-month period. Column 1 presents the estimated on first stage outcome with treatment assignment as the instrumental variable. All specifications are two-stage least square results using treatment assignment as the instrument. 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 A6: OLS Results on Main Outcomes

	<i>Dependent variable:</i>							
	Log Revenues (1)	Any Revenues (2)	Revenues (3)	Log # Visitors (4) (5)		Log # Buyers (6) (7)		Conversion (8)
Take-up Tasks	1.847*** (0.017)	0.223** (0.002)	0.771*** (0.019)	1.735** (0.014)	0.703*** (0.015)	0.825*** (0.009)	0.686*** (0.012)	0.005*** (0.0004)
Dep Var Mean	1.39	0.19	7.14	1.73	4.29	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 A7: Treatment Effect on Sellers' Ratings and Refunds

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

Notes: Table presents estimated coefficients β on treatment assignment dummy with specification 1. 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 assigns 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 B.1, 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.

Table A8: Treatment Effect Heterogeneity on Sellers' Basic Types

	<i>Dependent variable: Log Revenues</i>				
	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)
Dep Var Mean	1.39	1.39	1.39	1.39	1.39
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 A9: Treatment Effect Heterogeneity by Sellers' Registered Location

	<i>Dependent variable: Log Revenues</i>					
	Beijing Vicinity (1)	Resource-Oriented (2)	Northeast (3)	Coastal South (4)	Central (5)	West (6)
Treatment	0.015** (0.007)	0.018*** (0.006)	0.019*** (0.006)	0.014* (0.007)	0.015** (0.007)	0.019*** (0.007)
Location	-0.039*** (0.008)	-0.371*** (0.013)	-0.273*** (0.012)	0.470*** (0.006)	-0.128*** (0.007)	-0.642*** (0.007)
Treatment × Location	0.011 (0.017)	-0.042 (0.026)	-0.048* (0.025)	0.006 (0.013)	0.008 (0.014)	-0.013 (0.014)
Dep Var Mean	1.39	1.39	1.39	1.39	1.39	1.39
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, Hebei 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.

Table A10: Summary Statistics: Consumer-Search Session Sample

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Number of Sellers Browsed Per Session	1,381,273	4.892	2.691	3	3	6	55
Share of Treated New Sellers	1,381,273	0.016	0.065	0	0	0	1
Share of Control New Sellers	1,381,273	0.023	0.082	0	0	0	1
Purchase (Same Day)	1,381,273	0.188	0.390	0	0	0	1
Purchase (in 3 Days)	1,381,273	0.215	0.411	0	0	0	1
Purchase (in A Week)	1,381,273	0.231	0.422	0	0	0	1
Pay Amount	1,381,273	34.282	151.451	0	0	0	14,649
Order Size	1,381,273	40.865	227.007	0	0	0	88,000
Recent Spending	1,381,273	1,822.888	6,695.639	0	9.9	713	1,960,667
Recent Search	1,381,273	184.069	4,370.249	0	31	165	1,580,514
Consumers' Experience	1,381,273	4.983	2.378	-2	4	7	14
Sellers' Price Level	1,381,273	245.912	616.886	0.01	62.861	235.551	69,100
Number of Listed Products	1,381,273	1,196.689	11,734.010	0	100	831	4,216,488
Seller's Rating	1,381,273	13.032	2.916	-1	11.5	15	20

Notes: Table presents the summary statistics of main variables in the consumer-search session sample. Each observation is a consumer-search session.

Table A11: Summary Statistics: Consumer-Seller Sample

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Number of Sellers Browsed	300,273	4.685	2.621	1	3	5	55
Purchase (Same Day)	300,273	0.183	0.387	0	0	0	1
Purchase (in a Week)	300,273	0.219	0.414	0	0	0	1
Pay Amount	300,273	34.29	18.74	0	0	0	10,000
Refund	54,910	0.068	0.252	0	0	0	1
Return	54,910	0.023	0.150	0	0	0	1
Repeat Purchase	54,910	0.049	0.217	0	0	0	1
Recent Spending	300,273	4,886.792	9,656.899	0	141.8	5,691.6	260,495.5
Recent Search	300,273	160.37	194.731	0	48	211	5,133
Price Level (Seller)	300,273	202.96	392.37	0	53.2	221.5	80,000
Number of Listed Products (Seller)	300,273	527.033	7,681.935	0	13	299	3,359,564
Seller's Rating	300,273	11.209	4.059	-2	5	7	12

Notes: Table presents the summary statistics of main variables in the consumer-seller sample. Each observation is a consumer-seller pair where consumers purchased from some sellers during the search session within a week of visit.

Table A12: Visiting New Sellers and Consumers' Purchase

	<i>Dependent variable: Purchase</i>				
		Purchase		Log Spending	Log Order Size
	Same Day	In 3 Days	In a Week	Same Day	
	(1)	(2)	(3)	(4)	(5)
Treated Sellers	0.003** (0.001)	0.004** (0.002)	0.004** (0.002)	0.011* (0.007)	0.013* (0.007)
Control Sellers	0.004*** (0.001)	0.004*** (0.001)	0.005*** (0.001)	0.023*** (0.005)	0.023*** (0.006)
Incumbent Mean	0.18	0.21	0.22	0.76	0.78
Treatment - Control	-0.00035 (0.0017)	0.000088 (0.0018)	-0.0005 (0.0019)	-0.012 (0.0076)	-0.011 (0.0081)
Observations	1,381,273	1,381,273	1,381,273	1,381,273	1,381,273
R ²	0.680	0.668	0.657	0.698	0.691

Notes: All regressions include search query-date-size of consideration set fixed effects, consumer fixed effects and control for average sellers' price level, ratings and number of products offered as well as consumers' baseline characteristics. * significant at 10% level, ** significant at 5% level and *** significant at 1% level.

Table A13: 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 A14: Example of Tasks and Their Classifications

Task	Indicator	Area of Focus	Function	Type
Acquire customers' reviews	reviews	ratings	knowledge	outcome
Acquire free traffics	visitors from search channel	marketing	knowledge	outcome
Choose proper promotion products	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: inquiry	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 weitaio	followers' activities	customers	reminder	action
Improve per consumer spending	avg. order size	marketing	knowledge	outcome
Improve ratings on customer service	ratings	ratings	knowledge	outcome
Improve ratings on product quality	ratings	ratings	knowledge	outcome
Decorate store frontpage on app	decoration	basic	reminder	action
Improve tag/bookmark rates	bookmarked	marketing	knowledge	outcome
Imrpove ratings on delivery	ratings	service	knowledge	outcome
Optimize products' titles	traffics	marketing	knowledge	outcome
Participate in official sales events	sign-up	marketing	reminder	action
Pay security deposits	deposits	basic	reminder	action
Post products on store page	number of products	basic	reminder	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 paid "wangpu"	wangpu	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 inquiries	response time	service	knowledge	outcome
Upload videos for product descriptions	vidoes	basic	reminder	action

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 A15: Summary of Sellers' Survey

Category	Fraction	Category	Fraction
Respondent Chars		Business Chars	
<i>Education</i>		<i>Sources of Supply</i>	
Primary	2.8%	Own factory	32.5%
Middle School	29.0%	Offline wholesale markets	19.2%
High School	23.7%	Online wholesale markets	21.7%
Some College	28.0%	Distribution/brand subsidiary	19.5%
Bachelors	15.4%	Others	7.1%
Master's and Above	0.8%	<i>Number of Employees (inc. owners)</i>	
Professional Degrees (e.g. MBA)	0.3%	1 - 2 persons	71.9%
<i>Exp in Retail</i>		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%	<i>Total investments</i>	
More than 3 years	20.5%	<5k RMB	32.3%
<i>Exp in E-commerce</i>		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%
<i>Goal</i>			
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.

Figure A1: Quantile Treatment Effect on Revenue Over Time

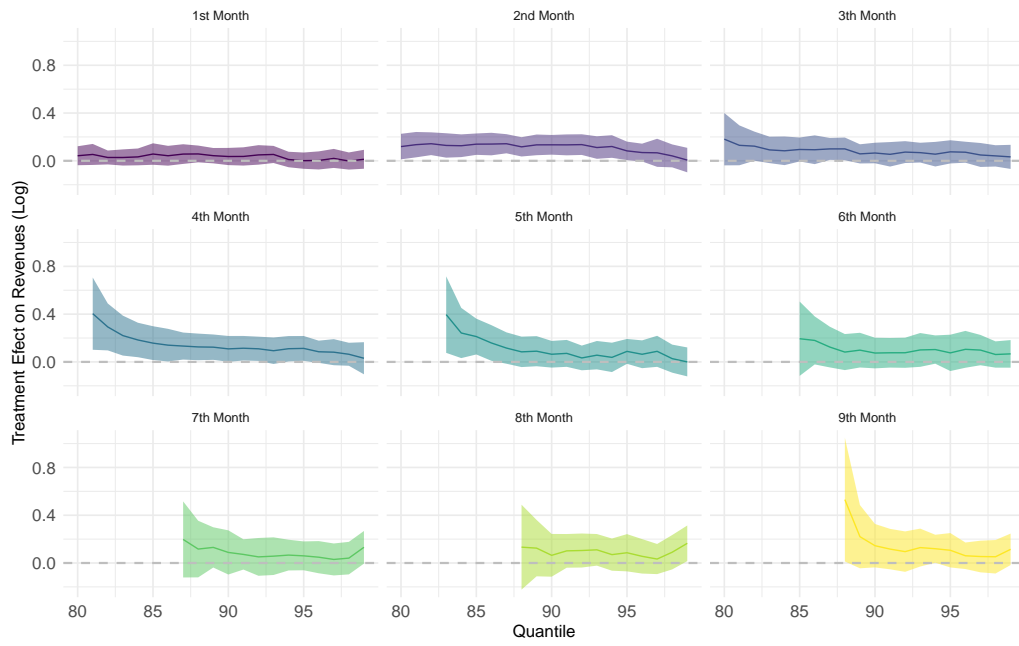


Figure A2: Long Term Treatment Effect on Main Outcomes

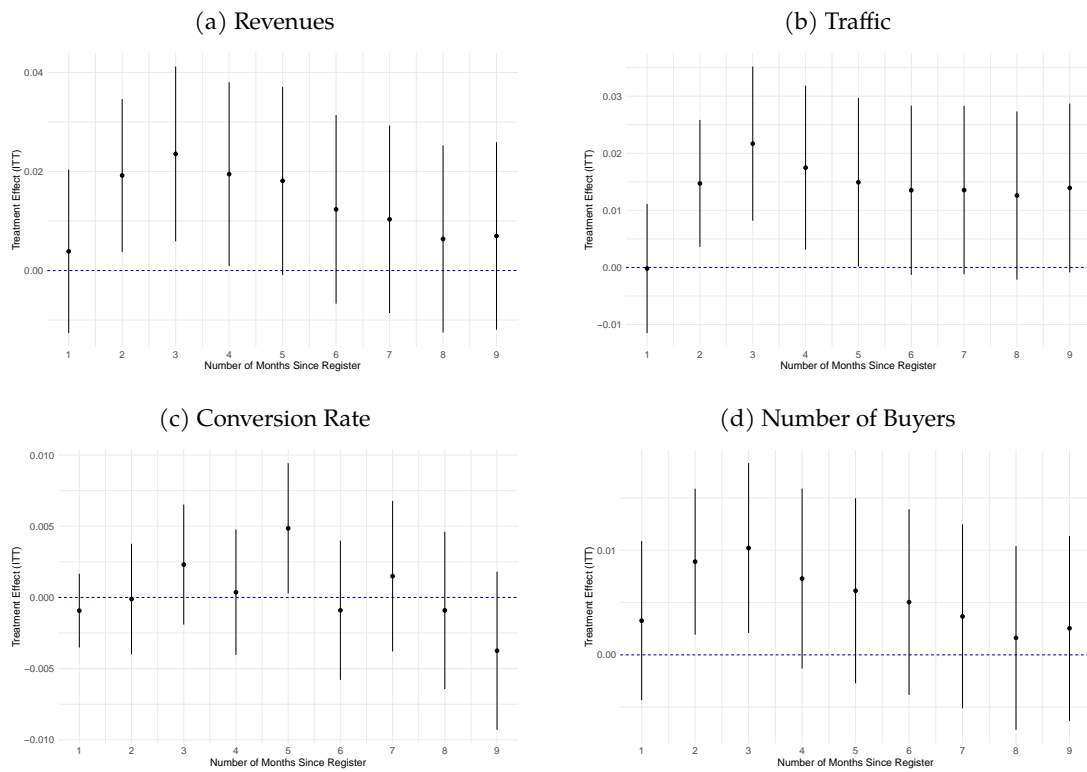


Figure A3: Procedures of Constructing the Consumer Sample

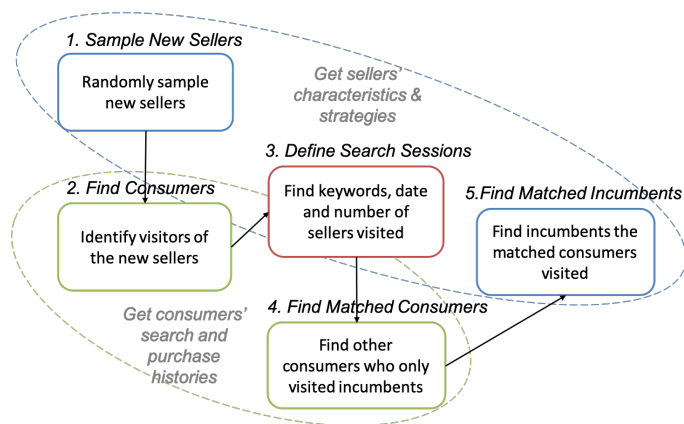


Figure A4: Distribution of Size of Search Sets

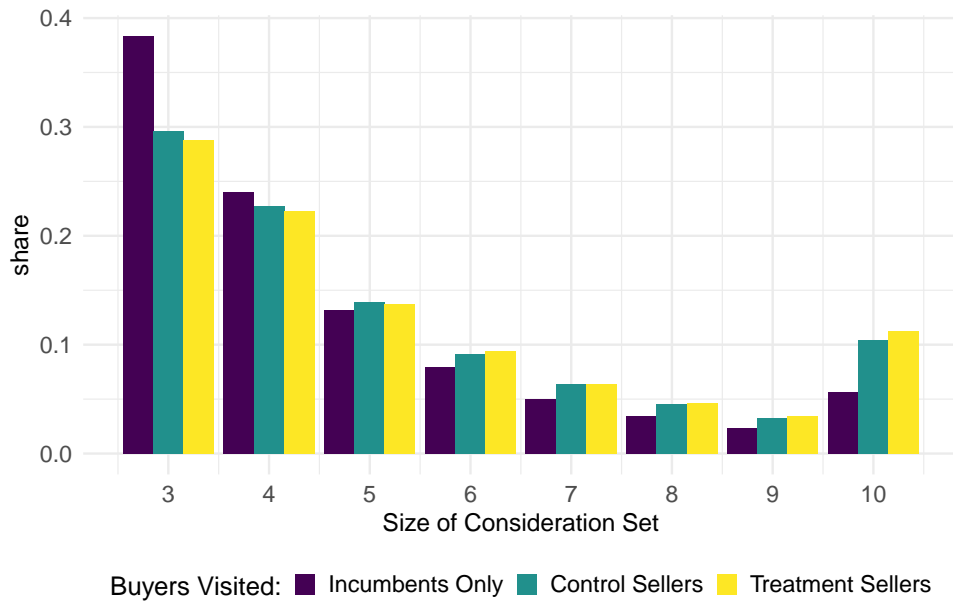


Figure A5: Distribution of Estimated Parameters

